Economica

Economica (2020) **87**, 52–80
doi:10.1111/ecca.12306

Exponential-growth Bias in Experimental Consumption Decisions

By MATTHEW R. LEVY† and JOSHUA TASOFF‡

†*London School of Economics* ‡*Claremont Graduate University*

Final version received 18 January 2019.

Exponential-growth bias (EGB) is the tendency to neglect the power of compounding interest, which 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 five comparative static predictions implied by EGB: (1) compound interest will increase consumption when the elasticity of intertemporal substitution is greater than 1; (2) higher interest rates lead to more compounding and hence increase consumption; (3) budget-neutral delays in income will increase consumption; (4) the person will exhibit a form of dynamic inconsistency that depends solely on the current account balance and is independent of time preferences; and (5) framing the frequency of interest in shorter units increases consumption. We test these predictions using an induced-value consumption–savings experiment in the laboratory, and find evidence in support of all predictions against the rational benchmark. We consider rules of thumb as alternative hypotheses and find that they cannot explain the results, although they contribute to some findings.

INTRODUCTION

How one consumes and saves over the lifecycle is one of the most consequential sets of financial decisions that 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 laboratory studies (see, for example, Wagenaar and Timmers 1979; Keren 1983; Benzion *et al*. 1992; MacKinnon and Wearing 1991; Eisenstein and Hoch 2007; McKenzie and Liersch 2011; Ensthaler *et al*. 2013) and field studies (see, for example, Almenberg and Gerdes 2012; Soll *et al*. 2011; Goda *et al*. 2014; Song 2015). 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 laboratory experiment. Outside the laboratory, 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 makes it difficult to draw clear conclusions from observational data. Instead, we use choices in an induced-value experiment as a model of real-world behaviour (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 2016). An agent with EGB views the period-T future value of a dollar at time $t \leq T$ as

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

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 $\alpha = 0$, 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 five main comparative static predictions that we test.

- 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.
- Higher interest rates lead to more compounding and hence a larger error for biased agents.
- 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 occurs.
- 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, then she will be surprised by how quickly her assets grow, and if her balance is negative, then 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.
- We predict that 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 *et al.* 2000; Noussair and Matheny 2000; Brown *et al.* 2009; Meissner 2016), with one (Johnson, Kotlikoff and Samuelson 1987) finding behaviour 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 five theoretical predictions has 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 behaviour is predicted to exhibit a systematic bias.

Recent studies in the field suggest that EGB may play an important role in real-world financial decision making. Goda *et al.* (2014) run a field experiment on university employees and show that an intervention that maps retirement contributions to projected retirement savings or retirement income significantly increases contributions. Likewise, Song (2015) conducts an experiment in China and shows that an advocacy campaign emphasizing the importance of exponential growth dramatically increases pension contributions. Levy and Tasoff (2016) find that an elicited measure of EGB predicts total assets in a representative sample of the USA. Recent work by Goda *et al.* (2015) find that an elicited measure of EGB predicts retirement savings, while controlling for cognitive ability, time preferences, risk preferences, and financial literacy. By turning to the laboratory, we are able to complement these field studies.

The experimental results strongly confirm four of the five theoretical predictions, while the fifth indicates that other factors must be considered in addition to EGB. First, compound interest causes subjects to increase their consumption early in the lifecycle. Second, higher interest rates lead to more overconsumption. 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. Our fifth hypothesis, that shifting income later in the lifecycle in a wealth-neutral way increases consumption, was not fully supported by the data. We find that a strong heuristic to avoid borrowing can dominate the predicted increase in consumption when moving from immediate to delayed income. However, the data support the prediction that moving from some delay to a greater delay will increase consumption, as the heuristic applies equally to both situations.

In addition to a borrowing-avoidance heuristic, we consider alternative rules of thumb such as perfectly smoothing consumption across all periods. While we find that our full set of results cannot be explained by EGB alone, likewise we find that alternative explanations cannot predict the pattern of results that we find due to EGB. That is, the results imply that the effects of EGB are large enough to dominate the effects stemming from rules of thumb in most circumstances. We conclude that underestimation of compound growth has a significant effect—both statistically and materially—on behaviour inside the laboratory, and therefore further exploration outside the laboratory may be warranted.

I. Conceptual Framework

We briefly present a model of EGB, and explain the main predictions of our paper. For concreteness, we present one possible parametrization of the bias, but note that all of the predictions that we test follow from the more general model in Levy and Tasoff (2016) 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

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

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 let $\vec{\imath} \in \mathbb{R}^T$ be a vector of interest rates.[3] The budget constraint can be written as

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

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

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

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 (1), which may be thought of as being enforced by other market actors.

In our experiments, we consider five 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, then 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 is greater than 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 that higher interest rates lead to more compounding, thus this should exacerbate any effects of EGB. With an EIS greater than 1, this too will lead to more overconsumption relative to the optimum.

The third 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 time 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 that 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 in each period she will experience a positive windfall. The fourth 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 that previously planned. 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 fifth 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 1 in the Appendix 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% per 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 favourable when presented as monthly, weekly or even daily rates. However, the effect goes only 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 $\ln(1+i_0)\approx i_0$ 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 behaviour when the budget constraint includes delayed income.

II. Experimental Tests of Exponential-growth Bias

We turn to the laboratory 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 that 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 behaviour, do changes in the economic environment or framing of the problem change behaviour 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 behaviour 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, then 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 laboratory 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.

Lifecycle consumption experiments

In all three experiments, 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 clear theoretical predictions on behaviour that would 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 experiments 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 that 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 to 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, an interest rate vector, and a number of periods. A consumption problem of length T required consumption choices in periods 0 to $(T-1)$, with any residual consumption allocated to period T automatically. After each period, subjects would be informed of their current balance (i.e. accumulated savings or debt). 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.

There were four types of consumption problems, each six periods long: $t \in \{0,\ldots,5\}$. The income and interest vectors for all problems are displayed in Table 1. The order in which the consumption problems were presented was random. Problems of type N (no compounding) have no compounding interest, so EGB is irrelevant. Behaviour on these problems established a benchmark level of imprecision and bias that is unrelated to EGB. The three other consumption problem types all had compounding interest: problems of type S (income at start), type M (income at middle), and type E (income at end).

The instantaneous utility function given to subjects was $u(x) = x^{1/2}$. 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.[5] 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 behaviour across different consumption problems without the confound of non-linear wealth effects. We kept the utility function constant throughout the experiments, including the training rounds, in order to focus just on the effects of manipulating the budget constraint. While EGB does make predictions regarding utility—for example, changing the risk-aversion parameter within the CRRA family—that is not our focus here.

After the experiment was completed, 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

TABLE 1
CONSUMPTION PROBLEMS IN EXPERIMENT 1

Problem type	Consumption problem	Interest vector	Income vector	Present value of wealth
No compounding	N_{LL}	⟨0%,0%,37%,0%,0%⟩	⟨100,0,0,0,0,0⟩	100
	N_{HL}	⟨0%,0%,75%,0%,0%⟩	⟨100,0,0,0,0,0⟩	100
	N_{LH}	⟨0%,0%,37%,0%,0%⟩	⟨340,0,0,0,0,0⟩	340
	N_{HH}	⟨0%,0%,75%,0%,0%⟩	⟨340,0,0,0,0,0⟩	340
Income at start	S_{LL}	⟨37%,37%,37%,37%,37%⟩	⟨100,0,0,0,0,0⟩	100
	S_{HL}	⟨75%,75%,75%,75%,75%⟩	⟨100,0,0,0,0,0⟩	100
	S_{LH}	⟨37%,37%,37%,37%,37%⟩	⟨340,0,0,0,0,0⟩	340
	S_{HH}	⟨75%,75%,75%,75%,75%⟩	⟨340,0,0,0,0,0⟩	340
Income at middle	M_{LL}	⟨37%,37%,37%,37%,37%⟩	⟨0,0,108,108,0,0⟩	100
	M_{HL}	⟨75%,75%,75%,75%,75%⟩	⟨0,0,195,195,0,0⟩	100
	M_{LH}	⟨37%,37%,37%,37%,37%⟩	⟨0,0,369,369,0,0⟩	340
	M_{HH}	⟨75%,75%,75%,75%,75%⟩	⟨0,0,663,663,0,0⟩	340
Income at end	E_{LL}	⟨37%,37%,37%,37%,37%⟩	⟨0,0,0,0,0,483⟩	100
	E_{HL}	⟨75%,75%,75%,75%,75%⟩	⟨0,0,0,0,0,1641⟩	100
	E_{LH}	⟨37%,37%,37%,37%,37%⟩	⟨0,0,0,0,0,1641⟩	340
	E_{HH}	⟨75%,75%,75%,75%,75%⟩	⟨0,0,0,0,0,5580⟩	340

utility, u_m the minimum possible utility, and $\overline{y}$ the maximum payment, a subject's additional payment was given by

$$\overline{y} - \overline{y} \cdot \left[1 - (u_a - u_m)/(u_o - u_m)\right]^{1/2} \text{ dollars.}$$

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, though it was emphasized that their payment was increasing in the total utility achieved.[6]

Experiment 1: Compounding and the timing of income

Design In Experiment 1 we manipulate the timing of income, the presence of compounding, the level of interest, and the present value of lifetime income. We use a constant interest vector of either 37% or 75%, and a present value of wealth of either 100 or 340. According to the classical model and our model of EGB, the wealth of a lifecycle should have no effect on wealth-normalized consumption. We select two wealth levels to test whether this is indeed the case.[7] We have a factorial design, yielding 16 consumption problems in total per subject, as shown in Table 1. The consumption problems were blocked 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.

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 an interest rate of 75% corresponds to 5.75% per year, which is near real historical long-run stock market returns. The low per-period interest rate of 37% corresponds to 3.2% per year, a return associated with a fairly conservative portfolio. One may worry that there is something particular about large interest rates that makes behaviour structurally different from lower rates. However, the theory actually predicts that economically equivalent formalizations of the same problem with more periods and lower rates will result in *greater* distortions of behaviour, not less. We address exactly this point later, in Experiment 3, and find that the results are consistent with theory: distortions are stronger with *lower* rates and more periods.

This online experiment was conducted with an Amazon MTurk subject pool. Subjects could earn up to $6.00 based on the quality of their responses, in addition to a $0.50 participation fee. The mean incentive payment was $2.20.

Hypotheses Given that the EIS is greater than 1, the model predicts that $c_t > c_t^*$, where c_t^* is the optimal consumption in any period $t<(T-1)$.[8] 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 from compounding) Normalized overconsumption, given by $\ln(c_t/c_t^*)$, in the consumption problems with compounding will be greater than overconsumption in consumption problems of type N.

Because higher interest rates lead to more compounding, the theory predicts that higher interest rates will lead to more overconsumption. Thus within consumption problems that compound (S, M and E), overconsumption should be greater when the interest rate is 75% compared to 37%.

Hypothesis 2 (Overconsumption from higher interest rates) Normalized overconsumption, given by $\ln(c_t/c_t^*)$, will be increasing in the interest rate in consumption problems with compounding, for a given problem type and present value of wealth.

Finally, shifting income to later periods in a wealth-preserving way should increase log-normalized consumption for a biased agent.[9] An unbiased agent, in contrast, would be unaffected by such shifts.[10] Consider the right-hand side of (1), 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 3 (Deferred income increases consumption) Normalized overconsumption, given by $\ln(c_t/c_t^*)$, should be greater for consumption problem E than for consumption problem M, and greater for consumption problem M than for consumption problem S.

Results The aggregate behaviour of subjects is shown in Figure 1, where we plot the mean consumption paths achieved by interest rate and present value of income across consumption problems as a proportion of the initial present value of income, as well as the optimal consumption profiles in problem types S, M and E.[11] Problem type N, without compounding, exhibits a similar smooth consumption pattern across all four panels. Subjects correctly respond by increasing consumption in period 3, after the positive interest rate, across all the consumption problems. Notably, people on average do react to the interest rate in the correct direction, but smooth consumption somewhat

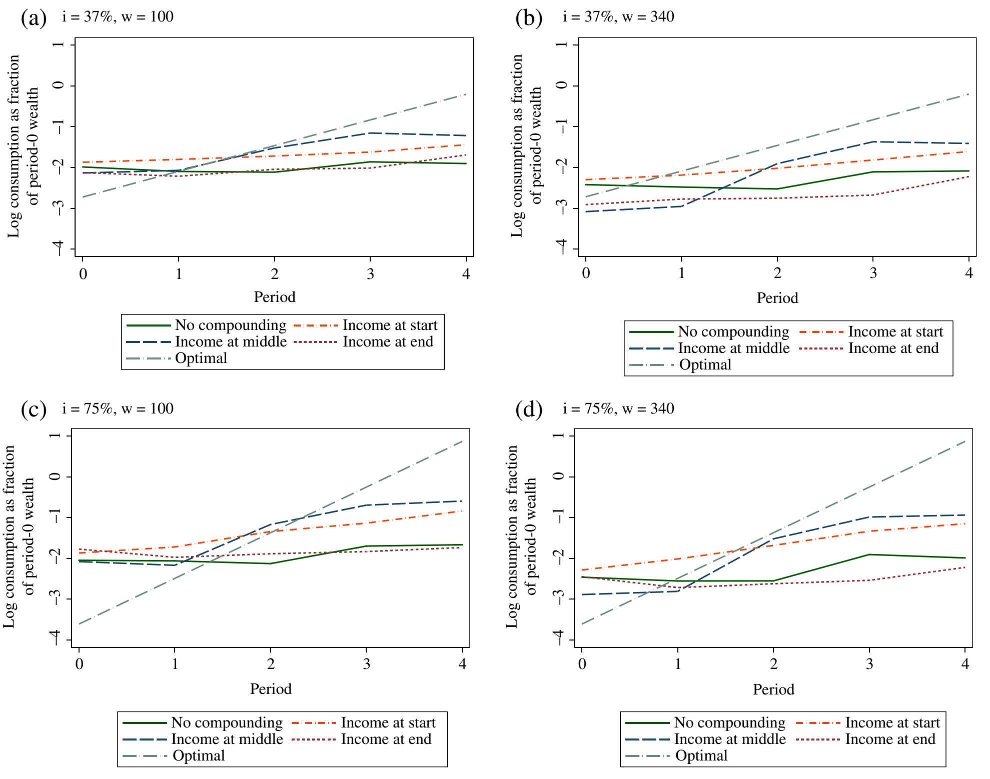

FIGURE 1. Consumption paths. *Notes*: Average log of consumption divided by lifetime wealth, by period, interest rate and present value of wealth. Consumption paths are constrained to be feasible under the true budget constraint by setting consumption to zero in periods after all resources are exhausted. [Colour figure can be viewed at wileyonlinelibrary.com]

too much in problem type 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 for a given interest rate vector, an unbiased economic agent should exhibit the exact same consumption profile across these problem types. Specifically, for an unbiased agent, all six paths in Figure 1 panels (a) and (b) should be the same, and all six paths in panels (c) and (d) should be the same. Examination of Figure 1 instead shows systematic differences. First, as we compare panel (a) to panel (b), and panel (c) to panel (d)—both comparisons representing an increase in the present value of income—we see a systematic downward shift in consumption. This phenomenon is not predicted by either the classical model or EGB, and instead suggests that there is an additional magnitude effect in subjects' choices. This point deserves emphasis as it shows that our main exercise is far from testing a tautology if subjects systematically err even here. Second, as we compare the low interest rate panels (a) and (b) to the high interest rate panels (c) and (d), the qualitative patterns are similar (note that the scale of the vertical axis is compressed for the higher interest rate panels), suggesting no major interactions between problem types and the interest rate. Scanning across all four panels, each problem type seems to maintain the same broad shape.

Checking for overconsumption in period 0, we find that problem types S, M and E show overconsumption in every panel except for (b). As we compare consumption in period 0 between problem types S, M and E, we find that the predicted ranking from Hypothesis 3—with E as the highest consumption and S as the lowest—is inconsistent with the empirical pattern. We will explore this in greater depth in the statistical analysis below.

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. Figure 2 plots the cumulative distribution function of the natural logarithm of the ratio of subjects' period-0 consumption to the optimal period-0 consumption for that consumption problem: $\ln(c_0/c_0^*)$. This provides a simple measure of overconsumption: subjects overconsume relative to the optimum when this variable is positive, and underconsume when it is negative. Panel (a) displays the distributions for the low interest rate, and panel (b) displays the distributions for the high interest rate.

As problem type N has only a single non-zero interest rate and therefore EGB does not predict overconsumption, we predict that the distribution of responses to these problems will be centred around zero. While the median value is slightly higher than zero, the overall distribution indicates only a minor tendency to overconsume in period 0. Any overconsumption in problem type N, however is dwarfed by the overconsumption from problem types S, M and E, which have more than one period of non-zero interest. As a test of Hypothesis 2, we compare each distribution in panel (a) of Figure 2 to its high-interest counterpart in panel (b). We can clearly see that for the problem types that have compounding (S, M and E), there is a clear shift of the distributions to the right as predicted by the theory. The prediction of Hypothesis 3 is that the distribution will be shifted to the right as income is delayed: S should be left of M, which should be left of E. However, this pattern does not appear to be evident.

We unpack the sources of overconsumption in Table 2, using OLS specifications in all columns, clustering standard errors by subject. Columns (1) and (2) test Hypothesis 1, by comparing the initial consumption in problem type N (which had no compounding) against problem type S (which featured compounding, but like N had all income arrive immediately), and restricting the sample to period $t = 0$ to control for any within-problem dynamic effects such as reinforcement learning. Column (1) shows the regression

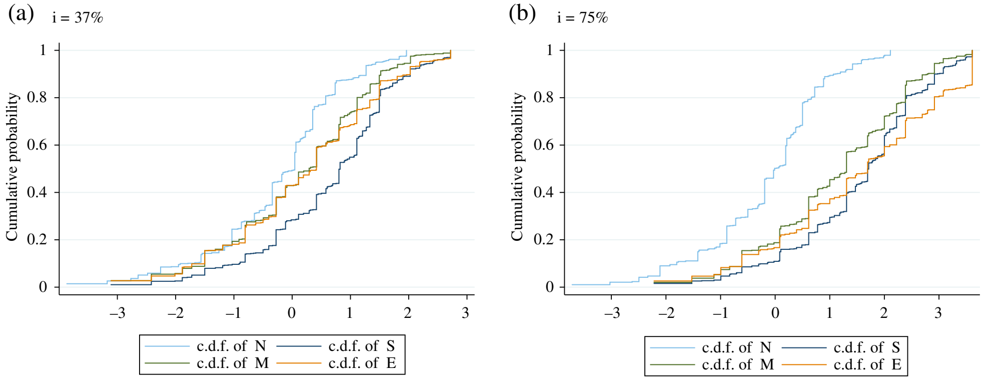

FIGURE 2. Overconsumption: c.d.f.s for all problem types. *Notes*: The graphs show the empirical distribution of $\ln(c_0/c_0^*)$: the natural logarithm of the ratio of subjects' choice of consumption for period 0 to the optimal consumption choice. Panel (a) restricts to the low (37%) interest rate problems, and panel (b) to the high (75%) interest rate problems. [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 2
OVERCONSUMPTION FROM COMPOUNDING

	i=37% (1)	i=75% (2)	i=37% (3)	i=75% (4)	Both (5)
Compounding	0.876***	1.678***	0.790***	1.600***	0.790***
	(0.044)	(0.045)	(0.042)	(0.045)	(0.042)
High wealth	−0.431***	−0.416***	−0.393***	−0.409***	−0.401***
	(0.042)	(0.042)	(0.034)	(0.035)	(0.029)
Question	−0.007	−0.014**	−0.011	−0.020**	−0.016**
	(0.007)	(0.007)	(0.007)	(0.008)	(0.007)
Period			−0.063***	−0.108***	−0.063***
			(0.014)	(0.016)	(0.014)
Compounding × Period			−0.287***	−0.513***	−0.287***
			(0.014)	(0.016)	(0.014)
i=75%					0.071**
					(0.031)
i=75% × Period					−0.045***
					(0.011)
i=75% × Compounding					0.810***
					(0.043)
i=75% × Compounding × Period					−0.226***
					(0.017)
Constant	0.042	0.197***	−0.052	0.115	−0.004
	(0.079)	(0.071)	(0.079)	(0.083)	(0.078)
Observations	1531	1537	7413	7428	14,841
Subjects	397	397	399	399	399

Notes Dependent variable is $\ln(c_t/c_t^*)$, i.e. the log of the ratio of actual consumption to optimal consumption given past behaviour. Columns (1) and (2) restrict to period 0 only. Columns (3)–(5) are for all periods (before t=5). All specifications restrict to consumption problems N and S. Standard errors clustered by subject.
*, **, ***indicate p<0.1, p<0.05, p<0.01, respectively.

results for the low interest rate consumption problems, and column (2) does the same but for high interest rate consumption problems. Columns (1) and (2) show that, averaged across all rounds, the inclusion of multiple non-zero interest rates (indicated by the variable ‘Compounding’) more than doubles subjects’ initial consumption choices as a percentage of the optimal level of consumption. This confirms Hypothesis 1. The doubling of the coefficient on EGB in column (2) compared to column (1) confirms Hypothesis 2: higher interest rates lead to more overconsumption. Finally, the negative coefficient on lifetime wealth confirms the magnitude effect that subjects exhibit in addition to EGB.

We next 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 actual consumption history to that point). Columns (3) and (4) of Table 2 repeat the analysis in columns (1) and (2) but use periods 0–4 and include a linear control for period and an interaction with the ‘Compounding’ indicator. The EGB effect remains large and positive. The coefficient on the ‘Compounding × Period’ interaction indicates that the EGB effect weakens in later periods exactly as theory would predict, since there is less compounding in later periods. The much higher coefficient on ‘Compounding’ in column (4) relative to column (3) again shows that higher interest rates lead to more overconsumption. Finally, the variable

'Question' is a linear control for question order. The coefficient is insignificant for the low interest rate and significant for the high interest rate. This suggests that there is some modest systematic learning over the course of the experiment. Across all specifications, the coefficient is quite small and we can rule out even moderately large effects. This negative coefficient, moreover, suggests little or no subject fatigue over the course of the experiments since responses are getting better on average. Using the coefficients in columns (3) and (4), the effect between the first and the last (16th) problem is as high as -0.165 for the low interest rate and -0.30 for the high interest rate—less than one-quarter of the main effect of 'Compounding'. Given the limited feedback and moderate stakes here, however, it is plausible that this understates the potential for learning in real life.

Finally, column (5) of Table 2 directly tests the effect of higher interest rates. The sample is restricted to problem types N and S, as in the previous columns, but now includes both levels of interest rate. The main effect of the high interest rate is significant and positive, though at 0.07 it suggests that the effect in the non-compounding problems is relatively limited. In contrast, the interaction of high interest with 'Compounding' is more than ten times greater, at 0.81. This large and significant coefficient confirms Hypothesis 2. Moreover, the triple interaction of high interest with 'Compounding' and 'Period' is negative, in line with the theoretical prediction, as the effect of a longer horizon is greater when there is greater compounding. Finally, we again find evidence of some limited learning across questions.

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 problem type N and problem types S, M and E, given the steep optimal consumption profile in the latter three. We show this in Section III.

From Hypothesis 3 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 problem types S, M and E. Hypothesis 3 says that overconsumption will be greater when all income is received later, and thus implies greater overconsumption in M than in S, and in E than in M. Oversmoothing 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).

In Table 3, we regress our measure of overconsumption, $\ln(c_t/c_t^*)$, on indicators for whether the income is received at the start (omitted category), received at the middle, or received at the end. Hypothesis 3 says that the coefficients should be positive and that the coefficient on E should be larger than the coefficient on M. In column (1) we restrict the sample to period 0 to control for any feedback effects. We find that the coefficient on E is significantly larger than the coefficient on M ($p<0.01$), in line with the income-delay prediction of Hypothesis 3. However, we surprisingly find that both coefficients are significantly negative. In column (2) we expand the sample to all periods with compounding for t periods 0–3, and we include a period control and interact it with the problem type indicators. The pattern in column (1) is reproduced in column (2) for $t = 0$ consumption. However, the interaction terms with period have opposite signs. For later periods, $t = 3$ for example, consumption is highest in M, then S, and E last.

This perplexing pattern led us to speculate whether people exhibited a strong borrowing aversion. For column (3) of Table 3, we include an indicator for periods in which consumption requires borrowing: periods 0 and 1 for problem type M, and periods

TABLE 3
EFFECT OF DEFERRED INCOME ON CONSUMPTION

	(1)	(2)	(3)
M	−0.466***	−0.615***	0.224***
	(0.052)	(0.056)	(0.058)
E	−0.235***	−0.185***	0.581***
	(0.055)	(0.056)	(0.070)
Period		−0.546***	−0.546***
		(0.017)	(0.017)
M × Period		0.236***	−0.070***
		(0.023)	(0.020)
E × Period		−0.297***	−0.297***
		(0.022)	(0.022)
Must borrow			−0.766***
			(0.058)
Constant	0.886***	1.025***	1.025***
	(0.048)	(0.046)	(0.046)
p-value: M = E	0.000	0.000	0.000
Observations	4307	16,812	16,812
Subjects	397	398	398

Notes
Dependent variable is $\ln(c_t/c_t^*)$, i.e. the log of the ratio of actual consumption to optimal consumption given past behaviour. Only problem types with compounding are included: S, M and E. Column (1) includes only observations from period 0. Columns (2) and (3) include all observations from periods 0–3. Standard errors clustered by subject.
*, **, *** indicate $p<0.1$, $p<0.05$, $p<0.01$, respectively.

0–4 for problem type E. The large negative coefficient reveals a strong borrowing-aversion heuristic in subjects' behaviour. When this effect is controlled, the coefficients on M and E become positive, with the coefficient on E larger than the coefficient on M, consistent with the theoretical predictions of EGB. This exploits within-problem variation in whether subjects have received that problem's endowment, and thus overcomes the borrowing aversion. Whether EGB or borrowing aversion dominates in the real world may depend on context. There is a large space of income vectors for which borrowing is not necessary to finance consumption, and we speculate that in these contexts, the effect of deferring income would play a more prominent role.[12]

Experiment 2: Dynamic inconsistency

Design The purpose of Experiment 2 is to test the effect of EGB on committed relative to uncommitted consumption. Intuitively, a biased agent will make worse choices when the planning horizon is longer, and therefore may adjust their plans in subsequent periods. Moreover, if provided with information about the actual compounding of their current assets or debts, they may further revise their plans. The theory makes a directional prediction here: a biased agent will change consumption plans in the direction of the current balance. Because biased agents are surprised by how rapidly debts or savings grow, the agent will update consumption plans in the direction of the balance. The consumption plan will be shifted upward in each period if there are savings, and it will be shifted downward in each period if there is debt. We test this hypothesis by comparing subjects who choose their consumption period-by-period to subjects who must state a committed consumption plan for the entire consumption problem. We use the two consumption problems in Experiment 2 as shown in Table 4.[13]

Experiment 2 assigned each subject to one of two arms (see Figure 3). 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 as in Experiment 1.[14] Such information is redundant to a classical decision-maker but is predicted to affect behaviour for a biased one.

In contrast to the dynamic arm, in the static arm subjects stated their consumption in each period without any feedback. Effectively, a subject in the static arm chose their consumption plan for all periods simultaneously.[15] 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 behaviour in the static arm elicits the subject's perceived optimal consumption plan at $t = 0$. Comparing behaviour 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 five rounds. The first round of each consumption problem is described in Table 4. 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 three-period plan given $\vec{\imath} = \langle 75\%, 75\% \rangle$ and $\vec{y} = \langle 100, 0, 0 \rangle$. In every round of consumption problems 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.

This laboratory experiment was conducted at the UCLA California Social Sciences Laboratory (CASSEL). Subjects were not provided with any tools, and calculators and mobile 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. 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.

TABLE 4
CONSUMPTION PROBLEMS IN EXPERIMENT 2

Consumption problem		Income vector	Interest vector
N	No compounding	⟨100,0,0,0,0,0⟩	⟨0,0,100%,0,0⟩
S	Income at start	⟨100,0,0,0,0,0⟩	⟨75%,75%,75%,75%,75%⟩
M	Income at middle	⟨0,0,100,100,0,0⟩	⟨75%,75%,75%,75%,75%⟩
E	Income at end	⟨0,0,0,0,0,500⟩	⟨75%,75%,75%,75%,75%⟩

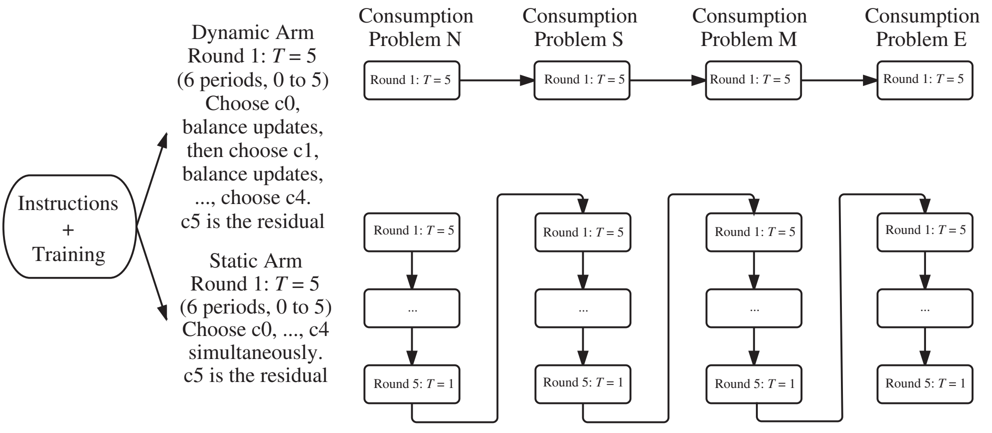

FIGURE 3. Experiment 2—design. *Notes*: Each round of length T has T+1 periods. Subjects in the static arm choose consumption vectors for a total of 20 rounds. Subjects in the dynamic arm choose consumption vectors with feedback on their current balance for a total of 4 rounds. The orders of consumption problems, and rounds within a consumption problem in the static arm, were randomized by subject.

Hypotheses Consumption plans should revise in the direction of the balance.[16] 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 behaviour 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 as a lump sum at $t = 0$, and this initial income finances all future consumption. In contrast, in consumption problem E the balance is always negative since income is received as a lump sum at $t = T$=5, 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 in the static arm (for the full-length $T = 5$ round), but the opposite should be true in consumption problem E.

Hypothesis 4 (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 less for the dynamic arm than the static arm in consumption problem E.

Results Although we will use only consumption problems S and E to assess subjects' dynamic inconsistency, we can first confirm the results of Experiment 1 using all four consumption problems. We find that overconsumption is again significantly higher in S relative to N, in M relative to S, and in E relative to M. However, because Experiment 2 does not control for lifetime wealth across consumption problems, we omit these results.

Because consumption problem S has all income received in period $t = 0$, subjects will always carry a weakly positive balance. Hypothesis 4 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 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 to 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 to c_4, with c_5 as a residual, without the opportunity for any revisions—and did so for each of five 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 5. 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 1 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 4. 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 × Period' interaction is large and positive, indicating that dynamic subjects' actual consumption choices in later periods are substantially

TABLE 5
DYNAMIC VS. STATIC PLANS

	Income at start (1)	Income at start (2)	Income at end (3)	Income at end (4)
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 × Period >0	0.649***		−0.575***	
	(0.244)		(0.178)	
Period		0.109**		0.085**
		(0.048)		(0.033)
Dynamic × 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)
Observations	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 $\ln(0)$ is undefined, these subjects are dropped. Standard errors clustered by subject.
*, **, *** indicate p<0.1, p<0.05, p<0.01, respectively.

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 two periods, the dynamic consumption has grown by $e^{0.257\times 2} \approx 67\%$ relative to the static consumption level). The opposite pattern emerges in column (3), where the 'Dynamic × (Period >0)' interaction is again large and significant, but now negative—indicating that dynamic subjects' 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 × 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. To test this, we can compare the fraction of lifetime wealth consumed by dynamic subjects in each period with the fraction that static subjects would have consumed if they were allowed to re-optimize their plans. That is, we construct a vector comprising the share of lifetime wealth consumed in period 0 by static subjects in the $T = 5$ lifecycle, the share consumed in period 0 in the $T = 4$ lifecycle, and so on. Theoretically, this is a valid approach because the objective function is homothetic. One constraint, however, is that we can construct this only 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 in order to infer their perceived lifetime wealth.

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 behaviour of the $T = 4$ round to compute c_1^s, which was 5% of wealth. Remaining wealth in the simulated consumption path is $0.98w(1+i_0) = 98w(1+0.75) = 1.715w$, and c_1^s is 5% of this which is $0.08575w$. A similar process is used to compute c_2^s to c_5^s. These values are all then normalized by the total wealth.

We can then compare the constructed vector of consumption shares for static subjects with the actual consumption decisions of dynamic subjects by regressing it on the period 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, has no impact on consumption beyond the prediction of Hypothesis 4. Whereas the previous results from Table 2 show that there is limited learning across the experiment, this result shows that there is no learning *within* a consumption problem.

Finally, we note that to earn a reasonable payment, subjects simply had to save enough so that they would have non-zero 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 that their debts were growing faster than they anticipated, in effect it was too late for them to get back on track.

Experiment 3: Compounding frequency

Design Our focus in Experiment 3 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 3 is very similar to that of 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 Experiment 1. The treatment presents problems that are equivalent, but rather than being framed as $T = 5$ periods long with an i% interest rate, problems are framed as $T = 50$ periods long with an $i^{1/10} = 5.75\%$ interest rate, and consumption opportunities every ten periods.[18] 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 ten-day segment. The treatment task is therefore a pure framing manipulation relative to the control: the incentives and choice set are identical.[19] The full instrument is available in the Online Appendix.

Because the information on the current balance is typical in consumption decisions, the design emulates Experiment 1 and the dynamic arm from Experiment 2. Table 6 lists all of the consumption problems in Experiment 3. The consumption problems were blocked 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.

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 and mobile 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 mean incentive payment was \$16.94; control subjects averaged \$16.44, while treatment subjects averaged \$17.26.

Hypotheses A biased agent's underestimation 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 1 in the Appendix). For example, a loan compounding at 1% interest per day may not sound like much to a biased person, but the equivalent loan at $1.01^{364} = 3,641\%$ interest per year may sound like a lot.

Hypothesis 5 (Framing) Normalized overconsumption, given by $\ln(c_0/c_0^*)$, will be greater in the treatment group than in the control group. The effect will be weakest when all income is immediate.

This effect is driven by both the price effect and the wealth effect of EGB. However, the wealth effect will be present only 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.

TABLE 6
CONSUMPTION PROBLEMS IN EXPERIMENT 3

Consumption problem	Income vector	Interest vector control
N_1	⟨100,0,0,0,0,0⟩	⟨0,0,100%,0,0⟩
N_2	⟨100,0,0,0,0,0⟩	⟨0,0,200%,0,0⟩
N_3	⟨0,0,0,0,0,900⟩	⟨0,0,100%,0,0⟩
N_4	⟨0,0,0,0,0,900⟩	⟨0,0,200%,0,0⟩
S_1	⟨100,0,0,0,0,0⟩	⟨75%,75%,75%,75%,75%⟩
S_2	⟨300,0,0,0,0,0⟩	⟨75%,75%,75%,75%,75%⟩
S_3	⟨500,0,0,0,0,0⟩	⟨75%,75%,75%,75%,75%⟩
S_4	⟨700,0,0,0,0,0⟩	⟨75%,75%,75%,75%,75%⟩
M_1	⟨0,0,100,100,0,0⟩	⟨75%,75%,75%,75%,75%⟩
M_2	⟨0,0,300,300,0,0⟩	⟨75%,75%,75%,75%,75%⟩
M_3	⟨0,0,500,500,0,0⟩	⟨75%,75%,75%,75%,75%⟩
M_4	⟨0,0,700,700,0,0⟩	⟨75%,75%,75%,75%,75%⟩
E_1	⟨0,0,0,0,0,500⟩	⟨75%,75%,75%,75%,75%⟩
E_2	⟨0,0,0,0,0,300⟩	⟨75%,75%,75%,75%,75%⟩
E_3	⟨0,0,0,0,0,800⟩	⟨75%,75%,75%,75%,75%⟩
E_4	⟨0,0,0,0,0,1200⟩	⟨75%,75%,75%,75%,75%⟩

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

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 c.d.f.s of our preferred measure of overconsumption, $\ln(c_0/c_0^*)$. Panel (a) shows significant overconsumption in both frames, which suggests that the eye-catching 75% interest rates are not misleading subjects significantly more than the 5.75% frame. On 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 5. A Kolmogorov–Smirnov test strongly rejects the null of identical c.d.f.s between control and treated subjects at $p = 0.018$.

In panels (b) and (c) of Figure 4, we disaggregate the distributions into problem type S and problem types M and E, respectively. These panels show that the first-order

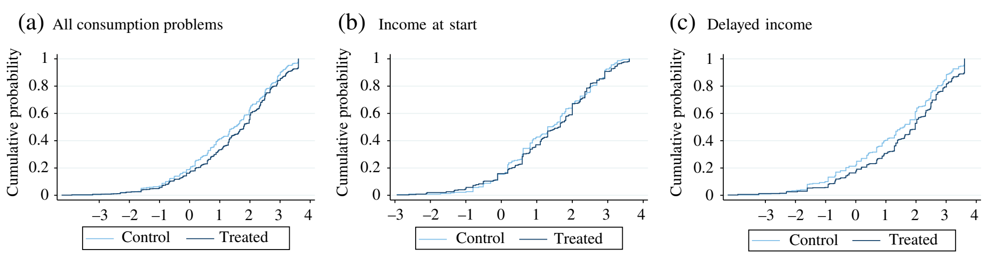

FIGURE 4. Overconsumption and framing. *Notes*: The graphs show the empirical distribution of $\ln(c_0/c_0^*)$: the natural logarithm of the ratio of subjects' choice of consumption for period 0 to the optimal consumption choice. Panel (a) shows all consumption problems; panel (b) restricts to consumption 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. [Colour figure can be viewed at wileyonlinelibrary.com]

stochastic dominance relationship is driven primarily by problem types 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]

In Table 7 we unpack the effect of the treatment on mean overconsumption. We now include all consumption choices, using our preferred measure of overconsumption, $\ln(c_t/c_t^*)$. 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 problem type S, where the framing effect is predicted to be weak. In column (2), we therefore include a dummy for problem type 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 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.

III. Alternative Explanations and Rules of Thumb

In this section we consider alternative explanations for the observed pattern of behaviour. 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, problem types S, M and E all have the same incentives, 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 (for a summary of the literature, see Winter *et al.* 2012; Hey and Knoll 2011), so we briefly consider some plausible candidates.

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 problem types N and S, *Never borrow* does not restrict behaviour, hence it does not predict any effect including effects from compounding (Hypotheses 1 and 2). In problem type M, *Never borrow* requires consuming zero until $t = 2$, and consuming weakly less than y_2. In problem type

TABLE 7
EFFECT OF FRAMING ON OVERCONSUMPTION

	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.197	0.269*	0.311**	0.351	1.406*	1.807**
	(0.158)	(0.152)	(0.145)	(0.225)	(0.746)	(0.877)
Treated × All at start		−0.156	−0.261		−0.343	−1.414
		(0.141)	(0.178)		(0.334)	(1.080)
Treated × Period =1		−0.014	−0.026		−1.074*	−1.524**
		(0.060)	(0.081)		(0.574)	(0.682)
Treated × Period =2		−0.140	−0.138		−1.181*	−1.616*
		(0.097)	(0.133)		(0.654)	(0.825)
Treated × Period =3		−0.103	−0.221		−1.323*	−1.897**
		(0.131)	(0.164)		(0.681)	(0.853)
Treated × Period =4		0.168	0.060		−1.151	−1.714*
		(0.166)	(0.196)		(0.708)	(0.877)
Treated × All at start			0.034			1.203
× Period =1			(0.096)			(0.782)
Treated × All at start			0.004			1.162
× Period =2			(0.163)			(0.984)
Treated × All at start			0.269			1.530
× Period =3			(0.180)			(1.044)
Treated × All at start			0.239			1.499
× Period =4			(0.206)			(1.066)
Observations	6489	6489	6489	7865	7865	7865
Subjects	100	100	100	100	100	100

Notes Dependent variable is $\ln\ (c_t/c^{\backslash\mathrm{label}}*_t)$ in columns (1)–(3), and c_t/c_t^* in columns (4)–(6). Domain dummies (when income is received × compounding) are included in all specifications.
*, **, *** indicate $p<0.1$, $p<0.05$, $p<0.01$, respectively.

E, *Never borrow* implies zero consumption until the final period. We find that only 5.8% of observations follow this rule in problem type E of Experiment 1. However, 10.0% of observations follow *Never borrow* in problem type M of Experiment 1. Thus *Never borrow* does seem to play some role. Likewise, as stated in the ‘Results’ subsubsection in the ’Experiment 1’ subsection of Section II, we find evidence of aversion to borrowing. A *Never borrow* heuristic would be an extreme version of aversion to borrowing. There we found that the borrowing-aversion effect was larger than the effect of delaying income. On net, the two effects would lead to a 17% decrease in consumption rather than the predicted increase.

Thus while we find that *Never borrow* helps to explain some of our findings, we interpret it as an additional factor influencing subjects’ choices on top of EGB. In particular, *Never borrow* does not predict overconsumption (Hypotheses 1 and 2) dynamic inconsistency (Hypothesis 4) or framing (Hypothesis 5), which are core to the perceptual error of EGB.

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 2—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.

Both Hypothesis 1 and Hypothesis 2 are consistent with this rule. The optimal consumption vector for consumption problem N is flatter than the optimal consumption vector for consumption problem S, so c_t^* is lower for S than for N. Hence a *Constant* rule will on average overconsume for N and overconsume even more for S. We drop all *Constant* subjects in Table 1 of the Online Appendix, and find that the increased consumption in consumption problem S relative to N persists even among subjects not employing this rule of thumb. In fact, the results barely change. Therefore a *Constant* rule cannot be the sole driver of our results.

Furthermore, this rule of thumb does not predict Hypotheses 3–5. Because the interest rates are the same for S, M and E, *Constant* predicts no difference in behaviour 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 behaviour. A person who applies *Constant* and suffers from EGB will still satisfy Hypothesis 3. Thus the predictions of the wealth effect of EGB are preserved even though the person may not dynamically optimize given perceptions. Similarly, for Hypothesis 4, dynamic inconsistency, if a *Constant* person estimates the budget properly, then there should be no change given that the rule is literally constant consumption. For Hypothesis 5, framing, a person who is unbiased in estimating the budget should not be affected by the framing manipulation, thus a *Constant* person should not exhibit a treatment effect. Thus while a significant number of subjects exhibit this rule of thumb, their behaviour still cannot be explained without the addition of EGB.

Save all

This is perhaps the easiest rule of thumb to implement, and in fact would have yielded a reasonable payment in our experiment. A subject using *Save all* will choose zero consumption every period and consume the residual in the final period. Among observations from Experiment 2, $T = 5$, the static arm—, only 1.3% of observations exhibit this behaviour. Moreover, *Save all* cannot be a confound because it does not imply the same behaviour as EGB. For all consumption problems, *Save all* implies constant behaviour. Thus there should be no effect of varying the consumption problems or of changing the framing, so this too cannot explain our main results.

Interest rate neglect

Perhaps subjects do not respond to interest rates at all. This does not appear to be the case. Using Experiment 2, $T = 5$, the static arm, consumption problem N, we divide planned consumption before the interest rate in period 3 with consumption after. People should respond by increasing consumption after the one period with interest. The average per-period consumption before is 21.16, and after is 36.17, which a t-test shows are significantly different at $p<0.003$. In the sample, 68% (26/38) have lower planned per-period consumption in first three periods then in the subsequent two periods. Among the other 32% there is no clear discernible pattern regarding their consumption. So there is at least some evidence that a substantial minority, approximately one-third, do not even respond to interest rates in the correct direction.

Even for those who respond to the interest rate in the correct direction, it may still be the case that subjects underestimate the effect of the interest rate on the budget. We cannot disentangle this from mistakes in optimizing the objective function. Subjects have a tendency to oversmooth, and this could be caused by either EGB or interest rate neglect.

IV. Discussion

In this paper, we tested the theoretical predictions of EGB in a set of controlled laboratory consumption–savings experiments. Four of the five main predictions—overconsumption from compound interest, higher consumption from higher compounding interest rates, dynamically inconsistent planning, and higher consumption from more frequent compounding—are supported against the null hypothesis of no effect. While we find evidence for our fifth prediction—the effect of higher consumption from delayed income—in the context of our experimental design this effect is confounded by other forces, most notably an aversion towards borrowing. These findings are of interest insofar as they provide insight into human judgment and decision-making, but also as they may apply more generally beyond the laboratory. While the environment is idealized, we expect that the comparative statics driving the hypotheses will be present broadly. We conclude by briefly describing some real-world applications of our findings, which may prove fertile ground for future research.

The increased consumption from deferred income (Hypothesis 3) could inflate demand for careers with back-loaded income trajectories. Biased people will overestimate the present value of future earnings. A biased person choosing between two income vectors of equivalent wealth will misperceive the one with later earnings as having higher wealth. Since advanced education leads to more back-loaded income trajectories (a delay in employment and debt for higher earnings later), EGB may play a significant role in human capital acquisition. Of course, many other factors play a role in education choice, but on the margin, a student who is close to indifferent between attending college or starting employment might choose the former route by not properly accounting for the present value of those higher future earnings and the cost of financing those college loans.

The dynamic inconsistency induced by EGB (Hypothesis 4) means that welfare losses are exacerbated if people commit their consumption. Generally, people get feedback on the repercussions of their consumption plans and can adjust accordingly: after a few months, a person sees how much their credit card debts have grown and can pay them off faster; after a few years, a person can see how well their stocks are performing and invest more; and so on. When consumption is committed, the initial plan, which exhibits the greatest degree of EGB, is implemented. Home ownership is a relevant application. While a person could in principle downsize by renting out rooms or expand through construction, both are limited and costly. Complex mortgage products could leave biased home owners with mortgages that grow more rapidly than expected. For example, biased borrowers will underestimate the jump in payments when an adjustable-rate mortgage experiences a rate adjustment. This is a particularly ripe domain for study given the long durations of the loans and hence high degree of compounding.

The effect of the frequency of framing in Hypothesis 5 has obvious implications. Framing a loan as compounding at a high-frequency low-interest rate will appear much more attractive than framing the loan as a low-frequency high-interest rate. This may be why payday loans typically advertise their products by emphasizing the monthly rate rather than the annual rate. In contrast, fund managers can better lure biased prospective investors by displaying aggregated returns.

Another concern regarding the relevance of these findings is whether individuals could simply learn away their bias. It is obvious that some people can and do learn away their bias. There are, however, both conceptual and empirical reasons to believe that most people do not fully learn away EGB. Conceptually, people live only one lifecycle—unlike in the experiment—thus the capacity to learn away EGB in a single lifetime is limited. People could in principle learn from the previous generation; however, the technological, financial and institutional contexts change greatly in generational timespans.[21] Moreover, research has shown that EGB is widespread. Levy and Tasoff (2016) and Goda *et al*. (2015) measure EGB in representative samples of the US population. Levy and Tasoff (2016) find that 33% of the US adult population can be be characterized as fully biased, and at least 80% have some bias. Goda *et al*. (2015) find that approximately 32% of the US adult population can be characterized as fully biased, and 39% as partially biased. Empirically, it seems that learning is quite limited.

As is often the case when behavioural 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.[22] The choice of compounding frequency used to present financial products is even more well-rooted in firm practice. A lender will want to use as high frequency a frame as possible, the 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.[23] 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.

APPENDIX: COMPOUNDING FREQUENCY FRAMING EFFECTS

Proposition 1 (Interest framing) Consider a single period of interest $i>0$, 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 $n>1$, we have the following:

(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 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:

$$\hat{V}_{0,n}(\vec{y},\vec{\imath};\alpha) = y_0 \cdot \left(\prod_{s=1}^{n}(1+\alpha i_n) + \sum_{s=1}^{n}(1-\alpha)i_n\right) < y_0 \cdot \left(\prod_{s=1}^{n}(1+i_n)\right) = y_0 \cdot (1+i_0) = \hat{V}_{0,n}(\vec{y},\vec{\imath};1).$$

Part (ii) also results from straightforward algebra:

$$\text{(A1)} \qquad \hat{V}_{0,n}(\vec{y}^{(n)},\vec{\imath}^{(n)};\alpha) \geq y_0 \cdot \left(1+\sum_{s=1}^{n} i_n\right) = y_0 \cdot \left[1+n\left((1+i_0)^{1/n}-1\right)\right].$$

We first note that the right-most term in (A1) is a strictly decreasing function in n for n>1, as

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

since

$$\left(1-\frac{\ln(1+i)}{n}\right)^{n} < \frac{1}{1+i} \qquad \text{for } i>0,\ n>1.$$

We thus take the limit to find a lower bound, and find

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

ACKNOWLEDGMENTS

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 Brian Leonard, Hisam Sabouni, 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), from 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’.

NOTES

1. Stango and Zinman (2009), Levy and Tasoff (2016), and Goda *et al.* (2015) find that EGB is correlated with several financial outcomes, but they note that this is a correlation that may be subject to non-causal interpretations.
2. We present the parametric model of Levy and Tasoff (2016), but note that all predictions follow from their general formulation of EGB.

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 for behaviour that are explored extensively elsewhere.
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. An experimental economics norm recommends avoidance of loaded labels for the legitimate concern that they may reduce the relative salience of monetary incentives. There are two reasons why this is not a concern here. First, all of our predictions are comparative static predictions. Since all conditions use the same labels, any effect appears under all conditions and is differenced out using the control condition. Second, the labels contain no connotations that are relevant to consumption or savings.
5. See Proposition 3 of Levy and Tasoff (2016), which states that a biased agent with EIS greater than 1 will overconsume in $t = 0$ for any income vector.
6. A programming error in our experiment code misspecified the optimal utility and minimum utility when calculating payments. In particular, the value u_o used in the calculation and final feedback screen always exceeded its true value. This had no effect on subjects' behaviour, however, as it came at the end of the experiment and, moreover, their payment was still strictly increasing in the utility that they achieved.
7. We thank an anonymous referee for this design suggestion. We chose the two wealth levels so that we could also make a direct comparison between a consumption problem with high and low interest rates in which the future value was fixed: i.e. E_{HL} and E_{LH} have the same income vector and the same future value but a different interest rate. This allows us to test the effect of compounding keeping future value fixed.
8. 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.
9. Below, for consistency with Hypotheses 1 and 2, we formulate the next hypothesis in terms of log-normalized *over*consumption; but it is equivalent to using log-normalized consumption.
10. This follows from Proposition 1 in Levy and Tasoff (2016).
11. 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 $\langle -2.0, -2.0, -2.0, -1.3, -1.3\rangle$ when $i = 37\%$ and $\langle -2.2, -2.2, -2.2, -0.8, -0.8\rangle$ when $i = 75\%$.
12. Moreover, the high levels of consumer debt in the USA suggest that perhaps borrowing aversion is not always relevant in real-world choices.
13. Experiment 2 is in fact the first experiment that we ran. We originally intended to test Hypotheses 1–3 using these data. However, the new Experiment 1 supersedes these data as it both has many more observations and experimentally controls for the present value of wealth. Indeed, contrary to theory, we found that there are strong wealth effects.
14. Figure 13 of the Online Appendix displays the user interface seen by subjects in the dynamic arm.
15. Figure 16 of the Online Appendix displays the user interface seen by subjects 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 explicate this relationship.
16. This follows from Proposition 4 in Levy and Tasoff (2016).
17. Each of the 82 subjects makes a decision for each of five periods (i.e. periods 0–4) for a given consumption problem. Observations with $c_0 = 0$ are dropped.
18. We again employed the metaphor of feeding a dog to increase subjects' comprehension. Treated subjects were told that the dog eats only every 10 days—i.e. on day 0, day 10, and so on—and that the problem ends on day 50.
19. The effect that we are testing is reminiscent of the subadditive discounting of Read (2001) in which an experiment shows that dividing a time period into more segments leads to greater total discounting. Our design is different conceptually since there is no discounting—only perception of compounding—and also operationally as we do not actually increase the number of periods in which a subject makes a decision. The difference between treatment and control is pure framing.
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.
21. For instance, technology has extended lifespans in each generation, new financial assets such as index funds, invented in the 1970s, dramatically changed the optimal personal investment strategy, and the replacement of defined benefit programmes with defined contribution programmes has required a new set of considerations for savers. Directly emulating successful people from previous generations is therefore suboptimal.

22. For example, Vanguard compares the total fees charged over 10 years on a $10,000 investment in their funds with the fees charged by the average fund in the category.
23. Although less popular in the USA, 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.

REFERENCES

Almenberg, J. and Gerdes, C. (2012). Exponential growth bias and financial literacy. *Applied Economics Letters*, **19**, 1693–6.

Anderhub, V., Güth, W., Müller, W. and Strobel, M. (2000). An experimental analysis of intertemporal allocation behavior. *Experimental Economics*, **3**, 137–52.

Benzion, U., Granot, A. and Yagil, J. (1992). The valuation of the exponential function and implications for derived interest rates. *Economic Letters*, **38**, 299–303.

Brown, A. L., Chua, Z. E. and Camerer, C. F. (2009). Learning and visceral temptation in dynamic saving experiments. *Quarterly Journal of Economics*, **124**, 197–230.

Eisenstein, E. M. and Hoch, S. J. (2007). Intuitive compounding: framing, temporal perspective, and expertise. Fox School of Business Working Paper, Temple University, Philadelphia, PA.

Ensthaler, L., Nottmeyer, O., Weizsäcker, G. and Zankiewicz, C.. (2013). Hidden skewness: on the difficulty of multiplicative compounding under random shocks. DIW Berlin Discussion Paper no. 1337.

Goda, G. S., Manchester, C. F. and Sojourner, A. (2014). What will my account be worth? Experimental evidence on how retirement income projections affect saving. *Journal of Public Economics*, **119**, 80–92.

———, Levy, M. R., Manchester, C. F., Sojourner, A. and Tasoff, J.. (2015). The role of exponential-growth bias and present bias in retirement savings decisions. NBER Working Paper no. 21482.

Hey, J. D. and Dardanoni, V. (1988). Optimal consumption under uncertainty: an experimental investigation. *Economic Journal*, **98**(390), 105–16.

——— and Knoll, J. A. (2011). Strategies in dynamic decision making—an experimental investigation of the rationality of decision behaviour. *Journal of Economic Psychology*, **32**, 399–409.

Johnson, S., Kotlikoff, L. J. and Samuelson, W. (1987). Can people compute? An experimental test of the life cycle consumption model. NBER Working Paper no. 2183.

Kahneman, D. and Klein, G. (2009). Conditions for intuitive expertise. *American Psychological Association*, **64**(6), 515–26.

Keren, G. (1983). Cultural differences in the misperception of exponential growth. *Perception and Psychophysics*, **34**(3), 289–93.

Levy, M. R. and Tasoff, J. (2016). Exponential-growth bias and lifecycle consumption. *Journal of the European Economic Association*, **14**(3), 545–83.

MacKinnon, A. J. and Wearing, A. J. (1991). Feedback and the forecasting of exponential change. *Acta Psychologica*, **76**, 177–91.

McKenzie, C. R. M. and Liersch, M. J. (2011). Misunderstanding savings growth: implications for retirement savings behavior. *Journal of Marketing Research*, **48**, S1–S13.

Meissner, T. (2016). Intertemporal consumption and debt aversion: an experimental study. *Experimental Economics*, **19**(2), 281–98.

Noussair, C. and Matheny, K. (2000). An experimental study of decisions in dynamic optimization problems. *Economic Theory*, **15**, 389–419.

Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty,* **23**(1), 5–32.

Smith, V. L. (1976). Experimental economics: induced value theory. *American Economic Review: Papers and Proceedings*, **66**(2), 274–9.

Soll, J. B., Keeney, R. L. and Larrick, R. P. (2011). Consumer misunderstanding of credit card use, payments, and debt: causes and solutions. Working Paper.

Song, C. (2015). Financial illiteracy and pension contributions: a field experiment on compound interest in China. SSRN Working Paper.

Stango, V. and Zinman, J. (2009). Exponential growth bias and household finance. *Journal of Finance*, **64**(6), 2807–49.

Wagenaar, W. A. and Sagaria, S. D. (1975). Misperception of exponential growth. *Perception and Psychophysics*, **18**(6), 416–22.

——— and Timmers, H. (1979). The pond-and-duckweed problem: three experiments on the misperception of exponential growth. *Acta Psychologica*, **43**, 239–51.

WINTER, J. K., SCHLAFMANN, K. and RODEPETER, R. (2012). Rules of thumb in life-cycle saving decisions. *Economic Journal*, **122**, 479–501.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

A Rules of Thumb—Sample Excluding *Constant* Observations
B Instrument
B.1 Experiment 2
B.2 Treatments