

Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?*

Ulrike Malmendier
UC Berkeley and NBER

Stefan Nagel
Stanford University and NBER

October 2009

Abstract

We investigate whether individuals' experiences of macro-economic outcomes have long-term effects on their risk attitudes, as often suggested for the generation that experienced the Great Depression. Using data from the Survey of Consumer Finances from 1960-2007, we find that individuals who have experienced low stock-market returns throughout their lives until the time of the survey report lower willingness to take financial risk, are less likely to participate in the stock market, and if they do, invest a lower fraction of their liquid assets in stocks, and are more pessimistic about future stock returns. Individuals who have experienced low bond returns are less likely to own bonds. All results are estimated controlling for age, year effects, and a broad set of household characteristics. Our estimates indicate that more recent return experiences have stronger effects, but older individuals are still influenced by experiences from several decades earlier. These experience effects can explain, for example, the relatively low stock-market participation of young households in the early 1980s, following the disappointing stock-market returns in the 1970s, and the relatively high participation of young investors in the late 1990s, following the boom years in the 1990s. The average investors' lifetime stock-market return experiences lines up well with the market's price/earnings ratio, which suggests that experience-driven variation in risk-taking could potentially provide an explanation for aggregate stock market fluctuations.

* We thank John Ameriks, Brad Barber, Robert Barro, Randy Cohen, Wayne Ferson, Tim Johnson, Larry Katz, John Sabelhaus, Karthryn Shaw, Stephen Zeldes, four anonymous referees, and seminar participants at Berkeley, BI Oslo, Bowling Green State, Copenhagen Business School, Erasmus University Rotterdam, Federal Reserve Bank of Chicago, Imperial College, Johns Hopkins, Harvard, Helsinki School of Economics, Hong-Kong University of Science and Technology, London School of Economics, Nanyang Technological University, National University of Singapore, New York University, NHH Bergen, Singapore Management University, Stanford, Stockholm School of Economics, Tilburg, University of British Columbia, University of Illinois Urbana/Champaign, University of Michigan, University of Southern California, the World Bank, University of Zurich, the American Finance Association Meeting, the Annual Meeting of German Economists Abroad, the European Symposium on Economics and Psychology, the NBER MIDM, Asset Pricing, Cohort Studies, and Economics of Household Savings meetings, as well as the Wharton Conference on Household Portfolio Choice and Financial Decision-Making for comments, and Sith Chaisurote and Nelli Oster for excellent research assistance. Ulrike Malmendier gratefully acknowledges financial support from NIA grant 5P30AG012839-14 and the Coleman Fung Risk Management Research Center at UC Berkeley.

I. Introduction

Does the personal experience of economic fluctuations shape individuals' risk attitudes? For the generation of "Depression Babies" it has often been suggested that their experience of a large macroeconomic shock, the Great Depression, had a long-lasting effect on their attitudes towards risk. In this paper, we ask more generally whether people who live through different macroeconomic histories differ in their level of risk taking.

Standard models in economics assume that individuals are endowed with stable risk preferences, unaltered by economic experiences. Standard models also assume that individuals incorporate all available historical data when forming beliefs about risky outcomes. In contrast, the psychology literature argues that personal experiences, especially recent ones, exert a greater influence on personal decisions than statistical summary information in books or via education (Nisbett and Ross 1980; Weber et al. 1993; Hertwig et al. 2004). Recent literature in economics suggests that the cultural and political environment in which individuals grow up affects their preference and belief formation, such as the level of trust in financial institutions, stock market participation, and preferences over social policies (Guiso, Sapienza, and Zingales 2004 and 2008; Osili and Paulson 2008; Alesina and Fuchs-Schündeln 2007).

We examine empirically whether individuals' willingness to take financial risks differ depending on the macroeconomic history they experienced over the course of their lives. In particular, we test whether individuals who experienced periods of low stock-market returns express a lower willingness to take financial risk, are less likely to participate in the stock market and invest less in stocks, and whether individuals who lived through periods of low bond returns are more wary of participating in the long-term bond market. Life-time experiences could affect risk-taking could arise because experiences affect risk preferences, or because they affect beliefs about future returns from risky asset investments. For the most part, we remain agnostic as to the channel through which experiences affect risk taking, but we offer some evidence suggesting that experiences affect expectations about future risky asset returns.

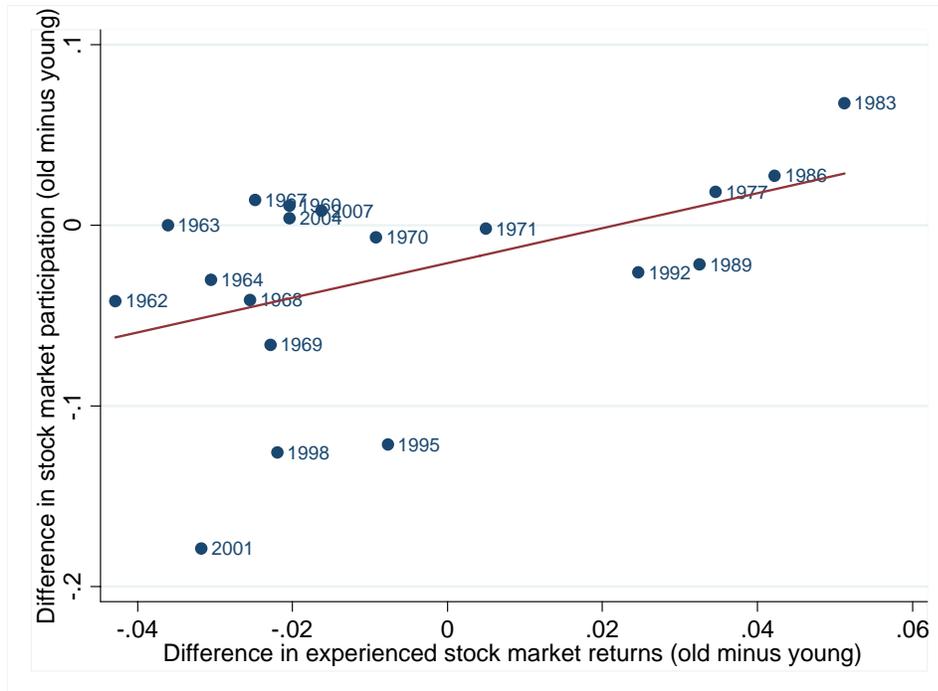


Figure 1: Differences in stock-market participation rates of old and young individuals plotted against differences in lifetime average stock-market returns. Stock market participation rates are the fraction of households who invest in stocks (including stock mutual funds and stocks held in retirement accounts). The y-axis shows the participation rate of old (household head age > 60 years) minus the rate of young (household head age ≤ 40 years) households. The x-axis shows the average real stock market return (S&P500 index) over the prior 50 years (as proxy for the return experienced by old households) minus the return over the prior 20 years (as proxy for the return experienced by young households). The years refer to the respective SCF survey waves. Observations are weighted with SCF sample weights.

A key implication of the experience hypothesis is that differences in the level of risk taking between individuals should be correlated with differences in their life-time experiences. For example, after years of low stock-market returns, e.g., after the recessions of the 1970s and early 1980s, the stock-market participation rate of young people should be lower relative to that of old people (who have also experienced better returns in their lifetime) than after boom years in the stock market, e.g., in the 1960s when older individuals at the time still had the memory of the Great Depression and hence a worse average experience than young investors in their lives so far. A simple scatter-plot of differences in stock-market participation between old and young against differences in experienced stock market returns (Figure 1) confirms this pattern in the raw data. In our main analysis, we test whether these differences persist when we use a broad range of risk-taking proxies, use a more sophisticated weighting scheme for

recent and distant experiences, and include a wide range of controls for demographics, wealth, income, and other variables.

We use repeated cross-section data on household asset allocation from the Survey of Consumer Finances (SCF) from 1960-2007, and construct four measures of risk-taking: (i) the responses to a survey question about individuals' willingness to take financial risk; (ii) stock-market participation; (iii) bond-market participation, (iv) the proportion of their liquid assets invested in stocks. All four measures are likely to reflect a mixture between risk preferences and beliefs about future payoffs on risky investments.

We relate these measures of risk-taking to households' experienced histories of stock and bond returns. For each household at each SCF survey date, we calculate the annual real returns of the U.S. stock-market and of long-term government bonds since the birth year of the household head. While individuals' true "experiences" of past returns presumably differ depending on previous investments, interest in economic matters, and other unobservables, stock-market and bond-market returns likely have substantial positive correlation with their experiences of financial asset payoffs. In our estimation, we allow recent observations and those early in life to carry different weights in influencing current risk-taking. In other words, we let the data simultaneously determine how individuals weight past observations and how strongly this weighted average of past return observations (which we label "experienced return") affects risk-taking.

We find that households' risk taking is strongly related to experienced returns. Households with higher experienced stock-market returns express a higher willingness to take financial risk, participate more in the stock market, and, conditional on participating, invest more of their liquid assets in stocks. In addition, households with higher experienced bond returns are more likely to participate in the bond market. The estimates of the weights that individuals apply to their past experiences are similar for all four risk-taking measures. More recent experiences always receive higher weights, and thus have a stronger influence on risk-taking than those early in life, but even returns experienced decades earlier still have some impact for older households. Our estimates suggest that young individuals, with short life-time histories, are particularly strongly influenced by recent data.

To obtain some insight to what extent the effect of experiences on risk-taking could reflect effects on beliefs about future risky asset payoffs rather than an effect on risk preferences, we examine micro data on stock return expectations from the UBS/Gallup survey from 1998 to 2007, again controlling for time effects and age effects. We find that an increase in the experienced stock return by 1 percentage point (pp) is associated with an increase in the stock return expected by the respondent during the next year of between 0.4 and 0.8 pp. Thus, the evidence is consistent with a beliefs channel, but it does not rule out that effects of experiences on risk preferences could exist as well.

All of our estimations control for year effects, age effects, wealth and income. Year effects remove time trends or any aggregate effects, in particular a mechanical positive relation between recent stock returns and households' stock allocation due to market clearing.¹ As illustrated in Figure 1, our identification of the experience effect comes from cross-sectional differences in risk-taking and in macroeconomic histories, and from changes of those cross-sectional differences over time. Age effects allow us to distinguish our results from life-cycle effects, e.g., possible increases in risk aversion with age or the effects of the absence of labor income in retirement. The inclusion of wealth and income controls addresses the possibility that a positive correlation between past returns and current wealth explains the relation between experienced returns and current risk taking if risk aversion is wealth-dependent. Moreover, to the extent that unobserved differences in wealth remain, they are unlikely to explain all four of our risk-taking measures. Prior literature finds significant wealth effects only for stock-market participation, (see, e.g., Vissing-Jorgensen 2003), but not for the risky asset share of stock-market participants (Brunnermeier and Nagel 2008) and elicited risk tolerance (Sahm 2007). Our results also hold when retirement account holdings are excluded from the asset holding measures.

A major advantage of our methodology is that we are able to simultaneously control for age and time effects. Previous work, which has tried to identify cross-cohort differences in risk-taking with cohort dummy variable regressions (see, e.g., Ameriks and Zeldes 2004) faced the problem that cohort effects

¹ Holding the supply of stocks fixed, the average portfolio share invested in stocks increases when aggregate stock market prices increase and, hence, past returns are high.

cannot be separated from age and time effects due to collinearity of age, time, and cohort (see, e.g., Heckman and Robb 1985, and Campbell 2001). Our identification strategy, in contrast, does not rely on estimating cohort effects. The experience hypothesis predicts a positive relationship between the experienced return variable and risk taking. Since experienced returns are not collinear with age and time effects, we can control for age and time effects simultaneously.²

The fall in the stock market in 2008 can be used to illustrate the economic magnitudes of the experience effects implied by our estimates. The real return of the S&P 500 index in 2008 was about -36%. These large negative returns strongly altered investors' (weighted) life-time average returns, and the effect was strongest for young investors. For example, compared with the counterfactual benchmark of a return equal to 8.2%, the 2008 downturn lowered the experienced return of a 30-year old by about 4.0 pp, while the experienced return of a 60-year old was lowered by roughly 2.0 pp. According to our estimates, this should lower 30-year olds participation rate, everything else equal, by about 10 pp (compared with an overall participation rate for this age in 2007 of about 54.6%), whereas the effect on the participation rate of 60-year olds should be half as big, approximately 5 pp.³ Our results also imply how long-lasting the effects of the crash will be. According to our estimates, the 2008 return receives a weight of about 8.9% in the experienced return of someone who is 30 years old in 2009. In 2019, when this individual is 40 years old, the weight on the 2008 return will be reduced to 4.0%, and a further 20 years later to 2.0%. Hence, after 30 years most of the effect has faded away.

Our paper connects to several strands of literature. A related idea in the literature on learning in games is the notion of reinforcement learning, which posits that subjects' choice of actions strongly depends on the payoffs they obtained from the same actions in the past, even if circumstances (beliefs about other players' behavior and hence predicted payoffs) have changed. Experimental evidence

² Moreover, since experienced returns vary not only across, but also within cohorts over time as each cohort experiences new returns every year, we can include an almost full set of cohort dummies and therefore control for any omitted variable that has cohort-level variation. See Supplementary Appendix G.

³ As a note of caution, hypothetical counterfactual of "no 2008 market crash" holds everything else equal and does not consider the effect on asset prices in general equilibrium that would arise if the level stock market participation changed. In particular, such changes could feed back into changes in participation rates.

indicates that reinforcement learning (Erev and Roth 1998; Camerer and Ho 1999) is important in understanding players' choices. Related evidence from financial decision making is provided by Kaustia and Knüpfer (2008) who find that the returns investors experience on their own investments in initial public offerings (IPO) are positively related to their future IPO subscriptions. Choi et al. (2009) report that high personally experienced returns in 401(k) accounts induce higher 401(k) savings rates in the future.

Other papers include circumstantial evidence that young people are influenced more strongly by recent experiences than older individuals, consistent with the life-time experience hypothesis. Greenwood and Nagel (2009) show that young mutual fund managers chose higher exposure to technology stocks in the late 1990s than older managers, consistent with our finding that young individuals' allocation to stocks is most sensitive to recent stock-market returns. In a similar vein, Vissing-Jorgensen (2003) shows that young retail investors with little investment experience had the highest stock-market return expectations during the stock-market boom in the late 1990s. Amromin and Sharpe (2009) analyze microdata on stock market return expectations and find that individuals expect higher returns in times of booms than in times of recessions. Piazzesi and Schneider (2006) report that after the high inflation years in the late 1970s, younger individuals expected higher future inflation than older ones. Malmendier and Tate (2005) find that corporate managers who are born in the 1930s ("depression babies") shy away from external sources of financing, and Graham and Narasimhan (2004) find that those who experienced the Great Depression as managers choose a more conservative capital structure with less leverage.

Our evidence that life-time macroeconomic experiences affect risk-taking at the micro-level suggests that movements in the average consumer's macroeconomic experiences could also affect risk-taking in the aggregate, and hence asset prices and the macroeconomy. Along these lines, Cogley and Sargent (2008) propose to explain the equity premium with a model that assumes that the Great Depression created a long-lasting shift towards pessimistic beliefs, as suggested by Friedman and Schwartz (1963).

II. Data and Methodology

The key variables for our analysis are several measures of risk-taking from household microdata as dependent variables and historical stock and bond market returns as explanatory variables. Since our household data extends back to the 1960s, and we include individuals up to age 74 in our sample, we need stock and bond return data stretching back to the late 19th century so that we can observe returns all the way back to birth for every individual. We obtain data on the annual real returns of the S&P500 stock market index going back to 1871 from Shiller (2005)⁴, and we calculate annual real bond returns from a total return index of 10-year U.S. Treasury bonds provided by Global Financial Data, and the CPI inflation rate from Shiller (2005). Unless otherwise noted, returns are always measured in real terms.

A. Survey of Consumer Finances

Our source of household-level microdata is the Survey of Consumer Finances (SCF), which provides repeated cross-section observations on asset holdings and various household background characteristics. Our sample has two parts. The first one is the standard SCF from 1983 to 2007, obtained from the Board of Governors of the Federal Reserve System and available every three years. The second source is the precursor of the “modern” SCF, obtained from the Inter-university Consortium for Political and Social Research at the University of Michigan. The precursor surveys start in 1947, partly annually, but with some gaps. The data before 1960 contains information in stock holdings in some years, but age is measured in 5 or 10-year brackets, which would make our measurement of experienced returns imprecise, particularly for younger individuals. For this reason, we start in 1960 and use all survey waves that offer stock-market participation information, i.e., the 1960, 1962, 1963, 1964, 1967, 1968, 1969, 1970, 1971, and 1977 surveys. We briefly describe the key variables here. More details are available in the Supplementary Appendix A.

⁴ The S&P index series consists of the S&P Composite index in the early part of the series and the S&P500 index in the later part. We thank Bob Shiller for providing the data on his website.

Our first risk-attitude measure is individuals' elicited willingness to take financial risk. In the 1983 and 1989-2007 survey waves, respondents are asked which of the following statements comes closest to describing the amount of financial risk that they are willing to take when they save or make investments: (1) not willing to take any financial risk; (2) take average financial risks expecting to earn average returns; (3) take above average financial risks expecting to earn above average returns; (4) take substantial financial risks expecting to earn substantial returns. We code the answer as an ordinal variable with integer values from 1 to 4, where a value of four indicates the highest risk tolerance. For ease of reference, we refer to the measure as "elicited risk tolerance," although one should not view this measure as a clean measure of risk tolerance (in the Arrow-Pratt sense) distinct from beliefs.⁵ We also note that we cannot interpret the measure in a cardinal sense since individuals may differ in how they interpret the available options quantitatively, e.g., "substantial" or "above average" risks and returns. The survey answers may also differ from interviewees' actual risky choices. Prior literature documents, however, that the measure predicts individual willingness to take risks, e.g., households' allocation to risky assets (Faig and Shum 2006) and differences in risky human capital investments and in wage growth (Shaw 1996). Dohmen et al. (2009) find that a similar financial risk tolerance measure in the German Socio-Economic Panel is strongly related to financial risk-taking, and a simple risk-taking measure of this kind also is a strong predictor of risky behavior in a real-stakes lottery field experiment. In our analysis, using both the elicited risk tolerance measure and direct measures of asset allocation ameliorates concerns about the connection between self-reported risk tolerance and actual behavior.

The second measure is a binary variable for stock-market participation, available from 1960-2007 in each survey wave of our sample. It indicates whether a household holds more than zero dollars worth of stocks. We define stock holdings as the sum of directly held stocks (including stock held through investment clubs) and the equity portion of mutual fund holdings. In our main tests, we also include stock holdings in retirement accounts (e.g., IRA, Keogh, and 401(k) plans). For this purpose, we need to impute

⁵ For example, an individual with optimistic beliefs about future risky asset returns might answer that she is willing to take substantial financial risk *because* she expects to earn very high returns.

the stock component of retirement account holdings from the total amount in these accounts in years 1983 and 1986. From 1989 to 2004, the SCF offers only very coarse information on the allocation of retirement assets (mostly stocks, mostly interest bearing, or split), and we follow the conventions of the SCF in assigning portfolio shares. Supplementary Appendix A provides more details, and it also reports robustness checks that exclude retirement account holdings from the analysis.

Our third measure of risk taking is a binary variable for bond-market participation, available from 1960-2007, with the exception of 1971, which indicates whether a household holds more than zero dollars worth of long-term bonds. Investments in bonds, even those in default-free government bonds, are risky in real terms because of unexpected inflation. We define bond holdings as the sum of direct holdings of government bonds and corporate bonds, tax-free mutual fund holdings, and, from 1989 onwards, the bond share of non-money market mutual funds. Our definition of bond holdings does not include retirement account holdings, because the SCF does not separate bonds from holdings of short-term instruments (e.g., money market funds) in retirement accounts.

Our fourth measure of risk taking is the fraction of liquid assets invested in stocks (directly held stocks plus the equity share of mutual funds), which can be calculated in all surveys from 1960-2007, with the exception of 1971. Liquid assets are defined as stock holdings plus bonds plus cash and short-term instruments (checking and savings accounts, money market mutual funds, certificates of deposit).

As a control variable for income we use total family income. All income, wealth, and asset holdings variables are deflated into September 2007 dollars using the consumer price index (CPI-U until 1997 and CPI-U-RS thereafter). Following previous SCF literature, we eliminate observations that are likely to be miscoded and households for which a meaningful asset allocation measure does not exist because they do not have any significant liquid asset holdings.⁶ Specifically, we require that households have at least \$100 of liquid assets outside of retirement accounts and annual family income greater than

⁶ For example, Dynan, Skinner, and Zeldes (2002) exclude households with income below \$1,000. Carroll, Dynan, and Krane (2003) exclude households in the top and bottom 0.1 percent of wealth and income.

\$100 (both in September 2007 dollars). We also require that the household head is more than 24 years and less than 75 years old. Our results are robust to using the full sample.

The 1983-2007 waves of the SCF oversample high-income households. The oversampling provides a substantial number of observations on households with significant wealth holdings, which is helpful for our analysis of asset allocation, but could also induce selection bias. In our main tests, we weight the data using SCF sample weights⁷ which undo the overweighting of high-income households and which also adjust for non-response bias. The weighted estimates are representative of the U.S. population.

We also adjust standard errors for multiple imputation. From 1989 onwards, the SCF employs a multiple imputation technique to impute missing values from other information in the survey, and to disguise observations that could potentially reveal the identity of the respondent (see Kennickell 2000). The data set contains five complete copies (“implicates”), and only imputed values vary across implicates to represent the sampling uncertainty inherent in the imputation. To adjust the standard errors for this uncertainty, we follow the method of Rubin (1987) which we describe in more detail in the Supplementary Appendix B.

B. Methodology

Our objective is to investigate the relationship between risk-taking and long-term return experiences. We want to allow for the possibility that experiences in the distant past have a different influence than more recent experiences. For example, the memory of past returns might fade away as time progresses. Alternatively, experiences at young age (perhaps conveyed by parents) might be particularly formative and have a relatively strong influence on individuals’ decisions even much later in life. We aim to allow for both possibilities. Such a flexible estimation faces some hurdles, however. If we simply included separate explanatory variables for each past year of return experience (back to the year of birth,

⁷ The SCF sampling weights are equal to the inverse of the probability that a given household was included in the survey sample, based on the U.S. population, adjusted for survey non-response. Following Poterba and Samwick (2001), we normalize the sample weights each year so that the sum of the weights in each year is the same.

for example), it would be impossible to estimate the large number of coefficients on those past returns with any meaningful precision. Moreover, the number of explanatory variables would differ across households depending on their age.

To solve both problems, we summarize experienced returns as a weighted average. We use a parsimonious specification of weights that introduces only one additional parameter but is flexible enough to allow the weights to decline, be constant, or increase with distance in time since the return was realized. In this way, we can let the data speak which weighting scheme works best in explaining households' risk-taking. Specifically, for each household i in year t , we calculate the following weighted average of past asset returns,

$$A_{it}(\lambda) = \sum_{k=1}^{age_{it}-1} w_{it}(k, \lambda) R_{t-k}, \text{ where } w_{it}(k, \lambda) = \frac{(age_{it} - k)^\lambda}{\sum_{k=1}^{age_{it}-1} (age_{it} - k)^\lambda}, \quad (1)$$

where R_{t-k} is the return in year $t-k$. In our main specification, we include returns as far back as the household head's birth year. The weights (w_{it}) depend on the age of the household head at time t (age_{it}), how many years ago the return was realized (k), and a parameter λ , which controls the shape of the weighting function. We estimate λ from the data.

To illustrate the shape of the weighting function, Figure 2 plots the weights $w_{it}(k, \lambda)$ for three values of λ for a household head who is 50 years old today as a function of how many years before a return was realized, i.e. as a function of k . If $\lambda < 0$, then the weighting function is always increasing and convex as the time lag k approaches age_{it} . In this case returns close to birth receive a higher weight than more recent returns. If $\lambda = 0$, we have constant weights and $A_{it}(\lambda)$ is a simple average of past returns since birth. With $\lambda > 0$ weights are decreasing in the lag k (concave for $\lambda < 1$, linear for $\lambda = 1$, and convex for $\lambda > 1$).

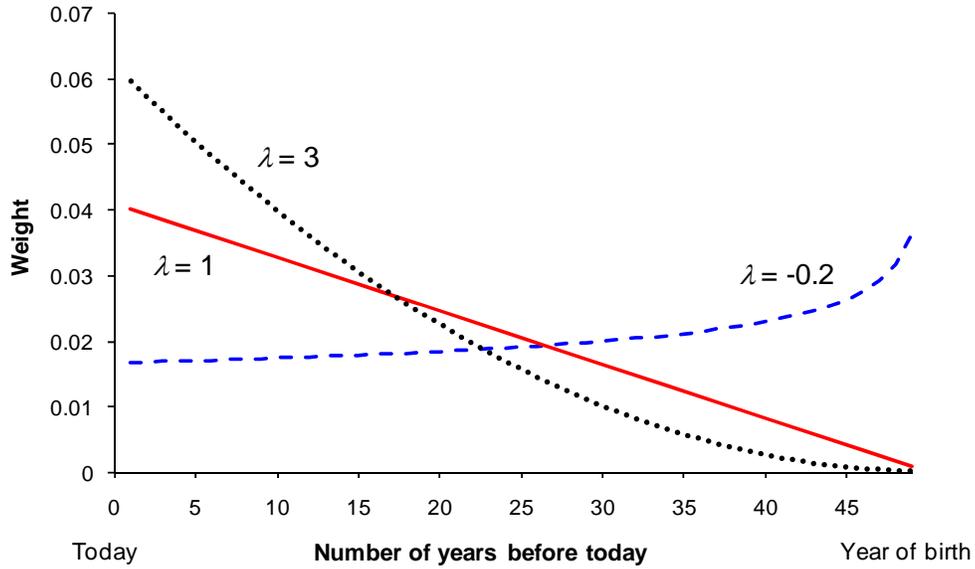


Figure 2: Weights on experienced returns implied by different values of λ for a 50-year old household head.

As the figure shows, the weighting function is quite flexible in accommodating different weighing schemes. The weights can be monotonically increasing, decreasing, or flat. We also experimented with quadratic weighting functions that allow “humps” or U-shaped weights, or a step function (see Supplementary Appendix F), but we did not find evidence that non-monotonicities are important. While the true weighting function may feature more complex weighting patterns, our restriction to a parsimonious one-parameter function biases the estimation against finding any significant effect of experienced returns on risk-taking.

As an illustration of how we simultaneously estimate the weights and individuals’ sensitivity to experienced returns calculated with those weights, consider the following generic regression model, with experienced returns $A_{it}(\lambda)$ and a vector of control variables x_{it} as the explanatory variables:

$$y_{it} = \alpha + \beta A_{it}(\lambda) + \gamma' x_{it} + \varepsilon_{it} \quad (2)$$

We simultaneously estimate β and λ . $A_{it}(\lambda)$ is a non-linear function of the weighting parameter λ , and hence non-linear estimation methods are required. For regression models, we choose β and λ to minimize

the sum of squared residuals; for Probit models, we choose them to maximize the likelihood. To ensure we are finding the global optimum, we first estimate the model on a tightly spaced grid of values for λ .⁸ We then choose the estimates that resulted in the lowest sum of square (or highest likelihood) as an initial guess for further numerical optimization.

The parameter β measures the partial effect of $A_{it}(\lambda)$ on y_{it} , i.e., conditional on the weighting parameter λ , it tells us how much y_{it} changes when $A_{it}(\lambda)$ changes, holding everything else equal. Given λ and the age of a household, one can calculate the weights $w_{it}(k, \lambda)$ as in Eq. (1). Multiplying weight $w_{it}(k, \lambda)$ with β yields, for a household of that age, the partial effect of a return experienced k years ago on the dependent variable. As an example, if $\lambda = 0$, then all returns in the household head's history since birth are weighted equally, and so their partial effects are all equal to their weight ($1/\text{age}_{it}$) times β .

Where we set the starting point for the experienced return calculation is of little importance for our results. If setting the starting point at birth is “too early” in the sense that individuals are not much influenced by experiences early in their lives, our weighting function can accommodate this with weights that decline relatively fast. If the starting point is “too late” in the sense that individuals are also influenced by observations realized prior to their birth (e.g., through their parents and social network), then setting the starting point earlier than birth could only improve the explanatory power of weighted average returns compared with our specification. The Supplementary Appendix reports some tests in which we vary the starting point to 10 years before or 10 years after birth, and find that this has little effect on our results.

C. Summary Statistics

Table I provides some summary statistics on our sample. Panel A includes all households that satisfy our sample requirements. Panel B restricts the sample to stock-market participants, i.e., households that have at least \$1 in stocks or mutual funds. Panel C restricts the sample to bond-market participants,

⁸ Given a value for the weighting parameter λ , the regression model is linear. (The probit model is still non-linear due to the non-linear transformation into probabilities.)

i.e., households that have at least \$1 directly invested in bonds. Comparing Panels A and B, we see that stock-market participants tend to be wealthier than the average household. For example, the median holding of liquid assets is \$11,642 in the full sample, but \$51,883 in the sample of stock-market participants. Panel C shows that bond-market participants are also wealthier, though less than stock-market participants, with median liquid assets of \$28,735. The pattern is similar for median income.

As Panel A shows, 38.4% of households participate on average in the stock market in the 1960-2007 period. These rates represent the U.S. population (not the SCF sample) since we apply the SCF sample weights.¹⁰ The bond-market participation rate is similar to the stock-market participation rate. The remaining two risk-taking measures show considerable dispersion across households. The proportion of liquid assets invested in stocks in Panel B has 10th and 90th percentiles of 7.1% and 90.2%. The 10th and 90th percentiles for elicited risk tolerance in Panel A are 1.0 and 3.0, respectively. It is noteworthy that mean elicited risk tolerance is higher for the stock-market participants in Panel B (2.132) than for the full sample in Panel B (1.890) and lies in the middle for bond market participants in Panel C (2.029). That is, the elicited risk-aversion measure is indeed correlated with households' actual attitudes towards financial risk-taking as revealed by their participation choices.

Our main question of interest is whether the variation in risk-taking measures across households is related to experienced stock and bond returns. To get a sense of the variation in these experienced returns for the households in our sample, we calculate the weighted average returns, $A_{it}(\lambda)$, from Eq. (1), for both stock and bond returns, setting $\lambda = 1.25$, which is in the ballpark of the estimates of λ that we find later. As Panel A shows, the 10th and 90th percentile for the experienced (real) stock return are 6.4% and 11.6% in the 1960-2007 sample. The 10th and 90th percentile for experienced (real) bond returns are -0.2% and 5.0%. Thus, over our sample period, experienced bond returns are as volatile in real terms as experienced stock returns. The amount of variation in experienced returns is similar for a range of values around the chosen values for λ . For example, with $\lambda = 1.00$ and $\lambda = 1.50$, boundaries of the interval that

¹⁰ The actual proportion of stock holders in the SCF is higher because high-income households are oversampled. This explains why the number of observations in Panel B is higher than 38.4% of the number of observations in Panel A.

contains most point estimates we obtain subsequently in our estimation, we get differences between the 10th and 90th percentile of 4.9 pp and 5.6 pp for real stock-market returns, respectively.

III. Results

A. Elicited risk tolerance

We start by relating experienced stock-market returns to elicited risk tolerance. We use y_{it} to denote the categorical SCF risk-aversion measure. It has four distinct categories, $y_{it} \in \{1, 2, 3, 4\}$. We model the cumulative probability of these ordinal outcomes with an ordered probit model

$$P(y_{it} \leq j | x_{it}, A_{it}(\lambda)) = \Phi(\alpha_j - \beta A_{it}(\lambda) - \gamma' x_{it}) \quad j \in \{1, 2, \dots, 4\}, \quad (3)$$

where $\Phi(\cdot)$ denotes the cumulative standard normal distribution function, the α_j denote the cutoff points that must be estimated ($\alpha_1 < \alpha_2 < \alpha_3 < \alpha_4 = \infty$), and x_{it} is a vector of control variables and includes income controls (log income, log income squared), demographics controls (the number of children and its square, dummies for retirement, completed high school education, completed college education, marital status, race, and for having a defined benefit pension plan), age dummies, and year dummies. We also control for the level of liquid assets held by the household (log liquid assets and log liquid assets squared, both interacted with year dummies to allow year-specific slopes). $A_{it}(\lambda)$ is the experienced stock-market return. Unlike the standard ordered probit model, $\Phi(\cdot)$ does not map a linear function of explanatory variables into the response probability P, because $A_{it}(\lambda)$ is a non-linear function of the weighting parameter λ .

We estimate the model with maximum likelihood to obtain estimates of β , λ , and γ and the cutoff points. The coefficient vector β does not have a direct economic interpretation. To interpret the results, we focus on the difference in fitted probabilities if we set the experienced return to its 10th and 90th percentile, leaving all other variables at their actual sample realizations. We calculate this difference in fitted probabilities for every observation, and then average across the whole sample. To aid in the interpretation of those differences in fitted probabilities, we will compare their magnitude to the unconditional frequencies with which individuals fall into the four elicited risk tolerance categories. As

shown in Table II, only few of them fall into the highest risk tolerance category 4, and the highest share of more than 40% is accounted for by category 2.

Before showing the results, it is useful to reiterate two identification issues. First, our method does not rely on estimating cohort effects. If we wanted to estimate unrestricted cohort effects, we would face the problem of non-separability of cohort, age, and time effects without further restrictions. Instead, the experience hypothesis predicts that a specific variable (experienced stock returns) is positively related to risk taking, allowing us to control for age and time effects at the same time. Moreover, this explanatory variable is predicted to generate variation in risk-taking not only across but also within cohorts as they experience new return realizations over time.

A second important identification issue is reverse causality. For example, if investors' risk aversion is time-varying for reasons other than variation in macroeconomic experiences, past stock market returns and current risk aversion could be mechanically correlated: stock prices rise when investors become less risk averse, and drop when investors risk aversion rises. This reverse-causality concern is addressed by our identification strategy. The year dummies absorb all aggregate time effects including variation in average risk aversion. The effect of experienced stock returns is therefore estimated from cross-sectional differences in risk taking and variation of those cross-sectional differences over time, but not from aggregate time-variation. For our other measures of risk-taking, which we consider below, year dummies also absorb all other unobserved aggregate factors that might affect stock and bond prices and, hence, simultaneously change past returns and investors' current aggregate allocation to stocks and bonds (through market clearing).

The inclusion of year dummies, and hence the focus on cross-sectional differences, also means that the null hypothesis implicit in our approach is that households are influenced in their risk-taking by all historical data and do not place higher weights on observations realized during their life-time than on data realized before they were born. In this case, cross-sectional differences in risk taking would not be correlated with cross-sectional differences in our experienced returns variable, which implies that $\beta=0$ under this null hypothesis.

Table II presents the results of the ordered probit model. We show the estimates of the parameters of interest (β and λ) at the top of the table, and the fitted probability differences for the experienced returns at the 10th and 90th percentile at the bottom.¹¹ Standard errors robust to misspecification of the likelihood function are shown in parentheses, and are adjusted for multiple imputation in the SCF.¹² Column (i), estimated on the 1983-2007 sample, shows that higher experienced stock-market returns increase the probability that risk tolerance is in the high categories (3 and 4), have little effect on the probability of being in category 2, and decrease the probability that the reported risk tolerance is in the lowest category (category 1). Thus, stock-market returns experienced in the past have a significant and positive effect on risk tolerance. As column (ii) shows, adding the liquid assets controls has little effect on the estimates.

The economic magnitudes are sizeable. For example, in column (ii), going from the 10th to the 90th percentile of experienced stock returns implies, on average, a 10.1 percentage points (pp) lower probability of being in the lowest risk tolerance category, and a correspondingly higher probability of being in the higher risk tolerance categories. Compared with an unconditional probability of 36.3% of being in the lowest risk tolerance category, this implied change by 10.1% pp is clearly economically significant.

The estimate of 1.470 (s.e. 0.294) for the weighting parameter λ in column (ii) implies that more recent returns are weighted more heavily, but also that even returns experienced many years in the past still affect households' level of risk tolerance (compare with Figure 2). Of course, there is a substantial standard error around the point estimate, but weights that are increasing with the time lag ($\lambda < 0$) are clearly ruled out. For older individuals, the estimates imply non-negligible weights of returns observed several decades earlier. Apparently, the memory of these early experiences fades away only slowly.

¹¹ The unreported coefficients of the control variables have the sign and magnitude that one would expect given the prior literature. We report the control variable coefficients in the Supplementary Appendix, Table A.2.

¹² See Section B in the Supplementary Appendix. Clustering by cohort or clustering by year does not have a material effect on our estimates.

B. Stock-market Participation

For our second estimation, the effect of life-time average returns on stock-market participation, we estimate the following probit model,

$$P(y_{it} = 1 | x_{it}, A_{it}(\lambda)) = \Phi(\alpha + \beta A_{it}(\lambda) + \gamma' x_{it}), \quad (4)$$

where the binary indicator y_{it} equals 1 if the stock holdings of household i at time t are greater than zero. We estimate the model with maximum likelihood. We are interested in the effect of experienced returns, $A_{it}(\lambda)$, on the probability of stock-market participation, and we focus again on average differences in fitted probabilities obtained from setting the experienced return variable to its 10th and 90th percentiles.

The vector x_{it} includes the same income and demographics controls as in the ordered probit model above. Controlling for liquid assets is particularly important in this context since a standard model with fixed per-period participation costs predicts that stock-market participation is positively related to the level of liquid assets (Vissing-Jorgensen 2003), and experienced stock returns are likely to be positively correlated with current liquid assets.

Columns (i) and (ii) in Table III report the estimates from our probit model. As shown in Column (ii), the life-time average returns have a positive and highly significant effect on stock-market participation. A change from the 10th to the 90th percentile of experienced stock implies an increase of about 14.6 pp in the probability that a household participates in the stock market. Thus, the stock-market return experience of different cohorts appears to have a large effect on stock-market participation. The fitted probability difference is quite similar in column (i) without the liquid assets controls.

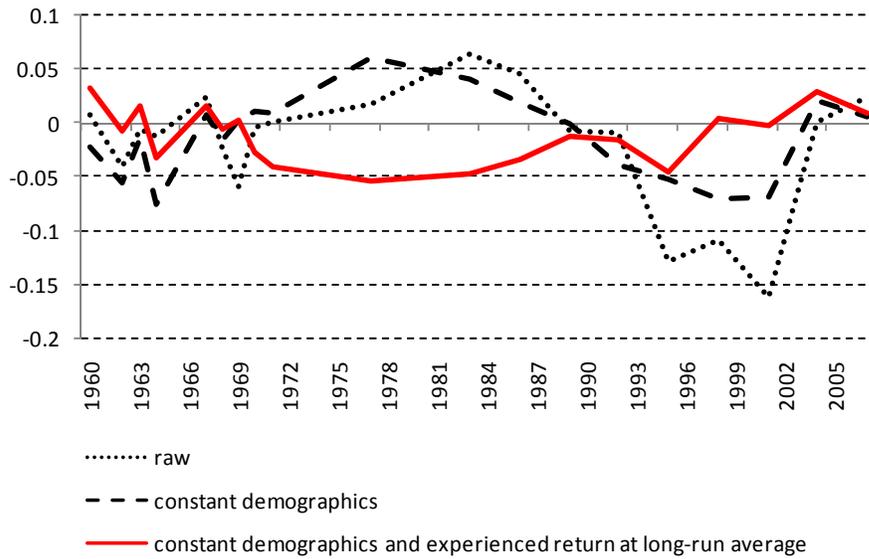
As with the previous measure, elicited risk tolerance, the estimate of 1.698 (s.e. 0.206) for the weighting parameter λ implies that households' stock-market participation decisions are affected by returns many years in the past, but rules out weights that are increasing with the time lag ($\lambda < 0$). The weighting parameter is remarkably similar to the estimate from the elicited risk-tolerance model in Table II, even though the latter one is a self-assessment by the respondent and thus not necessarily similar to one based on asset holdings. Yet, a significant part of the variation in both risk-taking measures can be

traced to variation in experienced real stock-market returns, with roughly similar weights on the history of past returns.

To provide additional perspective on the economic magnitudes, we conduct a simple counterfactual exercise. We take the point estimates from column (ii), and we calculate the fitted probabilities for each household in each survey year when all control variables, except age and year dummies, are set to their full sample averages in the corresponding age group (< 40 , $40-49$, $50-59$, ≥ 60). We label the (counterfactual) participation rates obtained from these fitted probabilities as the “constant demographics” participation rate. Next, we perform a similar calculation, but in addition to holding control variables constant over time, we also set experienced returns equal to the average stock market return since 1871, the first available year in our returns data set. This counterfactual exercise thus imagines households that consider the full return history, with equal weights for each year, going all the way back from the year prior to the survey year to 1871, without placing higher weight on life-time experiences.

Figure 3 presents the results. Panel (a) plots the stock market participation rates of the old (age ≥ 60) minus the participation rate of the young (age < 40) age group, for raw data (dotted), constant demographics (dashed), as well as constant demographics and experienced returns set to the long-term average since 1871 (solid). The raw data plot reveals big differences between participation rates of young and old in the early 1980s, when young households had much lower participation rates than the old age group, and in the late 1990s, when young households had much higher participation rates. Setting experienced returns equal to the long-term average since 1871 completely reverses the big difference in the early 1980s, and, together with constant demographics, eliminates the difference between young and old in the late 1990s. The figure only provides an incomplete picture as it shows the difference between only two relatively coarse age groups, but it provides at least rough illustration of the substantial magnitude of experienced return effects on stock market participation in the cross-section.

(a) Difference in stock market participation rates: Old (age ≥ 60) minus young (age < 40).



(b) Overall stock market participation rate

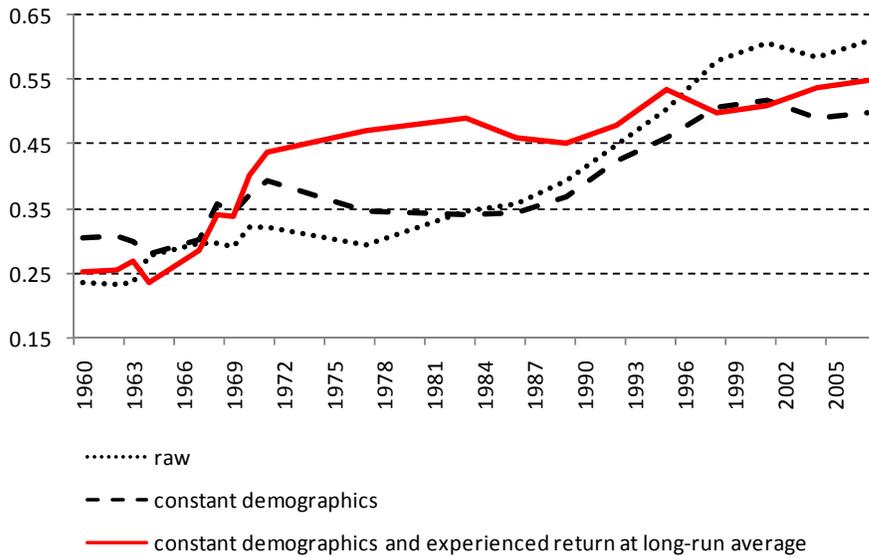


Figure 3: Counterfactual stock market participation rates. Figure (a) shows the difference in average fitted stock market participation probabilities between young (age < 40) and old (age ≥ 60) (old minus young) from the stock market participation probit model of Table III, column (ii), when all controls, including liquid assets and income, are set to their full-sample average within age groups (< 40 , 40-49, 50-59, ≥ 60) (dashed line) and when, in addition, the experienced stock market return is set to the average annual return since 1871 until the year prior to the survey year (solid line), compared with the participation rate in the raw data (dotted line). Figure (b) plots the average fitted participation probabilities in the whole sample. Observations are weighted with SCF sample weights.

Panel (b) plots the overall participation rates for the whole sample. As a note of caution, this counterfactual exercise for the overall participation rate is simplistic in that it does not consider equilibrium asset-pricing implications of changing demographics and experienced returns to counterfactual values. If these variables influence risk-taking, then changing them would presumably also influence asset prices, which might then also induce changes in overall stock market participation rates. One would need a full equilibrium model to conduct a true counterfactual investigation of aggregate effects. However, our simple calculation should at least provide some indication of the economic magnitudes of changes in aggregate risk-taking induced by experienced returns.

Comparing the effect of experienced returns (solid line relative to dashed line) with the effect of demographics, liquid assets, and income (dashed line relative to dotted line), it is evident that variations over time in experienced returns are associated with changes in stock market participation rates that are at least as big as changes induced by variations in demographics controls, including liquid assets and income. The biggest impact of experienced return variation appears in the early 1980s, when the participation rate based on the long-term average return since 1871 would have been more than 10 pp higher than with the actual experienced returns. Actual experienced returns were very low at the time, due to the poor real stock market returns in the 1970s.

C. Bond-market Participation

As our third measure of risk taking, we turn to investment in long-term bonds and test how participation in bond markets is related to experienced (real) returns on long-term government bonds. We estimate the same probit model as the one we used for stock-market participation. As column (iii) of Table III shows, experienced bond returns have a positive effect on bond-market participation, very similar to the effect of experienced stock returns on stock-market participation. Adding the liquid assets controls in column (iv) slightly increases the coefficient on experienced bond returns. A change from the 10th to the 90th percentile of experienced bond returns is associated with increase of about 15.3 pp in the probability that a household participates in the bond market. With 1.106, the point estimate for λ is a little

lower than in case of stock-market participation, but the pattern of implied weights is only marginally different. Thus, bond-market participation and stock-market participation both show positive correlation with the returns that individuals' experienced over their life-times in those markets.

D. Proportion of Liquid Assets Invested in Stocks

Table IV shows the estimated effect of experienced stock returns on the proportion of liquid assets that households invest in stocks. This measure allows us to control for fixed costs of stock-market participation, which are likely to affect stock-market participation but not the share of stocks conditional on participating. We use a non-linear regression model to estimate the effect of experienced returns,

$$y_{it} = \alpha + \beta A_{it}(\lambda) + \gamma' x_{it} + \varepsilon_{it} \quad (3)$$

where y_{it} refers to the proportion of liquid assets invested in stocks. The model is nonlinear, because the experienced stock-market return, $A_{it}(\lambda)$, is a nonlinear function of λ . We estimate the model with nonlinear least-squares. Unlike in the probit model, the partial effect of $A_{it}(\lambda)$ is now equal to the parameter β and so we can assess economic magnitudes directly by multiplying β with the variation in experienced returns. The control variables are the same as in Tables II and III.

In columns (i) and (ii), experienced returns are measured as real stock returns. As column (i) shows, without the liquid assets controls, the experienced stock return has only a statistically weak positive effect on the proportion of liquid assets invested in stocks. But when the liquid assets controls are added in column (ii), the effect is stronger, both in terms of statistical significance and economic magnitude. The point estimate of 1.476 (s.e. 0.445) implies that a change from the 10th to the 90th percentile of experienced stock returns (5.2 pp) leads to an increase of about $1.476 \times 5.2 \text{ pp} \approx 7.7 \text{ pp}$ in the allocation to stocks. This finding is remarkable since it is a common result in the empirical literature on household portfolio choice that, once one restricts the sample to stock-market participants, it is hard to find *any* household characteristics that have economically significant correlations with the portfolio risky asset share (see Curcuru, Heaton, Lucas, and Moore (2009), and Brunnermeier and Nagel (2008) for

recent evidence, and the control variable coefficients reported in Supplementary Appendix D.) In light of this evidence, experienced stock-market returns emerge as one of the major factors that influence a households' willingness to bear stock-market risk.

The point estimate for λ in column (i) is 0.923 (s.e. 0.323), which suggests weights that decline roughly linearly going back in time. This estimate for λ is approximately of the same magnitude as the λ -estimates in the elicited risk-aversion model in Table II and the stock and bond market participation models in Table III. The similarity of the estimates is noteworthy since elicited risk tolerance is a measure based on a very different approach (survey question versus investment choice) and financial market participation and choice of the risky asset share conditional on participation are possibly quite distinct decisions. The similarity is reassuring for our interpretation that the all of these variables capture a common attitude to financial risks and are subject to a common influence of macroeconomic experience.

We also test how the proportion of liquid assets allocated to stocks responds to the differential returns of stocks and bonds. Assuming the perspective of an investor choosing between investment in stocks and in bonds, the experience hypothesis predicts that only if stocks performed better than bonds over the lifetime of the investor, she will increase her investment in stocks relative to bonds. Columns (iii) and (iv) of Table IV repeat the regressions of columns (i) and (ii) with experienced excess returns, measured as stock-market returns in excess of long-term bond returns. We find that experienced excess returns explain household's allocation to stocks about as well as real stock returns. The point estimates for β in column (iv) are slightly higher than in column (ii), and the estimates for λ are moderately higher, too. The results are also similar if we restrict the sample to households that participate both in stock and bond markets, i.e. those that can presumably change their allocation to both stocks and bonds relatively flexibly, with participation cost already sunk.

It might seem that slow portfolio rebalancing of households in response to stock market movements could potentially provide an alternative explanation for the positive relationship between experienced returns and the percentage allocation to stocks. That households are slow to rebalance has

been documented, for example, in Brunnermeier and Nagel (2008) and Campbell, Calvet, and Sodini (2009). However, slow portfolio rebalancing cannot explain the positive coefficient on experienced returns in our regressions in Table IV. Section C in the Supplementary Appendix reports the results of simulations of an overlapping-generations model with agents that are slow to rebalance. We run regressions similar to those in Table IV on the simulated data, and we find that experienced returns do not receive a positive coefficient. The key is the inclusion of time dummies in our regression. Without time dummies there would indeed exist a mechanical positive relationship between experienced returns and the portfolio share of stocks. But when time dummies are included, and the regression thus focuses on cross-sectional variation, this positive relationship disappears.

E. Using Stock and Bond Returns Jointly to Explain Risk-taking

As an additional test of the experience hypothesis, we compare the relative predictive power of experienced stock returns and experienced bond returns for all of our risk-taking measures, we relate all four of our risk measures simultaneously to experienced stock returns and to experienced bond returns. That is, we re-run the specifications of Table II, column (ii), Table III, columns (ii) and (iv), and Table IV, column (ii) with both experienced real stock returns and experienced real bond returns as explanatory variables. Since an estimation of distinct weighting parameters for both stock and bond returns within the same model would be too demanding on the data and would not produce statistically reliable results, we fix the weighting parameters at the values obtained in the earlier specifications that had a single type of return.

Table V reports the results. In the first three columns, labeled “Full sample,” we use all the available data, as in Tables II and III. In column (iv), where the dependent variable is the percentage share invested in stocks, the sample is restricted to stock market participants, as in Table IV. In column (v) we also explore the share invested in bonds (using only non-retirement assets in the calculation of this share, as bond allocations in retirement accounts are not available), and we restrict the sample to households that participate in both stock and bond markets. For the probit models in columns (i) to (iii) we also report the

average of the fitted probabilities at the 90th percentile of experienced returns minus the fitted probabilities at the 10th percentile of experienced returns. To save space, we only report this average difference in fitted probabilities for category 1 (low risk tolerance) in column (i). The spread in fitted probabilities for the other three risk tolerance categories combined is of the same magnitude, but with opposite sign.

We find that elicited risk tolerance is positively related to experienced stock and bond returns according to the point estimates, but the standard errors are huge and so one cannot draw definitive conclusions from the results. Evidently, the short sample available for this risk-taking measure is not sufficient to disentangle the effects of stock and bond return experiences. Stock market participation, in column (ii), is more strongly related to stock market return experiences than to bond returns, while the opposite is true for the bond market participation measure. The percentage share allocated to stocks in column (iv) is positively related to experienced stock returns and negatively related to experienced bond returns. For the bond share in column (v) the opposite is true, but the coefficients are smaller, particularly the coefficient on bond returns, and are statistically not significantly different from zero.

The results help to further address concerns about unobserved wealth effects, i.e., the alternative interpretation that the correlation of return experiences with unobserved wealth components, coupled with wealth-dependent risk aversion, explains our results. Since both past stock and bond returns should be positively related to wealth, one would expect both stock and bond returns to predict each of the risk-taking measures with the same sign and magnitude. This is not the case.

Disentangling the joint roles of stock and bond returns also provides some hints on the question whether the life-time experiences we measure affect preferences or beliefs. The results are most easily reconciled with a belief-based story: If individuals' beliefs about future returns are positively related to their return experience with this *particular* asset class, stock returns should matter most for the stock-investment-based risk-taking measures, while bond returns should matter most for bond-market participation. A simple preference-based story, instead, in which individuals' level of relative risk aversion depends on past experiences of stock and bond returns, would not predict such differential

effects of stock and bond returns on the different risk-taking measures. Only more elaborate preference-based theories, where individuals' "tastes" for different asset classes depend on their return experiences with this particular asset class, could match the last set of results.

F. Experience Effects on Stock Market Return Expectations

In an effort to further disentangle the roles of beliefs and risk preferences, we investigate how past stock return experiences relate to expectations about future stock returns. We use data on stock return expectations of households from the UBS/Gallup survey, obtained from the Roper Center at the University of Connecticut. The data set covers the period from 1998 to 2007, mostly monthly, but with a few exceptions. Details on the data set are provided in the Supplementary Appendix A.

Using these data, Vissing-Jorgensen (2003) finds time-varying differences in expectations between different age groups. Young people have the highest stock return expectations, particularly around the time of the stock market peak in 2000. This pattern could be consistent with a positive effect of experienced returns on expectations. To find out, we examine whether our experienced return variable explains some of the variation in expectations. We use two types of expectations: The stock market return expected over the next 12 months, and the return that the respondent expects to earn on his or her own portfolio over the next 12 months.

Unfortunately, the UBS/Gallup survey covers only a much shorter time period than the SCF. This short time span makes it difficult to statistically disentangle the effect of experienced returns from age effects. To reduce the burden on the estimation, we work with the weighting parameter fixed at the values we obtained in the stock market participation ($\lambda = 0.923$, Table III) or percentage allocation to stocks ($\lambda = 1.698$, Table IV) baseline specification. We regress stock return expectations on experienced returns, where experienced returns are calculated for each individual, given λ and the individual's age, using real stock market returns up to the end of the year preceding the survey year, in the same way as in our previous analysis of risk-taking measures. We weight each observation with sample weights provided by

the UBS/Gallup survey. As in our earlier analysis of SCF data, we focus on individuals with age ranging from 25 to 74.

Table VI presents the results. Columns (i) and (ii) use the expected stock market return over the next 12 months as dependent variable. This variable is only available from 1998 to 2003. Despite the very short sample, the estimate for the coefficient on experienced returns in column (i) of 0.801 is more than two standard errors away from zero. With $\lambda = 1.698$ in column (ii), the point estimate is about half as big, and only marginally significant. The magnitudes of the point estimates are substantial. They imply that a 1 pp higher experienced return translates roughly into 0.4 pp to 0.8 pp higher expected return.

Columns (iii) and (iv) repeat this analysis, but now with the expected return on the respondent's own portfolio as the dependent variable. The benefit of using this variable is a substantially longer sample that extends until 2007. The downside of this variable is that it doesn't allow us to separate expectations from risk preferences as cleanly as the stock market return expectations variable does. It is perceivable that a highly risk averse investor might choose a portfolio with a low level of risk and hence expect a low return, while a risk tolerant individual might choose a high-risk portfolio and expect a high return. Thus, differences in people's expectations about their own future portfolio return could partly stem from differences in risk preferences. However, the point estimates for the coefficient on experienced returns are very similar to those in columns (i) and (ii), which suggests that there is not much room for a strong confounding effect of risk preferences on this variable. Due to the larger sample, standard errors are considerably smaller in columns (iii) and (iv).

The findings from expectations data lend further support to the view that experienced returns affect beliefs about future asset returns, and that this is the channel through which experienced returns affect risk-taking. This does not rule out that life-time return experiences could risk preferences as well, but it shows that, at a minimum, the beliefs channel accounts for an important part of the experience effects.

G. Methodological Variations and Robustness Checks

We check the robustness of our results to several further variations in methodology. All of these (and more) additional tests are reported in detail in the Supplementary Appendix.

Financial sophistication. We interact the experienced return with financial sophistication proxies: a dummy for high levels of liquid assets and dummy for completed college education. The results do not show a clear tendency of experience effects to be stronger or weaker with higher financial sophistication.

Non-monotonicities in the weighting function. To check whether our monotonic weighting function could be misspecified, we experiment with a step function, where the steps are defined over the first, middle, and most recent third of an individual's lifespan. The results, reported in Section F of the Supplementary Appendix, indicate a pattern of weights that closely resembles to the monotonic weights produced by our weighting function.

Excluding retirement assets. The allocation to stocks in retirement accounts is probably measured with considerable error since the SCF only provides coarse allocation brackets before 2004 and no allocation information before 1989, which necessitated an imputation of retirement allocations in 1983 and 1986. Repeating our baseline estimations with retirement account holdings completely excluded, we find that the estimates are generally very similar to our baseline estimates. Also, running the estimation with retirement accounts included, but excluding the years with imputed retirement allocations (1983 and 1986) has little influence on the estimates.

Variation in starting point. In our analyses above, the starting point for life-time experiences is set at birth. This should not be a crucial assumption because our weighting function can place low or high weight on returns experienced early in life. For example, if returns realized during the first 10 or 20 years of their life do not matter much, our weighting function should be able to approximately adapt to this with a relatively high value of λ . Consistent with this intuition, we find that our results do not change much if we set the starting point at 10 years after or at 10 years before birth.

Including cohort dummies. Our main explanatory variable, the life-time weighted average return is not a constant per cohort, but instead varies over time as the cohort experiences new return realizations.

This allows us to also include cohort dummies in our specifications (as many cohort dummies as possible up to the point that age, time, and cohort dummies are not perfectly collinear) to control for unobserved cohort effects. We find that the point estimates remain similar, with the exception of bond market participation, where the β coefficient drops considerably and standard errors are quite high.

Experienced Volatility. We also test whether experienced volatility affects households' risk-taking decisions. To calculate experienced volatility, we apply the same weighting function as before, but now to calculate the weighted standard deviation instead of the weighted average of returns. To limit the demands on the estimation, we fix the weighting parameter for both experienced return and experienced volatility at the point estimate for λ obtained in the baseline specifications. We do not find a consistent and strong effect of experienced volatility on risk-taking. Experienced volatility tends to be negatively associated with percentage allocation to stocks, but for elicited risk tolerance, as well as stock and bond market participation, the estimated effect is positive, but of smaller magnitude than the experienced return effect, and with relatively standard errors. Most importantly, however, the inclusion of experienced volatility has little effect on the coefficient on experienced returns. It is possible that experience of extreme events could perhaps affect risk-taking more strongly than experience of risk measured by standard deviations. But the rare nature of extreme events, combined with the difficulty in deciding what constitutes an extreme event, means that their effects are difficult to investigate empirically within our framework, and we leave an investigation of extreme events to future work.

IV. An Aggregate Perspective

Our estimation focused on focused on cross-sectional differences in risk-taking measures in order to absorb potential confounding macro and market-clearing effects with time dummies. This does not mean that our results only have cross-sectional implications. Differences in experienced returns exist not only between different age groups at a given point in time, but also between average experienced returns in the population at different points in time. Based on our estimates from microdata, one would expect

that this variation over time in the average experienced returns of the population should also influence risk-taking in aggregate. Such experience-induced variation in households' degree of risk-taking could help to explain the puzzling property of stock prices that valuation ratios like the price/earnings (P/E) ratio exhibit large variations over time, and that this variation is driven to a large extent by changes in expected returns rather than variation in expected cash flows. A full investigation of the asset-price effects of macroeconomic experiences is beyond the scope of this paper, but we carry out a plausibility check. For this theory to be plausible, the time series of experienced return of the average investor should line up well with the time series of the price/earnings ratio of the aggregate stock market. Periods of high experienced returns (and hence high levels of risk-taking due to low risk aversion and/or optimism) should coincide with periods of high price/earnings ratios.

We calculate experienced stock-market returns for each age from 25 to 74 in each year from 1946 to 2007 based on a weighting parameter of $\lambda = 1.25$, which is roughly the average parameter estimate across all specifications with liquid asset controls in Tables II-IV. We then average the experienced returns across all ages in each year and plot the resulting average experienced returns series against the annual price-to-earnings (P/E) ratio of the S&P 500 index from Shiller (2005).¹³ We also conducted a similar calculation with liquid assets-weighted averages of experienced returns and found similar results.

Figure 4 presents the results from this exercise. Each bar represents the aggregated experienced stock-market return of U.S. investors in the corresponding year, and the line shows the P/E ratio. The two series are highly positively correlated. Periods of high equity market valuations and low subsequent returns (the 1960s and 1990s) coincide with periods when investors have high experienced stock-market returns, and periods of low valuation and high subsequent returns (1940s and early 1980s) coincide with investors having low experienced stock-market returns. While this is not a definitive proof that variations in P/E ratios and expected returns are driven by experience effects, it underscores that a theory of experience-induced variation in the degree of risk-taking is at least a plausible explanation.

¹³ Shiller's P/E series uses a ten-year moving average of earnings in the denominator

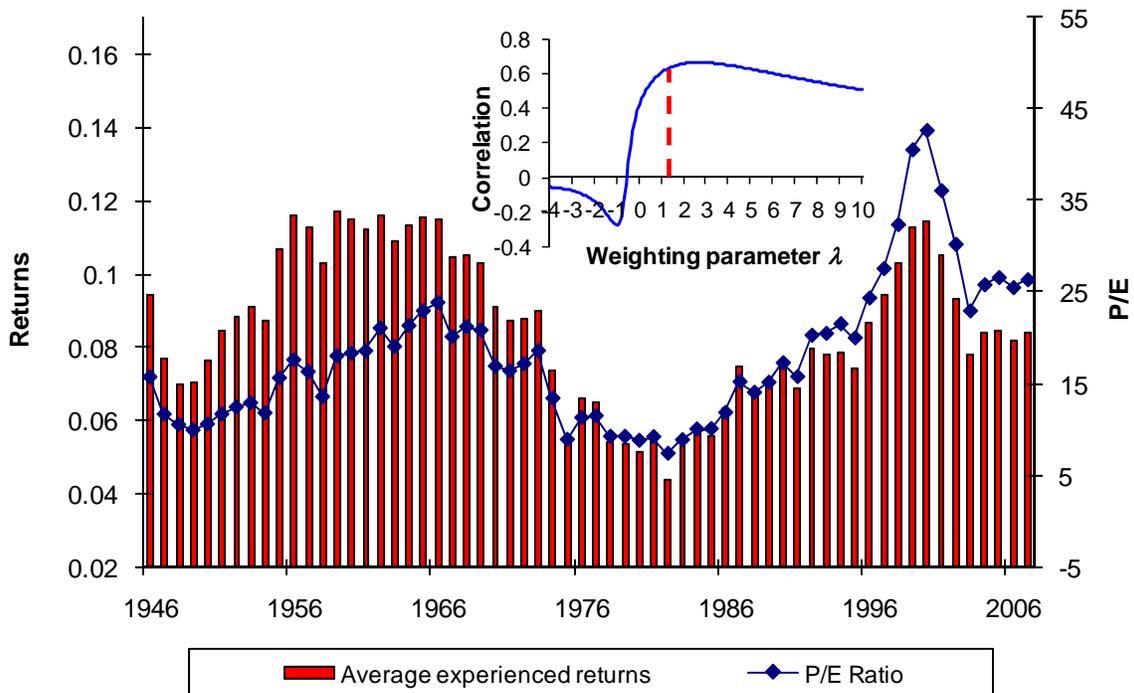


Figure 4: Average experienced real stock returns ($\lambda = 1.25$) and price-to-earnings ratio 1946-2007. The average experienced return series is calculated by averaging experienced returns across all ages at a given point in time. The small inset figure shows the correlation between the average experienced return series and the P/E ratio if the weighting parameter is varied from -4.0 to 10.0. The correlation for $\lambda = 1.25$ is indicated by the dashed line.

Note that this correlation does not mechanically reflect the well-known positive correlation between P/E ratios and past returns. We estimate the weighting parameter λ from *microdata*, where we exploit cross-sectional differences between investors' risk-taking measures. We do not use aggregate data in the estimation, and λ is not chosen to match movements in the P/E ratio over time. For example, the weighting parameter estimated from the microdata could have turned out to be strongly negative, which would mean that investors place a lot of weight on returns experienced early in life, but less on more recent returns. In that case, the average experienced return would have been uncorrelated with recent stock-market returns and the time pattern of the bars in Figures 4 would look very different. This point is demonstrated by the small inset figure in Figure 4. It plots the correlation between average experienced stock returns and the P/E ratio for different choices of the weighting parameter λ . The figure demonstrates

that the correlation between life-time average returns and the P/E ratios could easily have been smaller if the microdata-estimates of λ had turned out differently. The value of $\lambda = 1.25$ is actually close to the value that yields the maximum correlation. And the range of point estimates between 1.0 and 2.0 that we obtained in most of our estimated models all yield a high correlation of around 0.6.

The high correlation between aggregate experienced stock returns and stock-market valuation levels adds credibility to our microdata estimates, as the estimates imply plausible time-variation in aggregate demand for risky assets. Our results thus suggest the possibility that personally experienced risky asset returns affect asset prices via changes in investors' willingness to take risk. We leave a further exploration of such asset-pricing effects to future work, as the scope of the current paper is focused on estimating relationships in microdata.

V. Conclusion

Our results show that risky asset returns experienced over the course of an individual's life have a significant effect on the willingness to take financial risks. Individuals who have experienced high stock-market returns report lower aversion to financial risk, are more likely to participate in the stock market, and allocate a higher proportion of their liquid asset portfolio to risky assets. Individuals who have experienced high real bond returns are more likely to participate in the bond market. We find that individuals put more weight on recent returns than on more distant realizations, but experiences several decades ago still have some impact on current risk-taking of older households. The magnitudes of the effects are economically important. For example, an increase of 5 percentage points in the level of experienced real stock returns is associated with an increase in the probability of stock market participation of about 15 percentage points, and an increase in the percentage of liquid assets allocated to stocks of about 7 percentage points. Our results are consistent with the view that economic events experienced over the course of one's life have a more significant impact on risk taking than historical facts learned from summary information in books and other sources.

We also offer some evidence that that experiences influence risk taking seems to arise, at least partly, through a beliefs channel, as opposed to an effect on risk preferences. Using microdata on stock market return expectations, we find that higher experienced stock returns are associated with more optimistic beliefs about future stock returns. This suggest that the experience effects could be the result of individuals' attempts to learn from their experiences, albeit not by using all "available" historical data, as in standard rational and boundedly rational learning models, but by focusing on their life-time experiences. Consistent with this view, we show in follow-up work, Malmendier and Nagel (2009), that inflation expectations are influenced by individuals' inflation experiences in similar ways as risk-taking and stock return expectations are influenced by experiences of risky asset returns.

References

- Alesina, Alberto, and Nicola Fuchs-Schündeln, 2007. Good-bye Lenin (or Not?): The Effect of Communism on People's Preferences, *American Economic Review*, 97, 1507-1528.
- Ameriks, John, and Stephen Zeldes, 2004. How Do Household Portfolio Shares Vary With Age? *Columbia University Working Paper*.
- Amromin, Gene, and Stephen A. Sharpe, 2009. Expectations of Risk and Return Among Household Investors: Are Their Sharpe Ratios Countercyclical? *Board of Governors of the Federal Reserve System Working Paper*.
- Avery, Robert B., and Gregory E. Elliehausen, 1990. 1983 Survey of Consumer Finances: Technical Manual and Codebook, *Board of Governors of the Federal Reserve System Working Paper*.
- Brunnermeier, Markus, and Stefan Nagel, 2008. Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation, *American Economic Review* 98, 713-736.
- Camerer, Colin and Teck Ho, 1999. Experience-Weighted Attraction Learning in Normal Form Games, *Econometrica* 67, 827-874.
- Campbell, John Y., 2001. A Comment on James Poterba's "Demographic Structure and Asset Returns", *Review of Economics and Statistics* 83, 585-588.
- Campbell, John Y., Laurent E. Calvet, and Paolo Sodini, 2009. Fight or Flight? Portfolio Rebalancing By Individual Investors, *Quarterly Journal of Economics* 124, 301-348.
- Caroll, Christopher D., Karen Dynan, and Spencer D. Krane, 2003. Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets, *Review of Economics and Statistics* 85, 586-604.
- Cogley, Timothy, and Thomas J. Sargent, 2008. The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?, *Journal of Monetary Economics* 55, 454-476.
- Choi, James J.; Laibson, David; Madrian, Brigitte C., and Andrew Metrick, 2009. Reinforcement Learning and Savings Behavior, *Journal of Finance*, forthcoming.
- Curcuru, Stephanie, John Heaton, Deborah Lucas, and Damien Moore, 2009. Heterogeneity and Portfolio Choice: Theory and Evidence, in Y. Ait-Sahalia and L.P. Hansen, ed., *Handbook of Financial Econometrics* (North-Holland, Amsterdam, Netherlands), 337-382.
- Dohmen, Thomas; Falk, Armin; Huffman, David; Sunde, Uwe; Schupp, Jurgen; and Gert G. Wagner, 2009. Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences, *Journal of the European Economic Association*, forthcoming.
- Dynan, Karen, Jonathan Skinner, and Stephen P. Zeldes, 2002. The Importance of Bequests and Life-Cycle Saving in Capital Accumulation: A New Answer, *American Economic Review* 92, 274-278.

- Erev, Ido, and Alvin E. Roth, 1998. Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria, *American Economic Review* 88, 848-881.
- Faig, Miquel, and Pauline Shum, 2006. What Explains Household Stock Holdings?, *Journal of Banking and Finance* 30, 2579-2597.
- Friedman, Milton, and Anna J. Schwartz, 1963. *A Monetary History of the United States, 1867-1960*, Princeton University Press, Princeton, NJ.
- Graham, John R., and Krishnamoorthy Narasimhan, 2004. Corporate Survival and Managerial Experiences During the Great Depression, *Duke University Working Paper*.
- Greenwood, Robin, and Stefan Nagel, 2009. Inexperienced Investors and Bubbles, *Journal of Financial Economics* 93, 239-258.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2004. The Role of Social Capital in Financial Development, *American Economic Review*, 94, 526-556.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2008. Trusting the Stock Market, *Journal of Finance*, 63, 2557-2600.
- Hertwig, Ralph, Greg Barron, Elke U. Weber, and Ido Erev, 2004. Decisions from Experience and the Effect of Rare Events in Risky Choice, *Psychological Science* 15, 534-539.
- Kennickell, Arthur B., 1998. Multiple Imputation in the Survey of Consumer Finances, *Board of Governors of the Federal Reserve System Working Paper*.
- Kaustia, Markku, and Samuli Knüpfer, 2008. Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions, *Journal of Finance* 63, 2679-2702.
- Malmendier, Ulrike, and Stefan Nagel, 2009. Learning from Inflation Experiences, *Stanford University and UC Berkeley Working Paper*.
- Malmendier, Ulrike, and Geoffrey Tate, 2005. CEO Overconfidence and Corporate Investment, *Journal of Finance*, 60, 2661-2700.
- Nisbett, Richard E., and Lee D. Ross, 1980. *Human Inference: Strategies and Shortcomings of Social Judgment*, Englewood Cliffs, N.J.: Prentice Hall.
- Osili, Una Okonkwo, and Anna L. Paulson, 2008. "Institutions and Financial Development: Evidence from International Migrants in the United States," *Review of Economics and Statistics*, 90, 498-517.
- Piazzesi, Monika, and Martin Schneider, 2006. Inflation and the Price of Real Assets, *University of Chicago Working Paper*.
- Poterba, James M., and Andrew Samwick, 2001. Household Portfolio Allocations Over the Life-Cycle, in S. Ogura, T. Tachibanaki, and D. Wise, ed., *Aging Issues in the U.S. and Japan* (University of Chicago Press, Chicago), 65-103.

- Rubin, Donald B., 1987. *Multiple Imputation for Nonresponse in Surveys* (J. Wiley & Sons, New York).
- Sahm, Claudia R., 2007. Stability of Risk Preference, *Federal Reserve Board of Governors Working Paper*.
- Shaw, Kathryn L., 1996. An Empirical Analysis of Risk Aversion and Income Growth, *Journal of Labor Economics* 14, 626-653.
- Shiller, Robert J., 2005. *Irrational Exuberance* (second edition), Princeton University Press, Princeton NJ.
- Weber, Elke U., Ulf Böckenholt, Denis J. Hilton, and Brian Wallace, 1993. Determinants of Diagnostic Hypothesis Generation: Effects of Information, Base Rates, and Experience, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 1151-1164.
- Vissing-Jorgensen, Annette 2003. Perspectives on Behavioral Finance: Does “Irrationality” Disappear with Wealth? Evidence from Expectations and Actions, in M. Gertler and K. Rogoff, ed., *NBER Macroeconomics Annual* (MIT Press, Cambridge), 139-194.

Table I: Summary Statistics

	10 th pctile	Median	90 th pctile	Mean	Stddev	#Obs.
<i>Panel A: All households 1960 – 2007</i>						
Liquid assets	727	11,642	172,996	92,047	677,063	43,862
Income	17,049	48,718	109,336	65,764	182,221	43,862
Experienced real stock return ($\lambda = 1.25$)	0.064	0.091	0.116	0.090	0.021	43,862
Experienced real bond return ($\lambda = 1.25$)	-0.002	0.008	0.050	0.018	0.021	43,862
Stock market participation	0	0	1	0.384	0.484	43,862
Bond market participation	0	0	1	0.327	0.469	42,995
Elicited risk tolerance (1983-2007)	1	2	3	1.890	0.831	25,588
<i>Panel B: Stock market participants 1960-2007</i>						
Liquid assets	5,285	51,883	401,400	206,430	1,075,158	21,420
Income	28,370	66,525	158,828	96,813	285,908	21,420
Bond market participation	0	0	1	0.434	0.494	21,179
% Liquid assets in stocks	0.071	0.439	0.902	0.462	0.296	20,601
Elicited risk tolerance (1983-2007)	1	2	3	2.132	0.794	16,131
<i>Panel C: Bond market participants 1960-2007</i>						
Liquid assets	1,936	28,735	315,270	173,404	1,098,734	16,086
Income	24,637	58,783	134,110	84,700	264,119	16,086
Stock market participation	0	1	1	0.526	0.497	16,086
% Liquid assets in stocks	0	0.014	0.709	0.219	0.291	15,389
Elicited risk tolerance (1983-2007)	1	2	3	2.029	0.791	9,940

Notes: Stock returns and bond returns are real returns, deflated with CPI inflation rates. Wealth and income variables are deflated by the CPI into September 2007 dollars. Observations are weighted by SCF sample weights. The bond market participant sample in Panel C excludes the 1964 survey in which bond market participation information is not available.

Table II: Elicited Risk Tolerance

	(i)	(ii)
Experienced stock return coefficient β	5.378 (1.208)	6.619 (1.283)
Weighting parameter λ	1.719 (0.356)	1.470 (0.294)
Income controls	Yes	Yes
Liquid assets controls	-	Yes
Demographics controls	Yes	Yes
Age dummies	Yes	Yes
Year dummies	Yes	Yes
Average of fitted prob. at 90 th pctl. minus fitted prob. at 10 th pctl. of experienced stock return		
Risk tolerance = 1 (low)	-0.096	-0.101
[unconditional freq. = 36.3%]	(0.018)	(0.016)
Risk tolerance = 2	0.022	0.021
[unconditional freq. = 42.6%]	(0.004)	(0.006)
Risk tolerance = 3	0.050	0.052
[unconditional freq. = 16.7%]	(0.012)	(0.012)
Risk tolerance = 4 (high)	0.025	0.027
[unconditional freq. = 4.3%]	(0.009)	(0.010)
#Obs.	25,518	25,518
Pseudo R ²	0.07	0.09

Notes: Ordered probit model estimated with maximum likelihood. Sample period runs from 1983 to 2007 and excludes the 1986 survey (elicited risk tolerance not available). The experienced stock return is calculated from the real return on the S&P500 index. Liquid assets controls are log liquid assets and log liquid assets squared, both interacted with year dummies to allow for year-specific slopes. Demographic controls (coefficients not reported) include the number of children and number of children squared, as well as dummies for marital status, retirement, race, education, and for having a defined benefit pension plan. Observations are weighted with SCF sample weights. Standard errors shown in parentheses are robust to misspecification of the likelihood function and adjusted for multiple imputation.

Table III: Stock and Bond Market Participation

	Experienced stock returns and stock mkt. participation		Experienced bond returns and bond mkt. participation	
	(i)	(ii)	(iii)	(iv)
Experienced return coefficient β	6.944 (1.093)	10.139 (1.320)	8.936 (1.470)	9.488 (1.543)
Weighting parameter λ	1.900 (0.233)	1.698 (0.206)	1.323 (0.306)	1.106 (0.282)
Income controls	Yes	Yes	Yes	Yes
Liquid assets controls	-	Yes	-	Yes
Demographics controls	Yes	Yes	Yes	Yes
Age dummies	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
Average of fitted prob. at 90 th pctile. minus fitted prob. at 10 th pctile. of experienced return	0.128 (0.023)	0.146 (0.022)	0.156 (0.027)	0.153 (0.026)
#Obs.	43,660	43,660	42,793	42,793
Pseudo R ²	0.20	0.33	0.07	0.12

Notes: Probit model estimated with maximum likelihood. Sample period runs from 1960 to 2007 (excluding 1964 in the case of bond market participation). The experienced stock return is calculated from the real return on the S&P500 index. Liquid assets controls are log liquid assets and log liquid assets squared, both interacted with year dummies to allow for year-specific slopes. Demographic controls (coefficients not reported) include the number of children and number of children squared, as well as dummies for marital status, retirement, race, education, and for having a defined benefit pension plan. Observations are weighted with SCF sample weights. Standard errors shown in parentheses are robust to misspecification of the likelihood function and adjusted for multiple imputation.

Table IV: Percentage of Liquid Assets Invested in Stocks

	Experienced stock returns		Experienced excess returns of stocks over bonds	
	(i)	(ii)	(iii)	(iv)
Experienced return coefficient β	0.440 (0.395)	1.476 (0.445)	0.688 (0.397)	1.611 (0.439)
Weighting parameter λ	1.450 (1.372)	0.923 (0.323)	1.185 (0.775)	1.345 (0.391)
Income controls	Yes	Yes	Yes	Yes
Liquid assets controls	-	Yes	-	Yes
Demographics controls	Yes	Yes	Yes	Yes
Age dummies	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
#Obs.	20,247	20,247	20,247	20,247
R ²	0.07	0.10	0.07	0.10

Notes: Model estimated with nonlinear least squares on the sample of stock market participants. Sample period runs from 1960 to 2007, excluding the 1971 survey (percentage allocation not available). Experienced returns stock returns calculated from the real return on the S&P500 index and experienced excess return from the return on the S&P500 index minus the return on long-term U.S. Treasury bonds. Liquid assets controls are log liquid assets and log liquid assets squared, both interacted with year dummies to allow for year-specific slopes. Demographic controls (coefficients not reported) include the number of children and number of children squared, as well as dummies for marital status, retirement, race, education, and for having a defined benefit pension plan. Observations are weighted with SCF sample weights. Standard errors shown in parentheses are robust to heteroskedasticity and adjusted for multiple imputation.

Table V: Using Stock and Bond Returns Jointly

Dependent variable	Elicited risk tolerance	Stock mkt. participation	Bond market participation	% liquid assets in stocks	% liquid assets in bonds
Sample	Full	Full	Full	Stock market participation required	Stock and bond market participation required
Experienced stock return coeff. β_{stock}	3.422 (2.519)	9.050 (1.388)	0.829 (1.220)	1.565 (0.450)	-0.882 (0.576)
Weighting parameter for stocks λ_{stock}	1.470 [fixed]	1.698 [fixed]	1.698 [fixed]	0.923 [fixed]	0.923 [fixed]
Average of fitted prob. at 90 th pctile. minus fitted prob. at 10 th pctile. of experienced stock return	-0.052 (0.035)	0.131 (0.023)	0.016 (0.024)		
Probability of lowest risk tol. category		participation	participation		
Experienced bond return coeff. β_{bond}	8.026 (5.373)	6.490 (1.780)	9.238 (1.529)	-0.997 (0.582)	0.087 (0.606)
Weighting parameter for bonds λ_{bond}	1.106 [fixed]	1.106 [fixed]	1.106 [fixed]	1.106 [fixed]	1.106 [fixed]
Average of fitted prob. at 90 th pctile. minus fitted prob. at 10 th pctile. of experienced bond return	-0.141 (0.087)	0.085 (0.024)	0.149 (0.026)		
Probability of lowest risk tol. category		participation	participation		

Notes: Models and controls as in Table II, column (ii), Table III, columns (ii) and (iv), and Table IV, column (ii), but with experienced real stock and bond returns jointly included as explanatory variables and λ parameters fixed at the values obtained in those earlier regressions. The experienced stock return is calculated from the real return on the S&P500 index. The experienced bond return is calculated from the real return on long-term U.S. Treasury bonds. Observations are weighted with SCF sample weights. Standard errors shown in parentheses are robust to misspecification of the likelihood function/heteroskedasticity and adjusted for multiple imputation.

Table VI: Explaining Stock Market Return Expectations with Experienced Returns

	Expected stock market return over the next 12 months (1998 – 2003)		Expected return on own portfolio over the next 12 months (1998 – 2007)	
	(i)	(ii)	(iii)	(iv)
Experienced stock return coefficient β	0.801 (0.386)	0.393 (0.215)	0.800 (0.244)	0.480 (0.125)
Weighting parameter λ	0.923 [fixed]	1.698 [fixed]	0.923 [fixed]	1.698 [fixed]
Age dummies	Yes	Yes	Yes	Yes
Year-month dummies	Yes	Yes	Yes	Yes
#Obs.	40,145	40,145	72,631	72,631
R ²	0.08	0.08	0.06	0.06

Notes: Dependent variable is the expected return over the next 12 months from the UBS/Gallup survey. Estimation is done with least squares, weighted with sample weights. The λ parameter is fixed at the values obtained in column (ii) of Tables III and IV. The experienced stock return is calculated from the real return on the S&P500 index. Standard errors shown in parentheses are robust to heteroskedasticity.