Network for Studies on Pensions, Aging and Retirement

Netspar DISCUSSION PAPERS

Johannes Binswanger and Katherine Carman

The Miracle of Compound Interest

Does our Intuition Fail?

Discussion Paper 12/2010-079

The Miracle of Compound Interest: Does our Intuition Fail?*

Johannes Binswanger†

Katherine Grace Carman‡

December 10, 2010

Abstract

When it comes to estimating the benefits of long-term savings, many people rely on their intuition. Focusing on the domain of retirement savings, we use a randomized experiment to explore people's intuition about how money accumulates over time. We ask half of our sample to estimate future consumption given savings (the forward perspective). The other half of the sample is asked to estimate savings given future consumption (the backward perspective). From an economic point of view, both subsamples are asked identical questions. However, we discover a large "direction bias": the perceived benefits of long-term savings are substantially higher when individuals adopt a backward perspective. Our findings have important implications for economic modeling, in general, and for structuring advice and financial literacy programs, in particular.

Key words: Behavioral economics, financial intuition, financial literacy, compound interest, retirement saving.

JEL classification: D91, H55.

*We are very grateful to Doug Bernheim, Zvi Bodie, Xavier Gabaix, David Laibson, Arie Kapteyn, Miles Kimball, Peter Kooreman, Olivia Mitchell, Jan Potters, Josh Rauh, Arno Riedl, Arthur van Soest, as well as to seminar participants at the Boston Fed, Netspar, the SAVE Workshop, and the SITE Workshop on Psychology and Economics 8.0 for very helpful comments. We would like to thank Tim Colvin at the RAND institution for programming our questionnaire. Financial support from Netspar and the RAND institution is gratefully acknowledged. All errors are our own.

†Department of Economics, CentER and Netspar, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; J.Binswanger@uvt.nl.

‡Department of Economics, CentER and Netspar, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; K.G.Carman@uvt.nl.

1 Introduction

Some claim that Albert Einstein once said that compound interest was the most powerful force in the universe. While the attribution of this quote is suspect, the significance is clear. When you save, your money bears interest. If you save over a long horizon, the interest itself bears interest. As any economist (or Nobel prize winning physicist) knows, this effect gives rise to the miracle of compound interest: money accumulates exponentially, and the effect is more pronounced over a longer time horizon. Compound interest is particularly important when it comes to deciding whether to postpone saving; if individuals delay, they may forego this powerful force. Whether people benefit from this miracle when saving for their retirement depends on their perceptions of its magnitude.

As traditional defined benefit pensions become less common, people are becoming more responsible for their own retirement preparation. Thus, many households must decide for themselves whether they are saving the right amount or whether they should potentially save more. Tackling this question can be quite difficult (Skinner, 2007). It requires quantifying how an incremental amount of regular savings in working life, combined with the power of compound interest, would translate into additional consumption in retirement.

Individuals may address this problem in several ways: they can calculate, seek advice, or rely on intuition. Calculating the effect of current savings on future spending power requires an assumption about the expected returns on one's savings. In addition, it requires the ability to actually perform the implied compound interest calculations. In previous work (Binswanger and Carman, 2010a), we find that only roughly a third of individuals in our data do this. Previous findings in the literature suggest that many people have low financial literacy as well as low numeracy levels (Lusardi, 2008a, 2008b; Lusardi et al., 2009). Therefore, it is conceivable that many people will not be able to perform these calculations and will take a different approach, instead.

Second, people may seek advice from a specialist, a financial website, or financial planning software. However, about half of our sample indicate that they neither work out a plan themselves nor seek advice from an expert, website, or software. The fact that relatively few individuals seek advice is also highlighted by Lusardi (2008b). Furthermore, the quality of advice varies and there may be important agency problems involved (Mullainathan et al., 2010). In addition, those who seek advice may not be the people who most need help; our data indicate that people with high income, high education, and high financial literacy are significantly more likely to seek help for their retirement preparation (Binswanger and Carman, 2010b).

In fact, many people are likely to rely, in one way or the other, on their own intuition when it comes to quantifying how an incremental amount of savings translates into additional spending power in retirement.[1] In this paper, we therefore investigate the adequacy of people's intuition about the long-run returns to saving. We do so by administering a randomized experiment with the American Life Panel, an online Internet panel hosted at the RAND institution. The people in our sample are mostly above the age of 30 and have above average income. They are thus an ideal sample since they are likely to face retirement planning decisions of the sort that underlies our experimental question.

In our experiment, we ask individuals about their perceptions of the benefits of a marginal change in savings. We provide half of our sample with an amount of (regular) incremental savings. We then ask them what they think the resulting incremental spending power in retirement would be. For the other half of the sample, we switch perspectives. We provide an amount of incremental future spending power and ask them to estimate the amount of incremental savings required.

Individuals are assigned randomly to either treatment. In the first treatment, individuals adopt a *forward perspective* in that they start in the present with incremental savings and think about the resulting future spending power. In the second treatment,

[1] That people use their intuition for economic decision making is fully consistent with standard economic methodology. Milton Friedman (1953) famously made the point that a billiard player's skills derive from having a good intuition, not from mastering the differential equations describing the movements of the balls.

individuals adopt a *backward perspective.* They start with amounts of future spending and think about what savings are required now in order to finance this spending. From an economic perspective, respondents in both treatments face an identical question.

Perceptions of the trade-off between increased savings while working and increased consumption during retirement are likely to be important for several reasons. First, economic models of optimal long-term savings decisions typically consider this trade-off explicitly. The appropriateness of these models depends on the information and assumptions of decision makers. Measuring the perceptions of this trade-off contributes to evaluation of these models. Second, perceptions of the possible returns to saving are likely to influence the choice to save. Individuals who believe that returns are very high may save too little or postpone savings under the assumption that saving in the future will be sufficient. Alternatively, they may save too much because they perceive extraordinarily high benefits. Individuals who believe that returns are very low may reduce savings assuming that saving is not worth it, or may save too much to compensate for low perceived rates of return.

Our question focuses on individuals' *intuition* of the returns to saving, or their gut feeling about the rates of return. However, it is possible that intuition plays little role in determining actual behavior because people can turn to advisory services (such as financial planners) or calculate optimal savings rates. In past research, we have shown that calculations of optimal savings rates are uncommon (Binswanger and Carman, 2010a) and highly correlated with the use of advice (Binswanger and Carman, 2010b). Intuition may even affect the decision to make calculations or seek advice. Without calculations or advice, individuals are left with nothing but their intuition.

Our financial intuition question asks respondents to base their answer on their own personal situation. In particular, they should answer the question for their own age and expected retirement age. Furthermore, the incremental amounts of saving and spending in the question are determined by the respondents' income and set equal to between

5 and 10 percent of respondents' monthly income (depending on rounding). Overall, respondents thus face a situation that reflects their own circumstances when thinking about the benefits of incremental savings for themselves.

Our primary analysis focuses on whether the treatment a respondent is assigned to has a systematic effect on her answer. We hypothesize that, if individuals have a reliable intuition about the benefits of long-term savings, then we should not expect any systematic difference between answers across treatments. What we find, however, is a systematic and large difference. In particular, the perceived benefits of long-term savings are substantially higher when adopting the backward perspective. Based on median answers, this difference is equivalent to 3.5 percentage points of return per year. Over a horizon of 20 years, this amounts to a difference of more than 100 percent. Alternatively, this is equivalent to a difference in implied life expectancy of 23 years across treatments. By any measure, the treatment effect is very large.

Since assignment to a treatment is random, this effect cannot be explained by differences in age, perceived or experienced rates of returns, life expectancy, or any other fundamental factor. Our conclusion is, rather, that people are susceptible to a systematic framing effect that is triggered by either adopting a forward or a backward perspective on the evaluation of the benefits of long-term savings.[2] This effect has not, to our knowledge, been previously documented in the literature.

We investigate several potential explanations for this framing effect, such as anchoring and wording. Two possible explanations are compatible with our findings. First, individuals exposed to the framing of the backward treatment may, wrongly, take the amount of future spending that we provide as an upper bound for the amount of savings that are necessary for financing it. Second, the framing effect may be explained by a type of loss aversion that we describe in Section 5.2. Both effects may be intrinsically linked to how people think about the benefits of incremental savings. It is therefore important that

[2]See Della Vigna (2009) for a discussion of framing effects.

economists are aware of these effects.

We also analyze the variation of respondents' answers within a treatment. We conclude that it is too large to be justified by fundamental factors. Overall, we conclude that individuals' intuitive perceptions of the benefits of long-term savings are very unreliable. Given that many people are actually likely to rely, in one way or another, on their intuition, this finding has important implications for theory and practice.

Our paper is structured as follows. Section 2 discusses the related literature. Section 3 explains the data we use. Section 4 explains the mathematics behind our financial intuition question. Section 5 contains the results. Finally, Section 6 concludes.

2 Literature

Our paper is closely related to the literature on financial literacy. Lusardi and Mitchell (2007a, 2007b) investigate people's understanding of compound interest by asking members of the Health and Retirement Study (HRS) and the American Life Panel how money in an account accumulates over two consecutive years at a given interest rate. Knowing the answer to this question undoubtedly is an integral part of financial literacy. However, the perspective of this question is different from the one adopted in our paper, which considers the effect of an incremental amount of regular savings on future spending power. Financial literacy questions typically have a right and wrong answer, whereas our question measures perceptions. Our study is thus complementary to this work of Lusardi and Mitchell.

In a seminal study, Stango and Zinman (2009) document an exponential growth bias, i.e. a tendency to systematically underestimate the effect of compound interest. They show that this leads individuals to underestimate the future value of money when saving, and the costs of (short-term) loans when borrowing. Their study is based on the 1977 and 1983 waves of the Survey of Consumer Finances (SCF). They identify individual values of exponential growth bias by comparing two questions available for these SCF waves. The first asks about the future value of installment payments of a loan. In the second question, individuals are asked about the interest rate that would be equivalent to these payments. The size of the bias corresponds to the degree that both answers are inconsistent with each other.

Our analysis differs from Stango and Zinman's in several respects. First, our focus is on retirement savings, and therefore on long term effects of compounding. Second, individuals in our sample are likely to have had more exposure to estimating the benefits to long-term savings since defined contribution pensions and 401(k) plans have become much more prevalent since the 1980s. Thus, it is interesting to see whether the people in our sample have a reasonable intuition about the magnitude of the returns to long-term savings. Third, Stango and Zinman investigate whether people can correctly translate an interest rate into a stream of installment payments. In contrast, literally speaking, there is no right or wrong answer to our experimental question. Rather, we focus on individuals' intuition about the long-term returns to saving. Finally, the main conclusions of our analysis are based on an experimental manipulation of the perspective according to which individuals are made to think about the returns of long-term savings, either forward or backward.

Van Rooij et al. (2007) investigate the relationship between financial literacy and stock holdings. They find that that higher financial literacy is indeed associated with a higher likelihood of stock market participation and they provide evidence that this effect is causal. Lusardi and Mitchell (2008) show that there is a strong relationship between financial literacy and the propensity to plan for retirement among women. Behrman et al. (2010) show that financial literacy has a strong effect on wealth accumulation.

3 The Data

Our data come from a survey module conducted with the RAND American Life Panel (ALP). Our module was fielded in August 2008. The ALP is an internet-based platform for online surveys. The ALP provides internet access to households if needed.[3] It has previously been used for the study of financial literacy by Lusardi and Mitchell (2007b). One advantage of online surveys is that it is easy to tailor our questions to a particular respondent. In particular, our questions include numbers that are based on a respondent's income (see below). Furthermore, for online surveys it is straightforward to randomly assign individuals to different treatments. A final advantage of online surveys is that they avoid interviewer bias (Donkers et al., 2001).

Our sample only includes individuals who indicate that they are not (yet) retired. The ALP provides us with information on respondents' background characteristics such as age, gender, education, etc. Our measure of income is top-coded at $200,000; five percent of observations fall into this category. Table 1 includes information about the main demographic variables. Our sample is somewhat older and better educated than the general population. Thus, it is not fully representative. However, the advantage of this sample is that the vast majority are responsible for saving for their retirement and thus are likely to be personally confronted with the problem of quantifying the benefits of long-term savings. In particular, 90 percent of our sample is older than 30 (not shown in Table 1). Furthermore, 90 percent earn more than $30,000 (before taxes), and more than 90 percent have access to a defined-contribution pension plan, such as a 401(k) plan. This makes our sample particularly suitable for the purpose of our study.

We randomly assign individuals to one of two treatments. Respondent's assigned to our first treatment are asked the following question.

(Forward treatment) *Please don't engage in any type of calculation while*

[3] More information is available at http://www.rand.org/labor/roybalfd/american_life.html.

answering the following question. Just provide your best estimate. And please answer this question as if there were no price increases (i.e. no inflation).

Suppose that you save an extra \$[S_i] every month from now until you retire. About how much extra money would you get to spend each month during retirement?

Your best guess: $ _________

We dub this treatment the *forward* treatment since respondents are asked to start with current savings and then project the resulting future spending power.

For the second treatment, the question is:

> **(Backward treatment)** *Please don't engage in any type of calculation while answering the following question. Just provide your best estimate. And please answer this question as if there were no price increases (i.e. no inflation).*
>
> *Suppose that you need an extra \$[C_i] during retirement during each month. How much extra savings would you have to put aside each month from now until you retire, in order to achieve this?*
>
> *Your best guess: $ _________*

We dub this the *backward* treatment since individuals start with incremental future spending and then think about the implications for saving in the present.

Each respondent participates in only one treatment. Treatments are assigned randomly. The amounts S_i and C_i are roughly equivalent to 5 and 10 percent of total monthly household income, respectively, as detailed in Table 2. The amount C_i in the second treatment is two times the amount that would be shown for the same income range under the first treatment. The particular factor 2 is arbitrary, but results in higher amounts shown in the second treatment in order to roughly account for the fact that money grows over time. It is noteworthy that this cannot induce individuals to think that C_i *should* equal the double of S_i because individuals participate only in one treatment and see only one of those numbers. Information on the distribution of S_i and C_i is summarized in Table 2.

Our questions were answered by 869 individuals. Out of these, 408 fall into the *forward* treatment and 461 into the *backward* treatment. Table 3 provides summary statistics by treatments and shows that despite a slight overassignment to the *backward* treatment, the samples are randomly assigned. There are no significant differences in demographic characteristics across treatments, with the exception of college education where the difference is only significant at the 10-percent level.

In our experimental question, we ask individuals not to make any literal calculations. It is of interest to know whether individuals followed these instructions. We test this using data on the time a respondent spends answering this question. Interestingly, the distribution of response times is very similar for both treatments. In both treatments, the median time spent is 39 seconds and the 75th percentile 60 seconds. The 90th percentile is 92 seconds in the *forward* treatment and 86 seconds in the *backward* treatment. Based on these short response times, we find it plausible that at least three thirds, if not at least 90 percent of respondents did not use any help for answering this question. Importantly, the distribution of time spent is almost identical across treatments.

For our main analysis, we convert individual answers into implied interest rates, as detailed in the next section. For 36 observations, we cannot make this conversion due to inconsistent information.[4] Furthermore, for 4 observations, the implied annual interest rate exceeds 100 percent. These 4 outliers are also removed from the analysis. These outliers are, on average, 50.1 standard deviations above the next highest observation and have a maximum internal rate of return of 3,000 percent.

[4]For 6 observations, the answer to the *backward* treatment question is 0, which would imply an interest rate of infinity. Second, all respondents indicate at the beginning of the survey that they are not retired. However, 22 respondents supply us with an expected retirement age lower than their current age. Finally, for 8 observations, the expected retirement age exceeds our baseline measure of life expectancy (see below).

4 The Mathematics of Long-term Compounding

This section describes the equations to calculate a mathematically accurate answer to both the *forward* and the *backward* treatments. Furthermore, using these formulas we can infer each respondent's perceived rate of return or perceived life expectancy. Let S_i denote an amount of (incremental) monthly savings of a respondent i that are regularly put aside from current age t_i until retirement R_i. Let C_i denote the corresponding (incremental) consumption level that can be financed by these savings from retirement age R_i until death D_i. In the *forward* treatment, respondents are given an amount S_i and are asked to guess C_i. In the *backward* treatment, respondents are given an amount C_i and are asked to guess the corresponding S_i.

Given information about t_i, R_i, D_i, respondent i's' answer can be converted into an interest rate. This interest rate is simply the internal rate of return. This is the rate of interest that would make the respondent's answer exactly correct. Our main analysis is based on a comparison of these internal rates of returns across treatments.[5]

If all monetary amounts are expressed in dollars as of the last year of life, the internal rate of return r_i that is implicit in respondent i's answer is determined as follows.

$$\begin{aligned} S_i\,(1+r_i)^{D_i-t_i} &+ S_i\,(1+r_i)^{D_i-t_i-1} + \ldots + S_i\,(1+r_i)^{D_i-R_i+1} \\ &- C_i\,(1+r_i)^{D_i-R_i} - C_i\,(1+r_i)^{D_i-R_i-1} - \ldots - C_i = 0. \end{aligned} \tag{1}$$

[5]It is conceivable that individuals expect to earn higher returns during the accumulation phase than during the decumulation phase. If this is the case, a single constant rate represents an average of these two rates. We prefer focusing on one interest rate since this is more parsimonious. In fact, there are infinitely many possible pairs (or triples etc.) of interest rates that could be made compatible with a respondent's answer. Furthermore, a major part of our analysis is based on investigating whether the distribution of answers differs systematically across treatments. If there is a systematic effect of the treatment on age-invariant internal rates of returns, then there would also be a treatment effect for a given profile of age-dependent returns. As a result, the assumption of age-independent returns is not critical.

Let α_i denote the ratio C_i/S_i. Then (1) can be restated as

$$\begin{aligned}(1+r_i)^{D_i-t_i} + (1+r_i)^{D_i-t_i-1} + \ldots + (1+r_i)^{D_i-R_i+1} \\ - \alpha_i (1+r_i)^{D_i-R_i} - \alpha_i (1+r_i)^{D_i-R_i-1} - \ldots - \alpha_i = 0.\end{aligned} \tag{2}$$

This is a polynomial of degree $D_i - t_i$ in $1 + r_i$. Equation (2) may be rewritten as

$$\phi(t_i, R_i, D_i, \alpha_i, r_i) = 0. \tag{3}$$

In our sample, we observe t_i, R_i, α_i. We do not observe D_i and r_i. There are two possible strategies that can be pursued in order to analyze our data. First, we may make assumptions about D_i and then solve (2) or (3) for r_i, for each i. Alternatively, we may make an assumption about r_i (e.g. 3 percent) and then calculate the implied life expectancy D_i. In our analysis, we consider both.

For our baseline analysis, we use age-specific projections of the Social Security Administration for D_i and then solve for r_i.[6] For a robustness check, we also set $D_i = 100$ for all i. While life expectancy is certainly below 100 for most respondents in our sample, this may be interpreted as a planning horizon of $100 - t_i$. Individuals who are risk averse may choose a retirement planning horizon based on their perceptions of a maximum possible life expectancy. This allows them to avoid the possibility of running out of assets.

As an alternative approach, we assume a value of either 1, or 3 percent for r_i for all i and then calculate the distribution of life expectancies D_i that would be compatible with the above rates of return.[7]

From equation (2), it becomes apparent that the treatment an individual is assigned to should not influence the implied value of r_i or D_i. Equation (2) depends only on the ratio C_i/S_i, not on C_i and S_i individually. As a result, r_i should be determined only by

[6] We use life expectancies according to the "intermediate" scenario.

[7] We also use a rate of 6 percent that yields very similar conclusions. Most results for the 6-percent rate are not shown in the paper but are available upon request.

the "fundamentals" α_i, t_i, R_i, and D_i (or D_i only by r_i, t_i, R_i, α_i).

Figure 1 sheds some light on the mechanics of equation (2). The horizontal axis measures the ratio $\alpha_i = C_i/S_i$. The vertical axis measures the internal rate of return. The figure is drawn under the assumption that the retirement age is 65 and death age is 90. It shows how the relationship between α_i and r_i depends on current age t_i. The bottom (solid) curve refers to age 25, the middle (dashed) curve refers to age 40 and the top (dotted) curve to age 55. For a 25-year old, a ratio of 2 would imply an internal rate of return of about 1 percent. For a 40-year old, the corresponding rate would be about 3 percent. For a 55-year old, it would be about 12 percent. The figure can also be read in the direction from the y- to the x-axis. Suppose that you expect to earn an interest rate of 3 percent. For a respondent aged 25, α would then be approximately 4. At age 40, the corresponding α would be about 2. At age 55, α should be below one.

Figure 1 shows that internal rates of return can also be negative. This happens if both α_i and current age t_i are low, relative to each other. A negative internal rate of return may easily occur for numbers that may not look particularly "suspicious" to those not familiar with compound interest rate calculations. For instance, consider a respondent with age 50, expected retirement age of 65, and a life expectancy of 85. Suppose that we provide this respondent with a value for S_i of \$400. If the respondent's answer for C_i is \$100, then $\alpha_i = 0.25$ and the implied internal rate of return is minus 6 percent.

5 Results

5.1 Perspective Matters: An Analysis of the Treatment Effect

We now turn to the question of whether the distribution of answers differs systematically across the two treatments of our experiment. If individuals have a reasonable intuition about the benefits of long-term saving, we would not expect their perceptions to be influenced by the framing of the treatment. We examine differences in the ratio C_i/S_i, the implied internal rate of return (under different life expectancy assumptions), and the implied life expectancy (under different rate of return assumptions).

We start with the distribution of the raw ratio C_i/S_i to investigate whether the raw data exhibit a treatment effect. However, it should be noted that it is not straightforward to interpret a particular value of this ratio, e.g. 1.5, without taking into account additional information about a respondent, such as current age, expected retirement age and life expectancy. The first two columns in Table 4 show summary statistics for the ratio C_i/S_i for the *forward* and *backward* treatment, respectively. The median ratio is 1 in the *forward* treatment, and 1.6 in the *backward* treatment. The difference in medians is significant at the 1-percent level (see last row of Table 4). Since assignment to treatments is random, this represents a pure framing effect.

The upward shift is visible at all percentiles, with the exception of the tails of the distribution. In contrast, for the full sample, the mean is larger for the *forward* treatment, due to a few observations with very high values. If we cut 5 percent of observations at both the lower and upper tail of the distribution, separately for each treatment, then the mean is higher in the *backward* treatment (see second-to-last row of the upper panel in Table 4), in line with the median and most other percentiles. The p-value for a t-test of equality of means is 1.3 percent (see second row of lower panel).[8]

The results so far suggest whether a forward or backward perspective is adopted has a strong influence on perceived benefits of long-term savings. While a difference in medians of 1.6 suggests that the treatment effect is large, the meaning of this amount is not clear. It is therefore more appropriate to gauge the magnitude of the effect by looking at internal rates of return, as explained in Section 4 above. In particular, we assume life expectancies according to the Social Security Administration projections. The distribution of internal rates of returns across treatments is shown in the third and fourth

[8]The p-value for equality-of-means tests are higher than for the equality-of-medians test since the variance of individuals' answers is quite high, thus reducing the power of the t-test. All t-tests shown in this paper are for unequal variances.

columns of Table 4. Note that the numbers represent annual return rates (in percentages). In the *forward* treatment, the median rate is .4 percent, whereas it is 4.1 percent in the *backward* treatment. The difference is significant at the 1-percent level. Furthermore, the upward shift of returns in the *backward* treatment is pronounced at all percentiles except the extreme ones. The overall distributions are significantly different as shown by the Wilcoxon (Mann-Whitney) ranksum test. The same is true for means.

A perceived difference in returns of about 3.5 percent, at the medians, across treatments is both economically and statistically very large. Over a time horizon of 20 years, this implies a difference in total accumulation of more than 100 percent.[9] This means that, everything else equal, someone adopting the backward perspective expects to accumulate more than double the amount over a time of 20 years, compared to someone adopting the forward perspective.

As a robustness test, we also consider a life expectancy of 100 years (i.e. $D_i = 100$) for all individuals. The fifth and sixth columns of Table 4 show internal rates of returns under this assumption. The level of implied returns is now higher in both treatments. The reason is that a higher return is required to sustain a given ratio C_i/S_i until age 100. In this scenario, life expectancies are, on average, 16 years longer than in the first scenario. Apart from this general level effect, the pattern mimics the one in columns three and four very closely.

An alternative to assuming a particular life expectancy, D_i, and solving for the implied rate of return, r_i, is to assume a rate of return and solve for the implied life expectancy. The seventh and eighth columns of Table 4 show the resulting distributions of life expectancies if r_i is set to 3 percent for all individuals. The two final columns show life expectancies for an interest rate of 1 percent.

At an interest rate of 3 percent, for 30 percent of the observations the implied life

[9] When looking at the raw ratios C_i/S_i, the ratio is 60 percent higher in the *backward* treatment, whereas, for internal rates of returns it is more than 100. It is not possible to infer a difference in internal rates of return from a given difference of C_i/S_i. The reason is that internal rates of return depend on age, expected retirement and life expectancy. None of this information is contained in C_i/S_i.

expectancy is extremely high if not infinite. This arises if C_i/S_i is low for an individual, while the remaining accumulation horizon $R_i - t_i$ is relatively long. In this case, a given level of of savings of S_i may actually sustain a *low* level of consumption forever! This is the more likely to occur, the higher the assumed rate of return. Intuitively, if the interest rate is greater than the share of assets consumed each year the assets can last in perpetuity. In this case, the implied life expectancy according to equation (2) is infinite.[10] For the purpose of our analysis, we set the maximum implied life expectancy to 150.[11]

The numbers in the last four columns of Table 4 show that the implied life expectancy is substantially higher in the *forward* treatment. The intuition for this result is as follows. When adopting the forward perspective, individuals are more pessimistic and indicate a lower ratio C_i/S_i. If the actual return is fixed, then a lower C_i can be sustained by a given S_i over a longer horizon. Therefore, the implied life expectancy is higher. For an interest rate of 1 percent, the median difference in life expectancies is 9 years. For a rate of 3 percent it is 23 years. In fact, the treatment effect increases with higher rates of return at an increasing rate, due to the more pronounced non-linearity of cumulative returns at higher levels of returns. All differences of means and medians across treatments are significant at the 1 percent level.

The bottom line of the results shown so far is that, whatever perspective is chosen to analyze the data, the resulting treatment effect is consistently very large. A striking observation is that the implied life expectancies and implied internal rates of return are often outside the bounds that one might consider reasonable. These numbers could arise for several possible reasons. First, our assumptions of life expectancy (or return rates) may be different from those made by our respondents. However, while we would expect

[10] The implied life expectancy is infinite if $\sum_{k=1}^{R_i - t_i} (1+r)^k S_i/C_i \geq \sum_{k=0}^{\infty} (1/(1+r))^k$, were the right-hand side is equal to $(1+r)/r$. For instance, if $t_i = 35$, $R_i = 65$, $C_i = 100$, and $r = .03$, then the left-hand side exceeds the right-hand side if $S_i \geq 71$. In that case, savings support extra consumption of 100 in perpetuity.

[11] The mean life expectancy is determined by the value selected for the maximum implied life expectancy. Thus, the value of the mean is not meaningful per se. However, the comparison of means across treatments is still informative. It is noteworthy that, in our data, nearly all respondents with life expectancy greater than 149 have an implied life expectancy of infinity, in fact.

this type of error to lead to small differences for some individuals, on average, the Social Security life expectancy should be roughly correct. In particular, the errors cannot systematically differ across treatments since the latter are randomly assigned. Second, our results represent respondents' perceptions. Optimism or pessimism could lead to substantial variation of perceptions. However, it is unlikely that even the most optimistic individuals expect to live more than 110 years or earn rates of return over 25% per year. Finally, individuals may lack knowledge and intuition about compounding. Overall, we conclude that the variation is too large to be justified by only variation in fundamentals.

Next we consider whether a treatment effect is observed in multivariate regressions controlling for basic demographic characteristics and variables that broadly measure how skilled individuals are in making retirement savings decisions. The latter variables include dummy variables for self-assessed math skills, for retirement planning, using advice for retirement saving, understanding of compound interest, financial literacy, and for high numeracy. (See the Appendix for more information about those variables.) While the tests in Table 3 indicate that the two treatments groups are similar, these regression results will reassure us that a treatment effect exists. Table 5 shows the regression results for the two internal rates of return variables.[12]

The first two columns in Table 5 refer to the internal rates of return under Social Security life expectancies. In the first column, the whole sample is included. In the second column the top and bottom 5 percentiles are excluded. We show regressions for the limited sample in order to see to what degree the results are driven by extreme answers. With the Social Security life expectancy, we find a treatment effect equivalent to 3.1 percentage points in the whole sample. For the limited sample it amounts to 3.2 percentage points. Both are significant at the 1-percent level. Thus, also when controlling for other characteristics, those in the backward treatment have significantly higher internal rates of return. The last two columns in Table 5 refer to internal rates of returns with

[12]Because so many people have an implied life expectancy of infinity, it is not meaningful to consider a regression for implied life expectancy as the dependent variable.

100 years as the life expectancy. Here we find a treatment effect of 3.1 percentage points in the whole sample and also the limited sample.

Most of the other control variables in these regressions are insignificant. However there are a few exceptions. Individuals over the age of 55 give larger answers on average. In the full sample their responses are 6.7 or 7.6 percentage points higher, depending on the life expectancy assumptions. In the limited sample, these numbers become 3.6 and 4.0 percentage points. Despite the fact that we might expect older individuals to move their wealth to safer assets and thus have a lower rate of return, we find that their responses indicate a higher rate of return. Part of this is driven by the fact that older individuals are more likely to give answers in the extremes.[13]. We also find that, in some instances, individuals with the highest household incomes and individuals who have used some advisory services have higher internal rates of return, although this is not consistent across specifications. Those with a good understanding of compound interest and those with high numeracy levels have lower internal rates of return.

The next question that we address is whether there is any subgroup in our sample that is not susceptible to a treatment effect. We address this issue by comparing the two treatments for different subsets of the population. We consider subgroups based on age, education, income, retirement planning, self-assessed math ability, numeracy, financial literacy, compounding literacy and gender. These subgroups focus on characteristics that are the least likely to be susceptible to a treatment effect.[14]

Table 6 reports the mean and median internal rate of return by subgroup and by treatment and tests for differences in internal rates of return across treatments but within subgroups. We show results for internal rates of return based on Social Security life expectancies. Results for a life expectancy of 100 are very similar. In nearly all cases, we find significant treatment effects for medians, and for means when the 5 percent extreme answers are dropped. Inclusion of the 5 percent extremes substantially increases the

[13]This will be discussed further in Section 5.3

[14]See the Appendix for a definition of the respective variables.

variance. As a result, the t-test for equality of means has less power. However, in many cases, equality of means is also rejected in the whole sample. Overall, individuals in the *backward* treatment give significantly higher responses than those in the *forward* treatment. The only notable exception is the group of those over the age of 55. However, this group is particularly likely to give extreme answers (see columns labeled "fraction extreme" in Table 6 and the discussion of Table 8 in Section 5.3 below). Thus, we certainly cannot conclude that this group has a reasonable intuition about the returns to long-term savings.

Overall, we find no systematic evidence of there being any subgroup that is less susceptible to the treatment effect. This is true even for the highly numerate, those with high self-assessed math skills or high financial literacy. In further analysis, not shown here, we consider multivariate regression models to test whether there is a stronger treatment effect for particular subgroups of the population.[15] In these analyses, our results are consistent with the above conclusions.

5.2 Explaining the Treatment Effect

Our analysis has documented that adopting either the forward or the backward perspective for an intuitive estimate of the returns to long-term savings triggers a large framing effect. Furthermore, all subgroups of the population appear equally likely to be susceptible to this framing effect. Naturally, this finding raises the question: where does this framing effect come from? Below, we address several hypotheses. These relate to the wording of the questions, the plausible range of answers, rounding, anchoring, and loss aversion. We discuss each of them in turn.

First, we investigate whether it may be the wording of the question that makes individuals more pessimistic in the *forward* treatment. We chose the wording of our questions in a way that would make them sound as realistic as possible. However, vernacular lan-

[15] Results are available upon request.

guage often lacks the precision of scientific language and, in particular, may contain subtle cues that influence people's information processing. This may also be the case for our questions. In the *forward* treatment, we simply ask individuals to think about additional savings, without mentioning that they are intended to finance a particular target need during retirement. Since there is no future target value, individuals may be triggered to think that these savings would be invested in high-return/high-risk assets. In the *backward* treatment, we mention the word "need". As a result, individuals may apply a low risk-free discount rate to obtain their answer.

However, if this were the case, then the treatment effect would go exactly in the opposite direction from what we observe. We should then observe lower internal rates of return in the *backward* treatment and higher ones in the *forward* treatment. Hence, this wording effect cannot explain our findings.

The second hypothesis relates to rounding and the fact that the natural range of plausible answers differs across treatments. In the *backward* treatment, internal rates of return increase exponentially as the response gets closer to zero. If $100 of consumption per month can be sustained on only $1 of saving per month, the rate of return must be very high. However, because respondents tend to respond in multiples of 10, very high rates of return will be less likely in the *backward* treatment. In the *forward* treatment, internal rates of return increase (exponentially) as the response *increases* (exponentially); the responses are not bounded (psychologically) in any significant way, e.g. because respondents answer in multiples of 10. If this effect were important, then we should observe a downward shift of internal rates of returns in the *backward* treatment, relative to the *forward* treatment, at least in the right half of the distribution. However, what we observe is, again, exactly the opposite.

Alternatively, it is conceivable that individuals take the amount of consumption, C_i, in the *backward* treatment as an upper bound for their estimated savings, S_i. If so, there is reason for concern, since C_i should clearly not necessarily be an upper bound for S_i.

To see this, consider the case that the remaining accumulation time $R_i - t_i$ is short, while the retirement span $D_i - R_i$ is relatively long. In this situation, we clearly must have $S_i > C_i$ for reasonable rates of return.

To investigate whether this effect may explain our treatment effect we proceed as follows. We drop observations in the *backward* treatment for which $S_i < C_i$ but we expect to observe $S_i > C_i$, according to a benchmark value for S_i. For the calculation of the benchmark answer, we use equation (1), assuming an interest rate of 3 percent. Under this assumption, 82 observations wrongly take C_i as an upper bound for S_i. If we remove these observations from the *backward* treatment sample, the median in the *backward* treatment is reduced to 2.2 percent, compared to 4.1 percent for the full sample in this treatment. The difference in medians between the treatments is still significant at the 1-percent level. Differences in means are no longer significant after dropping these observations.

Overall, this analysis provides some evidence that part of the treatment effect may indeed be explained by the fact that up to 20 percent of individuals in the *backward* treatment wrongly take C_i as an upper bound for their answer. For these individuals, the internal rate of return is higher than it should be due to this misperception.[16]

A further issue to be addressed is whether the application of simple rounding heuristics could induce a systematic treatment effect. For instance, individuals may not indicate an answer of 233, but rather 200 or 250. In the *forward* treatment, the "rounding error" appears in the numerator of the ratio C_i/S_i, while, in the *backward* treatment, it would appear in the denominator. The question is whether a rounding error in the numerator would have a different effect from a rounding error in the denominator, such that a treatment effect would appear only for this reason. More abstractly, the question is whether adding or subtracting a given amount Δ to the "true" answer may induce a treatment effect.

[16] It is also conceivable that some individuals in the *forward* treatment wrongly take S_i as a lower bound for their answer. If so, their internal rates of returns are *higher* than they should be. Since returns in the *forward* treatment are actually more pessimistic than in the *backward* treatment, this cannot explain the treatment effect.

In order to explore this possibility, we calculate the benchmark answers using equation (1), assuming an interest rate of either 1, 3, or 6 percent. We then randomly add an amount of either 50, 75, or 100 to the benchmark answers of half of the sample. We subtract the same amount from the benchmark for the other half of the sample. We then use the perturbed answers to calculate internal rates of returns. We find that this procedure can indeed induce a treatment effect that increases with the size of the perturbation. However, the resulting treatment effect is fairly small and maximally amounts to half a percentage point. This is much smaller than the treatment effect observed in the data. In light of this, we rule out rounding as a main driving factor for our observed treatment effect.

We turn now to anchoring as a potential explanation of the treatment effect. In the simplest case of anchoring, respondents would simply reproduce the number provided by us. As a consequence, we would observe consumption equal to savings, $C_i = S_i$, for anchoring individuals. This, per se, would not induce a treatment effect since the ratio C_i/S_i would be 1 for all anchoring individuals, independent of treatment. A treatment effect would be induced if individuals were more likely to anchor to the number on the screen in one treatment than in the other. In our data, we find $C_i = S_i$ for 13 percent of respondents in the *forward* treatment, and 10 percent in the *backward* treatment. This difference cannot explain the treatment effect. If we exclude those individuals from the sample, we find a treatment effect that is as pronounced as when including them.

A more subtle form of anchoring would be the following. In the *backward* treatment, individuals are asked to estimate a value for savings, S_i, that would be sufficient to finance consumption, C_i. Some individuals may respond to this by simply anchoring to their own monthly savings ("well, it may just be about my monthly savings"). In the *forward* treatment, there is no such tangible anchor available. As a result, individuals' answers may simply depend, in one way or another, on the number provided by us.

If this form of anchoring were to occur, we should see a higher correlation between a respondent's answer and the number provided by us in the *forward* treatment. In our data, this correlation amounts to .01 in the *forward*, and .08 in the *backward* treatment. If we cut 5 percent of extreme values at the tails of the distribution, separately for each treatment, the correlation increases to .04 in the *forward*, and .14 in the *backward* treatment. This runs against the above hypothesis.

Our data contain information about monthly savings diverted to retirement preparation. The above hypothesis would require that there is a higher correlation between monthly savings and an individual's answer in the *backward* treatment. In the full sample, the correlation is -.05 in the *forward*, and -.01 in the *backward* treatment. When dropping extreme answers, we obtain -.07 for the *forward*, and .03 for the *backward* treatment, respectively. In fact, we find that the correlations are significant in the *forward*, but not in the *backward* treatment.[17] This is exactly the opposite of what would support the above anchoring hypothesis. In sum, we find no evidence for this type of anchoring.

A final explanation of the treatment effect may relate to a form of loss aversion.[18] Individuals may see savings as a loss, while current or future spending are seen as gains. According to loss aversion, losses loom larger and are more salient than gains. As a result, individuals may feel a psychological "need" to make perceived losses as small as possible, within the range of what appears still plausible to their intuition. In our *forward* treatment, their is no scope for making losses small. However, in the *backward* treatment, there is scope for reporting small amounts of savings that finance future levels of consumption, C_i. If respondents minimize their psychological suffering by reporting small amounts of S_i, the effect of this is, of course, that the ratios C_i/S_i are higher in the *backward* treatment, thus also the implied internal rates of return.

Of course, individuals could also foster their wellbeing in the *forward* treatment by

[17]To understand why this correlation can be significant, as in the *forward* treatment, recall that the amounts provided by us are calculated based on respondents' income. In particular, they are positively correlated with their income. If there is a correlation between an individual's answer and the number provided by us, on the one hand, and a correlation between income and savings, on the other, then the number provided by us may be correlated with savings.

[18]See Della Vigna (2009) and Goette et al. (2004) for a discussion of loss aversion.

increasing the amount C_i. But this means increasing a gain. Since gains are less salient then losses, this effect is likely to be less prevalent. In our opinion, this form of loss aversion is a plausible explanation of at least part of the treatment effect. With our data, we cannot formally test this hypothesis.

Overall, we find that the treatment effect can be partly explained by the fact that respondents in the *backward* treatment take the amount C_i provided by us as an upper bound for their answer for S_i. When controlling for this, the treatment effect is reduced by about half its size. Furthermore, we find it plausible that part of the treatment effect is explained by loss aversion. We do not find any evidence that wording, rounding or anchoring are the driving forces of our treatment effect.

5.3 What Drives Variation within a Treatment

In this subsection we aim to shed more light on what drives variation within a treatment. We do so by performing regression analyses, separately for each treatment. Since the magnitude and significance of regression coefficients are susceptible to extreme answers, we concentrate on results for which 5 percent of observations in either tail are dropped.[19]

The first two columns in Table 7 show regressions with the internal rate of return assuming Social Security life expectancies as the dependent variable. The first column refers to the *forward* treatment, while the second refers to the *backward* treatment. The third and fourth columns show regressions for the internal rate of return with a planning horizon of age 100 as the dependent variable, again for the *forward* and *backward* treatment, respectively.[20] A dummy for age greater than 55 is consistently significant across

[19] Qualitative results for the full sample are broadly in line with the results reported below, however standard errors are much larger due to large outliers. These results are available upon request.

[20] We do not show any regressions for implied life expectancies. Regression analysis is, by construction, an exercise based on means. As mentioned in Footnote 11, the implied life expectancy is infinite for many individuals. Regressions using implied life expectancies would thus explain either an average of infinity, or an average that is determined by whatever finite value we reset life expectancy for those with a value of infinity. Therefore, regression analysis is not meaningful for analyzing the variation of the data *within* a treatment.

regressions. Occasionally, belonging to the richest quartile, getting advice, and a high numeracy level are also significant. It is noteworthy that advice is only significant in the *backward* treatment and the sign goes in the same direction as the treatment effect. In other words, for individuals using advice, the treatment effect tends to be bigger. This means that individuals getting advice are particularly biased. While this may be the result of mediocre advice, it may also mean that individuals who know that they tend to make extreme estimates of the returns to savings are more likely to seek advice.

To complement this regression analysis, we also investigate who is most likely to give extreme answers that fall into either of the 5-percent tails of each treatment. In particular, Table 8 provides probit estimates for extreme internal rate of return under the assumed Social Security life expectancies. The first column shows the results for both treatments pooled. The second and third column show the results for the *forward* and *backward* treatment, respectively. The only variable that is significant in more than one specification and at a high level is the dummy for age greater than 55. While one may have assumed that these people do have the best intuition since retirement preparation is a more salient issue for them, we find the opposite. This suggests that it is important to include those close to retirement into financial literacy programs.

We also investigate to what degree individuals' answers are correlated with benchmark answers. To calculate the latter, we use again equation (1) and assume an interest rate of either 1, 3, or 6 percent. We then run regressions separately by treatment with an individual's answer (in dollar amounts) as the dependent variable. The explanatory variables are the calculated answers and a constant. We again omit answers that fall into the 5-percent tail of either treatment. The first three columns in Table 9 show the results for the *forward* treatment. While the calculated answer is always significant at the 10-percent level, the resulting R^2 are extremely low, ranging between .01 and .03. The last three columns refer to the *backward* treatment. There, the coefficients are significant at the 1-percent level. The R^2, all around .08, are still low, but consistently higher than for the *backward* treatment. This is a striking result.

Overall, this suggests that individuals cope better with the backward perspective than with the forward perspective. This result may be useful for financial advice. If advice is framed according to the logic of the backward perspective, individuals may get a better feeling of the returns to long-term savings.

As a final point, we investigate whether individuals' answers show any correlations with their actual retirement saving behavior. For a subset of our sample, we have information on their accumulated retirement wealth.[21] Table 10 shows regressions with the log of retirement wealth as the dependent variable and a range of standard explanatory variables. Again, observations from 5-percent tails of internal rates of returns are dropped. The explanatory variable of interest here is the internal rate of return.[22] In the first three columns, we only allow for a linear influence of the internal rate of return on retirement wealth. The first column refers to a sample where treatments are pooled. The second and third column refer to the *forward* and *backward* treatment, respectively. The internal rate of return is not significant in any of the regressions.

One reason for this result may be that the effect of perceived returns on savings may be ambiguous. In particular, a very low perceived return may induce people to save very little since it appears fruitless. On the other hand, a very high perceived return may induce people to postpone savings. Thus, the pattern of how return rates affect savings may be hump-shaped. To account for this, we include a fourth-order polynomial of the internal rate of return. The results are shown in the last three columns of Table 10. None of the coefficients of the internal rate of return terms are significant, either individually or jointly. We obtain similar results (not shown) when estimating regressions with monthly savings as the dependent variable. One reason for this lack of correlation may be that the location of the hump may differ across individuals. Alternatively, the income and substitution effects may work in opposite directions, making the sign of the relationship

[21]See Binswanger and Carman (2010a) for details about this variable.
[22]In Table 10, we assume Social Security life expectancies.

between the internal rate of return and wealth ambiguous. High rates of return may lead some people to postpone savings and others to save even more to capture the high returns.

6 Conclusion

In this paper, we have investigated people's intuition about the benefits of long-term saving. The motivation for this is, as we argued above, that many people are, in fact, likely to rely on their intuition when it comes to estimating the benefits of long-term savings. In particular, in our sample approximately 50 percent neither calculate the benefits of long-term savings themselves, nor get any help with this from advisors or financial websites. This is particularly noteworthy since our sample is relatively skilled, compared to the overall population, and 90 percent of our sample are older than 30. We thus expect that the fraction of people relying on their intuition may be even higher in the general population.

Our research is based on a randomized experiment. We ask individuals to estimate the returns on long-term savings under two conditions. In one condition, individuals take a forward perspective in that they start with a given amount of savings and estimate how much spending this would finance during retirement. In the other condition, individuals start with a given amount of spending during retirement and estimate the savings that are required to finance those spendings. They thus adopt a backward perspective.

From an economic point of view, both perspectives are equivalent. However, we discover a large "direction bias": individuals are far more pessimistic about the long-term benefits of savings when they adopt the forward perspective, compared to the backward perspective. The bias amounts to about 3.5 percentage points per year. Over a horizon of 20 years, this accumulates to more than 100 percent. Alternatively, we quantify the treatment effect to be equivalent to a difference in life expectancies of 23 years across treatments. Since assignment to treatments is randomized, this difference cannot be explained by fundamental factors but reflects a framing effect. We find this effect to be universal, i.e. there is no evidence that it varies with financial literacy, numeracy, education, age etc.

In our experiment, we ask individuals to think about the trade-off between consumption today and consumption tomorrow. If if individuals want to maximize their utility across the life-cycle this is exactly the trade-off they need to consider. Using a life-cycle model, we would expect people to try to smooth consumption over time. However, to do this they should have some information about this fundamental trade-off, whether it comes from formal calculations or merely intuition. Even if we reject the life cycle model, informed choices about savings require some knowledge about the returns to savings. Our experiment tests this fundamental assumption: whether individuals have a reasonable intuition about how their savings today will affect their consumption tomorrow. Our results suggest that many people fundamentally do not understand this trade-off. Thus they may be unable to make rational choices about saving for retirement.

As a consequence, our results have implications for economic theory as well as empirical analyses. We have pointed out that many people are likely to rely, in one way or the other, on their intuition or gut feeling. Since we find this gut feeling to be very unreliable, it may lead to many inadequate choices. As a result, people are unlikely to behave "as if" they knew how to perform compound interest calculations (Friedman, 1953). Thus, the Euler equations of the standard life cycle model may not apply. This is also important for empirical analyses, which often use predictions of the standard model as a starting point for deriving estimation equations.

On the practical side, our research implies that individuals need help when privately preparing for their retirement. A first important implication of our result is that the perception of the returns to long-term savings depends fundamentally on framing. It is important to be aware of the relevant effects when structuring advice or designing financial literacy programs. Furthermore, we provide evidence that individuals provide more reasonable answers when adopting the backward perspective. Our results suggest that this type of framing is preferable.

Overall, financial literacy programs may be a first important step to induce people to make more informed choices about retirement preparation (Bernheim and Garrett, 2003; Lusardi, 2008b). A potential difficulty associated with financial literacy programs is to encourage people to put the acquired knowledge into practice. People may have good intentions to promote their future well-being, but may be overconfident about how responsive their future selves are to those intentions (Choi et al., 2006; Della Vigna and Malmendier, 2006). In the context of retirement savings, improving people's understanding of the miracle of compound interest may help to minimize this problem; people may be more prone to delay because they do not fully understand the foregone benefits. Therefore, financial literacy programs may be more effective if they help people to understand the miracle of compound interest and the fundamental trade-off between saving today and consumption tomorrow. Furthermore, financial literacy, and financial advice in general, may make use of our finding that individuals cope better with a "backward perspective." What design of instruction or advice would be most effective remains an important topic for future research.

References

Ameriks, John, Caplin, Andrew, and Leahy, John (2003), "Wealth Accumulation and the Propensity to Plan," *Quarterly Journal of Economics*, 118(3), 1007-1047.

Behrman, Jere R., Mitchell, Olivia S., Soo, Cindy, and Bravo, David (2010), "Financial Literacy, Schooling, and Wealth Accumulation," NBER Working Paper 16452.

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

Binswanger, Johannes, and Carman, Katherine G. (2010a), "How Real People Make Long-Term Decisions: The Case of Retirement Preparation," Working Paper, Tilburg University.

Binswanger, Johannes, and Carman, Katherine G. (2010b), "The Role of Decision Making Processes in the Correlation Between Wealth and Health," Working Paper, Tilburg University.

Choi, James, Laibson, David, Madrian, Brigitte C., and Metrick, Andrew (2006) "Saving for Retirement on the Path of Least Resistance," in: McCaffrey, Edward J., and Slemrod, Joel (eds.), *Behavioral Public Finance: Toward a New Agenda*, New York: Russell Sage Foundation, 304-351.

Della Vigna, Stefano (2009), "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47(2), 315-372.

Della Vigna, Stefano, and Malmendier, Ulrike (2006), "Paying Not to Go to the Gym," *American Economic Review*, 96(3), 694-719.

Donkers, Bas, Melenberg, Bertrand, and van Soest, Arthur (2001), "Estimating Risk Attitudes using Lotteries: A Large Sample Approach," *Journal of Risk and Uncertainty*, 22(2), 165-195.

Goette, Lorenz, Huffman, David, and Fehr, Ernst (2004), "Loss Aversion and Labor Supply," *Journal of the European Economic Association,* 2(2-3), 216-228.

Frederick, Shane (2005), "Cognitive Reflection and Decision Making," *Journal of Economic Perspectives*, 19(4), 25-42.

Friedman, Milton (1953), "The Methodology of Positive Economics," in: *Essays in Positive Economics*, Chicago University Press, 3-43.

Lusardi, Annamaria (2008a), "Financial Literacy: An Essential Tool for Informed Consumer Choice?" NBER Working Paper 14084.

Lusardi, Annamaria (2008b), "Household Saving Behavior: The Role of Financial Literacy, Information, and Financial Education Programs," NBER Working Paper 13824.

Lusardi, Annamaria, and Mitchell, Olivia S. (2007a), "Baby Boomer Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth," *Journal of Monetary Economics*, 54(1), 205-224.

Lusardi, Annamaria, and Mitchell, Olivia S. (2007b), "Financial Literacy and Retirement Planning: New Evidence from the Rand American Life Panel," Working Paper.

Lusardi, Annamaria, and Mitchell, Olivia S. (2008), "Planning and Financial Literacy: How Do Women Fare?" *American Economic Review*, Papers and Proceedings 2008, 98(2), 413-417.

Lusardi, Annamaria, Mitchell, Olivia S., and Curto, Vilsa (2009), "Financial Literacy among the Young: Evidence and Implications for Consumer Policy," NBER Working Paper 15352.

Mullainathan, Sendhil, Nöth, Markus, and Schoar, Antoinette (2010), "The Market for Financial Advice: An Audit Study," Working Paper.

Skinner, Jonathan (2007), "Are You Sure You're Saving Enough for Retirement?" *Journal of Economic Perspectives,* 21(3), 59-80.

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

Van Rooij, Maarten, Lusardi, Annamaria, and Alessie, Rob (2007), "Financial Literacy and Stock Market Participation," NBER Working Paper 13565.

Appendix: Additional Variable Definitions

Here we describe the questions used to measure self-assessed math skills, planning, advice, financial literacy, compounding literacy, and numeracy. Self-assessed math skills, planning, and use of advice were measured at the same time as other questions used in this survey. Numeracy, financial literacy, compounding literacy were measured during other surveys, as described in more detail below.

Math skills is based on a question from Ameriks, Caplin and Leahy (2003). The corresponding survey question is:

> *I am highly confident in my mathematical skills. (Answers: Fully disagree; somewhat disagree; neither agree nor disagree; somewhat agree; fully agree.)*

In our analysis, we use the dummy variable *math skills* that takes on a value of 1 if the answer to the above question is *fully agree*, and zero otherwise. *Math skills* takes on a value of one for 24 percent of our sample.

Planner is based on the following two questions.

> **P1** *I've tried to determine my financial needs during retirement. (Answers: Fully disagree; somewhat disagree; neither agree nor disagree; somewhat agree; fully agree.)*
>
> **P2** *Have you ever tried to find out how much you should save in total today and in the coming years in order to finance your target needs during retirement? (Answers: yes; no.)*

The dummy variable *planner* takes on a value of one if a respondent at least somewhat agrees to P1, and if her answer to P2 is yes. 31 percent of our sample are classified as planners.

Advice is set equal to one if the respondent used any of three forms of advice for retirement saving: financial planning software, a financial calculator website, or a financial advisor or broker. The underlying survey question is:

For your retirement planning, do you rely on financial software, a website with a financial calculator, or a broker or financial advisor? You may check several answers. (Possible answers: Financial software; financial calculator; broker or financial advisor; I do not rely on any of them.)

44 percent of our sample use at least one form of advice.

Financial literacy and **compounding literacy** are measured based on questions used by Lusardi and Mitchell (2007a, 2007b, 2008) and van Rooij et al. (2007). These questions have been asked multiple times in the ALP: in Well Being 21 fielded from April 21, 2008 to September 10, 2009; in Well Being 64 fielded from March 5, 2009 to September 10, 2009; and in Well Being 5 fielded from May 8, 2006 to November 1, 2007. In order to maximize the sample for whom we can measure financial literacy, we used data from all three surveys. For respondents who answered multiple surveys, scores are consistent across surveys. But when multiple responses are available we use the responses from the date closest to the time our survey was fielded. Financial literacy is based on all 5 questions listed below and individuals are classified as financially literate if they get all 5 questions correct.

L1 *Suppose you had $100 in a savings account and the interest rate was 2 percent per year. After 5 years, how much do you think you would have in the account if you left the money to grow: more than $102, exactly $102, less than $102? (Answers: More than $102; exactly $102; less than $102; I don't know.)*

L2 *Suppose you had $100 in a savings account and the interest rate is 20 percent per year and you never withdraw money or interest payments. After 5 years, how much would you have on this account in total? (Answers: More than $200; exactly $200; less than $200; I don't know.)*

L3 *Imagine that the interest rate on your savings account was 1 percent per year and inflation was 2 percent per year. After 1 year, would you be able to buy more than, exactly the same as, or less than today with the money in this account? (Answers: More than today; exactly the same as today; less than today; I don't know.)*

L4 *Assume a friend inherits \$10,000 today and his sibling inherits \$10,000 but 3 years from now. Who is richer today because of the inheritance? (Answers: My friend; his sibling; they are equally rich; I don't know.)*

L5 *Suppose that in the year 2010, your income has doubled and prices of all goods have doubled too. In 2010, will you be able to buy more, the same or less than today with your income? (Answers: Buy more than today; buy the same as today; buy less than today; I don't know.)*

47.42 percent are financially literate. Because we are particularly interested in respondents' understanding of compounding, the variable *compounding literacy* is based on the first two questions only and we consider individuals to be literate if they get both L1 and L2 correct. The compounding literate group includes 74.47 percent of our sample.

Numeracy is measured using a cognitive reflection test, consisting of three questions. This test has previously been used by economists (see Frederick, 2005).[23] Information about numeracy is available for 744 respondents in our sample. These questions were fielded in a previous survey conducted by the ALP during June 2008, labeled as Well Being 32. The questions are as follows.

N1 *A bat and a ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost?*

N2 *If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?*

[23]Our results are robust to the use of other numeracy measures.

N3 *In a lake, there is a patch of lilypads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?*

We classify individuals are "numerate" if they got at least two questions right. This includes 25 percent of our sample.

Figure 1: Internal rates of return for $R = 65$, $D = 90$

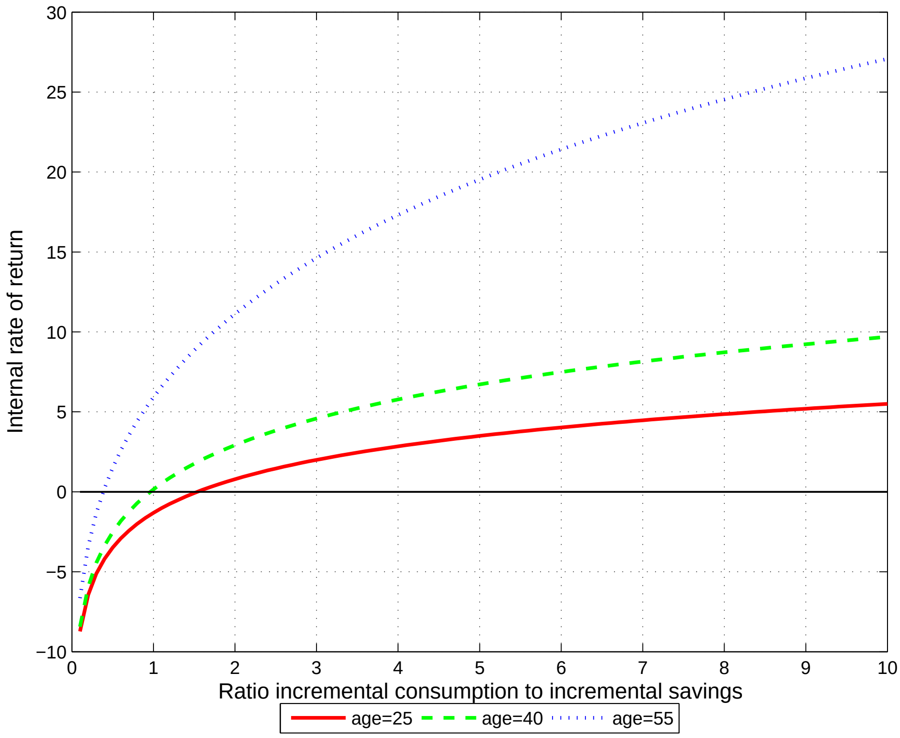

Table 1: Summary statistics

	Mean	Median	Standard deviation
Age	47.667	49	10.589
Male	0.480	0	0.500
Married	0.669	1	0.471
College	0.283	0	0.451
Advanced degree	0.262	0	0.440
Income	89,423	87,500	48,187
Income≥200,000	0.045	0	0.208
Home ownership	0.792	1	0.406
Access DC pension	0.908	1	0.289
Retirement wealth	190,076	62,500	365,268

Table 2: Saving/spending inputs for financial intuition question

	Forward		*Backward*	
Income Bracket	Input (S_i)	Frequency	Input (C_i)	Frequency
Less than 5,000	10	3	20	3
5,000 to 9,999	25	5	50	2
10,000 to 14,999	50	5	100	1
15,000 t0 34,999	100	49	200	38
35,000 to 59,999	200	80	400	103
60,000 to 74,999	300	58	600	55
75,000 or more	400	206	800	257

NOTE: The income brackets refer to annual income.

Table 3: Summary statistics by treatment

	Mean		Standard deviation		p-value equal means
	Forward	*Backward*	*Forward*	*Backward*	
Age	47.450	47.856	10.871	10.345	0.5762
Male	0.474	0.486	0.500	0.500	0.7277
Married	0.648	0.688	0.478	0.464	0.2058
College	0.310	0.259	0.463	0.439	0.0974*
Advanced degree	0.236	0.285	0.425	0.452	0.1014
Income	87,006	91,567	48,445	47,909	0.168
Income≥200,000	0.044	0.046	0.206	0.209	0.2788
Home ownership	0.768	0.813	0.422	0.391	0.1123
Access DC pension	0.912	0.904	0.283	0.295	0.6726
Retirement wealth	188,017	191,971	380,763	350,953	0.8906

Table 4: Treatment effect

	Ratio C_i/S_i		Internal rate of return Social security life expectancy		Internal rate of return Life expectancy age 100		Life expectancy Return=3%		Life expectancy Return=1%	
	Forward	*Backward*	*Forward*	*Backward*	*Forward*	*Backward*	*Forward*	*Backward*	*Forward*	*Backward*
5th percentile	0.050	0.032	-13.629	-17.649	-5.985	-7.343	63	61	63	61
10th percentile	0.125	0.333	-9.419	-6.443	-3.761	-1.033	66	65	66	64
25th percentile	0.500	0.800	-3.470	-0.560	-0.286	2.228	74	69	72	68
50st percentile	1.000	1.600	0.443	4.102	2.805	6.107	103	80	86	77
75th percentile	1.667	4.000	6.467	11.084	7.989	11.728	150	117	121	91
90st percentile	7.500	10.000	20.910	22.064	22.276	22.271	150	150	150	140
95th percentile	50.000	16.000	32.159	30.419	32.425	30.529	150	150	150	150
Mean	41.470	4.697	2.984	5.539	5.623	8.212	110.1679	94.62192	97.29262	85.50112
Mean excluding 5% tails	2.372	3.316	2.108	5.430	4.544	7.646	N/A	N/A	N/A	N/A
Standard deviation	310.909	10.847	16.473	14.372	15.440	11.526	36.0392	32.79626	31.1588	25.8885
p-value equality of means	0.0177**		0.0184**		0.0068***		0.000***		0.000***	
p-value equality of means excluding 5% tails	0.0133**		0.000***		0.000***		N/A		N/A	
p-value for Wilcoxon rank-sum (Mann-Whitney) test	0.000***		0.000***		0.000***		0.000***		0.000***	
p-value for equality of medians	0.000***		0.000***		0.000***		0.000***		0.000***	

NOTE: Cells showing N/A refer to means of implied life expectancies. Since more than 5 percent of observations take on the maximum value, excluding 5 percent of extreme values is not meaningful (see Footnotes 11 and 20). One, two, and three asterisks indicate a p-value less than .1, .05, and .01, respectively.

Table 5: Regressions for treatment effect

	Internal rate of return (Social Security life expectancy)		Internal rate of return (Life expectancy 100)	
	All	W/o extreme 5%	All	W/o extreme 5%
Backward treat.	3.112***	3.227***	3.087***	3.067***
	(1.104)	(0.680)	(0.986)	(0.560)
Age $\geq$55	6.735***	3.567***	7.559***	4.035***
	(1.455)	(0.892)	(1.234)	(0.728)
Male	-0.495	0.659	-0.446	0.871
	(1.112)	(0.697)	(1.052)	(0.575)
Married	0.391	0.091	0.973	0.010
	(1.319)	(0.800)	(1.156)	(0.655)
College	-0.134	-0.026	0.083	-0.135
	(1.436)	(0.806)	(1.197)	(0.661)
Advanced Degree	-0.143	0.317	-0.842	0.157
	(1.385)	(0.907)	(1.380)	(0.742)
Income$\geq$75th perc.	2.682**	1.280	1.862*	1.354**
	(1.334)	(0.819)	(1.122)	(0.679)
Math skills	-1.942	0.457	-1.226	0.030
	(1.275)	(0.851)	(1.069)	(0.713)
Planner	-0.437	-0.377	-0.090	-0.215
	(1.303)	(0.845)	(1.049)	(0.707)
Advice	1.034	1.528**	1.086	0.854
	(1.183)	(0.740)	(0.974)	(0.610)
Compounding literacy	-2.887*	-1.980*	-2.042	-1.916**
	(1.667)	(1.032)	(1.547)	(0.855)
Financial literacy	0.200	0.790	0.069	0.735
	(1.288)	(0.783)	(1.068)	(0.631)
Numerate	-2.189*	-2.012**	-1.342	-1.578**
	(1.196)	(0.781)	(0.971)	(0.625)
Constant	2.683*	1.339	4.013**	3.896***
	(1.610)	(0.999)	(1.624)	(0.824)
Observations	701	633	706	639
R^2	0.073	0.093	0.092	0.127

NOTE: Robust standard errors are indicated in parentheses. One, two, and three asterisks indicate a p-value less than .1, .05, and .01, respectively.

Table 6: Treatment effect for subgroups

	Sample Size		Median		Mean		Mean w/o extreme 5%		Fraction extreme 5%		Equality means	Equality means w/o extreme 5%	Wilcoxon rank-sum	Equality medians
	Forw.	*Backw.*	*Forw.*	*Backw.*	*Forw.*	*Backw.*	*Forw.*	*Backw.*	*Forw.*	*Backw.*				
Whole sample	406	459	0.443	4.102	2.984	5.539	2.108	5.430	0.094	0.096	0.0184**	0.000***	0.000***	0.000***
Age<55	294	333	-0.163	3.735	1.216	3.866	0.886	4.745	0.078	0.063	0.0157**	0.000***	0.000***	0.000***
Age≥55	112	126	4.805	7.108	8.233	10.160	6.059	7.628	0.134	0.183	0.4713	0.3207	0.206	0.533
Income<75th perc.	261	278	0.248	2.938	2.246	4.599	2.113	4.858	0.096	0.086	0.0948*	0.0009***	0.0002***	0.000***
Income≥75th perc.	145	181	1.012	4.848	4.275	6.970	2.100	6.324	0.090	0.110	0.1078	0.000***	0.0002***	0.000***
Low math skills	303	352	0.000	4.236	3.218	6.188	1.754	5.503	0.099	0.088	0.0207**	0.000***	0.000***	0.000***
High math skills	102	107	1.251	3.616	2.315	3.364	3.090	5.173	0.078	0.121	0.6002	0.0851*	0.1021	0.049**
Not planning	282	317	0.000	4.209	2.761	5.606	1.865	5.401	0.089	0.095	0.0313**	0.000***	0.000***	0.000***
Planning	124	142	1.251	4.055	3.488	5.389	2.670	5.494	0.105	0.099	0.3167	0.0136**	0.0152**	0.014**
Low Numeracy	252	297	0.455	4.884	3.080	6.772	2.332	5.935	0.095	0.063	0.0067*	0.000***	0.000***	0.000***
High Numeracy	95	100	0.443	2.336	1.616	3.113	1.362	3.939	0.105	0.090	0.4643	0.0304**	0.0799*	0.059*
Low fin. lit.	228	220	0.696	4.090	3.480	6.599	2.258	5.544	0.083	0.105	0.0417**	0.0006***	0.0005***	0.000***
High fin. lit.	170	234	0.000	4.445	2.343	4.637	1.854	5.438	0.112	0.090	0.1525	0.0001***	0.0001***	0.000***
Low comp. lit.	117	101	0.791	3.697	3.950	7.163	2.702	5.782	0.068	0.109	0.1467	0.0385**	0.0536*	0.121
High comp. lit.	281	353	0.365	4.491	2.604	5.149	1.835	5.408	0.107	0.093	0.048**	0.000***	0.000***	0.000***
No advice	225	262	0.000	2.732	3.672	4.420	1.947	4.464	0.098	0.115	0.6219	0.004***	0.0098***	0.001***
Advice	181	197	1.061	5.315	2.131	7.036	2.306	6.660	0.088	0.071	0.0014***	0.000***	0.000***	0.000***
Female	212	236	0.154	4.158	3.811	5.835	1.417	5.819	0.090	0.076	0.1878	0.000***	0.000***	0.000***
Male	191	223	0.997	4.090	2.041	5.227	2.907	5.000	0.099	0.117	0.0354**	0.0292**	0.0034***	0.002***
No college	184	209	0.806	4.090	3.418	6.305	2.245	5.340	0.099	0.100	0.0887*	0.0014***	0.001***	0.001***
College	126	119	0.000	3.752	1.933	5.323	1.640	5.374	0.111	0.109	0.0896*	0.0011***	0.0011***	0.004***
Advanced degree	96	131	0.710	4.558	3.539	4.494	2.444	5.619	0.063	0.076	0.6197	0.0151**	0.0119**	0.000***
Unmarried	143	143	1.011	2.699	3.896	3.808	2.590	4.249	0.084	0.070	0.9622	0.121	0.1102	0.033**
Married	263	316	0.164	4.605	2.507	6.328	1.854	5.992	0.099	0.108	0.004***	0.000***	0.000***	0.000***

NOTE: The two columns labeled "fraction extreme 5%" show the percentage of observations of each group falling into either the bottom or top 5 percent of the distribution. The last four columns show p-values for equality tests. One, two, and three asterisks indicate a p-value less than .1, .05, and .01, respectively.

Table 7: Regression analysis by treatment

	Internal rate of return (Social Security life expectancy)		Internal rate of return (Life expectancy 100)	
	Forward	*Backward*	*Forward*	*Backward*
Age ≥55	4.369***	2.835**	4.439***	3.808***
	(1.332)	(1.229)	(1.116)	(0.980)
Male	1.668*	-0.145	1.469*	0.390
	(1.004)	(0.996)	(0.831)	(0.819)
Married	-1.114	1.278	-1.213	1.186
	(1.163)	(1.115)	(0.967)	(0.906)
College	0.082	-0.210	0.180	-0.512
	(1.150)	(1.133)	(0.956)	(0.920)
Advanced Degree	0.371	0.121	0.218	-0.039
	(1.347)	(1.212)	(1.133)	(0.973)
Income≥75th perc.	0.851	1.417	0.802	1.720*
	(1.219)	(1.123)	(1.022)	(0.927)
Math skills	0.830	0.170	0.610	-0.404
	(1.239)	(1.176)	(1.043)	(0.983)
Planner	0.798	-1.394	0.360	-0.796
	(1.277)	(1.137)	(1.089)	(0.948)
Advice	0.419	2.455**	-0.000	1.582*
	(1.074)	(1.015)	(0.893)	(0.834)
Compounding literacy	-1.962	-1.747	-1.836	-1.919
	(1.353)	(1.591)	(1.139)	(1.297)
Financial literacy	0.427	0.758	0.635	0.518
	(1.144)	(1.087)	(0.948)	(0.867)
Numerate	-1.887*	-2.122*	-1.122	-1.921**
	(1.098)	(1.123)	(0.902)	(0.895)
Constant	1.710	4.219***	4.300***	6.646***
	(1.277)	(1.356)	(1.089)	(1.095)
Observations	292	341	297	342
R^2	0.086	0.064	0.106	0.098

NOTE: Robust standard errors are indicated in parentheses. One, two, and three asterisks indicate a p-value less than .1, .05, and .01, respectively.

Table 8: Probit regressions for extreme 5 percent (marginal effects)

	All	*Forward*	*Backward*
Age ≥55	0.093***	0.062	0.128***
	(0.028)	(0.040)	(0.039)
Male	0.016	0.011	0.030
	(0.021)	(0.033)	(0.027)
Married	0.021	0.026	0.012
	(0.023)	(0.034)	(0.030)
College	0.032	0.022	0.036
	(0.027)	(0.039)	(0.036)
Advanced Degree	-0.027	-0.037	-0.014
	(0.024)	(0.037)	(0.031)
Income≥75th perc.	0.015	-0.038	0.065*
	(0.024)	(0.035)	(0.034)
Math skills	0.015	-0.008	0.025
	(0.028)	(0.039)	(0.037)
Planner	0.020	0.038	0.017
	(0.026)	(0.042)	(0.032)
Advice	-0.028	-0.011	-0.043
	(0.022)	(0.034)	(0.027)
Compounding literacy	-0.008	0.021	-0.059
	(0.030)	(0.039)	(0.051)
Financial literacy	-0.005	0.003	0.001
	(0.025)	(0.038)	(0.033)
Numerate	-0.008	0.002	-0.013
	(0.024)	(0.039)	(0.029)
Observations	730	342	388

NOTE: The table shows marginal effects. Standard errors are indicated in parentheses. One, two, and three asterisks indicate a p-value less than .1, .05, and .01, respectively.

Table 9: Regressions for relationship between actual and benchmark answer

	Answer C (*forward*)	Answer C (*forward*)	Answer C (*forward*)	Answer S (*backward*)	Answer S (*backward*)	Answer S (*backward*)
Benchmark C 3%	0.190*					
	(0.101)					
Benchmark C 1%		0.279*				
		(0.154)				
Benchmark C 6%			0.092*			
			(0.047)			
Benchmark S 3%				0.496***		
				(0.184)		
Benchmark S 1%					0.376***	
					(0.145)	
Benchmark S 6%						0.704***
						(0.251)
Constant	407.269***	414.290***	407.036***	336.175**	328.651**	354.242***
	(72.618)	(72.306)	(70.363)	(140.079)	(148.289)	(128.779)
Observations	365	365	365	412	412	412
R^2	0.018	0.014	0.026	0.079	0.075	0.084

NOTE: Robust standard errors are indicated in parentheses. One, two, and three asterisks indicate a p-value less than .1, .05, and .01, respectively.

Table 10: Regressions for retirement wealth

	All	*Forward*	*Backward*	*All*	*Forward*	*Backward*
IRR	0.005	0.003	0.015	0.016	0.009	0.064
	(0.014)	(0.022)	(0.019)	(0.037)	(0.056)	(0.057)
IRR^2				0.001	0.002	-0.002
				(0.002)	(0.004)	(0.003)
IRR^3				-0.000	0.000	-0.000
				(0.000)	(0.000)	(0.000)
IRR^4				-0.000	-0.000	0.000
				(0.000)	(0.000)	(0.000)
Age	0.116	0.230	-0.026	0.105	0.204	-0.035
	(0.108)	(0.150)	(0.160)	(0.109)	(0.149)	(0.168)
Age^2	-0.001	-0.002	0.001	-0.001	-0.002	0.001
	(0.001)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
Log income	2.249***	2.207***	2.266***	2.244***	2.235***	2.229***
	(0.328)	(0.496)	(0.414)	(0.329)	(0.496)	(0.404)
Income$\geq$200,000	0.390	0.224	0.687	0.393	0.066	0.797
	(0.357)	(0.508)	(0.496)	(0.359)	(0.502)	(0.512)
House ownership	1.894***	1.899***	1.975***	1.913***	1.944***	1.967***
	(0.410)	(0.536)	(0.636)	(0.410)	(0.544)	(0.634)
Married	-0.486	-0.868*	-0.200	-0.483	-0.841*	-0.174
	(0.358)	(0.482)	(0.517)	(0.359)	(0.484)	(0.515)
Male	0.473*	0.344	0.606*	0.480*	0.331	0.604*
	(0.245)	(0.370)	(0.342)	(0.247)	(0.373)	(0.344)
College	0.592*	0.232	0.886**	0.586*	0.233	0.899**
	(0.312)	(0.468)	(0.406)	(0.314)	(0.471)	(0.413)
Advanced degree	1.149***	1.104***	1.116***	1.143***	1.095***	1.107***
	(0.287)	(0.388)	(0.410)	(0.289)	(0.389)	(0.411)
Constant	-20.052***	-21.267***	-17.935***	-19.822***	-21.117***	-17.392***
	(4.159)	(6.228)	(5.107)	(4.190)	(6.190)	(5.110)
Observations	556	261	295	556	261	295
R^2	0.356	0.360	0.368	0.357	0.366	0.371

NOTE: Robust standard errors are indicated in parentheses. One, two, and three asterisks indicate a p-value less than .1, .05, and .01, respectively.