Exponential Growth Bias and Household Finance*

Victor Stango
Graduate School of Management
University of California, Davis

Jonathan Zinman
Department of Economics
Dartmouth College

September 2008

Abstract

We document two widespread biases in how consumers perceive the costs and benefits of borrowing and saving, and explore the implications of these biases for household finance. *Payment/interest bias* is the systematic tendency to underestimate a loan interest rate given other loan terms. *Future value bias* is the systematic tendency to underestimate a future value given a present value, time horizon and periodic rate of return. We show that these biases may have a single cognitive source: *exponential growth bias*, the pervasive tendency to linearize exponential functions. Most importantly we show that these biases seem to affect household decisions and outcomes. Conditional on a rich set of household characteristics, a household-level metric of payment-interest bias is strongly correlated with borrowing, savings, portfolio choice, wealth and delegation. There is only weak evidence that our measure of bias merely proxies for broader financial sophistication. In all the results suggest that exponential growth bias represents a new class of behavioral biases that should be modeled in theoretical and empirical work on household finance.

JEL codes: D1, D9, G11

Keywords: exponential growth bias, household finance, intertemporal choice, psychology and economics, behavioral economics, behavioral finance, present-bias, financial sophistication, financial literacy

---

* Contacts: vstango@ucdavis.edu, jzinman@dartmouth.edu. Thanks to Jonathan Bauchet for research assistance and to Bob Avery and Art Kennickell for discussions on the 1983 Survey of Consumer Finances. Thanks to Andy Bernard, Dan Benjamin, James Choi, Xavier Gabaix, Al Gustman, David Laibson, Anna Lusardi, Ted O'Donoghue, Jesse Shapiro, Jon Skinner, Doug Staiger, and seminar/conference participants at the Yale SOM Behavioral Science Conference, Kellogg, Cornell, UC Davis, Georgetown, IZA, SITE, the AEA Annual Meetings, the Federal Reserve Banks of Boston, Chicago, and Philadelphia, the Federal Reserve Board, the Federal Trade Commission, the University of Michigan Retirement Research Center, the Dartmouth Economics Department and the Dartmouth Social Psychology Research Interest Group for comments. Special thanks to Chris Snyder for help with the math of exponential growth bias. Previous versions of this paper circulated with the titles "Fuzzy Math and Household Finance: Theory and Evidence", and "The Price is not Right…".

I. Introduction

What drives household financial decisions? The canonical economic model assumes that consumers choose to consume, borrow, or save based on their preferences, expectations, and the costs and benefits of borrowing and saving. A growing body of work applies specific insights from psychology to enrich specifications of three of the model's key pieces: preferences, expectations, and problem-solving conditional on parameter values.[1]

We bring psychological evidence to bear on a fourth specification issue: how consumers perceive the costs and benefits of borrowing and saving. We tie together existing and new evidence on these perceptions to show that most consumers err systematically when given information commonly available in the market. On the saving side, consumers display *future value bias*: a systematic tendency to underestimate a future value given a present value, time horizon and rate of return.[2] On the borrowing side, we present new evidence that consumers display *payment/interest bias*: a systematic tendency to underestimate a loan interest rate given a principal, monthly payment and maturity. The biases vary asymmetrically with maturity: future value bias increases with the time horizon, and payment/interest bias declines with maturity.

The striking thing about these perceptions of costs and benefits is not that consumers make *mistakes*, but that the mistakes are *biased* in particular ways. The "wisdom of crowds" fails here, and fails to a greater or lesser degree depending on side of the balance sheet and maturity. What explains this particular pattern? And is the pattern indicative of biases that affect actual decisions?

We show that future value bias and payment/interest bias are potentially linked by a single cognitive microfoundation: *exponential growth bias,* the tendency to linearize functions containing exponential terms when assessing them intuitively. A literature in cognitive psychology documents that individuals display exponential growth bias in a variety of contexts, and that the degree of exponential growth bias varies substantially in the cross-section. But economics has largely ignored the potential implications of exponential growth bias for household finance.[3]

The intuition for how exponential growth bias drives future value bias is clear: consumers underestimate how quickly a given yield compounds, so they underestimate the expected future

[1] We borrow this three-pronged taxonomy from DellaVigna's (2007) review of field evidence on psychology and economics. For a review focused on behavioral finance see Barberis and Thaler (2003).
[2] Future value bias is our term for the tendency documented most directly in Eisenstein and Hoch (2005).
[3] The cognitive psychology literature began with Wagenaar and Sagaria (1975); we provide a brief review in Web Appendix A. Economic applications of exponential growth bias to-date have been limited to perceptions about savings (Eisenstein and Hoch 2005) and inflation (Keren 1983; Jones 1984; Kemp 1984). Exponential growth bias does not appear in any of the many reviews of psychological evidence for economists; see., e.g., Rabin (1998), Glovich, Griffin, and Kahneman (2002), Kahneman (2003).

value for any given future date. Future value bias becomes more pronounced as the periodic return rises and the compounding horizon lengthens. On the borrowing side, exponential growth bias is mathematically equivalent to failing to account for the declining principal balance on an installment loan. So consumers *over*estimate how long they actually get to borrow the principal, thereby underestimating the true cost of borrowing. Payment/interest bias is more severe on short-term loans because principal balances on those loans decline faster than on long-term loans.[4]

We then examine the critical question for household finance: does exponential growth bias affect household balance sheets in the real world? We are not aware of any prior work on this question. To answer it we construct a household-level measure of payment/interest bias, and correlate it with a wide range of household financial outcomes. Bias matters: payment/interest bias is strongly correlated with more borrowing, less saving, portfolios tilted toward short-term installment debt and short-term assets, and lower net worth.[5] All of these results are conditional on controls for demographic and life-cycle factors, available resources, preferences, expectations and other decision inputs.[6] While our data lack a direct measure of future value bias, the pattern of results suggests that payment/interest bias captures future value bias as well; in particular, payment/interest bias is correlated with asset allocation conditional on the level of assets.

The findings motivate four follow-on questions. First, why doesn't consumer adaptation (learning, calculators, heuristics, etc.) render bias irrelevant? In fact we find that many consumers do effectively debias themselves by relying on outside financial advice. More biased households get more outside advice, all else equal, suggesting that many consumers are aware of their bias and/or its effects. And more-biased households who get outside advice are just as wealthy as the least-biased households. Yet our results also suggest that many biased households do not delegate, learn rapidly enough, or otherwise undo the effects of bias. Psychology again offers an

[4] We treat the link between exponential growth bias and borrowing cost perceptions formally in Web Appendix B, but for intuition the limiting case in the other direction is instructive. The formula for the interest rate on an infinite-maturity (interest-only) loan is $i=p/L$, the periodic payment divided by the principal; it does not involve any exponentiation and the principal balance never declines, so exponential growth bias (or failure to account for declining principal balances) is not an issue.

[5] Methodologically speaking, empirical work testing relationships between an individual-level measure of a potentially biased decision input and household/consumer financial choices is rare. Ashraf, Karlan, and Yin (2006) and Meier and Sprenger (2007) use survey questions to construct measures of time-inconsistent preferences and then examine relationships between preferences and saving or borrowing decisions. Puri and Robinson (2007) examine relationships between a measure of optimism based on life expectancy and financial decisions in the 1995-2001 Surveys of Consumer Finances. Graham, Harvey, and Puri (2007) summarize and extend the corporate finance literature on links between managerial attitudes (e.g., preferences and beliefs) and firm behavior.

[6] Our Web Appendix C details the full set of controls. We use the 1983 Survey of Consumer Finances because no more recent dataset has data on biased interest rate perceptions. In the Conclusion we note that the expansion and increased sophistication of retail financial markets may make biased perceptions even more relevant today (despite the growth of low-cost decision aids).

explanation: cognitive biases tend to persist and decision-making heuristics to fail when decisions are abstract and made infrequently (Stanovich 2003). Many borrowing, saving, and portfolio decisions in household finance seem to fit that description.

A second question is why supply-side factors or regulation fail to eliminate the effects of bias. We do find that credit constraints play a mitigating role, by preventing some biased households from borrowing as much as they would like. Truth-in-Lending law could make payment/interest bias irrelevant by forcing lenders to disclose an annual percentage rate (APR), but the APR disclosure mandated by Truth-in-Lending is imperfectly enforced. Many lenders use "monthly payment" marketing that shrouds or misrepresents interest rates, itself *prima facie* evidence that bias matters in the market given that violating Truth in Lending is costly.[7] Our companion paper examines the issue in further detail, and shows that consumers with greater payment/interest bias pay higher loan interest rates (Stango and Zinman 2007).[8] On the saving side, firms selling saving and investment products have incentives to debias consumers, but regulation may hinder them from highlighting returns over long horizons, where future value bias is most severe.[9]

A third question is whether our results reflect the specific effects of exponential growth bias, or whether bias is a measure of low financial sophistication defined more broadly.[10] On the asset side of the balance sheet, we conduct additional tests by estimating conditional correlations between our bias and standard indicators of sophistication, focusing on outcomes that would not necessarily be driven by exponential growth bias in its narrow form. The most-biased households are less likely to hold bonds, but the correlation is economically small. There is also some evidence of a relationship between bias and poor diversification. On two other indicators – holdings of own-company stock and frequent stock trading – we find no correlations with bias. On the liability side of the balance sheet, a standard hypothesis is that sophistication reduces the

[7] See Gabaix and Laibson (2006) for a model of a shrouding equilibrium.

[8] More specifically, we find that biased consumers pay higher interest rates on short-term installment loans, but only when borrowing from lenders that facing relatively weak Truth-in-Lending enforcement. Imperfectly enforced Truth-in-Lending may also have the perverse consequence of creating folk wisdom that using interest rates is the "right" way to make decisions and thereby nudging some biased consumers away from an effective decision rule: "Never try to infer an interest rate. Rather, make borrowing decisions based on other loan terms."

[9] E.g., SEC rule 230.482 requires mutual funds that advertise performance data to present 1-, 5-, and 10-year returns with equal prominence. A mutual fund that wishes to present returns earned over a longer horizon can do so, but only in addition to the 1-, 5-, and 10-year horizons, and with equal prominence. Our findings may also explain why mutual funds would highlight arithmetic rather than geometric mean fund returns. See Welch (2000) for a discussion of the difference.

[10] Several papers have found positive correlations between broader measures of financial sophistication (or planning, or cognitive ability) and stock market participation or wealth, e.g.: Ameriks, Caplin, and Leahy (2003); Lusardi and Mitchell, (2007); Benjamin, Brown, and Shapiro (2006); Christelis, Jappelli, and Padula (2006); van Rooij, Lusardi, and Alessie (2007).

participation cost of borrowing. Hence under this hypothesis one might expect our more biased households to hold *less* debt if unmeasured sophistication were driving our results.[11] We find little evidence of this pattern; short-term borrowing *increases* with bias, and long-term borrowing is uncorrelated. Overall then there seems to be a weak relationship between bias and lack of financial sophistication more broadly. Nevertheless the results do not rule out a link between bias and financial sophistication, and we hope that they will provoke further inquiry. Perhaps, for example, being aware of one's bias is a component of financial literacy.

Fourth and finally, it is possible that our measure of bias is correlated with unobserved elements of preferences or expectations. Our controls do include measures of time preference, risk aversion, and income expectations, making it unlikely that they are omitted variables driving the results. However, we lack measures of "behavioral" biases such as time-inconsistency, loss-aversion, or optimism. So it may be the case that individuals with exponential growth bias have biases in other dimensions as well, and that those other biases drive our observed relationships between payment/interest bias and financial decisions. This is a promising line of inquiry for future theoretical and empirical work. One intriguing possibility is that exponential growth bias is a tractable way to measure a portfolio of behavioral biases.

In all, the findings offer a new class of psychological biases that might affect household finance. Previous work has incorporated psychology-based specifications of preferences, expectations, and problem-solving.[12] But most work in household finance continues to assume that consumers correctly perceive the decline (increase) in future consumption that results from borrowing (saving) today. Our findings suggest that exponential growth bias leads consumers to get those assessments wrong, and to err systematically in particular directions that tilt portfolios toward short-term debt and away from long-term saving, increase borrowing and reduce saving, and depress overall wealth accumulation.

The paper proceeds as follows. Section II presents evidence showing that consumers display both future value bias and payment/interest bias. Section III shows that exponential growth bias can explain both biases, and also discusses some other explanations for the pattern of biases. Section IV describes our approach to estimating the link between payment/interest bias and

---

[11] Sophistication might instead push households to borrow less and save more; e.g., if sophisticates recognize subtle future risks (long-term care costs, reductions in social insurance) and others do not.

[12] For heuristic alternatives to dynamic optimization see, e.g., Lettau and Uhlig (1999), Hurst (2006), and Benartzi and Thaler (2007). There is also a related literature on financial planning; see e.g., Ameriks, Caplin, and Leahy (2003) and Lusardi (2003). For alternative formulations of beliefs see, e.g., Brunnermeier and Parker (2005) and Puri and Robinson (2007). For alternative formulations of preferences see, e.g., Angeletos, Laibson, Repetto, Tobacman, and Weinberg (2001), Barberis, Huang, and Santos (2001), and Gul and Pesendorfer (2004).

household financial outcomes, and also shows summary data on our outcomes and control variables. Section V presents our results. Section VI discusses complementary/alternative interpretations of the results. Section VII concludes.

II. Payment/Interest Bias and Future Value Bias: Evidence

In this section we discuss previous work showing empirical evidence of payment/interest bias and future value bias, present some new empirical evidence of the former, and summarize the stylized facts that one can draw from all of the work to-date.

A. Prior Work

Eisenstein and Hoch (2005) present lab data showing that most consumers display future value bias. Their study asks internet survey participants to estimate a future value given a present value, time horizon, and interest rate. Eisenstein and Hoch show that future value bias is prevalent (over 90% of respondents err on the low side), large on average, and increasing in the time horizon.[13] Respondents display a strong tendency to anchor on a linear forecast of the future value, and to ignore the returns provided by compounding.

On the borrowing side, several previous studies contain empirical evidence that consumers make mistakes when assessing interest rates.[14] Most studies establish this by asking respondents to estimate the interest rate implied by a given loan principal, maturity, and repayment stream. This work includes Juster and Shay (1964), National Commission on Consumer Finance (1972), Day and Brandt (1974), Parker and Shay (1974) and Kinsey and McAlister (1981). More recently, Bernheim (1995; 1998) and Moore (2003) find evidence consistent with limited understanding of loan terms, including interest rates.

The focus of prior work on the borrowing side is noteworthy; it primarily measures whether consumers make *mistakes* in assessing interest rates, rather than whether mistakes are *biased* in particular directions. The empirical implications of (presumably mean-zero) mistakes are different from the implications of bias, a point we elaborate on below. But despite the focus of

---

[13] Lusardi and Mitchell (2007) show that responses to a question on savings yields in the Health and Retirement Study (HRS) are consistent with the underestimation of compound yields. We note however that the HRS question does not necessarily capture a *bias* per se: its multiple choice format provides respondents with options that underestimate the yield implied by the question, but *not* with options that *overestimate* the yield.

[14] Studying how consumers infer rates from other loan terms was motivated by lender marketing practices that emphasized monthly payments and obscured or omitted interest rates (National Commission on Consumer Finance 1972). Policymakers view accurate and unbiased perceptions of interest rates as critical because rates potentially provide a standard unit of comparison for loans with different maturities, and for loans to savings instruments with returns stated as interest rates.

previous work on measuring mistakes, it is often easy to infer from summary data provided in the papers that consumers display payment/interest bias. Some papers do make more direct statements about bias; e.g., Parker and Shay (1974) note that consumers display "a strong tendency to underestimate annual percentage rates of charge by about one-half or more…"

B. New Evidence: Payment/Interest Bias on Hypothetical Loans

We build on this prior work in several ways. We start by presenting nationally representative empirical evidence on payment/interest bias from two previously untapped sources: the 1983 and 1977 Surveys of Consumer Finances. We use the 1983 SCF because it has the most recent (and, as far as we know, the only) nationally representative data on both payment/interest bias and household financial outcomes. We use the 1977 SCF because it contains richer data on payment/interest bias than the 1983 survey; the downside of the 1977 SCF is that is lacks comprehensive data on the household balance sheet. More recent SCFs lack any questions that elicit payment/interest bias and hence are not usable for our purposes.

We measure payment/interest bias using two hypothetical questions that appear in both the 1977 and 1983 SCFs.[15] The first question is:

> *"Suppose you were buying a room of furniture for a list price of $1,000 and you were to repay the amount to the dealer in 12 monthly installments. How much do you think it would cost in total, for the furniture after one year -- including all finance and carrying charges?"*

The response to this first question is a lump sum *repayment total* (e.g., $1200). Given the predefined maturity and principal amount, the repayment total yields i^*, the *implied APR*[16] per the respondent's self-supplied repayment total.[17] Figure 1a shows the distribution of the implied APR in the 1983 SCF across all households. The mean is 57 percent, which corresponds to a stream of

[15] The survey respondent is whomever was determined to be the "most knowledgeable about family finances." We use the terms "household," "individual", "consumer" and "borrower" interchangeably.

[16] See equation (4) below for a formal definition of the APR. Although the SCF question does not specify a particular definition of "rate of interest", we use the APR as our benchmark because: a) it has been the standard unit of comparison for borrowing costs in the U.S. since the enactment of Truth in Lending law in 1968; b) it is the rate respondents supply when asked about the most prevalent type of loan, home mortgages. Using alternative benchmarks such as the Effective Annual Rate, which is higher than the APR and may be a better measure of true borrowing costs, does not change the results because we use cross-sectional variation in perceptions; i.e., we use relative and not absolute bias (see equation (8)).

[17] We assume that the monthly installment payments are equal when calculating the implied APR. Different assumptions about payment arrangements do not change the qualitative results that respondents generally underestimate interest rates (even if we assume that the first eleven payments are zero, and the last completely repays the loan). More important, while such transformations change the level measure of misperception they do not alter the cross-sectional ranking in misperception. As noted directly above it is this ranking that provides identification in our empirical tests below.

payments over the year totaling roughly $1350. The modal implied APR is 35% ($1200), with other frequent rates corresponding to round repayment totals ($1300, $1400, $1100, etc.). The twenty-fifth percentile is 35% and the seventy-fifth is 81% ($1500).

The next question in the survey is:

"What percent rate of interest do those payments imply?"

This response is i^p, the *perceived APR*. Figure 1b shows the distribution of perceived APRs. The perceived rate distribution has a lower variance than the actual rate distribution but the perceived rate is still correlated with the actual rate; the correlation is 0.46 among those with implied APRs below the median.

Payment/interest bias is the difference between the perceived and implied APRs. Figure 2a presents a histogram of payment/interest bias in the 1983 SCF. Over 98% of respondents underestimate the actual rate. Roughly twenty percent of respondents give the "simple" or "add-on" rate (e.g., a repayment total of $1200 yields a perceived rate of 20%).[18] But responses are biased even relative to this rate; those who supply something other than the add-on rate tend to underestimate relative to the add-on (Figure 2b). The size of bias is quite striking, although not integral to our empirical work (which focuses on cross-sectional differences in bias). The median bias is -25 percentage points (2500 basis points), and the mean bias is -38 percentage points.[19] Table 1 shows tabular data on payment/interest bias in both the 1983 and 1977 SCFs. The data show that bias is similar in both surveys, although it is slightly smaller in the 1977 data. We stratify bias into the quintiles that we use to measure *relative* differences in bias for our analysis of whether bias affects decisions.

While we do not know of any more recent representative data measuring payment/interest bias, there is one bit of corroborating contemporary evidence. Following an internal presentation of this paper, a skeptical colleague gave an updated version of the SCF questions to students in a finance class that had recently covered discounting. Of thirty-seven students, all underestimated the APR: one gave a rate above the add-on rate, twelve gave the add-on rate, and the remainder underestimated relative to both the APR and the add-on rate.

---

[18] The add-on rate does not account for declining principal balances (i.e., for the fact that installment loan repayments include principal as well as interest), and therefore understates the APR. The add-on rate was the prevailing way of quoting loan terms before the enactment of the Truth-in-Lending Act (National Commission on Consumer Finance 1972). Our empirical work treats those supplying the add-on rate differently (by giving them their own intercept) conditional on the absolute level of bias

[19] The Juster and Shay (1964) results allow one to infer something about the size of payment/interest bias. Average bias in their sample of Consumers Union members is substantial (1500 basis points) but smaller than in our samples.

C. New Evidence: Bias on Actual Loans

Both the 1983 and 1977 SCFs also contain self-reported interest rates on *actual* loans: on all installment loans in the 1977 SCF, and on mortgages in the 1983 SCF. This is useful because with self-reported data on principal, maturity and payments, we can calculate the implied APR on each loan, assuming that consumers report non-interest loan terms accurately. This allows us to ask whether consumers also display payment/interest bias on actual loans, and moreover whether payment/interest bias varies with loan maturity (recall that the hypothetical question concerns only a one-year maturity).

Table 2 presents summary data on payment/interest bias on all actual non-mortgage installment loans in the 1977 SCF, and all actual mortgages in the 1983 SCF.[20] The data reveal substantial payment/interest bias on short-term loans; for the shortest-maturity loans actual rates average 30 percent while perceived rates average 13 percent. Payment interest bias on actual loans is positively correlated with payment/interest bias on hypothetical loans. This is evident from the bottom two panels of Table 2.

The other striking result is that bias falls with maturity, and is close to zero for the longest-maturity installment loans and mortgage loans (which themselves tend to have 15-30 year maturities). On mortgage loans virtually everyone is unbiased; 96% provide the correct APR.

D. Summary of the Evidence on Payment/Interest and Future Value Bias

There are three sets of stylized facts on how consumers intuitively perceive the costs and benefits of intertemporal tradeoffs. First, consumers systematically display future value bias in the lab. Second, they systematically display payment/interest bias on both hypothetical and actual loans. Third, the severity of each bias depends on the time horizon. Future value bias is more severe for long-term savings, while payment/interest bias is more severe on short-term loans.

When looking at these facts it is not surprising to see that consumers make mistakes, or even that they make large mistakes. The math of interest rates and future values is complex (as detailed in the next section). The striking thing is that consumers give answers that are *biased*: they almost always underestimate future values, and almost always underestimate loan interest rates. We now ask whether a common cognitive underpinning can explain not only payment/interest bias and future value bias, but also the relationship between each bias and the time horizon being considered.

---

[20] We discard installment loan responses from 1977 that imply negative interest rates; in all likelihood these are loans with balloon payments, which are not recorded. We also discard mortgages from 1977, because that survey does not identify the size of escrow payments for taxes and insurance in each household's mortgage payment, making calculation of the implied APR impossible.

III. Explaining Payment/Interest and Future Value Bias: Exponential Growth Bias and Other Possibilities

Here we consider several explanations for the pattern of payment/interest and future value bias documented above. In particular we show that exponential growth bias ("EG bias") provides a parsimonious explanation. EG bias is the tendency of individuals to systematically and dramatically underestimate the growth or decline of exponential series when asked to make intuitive assessments (without calculators).[21] Thirty years of research in cognitive psychology establishes that EG bias appears robustly across elicitation methods and contexts (see Web Appendix A for a review).

A. Exponential Growth Bias and Future Value Bias

It is intuitive that someone who underestimates exponential growth will display future value bias. Consider a consumer who saves a present value *PV* at a periodic interest rate *i* over time horizon *t*, with periodic compounding. The future value is:

$$FV = PV(1+i)^t \quad (1)$$

The term $f(i,t) = (1+i)^t$ is an exponential function, and an individual with EG bias will underestimate $(1+i)^t$. Because the future value is just a multiple of that term, there is a straightforward link between EG bias and future value bias. Even a mild degree of EG bias can lead to substantial future value bias; consider a consumer with the following form of EG bias:

$$f(i,t,\theta) = (1+i)^{(1-\theta)t} \quad (2)$$

The θ term parameterizes bias: Unbiased consumers have $\theta = 0$ and correctly perceive exponential growth, while those with $0 < \theta < 1$ have EG bias.[22] Figure 3 shows how an EG biased-consumer would perceive future values over different time horizons $t = [1,5,30]$, with bias on the interval $\theta \in [0,0.15]$.[23] Figure 3 uses i =7%, a benchmark return on equities. Perceived future values are calculated using:

[21] We focus on exponentiation rather than on the other mathematical operations involved in borrowing and savings calculations because there is little evidence of biases in basic arithmetic. For reviews of related evidence see Campbell and Xue (2001) and DeStefano and LeFevre (2004)

[22] Most research in psychology estimating EG bias uses this functional form, because it fits experimental data reasonably well using only one free parameter. See Appendix A for a discussion of more flexible approaches.

[23] This range of parameterized EG bias is actually small relative to that estimated by Eisenstein and Hoch for savings. Eisenstein and Hoch fit the slightly more flexible function $f(i,t,\alpha,\beta) = \alpha(1+i)^{\beta t}$ and estimate $(\hat{\alpha} \cong 0.45, \hat{\beta} \cong 0.50)$. We use smaller values for EG bias because they fit our loan data better. The median θ

$$FV = PV \cdot f(i,t,\theta) \quad (3)$$

The calculations use annual compounding and PVs that equalize the FV when $\theta = 0$, to facilitate comparison of perceived FVs as EG bias changes. Figure 3 illustrates that EG bias is essentially irrelevant over a one year horizon, and has large effects over a retirement planning (30-year) horizon. We show the effects for a single interest rate to conserve space, but it is evident from (1) and (2) above that the level effects of bias are increasing in the interest rate.

Another parameterization of EG bias is "linear bias," which is a useful benchmark because it describes a complete failure to account for compounding. The mathematical form for linear bias is $f(i,t) = 1 + it$, meaning that the perceived future value is linear in *t*. In lab experiments measuring EG bias, perceived future values are often closer to those implied by linear bias than to the true value.

B. Exponential Growth Bias and Payment/Interest Bias

Interest rate formulas also contain exponential functions. The formula relating a periodic interest rate *i* to a loan amount *L*, maturity *t*, and periodic payment *m* is:

$$m = Li + \frac{Li}{(1+i)^t - 1} \quad (4)$$

This equality contains the same exponential term that appears in the future value formula: $f(i,t) = (1+i)^t$.[24] There is no closed-form solution for the periodic rate; it is defined implicitly. If the period is one month, the annual percentage rate (APR) on the loan is equal to $12i$.[25]

Although the math is considerably more difficult than for future values, one can also show that EG bias produces payment/interest bias.[26] Web Appendix B presents a formal treatment of the issue, proving that EG bias produces payment/interest bias and showing conditions under which bias is greater for short-term loans.

Despite the subtlety of the math involved, the intuition for this result is straightforward. Payment/interest bias is a consequence of failing to account for *declining principal balances* on

---

implied by the hypothetical loan questions used in Figures 1a and 1b is 0.2, and the interquartile range is [0.14, 0.33]. The values implied by the actual loan questions from 1977 are smaller on average.

[24] There are many ways to write the formula in (4); we choose this one to make the link between the saving and borrowing calculations as clear as possible.

[25] The APR is not continuously compounded. The continuously compounded rate, $(1+i)^t$, is known as the Effective Annual Rate (EAR). It is not a widely used measure of borrowing costs.

[26] We discuss whether this accurately describes the inferences consumers actually make in the next section (or whether, for example, Truth in Lending forcing lenders to disclose APRs renders such inference unnecessary). Here our focus is simply on asking whether EG bias can explain payment/interest bias in the context of the questions in Section II.

installment loans. The most common incorrect answer on the hypothetical questions in Section IIC is the add-on interest rate, which represents the true cost of borrowing only if the borrower retains the loan principal for the entire loan term. But installment loans require borrowers to start repaying principal immediately, so given a fixed dollar amount of interest, the true cost of borrowing always exceeds the add-on rate.[27] A consumer who does not think about declining principal balances or underestimates their impact on borrowing costs will have the payment/interest bias we document in Sections IIB and IIC.

The mathematical correspondence between EG bias and failing to account for declining principal balances is best illustrated by the linear bias case. Suppose that instead of using the correct formula in (4) to infer the interest rate, a borrower with linear bias uses $f(i,t) = 1 + it$ to solve for the interest rate:

$$m = Li + \frac{Li}{(1+it)-1} \qquad (5)$$

In that case the closed form solution for the periodic rate is exactly the simple interest rate on the loan:[28]

$$i = \frac{mt - L}{Lt} \qquad (6)$$

Thus, having a form of EG bias that completely fails to account for compounding is mathematically equivalent to the intuitive effect of completely failing to account for declining principal balances.

EG bias can also produce the (perhaps less intuitive) result that payment/interest bias is more severe on short-term loans. This comes from the fact that principal balances decline less quickly as maturity increases. Consider the limiting case. As the maturity on the loan approaches infinity, the last term in equation (4) disappears and the formula becomes $m = Li$; the periodic payment equals the principal times the periodic rate. Because the exponential term disappears, even someone with severe EG bias will correctly infer the rate from a principal and payment. Put

---

[27] An alternative view of the intuition is that it reduces the effective loan principal on a short-term loan from *L* to something closer to *L/2*. That approximation generates a rule-of-thumb for short-term installment loan APRs, which is that they are roughly double the add-on rate. Only one percent of our sample supplies a perceived APR that is consistent with the use of this heuristic.

[28] Recall that this formula is for the periodic rate, and that the APR on a *t*-month loan is *it*. So, for a 12-month loan of $1000 with monthly payments of $100 and total payments of $1200 over the year, the formula yields a periodic rate of 1.67%, and a (misperceived) APR of 12*1.67%=20% , which is the simple interest rate.

another way, on an infinite-maturity (interest-only) loan, there is no declining principal balance to complicate inference about the interest rate.[29]

Numerical examples also illustrate these ideas, and show that even mild EG bias generates substantial payment/interest bias. Figure 4 compares the actual to perceived interest rates on 12-, 48- and 360-month installment loans where $t = [12,48,360]$, and $\theta \in [0,0.15]$. All calculations use an actual APR of 35% (to fit the modal rate implied by the questions we use to measure bias) and the same functional form for EG bias as in Figure 3, meaning that the perceived rate solves:[30]

$$m = Li + \frac{Li}{(1+i)^{(1-\theta)t} - 1} \tag{7}$$

Even relatively low levels of EG bias (i.e., of θ) lead to substantially lower perceived interest rates, and to payment/interest bias that is greater on the short-term loans. It is essentially irrelevant on the 30-year loan.

A final point to highlight is that EG bias can produce biased perceptions of borrowing cost and saving returns either directly or indirectly. The effect is direct if consumers actually try to (intuitively) solve the problems described above. The effect is indirect if EG bias leads consumers to adopt biased heuristics like linearizing yields or ignoring declining principal balances.

C. Other Explanations for Payment/Interest Bias and the Bias/Maturity Pattern

EG bias is an appealing explanation for the biases documented in Section II because it provides a simple and coherent explanation for the entire pattern of biases. It also has broad experimental support in other settings. But it is possible that there are other reasons for either the existence of payment/interest bias, or the bias/maturity relationship. We now discuss those other explanations, noting which can be dismissed and which cannot.

One set of explanations is for the existence of payment/interest bias (or future value bias) overall. At first blush one might think that payment/interest bias simply reflects math mistakes or uninformed guesses. It is certainly true that calculating an APR (or a future value) from other

[29] There is a limiting argument in the other direction as well, though it is looser. Suppose a consumer underestimates the exponential term in the denominator of equation (4). As maturity t falls the denominator approaches zero, increasing the value of the second term and requiring a lower perceived rate to make the equality hold (given a fixed loan principal and monthly payment). The statement is a bit imprecise because i itself appears in the exponential growth term, which motivates the more careful analysis in Web Appendix B, but the general thrust of the argument turns out to be correct.

[30] The functional form in our numerical examples has the advantage of simplicity, but can yield perceived rates that are zero or even negative if θ is large enough given the other parameters. We view that functional form as a useful approximation within the range of data that we observe, rather than a form that accurately models the effects of EG bias across a wide range of settings. Reassuringly, Appendix B shows that EG bias generates payment/interest bias under very general conditions on the form of bias.

information is extremely difficult, and the fact that people make mistakes is unsurprising. What is surprising is that *when people make mistakes, they err in a particular direction*. Most evidence suggests that even on difficult questions the "wisdom of crowds" centers the distribution of answers on the truth (Surowiecki 2005). Our findings suggest that something systematic moves the distribution of answers away from the truth.

Another version of this explanation is that something mechanical about the survey questions leads people to supply perceived rates below actual rates. But bias is evident on actual as well as hypothetical loans, and is robust to different frames. Moreover, a framing effect would imply that bias is spurious and should be uncorrelated with real-world outcomes; we test and reject that hypothesis below.

A more subtle mechanical influence might work through time-varying market rates. Say the SCF hypothetical elicits current market repayment totals (on average at least) and respondent perceptions of "normal" (rather than current) market rates. Then our measure will mechanically produce greater *levels* of "bias" during a period when there are high market rates, as was the case in 1983. But this will only induce an empirical relationship between bias and household finance if the propensity to mismatch current and normal rates is correlated with something else that drives decisions. If on the other hand the propensity to mismatch rates is uncorrelated with financial decisions, then we should see no relationship between what we call payment/interest bias and any outcome of interest.

In any case the perception data from 1977 belie this concern, because during that period rates were both stable and typical by historical standards.[31] Table 2 shows that hypothetical loan bias is not much different in 1977 than in 1983. Moreover bias is also prevalent on *actual* short-term loans in 1977.

A second set of alternative explanations concerns the bias/maturity relationship. In this case there are plausible, complementary explanations for one fact: that consumers correctly assess the interest rates on their long-term loans. One such explanation is that consumers learn and remember their mortgage rates because the stakes are high. Another explanation is that enforcement of the APR disclosure mandated by the Truth-in-Lending-Act (TILA) is effective for mortgages. Indeed our companion paper provides evidence consistent with this hypothesis; TILA has more bite for banks than nonbanks, and banks dominated the mortgage market during our sample period. So, both of these things could explain why consumers have accurate (and precise)

[31] The 10-year T-Bill rate was 7.21% on January 1, 1977. It averaged 7.24% and 6.71% during the previous 5- and 10-year periods.

knowledge of interest rates on long-term loans. Neither explains payment/interest bias on short-term loans, however.

In sum, EG bias or one of its intuitive analogues provides a coherent explanation for the pattern of misperceptions documented in Section II. No alternative explanation that we know of explains the entire pattern. We consider the possibility that our measure of bias is correlated with other factors that drive financial decisions in Section VI.

D. Could Bias Matter in the Market?

A final issue of interpretation is whether payment/interest bias or future value bias could matter in the market. As noted in the introduction there are several reasons to believe that bias might not be completely neutralized.

On the consumer side, the mitigating effects of learning, heuristics (including ignoring interest rates), and decision aids may be incomplete in a relatively abstract, low-feedback domain like household finance. On the supply side, incomplete Truth-in-Lending enforcement may allow many lenders to continue exploiting bias by shrouding interest rates, and SEC rules discourage the type of advertising that may be needed to de-bias consumers (namely, featuring expected future value earned over long investment horizons).

In the next two sections we detail our tests of the hypothesis that bias is neutralized or otherwise irrelevant in the market. We also test whether supply- or demand-side forces such as delegation and credit constrains counteract the effects of bias.

IV. Exponential Growth Bias and Household Finance: Empirical Strategy

We outline our empirical strategy in this section, and then detail our tests and report the main results in Section V.

A. Payment/Interest Bias and Household Finance in the Cross-section

Our empirical models estimate whether household-level differences in bias explain the cross-section of outcomes in household finance. We lack a household-level measure of future value bias, and therefore focus on our measure of payment/interest bias as the key explanatory variable. We use specifications of the form:

$$Outcome_h = f(Bias_h, Addon_h, X_h) \qquad (8)$$

Outcome varies in each model; we discuss the outcome measures below. *Bias* is a vector of payment/interest bias quintile indicators and a separate category for non-response to the hypothetical loan questions.[32] The latter allows for the possibility that non-responders are more biased than the least-biased (i.e., the quintile 1) households. *Addon* is a dummy variable equal to one if the perceived rate (used in the payment/interest bias measure) equals the add-on rate; including this as a covariate allows for the possibility that households providing the add-on rate interpret the payment/interest bias questions and/or behave differently, conditional on the degree of bias. *X* is a vector containing all of the controls detailed below.

Intuitively we would expect that payment/interest bias should increase short-term but not long-term borrowing, decrease savings rates, and decrease net worth.[33] If payment/interest bias is also a useful measure of future value bias (which would be true if both are driven by a common exponential growth bias), then *Bias* should be correlated with asset allocation as well as debt allocation, and exert stronger effects over the long term and for high-return assets. We therefore test whether bias decreases long-term (and high-yielding) but not short-term (and low-yielding) investing,

The null hypothesis in each case is that the *Bias* coefficients are zero (either because our bias measure is uninformative, or because bias is neutralized due to consumer or supplier adaptation). The least-biased households (quintile 1) serve as the omitted category.

B. Data, Outcomes and Control Variables

In the empirical work below we use data from the 1983 SCF rather than the 1977 SCF, because the 1983 survey covers the household balance sheet more comprehensively. The top panel of Table 3 shows our outcome measures, stratified by bias category. We discuss the outcomes in greater detail below, but they are: short-term debt/income, long-term debt/income, stock holdings as a share of total assets, CD holdings as a share of total assets, savings rates, and wealth accumulation. Not surprisingly, there are strong unconditional relationships between bias and all of the outcomes of interest. That highlights the need to control for other influences on household finance.

An advantage of the SCF is that the set of possible controls is extensive. Because minimizing omitted variable bias is critical, we take an approach that seeks to control for all factors that might

[32] We have used other functional forms for bias (a biased/unbiased indicator, the logarithm of bias, and a quadratic in the level of bias). The results are robust to these other functional forms, as well as to orthogonalizing bias by regress ln(bias) on the full set of controls, and then using the deviation of actual bias from its fitted value to reconstruct bias quintiles. See Web Appendix Table 1 for some results.

[33] There may be income or other effects as well, particularly if payment/interest bias is generated by exponential growth bias; e.g., EG bias may lead to lesser discounting of a given expected future income stream and thereby make consumers feel wealthier. In that case we view our empirical work as identifying the net effect of biased perceptions on our outcomes.

be correlated with both outcomes and our measure of payment/interest bias, hence erring on the side of "over-controlling."[34] Our controls include measures of preferences, expectations, available resources (including income, defined-benefit retirement wealth and credit constraints), claims on resources (including life-cycle factors), and problem-solving approaches (and financial sophistication more generally). We group them below for expositional purposes but emphasize that each of our empirical models includes *all* of the variables described below (and detailed completely in Web Appendix C).[35] Table 3 shows descriptive statistics on some key variables, by bias category. Table 4 (discussed in Section IV.C below) estimates multivariate correlations between bias and the control variables.

<u>Controls for available resources and claims on resources</u> include: total household labor income (dummies for the percentile), homeownership, pension coverage, pension wealth and Social Security wealth (we exclude this wealth from our left-hand side wealth measure since it was plausibly beyond the direct control of most households in 1983), number of members in the household, gender, education, race, age, marital status, health status, years with current employer, industry, and occupation (including business ownership or self-employment activity). We also observe two measures of credit constraints: whether a household has been denied credit or discouraged from applying in the "past few years", and whether it has a credit card.

<u>Controls for expectations</u> about lifetime wealth include measures of expected inheritance, expected tenure with current employer, and expected retirement age.

<u>Controls for preferences</u> include measures of risk preference, liquidity preference, and debt aversion; other work has shown that these are important determinants of household financial decisions. Risk preference is measured with the question: "Which of the following statements on this card comes closest to the amount of financial risk you are willing to take when you save or make investments?" Answers fall into four categories, ranging from "willing to take substantial financial risks to earn substantial returns" to "not willing to take any financial risks." Time preference or patience is measured with the question: "Which of the following statements on this card comes closest to how you feel about tying up your money in investments for long periods of time?" Answers range from "will tie up money in the long run to earn substantial returns" to "will not tie up money at all." Debt aversion (and perhaps an element of time preference) is measured with the question: "Do you think it is a good idea or a bad idea for people to buy things on the

[34] If there is a causal link between bias and any of these variables, we may underestimate the relationship between bias and our outcomes of interest (Angrist and Krueger 1999).
[35] See http://www.dartmouth.edu/~jzinman/Papers/Stango&Zinman_EG_Bias_HH_Finance_WebAppendix.pdf .

installment plan?" We consider the implications of possible omitted "behavioral" biases in preferences and expectations in Section VIB.

Controls for problem-solving approaches/overall financial sophistication include whether the respondent evaluates loan offers by focusing on APRs or other terms (e.g., monthly payment, available loan amount, down payment, collateral requirement). Focusing on payments or other terms may reflect a lack of financial sophistication, conditional on credit constraints.[36] Two other proxies for financial sophistication are ATM use (only 17% of our sample use ATMs at all), and of course education. We consider the possibility that elements of financial sophistication remain unmeasured, and the relationship between bias and financial sophistication more broadly, in Section VIA.

Controls for financial advice are categorical variables measuring whether households use external advice, and whether advice is from a professional, from friends and family, or from other sources.

C. Who is Biased? Multivariate Relationships

Table 4 sheds light on the conditional relationships between household characteristics and bias by presenting results from models with a measure of bias on the left-hand side, and sets of controls on the right-hand side. We use two functional forms for bias: bias quintile, and ln(Bias); the latter coefficients are easier to interpret. For each dependent variable, we present a parsimonious specification including a subset of important RHS variables, and a full specification with all of the controls listed above. One difference between these specifications and those we use below is that we use wage income quintile as a control here; below we use wage income percentile, which is much more flexible.

Not surprisingly, income and education are highly correlated with bias. These relationships hold up even conditional on all of our other controls. Gender and race are correlated with bias as well. Our preference variables, somewhat surprisingly, are not correlated with bias once we control for income and education; we consider possible omitted features of preferences in VIB.

A final point is that the fit is quite low. In the parsimonious specifications, 12-13% of the cross-sectional variation in bias is explained by other household characteristics. Even in the full models, household characteristics explain only 15-17% of the variation in bias. That is surprising given the sheer number of controls and the highly flexible functional form for many of them (the

[36] Focusing on non-interest terms may be rational for those who face binding liquidity constraints (Karlan and Zinman 2008; Adams, Einav and Levin forthcoming; Attanasio, Goldberg and Kyriazidou forthcoming).

total number of categorical variables is over 200). Bias is not easily explained by other household characteristics.

V. Results

This section reports results of our primary empirical tests from the multivariate model (8). Each multivariate specification conditions on the full set of control variables listed in Section IVB; some contain additional controls noted below. The tables suppress most of the control variable coefficients (including those on wage income percentile) to save space.

A. Liability Composition and Bias

Table 5 presents results of our test of whether payment/interest bias encourages short-term borrowing (by making it appear relatively cheap) but not long-term borrowing (since even individuals with severe payment/interest bias should accurately assess long-term interest rates).

Column 1 presents probit marginal effects from a model where the dependent variable equals one if the household used short-term installment debt to finance a recent large purchase (car, household item, or home improvement).[37] This model also controls for characteristics of the recent purchase (month/year, product purchased, and product price). The coefficient on each of the bias quintiles is positive, and households in quintiles 4, 5, and unknown are significantly more likely than the least-biased households (in the omitted quintile 1) to have used short-term debt for the purchase. Each of the bias coefficients implies economically large increases (11% to 43% of the sample mean). Columns 2 and 3 test whether more biased households have higher short-term debt-to-income ratios, conditional on having nonzero short-term debt.[38] Column 2 includes the entire sample of short-term borrowers, while Column 3 drops households that face relatively severe credit constraints (measured by recent credit denial or lacking a credit card). Overall, the results show a clear positive relationship between short-term borrowing and bias: all of the 10 bias coefficients are positive. The results are stronger statistically when constrained households are dropped from the sample, presumably because rationing prevents some households from borrowing as much as they would like. The point estimates suggest that bias increases the short-

[37] This result is conditional on having made a large purchase; we find no significant relationship between bias and the probability of purchase.

[38] We find no relationship between bias and having nonzero short-term debt. It may be that the extensive margin is not very elastic – short term installment debt is used primarily for financing vehicles and consumer durables, and the near-absence of second mortgage markets in 1983, along with small credit card credit lines, implies that savings was the main outside option for these types of purchase (responses to SCF question b5606, on the financing method for a large recent purchase, confirm this).

term debt-to-income ratio by from 23% to 54%. Households with a perceived rate that equals the add-on rate also carry significantly more short-term debt, conditional on the bias quintile.

Columns 4-5 show that payment/interest bias is not significantly correlated with greater long-term debt (mostly mortgages and lines of credit). Here the long-term debt-to-income ratio is the outcome of interest. Note that our usual set of control variables includes homeownership status. The sign pattern on the bias coefficients is less pronounced, and none of the bias quintile or add-on coefficients is significant, although the estimates are imprecise.

In sum, the evidence is consistent with exponential growth bias affecting borrowing choices via its effect on payment/interest bias.

B. Asset Composition and Bias

Table 6 examines the relationship between payment/interest bias and the composition of assets. If payment/interest bias is confined to the borrowing side, then asset choices should be unaffected by bias. If on the other hand payment/interest bias is a useful proxy for future value bias, then it should be negatively related to holdings of long-horizon, high-yielding assets. We proxy for long-term savings using the value of stock holdings (trading costs were high and mutual fund penetration was low in 1983), and for short-term savings using the value of CD holdings.[39] For each outcome we measure holdings as a share of total assets, and also as a share of financial assets (which exclude housing assets). These models condition on net worth decile in addition to our usual set of controls.

Columns 1 and 2 show that bias is negatively related to proportion of assets held in stocks. All of the bias coefficients are negative, and 8 out of 10 are statistically significant. The coefficients imply that bias induces large decreases in stockholding (of 18% to 55% of the sample means). The controls also enter the empirical relationship in meaningful ways. Higher education (shown in the table) and income (not shown because it is specified as a vector of percentile dummies) are both positively correlated with stock ownership. Our measures of risk aversion and impatience are both negatively correlated with stock ownership.

Columns 3 and 4 report the same share specifications for CDs. The pattern of signs (all positive) and significance (6 out of 10) suggests that biased households substitute short-term for long-term assets.

---

[39] The presumption here is that stocks have a higher historical return, and are typically held over longer time periods. There is a consensus that this is true; see Welch (2000; 2008) for discussions of consensus estimates of the equity premium. Of course, the early 1980s were a time of atypically high nominal interest rates, and if people use nominal returns to allocate assets that might weaken the relationship between EG bias and stock or CD holdings.

In sum, the results are consistent with exponential growth bias affecting asset composition via future value bias.

C. Saving, Wealth and Bias

Table 7 shows results for specifications correlating bias with savings rates and wealth accumulation. For savings we estimate an ordered probit where the categories are ranked (1=dissaved, 2=even, 3=saved), based on a question about household spending vs. income in the previous year. For wealth the dependent variable is the log of net worth.[40] As with the borrowing models, we also report estimates after dropping the most credit constrained households.

Columns 1 and 2 suggest that bias decreases the likelihood of saving. All of the 10 coefficients are negative, and 6 of the 10 are statistically significant. In the less constrained sample the results are stronger statistically and larger economically. The coefficients imply a large reduction in saving relative to the least-biased households.

We see a similar pattern of results with ln(wealth) as the dependent variable (Columns 3 and 4). All 10 bias coefficients are negative, and 8 of them are significant. The point estimates imply that more-biased households hold 7% to 33% less wealth than the least-biased households, conditional on our full set of controls. The controls matter in intuitive ways. Education and income are both very strongly and positively correlated with wealth. Risk aversion and impatience are both very strongly and negatively correlated with wealth.

In sum, the results on savings and net worth are consistent with exponential growth bias affecting decisions through payment/interest bias, future value bias, or both.

D. Bias and Delegation: The Role of External Advice

We examine the use and role of external financial advice in Table 8. While the specifications in Tables 4-7 include the use of advice as a control, Column 1 of Table 8 tests whether more biased households are more likely to get outside advice. The dependent variable is equal to one if the household obtains any external advice for financial decisions. The results show a strong positive relationship between bias and getting advice, and household in bias quintiles 3-5 and the nonresponse category are significantly more likely to get advice. The point estimates imply 17% to 22% increases relative to the sample mean. This result suggests that at least some households are aware of their bias and seek help.

---

[40] This drops the 7% of our sample with zero or negative wealth. We do not find any significant correlations between bias and having non-positive wealth. We have also used net worth percentile as the dependent variable and obtain nearly identical results.

The next three columns estimate the relationship between using advice and the bias/wealth relationship. Column 2 shows the coefficients on advice when those variables are used as controls. There is a strong and positive relationship between advice and wealth, conditional on bias and everything else. Columns 3 and 4 split the sample by advice and test whether advice is related to the bias/wealth relationship. The results are striking. Among households that do not use advice the bias/wealth relationship is negative, large and highly significant. In contrast, there is no significant relationship (or clear sign pattern) between bias and wealth for households that do use advice. This suggests that there is important heterogeneity in bias awareness or other mitigating strategies *within* bias quintiles; i.e., for a given degree of bias and rich set of household characteristics, some households effectively neutralize their bias (by delegating), while others do not.

Overall, the results in this section suggest that exponential growth bias matters in the market: our measure of payment/interest bias is correlated with portfolio allocation, savings and wealth in a way consistent with both a direct effect of payment/interest bias, and a correlation between payment/interest bias and future value bias. The market matters too: the effects of bias are mitigated by credit constraints on the supply side and delegation on the demand side.

VI. Interpreting the Results: Alternative and Complementary Explanations

The pattern of results is consistent with exponential growth bias affecting financial decisions in specific ways. The rich set of controls helps to rule out many sources of omitted variable bias. Payment/interest bias may nonetheless be correlated with some omitted variable that explains (or at least contributes to) the results.

A. Bias & Financial Sophistication

One possibility is that payment/interest bias is correlated with financial sophistication more broadly. There is growing evidence that financial sophistication is key to understanding equilibrium household finance (Campbell 2006). In many analyses the empirical role of financial sophistication is studied indirectly, through proxies like education, income (Calvet, Campbell and Sodini 2007), and ATM use (Attanasio, Guiso and Jappelli 2002). As noted above, we include controls for such proxies in our analysis. But at least one study has measured financial