LEARNING AND VISCERAL TEMPTATION IN DYNAMIC SAVINGS EXPERIMENTS[*]

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

Abstract

In experiments on savings with income uncertainty and habit formation, subjects should save early to create a buffer stock, to cushion bad income draws and limit the negative internality from habit formation. We find that 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 drink what they order with a time delay (as predicted by hyperbolic discounting and dual-self models). The estimated present-bias (beta), 0.7-0.8, is consistent with other studies.

---

[*] 18 April 2006. This research was supported by NSF grant SES-0078911. We thank Chris Carroll, Daniel Houser, Paul Kattman, George Loewenstein, Tanga McDaniel, and Nat Wilcox for comments. We also thank Julie Malmquist of the SSEL Caltech lab, Chong Juin Kuan (NUS), Hackjin Kim, Tony Bruguier, and especially Min Jeong Kang for her help in running the experiments.

I. Introduction

Do people save enough? High rates of consumer debt and personal bankruptcies in the United States, and the drop in consumption upon retirement, suggest people save too little. However, it is difficult to conclusively reject the basic premise of lifecycle saving, which is that current saving correctly anticipates future needs and smoothes consumption [Browning and Lusardi, 1996; Venti 2005]. 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. In these tests, apparent statistical evidence of undersaving[1] could easily be due to one or more econometric misspecifications.[2]

Experiments can create an economic environment whose structure matches the assumptions underlying standard theory precisely. In these experiments, whether subjects make optimal decisions or not provides clear evidence on whether theories are on the right track, in environments in which the assumptions of theory clearly hold. The subjects in our experiments are intelligent college students. Most of them will soon be making important savings decisions—establishing credit, and making retirement-fund investment allocations in their first jobs—so it is quite likely their behavior in these simple settings will generalize to their later behavior early in their savings lifecycle. The results can suggest boundary conditions under which saving is optimal, and when it is not, and might inspire new hypotheses. Of course, the experimental environment is much simpler than naturally-occurring environments and the stakes are lower.

---

[1] See Scholtz et al [forthcoming] for an argument that a generation of Americans saved "too much."

[2] 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.

But if theory or evidence suggests how adding naturally-occurring features to the experiment would change the results, those features can be easily included in more complex experiments.

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 the same 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 dynamic program). Optimal saving also requires subjects to save a lot in early periods for two different reasons: (i) To buffer against bad income shocks; and (ii) to avoid creating an early consumption "internality" that reduces utility from future consumption.

This paper reports two studies, using money rewards and beverage rewards.

The first study explores 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 can approximate optimization surprisingly well after learning over repeated lifecycles, or by learning "socially" from good and bad decisions by other subjects.

The second study explores the effect of visceral temptation when subjects are thirsty and their points are translated into sips of beverage (a cola soft drink). This visceral temptation is probably absent from the first study, but is present in food consumption, addiction, procrastination and perhaps in many kinds of spending and borrowing. This second study compares two conditions: An "immediate" condition in which period t decisions lead to immediate drinking; and a "delayed" condition in which the amount they decide to consume in period t is not delivered until period t+10. Subjects consume too much when rewards are immediate, and consumption is closer to optimal when rewards are delayed. The relative overconsumption in the more tempting immediate condition is consistent with models of hyperbolic discounting and dual-self conflict, which predict more consumption in this experiment when rewards are immediate rather than delayed [e.g., Bernheim and Rangel, 2005; Fudenberg and Levine, 2005; Loewenstein and O'Donoghue, 2005].

Together, the two experiments suggest that people will consume too much and save too little if they have not had a chance to learn from their own experiences (or from examples of others), or when rewards are viscerally tempting and immediate. But when subjects have a chance to learn, or the rewards are not consumed immediately, saving rates can be surprisingly close to optimal.

Do people save too little? Many people *say* they save too little. A 1997 survey found 76% of individuals believe they should be saving more for retirement [Angeletos et al, 2001]. Of the respondents who believed they should be "seriously saving already" 55% believed they were "behind" in their savings. Another survey found a difference of about 10% between the percentage of income individuals believed they should be saving for retirement and what they are actually saving [Angeletos et al, 2001].

If people are saving too little, and know that they are, they should seek external commitment devices. Financial advisors suggest setting aside money for savings each month. In fact, diverting a fraction of future pay raises into savings accounts through SaveMoreTomorrow™ plans, so that take-home pay does not fall in nominal terms when saving increases, does increase saving [Thaler and Benartzi, 2004]. Commitment mechanisms such as certificates of deposit, automated deposits and withdrawals, retirement accounts which penalize early withdrawal, and illiquid assets (housing) also "make up an overwhelming majority of assets held in the U.S. by the household sector." [Laibson, 1997]. The success of SMaRT, and the popularity of other mechanisms, is an index of how sophisticated consumers are about their inability to save without help.

Theories of optimal saving. Until the 1990's, most models of optimal saving assumed consumers solved a dynamic programming problem in which stochastic future income was replaced by a certainty-equivalent [e.g., Carroll, 2001]. Actual savings patterns are not consistent with many of the predictions of these models. For example, the median household under age fifty holds only a few *weeks* of income in asset savings (excluding housing), while the certainty-equivalent (CEQ) approach predicts an optimal savings stock of a few *months* of income. Under the CEQ model, households are not saving enough.

The newer "buffer stock" model uses dynamic optimization rather than the traditional quadratic utility functions and certainty-equivalence in the CEQ model [Zeldes, 1989; Deaton, 1991]. Under realistic parameter values, predictions of the buffer-stock model are closer to consumption levels observed empirically than the CEQ [Carroll 1997; Carroll 2001]. In the buffer stock model, individuals build up an ideal level of buffer stock, and then spend roughly their average income once their buffer stock is large enough.

Solving for optimal saving in the buffer stock model is computationally difficult. Until the 1990's, computers were not powerful enough to solve for optimal saving in realistic environments. Allen and Carroll [2001] shows that learning by simple reinforcement is far too slow to produce convergence to optimal saving in reasonable time scales. Therefore, a natural question is whether consumers figure out or learn to save optimally in these environments?

What have previous experiments shown? In investigating whether saving is optimal, a useful technique is to create an experimental environment that matches the assumptions of the theory. If intelligent subjects do not save optimally in these relatively simple experiments, then the results cannot be blamed on model misspecification. If experimental saving is suboptimal, then critics of the theories as descriptions of actual savings are entitled to be skeptical about whether average consumers save optimally in even more complex natural environments.

There have been few experimental tests of savings in lifecycle models, but all of them suggest that experimental subjects do not save enough. Kotlikoff et al [1987] found subjects “overdiscounted” future income in simple environments with no income uncertainty. Fehr and Zych [1998] complicate saving by introducing habit formation— that is, future utilities are lower if previous consumption is higher. They find that subjects do not seem to anticipate this “internality” (an internal intertemporal externality) and save too little. Ballinger et al [2003] divide subjects into social-learning groups of three where each subject in the group completes a savings task sequentially; subjects observe their predecessors and are encouraged to teach their successors. Consistent with social learning, the third subject in the group performs significantly better than the previous two.

Temptation. Our first experimental study uses small monetary rewards. After seeing that subjects in that study *can* learn to optimize rather impressively under some conditions, we wondered whether the abstract money rewards created the kind of visceral temptation present in dynamic optimization problems involving addiction, procrastination, credit card spending, and so forth. We therefore designed a second study to invoke temptation, by using beverage rewards when subjects are thirsty.

The beverage study is also a tool to study dynamic inconsistency. In the immediate condition, “savings” decisions in one period affect the amount of a beverage the thirsty subjects can drink right away. In the delayed condition, current decisions affect the amount of beverage delivered 10 periods later (about 10 minutes, which is a long time when you are thirsty). Under quasi-hyperbolic discounting, or dual-self models in which a "planner" spends cognitive resources suppressing the desires of a myopic "doer" [Shefrin and Thaler, 1988; Fudenberg and Levine 2005], subjects will drink more beverage in the immediate condition than in the delayed condition. Of course, small amounts of beverage are a pale imitation of more dramatic temptations like drug addiction, but the results give us a first contrast between money rewards and visceral temptations and can point to future research.

Laboratory experiments have some advantages for testing dynamic savings models and measuring temptation. The subjects are highly intelligent college students who are best equipped to make mathematical tradeoffs of the type assumed by the theories. And since they are young and soon to enter the workforce, the important first steps in saving for retirement are just ahead of them. Recent studies indicate that smarter students are less likely to violate normative assumptions such as dynamic consistency [Frederick, 2005; Benjamin and Shapiro, 2005]. Most of our subjects are very analytically skilled. If these students make simple mistakes in solving complex dynamic savings problems, then it is possible that average consumers do so too, unless informal and professional advice, or other institutional forces, correct their mistakes.[3]

This paper is organized as follows: Section 2 describes the standard model and the hyperbolic discounting and dual-self models. Section 3 describes the design and results of the first study, using money reward and exploring the power of individual and social learning. Section 5 describes the second study, using actual beverage rewards. Section 6 concludes.

II. The Buffer Stock Model

---

[3] Note also that average people exhibit financial illiteracy which would probably surprise professional economists. For example, when asked whether investing $100 at a 2% annual interest rate would yield more or less than $102 in five years, 22.2% of subjects gave the wrong answer [Lusardi, 2002]. Bernheim [1998] reports many similar statistics.

The experimental design implements the assumptions of the buffer stock savings model of Carroll, Overland and Weir [2000]. Agents earn income each period, subject to stochastic independent shocks from a distribution they know. In each period they choose how to divide their available cash— the previous buffer stock, plus new income-- between spending 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. Agents should maximize the discounted utility from consumption over the remainder of their lifetimes, which is a dynamic programming problem. The variables in this dynamic program are listed in Table I.

1. Exponential discounting

Assuming exponential discounting of future rewards, the consumer's maximization problem is

$$\max E_t\left[\sum_{s=t}^{T}\Delta^{s-t}u(\widetilde{C}_S,\widetilde{H}_{S-1})\right]. \tag{1}$$

The utility function incorporates constant relative risk-aversion (CRRA) and habit formation as follows:

$$u(C_t,H_{t-1}) = k + \frac{\theta}{1-\rho}\left(\frac{C_t+\hat{\varepsilon}}{H_{t-1}^{\gamma}}\right)^{1-\rho}. \tag{2}$$

The parameter ρ is the coefficient of relative risk aversion, and γ determines how strongly previous habitual consumption affects current utility (e.g., if γ=0 there is no effect of habit).[4]

The habit stock is given by $H_t = \lambda H_{t-1} + C_t$, where λ is a depreciation rate [as in Fehr and Zych, 1998]. Actual income each period is equal to permanent income multiplied by an income shock, $Y_t = P_t\eta_t$ where η_t is a random variable drawn from a distribution each period.

---

[4] 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 θ = *750* and *k=40.*

The value function for cumulative future utility, at period t, depends on three state variables—permanent income P_t, habit H_{t-1}, and available savings X_t. The optimal value function is

$$V_t(P_t, X_t, H_{t-1}) = \max_{C_t} \{u(C_t, H_{t-1}) + \Delta E_t [V_{t+1}(P_{t+1}, X_{t+1}, H_t)]\}$$

or, writing out the state variables,

$$V_t(P_t, X_t, H_{t-1}) = \max_{C_t} \{u(C_t, H_{t-1}) + \Delta E_t [V_{t+1}(GP_t, R[X_t - C_t] + \eta_{t+1} GP_t, \lambda H_{t-1} + C_t)]\} \tag{3}$$

subject to constraints

$$S_t = X_t - C_t \text{ (with } 0 \leq S_t \text{),} \tag{4}$$

$$X_{t+1} = S_t + \eta_{t+1} P_{t+1}, \tag{5}$$

$$H_t = \lambda H_{t-1} + C_t. \tag{6}$$

To make the problem easier to solve computationally, the state variable P_t can be eliminated by normalizing each variable by permanent income.[5]

In the experiment, most parameter values were roughly calibrated to those measured in actual savings data from the U.S. Carroll [1992] found income shocks η_t to be lognormally distributed with a mean value of one and a standard deviation of 0.2. We use η_t drawn from a lognormal distribution $\log \eta \sim N(-\frac{1}{2}, 1)$. We pick $\sigma = 1$ rather than .2, to create more income variation. This makes the problem more challenging for subjects, and therefore gives a better chance of observing a range of conditions under which performance is very bad or surprisingly good. Permanent income grows according to $P_{t+1} = (1.05) P_t$ with P_1=100. The discount factor and gross interest rate are both set equal to one (Δ=1, R=1). The risk-aversion coefficient is $\rho=3$, an estimate often used in consumption studies which seems to fit many types of aggregate data. For

---

[5] That is, $x_t = X_t / P_t$, $c_t = C_t / P_t$, $h_{t-1} = H_{t-1} / P_{t-1}$, $s_t = S_t / P_t$ and $\varepsilon_t = \hat{\varepsilon} / P_t$.

habit formation, $\gamma = 0.6$, depreciation $\lambda = 0.7$, and the initial habit is $H_0 = 10$. There are thirty periods in this experiment so $T = 30$.

Intuitively, an optimizing consumer has to spend very little and save a lot in early periods for two reasons: To build up a buffer stock to protect against bad future income draws (low values of η_t); and to limit the "internality" of current spending on future utility, through habit formation. Figure I below illustrates an optimal path of consumption, and cash-on-hand, given a particular lifecycle of income shocks. Savings is the gap between the black optimal consumption line and the gray cash-on-hand line. In this example, the optimal consumer should only spend more than current income in periods 6-7 (when income is unusually low). The optimal cash-on-hand steadily rises to 1500 in period 20, which is about six times the annual income at that point. Generally, consumers should brace themselves for a rainy day by saving until about period 20. In later periods, they should start to dissave by spending more than their current income and dipping into their cash-on-hand (i.e., the consumption line is usually *above* the dotted income line after period 20).

2. Quasi-Hyperbolic Discounting

Exponential discounting is dynamically consistent—the current tradeoff between two future points which is reflected in current decisions is preserved when those future points eventually arrive. Since Strotz [1956], the possibility that intertemporal preferences change to reflect a dynamic inconsistency has been on the back of economists' minds. Now that possibility is a lively topic of research [see Laibson, 1997; O'Donoghue and Rabin, 1999; Barro, 1999, etc]. We discuss two approaches— β-δ quasi-hyperbolic discounting, and a dual-self model in which a foresightful "planner" tries to restrain a myopic "doer" from spending too much.

In this section we show how these models work and contrast them with the exponential model. These models are also of special interest in study 2, which uses beverage rewards rather than money (as in study 1). In study 2, in one lifecycle thirsty subjects earn the same number of utility points as in the study with money, but x points are converted into x/2 milliliters of cola. In the immediate condition they drink the cola right away. In the delayed condition they "order in advance", so that spending decisions in period t determine the amount of cola that can be drunk in period t+10. Dynamically consistent subjects should make the same decisions in these immediate and delayed conditions. However, under β-δ discounting or dual-self models, subjects may "spend" more (i.e., earn more points which are converted to beverage) in the immediate condition than in the delayed condition.

For notational simplicity define ω,

$$\omega\left(\widetilde{C}_t, \widetilde{H}_{t-1}\right) = v\left(\frac{u(\widetilde{C}_t, \widetilde{H}_{t-1})}{2}\right) \tag{7}$$

where $\omega\left(\widetilde{C}_t, \widetilde{H}_{t-1}\right)$ is the beverage reward (in terms of a beverage utility function *v*) associated with consumption decision $\widetilde{C}_t$ and habit $\widetilde{H}_{t-1}$. In the quasi-hyperbolic discounting (or present-bias) model, current utilities have a weight of one, and utilities t periods in the future (t>1) have a weight of $\widetilde{\beta}\ \widetilde{\delta}^{t}$.

In the β-δ model, the implicit tradeoff between future periods is <u>not</u> necessarily the same as the tradeoff which is made when the future arrives. Therefore, the model requires a behavioral assumption about whether current agents are "naïve" or "sophisticated" about their own future behavior.

Naïve agents believe—incorrectly—that the weights they currently apply to future periods are the same as the (relative) weights they will apply when the future arrives. Intuitively, even though all future periods are discounted by a present bias β, the naïve discounter believes that in future evaluations there will no such present bias. This model corresponds to chronic procrastination which is justified by the hope that starting tomorrow, the activity that has been put off for so long will finally get done.

Sophisticated discounters have a present bias, but correctly realize that they will have a present bias in the future, too. A crucial difference is that sophisticated discounters will seek external commitment devices (to restrain the present bias they know they will have) while naïve discounters do not. Our view is that it is too early in the empirical literature to consider only one model, when both can be considered and compared, so we develop both here and calibrate them on the experimental data in section 4.

Sophistication: Optimal consumption can be determined by backward induction because we have a finite number of periods T. In the last period the subject will solve

(8) $$\ddot{V}_T\left(X_T, H_{T-1}, P_T\right) = \omega\left(C_T^{'}, \tilde{H}_{T-1}\right)$$

where

(9) $$C_T^{'} = \arg\max_{C_t}\left\{\omega\left(C_T, \tilde{H}_{T-1}\right)\right\}.$$

Assuming sophistication, the optimization problem can then be solved recursively using equations 10 and 11.

(10) $$C_t^{'} = \arg\max_{C_t}\left\{\omega\left(C_t, \tilde{H}_{t-1}\right) + \left(\tilde{\beta}\tilde{\delta}\right)E_t\left[\ddot{V}_{t+1}\left(GP_t, X_t - C_t + \eta_{t+1}GP_t, \lambda H_{t-1} + C_t\right)\right]\right\}.$$

(11) $$\ddot{V}_t\left(P_t, X_t, H_{r-1}\right) = \omega\left(C_t^{'}, \tilde{H}_{t-1}\right) + \tilde{\delta}E_t\left[\ddot{V}_{t+1}\left(GP_t, X_t - C_t^{'} + \eta_{t+1}GP_t, \lambda H_{t-1} + C_t^{'}\right)\right].$$

Notice that the function $\ddot{V}_t$ is different than a typical dynamic programming value function as in equation (3). In equation (3), the value function

$V_t(P_t, X_t, H_{t-1}) = \max_{C_t}\{u(C_t, H_{t-1}) + \Delta E_t[V_{t+1}(P_{t+1}, R[X_t - C_t] + \eta_{t+1}G_{t+1}P_t, \lambda H_{t-1} + C_t)]\}$ is the maximum of the current consumption utility plus the (discounted) continuation value function conditioned on that level of consumption. But the possibility of dynamic inconsistency requires us to create a pseudo-value function $\ddot{V}_t$ instead. Equation (10) dictates that a sophisticated agent will maximize utility consistent with her present preferences. However, the sophisticated hyperbolic knows that in the future she will not apply the same weights and make the same tradeoffs, so she needs a way to keep track of consumption utilities in future periods without aggregating them into a conventional value function. Here, $\ddot{V}_t$ is a pseudo-value function which is simply a sum of future utilities from consumption, discounted at the exponential rate $\tilde{\delta}$.

Naiveté: A naïve agent believes that her future decisions will be made as if she is an exponential discounter.[6] A naïve agent therefore creates a value function $\dddot{V}_t(P_{t,} X_t, H_{t-1})$ which exhibits present bias, but uses the exponential value function $\dddot{V}_t(P_{t,} X_t, H_{t-1})$ (a modified version of equation (3)) with $\Delta = \tilde{\delta}$ and $u = \omega$ in forecasting future utilities.

$$\begin{aligned} &C_t^*(P_t, X_t, H_{t-1}) = \\ &\arg\max_{C_t}\{\omega(C_t, H_{t-1}) + \tilde{\beta}\tilde{\delta}E_t[\dddot{V}_{t+1}(GP_t, X_t - C_t + \eta_{t+1}GP_t, \lambda H_{t-1} + C_t)]\} \end{aligned} \tag{10'}$$

where

$$\dddot{V}_t(P_t, X_t, H_{t-1}) = \max\{\omega(C_t, H_{t-1}) + \tilde{\delta}E_t[\dddot{V}_{t+1}(GP_t, X_t - C_t + \eta_{t+1}GP_t, \lambda H_{t-1} + C_t)]\}. \tag{11'}$$

---

[6] The reason is that her current weights on all future periods (for t>1), $\tilde{\beta}\,\tilde{\delta}^{t}$, imply relative tradeoffs in future periods in which the $\tilde{\beta}$ terms divide out for optimization.

Figure II shows an example consumption path which compares sophisticated and naïve hyperbolic consumption paths for $\delta=0.9$ and $\beta=0.8$, compared to the optimal path (with $\delta=\beta=1$) from Figure I.

As O'Donoghue and Rabin [1999] have stressed, present bias and sophistication can interact in interesting ways. Generally, a naïve person exhibits more present bias than a sophisticated. However, a sophisticated person who is sufficiently present-biased can succumb to temptation immediately because she knows her future self will too, while a naïve person might postpone temptation because she thinks in the future she will be more patient than she currently is. In the model, as Figure II indicates both sophisticated and naïve hyperbolic discounting cause an individual to overconsume relative to optimal in this experiment. The difference between the two paths is very small [as in Angeletos et al, 2002], but the naïve consumer does consume a little more than the sophisticated one in early periods.

Immediate and delayed beverage rewards: To compare the immediate and delayed beverage reward conditions, we first assume that the utility of beverage is linear in volume. Then $u(C_t, H_{t-1}) = \tilde{v}\left(\frac{u(C_t, H_{t-1})}{2}\right) = \tilde{\omega}(C_t, H_{t-1})$. For simplicity, we also assume that subjects do not satiate in beverage, and utilities are additively separable across periods. Even if these assumptions do not hold, there is no reason to think that they are violated more or less in the two conditions (immediate and delayed).

In the delayed condition subjects do not receive the beverage amount they decided upon in period t until period t+10. Since there is no immediate reward, all future consumption has a weight of $\widetilde{\beta}$ (along with discount factors $\widetilde{\delta}$) and the $\widetilde{\beta}$ terms divide out in optimization.[7] Then each subject will solve:

$$(3') \qquad V_t^{'}(P_t, X_t, H_{t-1}) = \max{}_{C_t}\left\{\widetilde{\omega}(C_t, H_{t-1}) + \widetilde{\delta}E_t\left[V_{t+1}(GP_t, X_t - C_t + \eta_{t+1}GP_t, \lambda H_{t-1} + C_t)\right]\right\}.$$

subject to constraints 4, 5, and 6. Notice that $V_t^{'}(P_t, X_t, H_{t-1}) = \frac{1}{2}V_t(P_t, X_t, H_{t-1})$ (from equation 3) if $\Delta = \widetilde{\delta}$. In this case both value functions will have the same optimal consumption path.

If we assume $\widetilde{\beta} < 1$ (present bias), subjects in the immediate condition will weigh payoffs in the earlier periods more heavily than subjects in the delayed condition weigh them, and will choose to consume more beverage in early periods; this spending over the total beverage-maximizing level will be greater in the earlier periods of the immediate condition. As a result, $\widetilde{\beta} < 1$ predicts that subjects in the immediate condition will receive a lower beverage total because they are deviating from the optimal beverage maximizing total.

3. A dual-self planner-doer model

Another way to model dynamic inconsistency is by positing two systems, or a “dual self”, which interact to create behavior [see Thaler and Shefrin, 1988; Bernheim and Rangel, 2005; Loewenstein and O’Donoghue, 1995]. For brevity, we focus on just one of these models, the Fudenberg-Levine [2005] approach. They assume long-run and short-run players, much as in Shefrin and Thaler’s earlier “planner-doer” model. For consistency with the hyperbolic discounting model, assume that the long-run player (L) has a linear discount factor for beverage, $\widetilde{\delta}$. L also knows that the short-run player (S) will spend all resources in a given period if L does

---

[7] See appendix 3 for this reasoning. The crucial assumption is that subjects discount rewards based on when they will actually receive them.

not exercise self-control. In equilibrium L will choose a strategy from histories $m \in M$ and states to actions $\sigma_{SC} : M \times Y \to A$ to maximize the following reduced form objective function:

(12) $$U_t = \sum_{s=t}^{T} \tilde{\delta}^{t-s} \int [\omega(y,0,a) - q(y,a)] d\pi_t(y(m)).$$

In our application the only actions are spending decisions, so replace a with $\tilde{C}_t$, and define the current state $y = (P_t, X_t, H_{t-1})$. The function $q(y,a)$ is self-control cost of the long-run player enforcing spending $\tilde{C}_t$ in state *y*. The function π_t is the measure associated with histories of previous short-run actions and a given state. Histories are irrelevant to the long-run player (except as summarized by the state variables); only the probabilities of other states are relevant.[8] Thus equation (12) can be rewritten as (13),

(13) $$U_t = \sum_{s=t}^{T} \tilde{\delta}^{t-s} \int \left[\omega\left(P_t, X_t, H_{t-1}, 0, \tilde{C}_t\right) - q\left(P_t, X_t, H_{t-1}, \tilde{C}_t\right)\right] d\mu\left(P_{t-1}, X_{t-1}, H_{t-2}\right)$$

where $d\mu(P_{t-1}, X_{t-1}, H_{t-2}))$ is the probability measure of states. If the subject is an expected utility maximizer, than her problem becomes very similar to the hyperbolic case. She will solve

(14) $$\max\left\{E_t\left[\sum_{s=t}^{T} \tilde{\delta}^{t-s}\left[\omega\left(P_t, X_t, H_{t-1}, 0, \tilde{C}_t\right) - q\left(P_t, X_t, H_{t-1}, \tilde{C}_t\right)\right]\right]\right\}.$$

In the delayed condition the short-run player has no control over how much utility she receives because the decision determining her current utility was made ten periods ago. According to assumption 4 of Fudenberg and Levine [2005] this means L can make all choices without exerting a self-control cost. So her optimization procedure will be identical to the traditional exponential-discounting case. L will maximize

---

[8] Fudenberg and Levine also ignore histories when modeling a simpler version of the consumption/savings problem.

(1")
$$\max\left\{ E_t\left[\sum_{s=t}^{T} \widetilde{\delta}^{\,s-t} \omega(\widetilde{C}_S, \widetilde{H}_{S-1}) \right]\right\}.$$

The optimal value function will be

(3")
$$V_t(P_t, X_t, H_{t-1}) = \max\left\{ \omega(C_t, H_{t-1}) + \widetilde{\delta} E_t\left[V_{t+1}(GP_t, X_t - C_t + \eta_{t+1} GP_t, \lambda H_{t-1} + C_t) \right]\right\}$$

subject to constraints 4, 5 and 6 as before.

In the immediate condition, if there are positive self-control costs in restraining S's spending in each period, the subject will consume more than is optimal (since S's myopic ideal is to consume everything).

The implication is that subjects in the immediate condition will spend more than is optimal in early periods, and will therefore spend more than the delayed-condition subjects (who, by assumption, optimize).

Hence the planner-doer and hyperbolic discount models both predict more early consumption, and less overall consumption, in the immediate condition compared to the delayed condition. To make a more precise prediction (and comparison between theories) requires a detailed specification of the utility costs of self-control, which is an important topic that lies beyond the scope of this paper.

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 Appendix 1 for details and instructional tables). To ease comprehension and avoid demand effects, economic jargon like 'income shocks', 'habit stock', and 'utility', were translated into plainer language-- 'adjustment factor', 'lifestyle index' and 'points' respectively.

The instructions explained that the adjustment factors (income shocks) followed a lognormal distribution, showed that distribution graphically, and gave 30-draw samples from that distribution to give participants a feel for how much their income could vary. 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. To ensure that participants understood the instructions, they were required to complete and correctly answer a quiz *before* they started their experiment. The quiz tested 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 arise from confusion.

Decisions were input to an Excel interface which displays the income obtained, the corresponding cash available, and the habit stock (see Figure III). 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. Therefore, 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 or '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 [9] plus a $5 showup fee.

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 they were given samples of what three actual subjects had done in the private learning condition.[10] The three samples were taken from the highest-earning subject, the lowest-earning subject, and from one subject chosen at random (and the social learning subjects knew how these three samples were chosen).

Of course, there are many other ways to implement social learning or imitation (e.g., Ballinger et al, in press, 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 an average Joe. 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 may overspend relative to the optimum; so it is not clear whether social learning will help or hurt, or have no effect.

Participants were 35 undergraduates from the National University of Singapore (NUS) and 37 undergraduates from California Institute of Technology[11]. These students are unusually adept at analytical thinking so they should represent an upper bound on how well average consumers can solve these intertemporal optimization problems in this experimental setting. The

---

[9] 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).

[10] 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.

[11] Experiments were conducted on 30/09/02 and 01/10/02 (NUS) and on 09/01/03 for Caltech students.

participants were recruited using the universities' mail servers. Half the participants (18 from each school) did the experiment with private learning and half (17 NUS, 19 Caltech) 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[12]). Most participants completed the instruction and seven lifecycles in about 90 minutes.

2. Basic results

Table II gives summary statistics of actual point outcomes in the two learning conditions. The first and second rows give the average of total lifecycle points across the 36 subjects 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.[13] The fifth row is the total income in each lifecycle (which gives an idea of whether deviations from optimality are due to bad decisions or 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 4 the average subject earns point totals within 80% of the optimum in this very difficult problem, and the variability across subjects shrinks.

Table II also shows (bottom panel) 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.

---

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

[13] 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.

3. Behavior relative to conditional optimization

The Table II 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 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 II but may be close to conditionally optimal overall.

For each participant, the average conditional deviation for each period is the difference between their actual spending and the optimum (conditioned on that participant's earlier decisions). Figure IV below 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 IV is the horizontal axis, the reader can judge at a glance whether deviations are significant by seeing whether the interval covers the zero line.

Figure IV confirms the conclusion from Table II: With only private learning, participants in lifecycle 1 are spending significantly more than optimal in early periods, until about period 20. However, the lifecycle 7 conditional deviations are never significantly different from zero, so learning is very effective over the seven lifecycles. (The actual spending path is insignificantly different from the conditional optima by lifecycle 4.)

Figure V shows the analogous data for the social learning condition. Deviations in lifecycle 1 are much smaller than the corresponding deviations from private learning, and are insignificantly different from zero in most periods. There is 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).[14]

Table III 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 (5 and 7). There is no significant effect of ethnicity and a small effect of gender (women deviate about 20% more).

IV. Study 2: Temptation and beverage rewards

1. Experimental design

Study 2 was identical to the first study with one large change and some small changes. The large change is that in one lifecycle, subjects received a beverage reward rather than money. Lifecycles 1-2 and 4-5 (with money rewards) were the same as in study 1. In lifecycle 3 subjects received a fixed monetary payment for their participation but did not earn any money. Instead, in each period they drank a sip of a beverage[15] based upon their spending decisions each period (1

---

[14] See Chua and Camerer [2004 URL] 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.

[15] 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, we did not think water was as motivating as

ml beverage/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.

To make this reward appealing, subjects were asked upon recruitment not to drink for four hours before the experiment began.[16] They also begin by eating some salty snacks, and 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. They likely do not satiate during the lifecycle. No subject received over 350 mL (< 12 oz, a can of coke) of beverage in that lifecycle.[17] The amounts are small enough within a period (a maximum of 20 ml/period (.7 oz)) that they are unlikely to satiate within a lifecycle. To induce separability and reduce satiation across periods, there was a one minute delay between subjects' decisions each period. Subjects were required to drink their entire beverage allotment in that one minute period.

---

colas, and because pilot subjects (including the middle coauthor) thought other beverages we tried (juices) were too filling and might induce satiation which complicates the analysis.

[16] We cannot verify whether they obeyed our request to show up thirsty. However, because assignment to the immediate and delayed conditions did not depend upon apparent thirst, any differences in pre-experiment thirst are just sources of sampling error in comparing the two groups.

[17] 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.

[18] Given the carbonation of some beverages, the amount of beverage received may vary from the number sent to the syringe pump. This occurs as carbon dioxide fills the syringe. Syringes were tilted downward to give subjects the most beverage possible.

A syringe pump with three syringes was used to deliver an exact[18] amount of beverage into a cup.[19] 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.

There were two different conditions in the beverage lifecycle in study 2. 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.[20] As shown in section 3 above, the quasi-hyperbolic or present-bias model, and a reasonable interpretation of the Fudenberg-Levine dual-self model, both predict that subjects will drink more beverage in the early periods of the immediate condition, and consume less total beverage.

Subjects were n=52 Caltech students.[21] Because a liquid-delivery apparatus was used, experiments were conducted in a single office rather than a computer lab.

---

[19] See Appendix 2 for a diagram of the beverage delivery apparatus.

[20] 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.

[21] 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.

[22] As we explained earlier, subjects can produce extremely negative point totals but those values cannot be reflected in beverage loss.

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 too much in early periods. This prediction is correct (see Table IV, row 1). The immediate-condition subjects drank less total beverage on average (179 ml, sd=84.6) than the delayed-condition subjects (226 ml, sd=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 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 that could be evidenced by differential performance in the four money lifecycles. To control for these skill differences, we estimate the regression

$$P_{it} = a + b_1 r_1 + b_2 r_2 + b_3 r_3 + b_4 r_4 + b_5 I + e_{it} \, . \qquad (15)$$

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.

Notice that point totals can be negative for the beverage lifecycle and total ml of beverage consumed cannot. As an explanatory example, consider a subject that produces a total of 300 points in periods 1-29, and -3000 points in period 30. The total points this subject would receive for the beverage lifecycle is -2700 but her total beverage received is 300/2=150 ml. Hence, there is a disparity between points and liquid received. Subjects are aware of this difference, so if a large deficit (>350) occurs in an earlier period, they will know that there is no amount of spending that will allow them to overcome that deficit. As a result, when subjects have large negative point totals they can become indifferent about future decisions (their marginal incentive is low). Indeed, some subjects did choose to produce very negative point totals when they got behind on points. The high deviations occurred disproportionately in the immediate condition, which then greatly overstates b_5 (see Table V). 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 their logarithms with their sign preserved (i.e., the dependent variable is $[|P_{it}|/P_{it}]\ln(P_{it})$).

Table V 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 (the weakest result is the beverage-conversion model, with p=.026 one-tailed). Since these results are stronger than the parametric t-test reported in Table IV, accounting for individual differences in skill by using the money-lifecycle results slightly enhances the significance of the immediate-delayed condition.

These analyses use the overall point totals in the lifecycle. It is also useful to examine conditional deviations. 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 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 0ml, we bounded all totals at 350 ml.[22]

Row 2 of Table IV shows the results. The average total expected beverage loss is much higher for the immediate condition, about twice as high as in the delayed condition.

2.3 Exploring the time series of overspending in early periods

We can compare each subjects' decisions to the optima conditioned on the values of the savings and habit state variables which a subject generates throughout the first 10 periods. Figure VI shows the average ratios of spending to conditionally optimal spending. (The optimal line is now just a flat line at a ratio of 1.) Figure VI confirms that even when conditioning on past decisions, the immediate-condition subjects are spending more in the first five periods.

We calculated the average expected loss for each overspending decision for each subject (Table IV, row 3) to explore whether the greater overspending by subjects in the immediate condition is responsible for their lower beverage reward totals.[23] The immediate group lost significantly more per overspending decision than the delayed group. Interestingly, the immediate group subjects made fewer overspending decisions than the delayed-condition subjects (41% to 51% of their decisions). But when they did overspend, the immediate immediate condition subjects spent much more than was optimal, leading to greater expected losses.

2.4 Calibrating hyperbolic discounting

The results presented to this point have supported the basic prediction of both hyperbolic discount and dual self models-- 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 $\tilde{\beta}$ and compare them with estimates from other studies.

---

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

[24] We thank Daniel Houser for this point.

The analysis is restricted to the first 10 periods. These periods are the most important in determining the spending path and give the most identification about $\tilde{\delta}$ and $\tilde{\beta}$. Periods 1-10 have the fewest deficits, which inhibit subjects from receiving an immediate reward and probably lower temptation and identification of $\tilde{\beta}$, and have the most similar state variables among subjects (since they all begin with the same initial state variables). Thus we are observing subject data when they are making decisions in the most similar circumstances.[24]

Since the discount factors on future periods multiply $\tilde{\beta}$ by a power of $\tilde{\delta}$, if $\tilde{\delta}$ is close to one, it is difficult to identify the two parameters separately by searching for a pair of best-fitting values (i.e., there are many pairs with low $\tilde{\beta}$ and high $\tilde{\delta}$ that fit about equally well). We therefore use a two-stage procedure to calibrate $\tilde{\delta}$ and $\tilde{\beta}$ for each subject Since behavior in the delayed condition gives no information about the present bias $\tilde{\beta}$, in theory, those data will be used to estimate $\tilde{\delta}$. Then we use the immediate-condition data, constraining $\tilde{\delta}$ for all subjects to equal the mean value derived from the delayed-condition, to estimate each subject's $\tilde{\beta}$ given the common $\tilde{\delta}$.

The objective function is the sum of squared deviations between the actual spending decisions and the conditionally optimal ones (given $\tilde{\beta}$ and $\tilde{\delta}$) for periods 1-10. The mean $\tilde{\delta}$ is estimated at a reasonable 0.904 (see Table VI). We then calculated spending paths for $\tilde{\delta}=0.904$ and various $\tilde{\beta}$ values and fit them to subjects in the immediate condition. Table VI shows the mean and standard deviation of $\tilde{\beta}$ for both sophisticated and naïve specifications. In the sophisticated model, the mean best-fitting $\tilde{\beta}$ is 0.758; in the naïve model the mean best-fitting $\tilde{\beta}$ is 0.865. Both values 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] ($\tilde{\beta}$=.74-.89) (from macroeconomic calibration, unemployment spells, and experiments in Vietnam, respectively). All values are significantly less than one at the 5% level, using a t-test with each subject's best-fitting estimate as a single datum.

In summed squared deviations, sophisticated hyperbolic discounting fits slightly better than the naïve model, and both models improve the fit of a no-discounting model by about 10% (see Table VI).

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]. 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).[25] Figure VII tests this hypothesis using study 1 data, by plotting nominal utility losses in each period from actual

---

[25] 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.

[26] 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.

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 VII 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, in press],[26] A piecewise-linear jackknife regression through the origin gives coefficients in the domain of positive and negative conditionally optimal utilities of .79 and .15, respectively. The ratio of these two slopes is 5.2.

Figure VIII shows a graph of the frequency of small actual and optimal utilities (between -10 and +10), across all subjects and periods. In the actual period-by-period utilities there is a huge spike on the slight positive side (which is not at all evident in the corresponding distribution of optimal utilities), indicating a strong preference for a small positive utility and aversion to loss.

Figures IX-X show actual-optimal utility scatter plots like Figure VII, for the money periods and beverage periods of study 2. For money (Figure IX), the jackknife regression gives positive and negative slopes of .88 and .10 (a ratio of 8.8). For beverage (Figure X), the slopes are .92 and .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 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 these Figures VII-X 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. 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].

VI. 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 many other unobservable variables. Experiments control for these assumptions, and generally show that experimental subjects save too little.

This paper seeks to establish conditions under which intelligent subjects at the start of their economic lifecycles make nearly-optimal or highly suboptimal decisions in a complex savings problem. The 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, with social learning from examples of successful, unsuccessful, and average their savings is quite close to optimal..

The fact that subjects can learn to save optimally for money rewards, from their own experience or examples of others, 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. These subjects who earn immediate rewards also overspend, compared to group of subjects who make decisions in one period but do not get to sip that period's beverage 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 optimum.

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 .758 for the sophisticated case and 0.865 for the naïve case (both immediacy parameters are significantly less than one with high significance, $p<0.05$). These values also close to values observed in some other studies using both calibration to aggregate data and direct experimental measurement.

In a sense, 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 experiment 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. [cf. Bernheim and Rangel, 2002].

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.

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. If young adults observe the savings decisions of a few elder family members, they may figure out how to solve these computationally difficult problems. Retirement advisors likely can solve these problems because they observe far more than one lifecycle as they make decisions for their clients. 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). The market may have solved the cognitive problem in savings models by having retirement advisors exist. The results of our learning study suggest that intelligent subjects are capable of approximating optimization, with enough private or social learning.

The unique feature of our study is adding visceral temptation, by rewarding subjects with beverages when they are quite thirsty. This design provides a model for studying highly tempting decisions like addiction, overeating, and perhaps spending splurges.

The fact that subjects in the delayed condition are able to resist temptation better (and drink more beverage as a result) corroborates the conclusion of models like Bernheim and Rangel's [2005], that creating a time wedge between "ordering" and consuming may be helpful to people. The next obvious step is to offer subjects a choice between the two conditions. Sophisticated subjects who know they benefit from external control will choose the deliberate delay. 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.

Appendix 1: Study 2 – Delayed Condition Instructions

INSTRUCTIONS (Version 2):

Before we begin, there are some rules in this experiment that are necessary for its validity. First we have asked you not to drink liquids for four hours before this experiment. We also will ask you not to drink during this experiment, except for the liquid rewards you receive. Additionally, you will be asked to consume a salty snack before this experiment begins. To ensure that you do not break these rules, if you need to leave the experimental room for a bathroom break, the experimenter or female assistant of the experimenter will monitor that you do not drink liquids during this time. If you feel at any time your health is at risk in this experiment, please tell the experimenter and the session will be stopped. You will receive all earnings up to that point and your show up payment. At the end of the experiment you will be asked to take a small questionnaire. At that time the experimenter will provide you with any beverages you may require.

What you need to know about this experiment.

In this experiment, we are interested in how you make your spending and saving decisions over a 30 period 'lifetime'. You will make these decisions for money and for a liquid reward. The instructions will explain how the computer interface works. It will also explain how the decisions you make determine the amount of money you will earn. The money for the experiment has been provided by a research foundation. If you follow the instructions, and think carefully before making your decisions, you can earn a considerable amount of money. This will be paid to you in cash at the end of the experiment.

You will play 5 sequences of a 30 period spending/saving game.

There will be 5 sequences of a 30 period game. In each sequence you will make 30 decisions in a row. The third sequence will be different from the other four. It will be for a liquid reward rather than cash. This round will be explained later. You will receive a fixed amount of money for participating in this round. The point totals from the other four sequences will be calculated to determine the total amount you will earn at the end of the experiment. The game will be played on a Microsoft Excel workbook. Table A below shows an example of what the **first** period of **one non-liquid reward** sequence might look like. An explanation of the liquid reward sequence will be given at the end of these instructions.

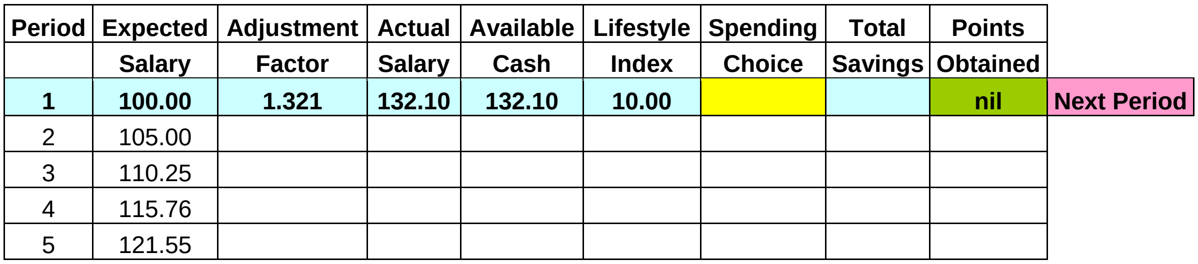

Period	Expected	Adjustment	Actual	Available	Lifestyle	Spending	Total	Points	
	Salary	Factor	Salary	Cash	Index	Choice	Savings	Obtained	
1	**100.00**	**1.321**	**132.10**	**132.10**	**10.00**			**nil**	**Next Period**
2	105.00								
3	110.25								
4	115.76								
5	121.55								

Table A (First period)

Appendix 1: Study 2 – Delayed Condition Instructions

In table A the row representing period 1 is highlighted because the computer is waiting for you to make a decision about what to do in period 1. After you make a decision in period 1 **(and click on the pink box labelled 'Next Period')**, the computer will record the decision and highlight the row for period 2. After you have made decisions for all 30 periods in a sequence, you will see your total point earnings for that sequence at the bottom right corner, and a pink box marked "Continue". When you click on the "Continue" box you will begin the **next** 30-period sequence.

Your total point earnings will be determined by a series of decisions about how much to **spend** from a sum of **available cash**. In each period you will have some cash available, which is the addition of what is left over from the previous period, and a new amount called "Actual Salary". The actual salary in each period is determined by multiplying two numbers—the **expected salary**, and a random **adjustment factor**. You will know the expected salary in advance (in fact, your computer screen will show the expected salary for all 30 periods in a sequence, and it is the same across all of the 5 sequences you will play). The **adjustment factor** will go up and down in each period in a way that we will explain about shortly.

Your expected salary grows by 5% each period.

The values of expected salary are shown for all thirty periods in the second column of the table. Expected salary increases at 5% each period. Therefore, if the first period's expected salary is 100 as shown in the table, then the second period's expected salary: 100 x 1.05 = 105.00. The third period's expected salary is: 105 x 1.05 = 110.25, and so forth for future periods.

Your expected salary is susceptible to adjustments.

The actual salary you get each period is determined by multiplying the expected salary by an a**djustment factor**. You will experience both good and bad adjustments to your expected salary, because the adjustment factor is often less than 1 (so that the actual salary is **less than** the expected salary), and the adjustment factor is often also greater than 1 (so that the actual salary is **more than** the expected salary). The exact adjustment factor in any one period is determined by a random draw from a probability distribution. The distribution is shown in the graph below, which may help you try to guess what adjustment factors are most likely to occur.

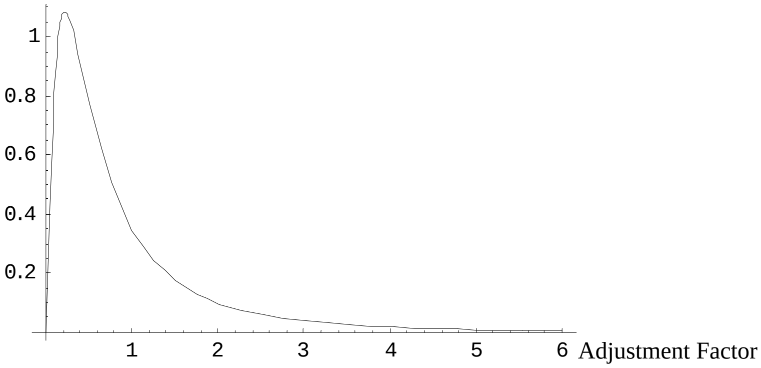

The x-axis of the graph shows the possible adjustment factors (where 0 is the lowest possible factor). The y-axis shows how likely it is that an adjustment factor on the x-axis will actually randomly occur. Notice that the **most common** adjustment factors are **less than one** (because the curve is very tall for values on the x-axis between 0 and 1); but some of the adjustment factors are very large. This means that **most** of the time, your actual salary will be below your expected salary, but sometimes your actual salary will be much larger than your expected salary. In fact, half the time the adjustment factor will be below 0.607 and half the time the adjustment factor will be above 0.607. About 10% of the time the adjustment factor will be very low, 0.168 or less, and about 10% of the time it will be very high, 2.185 or above. (In case you are curious, we can tell you that the statistical distribution of the adjustment factor is generated by taking a "normal" or "bell curve" distribution, then taking the mathematical constant "e" (which is roughly 2.718) raised to a power equal to the number drawn from the bell curve distribution.)

Note also that each adjustment factor is statistically independent of the factors in early periods. This means that whether the factor is particularly high or low does not depend on whether it was high or low in the previous periods.

Table B below shows three <u>example</u> sequences of 30 adjustment factors (Sequence A, B and C), randomly drawn using the above-described distribution.

	Sequence A		Sequence B		Sequence C
Period	Adjustment Factors		Adjustment Factors		Adjustment Factors
1	1.364		0.845		0.624
2	0.461		2.464		2.660
3	0.498		0.403		2.643
4	0.223		0.199		1.298
5	0.323		0.413		0.840
6	0.108		0.296		0.389
7	0.283		0.199		0.530
8	0.588		0.926		2.592
9	4.793		1.989		0.599
10	0.780		1.601		1.246
11	2.721		0.230		0.674
12	0.334		1.270		0.159
13	2.203		0.715		1.586
14	1.363		0.404		0.129
15	0.289		0.100		0.471
16	0.194		0.170		0.309
17	0.369		0.426		0.364
18	1.296		0.604		0.703
19	0.256		0.248		1.120
20	0.308		1.033		0.219
21	0.767		1.441		0.780
22	0.671		0.910		0.049
23	0.578		0.198		0.486
24	0.956		1.665		0.446

25	2.000		1.636		0.265
26	1.782		0.174		0.549
27	0.140		0.482		0.276
28	0.384		0.342		0.406
29	0.087		0.929		0.457
30	1.692		1.625		0.367

Table B

Please note that these sequences of adjustment factors are only examples of what a sequence of 30 adjustment factors might look like; the actual sequences of adjustment factors you will get in your experiment will be different, even though the underlying probability distribution from which it was drawn is the same.

Period	**Expected Salary**	**Adjustment Factor**	**Actual Salary**	**Available Cash**	**Lifestyle Index**	**Spending Choice**	**Total Savings**	**Points Obtained**	
1	100.00	1.321	132.10	132.10	10.00	60.00	72.10	38.49	
2	**105.00**	**0.345**	**36.23**	**108.33**	**67.00**			**nil**	**Next Period**
3	110.25								
4	115.76								
5	121.55								

Table C (Second Period)

<u>Actual Salary</u> = Expected Salary multiplied by Adjustment Factor.

Each period, the actual salary is equal to the expected salary in that period times the adjustment factor. For example, in table C, the actual salary in period 1 is given by: 100 x 1.321 = 132.10. A low adjustment factor in period 2 (0.345) means that the actual salary in that period is only 36.23, which is much lower than the expected salary of 105.00. Keep in mind that these adjustment factors are just examples. When you participate in the experiment and make your own decisions, the adjustment factors will probably be different.

<u>Available Cash</u> = Last Period's Savings + Current Period's Actual Salary.

Remember that the one decision you must make in each period is how much of your available cash to spend. In table C, suppose you decide to spend 60.00 in Period 1. The total savings for period 1 is then your available cash (equal to actual salary because there was no past savings before period 1) minus your spending choice, which is 132.10 – 60.00 = 72.10. Please note that you <u>do not</u> earn interest on savings.

In period 2 of the table above, your actual salary is 36.23. Therefore, your available cash for period 2 is your savings left over from period 1, which was 72.10, plus your actual salary in period 2, which is 36.23. The total is 72.10 + 36.23 = 108.33, which will be automatically calculated for you and shown in the available cash column.

Spending earns you points. Make your <u>Spending Choice</u> in the yellow box.

Enter your spending choice each period in the yellow box. For your spending decision, the corresponding **<u>Points Obtained</u>** will be shown in the green box.

In table D below, entering a spending decision of 60.00 in period 2 will show that you can get 25.18 points <u>for that period</u>. You can try out different levels of spending before you make your final decision, by entering different values in the yellow spending choice box. Every time you input a value and press "enter" the computer will calculate how many points obtained you would get from that spending choice.

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	100.00	1.321	132.10	132.10	10.00	60.00	72.10	38.49	
2	**105.00**	**0.345**	**36.23**	**108.33**	**67.00**	**60.00**	**48.33**	**25.18**	**Next Period**
3	110.25				106.90				
4	115.76								
5	121.55								

Table D

The number of points obtained also depends on your <u>Lifestyle Index</u>.

For each level of spending, the number of points you earn is dependent on your lifestyle index. At a higher lifestyle index, you will earn a smaller amount of points for a given level of spending than when you are at a lower lifestyle index. In table D for example, the lifestyle index is at a higher level in period 2 (67.00) than in period 1 (10.00), therefore, the same level of spending of 60.00 yields a lower level of points in period 2.

A point transformation table is placed on your desk. Table E below shows <u>part</u> of this point transformation table. At each level of lifestyle index, it displays the number of points you can get at different levels of spending.

		Lifestyle Index							
		10	20	50	100	150	200	250	300
	5	-60.2	-190.3	-651.5	-1548.7	-2544.4	-3609.9	-4730.7	-5897.4
	10	3.2	-44.7	-214.2	-544.0	-910.0	-1301.7	-1713.7	-2142.6
	20	28.5	13.5	-39.6	-142.8	-257.4	-380.0	-508.9	-643.2
	40	36.7	32.5	17.5	-11.7	-44.0	-78.7	-115.1	-153.1
	60	**38.5**	36.5	**29.6**	16.0	1.0	-15.0	-31.9	-49.5
	80	39.1	38.0	34.0	26.2	17.6	8.4	-1.4	-11.5
Spending	100	39.4	38.7	36.1	31.1	25.5	19.5	13.2	6.6
Choice	120	39.6	39.1	37.3	33.7	29.8	25.6	21.2	16.6
	140	39.7	39.3	38.0	35.4	32.5	29.4	26.1	22.7
	160	39.8	39.5	38.5	36.4	34.2	31.8	29.3	26.7
	180	39.8	39.6	38.8	37.2	35.4	33.5	31.5	29.5
	200	39.9	39.7	39.0	37.7	36.3	34.7	33.1	31.4
	220	39.9	39.7	39.2	38.1	36.9	35.6	34.3	32.9
	240	39.9	39.8	39.3	38.4	37.4	36.3	35.2	34.0
	260	39.9	39.8	39.4	38.6	37.8	36.9	35.9	34.9

Table E (Point Transformation Table)

As you can see, a lifestyle index of 10 and a spending choice of 60 gives you 38.5 points. However, if the lifestyle index is 50 and you spend the same level of 60, the points obtained will be at a lower level of 29.6.

You are advised to look up this table before you make your spending choice. Alternatively, you can find out how many points you can earn by inputting different values of spending in the yellow box.
The lifestyle index grows with spending.

In general, the lifestyle index for a period is calculated by taking the value of the index from the previous period times .70, and adding in the previous period's spending. For example, in table D, the lifestyle index for period 2 is calculated as shown:
0.7 * 10.00 *(1st Period Lifestyle Index)* + 60.00 *(1st Period Spending)* = 67.00.
Likewise, if spending is again 60.00 in the second period, the lifestyle index for period 3 is: 0.7 * 67.00 *(2nd Period Lifestyle Index)* + 60.00 *(2nd Period Spending)* = 106.90
When you enter a spending level each period, the lifestyle index for the **next** period will be automatically calculated and shown.

A lifestyle conversion table is also provided on your desk. It shows you how your lifestyle index in the next period is dependent on how much you spend in the current period. Table F below shows part of this lifestyle conversion table.

		Lifestyle Index, Current Period						
		10	**20**	**50**	**100**	**150**	**200**	**250**
	10	17	24	45	80	115	150	185
	20	27	34	55	90	125	160	195
	40	47	54	75	110	145	180	215
	60	**67**	74	95	130	165	200	235
	80	87	94	115	150	185	220	255
	100	107	114	135	170	205	240	275
Spending Level, Current Period	**120**	127	134	155	190	225	260	295
	140	**147**	154	175	210	245	280	315
	160	167	174	195	230	265	300	335
	180	187	194	215	250	285	320	355
	200	207	214	235	270	305	340	375
	220	227	234	255	290	325	360	395
	240	247	254	275	310	345	380	415
	260	267	274	295	330	365	400	435
	280	287	294	315	350	385	420	455
	300	307	314	335	370	405	440	475
	320	327	334	355	390	425	460	495

Table F (Lifestyle Conversion Table)

As you can see, choosing a spending level of 60, when lifestyle index is 10, will result in a lifestyle index in the **next period** of 67. If you decide to spend more, for example 140, then your lifestyle index for the **next period** will be at a higher level of 147.

Note that the more you spend in the current period the higher your lifestyle index will be in future periods. The **point transformation table** (table E) shows that for any particular level of spending, you earn fewer points if the lifestyle index is higher. So if you spend a lot in early periods, you will receive many points in those periods, but you also increase the lifestyle index for future periods, which will then reduce the points you obtain in future periods.

You cannot spend more than your available cash.

Each period, you are not able to spend more than the available cash you have. If you choose a spending level greater than the cash you have, the program will tell you to lower your spending.

Proceed to the next period when you have made your spending choice.

Once you have thought carefully about how much to spend each period, proceed to the next period by using your mouse to **click once** on the pink box labelled 'Next Period'. **Please note that the program prevents you from returning to earlier periods to change your spending choice.** Therefore, please be careful not to click the 'Next Period' box before you enter your spending decision, because you will not be able to return to change it.

Once you have completed each 30 period sequence, proceed to the next sequence of 30 periods by clicking the 'Continue' link, which will appear at the bottom right of your screen.

Please note that the sequence of adjustment factors will be **different** in each of the 5 sequences, but the overall statistical distribution of possible adjustment factors will be the same. Once you have completed all 5 sequences, a screen will appear to tell you your overall points obtained from all 5 sequences.

The computer will automatically spend <u>all</u> available cash in the <u>last</u> period of each sequence.

Available cash from one sequence will <u>not</u> be carried over to the next sequence. This means that the computer will be automatically spend <u>all</u> remaining available cash in period 30 of each sequence.

How your earnings are determined:

After you make your spending choice each period, the points you obtain that period, in addition to all points you obtain in previous periods will be tallied at the bottom of the screen. Some of the point outcomes each period will be negative but your total points from each sequence should be positive.

The total points you obtained from your four non liquid reward rounds will be calculated and will be converted to cash at a rate:
14 points = $0.03
466.67 points = $1
Your earnings from the money rounds, in addition to the **$5.00** show-up payment and a **$25.00** fixed payment for the liquid round, will be paid to you in cash when you leave the laboratory. They will be rounded up to the nearest dollar.

For the third sequence you will make decisions for a liquid reward.

The third sequence will be identical to the other sequences except your point totals will be converted to mL of liquid and dispensed in a cup to your right. Each mL of liquid will be equal to two points. Table G shows the first period of the liquid reward round:

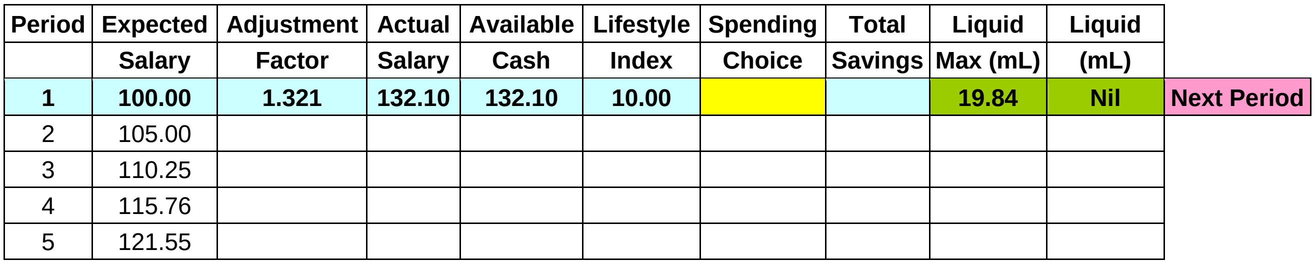

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Liquid Max (mL)	Liquid (mL)	
1	**100.00**	**1.321**	**132.10**	**132.10**	**10.00**			**19.84**	**Nil**	**Next Period**
2	105.00									
3	110.25									
4	115.76									
5	121.55									

The first eight columns are identical to those in the monetary reward rounds, and were explained in the instructions previously. ‘Liquid Max (mL)’ is the maximum amount of liquid that can be delivered at the end of the period. To achieve this amount you must spend all your available cash. Liquid (mL) is the actual amount of liquid that will be delivered to you at the end of the period. It is equal to half of ‘Total Points’ from previous rounds.

After you have made your spending choice and clicked ‘Next Period’ the program will ask you to close Microsoft Excel. You then will have a 60 second break, before the program opens. Beginning in period 11 you will begin to receive a liquid reward in your cup equivalent to your liquid reward decision from ten periods ago. This means if you decided on 10mL of liquid reward in period 5, you will receive it after you make your decision in period 15. Once you receive a reward, you must consume all of it during this break. You cannot save it for future periods.

After you have made your decision for period 30 the program will close. You will receive your rewards for periods 21-30 over the next ten minutes. In those ten minutes you will not need to interact with the computer.

If you sustain a negative result for any round, you will not receive liquid until you have produced enough positive periods to offset that result. For example, if you sustain a negative liquid mL total of -20 in period 8, 10 in period 9, and 15 in period 10, you will receive no liquid reward in period 8 or 9 but 5 mL of liquid reward in period 10.

Please note that if the liquid reward is a carbonated beverage, the volume of liquid may be slightly different than the value in the excel spreadsheet. This is due to the carbonation gas being measured as liquid in the syringe pump. It is unavoidable when using carbonated beverages.

Here is a brief summary of what you need to know.

You will be making decisions in 5 sequences of 30 periods. In each period you will have some available cash and will choose a level of spending. In the third sequence your decisions will determine the liquid rewards you receive over that round. All of the other sequences are important in determining your overall cash earnings, because your earnings will depend on the point total of sequences 1, 2, 4 and 5.

Expected salary grows at 5% each period. The actual salary that you get depends on a random adjustment factor that occurs during each period. These factors are randomly determined and the adjustment factor in one period does not depend on whether the previous period's adjustment factor was high or low. The available cash you have during each period is the actual salary you get in the current period plus the level of savings that was left over from the previous period.

The level of points (mL of liquid reward) you can get during each period depends on the level of spending you make, as well as your lifestyle index. More spending this period increases the lifestyle index for next period. A higher level of lifestyle requires a higher level of spending than before to obtain the same level of points. The point transformation table on your desk will give you a better idea on how this works.

Take as much time as you like to make your spending decision in each period. Please note that your spending level in each period cannot exceed the available cash you have. Remember that you cannot go back to earlier periods to change your spending level once you have clicked on the 'Next Period' box. Therefore, please make sure that you have correctly entered your final spending decision in the yellow spending choice box before proceeding to the next period.

The total points you have obtained for all four non-liquid reward sequences will be calculated and converted to cash. In the liquid reward round your decisions will determine how many mL of liquid you will receive between periods. Remember that the liquid reward round is very similar to the other rounds. The mL of liquid you receive at the end of each period is equivalent to half your period's points total in any other round.

If these instructions were not clear to you, or you have a question of any sort, please tell the experimenter now.

If you don't have any questions, please attempt the short quiz on the following page before you start the experiment. These questions will test whether you have fully understood the instructions. Once you are done with the questions, the experimenter will come by to check your answers. **If your answers are not right, the experimenter will give the correct answer and help you understand how the tables and instructions should enable you to give the correct answers.**

You can only start the experiment when all your answers are correct.

Quiz

Answer questions 1-7 for non-liquid reward sequences. (Sequences 1,2,4,5).

1) If you spend 60.00 this period, and your lifestyle index is 50.00, how many points will you obtain?

Ans: ________________

2) If you spend 80.00 this period, and your lifestyle index is 250, how many points will you obtain?

Ans: ________________

3) If you spend 450.00 this period, and your lifestyle index is 700, how many points will you obtain?

Ans: _________________

4) If you increase your spending level from 60.00 to 100.00, and your lifestyle index is 100.00, how many additional points will you get?

Ans: ________________

5) Your expected salary in period 2 is 150.00. The adjustment factor is 0.500 in the same period. Total savings from period 1 was 40.00. How much available cash do you have in period 2?

Ans: ________________

6) Your lifestyle index is 50 in period 1. If you decide to spend 60.00 in the same period, what would be the level of lifestyle index in period 2?

Ans: ________________

7) In period 20, your lifestyle index is 200.00. You decide to spend 120.00.
   a) How many points will you get?
   b) What will your lifestyle index be in period 21?

Ans: ____________________________________

8) Answer questions 1-4 again except assume they have been asked for the liquid reward sequence. That is substitute ‘mL liquid reward’ for ‘points’.

Ans: ________________
Ans: ________________
Ans: ________________
Ans: ________________

Appendix 2: Diagram of Beverage Delivery Apparatus

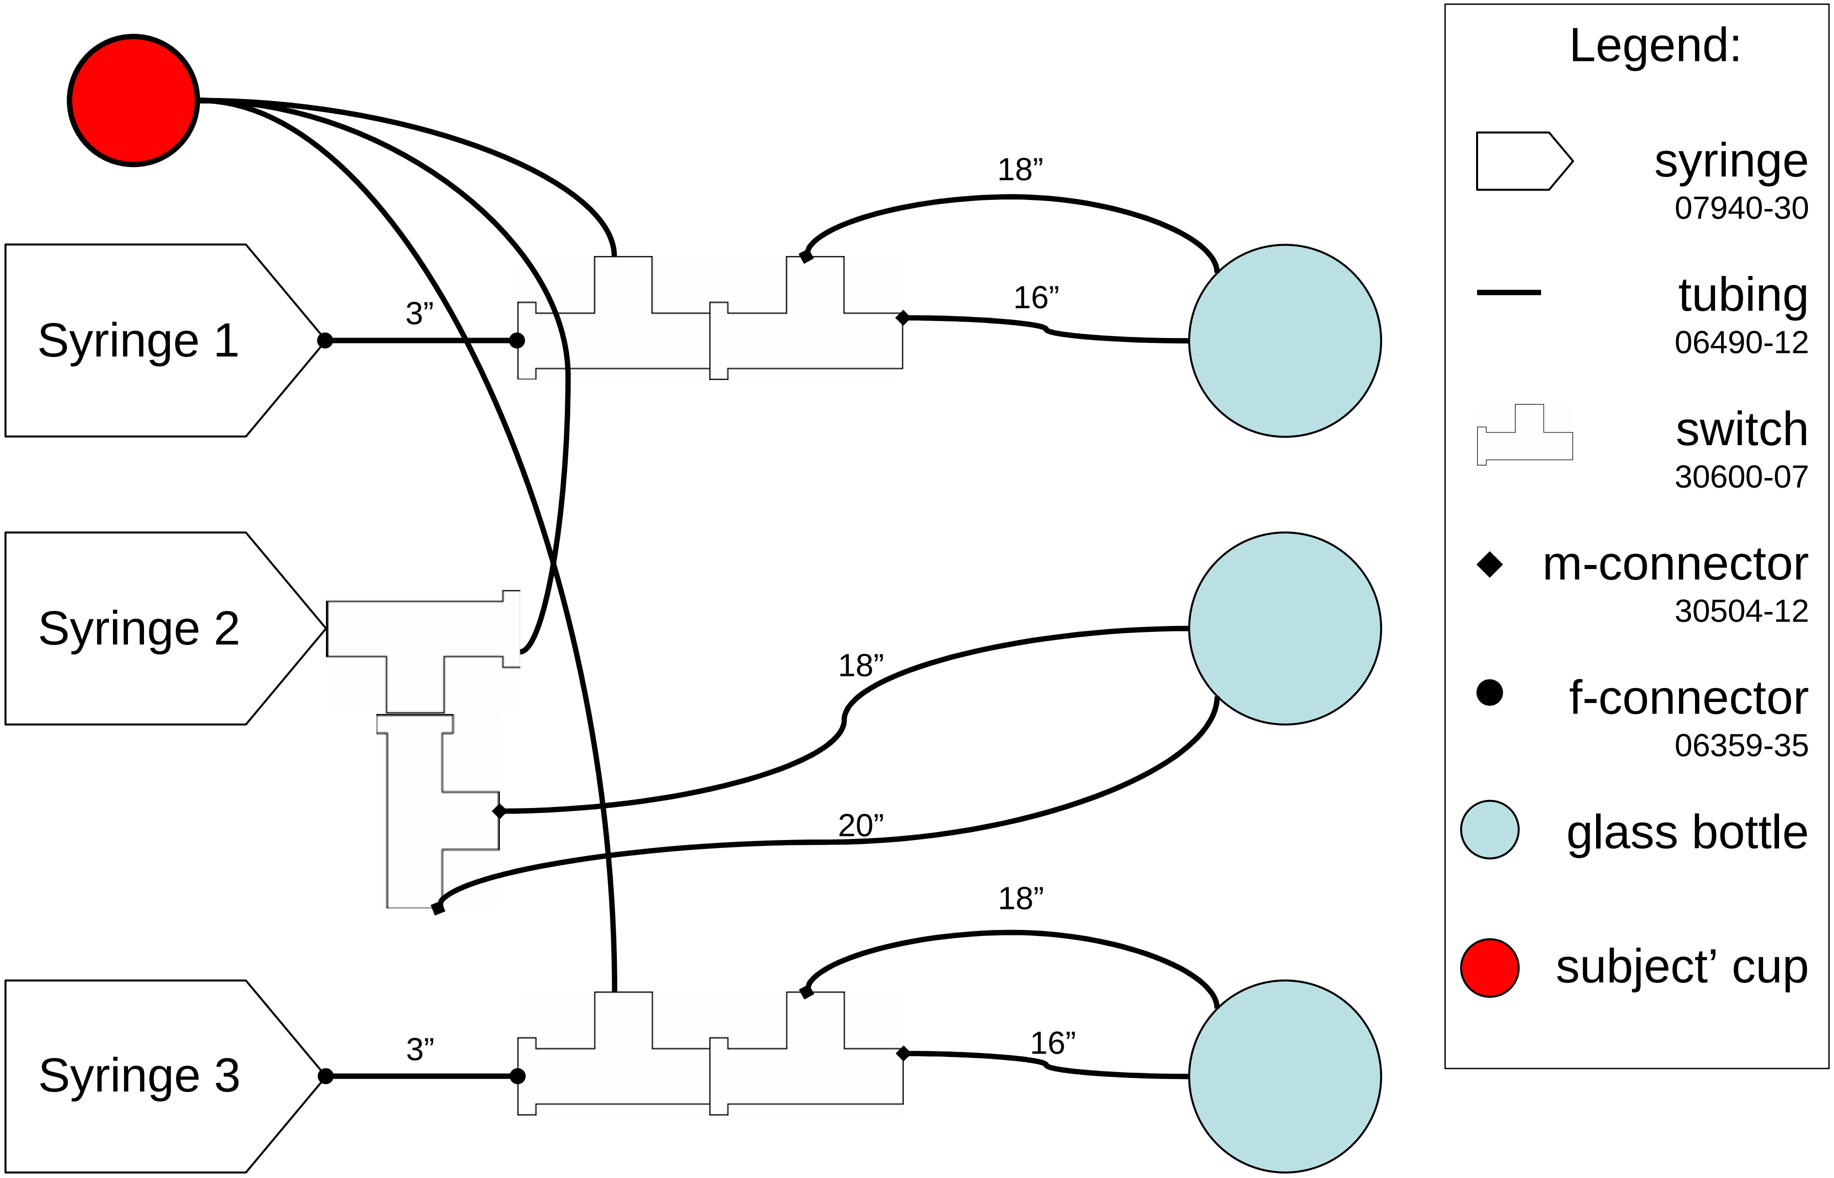

Appendix 3

Proposition 1 *With a ten minute delay and sophisticated preferences given by the hyperbolic discount model an agent has a value function that is equivalent to solving* $V_t^{'}\left(P_t, X_t, H_{t-1}\right) =$

$$\max\left\{\omega\left(C_t, H_{t-1}\right) + \widetilde{\delta} E_\eta\left[V_{t+1}^{'}\left(GP_t, X_t - C_t + \eta_{t+1} GP_t, \lambda H_{t-1} + C_t\right)\right]\right\}.$$

Proof. *Consider period 30, the last period. In this state liquid rewards for states 20-29 have been chosen, but not consumed. They can be modeled as state variables for this problem. Let* l_i *be the amount of liquid that was chosen to be consumed in period* i*. For notational simplicity let* $\overleftarrow{l_t} = \left\{l_{t-10}, l_{t-9}, l_{t-8}, l_{t-7}, l_{t-6}, l_{t-5}, l_{t-4}, l_{t-3}, l_{t-2}, l_{t-1}\right\}$ *for* $11 \leq t \leq 30$*. Then for period 30 we have 13 state variables. That is,* $s_{30} = \left\{\overleftarrow{l_{30}}, P_{30}, X_{30,} H_{29}\right\}$*. Our agent will solve* $C_{30}^{*}\left(s_{30}\right) =$

$$\underset{0\leq C_{30}\leq X_{30}}{\arg\max}\left\{\upsilon\left(l_{20}\right) + \widetilde{\beta}\sum_{i=21}^{29}\widetilde{\delta}^{i-20}\upsilon\left(l_i\right) + \widetilde{\beta}\widetilde{\delta}^{10}\omega\left(C_{30}, H_{29}\right)\right\} =$$

$$\underset{0\leq C_{30}\leq X_{30}}{\arg\max}\left\{\omega\left(C_{30}, H_{29}\right)\right\} = X_t = \underset{0\leq C_{30}\leq X_{30}}{\arg\max}\; V_{30}^{'}\left(P_{30}, A_{30}, H_{29}\right)$$

Then $\ddot{V}_{30}\left(s_{30}\right) = \sum_{i=20}^{29}\widetilde{\delta}^{i-20}\upsilon\left(l_i\right) + \widetilde{\delta}^{10}\omega\left(C_{30}^{*}, H_{29}\right) =$

$b_{30} + mV_{30}^{'}\left(P_{30}, X_{30}, H_{29}\right)$ *where* $b_t = \sum_{i=20}^{29}\widetilde{\delta}^{i-20}\upsilon\left(l_i\right)$ *and* $m = \widetilde{\delta}^{10}$*.*

Let us begin an inductive argument. Suppose for some $12 \leq t \leq 30$ *we have* $C_t^{*}\left(s_t\right) = \underset{0\leq C_t\leq X_t}{\arg\max} V_t^{'}\left(P_t, X_t, H_{t-1}\right)$ *and* $\ddot{V}_t\left(s_t\right) = b_t + mV_t^{'}\left(P_t, X_t, H_{t-1}\right)$ *where* $b_t = \sum_{i=t-10}^{t-1}\widetilde{\delta}^{i-(t-10)}\upsilon\left(l_i\right)$ *and* $m = \widetilde{\delta}^{10}$*. In period* $t-1$ *the agent solves* $C_{t-1}^{*}\left(s_{t-1}\right) =$

$$\underset{0\leq C_{t-1}\leq X_{t-1}}{\arg\max}\left\{\upsilon\left(l_{t-11}\right) + \left(\widetilde{\beta}\widetilde{\delta}\right)E_\eta\left[\ddot{V}_t\left(\overleftarrow{l_t}, GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}\right)\right]\right\}$$

where $l_{t-1} = \left(\frac{1}{2}\right)u\left(C_{t-1}, H_{t-2}\right)$*. Then* $C_{t-1}^{*}\left(s_{t-1}\right) =$

$$\underset{0\leq C_{t-1}\leq X_{t-1}}{\arg\max}\left\{\left(\widetilde{\beta}\widetilde{\delta}\right)mE_\eta\left[V_t^{'}\left(GP_{t-1}, X_{t-1} - C_{t-1} + \eta_{t-1}GP_{t-1}, \lambda H_{t-2} + C_{t-1}\right)\right] + \left(\widetilde{\beta}\widetilde{\delta}\right)b_t\right\} =$$

$$\underset{0\leq C_{t-1}\leq X_{t-1}}{\arg\max}\left\{mE_\eta\left[V_t^{'}\left(GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}\right)\right] + \sum_{i=t-9}^{t-1}\widetilde{\delta}^{i-(t-9)}\upsilon\left(l_i\right)\right\} =$$

$$\underset{0\leq C_{t-1}\leq X_{t-1}}{\arg\max}\left\{\widetilde{\delta}^{9}\upsilon\left(l_{t-1}\right) + \widetilde{\delta}^{10}E_\eta\left[V_t^{'}\left(GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}\right)\right]\right\} =$$

$$\underset{0\leq C_{t-1}\leq X_{t-1}}{\arg\max}\left\{\omega\left(C_{t-1}, H_{t-2}\right) + \widetilde{\delta}E_\eta\left[V_t^{'}\left(GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}\right)\right]\right\} =$$

$$\underset{0\leq C_{t-1}\leq X_{t-1}}{\arg\max}\; V_{t-1}^{'}\left(P_{t-1}, A_{t-1}, H_{t-2}\right).$$

Then $\ddot{V}_{t-1}\left(s_{t-1}\right) =$

$\upsilon\left(l_{t-11}\right) + \widetilde{\delta}E_\eta\ddot{V}_t\left(\overleftarrow{l_t^{*}}, GP_{t-1}, X_{t-1} - C_{t-1}^{*} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}^{*}\right)$ *where*

$l_{t-1}^{*} = \left(\frac{1}{2}\right)u\left(C_{t-1}^{*}, H_{t-2}\right).$

So $\ddot{V}_{t-1}\left(s_{t-1}\right) =$

$$\upsilon\left(l_{t-11}\right)+\widetilde{\delta}E_{\eta}\left[\widetilde{\delta}^{10}V_t^{'}\left(GP_{t-1},X_{t-1}-C_{t-1}^{*}+\eta_t GP_{t-1},\lambda H_{t-2}+C_{t-1}^{*}\right)+\sum_{i=t-10}^{t-1}\widetilde{\delta}^{i-(t-10)}\upsilon\left(l_i\right)\right] =$$

$$\upsilon\left(l_{t-11}\right)+\widetilde{\delta}E_{\eta}\left[\widetilde{\delta}^{9}\left(\upsilon\left(l_{t-1}\right)+\widetilde{\delta}V_t^{'}\left(GP_{t-1},X_{t-1}-C_{t-1}^{*}+\eta_t GP_{t-1},\lambda H_{t-2}+C_{t-1}^{*}\right)\right)\right]+$$

$$\widetilde{\delta}\sum_{i=t-10}^{t-2}\widetilde{\delta}^{i-(t-10)}\upsilon\left(l_i\right) =$$

$$\sum_{i=t-10}^{t-2}\widetilde{\delta}^{i-(t-11)}\upsilon\left(l_i\right)+\widetilde{\delta}^{10}E_{\eta}\left[\omega\left(C_{t-1}^{*},H_{t-2}\right)+\widetilde{\delta}V_t^{'}\left(GP_{t-1},X_{t-1}-C_{t-1}^{*}+\eta_t GP_{t-1},\lambda H_{t-2}+C_{t-1}^{*}\right)\right] =$$

$$b_{t-1}+m\left[\omega\left(C_{t-1}^{*},H_{t-2}\right)+\widetilde{\delta}E_{\eta}\left[V_t^{'}\left(GP_{t-1},X_{t-1}-C_{t-1}^{*}+\eta_t GP_{t-1},\lambda H_{t-2}+C_{t-1}^{*}\right)\right]\right] =$$

$$b_{t-1}+mV_{t-1}^{'}\left(P_{t-1},X_{t-1},H_{t-2}\right).$$

Then for all t *s.t.* $11 \leq t \leq 30$ *we have* $C_t^{*}\left(s_t\right) = \underset{0\leq C_t\leq X_t}{\arg\max}V_t^{'}\left(P_t,X_t,H_{t-1}\right)$ *and*

$\ddot{V}_t\left(s_t\right) = b_t+mV_t^{'}\left(P_t,X_t,H_{t-1}\right)$ *where* $b_t = \sum_{i=t-10}^{t-1}\widetilde{\delta}^{i-(t-10)}\upsilon\left(l_i\right)$ *and* $m = \widetilde{\delta}^{10}$.

Consider periods 1-10. Let $\overleftarrow{l_t} = \{l_1,...l_t\}$ *for* $1 \leq t \leq 10$. *For period 10 we have 12 state variables. That is,* $s_{10} = \left\{\overleftarrow{l_{10}},P_{30},X_{30,}H_{29}\right\}$. *Our agent will solve* $C_{10}^{*}\left(s_{10}\right) =$

$$\underset{0\leq C_{10}\leq X_{10}}{\arg\max}\left\{\left(\widetilde{\beta}\widetilde{\delta}\right)E_{\eta}\left[\ddot{V}_{11}\left(\overleftarrow{l_{11}},GP_{10},X_{10}-C_{10}+\eta_{11}GP_{10},\lambda H_9+C_{10}\right)\right]\right\}$$

So $C_{10}^{*}\left(s_{10}\right) = \underset{0\leq C_{10}\leq X_{10}}{\arg\max}\left\{E_{\eta}\left[b_{11}+mV_{11}^{'}\left(GP_{10},X_{10}-C_{10}+\eta_{11}GP_{10},\lambda H_9+C_{10}\right)\right]\right\} =$

$$\underset{0\leq C_{10}\leq X_{10}}{\arg\max}\left\{E_{\eta}\left[\sum_{i=1}^{10}\widetilde{\delta}^{i-1}\upsilon\left(l_i\right)+\widetilde{\delta}^{10}V_{11}^{'}\left(GP_{10},X_{10}-C_{10}+\eta_{11}GP_{10},\lambda H_9+C_{10}\right)\right]\right\} =$$

$$\underset{0\leq C_{10}\leq X_{10}}{\arg\max}\left\{E_{\eta}\left[\widetilde{\delta}^{9}\upsilon\left(l_{10}\right)+\widetilde{\delta}^{10}V_{11}^{'}\left(GP_{10},X_{10}-C_{10}+\eta_{11}GP_{10},\lambda H_9+C_{10}\right)\right]\right\} =$$

$$\underset{0\leq C_{10}\leq X_{10}}{\arg\max}\left\{\omega\left(C_{10},H_9\right)+\widetilde{\delta}E_{\eta}\left[V_{11}^{'}\left(GP_{10},X_{10}-C_{10}+\eta_{11}GP_{10},\lambda H_9+C_{10}\right)\right]\right\} =$$

$$\underset{0\leq C_{10}\leq X_{10}}{\arg\max}\ V_{10}^{'}\left(P_{10},A_{10},H_9\right)$$

Then $\ddot{V}_{10}\left(s_{10}\right) = \widetilde{\delta}E_{\eta}\left[\ddot{V}_{11}\left(\overleftarrow{l_{11}^{*}},GP_{10},X_{10}-C_{10}^{*}+\eta_{11}GP_{10},\lambda H_9+C_{10}^{*}\right)\right] =$

$$\widetilde{\delta}E_{\eta}\left[b_{11}^{*}+mV_{11}^{'}\left(GP_{10},X_{10}-C_{10}^{*}+\eta_{11}GP_{10},\lambda H_9+C_{10}^{*}\right)\right] =$$

$$\widetilde{\delta}E_{\eta}\left[\sum_{i=1}^{9}\widetilde{\delta}^{i-1}\upsilon\left(l_i\right)+\widetilde{\delta}^{9}\omega\left(C_{10}^{*},H_9\right)+\widetilde{\delta}^{10}V_{11}^{'}\left(GP_{10},X_{10}-C_{10}^{*}+\eta_{11}GP_{10},\lambda H_9+C_{10}^{*}\right)\right] =$$

$$\sum_{i=1}^{9}\widetilde{\delta}^{i}\upsilon\left(l_i\right)+\widetilde{\delta}^{10}\left[\omega\left(C_{10}^{*},H_9\right)+\widetilde{\delta}E_{\eta}\left[V_{11}^{'}\left(GP_{10},X_{10}-C_{10}+\eta_{11}GP_{10},\lambda H_9+C_{10}\right)\right]\right] =$$

$$b_{10}+\widetilde{\delta}^{10}\left[\omega\left(C_{10}^{*},H_9\right)+\widetilde{\delta}E_{\eta}\left[V_{11}^{'}\left(GP_{10},X_{10}-C_{10}+\eta_{11}GP_{10},\lambda H_9+C_{10}\right)\right]\right]$$

$b_{10}+mV_{10}^{'}\left(P_{10},X_{10},H_9\right)$ *where* $b_t = \sum_{i=1}^{9}\widetilde{\delta}^{i+(10-t)}\upsilon\left(l_i\right)$ *and* $m = \widetilde{\delta}^{10}$.

Let us begin an inductive argument. Suppose for some $2 \leq t \leq 10$ *we have* $C_t^{*}\left(s_t\right) = \underset{0\leq C_t\leq A_t}{\arg\max}V_t^{'}\left(P_t,X_t,H_{t-1}\right)$ *and* $\ddot{V}_t\left(s_t\right) = b_t+mV_t^{'}\left(P_t,X_t,H_{t-1}\right)$ *where*

$b_t = \sum_{i=1}^{t-1} \widetilde{\delta}^{i+(10-t)} \upsilon(l_i)$ *and* $m = \widetilde{\delta}^{10}$. *In period* $t-1$ *the agent solves* $C^*_{t-1}(s_{t-1}) =$
$\underset{0\le C_{t-1}\le X_{t-1}}{\arg\max} \left\{ \left(\widetilde{\beta}\widetilde{\delta}\right) E_\eta \left[\ddot{V}_t \left(\overleftarrow{l_t}, GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1} \right) \right] \right\} =$
$\underset{0\le C_{t-1}\le X_{t-1}}{\arg\max} \left\{ E_\eta \left[\sum_{i=1}^{t-1} \widetilde{\delta}^{i+(10-t)} \upsilon(l_i) + \widetilde{\delta}^{10} V_t^{'} (GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}) \right] \right\} =$
$\underset{0\le C_{t-1}\le X_{t-1}}{\arg\max} \left\{ E_\eta \left[\widetilde{\delta}^{9} \upsilon(l_{t-1}) + \widetilde{\delta}^{10} V_t^{'} (GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}) \right] \right\} =$
$\underset{0\le C_{t-1}\le X_{t-1}}{\arg\max} \left\{ \omega(C_t, H_{t-1}) + E_\eta \left[\widetilde{\delta}^{10} V_t^{'} (GP_{t-1}, X_{t-1} - C_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C_{t-1}) \right] \right\} =$
$\underset{0\le C_{t-1}\le X_{t-1}}{\arg\max} \left\{ V_{t-1}^{'} (P_{t-1}, A_{t-1}, H_{t-2}) \right\}$.
Then $\ddot{V}_{t-1}(s_{t-1}) =$
$\widetilde{\delta} E_\eta \left[b_t^* + m V_t^{'} \left(GP_{t-1}, X_{t-1} - C^*_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C^*_{t-1} \right) \right] =$
$\widetilde{\delta} E_\eta \left[\sum_{i=1}^{t-2} \widetilde{\delta}^{i+(10-t)} \upsilon(l_i) + \widetilde{\delta}^{9} \omega\left(C^*_{t-1}, H_{t-2}\right) + \widetilde{\delta}^{10} V_t^{'} \left(GP_{t-1}, X_{t-1} - C^*_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C^*_{t-1} \right) \right] =$
$\sum_{i=1}^{t-2} \widetilde{\delta}^{i+(10-t)+1} \upsilon(l_i) + \widetilde{\delta}^{10} \left[\omega\left(C^*_{t-1}, H_{t-2}\right) + \widetilde{\delta} E_\eta \left[V_t^{'} \left(GP_{t-1}, X_{t-1} - C^*_{t-1} + \eta_t GP_{t-1}, \lambda H_{t-2} + C^*_{t-1} \right) \right] \right] =$
$\sum_{i=1}^{t-2} \widetilde{\delta}^{i+(10-(t-1))} \upsilon(l_i) + m \left[V_{t-1}^{'} (P_{t-1}, X_{t-1}, H_{t-2}) \right] =$
$b_{t-1} + m \left[V_{t-1}^{'} (P_{t-1}, X_{t-1}, H_{t-2}) \right]$.
Then for all t *s.t.* $1 \le t \le 10$ *we have* $C^*_t(s_t) = \underset{0\le C_t\le X_t}{\arg\max} V_t^{'} (P_t, X_t, H_{t-1})$ *and*
$\ddot{V}_t(s_t) = b_t + m V_t^{'} (P_t, X_t, H_{t-1})$ *where* $b_t = \sum_{i=1}^{t-1} \widetilde{\delta}^{i+(10-t)} \upsilon(l_i)$ *and* $m = \widetilde{\delta}^{10}$.
Hence for all t *s.t.* $1 \le t \le 30$ *we have* $C^*_t(s_t) = \underset{0\le C_t\le X_t}{\arg\max} V_t^{'} (P_t, X_t, H_{t-1})$.
Thus our problem is equivalent to solving $V_t^{'} (P_t, X_t, H_{t-1}) =$
$\max \left\{ \omega(C_t, H_{t-1}) + \widetilde{\delta} E_\eta \left[V_{t+1}^{'} \left(GP_t, X_t - C_t + \eta_{t+1} GP_t, \lambda H_{t-1} + C_t \right) \right] \right\}$ *over* $1 \le$
$t \le 30$. ■

References

Allen W. Todd, and Carroll D., Christopher (2001): ‘Individual Learning about Consumption,’. *Macroeconomic Dynamics, 5(2),*

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*, 2001, 15(3), 47-68.

Ballinger, T. Parker, Palumbo, Michael and Wilcox, Nathaniel T., “Precautionary Saving and Social Learning Across Generations: An Experiment.” *Economic Journal*, 2003 Vol. 113, 920-947.

Ballinger, T. Parker, Hudson, Eric, Karkoviata, Leonie and Wilcox, Nathaniel T., “Saving Performance and Cognitive Abilities.” Unpublished working paper. Stephen F. Austin State University, 2005.

Barro, Robert J. (1999), "Ramsey meets Laibson in the neoclassical growth model," Quarterly Journal of Economics 114:1125-1152.

Benartzi, S. & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110(1) 73-92.

Benjamin, Daniel J., and Jesse M. Shapiro (2005). “Who is ‘Behavioral’? Cognitive ability and anomalous preferences.” Harvard University mimeo, July.

Bernheim, B. Douglas. (1998). “Financial illiteracy, education and retirement saving” In Olivia S. Mitchell and Sylvester J. Schieber (Eds.), *Living with Defined Contribution Pensions*, Pension Research Council, Wharton School, University of Pennsylvania, 38-68.

Bernheim, B. Douglas and Daniel M. Garrett. “The Effects of Financial Education in the Workplace: Evidence from a Survey of Households,” *Journal of Public Economics* 87(7-8), August 2003, 1487-1519.

Bernheim, B. Douglas, and Rangel, Antonio. “Addiction and Cue-Triggered Decision Processes.” *American Economic Review*, 2004, 94(5), 1558-1590

Bernheim, Douglas and Antonio Rangel (2002) “Cognition, Evaluation, and Hedonics: A Brain Processes Model of Decision-Making”, Stanford manuscript.

Bowman, David; Deborah Minehart; and Matthew Rabin. (1999). "Loss Aversion in a Consumption-Savings Model", Journal of Economic Behavior and Organization 38(2), 155-178.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1998): ‘Learning from the Behavior of Others: Conformity, Fads and Informational Cascades,’ *Journal of Economic Perspectives*, 12(3),151-170.

Browning, Martin and Annamaria Lusardi (1996). "Household Saving: Micro Theories and Micro Facts." Journal of Economic Literature **34**(4): 1797-1855.

Cagetti, Marco (2003): 'Wealth Accumulation over the Life Cycle and Precautionary Savings,' *Journal of Business and Economic Statistics*. 21,339-353.

Camerer, C. F., Babcock, L., Loewenstein, G. & Thaler, R. (1997). Labor supply of New York City cab drivers: One day at a time. Quarterly Journal of Economics, 111, 408-441.

Camerer, C. F. (2000). Prospect theory in the wild: Evidence from the field. In D. Kahneman & A. Tversky (eds.), Choices, values, and frames (pp. 288-300). Cambridge: Cambridge University Press.

Camerer, Colin F. (2003). *Behavioral Game Theory*. Princeton: Princeton University Press.

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

Carroll, Christopher D. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, 1997, 112(1), 1-56.

Carroll, Christopher D. 'A Theory of the Consumption Function, With and Without Liquidity Constraints.' *Journal of Economic Perspectives*, 2001, 15(3), 23-45.

Carroll, Christopher D., Jody Overland, and David N. Weil. "Saving and Growth with Habit Formation." *American Economic Review*, 2000, 90(3), 341-355.

Carroll, Christopher D.; Changyong Rhee, and Byungkun Rhee (1999). ``Does Cultural Origin Affect Saving Behavior? Evidence from Immigrants," *Economic Development and Cultural Change*, 48(1), 33-50,

Carroll, Christopher D., and David N. Weil (1994): *'Saving and Growth: A Reinterpretation.' Carnegie-Rochester Conference Series on Public Policy,* 40, 133-192.

Casella, George and Berger, Roger L. *Statistical Inference*. USA: The Duxbury Group, 2002.

Chua, Zhikang and Camerer, Colin F. "Experiments on Intertemporal Consumption with Habit Formation and Social Learning." Unpublished working paper. Caltech, 2003. Available at http://www.hss.caltech.edu/~camerer/savingsjpe7.doc.

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

Deaton, Angus (1992): *Understanding Consumption.(Clarendon Lectures in Economics).* Oxford University Press.

Deaton, Angus S. “Savings and Liquidity Constraints.” *Econometrica*, 1991, 59(5), 1221-1248.

Degeorge, F., Patel, J., & Zeckhauser, R. (1999). Earnings management to exceed thresholds. Journal of Business, 72, 1-33.

Della Vigna, Stefano. and M. Daniele Paserman (2005). "Job search and impatience." Journal of Labor Economics **23**(3): 527-588.

Dusenberry, James S. (1949). Income, Savings, and the Theory of Consumer Behavior. Cambridge MA: Harvard University Press.

Fehr, Ernst and Zych, Peter K. “Do Addicts Behave Rationally?” *Scandinavian Journal of Economics*, 1998, 100(3). 643-662.

Fehr and Zych. (in press) Intertemporal choice under habit formation. In C. Plott and V Smith (Eds) Handbook of Experimental Economics Results, http://www.iew.unizh.ch/wp/iewwp043.pdf

Fehr, Ernst and Jean-Robert Tyran (2003). Limited rationality and strategic interaction—the impact of the strategic environment on nominal inertia. Institute for Empirical Research in Economics, University of Zürich, Working Paper No. 130. http://www.iew.unizh.ch/wp/iewwp130.pdf

Frederick, Shane. “Cognitive Reflection and Decision Making.” *Journal of Economic Perspectives*, 2005 19(4), 24-42.

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

Fudenberg, Drew and Levine, David. “A Dual Self Model of Impulse Control.” Unpublished working paper. Harvard, 2004. Available at http://post.economics.harvard.edu/faculty/fudenberg/papers/dual_self.pdf.

Gale, Douglas (1996): ‘What have we learned from social learning?’ European Economic Review(40).

Genesove, D. & Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. Quarterly Journal of Economics, 116(4), 1233-1260.

Gourinchas, Pierre-Olivier, and Jonathan Parker (1999): 'Consumption over the Life Cycle,' Econometrica, 70, 47-89.

Gul, Faruk and Pesendorfer, Wolfgang. "Temptation and Self-Control." *Econometrica*, 2001, 69(6), 1403-1435.

Hall, Robert E., and Frederic Mishkin (1982): 'The sensitivity of Consumption to Transitory Income: Evidence from PSID Households,' *Econometrica*, 50, 461-481.

Houser, Daniel, Keane, Michael P, and McCabe, Kevin. "Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm." *Econometrica,* 2004, 72(3), 781-822.

Johnson, Eric J., Gaechter, Simon and Herrmann, Andreas, "Exploring the Nature of Loss Aversion" (March 2006). IZA Discussion Paper No. 2015 Available at SSRN: http://ssrn.com/abstract=892336

Kahneman, D., Knetsch, J. L. & Thaler, R. (1990). Experimental tests of the endowment effect and the Coase theorem. Journal of Political Economy, 98, 1325-1348.

Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47, 263-291.

Kandori, Michihiro (1997): *'Evolutionary Game Theory in Economics,*' Chapter 8 in *Advances in Economics and Econometrics: Theory and Applications, Volume I,* edited by David Kreps, and W. Wallis. Cambridge University Press, Cambridge.

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

Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 1997, 62(2), 443-78.

Lettau, Martin, and Harald Uhlig (1999): 'Rules of Thumb versus Dynamic Programming,' *American Economic Review* (89)(1) 148-174.

Loewenstein, George and O'Donoghue, Ted. "Animal Spirits: Affective and Deliberative Processes in Economic Behavior." Mimeo. Cornell University. May 1995.

Loewenstein, George and Prelec, Drazen. “Anomalies in Intertemporal Choice: Evidence and Interpretation.” *Quarterly Journal of Economics*, 1992, 107(2), 573-597.

Lusardi, Annamaria. Explaining why so many households do not save. Dartmouth working paper, December 2002. http://www.dartmouth.edu/~alusardi/donotsave.pdf

Mailath, George. (1998). Do people play Nash equilibrium? Lessons from evolutionary game theory. *Journal of Economic Literature*, 36, 1347-1374

Marimon, Ramon (1997): ‘Learning from Learning in Economics’ Chapter 9 in *Advances in Economics and Econometrics: Theory and Applications, Volume I*, edited by David Kreps, and W. Wallis. Cambridge University Press, Cambridge.

O'Doherty, J., Critchley, H.D., Deichmann, R.,Dolan R.J. Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *Journal of Neuroscience*, in press.

O'Donoghue, T. and M. Rabin (1999). "Doing it now or later." American Economic Review **89**(1): 103-124.

Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance*, 53(5), 1775-1798.

Read, Daniel, George Loewenstein, and Matthew Rabin. (1999). "Choice bracketing." Journal of Risk and Uncertainty **19**(1-3): 171-197.

Samuelson, Larry (1997): *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge.

Sargent, Thomas J. (1993): *Bounded Rationality in Macroeconomics, Chapter 5*. Oxford University Press, Oxford.

Sayman, Serdar. and Ayse Onculer (2005). "Effects of study design characteristics on the WTA-WTP disparity: A meta analytical framework." Journal of Economic Psychology **26**(2): 289-312.

Scholtz, J. K., A. Seshadri, et al. (forthcoming). "Are Americans Saving “Optimally” for Retirement?" Journal of Political Economy.

Schotter, Andrew and Barry Sopher. (in press). Social Learning and Coordination Conventions in Inter-Generational Games: An Experimental Study. Journal of Political Economy.

Shafir, E., Diamond, P. & Tversky, A. (1997). Money illusion. Quarterly Journal of Economics, 112(2), 341-74.

Shefrin, H. M. & Thaler, R. H. (1992). Mental accounting, saving, and self-control. In G. Loewenstein & J. Elster (eds.), Choice over time (287-329). New York: Russell Sage Foundation.

Souleles, Nicholas S. (1999): 'The Response of Household Consumption to Income Tax Refunds,' *American Economic Review*.89, 947-958.

Strotz, Robert H. (1956). "Myopia and Inconsistency in Dynamic Utility Maximization." The Review of Economic Studies **23**(3): 165-180.

Tanaka, T., C. F. Camerer, et al. (2006). Poverty, politics, and preferences: Experimental and survey data from Vietnam, California Institute of Technology. www.its.caltech.edu/~ttanaka/Vietnam.pdf

Thaler, Richard H. (1999). Mental accounting matters. Journal of Behavioral Decision Making, 12, 183-206.

Thaler, Richard H. and Shlomo Benartzi (2004). "Save More Tomorrow (TM): Using behavioral economics to increase employee saving." Journal of Political Economy **112**(1): S164-S187.

Thaler, Richard and Shefrin, H. M. "An Economic Theory of Self-Control." *Journal of Political Economy*, 89, 392-406.

Timmerman, Allan. (1993). How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices. Quarterly Journal of Economics, 108, 1135-1145.

Van De Stadt, Huib, A. Kapteyn, and S. Van De Geer (1985): 'The Relativity of Utility: Evidence from Panel Data.' *Review of Economics and Statistics*, 67(4).179-187.

Venti, Steven (2005), "Choice, Behavior and Retirement saving," forthcoming in G. Clark, A. Munnell and M. Orzsag (eds), *Oxford Handbook of Pensions and Retirement Income*, Oxford: Oxford University Press.

Zeldes, Stephen P. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics*, 1989, 104(2), 275-298.

Table I: Notation for the buffer stock model

Δ	Time preference factor (assumed constant)
X_S	Total cash/resources available in period *s* (“cash on hand”)
S_S	Savings in period *s* (portion of X_s not consumed)
C_S	Spending in period *s*
R	Gross interest rate each period
H_{S-1}	Habit stock from period *s-1*
$u(C_S, H_{S-1})$	Utility
Y_S	Actual income in period *s*
P_S	Permanent labor income in period *s*
$P_{s+1} = G P_s$	*G* is the growth rate of permanent income each period (assumed constant)
η_s	Income shock in period, a random variable, drawn from a distribution

Table II: Summary statistics of actual point outcomes

Lifecycle	1	2	3	4	5	6	7
Private Learning							
Mean Points $\sum_{t=1}^{T} u\left(C_t, H_{t-1}\right)$ (average over subjects)	118	-53	224	450	-65	435	440
Std. dev. Points	635	694	498	297	475	255	146
Deviation from Optimum $\sum_{t=1}^{T} u\left(C_t, H_{t-1}\right) -$ $\sum_{t=1}^{T} u\left(C_t^*, H_{t-1}^*\right)$ (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 III: 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 IV: 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 V: 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 VI: 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** (0.225)	0.904** (0.225)
mean $\tilde{\beta}$ of Immediate (std deviation)	1 n/a	0.758** (0.502)	0.865* (0.396)
total squared deviations periods 1-10, immediate subjects, from model using mean $\tilde{\beta}$ and $\tilde{\delta}$	34.16	31.32	31.87

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

Figure I: An Optimal Consumption Path

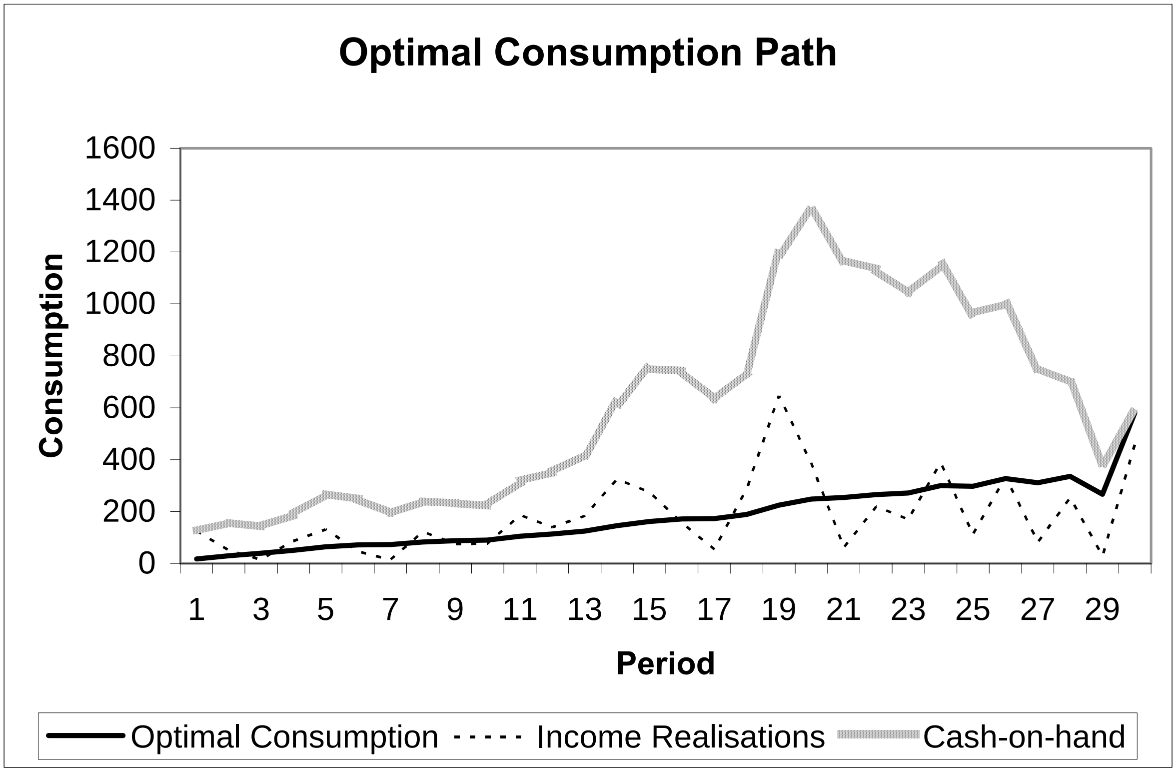

Figure 2: Quasi-hyperbolic consumption vs. optimal

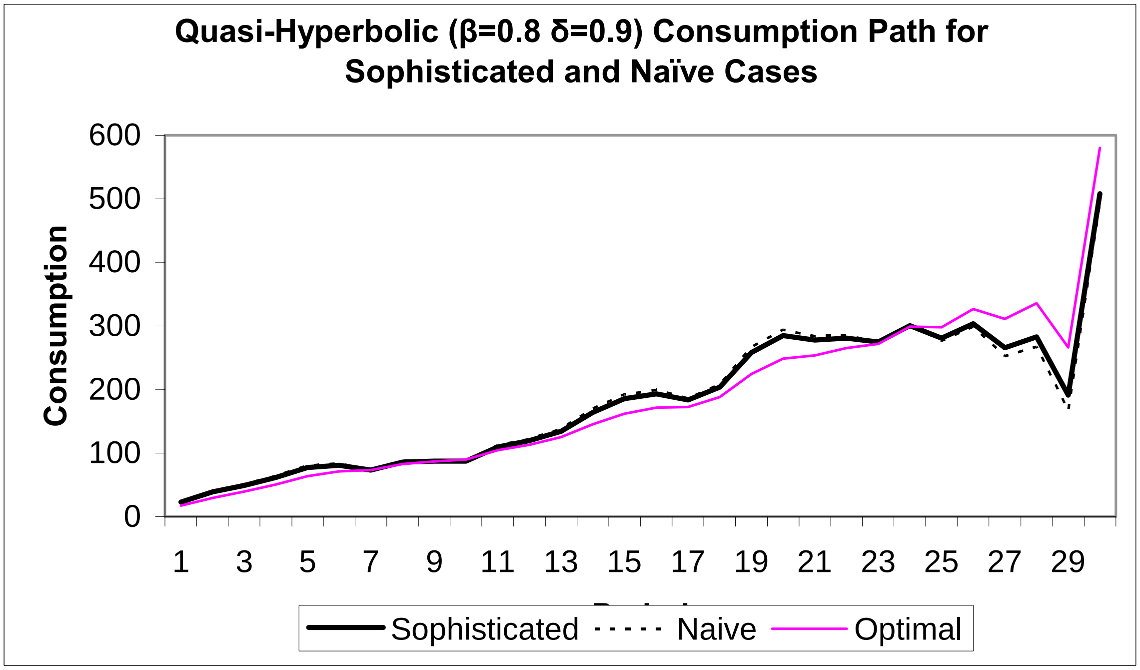

Figure III: Screenshot of Excel interface

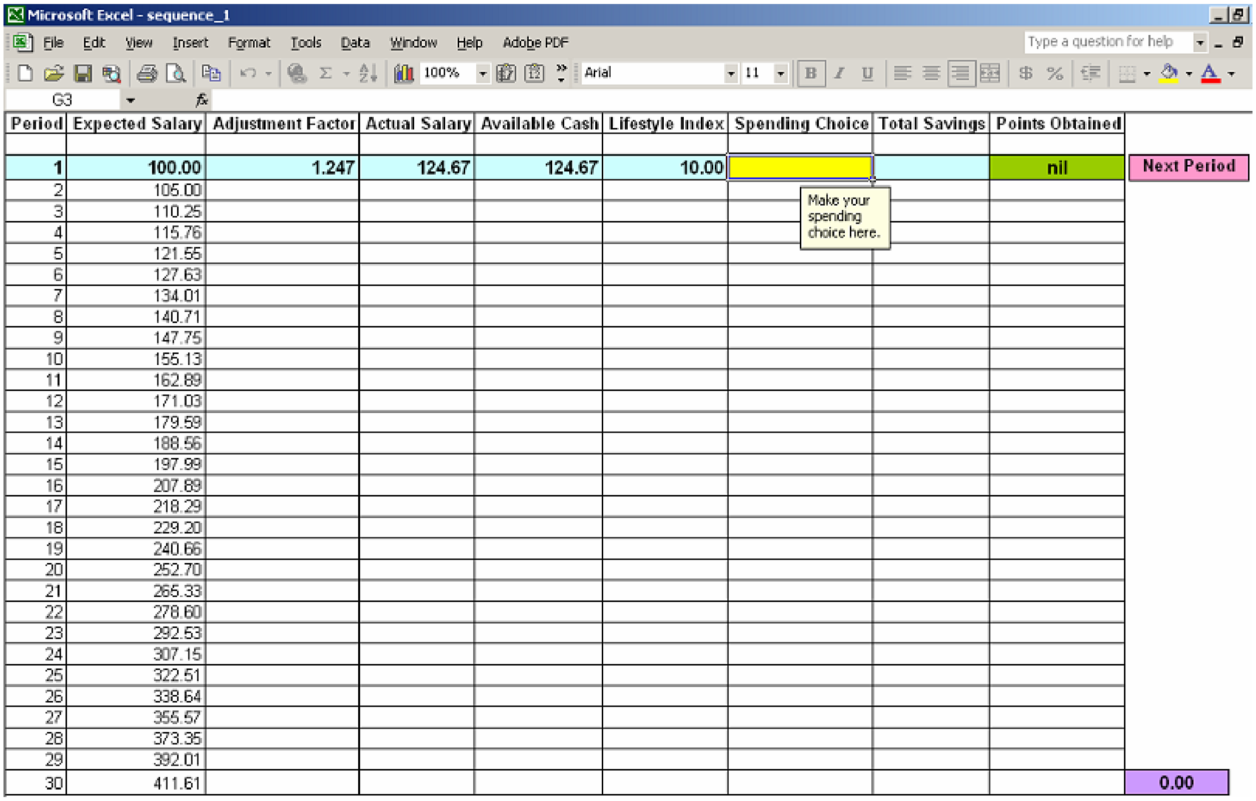

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	**100.00**	**1.247**	**124.67**	**124.67**	**10.00**			**nil**	**Next Period**
2	105.00								
3	110.25								
4	115.76								
5	121.55								
6	127.63								
7	134.01								
8	140.71								
9	147.75								
10	155.13								
11	162.89								
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								**0.00**

Make your spending choice here.

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

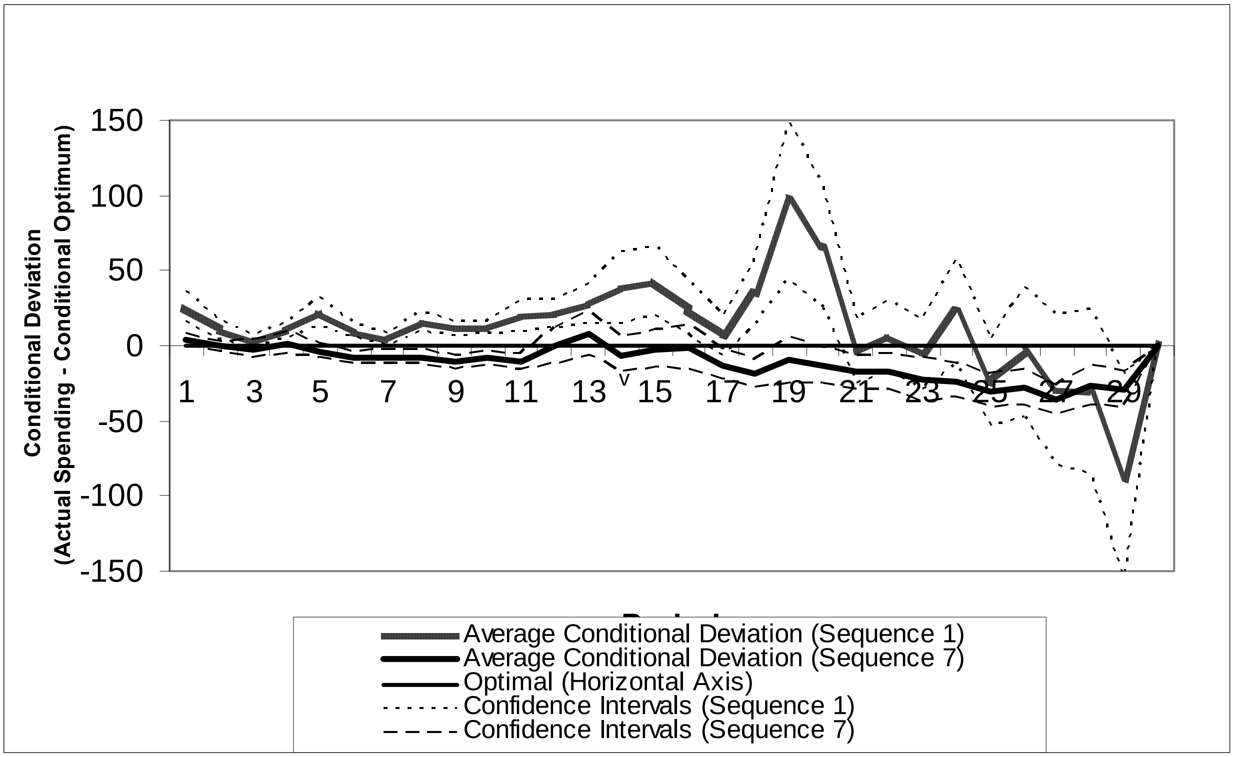

Figure V: Deviations from conditional optima, lifecycle 1 and 7, social learning

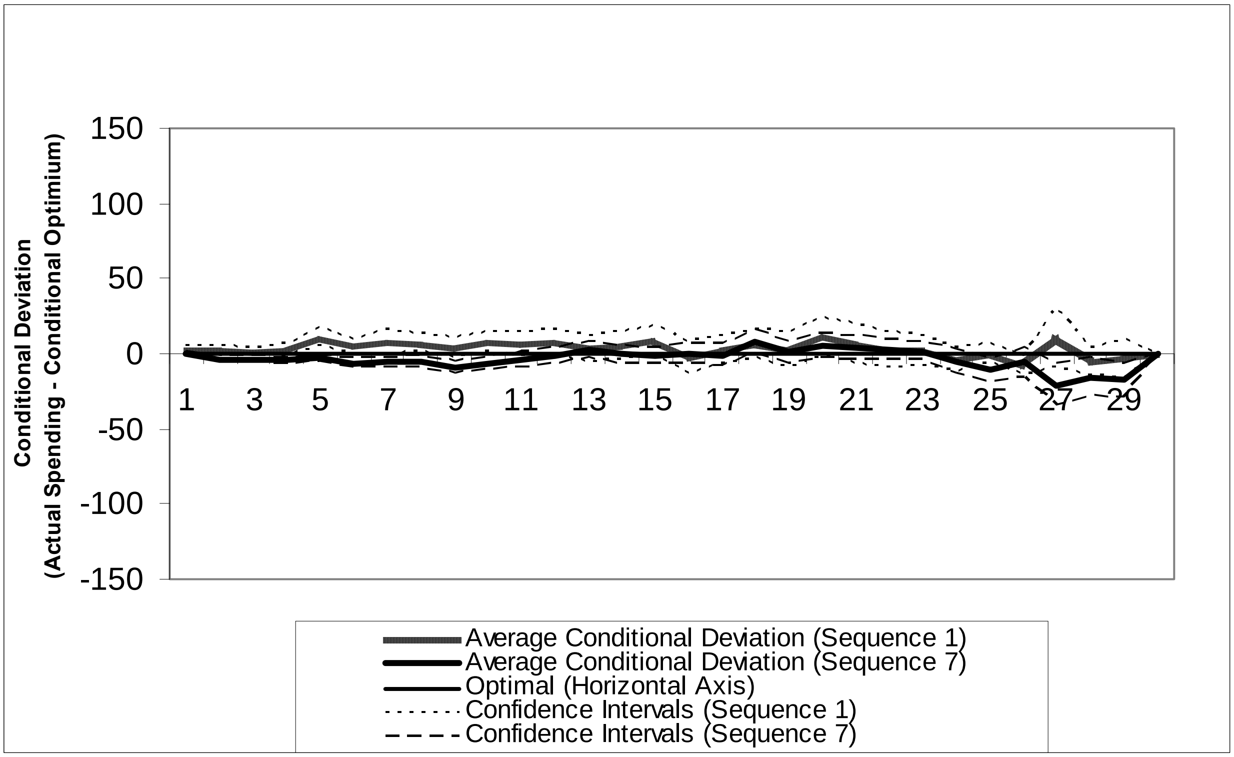

Figure VI: Ratio of average consumption to conditional optimal by condition, periods 1-10

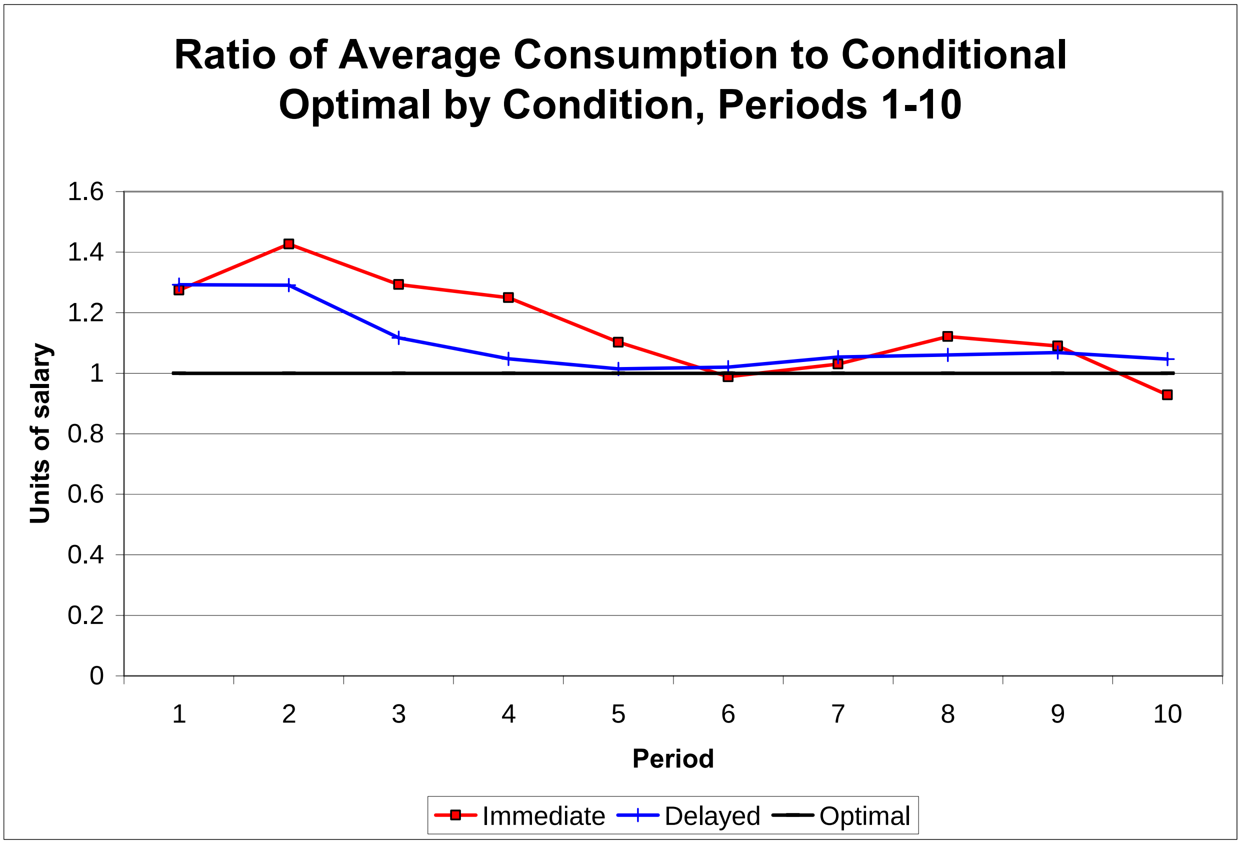

Figure VII: Actual (y) and conditionally optimal (x) utilities, observations between -50 and +50 (n=14,228)

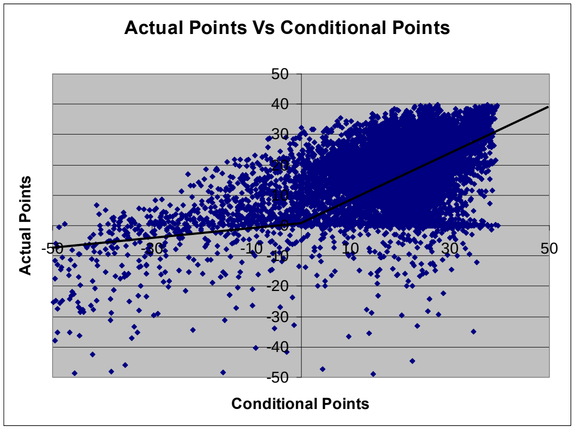

Figure VIII: Frequency of actual and conditionally optimal utilities, observations between -10 and +10 (n=14,228)

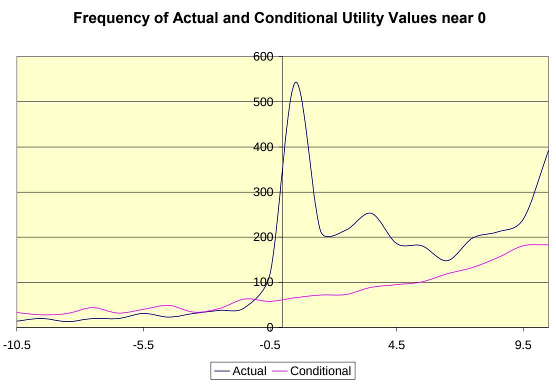

Figure IX: Actual (y) and conditionally optimal (x) utilities in study 2, observations between -50 and +50 (n=5,840)

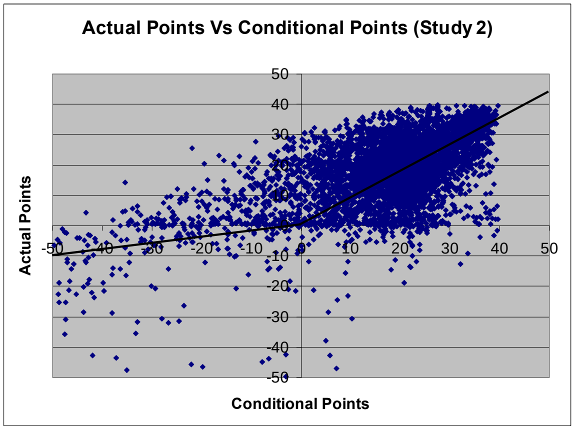

Figure X: Actual (y) and conditionally optimal (x) ml of beverage, observations between -25 and +25 (n=1346)