Exponential-Growth Bias in Experimental Consumption Decisions*

Matthew R. Levy
London School of Economics

Joshua Tasoff
Claremont Graduate University

September 3, 2016

Abstract

Exponential-growth bias (EGB) is the tendency to neglect the power of compounding interest, and has been found to be widespread in the population. A person with EGB will misperceive the intertemporal budget constraint, overestimating lifetime wealth and underestimating the differences in the cost of consumption across periods. We test four comparative static predictions implied by EGB: (1) compound interest will increase consumption, (2) budget-neutral delays in income will increase consumption, (3) the person will exhibit a form of dynamic inconsistency that depends solely on the current account balance and is independent of time preferences, and (4) framing the frequency of interest in shorter units increases consumption. We test these predictions using an induced-value consumption-savings experiment in the lab, and find evidence in support of all predictions against the rational benchmark. We consider two rules of thumb as alternative hypotheses and find that they cannot explain the results.

Keywords: exponential-growth bias, dynamic inconsistency, overconsumption, consumption-savings, personal finance
JEL: D03, D14, D91

*We would like to gratefully acknowledge the financial support of a Fletcher Jones Foundation Faculty Research Grant from Claremont Graduate University. We are grateful to Masyita Crystallin, Peiran Jiao, Andrew Royal, Quinn Keefer, and Oliver Curtiss for research assistance. We thank Alexander Brown, Ananda Ganguly, Matthew Rabin, Paige Skiba, Justin Sydnor, Charles Thomas, Jonathan Zinman, and various seminar participants for helpful comments. Institutional Review Board approval was obtained from OHRPP at UCLA [IRB #12-001092], CGU [IRB #1591], and from UC Irvine as a "typical experiment" under the ESSL IRB protocol narrative. Some results from Experiment 1 of this paper were formerly included in early working-paper versions of "Exponential-Growth Bias and Lifecycle Consumption".

1 Introduction

How one consumes and saves over the lifecycle is one of the most consequential financial decisions a person makes. Two fundamentally important features of the decision problem are the person's perception of prices over time and the person's perception of their total lifetime wealth. Much research has shown that people underestimate the speed at which compounding interest grows, a judgment bias referred to as exponential-growth bias (Wagenaar and Sagaria, 1975; Stango and Zinman, 2009). This tendency has been demonstrated in numerous lab (see e.g. Wagenaar and Timmers, 1979; Keren, 1983; Benzion, Granot and Yagil, 1992; MacKinnon and Wearing, 1991; Eisenstein and Hoch, 2007; McKenzie and Liersch, 2011; Ensthaler, Nottmeyer, Weizsäcker and Zankiewicz, 2013) and field studies (see e.g. Almenberg and Gerdes, 2012; Soll, Keeney and Larrick, 2011; Goda, Manchester and Sojourner, 2014; Song, 2012). Exponential-growth bias (EGB) leads a person to have biased perceptions of both prices over time and total lifetime wealth. These misperceptions can result in profound errors in the person's consumption-savings decisions.

In this paper, we test the comparative static predictions of EGB in a consumption-savings task in a controlled lab experiment. Outside the lab, people's income profiles, interest rates, and information are endogenous, so one cannot make causal inferences from such an analysis.[1] To experimentally manipulate people's income profiles, interest rates, and information over their lifetime and measure the impact on their lifecycle consumption would be both ethically and financially prohibitive. Moreover, unobserved heterogeneity in preferences, subjective life expectancy, and consumption shocks make it difficult to draw clear conclusions from observational data. Instead, we use choices in an induced-value experiment as a model of real-world behavior (Smith, 1976). In addition to exogenous manipulations, the experimental setting has the added advantage of perfectly controlling for credit constraints and uncertainty. The laboratory task therefore allows for an internally valid test of EGB on consumption-savings tasks.

The model of EGB specifies in a precise way the misperceptions of an agent that neglects compounding interest (Levy and Tasoff, 2015). An agent with EGB views the period-T future

[1] Stango and Zinman (2009), Levy and Tasoff (2015), and Goda et al. (2015) find that EGB is correlated with several financial outcomes, but they note that this is a correlation which may be subject to non-causal interpretations.

value of a dollar at time $t \leq T$ as $p_t(\vec{\imath}, T; \alpha) = \left(\prod_{s=t}^{T}(1+\alpha i_s) + (1-\alpha)\sum_{s=t}^{T} i_s\right)$, where $\vec{\imath}$ is the vector of interest rates and α is the degree of the person's accuracy.[2] When $\alpha = 1$, the agent's perceptions are accurate and when $\alpha < 1$ the agent neglects compound interest to some degree. When $\alpha = 1$, perceptions become fully linear. The agent then perceives that the intertemporal budget constraint is $\sum_{s=t}^{T} c_s \cdot p_s(\vec{\imath}, T; \alpha) \leq \sum_{s=t}^{T} y_s \cdot p_s(\vec{\imath}, T; \alpha)$ where c_s is consumption in period s and y_s is income in period s. On the left-hand side of the inequality, EGB leads to a misperception regarding the tradeoffs of consumption over time. This will lead to an erroneous income effect and an erroneous substitution effect. On the right-hand side of the inequality, EGB leads to a misperception of one's wealth.

There are four main comparative static predictions that we test. First, compound interest will lead biased agents to increase consumption. When a person has an elasticity of intertemporal substitution (EIS) greater than 1 the person will overconsume in the presence of compound interest, but not when the agent faces a similar choice set with no compound interest. Second, shifting income to later periods in a way that is wealth-neutral, thereby having no effect on the intertemporal budget constraint, will increase consumption. A biased agent will misperceive a change in wealth when no such change occurred. Third, a biased agent will exhibit a form of dynamic inconsistency in which she updates her consumption plans in the direction of her current balance. If her balance is positive, she will be surprised by how quickly her assets grow, and if her balance is negative she will be surprised by how quickly her debts grow. To an unbiased agent, the displayed balance of all her accounts provides only redundant information since she can compute how these grow over time, however to a biased agent the balance provides a partial computation that shifts perceptions more closely in line with reality. Fourth, we predict a biased agent is sensitive to the period length used in the description. An economic problem described with shorter periods of interest (e.g. days) will exacerbate EGB relative to a description of the same economic problem with longer periods (e.g. years). This effect is predicted to be stronger when income is received with a delay.

Although there have been other laboratory consumption-savings experiments (Hey and Dardanoni, 1988; Anderhub, Güth, Müller and Strobel, 2000; Noussair and Matheny, 2000; Brown, Chua

[2] We present the parametric model of Levy and Tasoff (2015), but note that all predictions follow from their general formulation of EGB.

and Camerer, 2009; Meissner, 2015), with one (Johnson, Kotlikoff and Samuelson, 1987) finding behavior consistent with an underestimation of compounding interest, ours is the first to focus on the diverse manifestations of EGB as a function of the economic environment. None of the four theoretical predictions have yet been tested. The existing experimental evidence suggests that subjects are on the whole reasonably competent at maximizing lifecycle consumption in laboratory settings. The purpose of the present study is to explore specific ways in which behavior is predicted to exhibit a systematic bias.

The experimental results strongly confirm all four theoretical predictions. First, compound interest causes subjects to increase their consumption early in the lifecycle. Second, shifting income later in the lifecycle in a wealth-neutral way increases consumption. Third, subjects consumption plans differ dramatically from their period-by-period updated consumption choices reflecting the theoretically predicted dynamic inconsistency. Fourth, dividing a timespan into more periods results in greater consumption.

We consider alternative hypotheses. Subjects may use rules of thumb and these may be the mechanisms that generate the observed behavior. We consider two specific rules of thumb: consume a constant quantity every period, and never borrow against future income. We find evidence that a small minority conform to each rule. However, these rules of thumb fail to predict the main predictions of EGB. Thus these rules of thumb cannot be alternative explanations for the results on their own. Rather, the results imply that the effects of EGB are large enough to dominate the effects stemming from rules of thumb. We conclude that underestimation of compound growth has a significant effect — both statistically and materially — on behavior inside the lab, and therefore further exploration outside the lab may be warranted.

2 Conceptual Framework

We briefly present a model of EGB, and explain the main predictions of our paper. For concreteness, we present one possible parameterization of the bias, but note that all of the predictions we test follow from the more general model in Levy and Tasoff (2015) that does not rely on a specific functional form. Let $p_t(\vec{\imath}, T; \alpha)$ be the agent's perception of the period-T value of one dollar invested at time t. The parameter α is the degree of accuracy of the agent's perceptions. Let

$$p_t(\vec{\imath}, T; \alpha) = \prod_{s=t}^{T-1}(1 + \alpha i_s) + \sum_{s=t}^{T-1}(1 - \alpha) i_s \tag{1}$$

When $\alpha = 0$ the agent perceives growth as linear as if all interest were simple. When $\alpha = 1$ the subject correctly perceives growth as exponential at the true growth rate. Values of $\alpha \in (0, 1)$ generate perceptions that are in between linear and exponential growth. A fraction α of the interest compounds and a fraction $(1 - \alpha)$ is perceived as simple.

The agent wishes to maximize total utility over the lifecycle. For simplicity we assume that instantaneous utility is additively separable over time, $u(\cdot)$ is smooth, concave, increasing, and satisfies the Inada conditions $u'(0) = \infty$, $\lim_{c\to\infty} u'(c) = 0$; though any utility model, including one with dynamically inconsistent preferences, may be substituted. Total utility at t is given by $U_s(\vec{c}) = \sum_{s=t}^{T} \delta^t u(c_s)$, where δ is the discount factor and $\vec{c} \in \mathbb{R}^{T+1}$ is the consumption vector. The agent has no issue with her discounting since she is fully aware of her own preferences. Her evaluation of her utility is not necessarily driven by conscious calculations even if they can be represented as such. In contrast, the intertemporal budget constraint may require conscious calculation. Let $\vec{y} \in \mathbb{R}^{T+1}$ be the income vector, and $\vec{\imath} \in \mathbb{R}^{T}$ be vector of interest rates.[3] The budget constraint can be written as,

$$\sum_{s=0}^{T} c_s \cdot p_s(\vec{\imath}, T; 1) \leq \sum_{s=0}^{T} y_s \cdot p_s(\vec{\imath}, T; 1). \tag{2}$$

Since the agent misperceives exponential growth, however, she perceives the budget constraint as:

$$\sum_{s=0}^{T} \hat{c}_s \cdot p_s(\vec{\imath}, T; \alpha) \leq \sum_{s=0}^{T} y_s \cdot p_s(\vec{\imath}, T; \alpha) \tag{3}$$

where $\hat{c}_s$ is the agent's planned consumption at time s. The agent is of course also subject to the true budget constraint in (2), which may be thought of as being enforced by other market actors.

[3] The purpose of the paper is to focus on the implications of EGB. Therefore we assume a certain environment with known interest rates and income isolating the effects of EGB from the effects of uncertainty. Of course uncertainty will have additional implications on behavior that are explored extensively elsewhere. There are also interaction effects and this is the topic of a paper in progress.

In our experiments, we consider four effects of EGB. The first is the effect of compounding versus no compounding. People systematically underestimate the value of assets that compound. If a person invests a dollar, the person will underestimate the value of the invested dollar in T periods in the future, and this leads to both an income effect (reducing perceived purchasing power over the time frame) and a substitution effect (reducing the perceived purchasing power in the future relative to the purchasing power today). Consequently if a person's EIS > 1, then the substitution effect dominates and EGB results in overconsumption in the present. This effect is driven by misperceptions on the left-hand side of the budget constraint, which we refer to as the *price effect of exponential-growth bias.*

The second effect is driven by misperceptions on the right-hand side of the budget constraint. The person misperceives total wealth, which we refer to as the *wealth effect of exponential-growth bias.* Income y received at t will be perceived as expanding the budget less than the equivalent income $y(1+i)^{\tau}$ received at a later time $t+\tau$. Consequently, wealth-neutral shifts of income to later periods will cause a biased person to believe her budget is larger causing her to consume more immediately.

As the person moves through time, she will be surprised by how quickly her balance grows. If she has positive savings then each period she will experience a positive windfall. The third prediction is that the person will revise her consumption plans in the direction of her balance. Thus if she has positive savings she will increase consumption relative to previously planned consumption. The revision is the exact opposite if she has debts. With non-zero debts, in each period the person will be surprised by how quickly her debts grow and the person will revise her consumption downward from her previously planned consumption. Observationally, this theory predicts that consumption patterns may dramatically differ based on the degree of flexibility and commitment inherent in consumption. A person who can update her consumption every period would revise her consumption in the direction of her balance, while a person with a committed consumption plan would like to do so but cannot.

The fourth prediction involves the frequency of compounding. A biased agent's perceptions can be manipulated through the framing of period length. Because shorter periods increase the degree of compounding, misperceptions will be exacerbated the shorter the period length. Proposition A.1 in Appendix A formalizes this intuition and shows that the distortion in the agent's beliefs is bounded. In the limit, as the period length approaches zero the perceived total interest over the original time length never falls below $\ln(1+i_0)$, where i_0 is the original total interest. For instance, a person offered a loan at 330% annual interest, will perceive the equivalent interest rate as much lower if it is framed as 0.4% a day of compound interest. As the period length approaches zero (i.e. continuous compounding) the periodic rate becomes infinitesimal but the number of periods explodes. In the limit, the biased agent will never perceive this loan as less than $\ln(1+3.3) = 146\%$ annually, regardless of the frame. This is a considerable difference compared to the original 330%, but clearly the extent of the framing manipulation is limited.

This proposition gives some insight into the types of frames often chosen in the market. Very large annual interest rates for loans (e.g. payday loans) will look much more favorable when presented as monthly, weekly, or even daily rates. However, the effect only goes so far. A firm cannot use this trick to make quadruple-digit annual interest rates appear to be double-digit or less. Nor can they greatly distort low interest rates since $\ln(1+i_0) \approx i$ when i_0 is small. Moreover, because it magnifies the wealth effect of EGB, this framing effect is predicted to have a larger impact on behavior when the budget constraint includes delayed income.

3 Experimental Tests of Exponential-Growth Bias

We turn to the lab to test for the theoretical predictions in the context of a well-controlled consumption-savings problem with random assignment of the income and interest vectors. We use an induced valuation design (Smith, 1976) in which we effectively give subjects a utility function to maximize and pay them based on their performance. The principle advantage of this approach is that it is a way to assign the preferences that our models assume, allowing for a test of theory under the assumption that subjects' true preferences are increasing in earnings. The method therefore isolates the perceptual aspects of the economic problem from individual heterogeneity in preferences. Our purpose is not to test whether people have additively separable concave instantaneous utility functions, the form of utility function we happen to use. Nor is it to test people's ability to solve consumption-savings problems more generally. Maximizing an intertemporal utility function is challenging even when interest rates do not compound. The purpose is to test *comparative static predictions* specific to the model: compared to a baseline behavior, do changes in the economic environment or framing of the problem change behavior in the ways that the theory predicts? The two assumptions required for our approach to be valid are (1) subjects' true preferences are increasing in their experimental earnings, and (2) any non-EGB optimization failures are present in the baseline (non-compounding) condition.

Given the control afforded by the laboratory, the model makes strong theoretical predictions. Still, obtaining the theoretical predictions in this experiment is not tautological. There are other cognitive factors in consumption-savings decisions that may bias behavior in other ways. Consumption-savings problems involve mathematical operations and estimations other than exponential growth, like arithmetic and the concept of consumption smoothing. If an agent has severe deficiencies in arithmetic this could in principle dominate any effects stemming from EGB, in which case the departures from optimality would just be noise.

One potential concern with a lab experiment, however, is that it may not capture opportunities for learning. The opportunities for learning when it comes to retirement savings are limited, but not non-existent. There are several reasons to believe that learning regarding EGB is particularly limited. First, even with feedback about the outcome of debts and assets, consumers may not know the source of the inaccuracy of their predictions. A consumer may not realize that the error stems from her own perceptual bias, and may instead infer that it stems from some other feature in the environment. Second, feedback will rarely be contemporaneous with choice, and the psychology literature suggests that learning is rare without immediate feedback (Kahneman and Klein, 2009). Our design, in fact, allows us to test for the effect of immediate feedback in a controlled environment, and we find no evidence of learning over the course of the experiment.

3.1 Experiment 1: Compounding, Timing, and Current Balance

3.1.1 Design

Subjects were given a task to maximize an explicitly stated utility function subject to an intertemporal budget constraint. By using the induced-value approach we can generate very clear theoretical predictions on behavior that would otherwise be impossible when people's preferences and constraints are unobserved. In addition to the explicit instantaneous utility function, subjects were assigned an income vector and interest rate vector, and had to choose how much income to consume each period. Savings and debts would accrue interest as a function of the interest rates.[4]

Since the purpose of the experiment was not to examine whether people are capable of maximizing arbitrary utility functions *per se*, we provided as many features as we could to help subjects understand the relationship between consumption and utility in the task, and hence their payment. We emphasized the concept of marginal utility and illustrated with several examples. We also provided on all screens the utility function itself, a graph of the utility function, and a calculator which took consumption as an input and gave utility as an output. Furthermore, subjects went through three training rounds with feedback about the optimal plan and how their choices fared relative to it. These training rounds each had a single positive interest rate, with the interest rate set to zero in other periods. This familiarized subjects with the task and helped train them on the basics of utility maximization without directly training them on compounding.

The main task consisted of four consumption problems: combinations of an income vector, interest rate vector, and number of periods. A consumption problem of length T required consumption choices in period 0 through $(T-1)$, with any residual consumption allocated to period T

[4] Given the abstractness and unfamiliarity of such a problem to a typical subject in our sample, we made an explicit design choice to use functional labels in which we framed the problem as a game to feed a "digital dog". Subjects were allotted "bucks" (income) on various "days" (periods) which could be used to purchase "dog food" (consumption). The dog food generated "wags" of the tail (utils). The object to be maximized was the sum of "tail wags" (utils) over the course of all the days. By using a familiar context we expressed the basic features in an easily digestible manner and kept subjects engaged with the task. A norm applied to many laboratory economic experiments is an avoidance of loaded labels for the legitimate concern that they may reduce the relative salience of monetary incentives thereby confounding the interpretation of the results. There are two reasons why this is not a concern here. The first is that all of our predictions are comparative static predictions. Since all conditions use the same labels, to the extent that the labels may have any effect, they appear under the control condition and are differenced out using the control group behavior. The second reason is that the labels, while facilitating comprehension and engagement, contain no connotations that are relevant to consumption or savings. If anything this should increase saliency.

automatically. Subjects' earnings were strictly increasing in the achieved utility, and all payments were made at the conclusion of the experiment. Thus regardless of subjects' own utility and time preferences, a preference for money means their optimal actions were to attempt to optimize the consumption problem as stated in the experiment.[5]

There were four consumption problems, each six periods long: $t \in \{0, \ldots, 5\}$. The income and interest vectors are displayed in Table 1. The order in which the consumption problems were presented was random. Consumption Problem N (non-compounding) has no compounding interest and so EGB is irrelevant. Behavior on this problem established a benchmark level of imprecision and bias that is unrelated to EGB. The three other consumption problems all have compounding interest. Consumption Problem S (income at Start), Consumption Problem M (income at Middle), and Consumption Problem E (income at End) all have the same interest vector. The only difference between these consumption problems is the time at which income is received.

Table 1: Consumption Problems in Experiment 1

Consumption Problem	Income Vector	Interest Vector
N: No compounding	$\langle 100, 0, 0, 0, 0, 0 \rangle$	$\langle 0, 0, 100\%, 0, 0 \rangle$
S: Income at Start	$\langle 100, 0, 0, 0, 0, 0 \rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\% \rangle$
M: Income at Middle	$\langle 0, 0, 100, 100, 0, 0 \rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\% \rangle$
E: Income at End	$\langle 0, 0, 0, 0, 0, 500 \rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\% \rangle$

We use high interest rates for statistical power in decisions involving just a few periods. Lower interest rates would require a longer experiment or more subjects. If one interprets each period in the consumption problem as a decade, however, then the annual interest rate in Consumption Problems S, M, and E corresponds to 5.75% per year, which is near real historical long-run stock market returns. One may worry that there is something particular about large interest rates that makes behavior structurally different from lower rates. However, the theory actually predicts that

[5] Because the experiment involves no uncertainty, risk preferences play no role for a neoclassical decision maker participating in the experiment. A participant who is unbiased but does not answer mathematical questions perfectly (perhaps due to inability, trembles, opportunity costs of computation, etc.) and is risk averse might wish to take a "safe" conservative approach to the problem. However, without knowing more concretely the distribution of one's errors, the "safe" conservative approach is not well defined.

economically equivalent formalizations of the same problem with more periods and lower rates will result in *greater* distortions of behavior, not less. We address exactly this point later in Experiment 2 and find that the results are consistent with theory: distortions are stronger with *lower* rates.

Experiment 1 assigned each subject to one of two arms (see Figure 1). In the dynamic arm, consumption was chosen sequentially each period. Subjects first chose the consumption in period 0. After submitting their answer they were informed of their current savings (or debt). They then chose consumption for the next period, and so on, receiving updated information about their current balance.[6] If chosen consumption in a given period would exceed the true budget, then the computer bounded consumption such that the budget was fully expended and set consumption in all remaining periods to zero. Similarly, any remaining wealth in the last period was automatically used for consumption. Thus the dynamic arm provides feedback about the current balance each period. Such information is redundant to a classical decision maker but is predicted to affect behavior for a biased decision maker.

Figure 1: Experiment 1 – Design

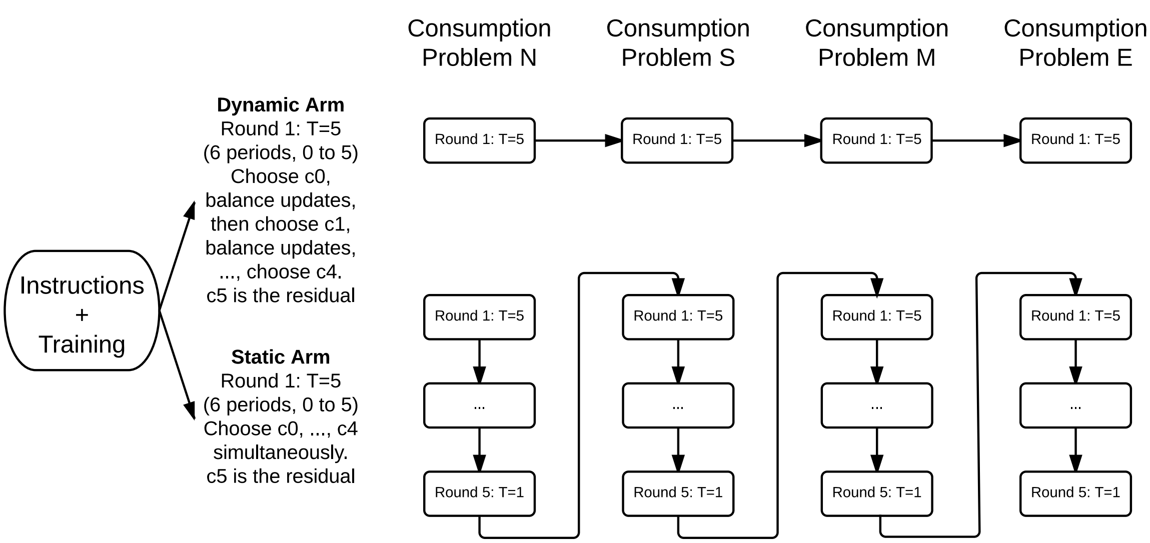

Notes: Each round of length T has $T+1$ periods. Subjects in the static arm choose consumption vectors for a total of 20 rounds. Subjects in the dynamic arm choose consumption vectors with feedback on their current balance for a total of 4 rounds. The order of Consumption Problems, and rounds within a consumption problem in the static arm, was randomized by subject.

In contrast to the dynamic arm, in the static arm subjects stated their consumption each period

[6] Figure D.13 displays the user interface subjects saw in the dynamic arm in Online Appendix D.

without any feedback. Effectively, a subject in the static arm chose their consumption plan for all periods simultaneously.[7] As in the dynamic arm, subjects could not exceed their budget. Subjects were informed that the computer would implement their consumption plan sequentially until all resources were exhausted, and any remaining money was automatically used for consumption in the final period $T = 5$. Thus behavior in the static arm elicits the subject's perceived optimal consumption plan at $t = 0$. Comparing behavior between the dynamic arm and the static arm allows one to test how actual consumption diverges from consumption plans.

In the static arm, each consumption problem was divided into 5 rounds. The first round of each consumption problem is described in Table 1. In subsequent rounds, each consumption problem has its first period progressively truncated. For example, in the fourth round of consumption problem S, subjects made a 3-period plan given $\vec{\imath} = \langle 75\%, 75\% \rangle$ and $\vec{y} = \langle 100, 0, 0 \rangle$. In every round of Consumption Problem N and S, the consumer received a lump sum in their first period; in every round of Consumption Problem E, the lump sum was received in the last period; and for Consumption Problem M the timing of income relative to the initial period varied as the first periods were progressively eliminated. We thus observe the full profiles that the agent expected to consume, period by period.

The instantaneous utility function given to subjects was $u(x) = x^{\frac{1}{2}}$. We chose this as the utility function for two reasons. Given that this is a constant relative risk aversion utility function (CRRA) with an elasticity of intertemporal substitution of 2, the theory predicts overconsumption for any income vector.[8] Furthermore, because CRRA utility functions are homothetic, the value of the income stream should not affect the proportion of the budget spent in each period for any agent, whether biased or unbiased. This allows for direct comparisons of behavior across different consumption problems without the confound of nonlinear wealth effects. We kept the utility function constant throughout the experiment, including the training rounds, in order to focus just on the effects of manipulating the budget constraint. While EGB does make predictions

[7]Figure D.16 in Online Appendix D displays the user interface subjects see in the static arm. The table at the top of the screen indicates the income and interest vectors. The entries below allow the subject to input their consumption plan. Below that is a calculator that allows the subject to compute how consumption relates to utility, and a table and graph that further explicates this relationship.

[8]See Levy and Tasoff (2015), Proposition 3 which states that a biased agent with EIS > 1 will overconsume in $t = 0$ for any income vector.

regarding utility, for example changing the risk-aversion parameter within the CRRA family, that is not our focus here.

3.1.2 Sample and Incentives

This laboratory experiment was conducted at the UCLA California Social Sciences Laboratory (CASSEL). Subjects were not provided with any tools, and calculators/cell phones were expressly forbidden. Subjects could earn up to \$35 based on the quality of their responses, in addition to a \$5 participation fee. After completing the experiment, one round was chosen at random by the computer. Subjects were paid based on how much additional utility their plan earned above the minimum achievable utility, as compared to how much additional utility would be achieved by the optimal plan. Letting u_a be the achieved utility, u_o the optimal utility, and u_m the minimum possible utility, a subject's additional payment was given by $\$35 - 35 \cdot [1 - (u_a - u_m)/(u_o - u_m)]^{\frac{1}{2}}$. This payment rule provides strong incentives for accuracy, as returns are increasing as one approaches the optimal utility. Subjects could click on a link to see this formula, but it was emphasized that their payment was increasing in the total utility. The mean incentive payment was \$18.46: dynamic subjects averaged \$17.68, while static subjects (who could be paid for rounds with fewer periods) averaged \$19.38.

3.1.3 Hypotheses

Given that the EIS>1, the model predicts that $c_t > c_t^*$, where c_t^* is the optimal consumption in any period $t < (T-1)$.[9] However, there may be other biases in people's problem-solving abilities that may cause them to overconsume. Thus we use the c_t in Consumption Problem N, which has no compounding, as our baseline. An appropriate comparison would be to use log-normalized consumption, comparing $\ln(c_t/c_t^*)$ since this weights proportional errors in a symmetric way.

Hypothesis 1 (Overconsumption) *Normalized overconsumption, given by normalized* $\ln(c_t/c_t^*)$*, in the consumption problems with compounding will be greater than overconsumption in Consump-*

[9] This is taking the consumption history as given. Overconsumption in earlier periods necessarily leads to underconsumption in later periods, compared to the ex-ante optimal plan. Conditional on the actual resources available in the period, however, a biased person will overconsume whenever there are at least two remaining periods.

tion Problem N.

Second, shifting income to later periods in a wealth preserving way should increase log-normalized consumption for a biased agent.[10] An unbiased agent in contrast would be unaffected by such shifts.[11] Consider the RHS of (2), and let $w \equiv \sum_{s=t}^{T} y_s \cdot p(\vec{\imath}, s; 1)$ for a given $\vec{y}$ and $\vec{\imath}$. Given a fixed w, the exact values of $\vec{y}$ and $\vec{\imath}$ are irrelevant to an unbiased agent. The biased agent responds not to w but to $\sum_{s=t}^{T} y_s \cdot p(\vec{\imath}, s; \alpha)$, which will vary with $\vec{y}$ and $\vec{\imath}$ for fixed w.

Hypothesis 2 (Deferred Income Increases Consumption) *Normalized overconsumption, given by* $\ln(c_t/c_t^*)$*, should be greater for Consumption Problem E than M, and greater for Consumption Problem M than Consumption Problem S.*

Third, consumption plans should revise in the direction of the displayed balance.[12] A biased agent will be surprised by how quickly her savings or debts have grown and revise her consumption in the direction of her balance. Thus a person in the dynamic arm who could revise her plans should exhibit behavior that is different from a person in the static arm, in which the initial plans are implemented up to a feasibility constraint. More specifically, in Consumption Problem S the balance is always positive since income is received lump sum in $t = 0$ and this initial income finances all future consumption. In contrast, in Consumption Problem E the balance is always negative since income is received lump sum in $t = T = 5$, and so debts are required in all periods $t < T$ to finance consumption. Thus for Consumption Problem S, normalized consumption in $0 < t < 5$ should be greater in the dynamic arm than the in the static arm (for the full length $T = 5$ round), but the opposite should be true in Consumption Problem E.

Hypothesis 3 (Dynamic Inconsistency) *Normalized consumption, given by* $\ln(c_t/W_0)$*, in* $0 < t < 5$ *should be greater for the dynamic arm than the static arm in Consumption Problem S, and should be lesser for the dynamic arm than the static arm in Consumption Problem E.*

[10] Below, for consistency with Hypothesis 1, we formulate this in terms of log-normalized *over*consumption; but it is equivalent to using log-normalized consumption.

[11] This follows from Proposition 1 in Levy and Tasoff (2015)

[12] This follows from Proposition 4 in Levy and Tasoff (2015).

Hypothesis 1 is driven by the price effect of EGB and hence relies on subjects' recognition of the intertemporal tradeoffs of consumption given the induced-value utility function. Since this may be quite challenging we expect that this hypothesis is less likely to obtain than Hypotheses 2 and 3, which are driven largely by the wealth effect of EGB. As long as subjects increase consumption c_0 when perceived wealth increases, Hypotheses 2 and 3 should obtain. This is true even if subjects are not trying to optimize at all but rather use various rules of thumb. We explore this in greater depth in Section 4.

3.1.4 Results

The aggregate behavior of subjects is shown in Figure 2, where we plot the mean consumption paths achieved by static and dynamic subjects across consumption problems as a proportion of the initial present value of income, as well as the optimal consumption profile in Consumption Problems S, M, and E.[13] For the static subjects in panel (a), this requires constraining their stated consumption plans to be feasible by setting consumption to zero in periods after all resources are exhausted. Consumption Problem N, without compounding, has the lowest initial consumption, even though the optimal consumption would be slightly higher than for the other Consumption problems. Notably, people on average do react to the interest rate in the correct direction, planning to consume more in later periods. However, subjects' smooth consumption somewhat too much in N. We therefore use this consumption problem to control for subjects' performance in the absence of exponential growth, which raises the hurdle for our tests of EGB.

Because Consumption Problems S, M, and E all have the same optimal consumption path, an unbiased economic agent should exhibit the exact same consumption profile across these problems. Examination of Figure 2 instead shows systematic differences in the direction predicted by Hypothesis 2. As we compare consumption between Consumption Problems S and E we find that the more delayed the income the greater the initial consumption. The difference appears much less pronounced when comparing S to M. The consumption patterns also clearly demonstrate dynamic inconsistency. As we compare S across fixed plans, panel (a), to flexibility, panel (b), we find that

[13] We do not superimpose the optimal consumption paths for Consumption Problem N given the already complex figure, but note for reference that the optimal path would be $\langle -2.2, -2.2, -2.2, -0.8, -0.8\rangle$

people consume much more than they initially planned. People are surprised by how quickly their savings grow and so they shift their consumption up. Consumption in E across the static and dynamic arms appears similar but this is largely due to the fact that many subjects exhaust their full budget within the first couple of periods and thus have little room to shift their consumption down. In the remainder of this section we show that dynamic subjects shift their consumption downwards relative to static subjects. People are surprised by how quickly their debts grow and so they shift their consumption down.

Figure 2: Consumption Paths

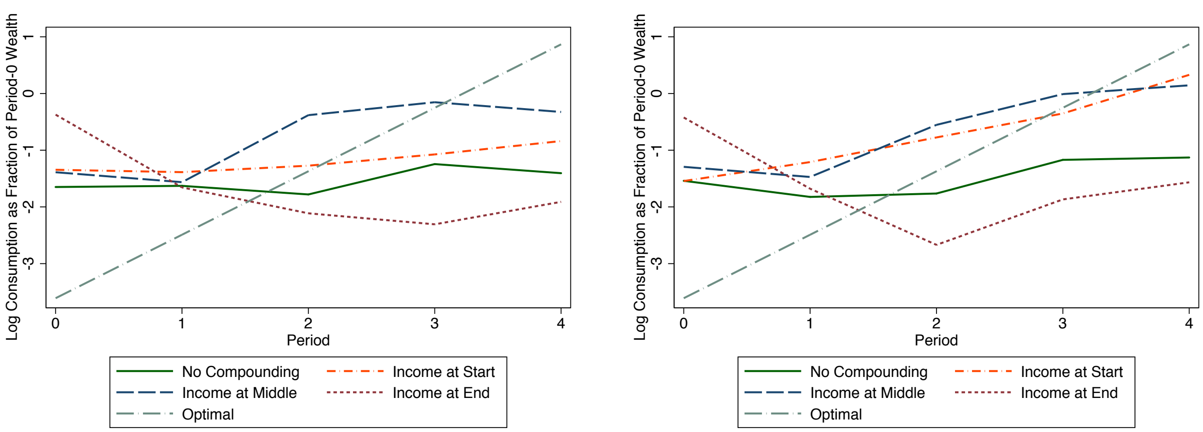

(a) Static Subjects

(b) Dynamic Subjects

Notes: Mean log consumption paths, by subject type. Panel (a) shows log consumption for subjects in the static arm of the experiment; panel (b), consumption of subjects in the dynamic arm of the experiment. Consumption paths are constrained to be feasible under the true budget constraint by setting consumption to zero in periods after all resources are exhausted.

We begin by testing Hypothesis 1, that compounding leads to overconsumption in early periods relative to the optimal spending path. We first focus on consumption in the initial period for simplicity and to eliminate the effect of feedback in later periods for subjects in the dynamic arm. Figure 3 plots the distribution of the natural logarithm of the ratio of subjects' period-0 consumption to the optimal period-0 consumption for that consumption problem: $\ln(c_0/c_0^*)$.[14] This provides a simple measure of over-consumption: subjects overconsume relative to the optimum when this variable is positive, and underconsume when it is negative. For simplicity, we pool static

[14] There are 38 subjects in the static arm each generating a consumption vector for each of 5 rounds for each of 4 consumption problems, and 44 subjects in the dynamic arm each generating a consumption vector for each of 4 consumption problems. We drop the observations for which $c_0 = 0$, and address zeroes in Online Appendix B.

and dynamic subjects' answers, restricting to the $T = 5$ rounds for static subjects for comparability. In panel (a) of Figure 3, we plot the distribution of the log-ratio for Consumption Problem N, which has only a single nonzero interest rate and therefore EGB does not predict overconsumption. While the median value is higher than zero, the overall distribution indicates only a minor tendency to overconsume in period 0. Any overconsumption in panel (a), however, is dwarfed by the overconsumption shown in panels (b)–(c), where we present the distribution of the log-ratio for Consumption Problems S, M, and E, which have more than one period of non-zero interest. The distribution is clearly shifted to the right in each successive panel (the median from the preceding panel is indicated to facilitate comparisons), providing strong evidence both that compounding increases overconsumption and that this effect is further exacerbated as income is delayed. In all three panels, both the mean and median are significantly more positive than without compounding.

We unpack the sources of overconsumption in Table 2. We now expand the analysis to consider $\ln(c_t/c_t^*)$ across all periods (where c_t^* is the optimal consumption in that period conditional on the consumption history), and use OLS specifications in all columns. Columns (1) and (2) of Table 2 directly test Hypothesis 1 by comparing the initial consumption in consumption problem N (which had no compounding) against consumption problem S (which featured compounding, but like N had all income arrive immediately). Column (1) shows that, averaged across all rounds, the inclusion of multiple non-zero interest rates more than doubled subjects' initial consumption choices as a percentage of the optimal level of consumption. Column (2) tests the prediction that this effect ought to be driven by choices with the longest horizons, including the period number and interacting this with the compounding dummy. For $t = 0$, the effect of compounding is large and highly significant. The effect falls rapidly, however, and the period advances (and the number of remaining periods therefore falls). The interaction of $Compounding \times Period$ is significantly negative, and implies that there is no remaining effect of compounding by period 4 — as predicted by the model.[15] There is also a significant, though very small, direct effect of the period on overconsumption, suggesting that effects other than EGB may be contributing to behavior in minor ways.

[15] Because the decision in period 4 involves just two periods, 4 and 5, there is no compounding and therefore ought to be no difference between the lifecycles.

Figure 3: Overconsumption

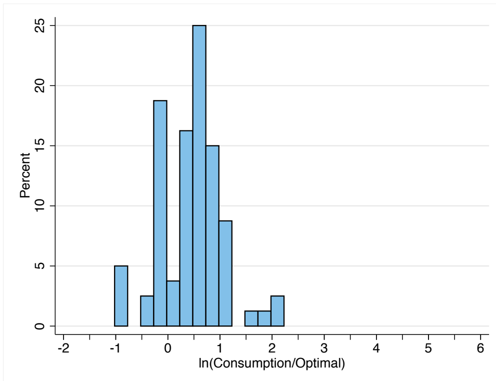

(a) No Compounding

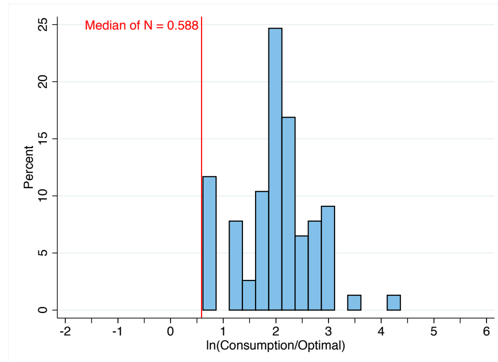

(b) Income at Start

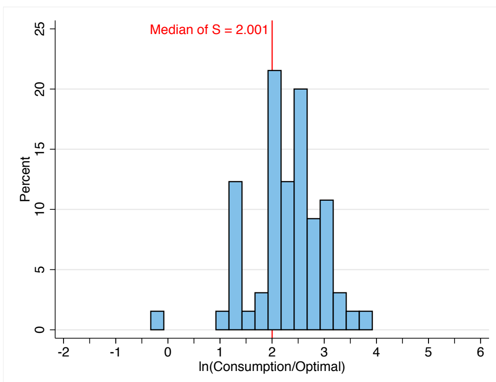

(c) Income at Middle

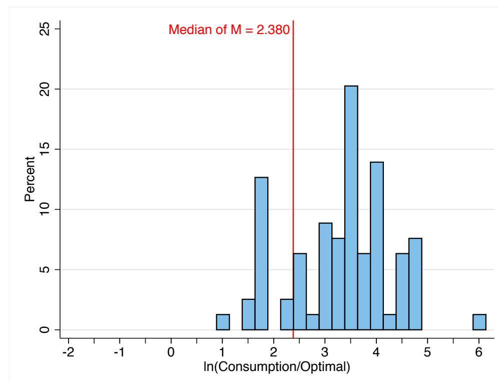

(d) Income At End

Notes: Figures show the empirical distribution of $\ln(c_0/c_0^*)$: the natural logarithm of the ratio of subjects' choice of consumption for period 0 to the optimal consumption choice. Panel (a) restricts to Consumption Problem N, which did not feature compounding interest rates and is therefore not affected by EGB. Panel (b)–(d) plot Consumption Problems S, M, and E, respectively, which do feature compounding interest rates. The median of the previous panel is plotted for panels (b)–(d).

We repeat the analysis in Columns (3) and (4), including the remaining consumption problems M and E in the sample. Compared against consumption problem N, the three problems featuring compounding exhibit significantly greater initial consumption choices. The coefficients are greater than those in the preceding two columns, as the added observations include not just the price effect from consumption problem S, but also include the wealth effect of EGB. We will return to this comparison later, when testing Hypothesis 2.

Finally, Column (5) of Table 2 shows whether subjects learn to correct their bias over the course of the experiment. We restrict the sample to those consumption problems with multiple non-zero interest rates, and include all choices for both static and dynamic subjects. Because the

Table 2: Effect of Compounding on Overconsumption

	(1) N vs S	(2) N vs S	(3) All	(4) All	(5) All EGB
Compounding	0.681***	1.484***	1.067***	2.056***	
	(0.069)	(0.087)	(0.062)	(0.079)	
Period		-0.080***		-0.059***	-0.668***
		(0.024)		(0.022)	(0.040)
Compounding X Period		-0.432***		-0.603***	
		(0.034)		(0.028)	
Dynamic					-0.365
					(0.385)
Question					-0.005
					(0.013)
Question X Dynamic					-0.019
					(0.077)
Question X Period					0.004
					(0.003)
Constant	0.215***	0.400***	0.236***	0.380***	2.682***
	(0.043)	(0.084)	(0.047)	(0.081)	(0.168)
N	867	867	1,433	1,433	933
Subjects	82	82	82	82	82

Notes: Dependent variable is $\ln(c_t/c_t^*)$; all columns are OLS regressions. Columns 1 and 2 restrict to Consumption Problems N and S (i.e. testing the price effect of EGB), and Columns 3-6 include all Consumption Problems. Extra periods is the number periods by which T > 1. Question refers to the order in which the subject answered the question; subjects in the static arm answered 20 questions excluding the training, while subjects in the dynamic arm answered 4 excluding training. Standard errors clustered by subject. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

consumption problems were presented in a random order, as were the multiple rounds within a consumption problem for static subjects, we can include a variable *Question* which records the order in which a subject actually answered that particular round. We would not expect any learning among static subjects, since they received no feedback, but it is possible that dynamic subjects would learn about their EGB and revise their behavior accordingly. This would yield a negative coefficient on the $Question \times Dynamic$ interaction. However, neither *Question* nor any of its interaction terms are significantly different from zero, suggesting that subjects did not learn to perform better over the course of the experiment. Given the limited feedback and moderate stakes here, however, it is plausible that this understates the potential for learning in real life.

The preceding results showed significant overconsumption, but we acknowledge that EGB is not the only possible explanation. For example, a simple rule-of-thumb to smooth consumption equally across periods would also produce a difference between Consumption Problem N and Consumption Problems S, M, and E given the steep optimal consumption profile in the latter three. We show this in Section 4.

From Hypothesis 2 we predict that delaying the timing of income will exacerbate overconsumption of EGB individuals. The simplest test of this prediction is to compare the degree of overconsumption in Consumption Problem S, where all income is received immediately, to the degree of overconsumption in Consumption Problem E, where all income is received in period T. Hypothesis 2 says that overconsumption will be greater when all income is received at the end of the consumption problem. Over-smoothing plays no role in this comparison, since the optimal consumption profiles are identical, allowing us to focus entirely on differences in subjects' perceived lifetime wealth (the wealth effect of EGB).

Table 3: Effect of Delayed Income on Overconsumption

	(1) All	(2) All	(3) T=5	(4) T=5
Income at End	1.216*** (0.192)	1.570*** (0.132)	1.316*** (0.124)	1.280*** (0.184)
Income at End X Period		-0.471*** (0.088)		
Period		-0.512*** (0.036)		
Income at Middle			0.288** (0.113)	0.125 (0.138)
Dynamic				-0.212 (0.166)
Income at End X Dynamic				0.065 (0.250)
Income at Middle X Dynamic				0.306 (0.220)
Constant	0.896*** (0.072)	1.885*** (0.082)	2.009*** (0.083)	2.122*** (0.124)
N	647	647	221	221
Subjects	82	82	82	82

Notes: Dependent variable is $\ln(c_t/c_t^*)$; all columns are OLS regressions. For columns 1-2 sample comprises those lifecycles where the entire endowment is received either in period 0 (Income At Start) or period T (Income At End). For columns 3-4 all lifecycles of length T = 5 are used. Standard errors clustered by subject. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

In Table 3, we regress our measure of overconsumption, $\ln(c_t/c_t^*)$, on an indicator for whether income is delayed, using rounds of all lengths from Consumption Problems S and E and pooling dynamic and static subjects together.[16] In Column (1), the estimate of 1.216 on the delayed income dummy variable is highly significant and confirms the prediction of the theory. Delaying income causes an approximate increase in consumption by about 250%. Column (2) interacts the delayed income dummy with the current period number. The coefficient of -0.512 on *Period* reiterates the finding from Table 2 that overconsumption is increasing in the number of remaining periods (as both consumption problems S and E feature compounding). The interaction of the delayed income dummy with the period number is also negative at -0.451 and highly significant, confirming that the wealth effect is also larger when income is delayed by more.

In Columns (3) and (4) of Table 3, we restrict the sample to those rounds with T=5, and add Consumption Problem M where all the income is received in the middle periods. This ensures that both static and dynamic subjects have two observations each, and confirms that the effect is largest in the longest rounds. This is a very strong test of Hypothesis 2: the coefficient on *Income at Middle* should be positive and less than the coefficient on *Income at End*, as the perception of the value of income received three periods in the future is less distorted by EGB than that of income received five periods in the future. We confirm that overconsumption is indeed greater in consumption problem M than in S ($p < 0.05$), and greater in consumption problem E than in M ($p < 0.01$). The last column confirms that there is no systematic difference between subjects in the two arms: neither the dynamic dummy nor its interaction with the delayed income dummy enters significantly.

Finally, we address our prediction of dynamic inconsistency from Hypothesis 3. Once again it is easiest to test this hypothesis by restricting attention to Consumption Problems S and E. Because Consumption Problem S has all income received in period t=0, subjects will always carry a weakly positive balance. Hypothesis 3 says that dynamic subjects will be surprised at how quickly their balance grows, and will revise their consumption upwards relative to their initial plan in later

[16] There are 38 subjects in the static arm each generating a consumption vector for each of 4 rounds (the shortest 5th round does not have compounding and is excluded) for each of 2 consumption problems (S and E). A truncated 3-period round of Consumption Problem M also has all income at the start so these are also included. In the dynamic arm there are 44 subjects each generating a consumption vector for each of 2 consumption problems

periods. Conversely, subjects in Consumption Problem E must borrow against their income from the final period to finance early consumption, and should revise their consumption downwards in later periods. We can test these predictions by using static subjects' plans from the full length $T = 5$ rounds as proxies for dynamic subjects' initial plans. That is, we observe dynamic subjects' initial choice only for c_0, but not their planned choices for c_1 through c_4 (and the implied c_5 as a residual) that led them to believe that this choice of c_0 was optimal. Static subjects, however, were asked for a full consumption plan, c_0–c_4 with c_5 as a residual, without the opportunity for any revisions – and did so for each of 5 rounds per consumption problem. We can therefore compare dynamic subjects' *actual* consumption to the committed consumption plans of static subjects in order to measure the effect of feedback about the remaining budget.

We use OLS regressions to perform these comparisons in Table 4. As we are interested in deviations from plans, the dependent variable is now the natural logarithm of period-t consumption as a fraction of consumption problem wealth, W_0.[17] Column (1) restricts the sample to Consumption Problem S in which all the income is received at the start. The model regresses normalized consumption on: an indicator for dynamic subjects, an indicator that takes the value one if the period is not $t = 0$, and an interaction term between the two. The coefficient on the interaction is highly significant, confirming the prediction of Hypothesis 3. Consumption is about 91% higher in the dynamic arm over the static arm in periods 1-4. Reassuringly, the groups are statistically indistinguishable initially, indicated by the insignificant coefficient on *Dynamic*. Column (2) interacts *Dynamic* with *Period*. The Dynamic indicator is again insignificant in Column (2), as dynamic and static subjects do not differ in their plans for c_0. The coefficient on the $Dynamic \times Period$ interaction is large and positive, indicating that dynamic subjects' actual consumption choices in later periods are substantially higher than static subjects' consumption plans. In each subsequent period, dynamic consumption increases relative to the static consumption of the same period by about 29% (e.g. after 2 periods the dynamic consumption has grown $e^{0.257\times 2} \approx 67\%$ relative to the static consumption level). The opposite pattern emerges in Column (3), where the $Dynamic \times (Period > 0)$ interaction is again large and significant, but now negative – indicating that dynamic subjects'

[17] Each of the 82 subjects makes a decision for each of 5 periods (i.e. periods 0–4) for a given consumption problem Observations with $c_0 = 0$ are again dropped.

actual consumption choices are revised substantially downwards from static subjects' plans at t=0. Here, dynamic consumption decreases relative to the static consumption in periods 1-4 by about 44%. In Column (4) we interact $Dynamic \times Period$ and find that dynamic consumption decreases relative to the static consumption of the same period by about 21%.

Strikingly, this revision by dynamic subjects is no more than we would have expected from static subjects' choices in shorter rounds. One may expect a difference if the feedback in the dynamic arm causes learning *within* the consumption problem. We test this by using the consumption plans of the static arm to construct a counterfactual consumption vector that represents what that subject *would* have done if her consumption plans were non-binding. Denote this by c^s. We construct this "simulated consumption" vector by using the subject's choice of immediate consumption in each round as if it were the continuation of her choices in longer rounds. We begin with the actual choice of c_0 from the full $T = 5$ round as c_0^s. We then calculate the proportion of wealth actually consumed as c_0 from the $T = 4$ round, and adjust it by the wealth which would remain after consuming c_0^s to produce c_1^s. This step relies on the homotheticity of the objective function. We then repeat for the $T = 3$ round to produce c_2^s, etc., to generate the full vector of simulated consumption.[18] We can only use this for consumption problems that have all the income received at the start; otherwise the comparison would require us to know the subject's exact degree of EGB to infer their perceived lifetime wealth. We therefore only simulate for Consumption Problems N and S.

While we can only simulate how static subjects would have re-optimized in Consumption Problem S, this simulated static consumption without commitment is indistinguishable from the dynamic consumption. Using the simulated static consumption without commitment and dynamic consumption data, we can regress consumption on a set of time dummies interacted with the *Dynamic* dummy. We do not reject that the interacted dummies are jointly different from zero (F(5, 81) = 0.41, p=0.840). This implies that feedback on one's balance, as in the dynamic arm,

[18] For example, suppose in the $T = 5$ round with $\vec{y} = \langle 100, 0, 0, 0, 0, 0\rangle$ and $\vec{\imath} = \langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$, a subject chooses $\vec{c} = \langle 2, 3, 4, 6, 18\rangle$, with c_5 being the residual, and in the $T = 4$ round with $\vec{y} = \langle 100, 0, 0, 0, 0\rangle$ and $\vec{\imath} = \langle 75\%, 75\%, 75\%, 75\%\rangle$, a subject chooses $\vec{c} = \langle 5, 7, 12, 21\rangle$, with c_4 being the residual. Since in period 0 of the $T = 5$ round the person consumed 2% of wealth, $c_0 = 0.02w$ where w represents the $t = 0$ value of wealth. We use the c_0 behavior of the $T = 4$ round to compute c_1^s which was 5% of wealth. Remaining wealth in the simulated consumption path $0.98w(1 + i_0) = 98w(1 + 0.75) = 1.715w$, and c_1^s is 5% of this which is 0.08575w. A similar process is used to compute c_2^s through c_5^s. These values are all then normalized by the total wealth W_0.

has no impact on consumption beyond the prediction of Hypothesis 3. Whereas the previous results from Table 2 show that there is no learning across the experiment (one could not reject the null that the coefficient on *Question* was zero), this result shows that there is also no learning *within* a consumption problem.

Table 4: Dynamic vs. Static Plans

	(1) Income at Start	(2) Income at Start	(3) Income at End	(4) Income at End
Dynamic	-0.212	-0.211	-0.147	-0.138
	(0.165)	(0.155)	(0.228)	(0.250)
Period> 0	0.208		0.165	
	(0.142)		(0.104)	
Dynamic X Period> 0	0.649***		-0.575***	
	(0.244)		(0.178)	
Period		0.109**		0.085**
		(0.048)		(0.033)
Dynamic X Period		0.257***		-0.242***
		(0.076)		(0.077)
Constant	-1.488***	-1.540***	-0.208	-0.245
	(0.123)	(0.117)	(0.180)	(0.181)
N	398	398	365	365
Subjects	81	81	80	80

Notes: Dependent variable is $\ln(c_t/W_0)$, i.e. consumption normalized by actual starting wealth. The sample comprises the compounding T=5 consumption problems with all income received either in period 0 or in period 5. One subject planned consumption = 0 for periods 1-4 when income at start; two (different) subject planned consumption = 0 for periods 1-4 when income at end. Since the ln(0) is undefined these subjects are dropped. Standard errors clustered by subject. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Finally, we note that to earn a reasonable payment, subjects simply had to save enough so that they would have nonzero consumption every period. It was on this front that both dynamic and static subjects failed. Largely due to their c_0 choice, 36.0% of static subjects and 34.8% of dynamic subjects' $T = 5$ plans from Consumption Problems S, M, and E led them to exhaust all their resources early (i.e. $c_5 = 0$). This rose to 73.7% and 68.2% in Consumption Problem E, of which 63.2% and 59.1% respectively occurred by the end of the second period! While dynamic subjects realized their debts were growing faster than they anticipated, in effect it was too late for them to get back on track.

For robustness, all tables in this section are replicated in Appendix B to control for drops stemming from $c_t = 0$. The appendix includes analysis using the ratio c_t/c_t^* as the outcome and an analysis using $\ln(c_t/c_t^*)$ as the outcome with Heckman two-step selection correction. The tests are all consistent with the results in this section.

3.2 Experiment 2: Compounding Frequency

3.2.1 Design

Our focus in Experiment 2 is to test the framing effect of period length on consumption. It may be intuitive to think that the results of Experiment 1 were driven by the high interest rates presented to subjects, and that they would not extend to lower rates. The theory, however, implies the exact opposite: the use of high interest rates with few periods, as in Experiment 1, should *reduce* the effects of EGB relative to identical choices framed as low interest rates with many periods. More frequent compounding should result in a greater bias, and consequently we should observe more overconsumption. The design of Experiment 2 is very similar to Experiment 1, but now our manipulation is the frequency with which interest is presented. The control presents problems that are very similar or identical to those in the dynamic arm of Experiment 1. The treatment presents problems that are equivalent, but rather than being framed as $T = 5$ periods long with a i% interest rate, problems are framed as $T = 50$ periods long with a $i^{\frac{1}{10}} = 5.75\%$ interest rate, and consumption opportunities every 10 periods.[19] Thus the number of consumption periods is held constant, and the problems are isomorphic to the control group's. The rate was explained as a daily rate that applies over a given 10-day segment. The treatment task is therefore a pure framing manipulation relative to the control: the incentives and choice set are identical. The full instrument is available in the Online Appendix.

Because the information on the current balance is typical in consumption decisions, the design emulates the dynamic arm from Experiment 1. Since the dynamic arm was much shorter than the static arm more consumption problems could be included in a single experimental session. Table 5 lists all of the consumption problems in Experiment 2. The consumption problems were blocked

[19] We again employed the metaphor of feeding a dog to increase subject comprehension. Treated subjects were told that the dog only eats every 10 days — i.e. on day 0, day 10, etc. — and that the problem ends on day 50.

within a consumption problem type (N, S, M, E). Within the block, the order of the consumption problems was randomized and the order of blocks was randomized as well.

Table 5: Consumption Problems in Experiment 2

Consumption Problem	Income Vector	Interest Vector Control
N_1	$\langle 100, 0, 0, 0, 0, 0\rangle$	$\langle 0, 0, 100\%, 0, 0\rangle$
N_2	$\langle 100, 0, 0, 0, 0, 0\rangle$	$\langle 0, 0, 200\%, 0, 0\rangle$
N_3	$\langle 0, 0, 0, 0, 0, 900\rangle$	$\langle 0, 0, 100\%, 0, 0\rangle$
N_4	$\langle 0, 0, 0, 0, 0, 900\rangle$	$\langle 0, 0, 200\%, 0, 0\rangle$
S_1	$\langle 100, 0, 0, 0, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
S_2	$\langle 300, 0, 0, 0, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
S_3	$\langle 500, 0, 0, 0, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
S_4	$\langle 700, 0, 0, 0, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
M_1	$\langle 0, 0, 100, 100, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
M_2	$\langle 0, 0, 300, 300, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
M_3	$\langle 0, 0, 500, 500, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
M_4	$\langle 0, 0, 700, 700, 0, 0\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
E_1	$\langle 0, 0, 0, 0, 0, 500\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
E_2	$\langle 0, 0, 0, 0, 0, 300\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
E_3	$\langle 0, 0, 0, 0, 0, 800\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$
E_4	$\langle 0, 0, 0, 0, 0, 1200\rangle$	$\langle 75\%, 75\%, 75\%, 75\%, 75\%\rangle$

Notes: Interest rates for treated subjects were presented as $(1+i)^{\frac{1}{10}} - 1$, compounding ten times between each consumption period. For example, 75% was presented as 5.75% compounding ten times.

3.2.2 Sample and Incentives

This laboratory experiment was conducted at the UCI Experimental Social Sciences Laboratory (ESSL); 100 subjects participated. Subjects were not provided with any tools, and calculators/cell phones were expressly forbidden. Subjects could earn up to \$30 based on the quality of their responses, in addition to a \$7 participation fee. The payment scheme was identical to Experiment 1 except show up pay was slightly higher and incentive pay was scaled to a maximum of \$30 to compensate. A subject's additional payment was given by $\$30 - 30 \cdot [1 - (u_a - u_m)/(u_o - u_m)]^{\frac{1}{2}}$.

The mean incentive payment was \$16.94: control subjects averaged \$16.44, while treatment subjects averaged \$17.26.

3.2.3 Hypotheses

A biased agent's under-estimation of the present value of a financial product is exacerbated when it is described in terms of shorter time periods and hence more compounding (see Proposition A.1). For example, a loan compounding at 1% interest per day may not sound like much to a biased person but the equivalent loan at 3,355% interest per year may sound like a lot.

Hypothesis 4 (Framing) *Normalized overconsumption, given by normalized* $\ln(c_0/c_0^*)$*, in the treatment group will be greater than the control group. The effect will be weakest when all income is immediate.*

This effect is driven by both the price effect and wealth effect of EGB. However the wealth effect will only be present in consumption problems in which some income is received in the future $y_t > 0$ for $t > 0$, because a biased agent correctly perceives the present value of income in the present. Thus we should expect that the framing effect will be stronger in Consumption Problems M and E for which income is received after the initial period.

3.2.4 Results

We first present evidence using just subjects' initial choices, which have the most straightforward interpretation. We will utilize all choices in the regression analysis. Figure 4 displays the CDF of our preferred measure of overconsumption, $\ln(c_0/c_0^*)$. Panel (a) shows significant over-consumption in both frames, which suggests that the eye-catching 75% interest rates are not misleading subjects significantly more than the 5.75% frame. To the contrary, we can see that overconsumption in the treatment first-order stochastic dominates overconsumption in the control, in line with the prediction of Hypothesis 4. A Kolmogorov-Smirnov test strongly rejects the null of identical CDFs between control and treated subjects at $p = 0.018$.

In panels (b) and (c) we disaggregate the distributions into Consumption Problem S and Consumption Problems M and E. These panels show that the first-order stochastic dominance relationship is driven primarily by Consumption Problems M and E. The stochastic dominance is again apparent in panel (c), and once again a Kolmogorov-Smirnov test rejects equal distributions at $p = 0.039$. In contrast, panel (b) shows that the distributions are very similar between the control and treatment in Consumption Problem S. Confining the sample to Consumption Problem S, where all income is received at the start, the Kolmogorov-Smirnov test no longer rejects equal distributions.[20]

Figure 4: Overconsumption And Framing

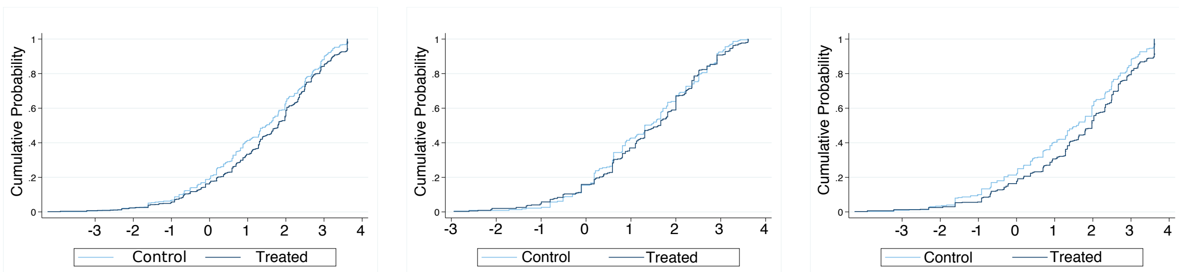

(a) All Consumption Problems (b) Income at Start (c) Delayed Income

Notes: Figures show the empirical distribution of $\ln(c_0/c_0^*)$: the natural logarithm of the ratio of subjects' choice of consumption for period 0 to the optimal consumption choice. Panel (a) shows all Consumption Problems; Panel (b) restricts to Consumption Problem S, where all income was received in the initial period and EGB operates only through the price effect; Panel (c) plots Consumption Problems M, and E, which allow the wealth effect as well.

In Table 6 we unpack the effect of the treatment on mean overconsumption. We now include all consumption choices, using our preferred measure of overconsmption, $\ln(c_t/c_t^*)$. In column (1), we include just an indicator for the frequent-compounding treatment. Unsurprisingly given the results in Figure 4, we find a small but statistically insignificant effect overall. This is because the coefficient includes the effect both on late periods and on choices in Consumption Problem S, where the framing effect is predicted to be weak. In column (2), we therefore include a dummy for Consumption Problem S (All at Start) and for the period number, as well as their interactions with the treatment. The coefficient on Treated is therefore interpreted as the effect of the framing treatment on a period-0 choice in a problem featuring delayed income. It is positive and marginally significant ($p = 0.08$), and its magnitude of 0.269 implies that treated subjects consumed $e^{0.269} - 1 = 31\%$ more than control subjects initially. Our preferred specification is in column (3), and additionally

[20] This pattern of significance is replicated if we replace the Kolmogorov-Smirnov test with a Mann-Whitney test. The Mann-Whitney test is significant for the pooled data (p=0.007) and for the delayed-income Consumption Problems M and E (p = 0.004), but not for the immediate-income Consumption Problem S.

includes treated × all at start × period interactions. Again, the interpretation of the coefficient on Treated is that a treated subject in period 0 of a problem featuring delayed consumption consumed $e^{0.311} - 1 = 36\%$ more than a control subject ($p = 0.035$). As a robustness exercise, we repeat this analysis in columns (4)–(6) using c_t/c_t^* as the dependent variable and find very similar results. The Treated × Period interactions are now marginally significant, providing better evidence that the framing effect is strongest at the beginning of the problem, as predicted.

Table 6: Effect of Framing on Overconsumption

	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.197	0.269*	0.311**	0.352	1.406*	1.807**
	(0.158)	(0.152)	(0.145)	(0.225)	(0.746)	(0.877)
Treated x All At Start		-0.156	-0.261		-0.343	-1.414
		(0.141)	(0.178)		(0.334)	(1.080)
Treated x Period=1		-0.014	-0.026		-1.074*	-1.524**
		(0.060)	(0.081)		(0.574)	(0.682)
Treated x Period=2		-0.140	-0.138		-1.181*	-1.616*
		(0.097)	(0.133)		(0.654)	(0.825)
Treated x Period=3		-0.103	-0.221		-1.323*	-1.897**
		(0.131)	(0.164)		(0.681)	(0.853)
Treated x Period=4		0.168	0.060		-1.150	-1.713*
		(0.166)	(0.196)		(0.708)	(0.877)
Treated x All at Start x Period=1			0.034			1.203
			(0.096)			(0.782)
Treated x All at Start x Period=2			0.004			1.162
			(0.163)			(0.984)
Treated x All at Start x Period=3			0.269			1.530
			(0.180)			(1.044)
Treated x All at Start x Period=4			0.239			1.498
			(0.206)			(1.066)
N	6,489	6,489	6,489	7,864	7,864	7,864
Subjects	100	100	100	100	100	100

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Dependent variable is $\ln(c_t/c_t^*)$ in columns 1-3 and (c_t/c_t^*) in columns 4-6. Domain dummies (when income is received X compounding) are included in all specifications

4 Alternative Explanations and Rules of Thumb

In this section we consider alternative explanations for the observed pattern of behavior. It is worth noting that subjects' own preferences over their earnings from the experiment do not provide a confound, provided that they prefer more money to less. Time preferences cannot drive the results since subjects receive all earnings as a lump sum at the end of the experiment. Risk preferences also cannot explain the results since there is no uncertainty in the experiment. On the other hand, bounded cognition makes the optimal solution to the consumption problems unknown and hence potentially risky from the subject's perspective. Nonetheless, Consumption Problems S, M, and E all have the same incentives and therefore non-EGB models of bounded rationality do not predict comparative statics differences across these problems. However, there is existing evidence that suggests that people do use rules of thumb for complicated tasks (see Winter, Schlafmann and Rodepeter, 2012; Hey and Knoll, 2011, for a summary of the literature), and so we briefly consider some plausible candidates.

4.1 Constant Consumption

A reasonable rule of thumb to employ is to consume a constant amount each period across the consumption problem. This approximates (though overdoes) the consumption smoothing motive and is cognitively easier to implement than estimating an appropriately scaled increasing consumption vector. The data from Experiment 1, the $T = 5$ consumption problems of the static arm are the best data for observing rules of thumb since we can observe the subjects' full plans. Indeed, we find that 19.1% of observations in this sample appear to follow this rule.

The one hypothesis consistent with this rule is Hypothesis 1: overconsumption from compounding. The optimal consumption vector for Consumption Problem N is flatter than the optimal consumption vector for S, and so c_t^* is lower for S than N. Hence a *Constant* rule will on average overconsume for N and overconsume even more for S. We replicate the results of Table 2 dropping all *Constant* subjects in Table C.8 in the Appendix, and find that the increased consumption in Consumption Problem S persists even among subjects not employing this rule of thumb. Therefore a *Constant* rule cannot be the sole driver of our results.

Furthermore, this rule of thumb does not predict Hypotheses 2-4. Because the interest rates are the same for S, M, and E, *Constant* predicts no difference in behavior across these problems if agents do not exhibit a wealth effect of EGB. A *Constant* subject would estimate the budget without bias and then choose the appropriate constant consumption vector. Therefore it predicts no difference in behavior. A person who applies *Constant* and suffers from EGB will still satisfy Hypothesis 2. Thus the predictions of the wealth effect of EGB are preserved even though the person may not dynamically optimize given perceptions. Similarly, for Hypothesis 3, dynamic inconsistency, if a *Constant* person estimates the budget properly there should be no change given that the rule is literally constant consumption. For Hypothesis 4, framing, a person who is unbiased in estimating the budget should not be affected by the framing manipulation and thus a *Constant* person should not exhibit a treatment effect. Thus while a significant number of subjects exhibit this rule of thumb, their behavior still cannot be explained without the addition of EGB.

4.2 Never Borrow

Avoiding borrowing from future earnings may be a sensible rule to follow, especially when interest rates are high. It is also easy to implement. In Consumption Problems N and S, *Never Borrow* does not restrict behavior, hence it does not predict any effect including an effect from compounding (Hypothesis 1). In Consumption Problem M, *Never Borrow* requires consuming zero until $t = 2$ and consuming weakly less than y_2. In Consumption Problem E *Never Borrow* implies zero consumption until the final period. We find that only 2.6% follow this rule in (Consumption Problem E of Experiment 1, $T = 5$). However 18.4% follow *Never Borrow* in Consumption problem M. Thus *Never Borrow* does seem to play some role in at least one consumption problem. However when comparing M to E, *Never Borrow* implies earlier consumption in M, and when comparing S to M it again implies higher early consumption in S. This is the opposite of Hypothesis 2 and the empirical findings. Likewise, *Never Borrow* does not predict dynamic inconsistency (Hypothesis 3) or framing (Hypothesis 4).

5 Discussion and Conclusion

In this paper, we tested the framing effects of EGB in a set of controlled laboratory consumption-savings experiments. All four of the main predictions — sensitivity to compound interest, sensitivity to timing of income, dynamically inconsistent planning, and sensitivity to compounding frequency — are supported against the null hypothesis of no effect.

Nevertheless, there may be much value-added in testing whether the results presented in this paper can be identified in more naturalistic economic environments. Both natural experiments and randomized controlled trials in the field have the potential to say much on this question. This would also address a second limitation of our study. By using the induced-valuation method we eliminated the role of preference heterogeneity. This makes for a tightly controlled test of theory, but it precludes any comparison between the relative importance of preferences and perceptions in consumption-savings. Our experiments show that (mis-)perceptions have the potential to have a large impact on choices, and so it is natural to ask outside of the lab are most people's consumption-savings decisions driven largely by perceptions or by preferences? We have isolated mis-perceptions here, but the effects we have demonstrated may sometimes be mitigated and sometimes exacerbated by the interaction with preferences. For example, our study does not inform how much of payday loan usage is driven by high (or present-biased) discount rates and how much by misperception of the cost of loans. However, it does suggest that analyses which ignore EGB may yield biased results.

As is often the case when behavioral economists make new predictions, it seems that market actors already reflect many of these insights in their current practice. For example, low-cost mutual fund indexers already recognize that the difference between a 1% expense ratio and a 0.2% expense ratio may not be salient to many clients. When communicating the cost advantages of their fund, they choose a non-compounding frame by estimating the difference in total final portfolio levels.[21] The choice of compounding frequency used to present financial products is even more firmly rooted in firm practice. A lender will want to use as high-frequency a frame as possible, the

[21] For example, Vanguard compares the total fees charged over 10 years on a $10,000 investment in their funds with the fees charge by the average fund in the category.

leading example being payday lenders who often quote daily or fortnightly rates. Conversely, a firm seeking investment will want to frame as broadly as possible.[22] We have demonstrated that these practices which firms have developed are in fact best-responses to a simple theory of consumer misperceptions, and there are clearly further ramifications both for firms and for policymakers.

References

Almenberg, Johan and Christer Gerdes, "Exponential Growth Bias and Financial Literacy," *Applied Economics Letters*, 2012, *19*, 1693–1696.

Anderhub, Vital, Werner Güth, Wieland Müller, and Martin Strobel, "An Experimental Analysis of Intertemporal Allocation Behavior," *Experimental Economics*, 2000, *3*, 137–152.

Benzion, Uri, Alon Granot, and Joseph Yagil, "The Valuation of the Exponential Function and Implications for Derived Interest Rates," *Economic Letters*, 1992, *38*, 299–303.

Brown, Alexander L., Zhikang Eric Chua, and Colin F. Camerer, "Learning and Visceral Temptation in Dynamic Saving Experiments," *Quarterly Journal of Economics*, 2009, *124*, 197–230.

Eisenstein, Eric M. and Stephen J. Hoch, "Intuitive Compounding: Framing, Temporal Perspective, and Expertise," December 2007. WorkingPaper.

Ensthaler, Ludwig, Olga Nottmeyer, Georg Weizsäcker, and Christian Zankiewicz, "Hidden skewness: On the difficulty of multiplicative compounding under random shocks," November 2013. DIW Berlin Discussion Paper No. 1337.

Goda, Gopi Shah, Colleen Flaherty Manchester, and Aaron Sojourner, "What Will My Account Be Worth? Experimental Evidence on How Retirement Income Projections Affect Saving," *Journal of Public Economics*, 2014, *119*, 80–92.

_ **, Matthew R. Levy, Colleen Flaherty Manchester, Aaron Sojourner, and Joshua Tasoff**, "The Role of Exponential-Growth Bias and Present Bias in Retirement Savings Decisions," 2015. Working Paper.

Hey, John D. and Julia A. Knoll, "Strategie in Dynamic Decision Making — An Experimental Investigation of the Rationality of Decision Behaviour," *Journal of Economic Psychology*, 2011, *32*, 399–409.

_ **and Valentino Dardanoni**, "Optimal Consumption Under Uncertainty: An Experimental Investigation," *Economic Journal*, 1988, *98* (390), 105–116.

Johnson, Stephen, Laurence J. Kotlikoff, and William Samuelson, "Can People Compute? An Experimental Test of the Life Cycle Consumption Model," March 1987. NBER Working Paper No. 2183.

Kahneman, Daniel and Gary Klein, "Conditions for Intuitive Expertise," *American Psychological Association*, 2009, *64* (6), 515–526.

Keren, Gideon, "Cultural Differences in the Misperception of Exponential Growth," *Perception and Psychophysics*, 1983, *34* (3), 289–293.

Levy, Matthew R. and Joshua Tasoff, "Exponential-Growth Bias and Lifecycle Consumption," *Journal of the European Economic Association*, 2015.

MacKinnon, Andrew J. and Alexander J. Wearing, "Feedback and the Forecasting of Exponential Change," *Acta Psychologica*, 1991, *76*, 177–191.

McKenzie, Craig R.M. and Michael J. Liersch, "Misunderstanding Savings Growth: Implications for Retirement Savings Behavior," *Journal of Marketing Research*, 2011, *48*, S1–S13.

[22] Although less popular in the US, so-called structured deposits in the UK repackage financial derivatives and commonly use a multi-year frame to describe possible returns, for example, "35.95% after 6 years" rather than the equivalent 5.25% annual return.

Meissner, Thomas, "Intertemporal consumption and debt aversion: an experimental study," *Experimental Economics*, June 2015, pp. 1–18.

Noussair, Charles and Kenneth Matheny, "An experimental study of decisions in dynamic optimization problems," *Economic Theory*, 2000, *15*, 389–419.

Smith, Vernon L, "Experimental Economics: Induced Value Theory," *American Economic Review*, 1976, pp. 274–279.

Soll, Jack B., Ralph L. Keeney, and Richard P. Larrick, "Consumer Misunderstanding of Credit Card Use, Payments, and Debt: Causes and Solutions," 2011. Working Paper.

Song, Changcheng, "Financial Illiteracy and Pension Contributions: A Field Experiment on Compound Interest in China," January 2012. Working Paper.

Stango, Victor and Jonathan Zinman, "Exponential Growth Bias and Household Finance," *Journal of Finance*, December 2009, *64* (6), 2807–2849.

Wagenaar, Willem A. and Han Timmers, "The Pond-and-Duckweed Problem: Three Experiments on the Misperception of Exponential Growth," *Acta Psychologica*, 1979, *43*, 239–251.

— **and Sabato D. Sagaria**, "Misperception of Exponential Growth," *Perception and Psychophysics*, 1975, *18* (6), 416–422.

Winter, Joachim K., Kathrin Schlafmann, and Ralf Rodepeter, "Rules of Thumb in Life-Cycle Saving Decisions," *Economic Journal*, May 2012, *122*, 479–501.

A Compounding-Frequency Framing Effects

Proposition A.1 (Interest Framing) *Consider a single period of interest $i_0 > 0$, and a single cash flow y_0. Let $i_n = (1+i_0)^{1/n} - 1$ denote the equivalent rate when interest is received at a frequency of $n > 1$. Then for all $\alpha < 1$ and all $n > 1$:*

(i) The perceived value of $|y_0|$ after n periods of interest i_n is strictly smaller than the perceived value of $|y_0|$ after 1 period of interest i_0

(ii) The perceived value of $|y_0|$ after n periods of interest i_n is strictly greater than the perceived value of $|y_0|$ after 1 period of interest $ln(1+i_0)$

Proof of Proposition A.1 Define the t-period perception of the period $\tau > t$ value of all future income through period τ as $\hat{V}_{t,\tau}(\vec{y}_t, \vec{\imath}; \alpha) = \sum_{s=t}^{\tau} y_s \cdot p_s(\vec{\imath}, \tau; \alpha)$. Part (i) is an immediate result:

$$\begin{aligned} \widehat{V}_{0,n}(\vec{y}, \vec{\imath}; \alpha) &= y_0 \cdot \left(\prod_{s=1}^{n} (1 + \alpha i_n) + \sum_{s=1}^{n} (1-\alpha) i_n \right) \\ &< y_0 \cdot \left(\prod_{s=1}^{n} (1 + i_n) \right) = y_0 \cdot (1 + i_0) = \widehat{V}_{0,n}(\vec{y}, \vec{\imath}; 1) \end{aligned} \tag{4}$$

Part (ii) also results from straightforward algebra:

$$\widehat{V}_{0,n}(\vec{y}^{(n)}, \vec{\imath}^{(n)}; \alpha) \geq y_0 \cdot \left(1 + \sum_{s=1}^{n} i_n \right) = y_0 \cdot \left[1 + n \left((1+i_0)^{1/n} - 1 \right) \right] \tag{5}$$

We first note that the right-most term in (5) is a strictly decreasing function in n for $n > 1$, as:

$$d/dn \left[1 + n \left((1+i_0)^{1/n} - 1 \right) \right] = (1+i_0)^{1/n} \left(1 - \frac{\ln(1+i_0)}{n} \right) - 1 < 0 \tag{6}$$

since $(1 - \frac{\ln(1+i)}{n})^n < \frac{1}{1+i}$ for $i > 0$, $n > 1$. We thus take the limit to find a lower bound, and find:

$$\lim_{n \to \infty} y_0 \left[1 + n \left((1+i_0)^{1/n} - 1 \right) \right] = y_0 \left[1 + \ln(1+i_0) \right] = \widehat{V}_{0,1}(\vec{y}, \ln(1+i_0); \alpha) > y_0 \tag{7}$$

B [Online Publication Only] Including Zero Consumption

B.1 Using c_0/c_0^* as the Outcome

This section replicates Tables 2, 3, 4, and 6 using c_t/c_t^* as the outcome instead of $ln(c_t/c_t^*)$. By using the raw ratio, observations for which $c_t = 0$ are included in the sample. The conclusions remain the same. Table B.1 shows that compounding increases overconsumption. Table B.2 shows that shifting income later increases consumption. Table B.3 shows dynamic inconsistency.

Table B.1: Effect of Compounding on Overconsumption

	(1) N vs S	(2) N vs S	(3) All	(4) All	(5) All EGB
Compounding	2.493***	5.725***	8.565***	17.920***	
	(0.271)	(0.581)	(1.210)	(2.351)	
Period		-0.189**		-0.041	-4.719***
		(0.084)		(0.060)	(0.968)
Compounding X Period		-1.733***		-5.588***	
		(0.182)		(0.698)	
Dynamic					-5.375
					(8.915)
Question					0.014
					(0.321)
Question X Dynamic					0.695
					(1.681)
Question X Period					-0.119
					(0.100)
Constant	1.568***	2.003***	1.757***	1.856***	20.926***
	(0.159)	(0.346)	(0.325)	(0.232)	(5.703)
N	890	890	1,526	1,526	1,008
Subjects	82	82	82	82	82

Notes: Dependent variable is (c_t/c_t^*); all columns are OLS regressions. Columns 1 and 2 restrict to Consumption Problems N and S (i.e. testing the price effect of EGB), and Columns 3-6 include all Consumption Problems. Extra periods is the number periods by which T > 1. Question refers to the order in which the subject answered the question; subjects in the static arm answered 20 questions excluding the training, while subjects in the dynamic arm answered 4 excluding training. Standard errors clustered by subject. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table B.2: Effect of Delayed Income on Overconsumption

	(1) All	(2) All	(3) T=5	(4) T=5
Income at End	20.946***	32.900***	33.196***	42.307***
	(4.376)	(5.611)	(4.943)	(9.304)
Income at End X Period		-10.844***		
		(1.400)		
Period		-2.212***		
		(0.260)		
Income at Middle			0.641	-1.219
			(1.103)	(1.652)
Dynamic				-2.707
				(2.173)
Income at End X Dynamic				-16.980
				(10.237)
Income at Middle X Dynamic				3.466
				(2.201)
Constant	5.625***	8.858***	9.150***	10.603***
	(0.652)	(0.992)	(1.045)	(1.942)
N	430	430	246	246
Subjects	82	82	82	82

Notes: Dependent variable is (c_0/c_0^*) ; all columns are OLS regressions. For columns 1-2 sample comprises those lifecycles where the entire endowment is received either in period 0 (Income At Start) or period T (Income At End). For columns 3-4 all lifecycles of length T = 5 are used. Standard errors clustered by subject. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table B.3: Dynamic vs. Static Plans

	(Inc At Start)	(Inc At Start)	(Inc At End)	(Inc At End)
Dynamic	-0.073	-0.179**	-0.533*	-0.448
	(0.058)	(0.074)	(0.305)	(0.309)
Period> 0	0.054		0.288**	
	(0.065)		(0.138)	
Dynamic X Period> 0	0.446***		-0.588***	
	(0.102)		(0.176)	
Period		0.045*		0.173**
		(0.025)		(0.078)
Dynamic X Period		0.231***		-0.277***
		(0.057)		(0.090)
Constant	0.287***	0.240***	1.431***	1.315***
	(0.052)	(0.043)	(0.285)	(0.285)
N	410	410	410	410
Subjects	82	82	82	82

Notes: Dependent variable is c_t/W_0, i.e. consumption normalized by actual starting wealth. The sample comprises the compounding T=5 lifecycles with all income received either in period 0 or in period 5. Standard errors clustered by subject.
$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

B.2 Heckman Two-Step Analysis

This section replicates Tables 2, 3, 4, and 6 using the Heckman two-step selection correction to account for data dropped in the log-transform because of a zero-response. The interpretation of such responses is ambiguous. While they may represent an earnest attempt to maximize utility, they are clearly sub-optimal given the infinite marginal utility of consumption at zero. If subjects are instead using a rule of thumb, or are using a zero to indicate they do not know how to optimize a given problem, then zeros may indicate that a subject's "true" response – i.e. the response they would give if forced to give a nonzero response – is unobserved.

Under such an interpretation, our OLS estimates in the main text are unbiased if subjects are missing-at-random. However it is possible that either more- or less-biased subjects systematically are more likely to be unobserved. In this section, we account for selection on observables. In all tables, we report the results of a Heckman two-step estimator, using consumption problem dummies and a subject's average overconsumption on problems for which they gave non-zero answers as the selection equation. The conclusions remain the same as the main analysis.

Table B.4: Effect of Compounding on Overconsumption

	(1) All	(2) All	(3) T=5	(4) T=5	(5) All EGB
Compounding	1.021***	2.009***	1.025***	2.150***	
	(0.064)	(0.081)	(0.084)	(0.102)	
Extra Periods		-0.060***			-0.661***
		(0.022)			(0.040)
Compounding X Extra Periods		-0.598***			
		(0.028)			
Dynamic				-0.331***	-0.405
				(0.118)	(0.404)
Compounding X Dynamic				-1.349***	
				(0.135)	
Question					-0.006
					(0.012)
Question X Dynamic					-0.031
					(0.082)
Question X Extra Periods					0.004
					(0.003)
Constant	0.193***	0.349***	0.071	0.360***	2.607***
	(0.046)	(0.077)	(0.055)	(0.103)	(0.160)
N	1,640	1,640	1,032	1,032	1,116
Subjects	82	82	82	82	82
lambda	0.400	0.328	0.630	0.561	0.515
	0.162	0.165	0.184	0.145	0.157

Table B.5: Effect of Delayed Income on Overconsumption

	(1) All	(2) All	(3) T=5	(4) T=5
Income at End	1.241***	1.689***	2.210***	1.194***
	(0.195)	(0.166)	(0.486)	(0.206)
Income at End X Extra Periods		-0.483***		
		(0.088)		
Extra Periods		-0.511***		
		(0.036)		
Income at Middle			0.588***	0.102
			(0.185)	(0.137)
Dynamic				-1.500***
				(0.161)
Income at End X Dynamic				-0.240
				(0.441)
Income at Middle X Dynamic				0.224
				(0.203)
Constant	0.899***	1.897***	0.927***	2.118***
	(0.074)	(0.082)	(0.097)	(0.124)
N	782	782	774	774
Subjects	82	82	82	82
lambda	-0.065	-0.268	-1.299	0.200
	0.209	0.206	0.388	0.227

Notes: Dependent variable is $\ln(c_0/c_0^*)$; all columns are OLS regressions. For columns 1-2 sample comprises those lifecycles where the entire endowment is received either in period 0 (Income At Start) or period T (Income At End). For columns 3-4 all lifecycles of length T = 5 are used. Standard errors clustered by subject. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table B.6: Dynamic vs. Static Plans

	(Inc At Start)	(Inc At Start)	(Inc At End)	(Inc At End)
Dynamic	-0.260	-0.251	0.052	0.026
	(0.171)	(0.164)	(0.238)	(0.263)
Period> 0	0.188		0.198**	
	(0.147)		(0.099)	
Dynamic X Period> 0	0.689***		-0.467***	
	(0.236)		(0.164)	
Period		0.110**		0.104***
		(0.050)		(0.036)
Dynamic X Period		0.267***		-0.177**
		(0.077)		(0.081)
Constant	-1.419***	-1.486***	-0.167	-0.216
	(0.138)	(0.127)	(0.177)	(0.183)
N	410	410	410	410
Subjects	82	82	82	82
lambda	-0.675	-0.717	-1.152	-1.141
	0.200	0.163	0.102	0.109

Notes: Dependent variable is $\ln(c_t/W_0)$, i.e. consumption normalized by actual starting wealth. The sample comprises the compounding T=5 lifecycles with all income received either in period 0 or in period 5. Standard errors clustered by subject.
$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Table B.7: Effect of Framing on Overconsumption (Logs)

	(1)	(2)	(3)
Treated	0.178	0.240	0.280**
	(0.158)	(0.150)	(0.143)
Treated x All At Start		-0.140	-0.240
		(0.139)	(0.176)
Treated x Period=1		-0.014	-0.025
		(0.060)	(0.080)
Treated x Period=2		-0.136	-0.131
		(0.097)	(0.132)
Treated x Period=3		-0.101	-0.213
		(0.131)	(0.164)
Treated x Period=4		0.171	0.068
		(0.166)	(0.195)
Treated x All at Start x Period=1			0.033
			(0.095)
Treated x All at Start x Period=2			-0.004
			(0.163)
Treated x All at Start x Period=3			0.261
			(0.180)
Treated x All at Start x Period=4			0.232
			(0.206)
N	8,000	8,000	8,000
Subjects	100	100	100

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Dependent variable is $\ln(c_t/c_t^*)$Domain dummies (when income is received X compounding) are included in all specifications

C [Online Publication Only] Rules of Thumb – Sample Excluding *Constant* Observations

Table C.8: Effect of Compounding on Overconsumption (Sample of non-*Constant* Subjects)

	(1) N vs S	(2) N vs S	(3) All	(4) All	(5) All EGB
Compounding	0.864***	1.449***	1.384***	1.950***	
	(0.077)	(0.091)	(0.075)	(0.075)	
Period		-0.059*		-0.059*	-0.468***
		(0.034)		(0.034)	(0.076)
Compounding X Period		-0.394***		-0.507***	
		(0.053)		(0.041)	
Dynamic					0.128
					(0.419)
Question					0.002
					(0.016)
Question X Dynamic					0.000
					(0.083)
Question X Period					-0.006
					(0.005)
Constant	0.439***	0.496***	0.439***	0.496***	2.368***
	(0.080)	(0.079)	(0.080)	(0.079)	(0.216)
N	356	356	593	593	427
Subjects	81	81	82	82	82

Notes: Dependent variable is $\ln(c_t/c_t^*)$; all columns are OLS regressions. Columns 1 and 2 restrict to Consumption Problems N and S (i.e. testing the price effect of EGB), and Columns 3-6 include all Consumption Problems. Question refers to the order in which the subject answered the question; subjects in the static arm answered 20 questions excluding the training, while subjects in the dynamic arm answered 4 excluding training. Standard errors clustered by subject.
$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

D [Online Publication Only] Instrument

D.1 Control

Figure D.1: Opening Screen

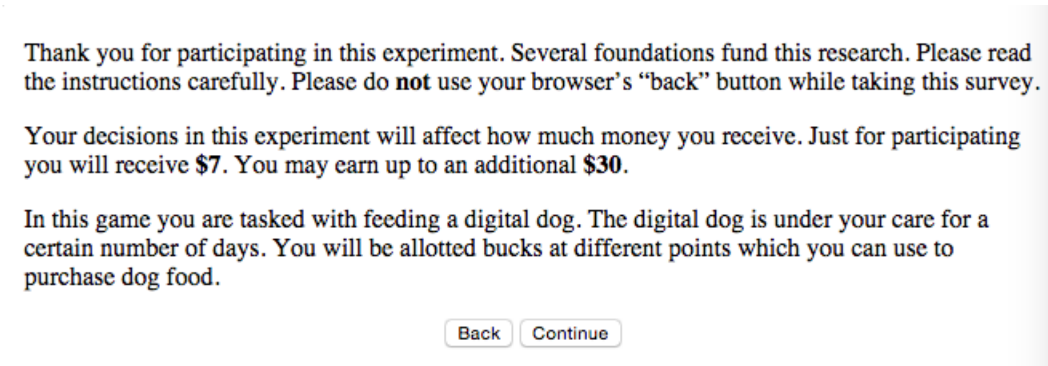

Thank you for participating in this experiment. Several foundations fund this research. Please read the instructions carefully. Please do **not** use your browser's "back" button while taking this survey.

Your decisions in this experiment will affect how much money you receive. Just for participating you will receive **$7**. You may earn up to an additional **$30**.

In this game you are tasked with feeding a digital dog. The digital dog is under your care for a certain number of days. You will be allotted bucks at different points which you can use to purchase dog food.

Back Continue

Figure D.2: Instructions 1

Dog: The dog likes to eat. The more the dog eats, the happier the dog is. The dog's happiness is measured in terms of tail wags. Specifically, if the dog eats *x* units of dog food, the dog has *x*^(1/2) wags. The table below gives some sample values for food and wags. The graph plots the number of wags as a function of food.

How food affects wags

Food	0	5	10	25	50	100	150	200	300	500
Wags	0	2.24	3.16	5	7.1	10	12.2	14.1	17.3	22.4

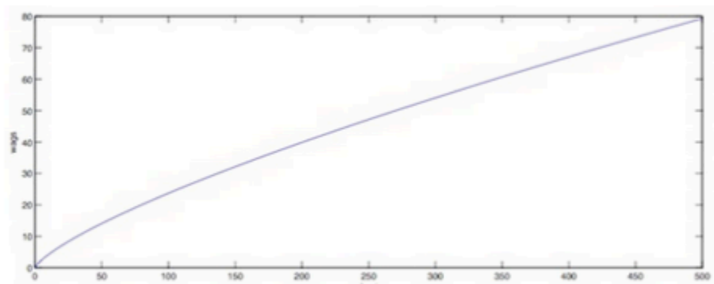

Notice that each unit of food gives the dog more wags when the dog eats little than when the dog eats a lot. For example. If the dog eats 5 units then it is 5/2.24 = 2.23 units of food per wag. But if the dog eats 100 units, then it is 100/10 = 10 units of food per wag. So in other words, the first units of food give the hungry dog lots of additional happiness but there are diminishing returns. The 100th unit of food provides incrementally less additional happiness.

Back Continue

Figure D.3: Instructions 2

Savings & Loans: Each day you will receive a certain number of bucks. A single buck buys a single unit of food. All food purchased that day is eaten by the dog (you may not save food). However, you may save bucks for future days. Whenever you spend less in a day than you are allotted, the difference is automatically saved into the next day. Saved bucks accumulate interest at a rate of *i* that is specified for every day. So for example, if *i* = 60% then saving 1 buck will add 1.6 bucks to your allotment in the next day. If *i* = 100% then saving 1 buck adds 2 bucks in the next day.

You may also take out loans to pay for food in early days. Whenever you spend more in a day than you are allotted, the difference is automatically debited against the next day. Debts accumulate interest at a rate *i* that is specified for every day. So for example, if *i* = 60% then debiting 1 buck will decrease the allotment in the next day by 1.6 bucks. If *i* = 100% then debiting 1 buck will decrease the next day's allotment by 2 bucks. You may sometimes borrow so much that your allotment in the next day is negative, and you must then borrow again the next day to cover the debt. However, you may not borrow so much that your allotment in the final day becomes negative. The computer will automatically limit your actual spending to the amount you could afford given your allotments – if you reach this amount before the final day, the computer will automatically set spending in all future days to zero. You will automatically spend your full final-day allotment on dog food.

Back Continue

Figure D.4: Instructions 3

At the end of all the days, the number of wags will be summed up. This is called a single round. You will have two practice rounds and 16 actual rounds. At the end of the 16 actual rounds, one actual round will be randomly chosen. You will be paid based on the sum of tail wags in that round. The more tail wags you have the more money you make. To see the exact formula that describes your earnings click here.

Back Continue

Figure D.5: Example 1

Let's run through some examples.

In each round, at the beginning of the round, you will choose how much the dog eats on the first day. When you click to continue you will see how much money you have. You will then choose how much the dog eats on the second day. This process iterates until you complete the second to last day. Any remaining wealth is used to purchase dog food in the last day.

EXAMPLE 1

Suppose there are *three* days, $i = 0$, and you receive 10 bucks on Day 0 and 0 in all other days. At this point, you specify how much food you wish to purchase for the dog on Day 0. Suppose you type:

- Day 0: 6 units.

The next screen will indicate that you have 4 bucks to spend. Since no income is received in future days you cannot take a loan since you wouldn't be able to pay it off. Suppose you type:

- Day 1: 4 units

Since there are no bucks remaining, 0 units of dog food are eaten in Period 2. The total number of tail wags generated is 6^(1/2) + 4^(1/2) = 4.4

Back Continue

Figure D.6: Example 2

Let's run through some examples.

EXAMPLE 2:
Suppose you have the dog for *two* days, *i* = 100%, and you receive 30 bucks on Day 1 and 0 on Day 0. Suppose you list the following amounts of food:

- Day 0: 10 units
- Day 1: The remaining money will be spent on food in day 1.

10 bucks were spent on Day 0, which was more than your allotment of 0 bucks for that day. This means that (10-0)=10 bucks were automatically borrowed. The interest rate is 100%, so your borrowing decreased your Day 1 allotment by (10 × 2) = 20 bucks. This means that there are (30-20)=10 bucks that are automatically spent on food on Day 1. The total number of tail wags generated is 10^(1/2) + 10^(1/2) = 6.3

Make sure you understand these examples before continuing
Click below to begin the training session.

Back Continue

Figure D.7: Training 1

TRAINING SESSION (1)

You are tasked with feeding the dog for five days. The table below gives your income and the interest rates in each of the five days. Remember, any amount of money that is left over after the second to last day is automatically spent on the last day.

Day	Day 0	Day 1	Day 2	Day 3	Day 4	
Income	50	50	50	50	50	
Daily Interest Rate	0%	0%	100%	0%	-	
Your Spending	-	-	-	-	-	-

(Interest is applied to savings and debts at the end of that day.)

Based on your choices so far, you are carrying a balance of **$50.0** bucks over into today. You can borrow against your endowment from future days, or save some bucks to spend later. Remember: if you ever run out of bucks, you will automatically buy zero units of food on subsequent days.

Below, please specify how much dog food you wish to purchase.

Food for Day 0:

Submit

If you'd like to calculate the number of wags you will receive for some dollar investment, enter the dollar amount below.

Dollars: 0 **Wags:** 0

How food affects wags

Food	0	5	10	25	50	100	150	200	300	500
Wags	0	2.24	3.16	5	7.1	10	12.2	14.1	17.3	22.4

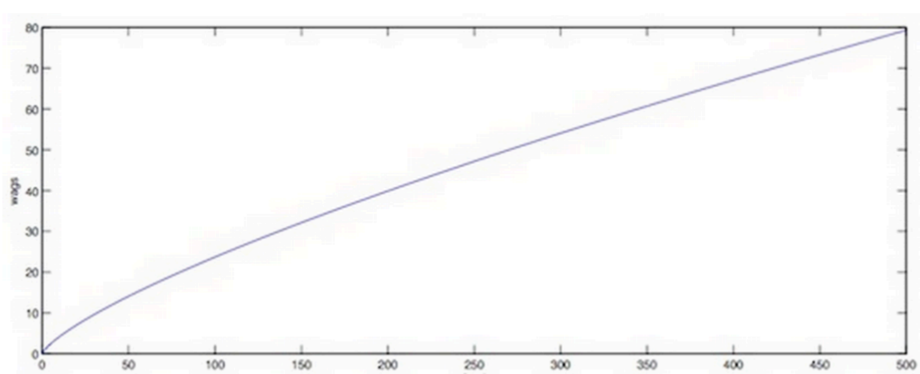

Figure D.8: Review

Here are your responses, the amount of food the dog ate each day, and the optimal solution to the problem:

Day	You chose...	The dog wagged...	The optimal allocation was...	Making the dog wag...
Day 0	10.0	3.16	28.57	5.35
Day 1	50.0	7.07	28.57	5.35
Day 2	40.0	6.32	28.57	5.35
Day 3	100.0	10.0	114.28	10.69
Day 4	100.0	10.0	114.28	10.69
Wags earned	-	36.55	-	37.43

When you are done reviewing, press 'Continue' below to take another practice problem.

Continue

Figure D.9: Training 2: This displays the $t = 0$ screen.

TRAINING SESSION (2)

You are tasked with feeding the dog for five days. The table below gives your income and the interest rates in each of the five days. Remember, any amount of money that is left over after the second to last day is automatically spent on the last day.

Day	Day 0	Day 1	Day 2	Day 3	Day 4	
Income	50	50	50	50	50	
Daily Interest Rate	0%	0%	75%	0%	-	
Your Spending	-	-	-	-	-	-

(Interest is applied to savings and debts at the end of that day.)

Based on your choices so far, you are carrying a balance of **$50.0** bucks over into today. You can borrow against your endowment from future days, or save some bucks to spend later. Remember: if you ever run out of bucks, you will automatically buy zero units of food on subsequent days.

Below, please specify how much dog food you wish to purchase.

Food for Day 0:

Submit

If you'd like to calculate the number of wags you will receive for some dollar investment, enter the dollar amount below.

Dollars: **Wags:**

How food affects wags

Food	0	5	10	25	50	100	150	200	300	500
Wags	0	2.24	3.16	5	7.1	10	12.2	14.1	17.3	22.4

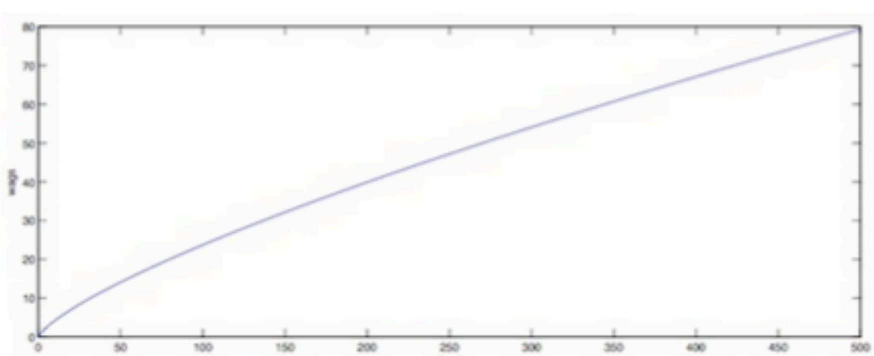

Figure D.10: Training 2: This displays the $t = 2$ screen

TRAINING SESSION (2)

You are tasked with feeding the dog for five days. The table below gives your income and the interest rates in each of the five days. Remember, any amount of money that is left over after the second to last day is automatically spent on the last day.

Day	Day 0	Day 1	Day 2	Day 3	Day 4	
Income	50	50	50	50	50	
Daily Interest Rate	0%	0%	75%	0%	-	
Your Spending	30.0	40.0	-	-	-	-

(Interest is applied to savings and debts at the end of that day.)

Based on your choices so far, you are carrying a balance of **\$80.0** bucks over into today. You can borrow against your endowment from future days, or save some bucks to spend later. Remember: if you ever run out of bucks, you will automatically buy zero units of food on subsequent days.

Below, please specify how much dog food you wish to purchase.

Food for Day 2:

Submit

If you'd like to calculate the number of wags you will receive for some dollar investment, enter the dollar amount below.

Dollars: 0 **Wags:** 0

How food affects wags

Food	0	5	10	25	50	100	150	200	300	500
Wags	0	2.24	3.16	5	7.1	10	12.2	14.1	17.3	22.4

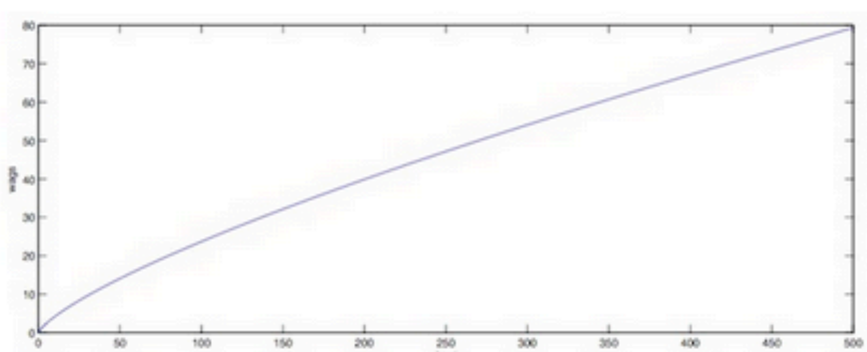

Figure D.11: Review

Here are your responses, the amount of food the dog ate each day, and the optimal solution to the problem:

Day	**You chose...**	**The dog wagged...**	**The optimal allocation was...**	**Making the dog wag...**
Day 0	30.0	5.48	32.4	5.69
Day 1	40.0	6.32	32.4	5.69
Day 2	70.0	8.37	32.4	5.69
Day 3	65.0	8.06	99.22	9.96
Day 4	52.5	7.25	99.22	9.96
Wags earned	-	35.48	-	36.99

When you are done reviewing, press 'Continue' below to begin the experiment!

Continue

Figure D.12: Training Complete

Training Complete

The training session is over. All subsequent rounds are for real stakes. One round will be randomly chosen for payment.

Continue

Figure D.13: Experiment 2 Lifecycle E_4 – This shows how the lifecycle tasks are displayed in the control (dynamic arm).

Day	Day 0	Day 1	Day 2	Day 3	Day 4	Day 5
Income	0	0	0	0	0	1200
Daily Interest Rate	75%	75%	75%	75%	75%	-
Your Spending	-	-	-	-	-	-

(Interest is applied to savings and debts at the end of that day.)

Based on your choices so far, you are carrying a balance of **$0.0** bucks over into today. You can borrow against your endowment from future days, or save some bucks to spend later. Remember: if you ever run out of bucks, you will automatically buy zero units of food on subsequent days.

Below, please specify how much dog food you wish to purchase.

Food for Day 0:

Submit

If you'd like to calculate the number of wags you will receive for some dollar investment, enter the dollar amount below.

Dollars: 0 **Wags:** 0

How food affects wags

Food	0	5	10	25	50	100	150	200	300	500
Wags	0	2.24	3.16	5	7.1	10	12.2	14.1	17.3	22.4

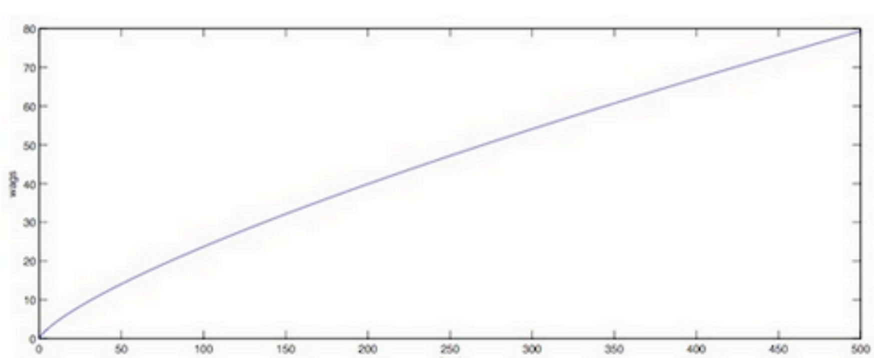

Figure D.14: Payment

We have randomly decided to pay you based on your performance in the 5th trial. To recap, your responses were:

Day	You chose…	The dog wagged…	The optimal allocation was…	Optimal wags
Day 0	20.0	4.47	4.11	2.03
Day 1	30.0	5.48	12.59	3.55
Day 2	30.0	5.48	38.55	6.21
Day 3	170.0	13.04	118.05	10.87
Day 4	170.0	13.04	361.53	19.01
Day 5	52.77	7.26	1107.19	33.27
Wags earned	-	48.77	-	74.94

The minimum wags is 10.0.

Based on how many wags you received, you have earned $17.96.

The experiment is complete! Thank you for participating!

Figure D.15: Thank You

Thank You!

Thank you for participating in the experiment! Your results have been saved. If you would like to finish the test at a later point in time, you may do so by logging in again.

Please be aware that until you finish the test, we cannot pay you for your contribution.

D.2 Static Arm

Figure D.16: Experiment 1 Lifecycle E – This shows how the lifecycle tasks are displayed in static arm.

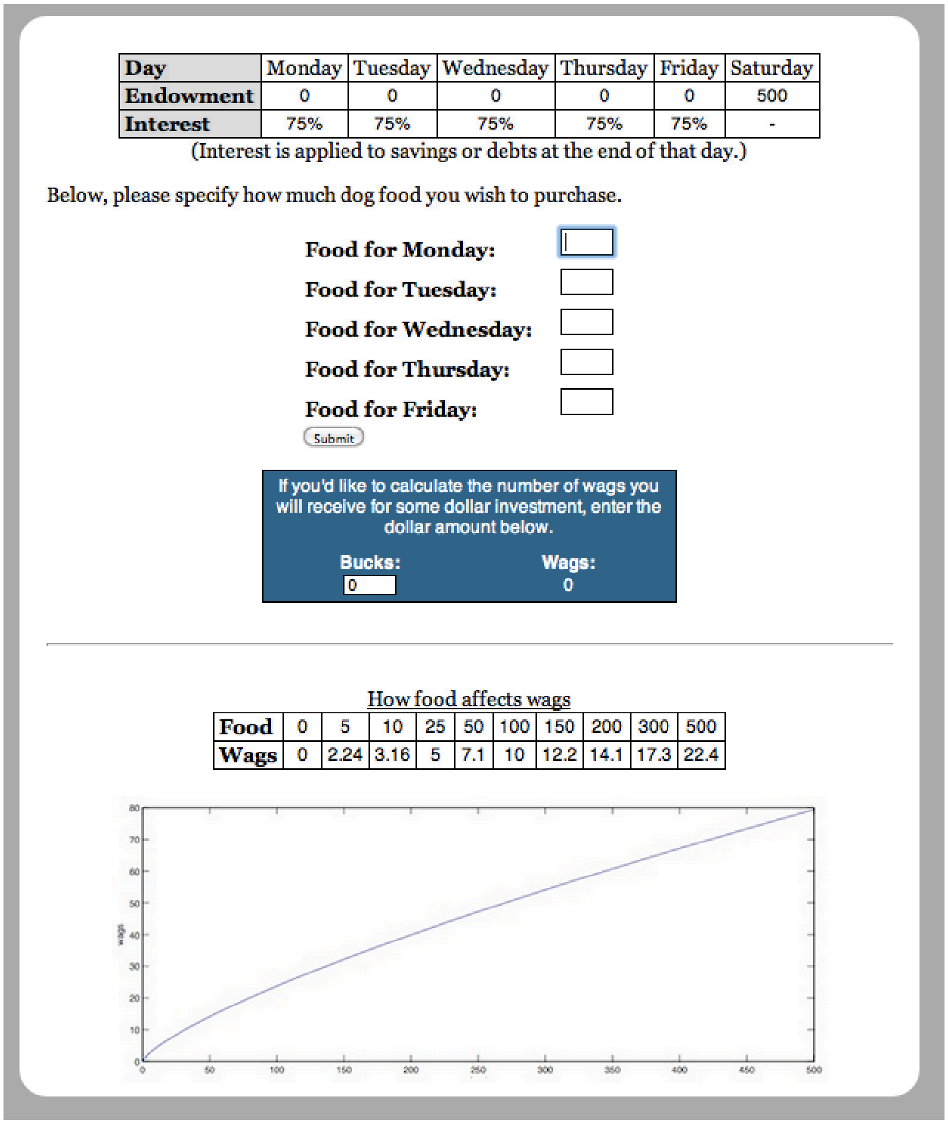

D.3 Treatment

Figure D.17: Experiment 2 Lifecycle E_4 – This shows how the lifecycle tasks are displayed in the treatment.

Day	Day 0	Day 10	Day 20	Day 30	Day 40	Day 50
Income	0	0	0	0	0	1200
Daily Interest Rate	5.75% in days 0-9	5.75% in days 10-19	5.75% in days 20-29	5.75% in days 30-39	5.75% in days 40-49	-
Your Spending	-	-	-	-	-	-

(Interest is applied to savings and debts at the end of that day.)

Based on your choices so far, you are carrying a balance of **$0.0** bucks over into today. You can borrow against your endowment from future days, or save some bucks to spend later. Remember: if you ever run out of bucks, you will automatically buy zero units of food on subsequent days.

Below, please specify how much dog food you wish to purchase.

Food for Day 0:

Submit

If you'd like to calculate the number of wags you will receive for some dollar investment, enter the dollar amount below.

Dollars: 0 **Wags:** 0

How food affects wags

Food	0	5	10	25	50	100	150	200	300	500
Wags	0	2.24	3.16	5	7.1	10	12.2	14.1	17.3	22.4