Exponential-Growth Bias and Lifecycle Consumption*

Matthew R. Levy
London School of Economics

Joshua Tasoff
Claremont Graduate University

September 10, 2012

Abstract

Exponential-growth bias (EGB) is the tendency for individuals to underestimate exponential growth due to the neglect of the role of compounding. We develop a theoretical model in which EGB causes consumers to mis-perceive the budget constraint. Because consumers dynamically misperceive prices, consumers will over-consume early in their lifecycle if the elasticity of intertemporal substitution is sufficiently large. This over-consumption is exacerbated when income is received later in the lifecycle, as biased agents will additionally mis-perceive the present value of their earnings. These effects can lead to over-indebtedness both for early consumption and for over-investment in human capital. Finally, we show that EGB leads to a form of dynamic inconsistency in which planned future consumption increases when the agent's asset balance is positive and decreases when the balance is negative. This inconsistency can lead to significant welfare losses if the consumers are able to commit lower bounds on future consumption, such as consumption commitments in housing – commitment which other forms of dynamic inconsistency, such as present-biased preferences, would predict to be welfare-enhancing. We directly test our theoretical predictions in a simulated lifecycle-consumption task, and confirm all the main predictions of the theory. We find strong evidence of over-consumption early in the lifecycle and this over-consumption is exacerbated when income is received later in the lifecycle. There is no evidence for learning over the experiment and continual dynamic feedback on one's current balance has virtually no effect. We then run the first incentivized study to measure the extent of exponential-growth bias in a representative sample of the U.S. population. We find substantial EGB, with approximately one-third of subjects estimated to be the fully-biased type. The magnitude of the bias is negatively associated with total asset accumulation, even controlling for income and education. We find that the bias is robust to a simple graphical intervention, and test our model of EGB against potential alternative mistakes. Finally, in a laboratory study with a sample of elite college students we again find systematic bias, but show that subjects are largely unaware and unwilling to pay for aid. This overconfidence suggests that markets will not correct for the bias, as consumers will not seek enough financial advice.

*We would like to gratefully acknowledge the financial support of the Flectcher Jones Foundation and Time-sharing Experiments for the Social Sciences (TESS). We are grateful to Masyita Crystallin, Peiran Jiao, Andrew Royal, Quinn Keefer, and Oliver Curtiss for research assistance.

1 Introduction

Many economists have observed that the savings rate in the U.S. is too low and personal debt too high. The average personal savings rate in the U.S. has ranged between 1% and 4% during 2004–2010,[1] and more than 40% of households are saving insufficiently to maintain their standard of living into retirement (Munnell et al., 2006). At the same time, revolving debt (predominantly credit card debt) in the United States reached \$827 billion as of May 2012, or roughly \$2,600 per capita[2], while payday loans with APRs often exceeding 400% have grown significantly.[3] While behavioral economists have proposed many compelling explanations for these behaviors, we suggest that a parsimonious explanation comes from "Exponential-growth bias" (EGB) – the psychological tendency for individuals to underestimate exponential growth due to the neglect of the role of compounding.

In this paper, we construct a model of EGB and derive and test several novel predictions. The consumer is modeled as believing her asset is divided into two accounts. She perceives a fraction α grows with compounding interest, and a fraction $1-\alpha$ grows with simple interest. Thus when $\alpha = 1$ the consumer has correct perceptions and when $\alpha = 0$ the consumer believes an asset with compounding interest grows linearly. This leads the consumer to make two fundamental errors regarding her intertemporal budget constraint. First, the consumer misperceives the value of her income over time. With positive interest rates this causes the consumer to overestimate the value of future income relative to present income. We call this the *budget effect* of exponential-growth bias. Consequently, shifting income to later periods in a way that preserves lifetime income will increase consumption in the present. Second, the consumer misperceives the relative prices of consumption over time. With positive interest rates this causes the consumer to overestimate the price of future consumption relative to present consumption. We call this the *price effect* of exponential-growth bias. We derive sufficient conditions under which the consumer will over-consume in the present for any positive income vector. Because the perception of future prices and lifetime wealth changes each period, the consumer will behave in a dynamically inconsistent manner that is distinct from the pattern generated by dynamically inconsistent time preference. Specifically, the consumer maintains a negative (positive) balance she underestimates the growth of her debts (assets) and so she revises downward (upward) her consumption plans.

While there are folk stories illustrating people's underestimation of exponential growth going back millennia,[4] to our knowledge, Wagenaar and Sagaria (1975) conducted the first published experiment demonstrating this phenomenon in the Psychology literature. Subsequent studies found the same pattern of underestimation (Wagenaar and Timmers, 1979; Keren, 1983; Benzion et al., 1992; Almenberg and Gerdes, 2011). Wagenaar and Sagaria (1975) wrote an early model of EGB that used two parameters in which an exponential function of the form $x(t) = ac^t$ is perceived as $\hat{x}(t) = \alpha a c^{\beta t}$. The model was subsequently used by Stango and Zinman (2009a) and Goda et al. (2012). This model captured some basic implications of underestimation, but had several undesirable properties. First, when taken literally, it implies that a biased consumer will

[1] http://www.bea.gov/briefrm/saving.htm

[2] Board of Governors of the Federal Reserve System website accessed on July 16, 2012: http://www.federalreserve.gov/releases/g19/HIST/cc_hist_r_levels.html and authors' calculation.

[3] The payday loan industry is so successful that U.S. brick-and-mortar payday loan locations exceed the number of McDonalds and Starbucks combined (Skiba and Tobacman, 2011).

[4] According to legend, the ruler of an Indian kingdom granted the inventor of chess a single boon. The inventor requested a quantity of rice that doubled for every square on the chessboard, starting with a single grain. The ruler quickly accepted the request only to later discover that the sum exceeded the kingdom's entire store.

underestimate exponential growth even after one period when interest has not yet compounded. Second, this model did not nest full neglect of compounding (misperceiving compound interest as simple interest), which we observe in about one third of our sample. Third, this model predicted that a biased individual will overestimate the value of an asset that depreciates with a fixed interest rate, whereas we find evidence of the opposite. Fourth, this model did not prescribe inclusion of interest rates that vary over time. Our new model satisfies this desiderata and is motivated by plausible cognitive considerations.

We test our predictions in a laboratory lifecycle-consumption experiment. Although there have been other laboratory lifecycle-consumption experiments (Johnson et al., 1987; Hey and Dardanoni, 1988; Brown et al., 2009), ours is the first to focus on the role of exponential-growth perceptions. Subjects make consumption choices in a life-cycle consumption environment with an explicit utility function and earn rewards as a function of their performance. All of the main predictions of the theory are confirmed. Subjects over-consume early in their lifecycle and go bankrupt at a substantial rate; over-consumption is dramatically exacerbated when income is received later in the lifecycle; and with commitment, consumption differs starkly from consumption without commitment due to EGB-induced dynamic inconsistency. We find no evidence for learning during the course of the experiment.

Of course the presence of EGB in the lab does not necessarily imply errors in the market. Conditions in the lab differ from those in which consumers make financial decisions: consumers have more time and tools to make the appropriate calculations and they have incentives to do so. Moreover, consumers can request or pay for advice in the market. Thus even though consumers may make errors when confronted with hypothetical financial questions, it need not be the case that they make these mistakes in practice. However Stango and Zinman (2009a) find that there is substantial bias in the 1977 and 1983 Survey of Consumer Finances, which they diagnose from one un-incentivized survey question phrased in natural terms. They find that those with higher bias have higher short-term debt to income ratios, proportionately less stock as a share of their assets, and lower net worth.

Experiments 2 and 3 establish that EGB is widespread, severe, and robust to the presence of graphical displays and the availability of costly tools. In our second experiment we correlate EGB with financial outcomes in a representative sample of the population. This is the first study to measure EGB using incentivized elicitations and the first to estimate a single model. Subjects answer quiz-like financial questions and are paid for accuracy. Using our model we estimate the accuracy α by subject and find that about one third of the population is fully-biased with $\alpha = 0$. The median bias is 0.53 and 96% of subjects are estimated to have an $\alpha < 1$ (i.e. underestimate compound growth). This is in spite of the fact that subjects participated online and had access to whatever tools (e.g. financial calculators, help from friends) that they chose to use. Various questions also produce "fingerprinted" EGB responses that are predicted by our model but not by the Wagenaar-Sagaria (WS) model: subjects correctly estimate exponential growth when there is only one period (the WS model predicts undertestimation), subjects underestimate the value of a depreciating asset (the WS model predicts overestimation), and subjects use the arithmetic mean of the interest rate to estimate the value of an asset rather than the geometric mean. Regressing log savings on α, we find that it enters positively and significantly while controlling for income, education, age, and other covariates. Going from full bias to full accuracy is associated with a ceteris paribus 62% increase in life savings. The evidence suggests that EGB is pervasive in the population, and an important predictor of savings behaviors.

A natural first reaction is to ask whether consumers can be de-biased. Several experimental interventions

have tried this with mixed results (MacKinnon and Wearing, 1991; Eisenstein and Hoch, 2007; McKenzie and Liersch, 2011; Soll et al., 2011; Goda et al., 2012). In contrast to previous experiments, we provide incentives for accuracy, and allow the free use of tools such as spreadsheets, financial calculators, or the opportunity to seek advice from friends. These are important distinctions given that solving financial problems is effortful, and for many people, quite dull, and many tools are available in the economic environment, and may have a large impact on behavior. We test the effect of presenting a graph of \$100 growing at the relevant interest rate over time. The intervention had no effect on performance suggesting the robustness of the bias.

Given the widespread availability and low costs of financial tools, such as financial calculators and financial advisors, it is surprising that this error seems to persist in the marketplace. A consumer need not be skilled in personal finance as long as she can outsource financial decisions to an expert. However, if consumers are overconfident in their financial acumen then they may not obtain financial advice when they should. Several hundred studies demonstrate overconfidence in various domains (Camerer and Lovallo, 1999; Moore and Healy, 2008). However, it has not been demonstrated that people are overconfident about their mathematical ability in the domain of financial decisions. The overconfidence results are tempered by the finding that difficult problems often lead to underconfidence (Moore and Healy, 2008). Our paper explicitly addresses this issue. In a laboratory experiment we elicit subjects willingness to pay (WTP) for a spreadsheet and WTP for the correct answer while answering incentivized quiz-like financial questions. A risk-neutral subject who expects to lose x on a question should be willing to pay up to x for the correct answer. The financial error on average costs subjects \$13.94, but their WTP for the correct answer is only \$5.76 suggesting that subjects believed they would earn 74% more than they actually did. Thus overconfidence acts as a significant damper on the market's ability to correct EGB through professional advice.

The next section presents the model. Section 3 describes the laboratory lifecycle-consumption experiment that tests the novel predictions of the theory. The overconfidence experiment is in Section 5. The paper concludes with Section 6.

2 Theory

2.1 Model

For ease of exposition, we will refer to an agent who exhibits exponential-growth bias as Eddie. Eddie faces a financial asset with an initial value P_0, a vector of interest rates $\vec{\imath} =< i_1, \ldots, i_T >$, and a vector of future cash flows (income, contributions, debits, etc.) $\vec{y} =< y_1, \ldots, y_T >$. Eddie correctly perceives the cash flows and interest rates, but mistakenly perceives the growth of the asset to be as though it were divided into two accounts. A fraction α of the interest accumulates in an account that will grow in future periods with the interest rate αi, and the remaining fraction $1 - \alpha$ of the accumulated interested is sequestered to a non-growing account (e.g. placed under Eddie's mattress). The parameter α thus denotes the accuracy of Eddie's perception. When $\alpha = 1$, Eddie correctly perceives the asset growing exponentially. When $\alpha = 0$, Eddie is fully biased and perceives the asset growing linearly according to simple interest. Thus both the correct model and full neglect of compounding interest are nested in the model.

Eddie's perception of the value of the asset is given by

$$\hat{V}_t(P_0, \vec{\imath}, \vec{y}) = \overbrace{P_0 \left(\prod_{s=1}^{t} (1 + \alpha i_s) \right)}^{\text{compounding interest and principal}} + \overbrace{\left(\sum_{s=1}^{t} y_s \prod_{r=s}^{t} (1 + \alpha i_r) \right)}^{\text{compounding interest and cash flows}}$$
$$+ (1 - \alpha) \left[\underbrace{\left(\sum_{r=1}^{t} i_r \right) P_0}_{\text{simple interest of the principal}} + \underbrace{\sum_{s=1}^{t} \left(\sum_{r=s}^{t} i_r \right) y_s}_{\text{simple interest of the cash flows}} \right] \tag{1}$$

When an asset receives no additional cash flows before it matures, $\vec{y} =< 0, \ldots, 0 >$, the formula simplifies considerably. Then equation (1) simplifies to:

$$\hat{V}(P_0, \vec{\imath}, \vec{0}) = P_0 \cdot \left[\prod_{t=1}^{T} (1 + \alpha i_t) + \sum_{t=1}^{T} (1 - \alpha) i_t \right].$$

When the interest rate is constant over time $\vec{\imath} =< i, \ldots, i >$, and $|\vec{\imath}| = T$, then

$$\hat{V}(P_0, \vec{\imath}, \vec{0}) = P_0 \cdot \left[(1 + \alpha i)^T + (1 - \alpha) i T \right].$$

A few observations are worth noting. The first is that Eddie should correctly predict the value of an asset after one period no matter his bias. Second, Eddie underestimates the value of an asset that depreciates at a constant rate. For example, an asset that depreciates at 10% a year would be worth about 59% of its original value after 5 years. However, fully biased Eddie would perceive the asset's value to be exactly 50% after 5 years. Furthermore, Eddie's perception of the mean interest rate from a vector of interest rates is not the geometric mean. When Eddie is fully biased his perception of the mean interest rate is the arithmetic mean, and for intermediate levels of bias his perception is in-between the geometric and arithmetic means (there is no closed-form expression). In the long run, Eddie's perception of the growth of the asset will be dominated by the exponential term, but with growth rate αi instead of the correct growth rate i.

2.2 Theoretical Results

In this paper we explore Eddie's behavior in a lifecycle-consumption environment. Suppose Eddie has an instantaneous utility function over consumption $u(c_t)$ that is continuously differentiable and strictly concave. To guarantee that the first-order condition is sufficient to characterize behavior, utility satisfies the Inada conditions: $u'(0) = \infty$ and $\lim_{c_t \to \infty} u'(c_t) = 0$. Eddie is born in period 0 and dies in period $T > 1$, and he must choose his consumption in each period subject to an intertemporal budget constraint. In each period he receives a (possibly negative) cash flow y_t and he may save or borrow at that period's weakly positive interest rate $i_t \geq 0$ with the inequality strict for at least two periods.[5] Eddie discounts future utility exponentially by the discount factor $\delta \leq 1$.[6]

[5] If the interest rate does not differ from zero for at least two periods, then compounding plays no role in the agent's optimization and our model of EGB does not come into play. We focus on positive interest rates for reasons of exposition and economic relevance, but the model and propositions are straightforwardly extended to allow negative rates.

[6] Although we represent the agent's time preferences using the conventional exponential form, we do not model Eddie as applying his EGB to his discount function which is a component of his preferences. He accurately perceives his future preferences;

Thus Eddie's objective is

$$\max_{\vec{c}} \sum_{t=0}^{T} \delta^t u(c_t) \tag{2}$$

subject to the budget constraint written in terms of the period T-value of money,

$$\sum_{t=0}^{T} c_t \prod_{j=t}^{T}(1+i_j) = \sum_{t=0}^{T} y_t \prod_{j=t}^{T}(1+i_j). \tag{3}$$

However, since Eddie misperceives exponential growth, he believes that his budget constraint is instead:

$$\sum_{t=0}^{T} c_t \cdot p(t,\alpha) = \sum_{t=0}^{T} y_t \cdot p(t,\alpha) \tag{4}$$

$$\text{where} \quad p(t,\alpha) \equiv \prod_{j=t}^{T}(1+\alpha i_j) + \sum_{j=t}^{T}(1-\alpha) i_j.$$

Fundamentally, the interest rate gives the price of present consumption in terms of the price of future consumption. The term $p(t,\alpha)$ can be thought of as Eddie's perceived price of period-t consumption in terms of the period-T price. Thus, one unit of consumption in period t is perceived to cost $p(t,\alpha)$ units of consumption in period T.

Solving for Eddie's perceived Euler equation,

$$u'(\hat{c}_t) = \left(\frac{u'(c_0)p(s,\alpha)}{p(0,\alpha)\delta^t} \right) \tag{5}$$

where $\hat{c}_t$ is Eddie's planned consumption in period t at time 0. Define $\widehat{W}_{t,T}(\vec{y}) \equiv \sum_{s=0}^{T} p(s,\alpha) y_s$, as the period-$T$ equivalent of lifetime wealth from income stream $\vec{y}$, as (mis)perceived in period t. Substituting this into the budget constraint,

$$c_0 p(0,\alpha) + \left[\sum_{s=1}^{T} p(s,\alpha) u'^{-1} \left(\frac{p(s,\alpha)u'(c_0)}{p(0,\alpha)\delta^s} \right) \right] \equiv \widehat{W}_{0,T}(\vec{y}) \tag{6}$$

The intuition for the theoretical results of this paper lies in Equation (4). Eddie's bias leads to two perceptual errors, one for each side of the equation. On the left-hand side of the equation, Eddie misperceives the prices of consumption over time. If interest rates are always positive and $\alpha < 1$ then $p(t,\alpha)$ will be too low. Thus Eddie will perceive the price of future consumption to be relatively too high, and the price of present consumption relatively too low. We henceforth refer to this as the *price effect* of exponential-growth bias. From consumer theory, a change in prices (albeit a misperception in this case) leads to an income effect and a substitution effect. Since future prices are perceived to be higher than they actually are ($1 today is perceived to buy less in the future than it actually can), income is perceived to be lower. This force will generally decrease planned consumption in all periods. But since the *relative* prices of early periods are

he only misperceives the growth of assets.

perceived to be lower than they actually are, this will cause more planned consumption in early periods and less planned consumption in later periods. The net change in immediate consumption is therefore ambiguous. It depends on the elasticity of intertemporal substitution, which we turn to in Proposition 2.

On the right-hand side of Equation (4) Eddie misperceives the future value of his income. If $\alpha < 1$ then $p(t, \alpha)$ will be too low and so Eddie will underestimate the future value of his present income and overestimate the present value of future income. In other words, Eddie underestimates his budget when income is received early in life but he overestimates income when he receives it late in life. We henceforth refer to this as the *budget effect* of exponential-growth bias.

The budget effect leads to our first result. Because Eddie underestimates the value of early income relative to later income, Eddie will perceive an income stream that delays income as more valuable than an equally valuable income stream that expedites income. For example, suppose the interest rate is $i = 9\%$ and $T = 10$. The value of \$1 in $t = 0$ is worth \$2.37 in $t = 10$. A fully biased Eddie, however, will perceive the \$1 in $t = 0$ as worth \$1.90 in $t = 10$. Consider an income stream in which \$100 is received in $t = 0$ and nothing in all other periods, and a second stream in which \$237 is received in $t = 10$ and nothing in all other periods. An unbiased decision maker would be indifferent between these two income streams. However Eddie would gladly choose the latter because he overestimates its present value.

Lemma 1 (Income Deferment) *When $\alpha \in [0, 1)$, given income streams $\vec{y}$ and $\vec{z}$ and interest rates $\vec{\imath}$ such that:*

$$(i) \qquad \sum_{s=0}^{T} \left(\Pi_{j=s}^{T}(1+i_j)\right) y_s = \sum_{s=0}^{T} \left(\Pi_{j=s}^{T}(1+i_j)\right) z_s > 0$$

$$(ii) \qquad \sum_{s=0}^{t} \left(\Pi_{j=s}^{t}(1+i_j)\right) y_s > \sum_{s=0}^{t} \left(\Pi_{j=s}^{t}(1+i_j)\right) z_s \quad \forall t \in \{0, ..., T-2\},$$

then $\widehat{W}_{0,T}(\vec{y}) < \widehat{W}_{0,T}(\vec{z})$.

All proofs are relegated to the appendix.

The lemma states that (i) given two income streams of the same (actual) discounted value in which (ii) the value of the income received from $\vec{y}$ up to any point $t < T - 1$ exceeds the value of the income received thus far from $\vec{z}$, Eddie will perceive $\vec{z}$ as having higher value than $\vec{y}$. The second condition is similar to the notion of first-order stochastic dominance. Any delay of income will be perceived as more valuable from a $t = 0$ perspective, and hence any income distributions that have proportionately more mass of their present value later in life will be perceived as more valuable from a $t = 0$ perspective.

The direct implication is that Eddie will prefer jobs and career paths that have delayed income. This suggests that the preference for advanced education may partially be driven by error. Suppose Eddie has no consumptive value for college and he will only attend if it increases his lifetime earnings. Further suppose Eddie is the marginal student who makes exactly the same earnings from college and no-college. Since four years of college result in four years of zero income, high initial debt, but higher future wages, college would satisfy condition (ii) in Proposition 1. If Eddie had higher aptitude, college would generate higher lifetime income, and if he had lower aptitude no-college would generate higher lifetime income. In this case Eddie will prefer college because he misperceives the income stream with college as more valuable. Consequently,

if Eddie's aptitude were lower but not too low, Eddie would prefer college when his lifetime earnings would actually be higher without college.

The attentive reader may object that this pattern does not accord well with the college wage premium, which is large and positive (Avery and Turner, 2012). However, these are non-causal correlates. Second, even if this premium were causal, it is the difference between averages and not the return to a marginal student. The prediction of our model is that the marginal college prospective gets a negative premium whereas the standard model would predict a zero premium. We note that the bite of EGB does not come from low returns — Card (1999) reports average returns between 6 and 13% — but from under-appreciating the debt burden and opportunity cost of attendance. Avery and Turner (2012) find that 47.6% of college students in the U.S. from 2003-2004 who intend to get a bachelor's degree do not earn one within six years. Furthermore, of those who attain no degree, 51.3% have student loans with a mean loan amongst borrowers of $14,457. It remains an open question how many students would have been better off ex ante if they had not made the attempt, and there are likely other economic factors at work here including possibly overconfidence in aptitude and self-control problems. EGB may be one of several factors that result in inefficient human capital investment.

Since income received in later periods is overvalued and income received in the present is undervalued, delaying income will lead to the misperception that the value of lifetime income is greater than it actually is.

Proposition 1 (Deferred Income Increases Consumption) *When $\alpha \in [0, 1)$, delaying income from period $t < T - 1$ to $\tau > t$ in a manner that keeps (true) lifetime income unchanged will cause the agent to increase consumption in period 0. Shifting money between periods (T-1) and T will have no effect.*

The implication is that when Eddie receives compensation stated nominally and received in the future, he will overestimate his budget and overconsume in the present. The larger the delay, the larger Eddie's error.

The most pressing question regarding Eddie's behavior in this economic environment is whether he consumes too much or too little early in life. From Proposition 1, if income is sufficiently delayed, interest rates are sufficiently high, and α is sufficiently low, Eddie will overconsume in $t = 0$ relative to the optimum. However, the price effect may drive Eddie to underconsume if the income effect is too large. Below is a condition on observable behavior that indicates whether EGB causes Eddie to overconsume or underconsume.

Proposition 2 (Overconsumption) *If an agent faces a vector of weakly-positive cash flows $\vec{y}$ ($y_t \geq 0 \, \forall t$, $\exists s$ s.t. $y_s > 0$), then period-0 consumption is greater than that of an unbiased agent if $-\frac{u'(c)}{u''(c)c} > 1$, i.e. the elasticity of intertemporal substitution is greater than one.*

Proposition 2 states that a sufficient condition for Eddie to over-consume in the present relative to the optimal consumption path is that the elasticity of intertemporal substitution is greater than 1. Intuitively, the budget effect of EGB always leads to planned over-consumption, as EGB causes agents to systematically over-estimate the real value of their income. Indeed, the budget effect on its own could lead to over-consumption as the number of periods grows arbitrarily large. It is therefore sufficient that the price effect of EGB generate over-consumption. Because EGB raises the perceived price of future consumption in terms of current consumption, it will lead to over-consumption in the current period whenever the substitution effect of this mis-perception dominates the income effect – i.e., when the elasticity of intertemporal substitution exceeds 1.

Present-biased preferences, in which consumers overweight the present relative to the future in a dynamically inconsistent manner (Laibson, 1997; O'Donoghue and Rabin, 1999), are often invoked to explain apparent overconsumption. While this theory is remarkably successful in many domains, it may not tell the full story when applied to financial decisions. First, in order to explain take-up of loans at triple-digit interest rates, consumers would need to have calibrationally extreme discount factors (an annual $\beta \cdot \delta \leq 0.01$). Second, present-biased preferences are often invoked to explain a preference for smaller amounts of income sooner which leads to well-known estimation difficulties in incorporating risk preferences.[7] However, a present-biased person with access to credit (and no commitment devices) would choose to receive income to maximize her intertemporal budget constraint, just like an exponential discounter. Thus if the implied interest rate were higher than her outside borrowing rate, even a severely present-biased consumer should prefer the income later than sooner. She could then finance more current consumption by taking on greater debt.

Rather than impose a condition on the utility function, we can also rely on the budget effect as a sufficient condition for over-consumption.

Proposition 3 (Overconsumption From Future Wealth) *If an agent receives a lump-sum of wealth in period $T > 1$ and borrows to finance consumption in earlier periods, then period-0 consumption is decreasing in α (i.e. increasing in EGB).*

Thus even if Eddie is highly inelastic (i.e. his instantaneous utility function is very concave), he will still overconsume in period 0 if his wealth is fully received in the last period. Moreover, the magnitude of Eddie's error is both increasing in T and increasing in the interest rates i_t. Eddie's error can be made arbitrarily large and hence as a consequence Eddie may exhaust all his lifetime wealth before the final period. Suppose the agent has access to perfectly informed (and unbiased) credit markets, so that the true budget constraint must also be satisfied for actual (though not planned) consumption.

Corollary 1 (Bankruptcy) *If $\alpha < 1$ and then there exists $(\vec{\imath}, \vec{y})$ such that $c_t = 0$ for all $t > 0$.*

Corollary 1 states that any biased individual could be made to go bankrupt in the very first period, if income arrives late and interest rates are large and therefore the effect of EGB is sufficiently high.

The magnitude of Eddie's misperception is a function of timespan (e.g. Eddie makes no error regarding the present value of money from next period since compounding only occurs after spans greater than one), and so Eddie will generally behave in a dynamically inconsistent manner. Conceptually, this can be distinguished from other varieties of dynamic inconsistency that are preference-based (Strotz, 1956; Loewenstein, 1987; Laibson, 1997; O'Donoghue and Rabin, 1999; Loewenstein et al., 2003; Kőszegi and Rabin, 2006), since this dynamic inconsistency is generated instead by perceptual errors regarding compounding interest. Thus the welfare consequences are quite clear since the standard model's optimal consumption path is still Eddie's optimum. Of particular interest is the predictable pattern in which the dynamic inconsistency manifests.

Proposition 4 (Dynamic Inconsistency) *If the agent has a negative level of savings at the end of period $t < T - 1$,*

$$\sum_{s=0}^{t} (y_s - c_s) \Pi_{j=s}^{t} (1 + i_j) < 0 \tag{7}$$

[7] See, for example, Andreoni and Sprenger (2012) for a recent entry in this literature.

then the agent's period-t plan of consumption will exceed the period-(t+1) plan in all periods. Formally $\widehat{c}_{t,\tau} > \widehat{c}_{t+1,\tau}\ \forall \tau > t$. *If the inequality in expression (7) is reversed then planned consumption in* $t+1$ *will increase for all periods,* $\widehat{c}_{t,\tau} < \widehat{c}_{t+1,\tau}\ \forall \tau > t$, *and if the balance equals zero planned consumption in* $t+1$ *will be unaffected.*

The proposition can be explained intuitively. Since Eddie underestimates exponential growth, each period he will underestimate the change in his balance. If his balance is positive then he gets an unexpected windfall and if the balance is negative he gets an unexpected loss. Eddie's perception of the period-T value of income received at some intermediate future period τ does not change over time (he still makes an error on this but the error is consistent over time). So the only thing that changes for the perception of the budget is the growth of the current balance. These surprise changes to his current wealth cause him to shift his planned consumption vector in the same direction as the change.

For example, suppose Eddie receives his full income lump-sum in period T. Since Eddie must borrow to finance his current consumption his balance will always be negative. Hence his planned consumption for any future period τ will decrease each period. Therefore Eddie will overestimate consumption in all future periods and the magnitude of his error for consumption on a particular period τ will decrease as τ approaches.

If, on the other hand, Eddie receives his full income lump-sum in period $t = 0$, the exact opposite occurs. Eddie will maintain a positive balance and hence his planned consumption for any future period τ will increase every period.

This proposition implies that if Eddie receives most of his income late in life, his projected consumption plans will always exceed his actual consumption. This can be particularly dangerous for Eddie if he has the option to commit lower bounds on his future consumption. This may manifest in the housing market, where housing is a consumption commitment. A homeowner will find it costly and difficult to decrease her housing consumption in the next period since this generally requires selling the home (a long and arduous process). Dynamic inconsistency generated from EGB may lead to excessive commitments to housing, and this can be distinguished from a similar form of present-biased dynamic inconsistency (Laibson, 1997; O'Donoghue and Rabin, 1999). The latter will lead to excessive housing commitments only if the consumption begins in the present. In contrast, dynamic inconsistency can result in excessive commitments in which the consumption begins in the future. In fact, the later the consumption begins, the larger Eddie's error. So for example, Eddie may commit his funds to purchase a house beyond his means that will not be available for living until several years hence. A present-biased but otherwise financially sophisticated consumer would only overconsume if the house were available for immediate use.

Additionally, Eddie will underestimate the costs of debt. Eddie underestimates the speed at which a debt grows. Consequently he will underestimate the number of payments and the magnitude of the payments necessary to amortize a debt.

Proposition 5 (Debt Repayment) *When* $\alpha \in [0, 1)$, *the agent will underestimate the periodic payment required to fully amortize a given debt in a given amount of time. The magnitude of the error is monotonic in* α.

Corollary 2 *The agent underestimates the number of periodic payments of any given size needed to repay a given debt, and therefore the total amount spent on debt servicing.*

The proposition above formally states the commonplace intuition that EGB can lead to excessive leverage. Quite simply by underestimating compounding interest, Eddie will underestimate the costs of holding debt leading to the various puzzles discussed earlier.

3 Experiment 1: Lifecycle Consumption

The theoretical results in Section 2 predict that EGB can cause overconsumption and dynamic inconsistency in predictable ways. Experiment 1 is designed to test the predictions of the theory in a controlled simplified lifecycle-consumption environment. The parameters of the environment are set such that the main predictions are threefold: subjects will overconsume in early periods (Proposition 2), delaying the receipt of income increases overconsumption (Proposition 1 , and subjects will plan to consume more in future periods than they actually will (Proposition 4) if and only if their current assets are positive.

3.1 Design

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

The task was to choose a consumption vector for a given intertemporal utility function subject to a budget constraint. To keep subjects' interested, the problem was framed as a game to feed a simulated dog over several days. Subjects were provided with an income vector, an interest rate vector, and the instantaneous utility function for the "dog". Subjects could save and borrow freely at the given interest rates, although actual consumption was constrained to be feasible given the true budget constraint. Utility was described as "tail wags", and subjects earnings' were strictly increasing in the dog's total tail wags. Since the purpose of the experiment is to test the role of EGB in a lifecycle-consumption environment, and not to examine whether people are capable of maximizing arbitrary utility functions, we provided as many features as we could to help subjects understand the relationship between food and tail wags. We provided the explicit utility function, a graph of the utility function, and a calculator that took food as inputs and gave wags as output. This information was present on every screen in which subjects gave a response. Furthermore, subjects went through three training rounds with feedback about the optimal plan and how their choices fared. These training rounds each had a single day with a positive interest rate and the interest rate was zero on all other days. This familiarized subjects with the task and helped train them on the basics of utility maximization without unduly training them to detect exponential growth.

The instantaneous utility function of the dog was $u(x) = x^{\frac{1}{2}}$. We chose this as the utility function for two reasons. Given that this is a constant relative risk aversion utility function (CRRA) with an elasticity of intertemporal substitution of 2, Proposition 2 predicts overconsumption for any income vector. Furthermore, because CRRA utility functions are homothetic, the value of the income stream should have no effect on the proportion of the budget spent in each period for both an accurate agent nor a biased one. This allows for direct comparisons of behavior across different lifecycles without the confound of nonlinear wealth effects.[8]

The experiment had two arms. In the static arm, subjects chose their consumption plan vector all at once. However, actual consumption was implemented sequentially by the computer. Purchases were debited from

[8] Eliminating wealth effects for a non-homothetic utility function is not an option because a subject facing lifecycles with the same lifetime wealth but different income vectors will perceive different lifetime wealth due to EGB.

the subject's budget in chronological order. If consumption in a given period would exceed the overall budget, then the computer bounded consumption such that the budget was fully expended and set consumption in all remaining periods to zero. All left over money on the last period was automatically used for consumption in that period. In the dynamic arm, consumption was chosen sequentially. Subjects first chose the consumption in the first period. After submitting their answer they were told their current savings (or debt). They then chose consumption for the second period, and so on, receiving updated information about their asset position. As before, subjects could not exceed their budget, and any remaining money was used for consumption in the last period. Comparing the consumption across the two arms identifies the role of feedback about one's balance on consumption. We posit two potential effects of feedback. First, Proposition 4 predicts that the feedback will cause subjects in the dynamic arm to re-optimize their consumption within a lifecycle in a predictable way. Second, subjects may learn about their bias across lifecycles and improve their choices later in the experiment.

After the training sessions, subjects faced the main session. There were four lifecycles in the main session. Lifecycle 1 had $T = 5$, $\vec{y} =< 100, 0, 0, 0, 0, 0 >$ and $\vec{\imath} =< 0, 0, 100\%, 0, 0 >$. There was no compounding interest in this lifecycle and so EGB is irrelevant. Behavior on this problem establishes a baseline under incentives for the subject's performance. Lifecycle 2 had the same income vector as Lifecycle 1 and $\vec{\imath} =< 75\%, \ldots, 75\% >$. Lifecycle 3 had $\vec{\imath} =< 75\%, \ldots, 75\% >$ and $\vec{y} =< 0, 0, 0, 0, 0, 500 >$. Lifecycle 4 again uses the same interest vector $\vec{\imath} =< 75\%, \ldots, 75\% >$, but had an income vector $\vec{y} =< 0, 0, 100, 100, 0, 0 >$. The order in which the lifecycles were presented was random. Prop 1 predicts that overconsumption will be greatest for Lifecycle 3 and least for Lifecycle 2.

In the static arm, each lifecycle has 5 rounds. The first round is described above. In the second round, the consumption vector is chosen as if the consumer were making the decision in period 2 (period 1 has past and is fully sunk). Under Lifecycle 1 and 2, the consumer gets a lump sum in their first period, under Lifecycle 3, the lump sum is still received in the last period, and in Lifecycle 4 income is received in the second and third periods. The third round begins as if the consumer were allocating in the third period, etc. The homotheticity allows us to compare consumption vectors chosen within and across lifecycles: c_0 in the full 6-period round gives consumption in $t = 0$, c_0 in the 5-period round implies the fraction of wealth that would be consumed in $t = 1$, c_0 in the 4-period round implies the fraction of wealth that would be consumed in $t = 2$, and so on.[9] This simulated consumption vector predicts what the subject would do if her consumption were not committed but instead re-optimized in every period.

The dynamic arm has no rounds within a lifecycle. Instead, subjects merely give their consumption plan for each period iteratively and receive feedback on the current value of their savings or debt. Dynamic subjects' choices are thus analogous to the simulated static consumption without commitment, but with the addition of feedback. Subjects may become aware of their errors and use this information to improve their responses on the subsequent periods, or lifecycles. Thus we interpret any difference between the dynamic consumption and the simulated static consumption without commitment to be the effect of this feedback over the lifecycle.

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

[9] We can only use this for lifecycles that have all the income received at the start; otherwise the comparison would require us to know the subject's α to infer their perceived lifetime wealth. Among lifecycles with exponential growth, we can only simulate for Lifecycle 2.

based on how much additional utility their plan earned above the minimum achievable utility, as compared to how much additional utility would be achieved by the optimal plan. Letting u_a be the achieved utility, u_o the optimal utility, and u_m the minimum possible utility, a subject's additional payment was given by $\$35 - 35 \cdot [1 - (u_a - u_w)/(u_o - u_w)]^{\frac{1}{2}}$. This payment rule yields increasing returns as subjects approach the optimal utility. The mean incentive payment was \$18.46: dynamic subjects averaged \$17.68, while static users (who could be paid for rounds with fewer periods) averaged \$19.38.

3.2 Results

The overall performance of subjects is shown in Figure 1, where we plot the median consumption paths achieved by static and dynamic subjects by lifecycle. Panel (a) shows that subjects generally understood the task in the non-compounding Lifecycle 1, although there was a tendency to smooth somewhat too much. We therefore use this lifecycle to control for subjects' performance in the absence of exponential growth. Panel (b) shows that subjects over-consumed relative to the optimum in period-0, and that both dynamic and static subjects revised consumption up relative to this initial plan. Comparing panel (c) and (d) to panel (b) shows that delays in income result in dramatic increases in period-0 overconsumption. In the remainder of this section, we will address these features in more detail.

We begin by testing Proposition 2, that exponential-growth bias leads to overconsumption in early periods relative to the optimal spending path. Figure 2 plots the distribution of the natural logarithm of the ratio of subjects' period-0 consumption to the optimal period-0 consumption for that lifecycle. This is a simple normalization that allows a simple measure of overconusmption: subjects have over-consumed relative to the optimum when the variable is positive, and under-consumed when it is negative. For simplicity, we combine both static and dynamic subjects' answers, restricting to the $T = 5$ rounds for static users. In panel (a) of Figure 2, we plot the distribution of the logratio for Lifecycle 1, which has only a single nonzero interest rate and therefore EGB does not predict overconsumption. While the median value is higher than zero, the overall distribution indicates only a minor tendency to over-consume in period 0. As there is one non-zero interest rate, this suggests that subjects are slightly less price-responsive than is necessary to maximize utility. Any over-consumption in panel (a), however, is dwarfed by the over-consumption shown in panel (b), where we present the distribution of the logratio for Lifecycles 2–4, which have more than one period of non-zero interest. The bias towards over-consumption here is apparent, and both the mean and median are highly significantly positive.

We unpack the sources of over-consumption in Table 1. We use the same dependent variable, $\ln(c_0/c_0^*)$ as in Figure 2, and use OLS specifications in all columns.[10] Column 1 of Table 1 confirms that the presence of multiple non-zero interest rates significantly increases the initially mild over-consumption of subjects in period 0 of a round. In column 2, we see that this over-consumption is driven by the longer rounds. The variable "Extra Periods" records the number of periods by which the round is longer than 2. With just 2 periods (T=1) there is of course no possibility for EGB to distort decision-making, and we find that the "compounding" indicator is indeed not significantly different from zero. As the number of periods grows, there is small but significantly positive increase in the level of overconsumption. The main, effect, however, is the interaction of "Extra Periods" with the compounding dummy, which indicates that adding an additional

[10] LAD estimates for the effects on the median, as well as regressions of levels rather than logs, are available for all tables in this section, and do not substantively alter any results.

Figure 1: Feasible Consumption Paths

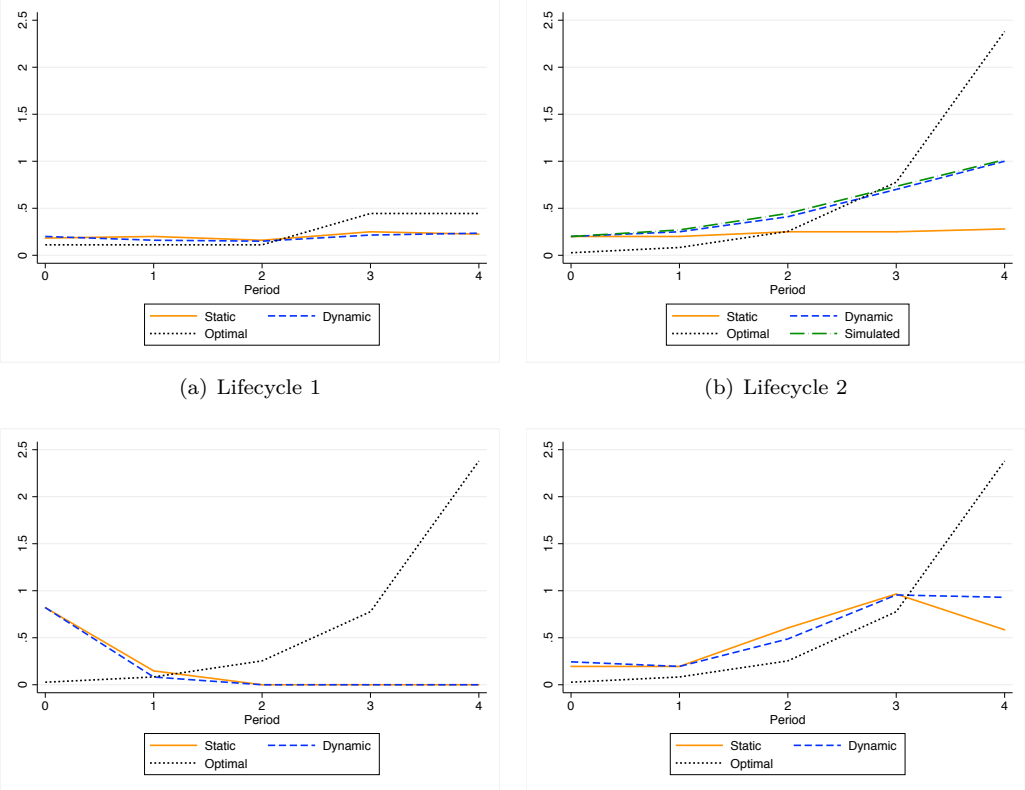

(a) Lifecycle 1

(b) Lifecycle 2

(c) Lifecycle 3

(d) Lifecycle 4

Notes: Median consumption paths achieved (i.e. subjects' plans restricted to feasible consumption), by lifecycle. Static refers to the feasible plan implemented by the $T = 5$ plan of static users. In panel (b), simulated refers to our simulation of static users' re-optimization based on their choices in all rounds.

Figure 2: Overconsumption

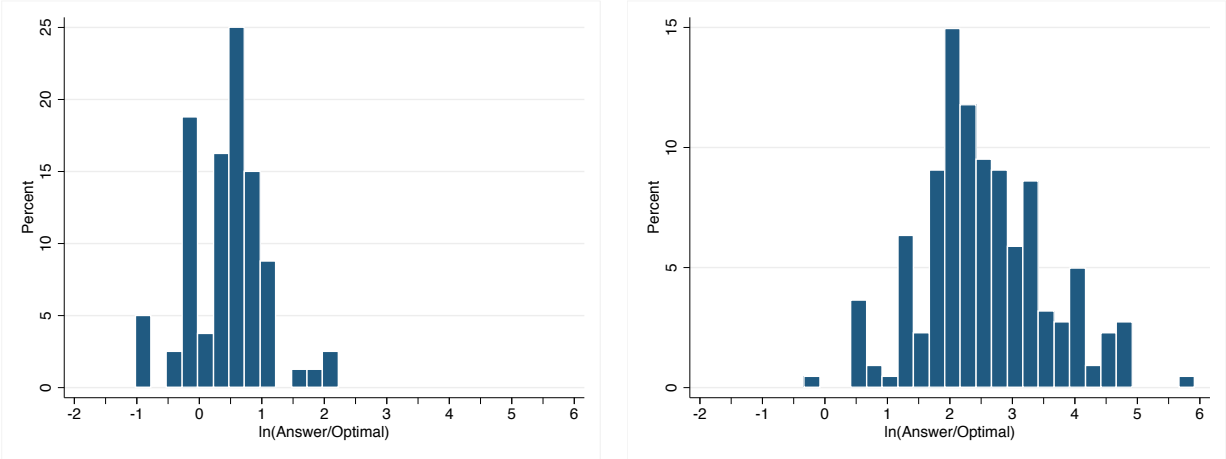

(a) Non-Compounding Interest

(b) Compounding Interest

Notes: Figures show the empirical distribution of $\ln(c_0/c_0^*)$: the natural logarithm of the ratio of subjects' choice of consumption for period 0 to the optimal consumption choice. Panel (a) restricts to Lifecycle 1, which did not feature compounding interest rates and is therefore not affected by EGB. Panel (b) includes Lifecycles 2–4, which do feature compounding interest rates.

decision-making period in a lifecycle with compounding interest rates of 75% leads to roughly a 50% increase in the ratio of period-0 consumption to its optimal value. This is exactly the prediction of the theory, as subjects' perceptions of the discount factor to be applied to period-t consumption are increasingly distorted for larger t. In column 3, we restrict attention to the longest rounds (T=5) and find that the effect of non-zero interest rates is indeed significantly greater than when the shorter rounds are included.

Restricting to the T=5 rounds also means that we have one observation per lifecycle for both static and dynamic subjects. In column 4, we investigate whether the two groups behave differently in their choices of c_0. Because dynamic subjects have not received any additional feedback until after they have chosen c_0, we would not predict any systematic difference. Finding one would suggest that subjects are approaching the tasks in fundamentally different ways, and would call into question the validity of our later comparisons between the groups. However, neither the indicator for Dynamic nor the interaction of that indicator with the dummy for multiple non-zero interest rates is significantly different from zero.

Finally, column 5 of Table 1 asks whether subjects learn to correct their bias over the course of the experiment. We restrict the sample to those lifecycles with multiple non-zero intersest rates, but include the shorter rounds for static users. Because the lifecycles were presented in a random order, as were the multiple rounds within a lifecycle for static users, we can include a variable "Question" which records the order in which a subject actually answered that particular round. We would not expect any learning amongst static subjects, since they received no feedback, but it is possible that dynamic users would learn about their EGB and revise their behavior accordingly. However, neither the Question variable nor any of its interaction terms are significantly different from zero, suggesting that subjects did not learn over the course of the experiment.

We next turn to our second hypothesis for this experiment. From Proposition 1, we have that delaying the timing of income will exacerbate overconsumption of EGB individuals. The simplest test of this prediction,

Table 1: Effect of Compounding on Overconsumption

	(All)	(All)	(T=5)	(T=5)	(All EGB)
Compounding	1.503***	-0.016	2.118***	2.238***	
	(0.059)	(0.090)	(0.069)	(0.101)	
Extra Periods		0.063***			0.538***
		(0.023)			(0.079)
Compounding X Extra Periods		0.516***			
		(0.030)			
Dynamic				0.109	0.086
				(0.133)	(0.410)
Compounding X Dynamic				-0.225	
				(0.137)	
Question					-0.009
					(0.011)
Question X Dynamic					-0.020
					(0.080)
Question X Extra Periods					0.004
					(0.005)
Constant	0.347***	0.247***	0.447***	0.388***	0.331*
	(0.055)	(0.061)	(0.066)	(0.099)	(0.189)
N	883	883	301	301	544
N_Clust	82	82	82	82	82

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

Notes: Dependent variable is $\ln(c_0/c_0^*)$; all columns are OLS regressions. Extra periods is the number periods by which T > 1. Question refers to the order in which the subject answered the question; subjects in the static arm answered 20 questions excluding the training, while subjects in the dynamic arm answered 4 excluding training. Standard errors clustered by subject.

Table 2: Effect of Delayed Income on Overconsumption

	(All)	(All)	(T=5)	(T=5)	(T=5)
All At End	1.293***	0.392**	1.316***	1.316***	1.280***
	(0.097)	(0.187)	(0.124)	(0.124)	(0.184)
At End X Extra Periods		0.248***			
		(0.058)			
Extra Periods		0.447***			
		(0.036)			
All At Middle				0.288**	
				(0.113)	
Dynamic					-0.212
					(0.166)
At End X Dynamic					0.065
					(0.250)
Constant	1.288***	0.198*	2.009***	2.009***	2.122***
	(0.071)	(0.114)	(0.083)	(0.083)	(0.124)
N	450	450	156	221	156
N_Clust	82	82	82	82	82

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

Notes: Dependent variable is $\ln(c_0/c_0^*)$; all columns are OLS regressions. The sample for columns 1-3 and 5 comprises those lifecycles where the entire endowment is received either in period 0 (All At Start) or period T (All At End). The sample for column 4 further includes Lifecycle 4 where all the income is received in the middle. Standard errors clustered by subject.

then, is to compare the degree of overconsumption in period 0 from Lifecycle 2, where all income is received in period 0, to that from Lifecycle 3, where all income is received in period T. Our hypothesis is then simply that overconsumption will be much greater when all income is received at the end of the lifecycle.

In Table 2, we regress our measure of overconsumption, $\ln(c_0/c_0^*)$, on an indicator for whether income is delayed, using all rounds from Lifecycles 2 and 3 and pooling dynamic and static users together. In the first column, the estimate of 1.293 on the delayed income dummy variable is highly significant and confirms the prediction of the theory. The second column interacts the delayed income dummy with the number of periods in the round. The coefficient of 0.447 on the Extra Periods variable reiterates the finding from Table 1 that overconsumption is increasing in the number of periods. The interaction of the delayed income dummy with the number of periods, however, is also positive at 0.248 and highly significant. It is straightforward to extend Proposition 1 to generate exactly this prediction.

In column 3 of Table 2, we restrict the sample to those rounds with T=5. This confirms that the effect is largest in the longest rounds, and ensures that both static and dynamic subjects have exactly two observations each. Column 4 includes Lifecycle 4 where all the income is received in the middle periods. That the coefficient on All At Middle should be positive and less than the coefficient on All At End is a stronger test of Proposition 1, and is confirmed (we reject equality of the coefficients at $p < 0.01$). The next column confirms that there is no systematic difference between subjects in the two treatments: neither the dynamic dummy nor its interaction with the delayed income dummy enters significantly.

Finally, we address our prediction of dynamic inconsistency from Proposition 4. Once again it is easiest to test this hypothesis by restricting attention to Lifecycles 2 and 3. Because Lifecycle 2 has all income received in period t=0, subjects will always carry a (weakly) positive balance. Proposition 3 suggests that dynamic users 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 Lifecycle 3 must borrow against their income from the final period to finance early consumption. Here, dynamic users will be surprised by how quickly their debts grow and should revise their consumption downwards in later periods. We can test these predictions by using static subjects' plans from the T=5 rounds as counterfactuals for dynamic subjects' initial plans. That is, if a dynamic subject were asked to enumerate their full consumption plan at $t = 0$ without any feedback, they would be in exactly the static treatment. We can therefore use static subjects' answers from $T = 5$ rounds to stand in for dynamic subjects' initial plans. We then compare the dynamic subjects' *actual* consumption to these initial plans on a period-by-period basis to explore the dynamic inconsistency induced by re-evaluation of the budget constraint.

We use OLS regressions to perform these comparisons in Table 3. The dependent variable is now the natural logarithm of period-t consumption normalized by the period-0 value of the income stream, to account for the higher attainable consumption in Lifecycle 2. The first column pools Lifecycles 2 and 3, and finds no differences between the static and dynamic users in the aggregate. The coefficient of 0.168 on the period variable indicates that subjects correctly plan for an upward-sloping consumption profile, although the trajectory is flatter than would be optimal. Columns 2 and 3, which separate out Lifecycles 2 and 3 respectively, confirm the predictions of Proposition 3. The Dynamic indicator is again insignificant in column 2, as dynamic and static subjects do not differ in their plans for c_0. In later periods, however, dynamic users are surprised to learn how much their savings have grown. The coefficient on the Dynamic X Period interaction is large and positive, indicating that dynamic users' actual consumption choices in later periods are revised substantially upwards from a static users' plan at t=0. The opposite picture emerges in column 3, where Lifecycle 3 dynamic users are surprised by how little they have left to spend. The Dynamic X Period interaction is again large and significant, but now negative – indicating that dynamic users' actual consumption choices are revised substantially downwards from static users' plans at t=0.

We note, however, that this revision by dynamic subjects is no more than we would have expected from static subjects' choices in shorter rounds. While we can only simulate how static subjects would have re-optimized in Lifecycle 2, the simulated static consumption without commitment is indistinguishable there from the dynamic consumption. This was seen clearly in Figure 1, panel (b).[11] This implies that feedback on one's balance has no impact on consumption.

We also note that, while subjects in the dynamic arm of the experiment did revise their consumption plans relative to static subjects' $T = 5$ plan, the overall welfare differences between static and dynamic subjects – the proportion of the utility from the optimal plan they actually achieved – were quite small. Across all the $T = 5$ with multiple non-zero interest rates (Lifecycles 2–4), subjects in the static arm earned an average of 64.2% of the maximum utility, while subjects in the dynamic arm earned an average of 63.4%. This welfare result is largely due to our choice of utility function: with CRRA function with an EIS of 2, simply delaying all consumption until the final period would earn 66% of the maximum utility. To earn

[11] Using the simulated static consumption without commitment and dynamic consumption data, we can regress consumption on a set of time dummies interacted with a "dynamic" dummy. We cannot reject that the time dummies interacted with dynamic are jointly different from zero ($F(5, 81) = 0.41$, p=0.840).

Table 3: Dynamic vs. Static Plans

	(All)	(Inc At Start)	(Inc At End)
Dynamic	-0.224	-0.204	-0.135
	(0.167)	(0.156)	(0.250)
Period	0.168***	0.307***	0.046
	(0.029)	(0.048)	(0.033)
Dynamic X Period	0.035	0.254***	-0.245***
	(0.057)	(0.076)	(0.078)
Constant	-1.109***	-2.135***	-0.128
	(0.128)	(0.117)	(0.180)
N	763	398	365
N_Clust	82	81	80

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

Notes: Dependent variable is $\ln(c_t/W_0)$, i.e. consumption normalized by actual starting wealth. The sample comprises the compounding T=5 lifecycles with all income received either in period 0 or in period 5. One subject planned c0-c4 = 0 when income at start; two (different) subjects planned c0-c4 = 0 when income at end. Standard errors clustered by subject.

a reasonable payment, subjects simply had to avoid an early bankruptcy. 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 Lifecycles 2–4 led them to bankruptcy (i.e. $c_5 = 0$). This rose to 73.7% and 68.2% in Lifecycle 3, of which 63.2% and 59.1% respectively occurred by the end of the second period. While dynamic subjects realized their debts were growing faster than they anticipated, in effect it was too late for them to get back on track.

4 Experiment 2: EGB in the U.S. Population

In our second experiment, we examine the external validity of EGB by measuring the distribution of exponential-growth bias in the population, and examining the correlation of the bias with individuals' financial outcomes. We also test whether the bias is robust to a simple graphical intervention. We shift from a lifecycle consumption paradigm to a direct belief elicitation in order to directly estimate the EGB parameter α with the fewest potential confounds.

4.1 Design/Method

4.1.1 Design

Subjects faced a series of questions describing two assets, at least one of which involved a form of exponential growth, and were asked to indicate the initial value for one asset which would equate the assets' final values after an indicated length of time. For example, the first question is a choice between "Asset A that has an initial value of $100 and grows at an interest rate of 10% each period; Asset B has an initial value of $X and does not grow." Subjects were asked for the value of X which would make the two assets equal value after

20 periods. Appendix Table B.1 displays the full list of questions presented to subjects.

Questions 1–10 are our primary focus in the analysis, and the order of presentation was randomized first at the domain level and then within domain at the question level. The exponential domain comprised four questions similar to the example above. The fluctuating-interest domain comprised three questions involving fluctuating interest rates, of the form: "Asset A has an initial value of $\$P$ and grows at an interest rate of $i\%$ in odd periods (starting with the first) and at $j\%$ in even periods; Asset B has an initial value of $\$X$ and does not grow; What value of X will cause the two assets to be of equal value after T periods?". The catch-up savings domain comprised three questions which varied the maturity of the assets, of the form: "Asset A has an initial value of $\$P$ and grows at an interest rate of $i\%$ each period; Asset B has an initial value of $\$X$ and grows at an interest rate of $i\%$ each period; What value of X will cause the two assets to be of equal value after Asset A grows for T periods and Asset B grows for S periods?"

In order to incentivize subjects to answer the questions correctly, they were informed that each answer would receive a payment based on its accuracy. The payment rule was piecewise-linear in the percentage error: each answer within 10% of the truth would receive \$0.80; each answer within 25% would receive \$0.60; each answer within 50% would receive \$0.20; and each answer less than 50% of or more than 150% than the truth would not receive a payment. In addition to the incentive payments, subjects received \$5.00 for completing the entire experiment. Subjects had a week to do the experiment at their leisure. All payments were made through Knowledge Networks' internal payment mechanism, which subjects were already experienced with, and were usually paid within a week of completion.

The experimental instrument intentionally did not mention the use of tools for answering the questions. Subjects could potentially use whatever tools that they had access to: from nothing, to advice from friends or financial calculators. This design neither discouraged subjects' natural tendency to use tools nor did it prime subjects to use them. Although the incentives are not nearly as large as they would be for making actual financial decisions in the marketplace, a subject that exploited tools could earn substantially more for their time spent.

After completing the primary experiment, subjects were randomly assigned into a control (N=384) and a treated (N=185) group to test the effect of a simple information presentation or "nudge" on a second set of questions. The intervention, shown in Appendix Figure B.3, shows the growth of \$100 at one or more relevant interest rates, and allowed the user to specify the time horizon plotted. Treated subjects were shown this intervention beneath each question, while control subjects were not. Subjects answered an additional 16 questions, 10 from the original three diagnostic domains and 3 each from the periodic savings and portfolio domains. These latter two domains are qualitatively more complicated to solve.[12] The periodic savings domain has questions of the form, "Asset A: At the beginning of each period, receives a \$10 contribution. These contributions earn 7% interest every period, and Asset A includes both the contributions and the interest earned at the end. Asset B returns a fixed amount of \$X at the end. What value of X will cause the two assets to be of equal value after 40 periods?" The portfolio domain asks question of the form, "Asset A has an initial value of \$100, and grows at an interest rate of 10% each period. Asset B consists of two pieces. One piece has an initial value of \$50, and grows at an interest rate of 5% each period. The other piece has an initial value of \$X, and grows at an interest rate of 15% each period. What value of X will cause the two

[12] We do note formalize a metric of "difficulty", but note that, for example, Catch-up Savings questions require a subject to sum T separate Exponential questions.

assets to be of equal value after 40 periods?" We exclude these additional questions from all analysis until Sections 4.2.3 and 4.2.4, focusing just on the 10 pre-intervention questions where all subjects faced identical circumstances.

4.1.2 Sample

Experiment 2 consists of an incentivized online experiment conducted on a nationally-representative sample. Participants were recruited through Knowledge Networks, which maintains a non-volunteer panel of U.S. households. Participant households are selected randomly by Knowledge Networks based on their address, and are provided with a laptop and free internet access if necessary.[13] A weighted sample of subjects from the Knowledge Networks panel were invited to participate in our study. Subjects logged into their accounts through the Knowledge Networks portal, and were automatically transferred to an external website where our study was hosted.

Table 4 shows summary statistics for our sample. Column 1 shows the characteristics of all 990 KN panelists who were invited to participate, while Column 2 comprises the 569 subjects who answered or skipped all the questions in our study. Men were significantly more likely to complete the study (63% vs. 52% for women, $p<0.01$), so that 46% of the final sample were women. The average age of those opting to complete the study was also somewhat lower than those opting out, although this result was largely driven by a very high completion rate amongst 18–21 (i.e. college-aged) panelists. Race and education characteristics did not predict study completion, with 28% of subjects having only a high school degree, 29% some college or an associate's degree, and 37% having a bachelor's degree or more.

For some of the analysis, we merge our experimental dataset with an external dataset containing subjects' financial characteristics. Participants in the Knowledge Networks panel are regularly asked about their income and assets, and we will use this information to investigate the effect of exponential-growth bias on savings. These data are only available for a subset of subjects (the others either being ineligible or refusing to answer), and we present them in the third column of Table 4. Unsurprisingly, this subsample tends to be older and better educated than those for whom financial data are unavailable: the mean age is 50.02, and 53% have at least a bachelor's degree. The higher education attainment rate is also reflected in the high average household income of $90,257 among this group. This group also had significant investible assets — a mean of $241,055 — suggesting that they may overstate the degree of financial sophistication relative to a poorer, less well-educated population.

4.2 Results

We begin this section by showing evidence that subjects in Experiment 2 were systematically biased in their answers, and in the direction predicted by exponential-growth bias. We then estimate an individual-level bias parameter based on the model presented in Section 2, and investigate its distribution and correlation with household finances.

[13] More details on the KnowledgePanel sampling methodology are available at http://www.knowledgenetworks.com/knpanel/KNPanel-Design-Summary.html

Table 4: Summary Statistics — Experiment 2

	Initial Sample	Study Completers	Complete Data
Demographics			
Female	0.51	0.46	0.46
Age	47.38	44.73	50.02
Education			
Completed High School	0.27	0.28	0.19
Some College	0.30	0.29	0.26
Bachelor's Degree +	0.36	0.37	0.53
Race/Ethnicity			
Black, Non-Hispanic	0.07	0.07	0.06
Other, Non-Hispanic	0.05	0.05	0.07
Hispanic	0.13	0.13	0.10
2+ Races, Non-Hispanic	0.03	0.03	0.03
Financial Products			
Has credit card		0.83	0.94
Has used payday loan		0.07	0.05
Has car loan		0.61	0.74
Has mortgage		0.60	0.79
Has second mortgage		0.16	0.23
Assets/Income			
Non-Housing Assets			241055.74
Income			90257.60
Observations	990	569	296

4.2.1 Bias

For each subject i and question j, we first calculate the natural logarithm of the ratio of the given answer to the correct answer. Let a subject's responses on question $j \in \{1, ..., J\} = \mathcal{J}$ be denoted by r_{ij}, and the correct response be given by c_j. We calculate $e_{ij} = \ln(r_{ij}/c_j)$. We choose this measure to provide a consistent measure across questions that may have answers that differ by several orders of magnitude. Were all subjects to answer exactly correctly, this statistic would be exactly zero. As subjects were not prohibited from using calculators, spreadsheets, and online tools to help them answer the questions, such an outcome would not strain credulity. More likely, subjects would not be exactly correct, but incorporate some noise into their responses – particularly if they were making an informal estimate. If subjects are unbiased, then the mistakes should be evenly distributed about zero. Specifically, if errors on an absolute or percentage basis are symmetrically distributed around zero, then the median of $\ln(r_{ij}/c_j)$ should also be zero. More in the spirit of exponential-growth bias, if subjects' answers are a power of the correct answer, $r_{ij} = c_j^{1+\varepsilon}$ where $\varepsilon \sim \mathcal{N}(0, \sigma^s)$, then the log-ratio should be normally distributed about zero.

Instead, we find a systematic bias in the error, the sign of which depends on whether exponential-growth bias predicts that subjects should over- or under-predict on that question. Figure 3 plots the distribution of log errors at the question–subject level. As expected, the modal error is zero – the likeliest interpretation is that a large mass of subjects are able to use calculators to get the answer exactly correct. This holds both for questions where under-estimation is predicted as well as those predicting over-estimation. Apart from the zeros, however, the differences are stark. Where under-estimation is predicted, the distribution is shifted sharply to the left. Both the median (-0.349) and mean (-0.554) are significantly negative ($p<0.01$). The pattern is reversed where exponential-growth bias predicts over-estimation: the distribution is shifted sharply to the right and the mean answer (0.209) is now significantly positive ($p<0.01$), although the median answer in this case is zero.

Turning from question-level mistakes to subject-level mistakes, the pattern becomes even more clear. We calculate subject-level averages of the above log-ratio:

$$\bar{e}_i = \frac{1}{10} \sum_{j=1}^{10} \ln(answer_{ij}/correct_j)$$

If subjects are making unbiased errors, then the above results for the means and medians hold. Moreover, if r_{ij}/c_j is i.i.d. lognormal, then the averaging should cause the distributions to collapse towards zero. Figure 8 plots the distribution of the subject-level averages. Rather than converging towards zero, the subject averages even more closely follow the pattern predicted by exponential-growth bias. Averaged across questions where under-estimation is predicted, the modal bin is now negative rather than zero, and both the median (-0.507) and mean (-0.602) are more negative the before. When averaged across questions where over-estimation is predicted, mass also shifts away from zero — and while the mean (0.400) is again positive, so too now is the median (0.405).

4.2.2 Estimating Alpha

Let $\vec{a}(\alpha) : \mathbb{R} \to \mathbb{R}_+^J$ be a function that generates the J answers consistent with any level of α. Thus $\vec{a}(1)$ would be a vector containing the J correct answers. For every subject, we calculate the value of α which

Figure 3: Question-Level Mistakes

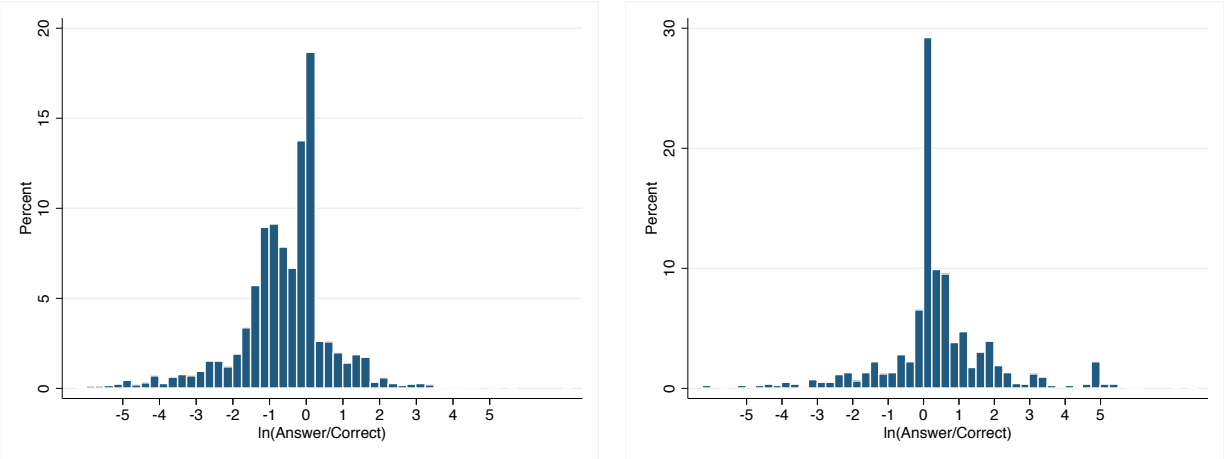

(a) Under-Estimation (b) Over-Estimation

Notes: Underestimation based on the questions for which EGB predicts a downward-biased answer; overestimation from those where an upward bias is predicted by the theory. Panels show the distribution of errors in predicted asset growth, and should be symmetric about zero if subjects' errors on a percentage basis are symmetric about zero. The means of both distributions are significantly different from zero (significant at $p < 0.01$).

Figure 4: Subject-Level Mistakes

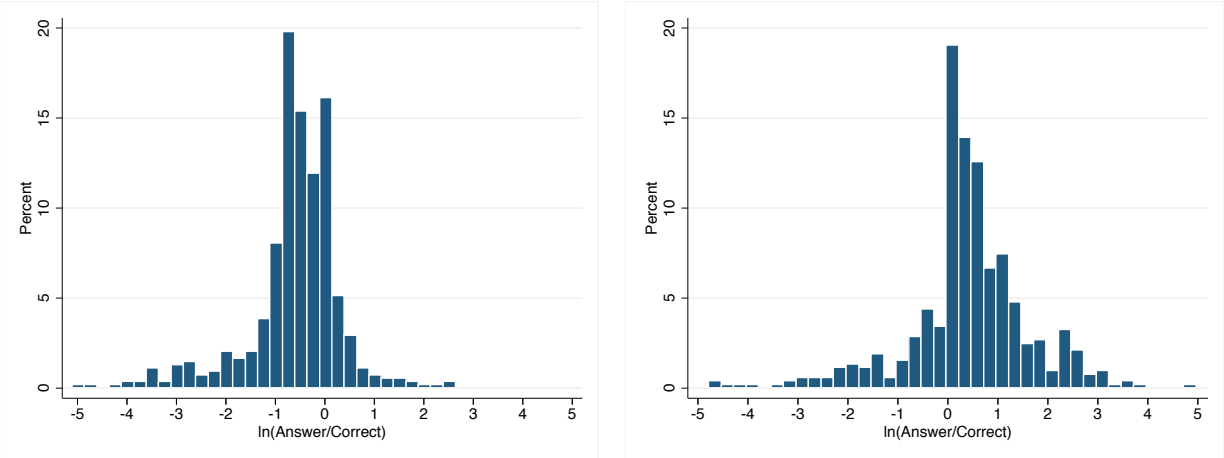

(a) Under-Estimation (b) Over-Estimation

Notes: Underestimation based on the questions for which EGB predicts a downward-biased answer; overestimation from those where an upward bias is predicted by the theory. Panels compute the mean of $\ln(answer/correct)$ at the subject level, and should converge to a point mass at zero in the absence of systematic bias. The means of both distributions are significantly different from zero (significant at $p < 0.01$).

minimizes the mean squared error of the model against their actual answers, with each question normalized by the correct answer. This normalization avoids having those questions which contain large values for the solution arbitrarily dominate the estimation procedure. That is, we estimate:

$$\widehat{\alpha}_i = \arg\min_{\alpha} \frac{1}{|J|} \sum_{j \in \mathcal{J}} \left(\frac{r_{ij} - a_j(\alpha)}{a_j(1)} \right)^2 \tag{8}$$

The estimator described by (8) is not constrained to values lying within the unit interval. Values of α greater than one are simply interpreted as an individual who over-estimates the rapidity of exponential growth. Values less than zero are less intuitive, but represent individuals who estimate growth to be slower than linear. We perform an unconstrained numerical optimization to estimate an $\widehat{\alpha}_i$ for each subject.

Figure 5 plots the cumulative distribution of our estimates of α, using our full sample of completers. We characterize 85% of the population with an α between $[0, 1]$. The median α is 0.53, and the mean is 0.60. Moreover, we have a large number of people who are completely, or nearly completely, fully biased: 33% of subjects (184/561) have an alpha of "exactly" zero (i.e. within $[-0.001, 0.001]$). In contrast, only 4% (23/561) are completely correct (even using a more generous definition of [0.99,1.01]). Based on our bootstrapping procedure, we can reject that the 80th percentile has $\alpha = 1$ at 95% confidence. Similarly, we cannot reject that the 37th percentile has $\alpha = 0$.

Figure 5: Population Distribution of Alpha

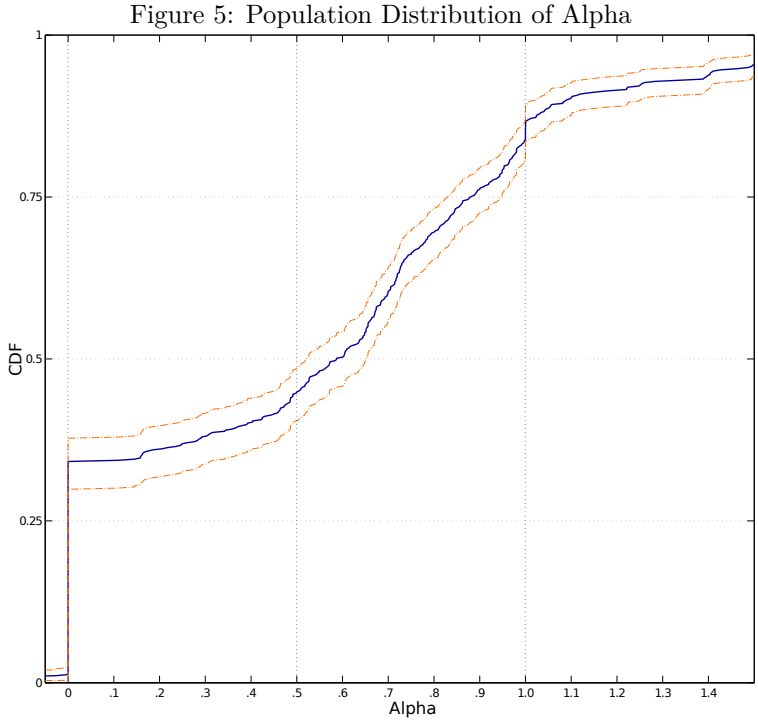

Notes: Cumulative distribution of alpha, based on full (unweighted) estimation sample. Dashed lines indicate 95% confidence interval for percentiles of the distribution, based on 5000 bootstrap replications. Confidence intervals should be read horizontally, e.g. the median alpha is estimated to lie between 0.520 and 0.651 while the 25th percentile is estimated to be exactly zero.

Our measure of α is uncorrelated with education, age, race, and sex. Unsurprisingly, the 6% of subjects

who reported an online calculator perform substantially better than the rest of the population. The mean α in this group is 0.84 (0.32 higher than those who do not use financial calculators) and with a median α of 0.96 (relative to a median amongst those who do not use financial calculators of 0.56). There may be both a causal and self-selection effect in this population, which our research design does not distinguish.

We can also relate α to numeracy. In our post-experiment survey we asked subjects what math was required to provide the correct response: simple arithmetic (addition, subtraction, multiplication, division), advanced arithmetic (exponentiation, logarithmic operations), pre-calculus (trigonometric operations), calculus and other advanced math. The correct answer is the first two, and only 23% of subjects answered this correctly. However answering this correctly had no statistically significant association with α. This is is also consistent with our finding that education was not statistically associated with α either. Question 3 also gives us a test of basic interest rate numeracy. The question asks for the value of an asset after it grows for only one period ($P_0 = 100$, $i = 4\%$), and 71% of subjects answered correctly. We note that our model, unlike some alternatives, predicts a correct answer to this question regardless of the degree of bias. Moreover, mistakes on this question are uncorrelated with our measure of α, which provides reassurance that we are estimating a systematic bias rather than noise. Dropping subjects who fail to answer this question correctly does not substantively change any of the remaining analysis.

With estimates of individuals' α we can address the question of central importance: the relationship between EGB and long-run financial outcomes. Proposition 1 states that biased agents will systematically over-consume in early periods relative to the optimal consumption path when the elasticity of intertemporal substitution is greater than unity. This amounts to a strong empirical claim: biased agents will have lower levels of savings than otherwise-identical unbiased agents. This posits exponential-growth bias as a partial explanation for the high degree of variation in retirement savings within income and education categories found by Bernheim et al. (2001).

We are able to match financial records from an external survey to 296 of our 569 experimental subjects. As the unmatched cases correspond to refusals or ineligible cases (often college students), this leaves us with a slightly older and better-educated subsample. We then estimate the relationship between our estimate of α and investible assets. We perform a linear regression in ln(assets) of the form:

$$\ln(assets_i) = \theta_1 + \theta_2 \cdot \alpha_i + \theta_3 \cdot \ln(income_i) + \theta_4 \cdot age_i + \theta_5 \cdot female_i + \theta_6 \cdot hhsize_i + \theta_7 educ_i + \varepsilon_i$$

The results of this regression are shown in Table 5. Column 1 reports the main specification. Unsurprisingly, older people have accumulated more assets, and a 1% increase in income tends to be associated with slightly more than 1% higher level of savings. Our coefficient of interest, α, enters significantly positively at 0.485. The estimated magnitude of the effect is large: all else equal, a fully-biased agent is expected to have accumulated just 62% of the assets of an un-biased agent.

The other columns perform some robustness checks on this result. Omitting income as a regressor in column (2) actually lowers the coefficient on α, as does omitting both income and education in column (3), suggesting that this is not just a case of higher ability agents both earning more and making better investment decisions. Indeed, Lemma 1 implies that exponential-growth bias may actually lead agents to over-invest in human capital. Column (4) breaks education apart into categories rather than a continuous variable, and does not significantly change any estimates. Column (5) omits α as a regressor, and does not change our estimates of the other coefficients.

Table 5: Alpha and Savings

	(1)	(2)	(3)	(4)	(5)
Alpha	0.485***	0.390**	0.361**	0.466***	
	(0.166)	(0.171)	(0.176)	(0.170)	
Ln(Income)	1.173***			1.184***	1.158***
	(0.153)			(0.157)	(0.161)
Education	0.076	0.219***			
	(0.056)	(0.060)			
Age	0.057***	0.063***	0.062***	0.057***	0.057***
	(0.006)	(0.007)	(0.007)	(0.006)	(0.006)
Female	-0.046	-0.133	-0.175	-0.045	-0.067
	(0.190)	(0.218)	(0.215)	(0.188)	(0.193)
Household Size	-0.131	-0.012	-0.043	-0.138	-0.125
	(0.088)	(0.080)	(0.075)	(0.091)	(0.093)
High School				-0.095	-0.325
				(1.050)	(1.034)
Some College				0.069	-0.183
				(1.059)	(1.037)
College Degree +				0.245	0.050
				(1.055)	(1.042)
Constant	-5.629***	5.412***	8.042***	-5.033***	-4.308**
	(1.624)	(0.850)	(0.462)	(1.763)	(1.726)
N_Clust	296	296	296	296	296

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

Dependent variable is ln(investible assets)

4.2.3 Domain-Specific Predictions and Fingerprints

The domains were carefully designed not only for aggregate estimation of α, but also to test subtle and specific predictions of the model. Each domain demonstrates how EGB manifests in a simple and common financial problem. Questions 4, 13, and 14 of the "exponential" domain give the subjects negative interest rates. This allows us to test our model against the WS model described in the introduction. The WS model predicts that exponential growth is attenuated, and thus biased individuals will *overestimate* the value of a depreciating asset. In contrast, our model predicts that subjects will *underestimate* the value of the asset due to neglecting compounding (i.e. a shrinking principal). Amongst the three questions, 19% of answers are correct and therefore consistent with any model that nests an unbiased case. Of the remaining answers that are equal or below the principal, twice as many under-estimate the asset as over-estimate it. Moreover, 21% of responses on these three questions are exactly what our model predicts for the fully biased type. Thus a fully biased agent leaves a fingerprint that allows for easy identification. Across the whole domain, 33% of responses can be fingerprinted as full-bias. The comparison favors our model.

Perhaps an even simpler test of our model against alternatives comes from Question 3, which asks subjects about the value of a \$100 asset which grows for exactly one period at 4% interest. Given the lack of compounding, our model predicts that subjects should not make any error. Previous models of EGB would predict an error even on this question, as would a broader lack of numeracy in the population. In our sample, 75.3% of subjects were able to correctly identify the correct value of \$104, and a further 3.96% answered with the interest-only value of \$4. While this is not the 100% predicted by the theory, it is sufficiently far above the 17.0% rate of correct answers on other questions in this domain that we take it as evidence that compounding is genuinely at the heart of people's errors.

The fluctuating-interest domain and the portfolio domain both demonstrate people's tendency to take the arithmetic mean when combining multiple interest rates. A fully biased agent will use the arithmetic mean to determine the principal on both of these problems. The arithmetic mean will overweight the impact of the higher interest rate on the fluctuating-interest domain, since the geometric mean is more conservative. In contrast, the arithmetic mean will underweight the impact of the higher interest rate in the portfolio domain, since in the long run the asset with the highest interest rate will dominate. We find exactly this pattern amongst our subjects. For example, in question 7 of the fluctuating-interest domain, the mean growth is exactly zero, but the arithmetic mean of the interest rates is positive. On this question, 58% of subjects believed the asset would increase while only 22% believed it would decrease (and 20% got it exactly right). On question 16 of the same domain, the mean growth was negative but the arithmetic mean is exactly zero: 32% of subjects fingerprinted themselves as full-bias by responding with a zero change in the asset. Over the domain as a whole 17% of responses left a full-bias fingerprint. On the portfolio domain, in question 24, the question with the simplest computation for a full-biased agent, only 10% overestimated the impact of the high interest rate (these subjects chose a principal lower than the correct value), and 74% underestimated the impact of the high interest rate (these subjects chose a principal above than the correct value) as predicted; 15% left a full-bias fingerprint. In the domain as a whole, 9% of responses left a full-bias fingerprint.

These two domains show how Eddie will incorrectly combine interest rates and as a consequence exhibit *as-if risk preferences*. Keep in mind that there is no uncertainty in these problems. But an economist who thinks that uncertainty generates the fluctuating-interest and believes Eddie to be unbiased, would infer that Eddie is risk-seeking. In contrast an economist who thinks that Eddie's portfolio faces risk would observe

that Eddie is heavily invested in low-return assets. If the economist assumed a risk-return tradeoff, she would infer that Eddie is quite risk averse.[14]

The catch-up savings and periodic savings domains are significantly more complicated to think about, even if one does take a full-biased approach. As a result 0% of subjects leave a full-bias fingerprint. Nonetheless, the directional prediction holds true: subjects vastly underestimate the value of an asset that has grown over a long duration. Thus subjects underestimate how much they need to contribute to catch-up (70% underestimate, 21% overestimate) and underestimate the value of periodic savings (86% underestimate, 11% overestimate). The lesson from these more complex domains is that the model does well in predicting direction but loses precision. The complexity of these questions provides ample opportunity for other mathematical errors. Indeed, 33% of responses in the periodic savings domain are below the sum of the contributions!

Fundamentally, the model presumes that Eddie broadly brackets each problem.[15] For example, Question 8 can be simplified to solving for a principal of \$100 growing for five periods at 13% interest. More generally, since Eddie gets one round of interest exactly right, if he were to break down a problem into a sequence of iterated one-period problems, he would make no mistake. Question 10 in the catch-up savings domain was designed to directly address this issue; it asks what principal is needed for a one period delay in savings (the answer is $\frac{1}{1+i}P_0$). Subjects do not seem to simplify this into a one period problem: 19% got the answer correct which is about their accuracy on other problems, and 67% respond with a principal in the predicted direction.

4.2.4 Stability and De-Biasing

Finally, we use the fact that subjects repeated the experiment on a second set of questions in order to address the stability of our parameter estimates, and to examine whether the bias is robust to a simple graphical intervention.

To address the stability of our α parameter estimates, we re-estimate equation (8) using only subjects' responses to the second set of 10 questions and then compare our two estimates of α within subjects. We are most interested in whether subjects identified as the "extreme" types – that is, with $\alpha \in \{0, 1\}$ are consistent. This does appear to be the case. Of 126 control subjects identified as having $\alpha = 0$ on the first set of questions, 78 (61.9%) yielded an estimate of $\alpha = 0$ on the second set.[16] In a linear probability model (available in the online appendix), we find that having $\alpha = 0$ in the first set of questions raises the probability of having $\alpha = 0$ in the second set by 22.28 percentage points (s.e. 4.41), while having initially had $\alpha = 1$ lowers the probability by 41.61 percentage points (s.e. 9.21). We conclude that there is substantial persistence in subjects' measured exponential-growth bias.

We also conclude that the bias is robust to the provision of information. This may be surprising, as the intervention made calculating the correct answer in the Exponential Domain all but trivial, but is not unreasonable – subjects were already free to use whatever tools they wanted to help them, including ones far more sophisticated than a simple graph. We find that subjects in the control and treated groups were

[14] Eddie's as-if risk preferences and behavior under uncertainty is the subject of a sister paper in progress.

[15] For more on choice bracketing see Read et al. (1999).

[16] We focus on control subjects to separate out any effect of the graphical intervention. It is not surprising, however, that a comparable 34 of 58 (58.6%) treated subjects identified as having zero in the first set of questions were also identified as having zero in the second set.

statistically indistinguishable both in the pre-intervention and post-intervention phase. A Kolmogorov-Smirnov test of equal distributions in α values calculated from pre-intervention data fails to reject at a significance of p=0.802. The same test on α values calculated from post-intervention data fails to reject at p=0.618. Thus exponential-growth bias is unlikely to be eliminated by simple "nudges".

5 Experiment 3: Overconfidence about Exponential-Growth Bias

While Experiment 2 demonstrated the prevalence of exponential-growth bias in a representative population and its predictive power for asset accumulation, the third experiment demonstrates that people are substantially unaware of their bias. This unawareness is not tautological – while the law of iterated expectations requires that people cannot believe they hold biased beliefs, it is still possible that people would rationalize their systematic bias as mean-zero noise (which to an objective observer could be predicted using the model of exponential-growth bias). If people were aware of their poor decision-making over questions involving exponential growth, the bias would not long survive in the market. Tools such as calculators and spreadsheets are readily available, as are more purpose-designed online calculators. Moreover, a market for expert financial advice could easily fully de-bias consumer decisions. In this section, we show that people are unaware of their bias, and are unwilling to pay for de-biasing.

5.1 Design

This study was conducted through the Center for Neuroeconomic Studies (CNS) at Claremont Graduate University on a sample of Claremont Colleges students. It was necessary to conduct the study under laboratory conditions rather than using an online panel, in order to exercise experimental control over subjects' problem-solving resources. Subjects were not provided with any tools, and calculators/cell phones were expressly forbidden.

As with Experiment 2, subjects faced a series of questions relating the growth of two hypothetical assets. Subjects were informed that one question would be chosen randomly by computer, and that they would receive an incentive payment based on the accuracy of their response to that question (in addition, subjects received a show-up fee of \$5.00). The incentive payment was quadratic in accuracy, bounded below by zero. That is, if a subject responded r_j to a question on which the correct answer was c_j, then their payment would be $\pi = \max\left\{25 - 100 \cdot \left(1 - \frac{r_j}{c_j}\right)^2, 0\right\}$. Subjects were given examples of this payment rule in the instructions, and were provided with a table of payments corresponding to different percentage errors alongside every question.

Prior to giving their final answers, subjects indicated their willingness-to-pay (WTP) to receive the use of a spreadsheet and their WTP for the correct answer, on a question-by-question basis. The elicitation procedure was based on a modified Becker-DeGroot-Marshack mechanism. Subjects were told they would be randomized into one of four treatments. Subjects in the control group were not given the spreadsheet nor the correct answer, regardless of their WTP. Subjects in the spreadsheet treatment, were given the spreadsheet, regardless of their WTP. Subjects in the answer group would receive the correct answer, regardless of their WTP. And subjects in the incentive-compatible group would purchase the spreadsheet and the correct answer at a randomly-drawn price X if it were below their indicated WTP, and would not receive the correct answer

and would not pay anything if X were above their WTP. The first treatment allowed confirmation that subjects given the correct answer would actually use it, while the second group allowed observation of final answers among all subjects (not just those least willing to pay for help). The third group enforced incentive compatibility. Subjects were not told the distribution of treatments, or the distribution from which X was drawn. The research design is presented in Figure 6.

Figure 6: Experiment 3 Design

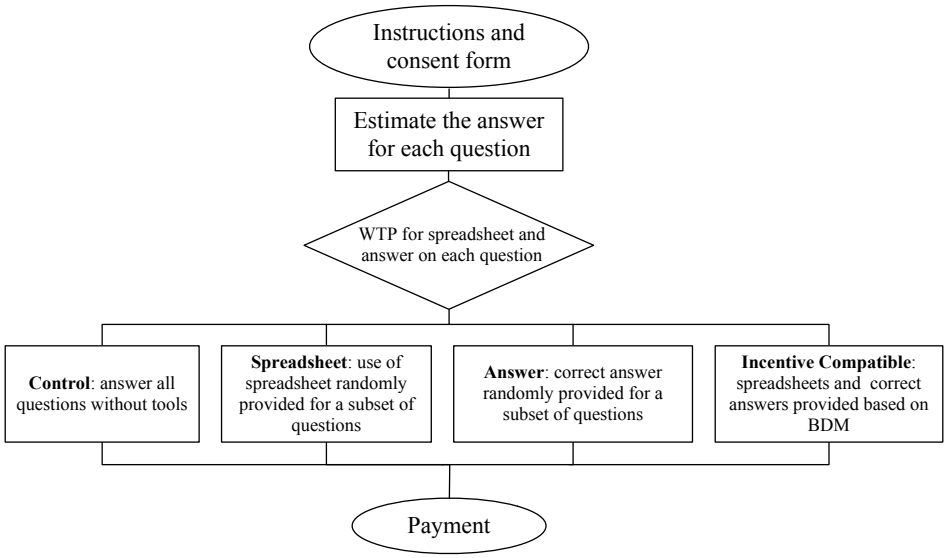

There is every indication that subjects understood the incentive-compatibility of the willingness-to-pay elicitation mechanism. The instructions explicitly stated that the optimal strategy was to enter the amount by which they expected having the correct answer would increase their earnings if the question were chosen, i.e. $\$25 - \mathrm{E}\left(100 \cdot \left(1 - \frac{r_j}{c_j}\right)^2\right)$, and were given examples of how under-bidding and over-bidding were dominated strategies. Moreover, subjects were asked what bid would maximize their expected earnings if they thought their answer with no help would earn \$9.50. Only once they answered that a WTP of \$15.50 would maximize expected earnings were they allowed to exit the instructions and proceed to the experiment. We would expect that, if anything, this would anchor their stated willingness-to-pay at \$15.50 if subjects interpreted this example as containing information about their expected performance.

Subjects were presented 32 questions across four domains, randomized first at the domain level and then within-domain. A list of all 32 questions is given in Appendix Table B.2. The domains comprised the exponential, fluctuating-interest, and periodic savings domains from experiment 1, and a new domain that presents consumers with the problem to identify the maximum front-end load on an asset that makes it of equal value to an asset that grows at a lower interest rate. Questions in this domain took the form, "Asset A has an initial value of \$100, and grows at an interest rate of 5% each period. Asset B has an initial value equal to $\$X$, and grows at an interest rate of 8% each period. What \$X will make the value of asset A and B equal at the end of 25 periods?" Subjects were asked to indicate a preliminary guess for the correct answer – with the goal that they would focus on the details of the particular question – but were told this guess

would not count for payment. They next indicated their willingness to pay for help: if they were in the IC treatment, the computer would randomly draw a number for each question. If the number was below the stated threshold, subjects would receive the tool for that question and pay the realized draw if the question were selected for payment. Otherwise they would not receive a tool and not pay if the question were selected for payment.

After subjects completed the willingness to pay elicitation task for all questions, the computer randomized them first into the four treatment groups and then drew values of X for the BDM randomizer from a uniform [0,25] distribution. There were 49 subjects who were allocated to the first treatment group, which received no help on any questions, and we will focus on these subjects in the analysis. Subjects then gave final (i.e. for-payment) answers to all 32 questions, again in an order randomized first across and then within domains. Subjects in the other treatment groups first gave final answers to questions for which they did *not* receive aid and then the questions for which they *did*. The spreadsheet group had 38 subjects and the remaining 6 subjects were assigned to the incentive-compatible group.

5.2 Results

5.2.1 Bias

As in Experiment 2, we calculate a measure of bias for each subject i and question j, namely the natural logarithm of the ratio of the given answer to the correct answer. We replicate the earlier finding that subjects are systematically biased in the direction predicted by exponential-growth bias. Figure 7 plots the distribution of log errors at the question–subject level for each of the 49 subjects' responses to each of the 32 questions. We are left with 1481 subject-question observations after dropping skips. The by-now familiar pattern emerges that subjects systematically answer too low where exponential-growth bias predicts under-estimation, and too high where it predicts over-estimation. In panel (a), where under-estimation is predicted, both the mean (-0.51) and median (-0.34) are significantly negative ($p<0.01$). Similarly, both the mean (0.26) and median (0.19) of the distribution in panel (b) are signficiantly positive ($p<0.01$). Table 6 displays under-estimation on the questions predicted to be under-estimated.

The modal error is zero, which holds both for questions where under-estimation is predicted as well as those predicting over-estimation. Apart from the zeros, however, the differences are stark. Where under-estimation is predicted, the distribution is shifted sharply to the left. Both the median (-0.349) and mean (-0.554) are significantly negative ($p<0.01$). The pattern is reversed where exponential-growth bias predicts over-estimation: the distribution is shifted sharply to the right and the mean answer (0.209) is now significantly positive ($p<0.01$), although the median answer in this case is zero.

In Figure 8, we present the distribution of subject-level averages of the above log-ratio: $\bar{e}_i = \frac{1}{32}\sum_{j=1}^{32}\ln(r_{ij}/c_j)$. If prediction errors are log-normally distributed, these distributions should collapse to unit masses at zero as the number of questions becomes very large. Once again, however, the directional bias is preserved. In panel (a), the median subject's average log-ratio was -0.42, while in panel (b) the median average log-ratio was 0.25, with both the means and medians statistically different from zero at $p<0.01$.

Figure 7: Question-Level Mistakes

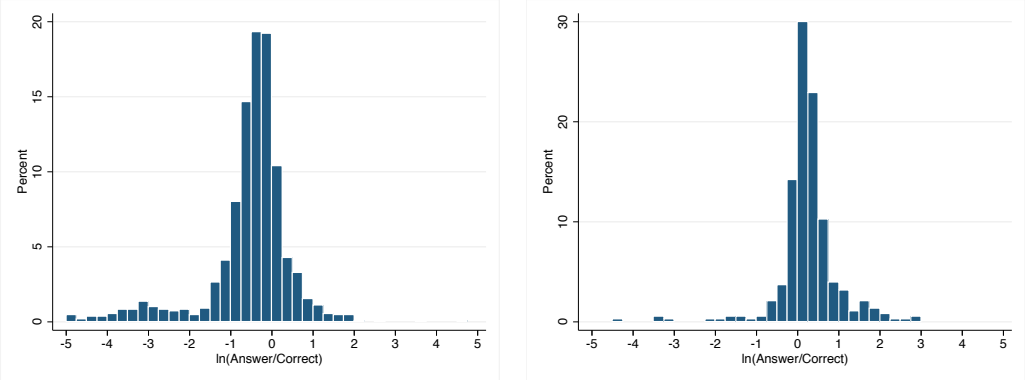

(a) Under-Estimation (b) Over-Estimation

Figure 8: Subject-Level Mistakes

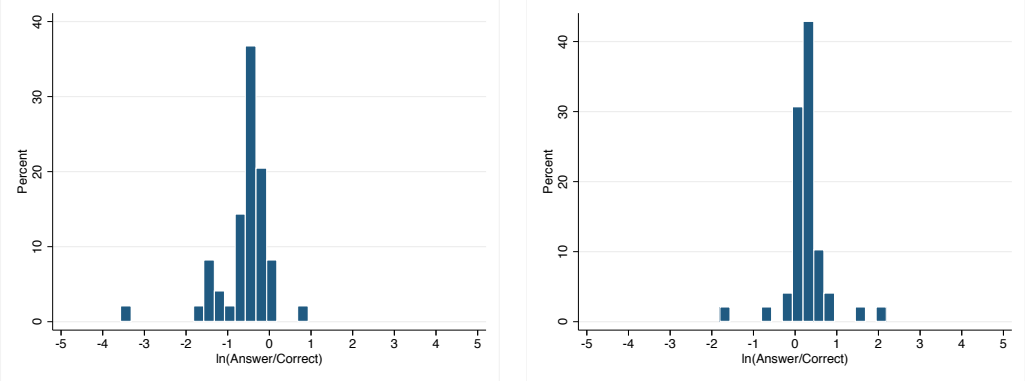

(a) Under-Estimation (b) Over-Estimation

Notes: Underestimation based on the questions for which EGB predicts a downward-biased answer; overestimation from those where an upward bias is predicted by the theory. Figure 7 shows the distribution of errors in predicted asset growth, and should be symmetric about zero if subjects' errors on a percentage basis are symmetric about zero. The means of both distributions are significantly different from zero (significant at $p < 0.01$). Figure 8 computes the mean of $\ln(answer/correct)$ at the subject level, and should converge to a point mass at zero in the absence of systematic bias. The means of both distributions are significantly different from zero (significant at $p < 0.01$).

Table 6: Underestimation

	(1)	(2)	(3)	(4)	(5)
Constant	-0.515***	-0.777***	-0.477***	-0.712***	-0.728***
	(0.069)	(0.182)	(0.078)	(0.148)	(0.148)
Has Taken Adv. Math		0.193		0.503***	0.499***
		(0.187)		(0.154)	(0.152)
Owns Stocks		0.319		0.023	0.050
		(0.213)		(0.187)	(0.184)
No Credit Card Balances		0.239		0.115	0.116
		(0.182)		(0.153)	(0.151)
Received Excel			-0.127	-0.162	-0.166
			(0.132)	(0.128)	(0.127)
WTP for Excel					0.002***
					(0.001)
WTP X Received Excel					0.000
					(0.002)
N	1,098	1,098	2,301	2,301	2,290
N_Clust	49	49	96	96	96

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

Notes: Dependent variable is $\ln(answer_{ij}/correct_j)$. Standard errors clustered by subject.

5.2.2 Overconfidence

The preceding section established that laboratory subjects were systematically biased in their responses. This section seeks to establish that subjects over-estimated both their accuracy and their precision. We begin by demonstrating that subjects systematically stated a willingness to pay for the correct answer that was below the ex post optimal level. We then show that the elicited WTP measures are too low to be justified even by the observed level of precision.

We first calculate for every question the payment that a subject would have earned had that question been chosen for implementation, according to the quadratic payment rule: $p_{ij} = \max\left\{25 - 100 \cdot \left(1 - \frac{r_{ij}}{c_j}\right)^2, 0\right\}$. Subjects answering exactly correctly would have an associated payment of \$25, while responses more than 50% from the correct answer would receive zero. The average associated payment across all 1481 subject-question pairs was \$11.28 (s.d. 10.25).

If agents are risk-neutral over \$25 stakes, then the optimal strategy would be to state a willingness to pay for the correct answer of $WTP_{ij} = \$25 - E(p_{ij})$. Any concavity in utility would set this as a lower bound, as paying for the correct answer can be viewed as providing insurance for the earnings.[17] A simple test of whether subjects accurately predicted their performance is to compare the actual willingness to pay against this bound. Subjects may under- or over-pay on some questions, but by the law of large numbers the average willingness to pay across all questions should converge to $(\$25 - \bar{p}) = \$13.94$. Instead, the mean willingness to pay is significantly lower at \$5.76 (p<0.01). That is, subjects on average expect their answers

[17] Depending on the reference point, loss aversion would also predict this to be a lower bound.

to earn at least 40% more than they actually do.

Panel (a) of Figure 9 plots the distribution of overconfidence at the subject-question level. The depicted variable is the difference between the ex-post 'optimal' WTP (i.e. \$25 less the actual associated payment) and the stated willingness to pay for the answer, normalized by \$25. Thus a value of 1 indicates that a subject would pay \$0 for an answer to a question on which they would have earned no payment, and a value of -1 indicates that a subject would pay \$25 for an answer to a question on which they would have earned the full payment. This variable should be distributed about zero if subjects are risk-neutral, or some negative number if they are risk-averse. Instead, the distribution has a positive mean (0.318), and is skewed highly positive.

The second panel of Figure 9 helps establish that this result is driven by a large fraction of subjects being systematically overconfident across all questions. Panel (b) computes the mean of the under-payment variable from panel (a) at the subject level, and plots the distribution of this subject-level outcome. A subject who over-pays on some questions but under-pays on others would of course converge towards zero as we average over a large number of questions. Instead we find that both the mean (0.31) and median (0.28) are significantly over-confident (p<0.01).

Figure 9: Overconfidence

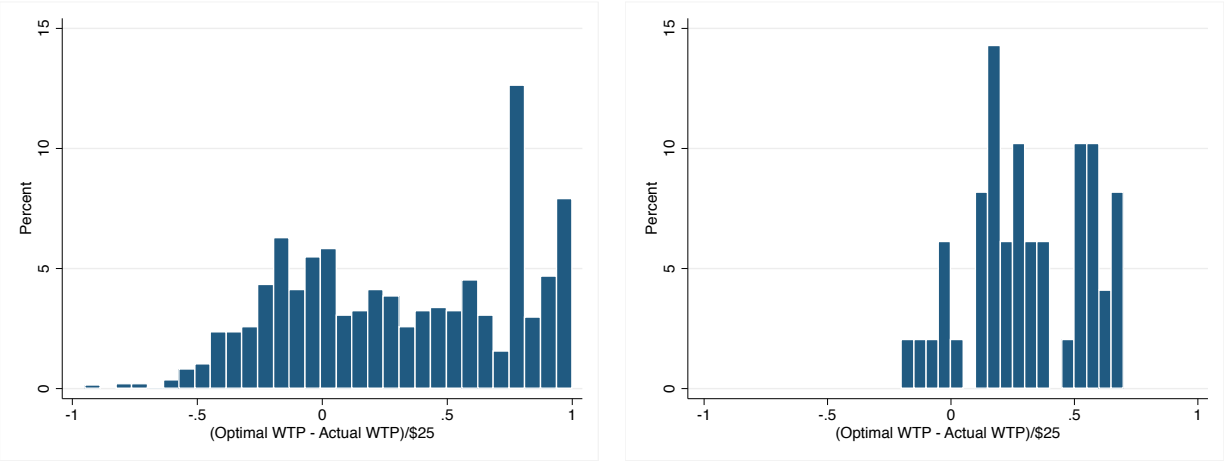

(a) Question x Subject Level

(b) Subject Averages

Notes: "Optimal WTP" is defined as \$25 less a subject's actual earnings on a question, and is therefore *ex post* optimal. Panel (a) shows the distribution of under-payment, and the mass weighted by the squared error should be equal on either side of 0 in the absence of systematic bias (or about some negative amount if subjects are risk-averse over \$25 stakes). Panel (b) computes mean under-payment at the subject level, and should converge to a point mass at zero in the absence of systematic bias (or a mass at some negative amount if subjects are risk-averse).

We next ask whether the overconfidence comes only from the requirement that subjects cannot be aware of their systematic bias, or whether they are also over-confident about the precision of their errors even conditional on there being no systematic error. Suppose an agent believes that her responses are noisy, so that $r_{ij} = (b + \eta_{ij}) \cdot c_j$ for some η_{ij} drawn i.i.d. drawn from an exponential distribution: $F_\eta(y) = 1 - e^{-\lambda y}$. An agent believing herself to be unbiased must have expectations $E[b+\eta] = 1$, and thus believe $b = 1 - 1/\lambda$. Note that as $\lambda \to 1$, any positive r_{ij} becomes possible, although realizations close to zero remain unlikely. This condition does not hold for a biased agent, and so we use the variance of the empirical distribution of

Table 7: Overconfidence

	(1)	(2)	(3)	(4)	(5)
Constant	0.318***	0.451***	0.105***	-0.142***	-0.103
	(0.034)	(0.080)	(0.035)	(0.030)	(0.062)
Has Taken Adv. Math		-0.195**			-0.055
		(0.087)			(0.060)
Owns Stocks		-0.140*			-0.162**
		(0.073)			(0.072)
No Credit Card Balances		-0.093			-0.002
		(0.080)			(0.056)
| ln(Answer/Correct) |			0.320***	1.052***	1.057***
			(0.035)	(0.070)	(0.064)
ln(Answer/Correct)^2				-0.211***	-0.213***
				(0.019)	(0.018)
N	1,481	1,481	1,478	1,478	1,478
N_Clust	49	49	49	49	49

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

Notes: Dependent variable is question-level overconfidence, defined as $(25 - p_{ij} - WTP_{ij})/25$. Standard errors clustered by subject.

a subjects' realizations of r_{ij}/c_j to estimate λ_i. The mean value across all subjects for $1/\lambda$ is 1.07. We can then simulate the subject's earnings under the counter-factual restriction that η_{ij} exponentially distributed according to λ_i, but imposing the restriction that $b = 1 - 1/\lambda_i$ to simulate an unbiased agent.

We perform this simulation exercise separately for questions on which exponential-growth bias predicts a positive and a negative bias. In both cases, the simulated responses are associated with higher earnings than subjects' actual answers: \$13.87 (s.d. 0.23) and \$14.23 (s.d. 0.38), respectively, as compared to actual means of \$10.93 and \$12.28. Subjects who were aware of the noise in their answers, but not the systematic bias, therefore ought to have a willingness to pay for the correct answer of between \$10.77 and \$11.13. This is still substantially above the observed willingness to pay of our subjects, which indicates that they are overoptimistic about the precision of their answers in addition to being unaware of their bias. Indeed, the low willingness to pay is rationalized only if the variance of η_i is one-quarter of its true value.

It is worth noting that overconfidence is increasing in the absolute value of the logratio as indicated in columns 3-5 of Table 7. This need not necessarily have been the case, if for example those who underpay the most are those who make small errors, while those who make large errors are self-aware and thus have high WTP. The finding that the the most error-prone subjects are the most overconfident is a common finding in the overconfidence literature. This suggests a pathological selection in the market, whereby the least competent avoid advice the most.

Subjects were overconfident with their answers and they were also overconfident in their ability to use a spreadsheet. Columns 3–5 of Table 6 indicate that receiving the spreadsheet had no effect on the logratios. Unlike the overconfidence analysis, we cannot measure a within-subject overconfidence in spreadsheet

ability.[18] However we can estimate a between subject treatment effect of the spreadsheet on hypothetical earnings. If the spreadsheet group had indeed paid their WTP for the spreadsheet, their earnings averaged over all questions would have been \$5.40 (p=0.003) less than the control group's earnings.

6 Conclusion

While the unintuitively rapid growth of exponential functions has been observed for ages, the economic implications have only been considered recently. We develop a model to captures all of the relevant features of EGB, and embed it in a lifecycle-consumption environment. We derive and test the implications that consumers will make very specific – and very large – errors in their consumption plans, and the results of the first experiment overwhelmingly support this. Moreover, since the bias is fundamentally about the budget constraint, the model is modular and can thus be easily married to other economic settings or extensions.

Moreover, the bias seems to prevail in the population as a whole and is a strong predictor of saving behavior even after controlling for the standard explanatory factors. And just as importantly, our second two experiments indicate that the bias will not be eliminated by competition in the marketplace. The bias was robust to an intervention designed to make exponential growth more salient, and which could be used to obtain the correct response in some domains. Perhaps more importantly, people are unlikely to seek out help since they are significantly overconfident about the magnitude of their errors.

Although laboratory experiments can be very valuable in identifying the existence and mechanisms that underly the bias (as well as pre-testing efficacious interventions), ultimately field experiments with large-stake financial decisions are needed for ecological validity — since ultimately these are the target applications of interest.[19]

Additional research on the efficacy of interventions to combat EGB in the field are necessary. While EGB was robust to the intervention in our Experiment 2, it is possible that other "nudges" designed around the predictions of our model could help improve welfare. The potential welfare consequences are quite large. For example, a back-of-the-envelope calculation suggests that EGB has a deleterious welfare impact on the order of 10% of the median consumer's income for the median value in our data of $\alpha = 0.6$.[20] An important literature has focused on failures of consumers to perceive their true (future) preferences, to which we add failure to perceive the true budget constraint as a significant additional force.

[18] Subjects reported estimates on all problems which were intended to estimate within-subject treatment effects of the spreadsheet. However, these first round estimates appear substantially lower than final responses in both the control and spreadsheet groups. Perhaps not so surprisingly, incentives matter.

[19] That is, while our Experiment 2 used a nationally representative sample, the data are still "artefactual".

[20] We take a representative agent earning the median personal income between ages 20 and 65, who then retires until death at 78. We assume discounting is already reflected in interest rates, so the risk-free rate (post-retirement) is zero and risky rate (pre-retirement) is 5%, and use log-utility. We then calibrate the scale factor by which income must increase such that an $\alpha = 0.6$ biased type achieves the same lifetime utility as an $\alpha = 1$ unbiased type.

References

Almenberg, Johan and Christer Gerdes, "Exponential Growth Bias and Financial Literacy," June 2011. Working Paper, IZA DP No. 5814.

Andersen, Steffen, Glenn Harrison, Morten Lau, and E. Elisabet Rutström, "Elicitation using multiple price list formats," *Experimental Economics*, 2006, *9*, 383–405.

Andreoni, James and Charles Sprenger, "Estimating Time Preferences from Convex Budget Sets," *American Economic Review*, Forthcoming 2012.

Andrews, Rodney J., Jing Li, and Michael F. Lovenheim, "Quantile Treatment Effects of College Quality on Earnings: Evidence from Administrative Data in Texas," 2012. NBER Working Paper 18068.

Avery, Christopher and Sarah Turner, "Student Loans: Do College Students Borrow Too MuchÑOr Not Enough?," *Journal of Economic Perspectives*, 2012, *26* (1), 165–192.

Banks, James and Zoë Oldfield, "Understanding Pensions: Cognitive Function, Numerical Ability and Retirement Saving," *Fiscal Studies*, June 2007, *28* (2), 143–170.

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

Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg, "What Accounts for the Variation in Retirement Wealth among U.S. Households?," *American Economic Review*, 2001, *91* (4).

Bertrand, Marianne and Adair Morse, "Information Disclosure, Cognitive Biases, and Payday Borrowing," *Journal of Finance*, 2011, *66* (6), 1865–1893.

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

Camerer, Colin and Dan Lovallo, "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review*, 1999, *89* (1), 306–318.

Card, David, "The Causal Effect of Education on Earnings," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 3, Elsevier Science B.V., 1999, chapter 30, pp. 1801–1859.

Cole, Shawn and Gauri Kartini Shastry, "Smart Money: The Effect of Education, Cognitive Ability, and Financial Literacy on Financial Market Participation," February 2009. Working Paper.

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

Gerardi, Kristopher, Lorenz Goette, and Stephan Meier, "Financial Literacy and Subprime Mortgage Delinquency: Evidence from a Survey Matched to Administrative Data," April 2010. Federal Reserve Bank of Atlanta, Working Paper 2010-10 April 2010.

Goda, Gopi Shah, Collen Flaherty Manchester, and Aaron Sojourner, "What Will My Account Really Be Worth? An Experiment on Exponential Growth Bias and Retirement Saving," March 2012. Financial Literacy Center Working Paper.

Gustman, Alan L., Thomas L. Steinmeier, and Nahid Tabatabai, "Financial Knowledge and Finacial Literacy at the Household Level," October 2010. NBER Working Paper 16500.

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

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

Jones, Gregory V., "Polynomial Perception of Exponential Growth," *Perception and Psychophysics*, 1976, *21* (2), 197–200.

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

Kemp, Simon, "Estimation of Past Prices," *Journal of Economic Psychology*, 1987, *8*, 181–189.

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

Kőszegi, Botond and Matthew Rabin, "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 2006, *121* (4), 1133–1166.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

Loewenstein, George, "Anticipation and the Valuation of Delayed Consumption," *The Economic Journal*, 1987, *97* (387), 666–684.

_ , Don A. Moore, and Roberto A. Weber, "Misperceiving the value of information in predicting the performance of others," *Experimental Economics*, 2006, *9*, 281–295.

_ , Ted O'Donoghue, and Matthew Rabin, "Projection Bias in Predicting Future Utility," *The Quarterly Journal of Economics*, November 2003, pp. 1209–1248.

Lusardi, Annamaria and Olivia S. Mitchell, "How Ordinary Consumers Make Complex Economic Decisions: Financial Literacy and Retirement Readiness," September 2009. Working Paper.

_ and Peter Tufano, "Debt Literacy, Financial Incentives, and Overindebtedness," April 2009. Working Paper.

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

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

Meier, Stephan and Charles Sprenger, "Present-Biased Preferences and Credit Card Borrowing," *American Economic Journal: Applied Economics*, 2010, *2* (1), 193–210.

Moore, Don A. and Paul J. Healy, "The Trouble with Overconfidence," *Psychological Review*, 2008, *115* (2), 502–517.

Munnell, Alicia, Anthony Webb, and Luke Delorme, "Retirements at Risk: A New National Retirement Index," June 2006. Working Paper.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *The American Economic Review*, Mar 1999, *89* (1), 103–124.

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

Shui, Haiyan and Lawrence M. Ausubel, "Time Inconsistency in the Credit Card Market," January 2005. Working Paper.

Skiba, Paige Marta and Jeremy Tobacman, "Payday Loans, Uncertainty, and Discounting: Explaining Patterns of Borrowing,Repayment, and Default," 2007. Working Paper.

_ **and** _ , "Do Payday Loans Cause Bankruptcy?," 2011. Vanderbilt Law School Working Paper 11-13.

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

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

_ **and** _ , "Fuzzy Math, Disclosure Regulation and Credit Market Outcomes: Evidence from Truth-in-Lending Reform," June 2009. Working Paper.

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 1956, *23*, 165–180.

Sunstein, Cass and Richard Thaler, "Libertarian Paternalism," *American Economic Review*, 2003, *93* (2), 175–179.

Timmers, Hans and Willem A. Wagenaar, "Inverse Statistics and Misperception of Exponential Growth," *Perception and Psychophysics*, 1977, *21* (6), 558–562.

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

Wagenaar, William A. and Sabato D. Sagaria, "Mispereeption of Exponential Growth," *Perception and Psychophysics*, 1975, *18* (6), 416–422.

A Proofs

Proof of Lemma 1 Define $\vec{y}^{(1)} = \vec{y} + <(z_0 - y_0), (y_0 - z_0)(1+i_0), 0, ..., 0>$. By (ii), $y_0 > z_0$. Note that $(1+i_s)p(s+1,\alpha) > p(s,\alpha)$, since:

$$
\begin{aligned}
(1+i_s)\left[\prod_{j=s+1}^{T}(1+\alpha i_j) + (1-\alpha)\sum_{j=s+1}^{T} i_j\right] &\geq \prod_{j=s}^{T}(1+\alpha i_j) + (1-\alpha)\sum_{j=s}^{T} i_j \\
(1-\alpha)i_s \prod_{j=s+1}^{T}(1+\alpha i_j) &\geq (1-\alpha)\sum_{j=s}^{T} i_j - (1+i_s)(1-\alpha)\sum_{j=s+1}^{T} i_j \\
i_s \prod_{j=s+1}^{T}(1+\alpha i_j) &\geq \sum_{j=s}^{T} i_j - \sum_{j=s+1}^{T} i_j - i_s\sum_{j=s+1}^{T} i_j = i_s - i_s\sum_{j=s+1}^{T} i_j \\
\prod_{j=s+1}^{T}(1+\alpha i_j) + \sum_{j=s+1}^{T} i_j &\geq 1
\end{aligned}
$$

since $i_j > 0$. Then $\widehat{W}_{0,T}(\vec{y}^{(1)}) - \widehat{W}_{0,T}(\vec{y}) = ((1+i_0)p(1,\alpha) - p(0,\alpha))(y_0 - z_0) > 0$.

Similarly, for $s = 2, ..., T$, recursively define $\vec{y}^{(s)} = \vec{y}^{(s-1)} + < 0, ..., (b_s - a_s^{(s-1)}), (a_s^{(s-1)} - b_s)(1 + i_{s-1}), ..., 0 >$. That is, by shifting $(y_s^{(s-1)} - z_s)$ from period $s-1$ to period s, at the interest rate i_{s-1}. By (ii), $(y_s^{(s-1)} - z_s) > 0$, and so $\widehat{W}_{0,T}(\vec{y}^{(s-1)}) - \widehat{W}_{0,T}(\vec{y}^{(s)}) = [p(s,\alpha)(1+i_{s-1}) - p(s-1,\alpha)]\,(y_{s-1} - z_{s-1})$ for all $s < T$, and equal to zero for $s = T$. From (i), however, we have that $\vec{y}^{(T)} = \vec{z}$. Thus $\widehat{W}_{0,T}(\vec{y}) < \widehat{W}_{0,T}(\vec{y}^{(1)}) < \ldots < \widehat{W}_{0,T}(\vec{y}^{(T-1)}) = \widehat{W}_{0,T}(\vec{y}^{(T)}) = \widehat{W}_{0,T}(\vec{z})$.

Proof of Proposition 1 From Equation (6), we know that we can write $c_0 p(0,\alpha) + g(c_0,\alpha) = \widehat{W}_{0,T}(\vec{y})$, where $g(c_0,\alpha)$ represents the bracketed term of (6). $\frac{\partial g(c_0,\alpha)}{\partial c_0} > 0$ because utility is concave and increasing. Now suppose we reduce period-s income by ϵ and increase period (s+1) income by $(1+i_s)\epsilon$. By Lemma 1, $\Delta\widehat{W}_{0,T}(\vec{y}) > 0$ if $s < T-1$, and $\Delta\widehat{W}_{0,T}(\vec{y}) = 0$ if $s = T-1$. Thus c_0 strictly increases for $s < T-1$ and is unchanged for $s = T-1$.

Proof of Proposition 2 By Proposition 1 it is sufficient to show that an agent will over-consume when all their income is included in their period-0 endowment; any deferment will exacerbate the over-consumption so long as all cash flows are weakly positive.

Suppose $c_0 \leq c_0^*$ and $-\frac{u'(c)}{u''(c)c} > 1$. Then from the agent's Euler condition,

$$
\begin{aligned}
u'(\hat{c}_s)\frac{p(0,\alpha)}{p(s,\alpha)}\delta^s = u'(c_0) \geq u'(c_0^*) = u'(c_s^*)\frac{p(0,1)}{p(s,1)}\delta^s \\
\Rightarrow \hat{c}_s < c_s^* \quad \text{since } \frac{p(0,\alpha)}{p(s,\alpha)} < \frac{p(0,1)}{p(s,1)}
\end{aligned}
$$

Now suppose that $\hat{c}_s \frac{p(s,\alpha)}{p(0,\alpha)} \geq c_s^* \frac{p(s,1)}{p(0,1)}$. Substituting in the Euler condition yields

$$
\begin{aligned}
\hat{c}_s \frac{u'(\hat{c}_s)}{u'(c_s^*)}\frac{p(s,1)}{p(0,1)} &\geq c_s^* \frac{p(s,1)}{p(0,1)} \\
\hat{c}_s u'(\hat{c}_s) &\geq c_s^* u'(c_s^*) \\
\Rightarrow \hat{c}_s &\geq c_s^*
\end{aligned}
$$

As $-\frac{u'(c)}{u''(c)c} > 1$ implies that $c \cdot u'(c)$ is an increasing function. This is a contradiction, as $\hat{c}_s < c_s^*$, and thus

$\hat{c}_s \frac{p(s,\alpha)}{p(0,\alpha)} < c_s^* \frac{p(s,1)}{p(0,1)}$ The budget constraint then implies:

$$c_0 + \sum_{s=1}^{T} \frac{p(s,\alpha)}{p(0,\alpha)} \hat{c}_s < c_0^* + \sum_{s=1}^{T} \frac{p(s,1)}{p(0,1)} c_s^* = y_0$$

Which is a violation of Walras' law, and therefore c_0 cannot be optimal. Thus $c_0 > c_0^*$.

Proof of Proposition 3 Differentiating Equation (6) w.r.t α yields:

$$p(0,\alpha)\frac{dc_0}{d\alpha} + c_0 \frac{\partial p(0,\alpha)}{\partial \alpha} + \frac{\partial g(c_0,\alpha)}{\partial c}\frac{dc_0}{d\alpha} + \frac{\partial g(c_0,\alpha)}{\partial \alpha} = 0 \tag{9}$$

Where again $g(c_0,\alpha)$ is the bracketed term from (6). The RHS is zero because the wealth is received lump sum in T and so there is no misperception about the wealth available in period T units.

By assumption, $p(s,\alpha) > 0$, $c_t \geq 0$. Straightforwardly,

$$\frac{\partial p(s,\alpha)}{\partial \alpha} = \sum_{j=s}^{T} \Big[i_j \Big(\prod_{\substack{k=s \\ k \neq j}}^{T} (1+\alpha i_k) \Big) - i_j \Big] \geq 0$$

Where the inequality is strict when $\alpha > 0$ and $i_t > 0$ for some $t \in [s,T]$. It remains to show that $\frac{\partial g(c_0,\alpha)}{\partial \alpha} > 0$.

First we show that $\frac{\partial}{\partial \alpha}\left[\frac{p(s,\alpha)}{p(0,\alpha)}\right] < 0$, $\forall 0 < s < T-1, i_j > 0$:

This derivative is negative if $p_\alpha(s,\alpha)p(0,\alpha) < p_\alpha(0,\alpha)p(s,\alpha)$.As shown in the proof of Lemma 1, $(1+i_s)p(s+1,\alpha) > p(s,\alpha)$, so it is sufficient to show $\prod_{j=0}^{s-1}(1+i_j)\cdot p_\alpha(s,\alpha) < p_\alpha(0,\alpha)$

We first show that $p_\alpha(s,\alpha) > (1+i_s)p_\alpha(s+1,\alpha)$:

$$p_\alpha(s,\alpha) > (1+i_s)p_\alpha(s+1,\alpha)$$

$$\sum_{j=s}^{T}\left[i_j \left(\prod_{k=s,k\neq j}^{T} (1+\alpha i_k) \right) - i_j \right] > (1+i_s) \sum_{j=s+1}^{T}\left[i_j \left(\prod_{k=s+1,k\neq j}^{T} (1+\alpha i_k) \right) - i_j \right]$$

$$\sum_{j=s}^{T}\left[i_j \left(\prod_{k=s,k\neq j}^{T} (1+\alpha i_k) \right)\right] - \sum_{j=s}^{T} i_j > (1+i_s)\left(\sum_{j=s+1}^{T}\left[i_j \frac{\left(\prod_{k=s,k\neq j}^{T}(1+\alpha i_k)\right)}{(1+\alpha i_s)} \right] - \sum_{k=s+1}^{T} i_j \right)$$

And because $(1+i_s)/(1+\alpha i_s) > 1$ it is sufficient that

$$\sum_{j=s}^{T}\left[i_j \left(\prod_{k=s,k\neq j}^{T} (1+\alpha i_k) \right)\right] - \sum_{j=s}^{T} i_j > \sum_{j=s+1}^{T}\left[i_j \left(\prod_{k=s,k\neq j}^{T} (1+\alpha i_k) \right)\right] - (1+i_s)\sum_{k=s+1}^{T} i_j$$

$$i_s \left(\prod_{k=s+1}^{T} (1+\alpha i_k) \right) - i_s > -i_s \sum_{k=s+1}^{T} i_j$$

Thus $p_\alpha(s,\alpha) > (1+i_s)p_\alpha(s+1,\alpha)$, if $j < s \to i_j \geq 0$. Then $p_\alpha(0,\alpha) > (1+i_0)p_\alpha(1,\alpha) > ... > \prod_{j=0}^{s-1} p_\alpha(s,\alpha)$.

Thus $\frac{\partial}{\partial \alpha} u'^{-1}\left(\frac{p(s,\alpha)u'(c)}{p(0,\alpha)\delta^s}\right) > 0$. Since $\frac{\partial p(s,\alpha)}{\partial \alpha} > 0$, we have term-by-term that $\frac{\partial g(c_0,\alpha)}{\partial \alpha} > 0$. For (9) to hold, then, we require $dc_0/d\alpha < 0$. Since c_0 is optimal when $\alpha = 1$, we have over-consumption for all $\alpha < 1$ and the magnitude of over-consumption is decreasing in α.

Proof of Corollary 1 Let $\vec{y} =< 0,\ldots,0,y_T >$. As given in equation (6) the perceived budget constraint

is

$$c_0 p(0,\alpha) + g(c_0,\alpha) = W(\vec{y},\alpha) = z_T.$$

The agent exhausts his wealth in the first period if $c_0 p(0,1) \geq z_T$, which from the budget constraint we re-write as $z_T \frac{p(0,1)}{p(0,\alpha)} \geq z_T + g(c_0,\alpha)\frac{p(0,1)}{p(0,\alpha)}$. Since $z_T > g(c_0,\alpha)$ whenever $c_0 > 0$ (guaranteed by Inada conditions), this will be satisfied if $\frac{p(0,1)}{p(0,\alpha)}$ is sufficiently large. Let $i_t = i$, and let $i \to \infty$. Since $\lim_{i\to\infty} \frac{p(0,1)}{p(0,\alpha)} = \lim_{i\to\infty} \frac{(1+i)^T}{(1+\alpha i)^T + (1-\alpha)iT} = \infty$, there will be some i_B such that the agent will exhaust all resources in period 0.

Proof of Proposition 4 Let $\hat{c}_{t,\tau}$ denote the agent's period-t expectation of consumption in period $\tau > t$, and let the period-t perception of final wealth net of obligations incurred in periods $\{0,...,t-1\}$ be given by:

$$\widehat{W}_{t,T} = \left[\sum_{s=0}^{t-1}(y_s - c_s)\prod_{j=s}^{t-1}(1+i_j)\right] p(t,\alpha) + \sum_{s=t}^{T} p(s,\alpha) y_s$$

We note that consumption in every period is a normal good, and from the perceived budget constraint

$$\widehat{c}_{t,t} p(t,\alpha) + \sum_{s=t+1}^{T} p(s,\alpha)\widehat{c}_{t,s} = \widehat{W}_{t,T}$$

one can see that if $\widehat{W}_{t,T}$ increases, $\widehat{c}_{t,s}$ must also increase for all $s \in \{t,...T\}$.

At time t, the budget constraint yields

$$\sum_{s=t+1}^{T} \widehat{c}_{t,s} p(s,\alpha) = \widehat{W}_{t,T} - c_t p(t,\alpha)$$

At time $t+1$,

$$\widehat{W}_{t+1,T} = \left[\sum_{s=0}^{t-1}(y_s - c_s)\prod_{j=s}^{t-1}(1+i_j)\right] p(t+1,\alpha)(1+i_t) + \left(\sum_{s=t+1}^{T} p(s,\alpha) y_s\right) + (y_t - c_t)p(t+1,\alpha)(1+i_t)$$

and the budget constraint yields

$$\sum_{s=t+1}^{T} \widehat{c}_{t+1,s} p(s,\alpha) = \widehat{W}_{t+1,T}.$$

Thus the perceived budget will decrease only if

$$\widehat{W}_{t,T} - c_t p(t,\alpha) > \widehat{W}_{t+1,T}$$

$$\left[\sum_{s=0}^{t-1}(y_s - c_s)\prod_{j=s}^{t-1}(1+i_j) + (y_t - c_t)\right][p(t,\alpha) - (1+i_t)p(t+1,\alpha)] > 0$$

$$\sum_{s=0}^{t}(y_s - c_s)\Pi_{j=s}^{t-1}(1+i_j) < 0$$

since $p(t,\alpha) < (1+i_t)p(t+1,\alpha)$. Thus $\sum_{s=t+1}^{T} \widehat{c}_{t+1,s} p(s,\alpha) < \sum_{s=t+1}^{T} \widehat{c}_{t,s} p(s,\alpha)$, and from the Euler equation each term in the sequence $c_{t+1,s} < c_{t,s}$ giving the desired result.

Proof of Proposition 5 The perceived periodic payment required per unit of initial debt is given by:

$$a(i,T,\alpha) = \frac{\left[(1+\alpha i)^T + (1-\alpha)iT\right]}{\sum_{k=1}^{T}\left[(1+\alpha i)^k + (1-\alpha)ik\right]}$$

Differentiating wrt α yields

$$\frac{\partial a}{\partial \alpha} > 0 \iff \frac{T(1+\alpha)^{T-1} - T}{\sum_{k=1}^{T}\left[k(1+\alpha i)^{k-1} - k\right]} \quad > \quad \frac{(1+\alpha i)^T + (1-\alpha)iT}{\sum_{k=1}^{T}\left[(1+\alpha i)^k + (1-\alpha)ik\right]}$$

Multiplying through and simplifying yields:

$$T(1+\alpha i)^{T-1}\sum_{k=1}^{T}(1-\alpha)ik + T(1+\alpha i)^{T-1}\sum_{k=1}^{T}(1+\alpha i)^k + (1+\alpha i)^T\sum_{k=1}^{T}k$$

$$> (1-\alpha)iT\sum_{k=1}^{T}k(1+\alpha i)^{k-1} + (1+\alpha i)^T\sum_{k=1}^{T}k(1+\alpha i)^{k-1} + T\sum_{k=1}^{T}(1+\alpha i)^k$$

On inspection, the first term on the left-hand side is greater than the first term on the right-hand side. Similarly,

$$\begin{aligned}\sum_{k=1}^{T}(1+\alpha i)^T\left[(1+\alpha i)^{k-1}T + k\right] &> \sum_{k=1}^{T}(1+\alpha i)^T\left[(1+\alpha i)^{k-1}k + T\right] \\ &> \sum_{k=1}^{T}\left[(1+\alpha i)^T(1+\alpha i)^{k-1}k + (1+\alpha i)^k T\right]\end{aligned}$$

B Additional Figures and Tables

Table B.1: Experiment 2 Questions

Domain	Question	Asset A: T	Asset B: T	Asset A: P_0	Asset B: P_0	Asset A: i	Asset B: i	Correct X	Full-bias X
Exponential	1	20	20	100	X	10%	0%	672.75	300
	2	50	50	100	X	5%	0%	1146.74	350
	3	1	1	100	X	4%	0%	104	104
	4	10	10	500	X	-8%	0%	217.19	100
	11	2	2	100	X	50%	0%	225	200
	12	35	35	100	X	7%	0%	1067.66	345
	13	20	20	1000	X	-4%	0%	442.00	200
	14	10	10	1000	X	-10%	0%	348.68	0
Fluctuating i	5	20	20	100	X	0% in odd; 21% in even	0%	672.75	310
	6	24	24	100	X	4% in odd; 22% in even	0%	1740.68	412
	7	24	24	500	X	-20% in odd; 25% in even	0%	500	800
	15	16	16	100	X	0% in odd; 40% in even	0%	1475.79	420
	16	20	20	500	X	-30% in odd; 30% in even	0%	194.71	500
	17	36	36	100	X	-28% in odd; 59% in even	0%	1140.60	658
Catch-up Savings	8	20	15	100	X	13%	13%	184.24	122.03
	9	10	24	500	X	8%	8%	170.23	308.22
	10	10	9	500	X	25%	25%	625	538.46
	18	40	20	100	X	7%	7%	386.97	158.33
	19	20	10	100	X	13%	13%	339.46	156.52
	20	14	40	100	X	5%	5%	28.12	56.67
Periodic Savings	21	40	40	+10/period	X	7%	0%	2136.10	974
	22	20	20	+10/period	X	10%	0%	630.02	410
	23	30	30	+10/period	X	12%	0%	2702.93	692
Portfolio	24	40	40	100	(50,X)	10%	(5%,15%)	15.58	50
	25	36	36	100	(50,X)	4%	(6%,1%)	2.12	100
	26	30	30	100	(50,X)	12%	(8%,4%)	768.60	200

Figure B.1: Example of Static Task for Experiment 1

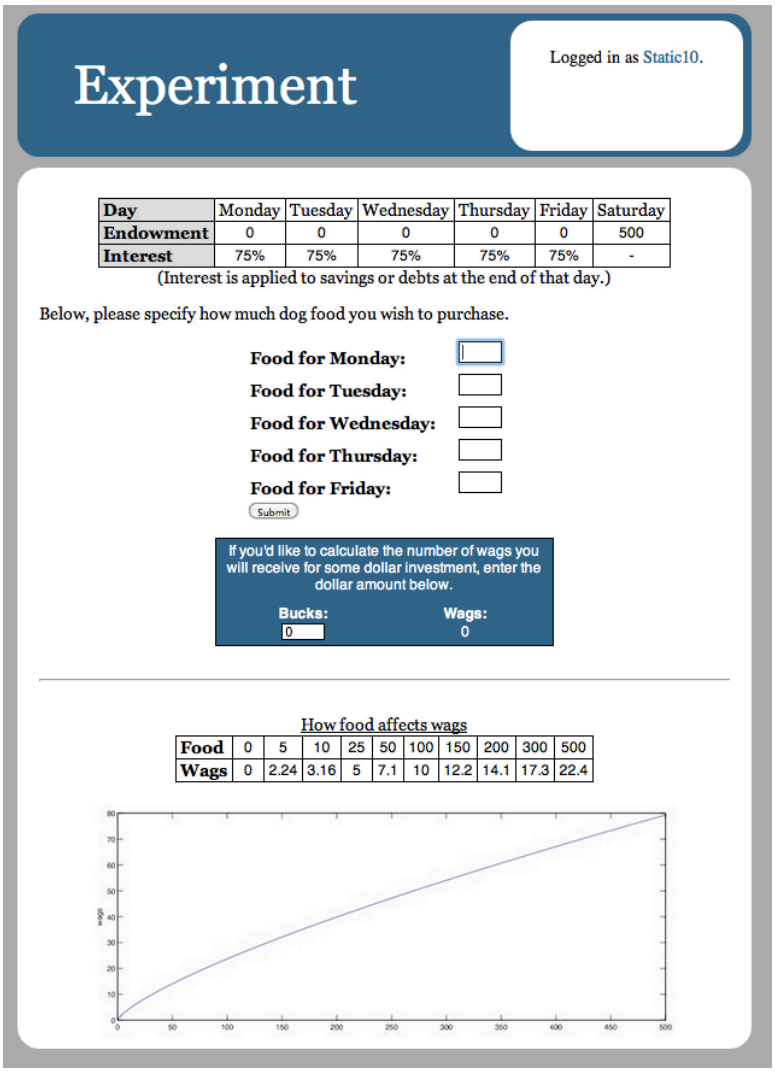

Experiment

Logged in as Static10.

Day	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Endowment	0	0	0	0	0	500
Interest	75%	75%	75%	75%	75%	-

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

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

Food for Monday:
Food for Tuesday:
Food for Wednesday:
Food for Thursday:
Food for Friday:

Submit

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

Bucks: 0 **Wags:** 0

How food affects wags

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

Figure B.2: Example of Dynamic Task for Experiment 1

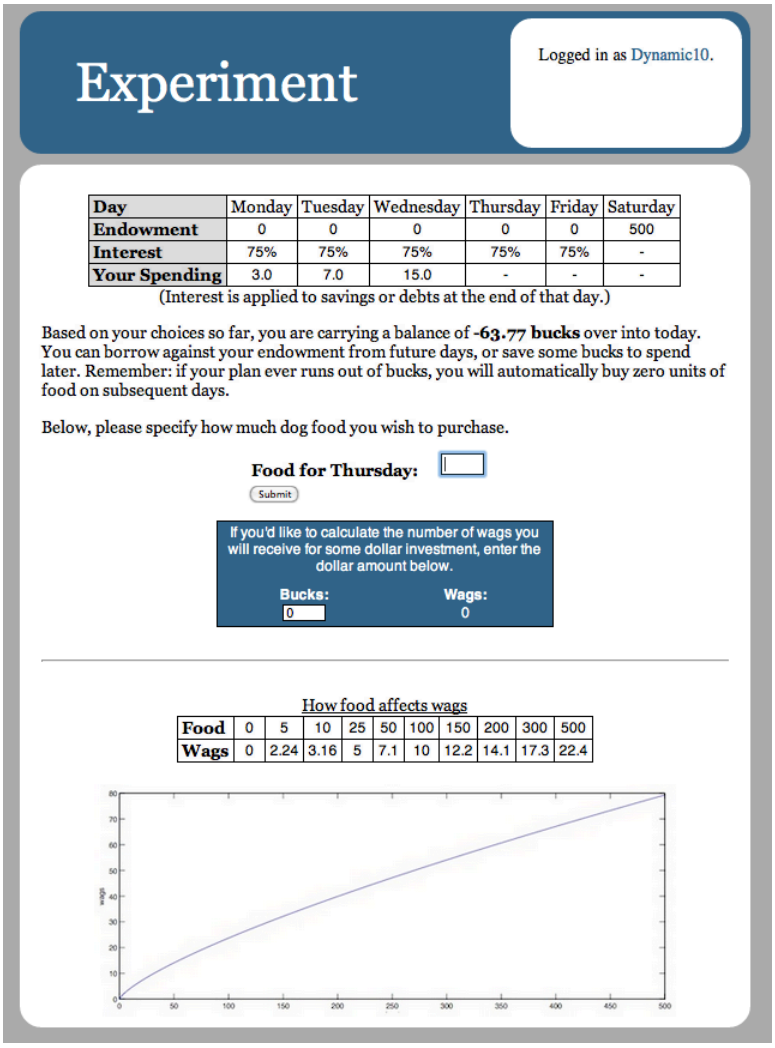

Experiment

Logged in as Dynamic10.

Day	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Endowment	0	0	0	0	0	500
Interest	75%	75%	75%	75%	75%	-
Your Spending	3.0	7.0	15.0	-	-	-

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

Based on your choices so far, you are carrying a balance of **-63.77 bucks** over into today. You can borrow against your endowment from future days, or save some bucks to spend later. Remember: if your plan ever runs out of bucks, you will automatically buy zero units of food on subsequent days.

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

Food for Thursday:

Submit

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

Bucks: 0 **Wags:** 0

How food affects wags

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

Figure B.3: Screenshot of the Graphical Intervention for Experiment 2

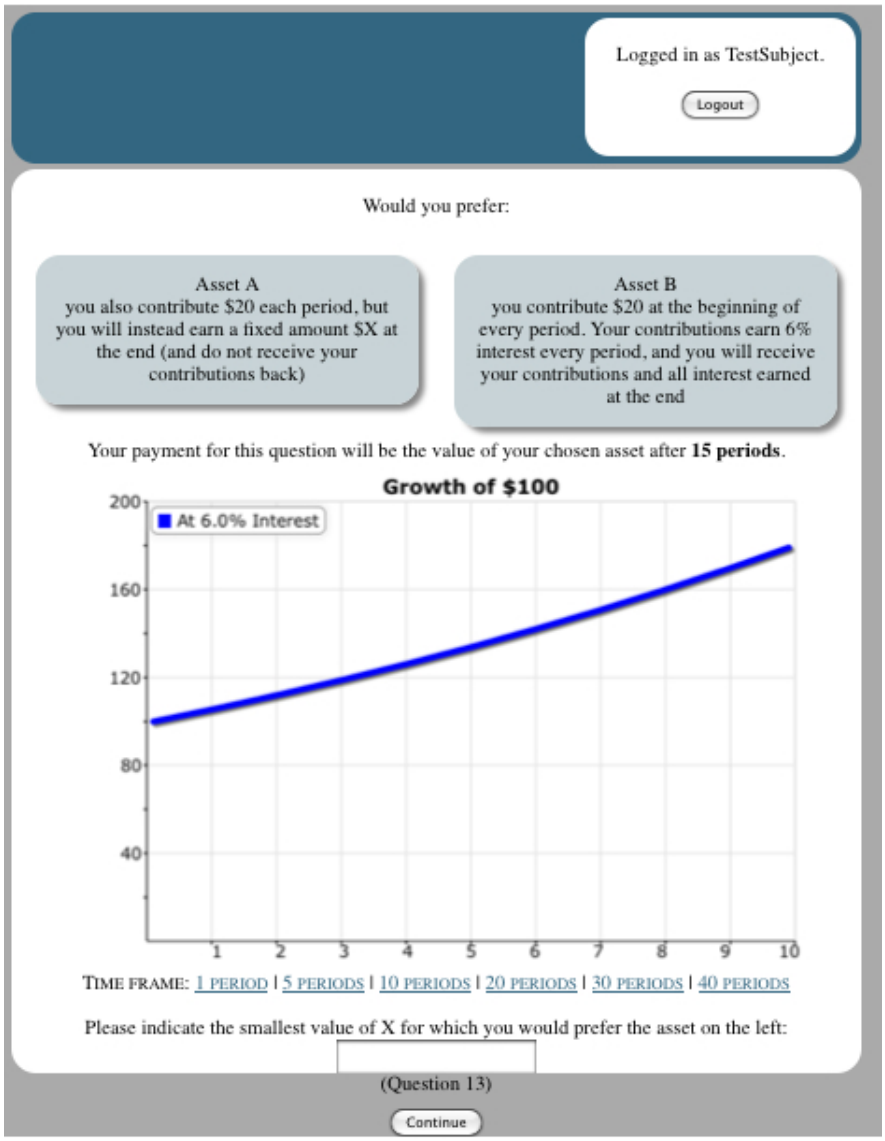

Table B.2: Experiment 3 Questions

Domain	Question	Asset A: T	Asset B: T	Asset A: P_0	Asset B: P_0	Asset A: i	Asset B: i
Exponential	1	10	10	100	X	10%	0%
	2	20	20	200	X	5%	0%
	3	10	10	90	X	25%	0%
	4	35	35	30	X	7%	0%
	5	30	30	120	X	4%	0%
	6	10	10	500	X	-8%	0%
	7	20	20	1000	X	-4%	0%
	8	8	8	1000	X	-15%	0%
Periodic Savings	9	30	30	+10/period	X	5%	0%
	10	30	30	+15/period	X	2.5%	0%
	11	20	20	+12/period	X	10%	0%
	12	20	20	+20/period	X	5%	0%
	13	15	15	+20/period	X	6%	0%
	14	20	20	+20/period	X	6%	0%
	15	25	25	+4/period	X	5%	0%
	16	50	50	+6/period	X	5%	0%
Front-end Load	17	25	25	100	X	5%	8%
	18	25	25	100	X	5%	10%
	19	20	20	100	X	10%	13%
	20	10	10	200	X	10%	20%
	21	25	25	40	X	8%	5%
	22	25	25	100	X	8%	5%
	23	20	20	50	X	13%	10%
	24	30	30	100	X	5%	4%
Fluctuating i	25	10	10	100	100	40% in odd; 0% in even	X%
	26	14	14	50	50	30% in odd; 0% in even	X%
	27	10	10	100	100	50% in odd; 0% in even	X%
	28	6	6	50	50	100% in odd; 0% in even	X%
	29	20	20	100	100	10%	-20% in odd; X% in even
	30	6	6	300	100	0%	-40% in odd; X% in even
	31	30	30	30	30	10%	-10% in odd; X% in even
	32	20	20	20	20	20%	-20% in odd; X% in even