

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

Ulrike Malmendier
UC Berkeley and NBER

Stefan Nagel
Stanford University and NBER

March 2008
First draft: December 2006

Abstract

We investigate whether individuals' experiences of macro-economic shocks 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 1964-2004, we find that individuals that have experienced high stock-market returns throughout their lives report lower risk aversion, are more likely to be stock-market participants, and, if they participate, invest a higher fraction of liquid wealth in stocks. At the same time, individuals that have experienced high inflation are less likely to invest their (non-stock) assets in bonds and favor inflation-proof cash-like investments. All results are estimated controlling for age, year effects, and a broad set of household characteristics. Our estimates indicate that the most recent returns and inflation rates have the strongest effect, but experiences earlier in life still have some influence, even several decades later. Thus, the experience of risky asset payoffs over the course of an individual's life affects subsequent risk-taking. Our results can explain, for example, the relatively low rates of stock-market participation among young households in the early 1980s (following the disappointing stock-market returns in the 1970s depression) and the relatively high participation rates of young investors in the late 1990s (following the boom years in the 1990s).

* We thank John Ameriks, Randy Cohen, Wayne Ferson, Tim Johnson, Karthryn Shaw, and seminar participants at Berkeley, Bowling Green State, Harvard, Hong-Kong University of Science and Technology, Nanyang Technological University, National University of Singapore, Singapore Management University, Stanford, University of Illinois Urbana/Champaign, University of Michigan, University of Southern California, the American Finance Association Meeting, the Annual Meeting of German Economists Abroad, and the NBER MIDM meeting for comments, and Sith Chaisurote and Nelli Oster for excellent research assistance.

“I don’t know about you, but my parents were depression babies, and as a result, avoided the stock market and all things risky like the plague.”

Source: moneytalks.org (“Investing: The Basics”)

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 had a long-lasting effect on their attitudes towards risk.

We ask more generally whether people who live through different macroeconomic histories make different risky choices. 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 (e.g., in books or via education) (Nisbett and Ross, 1980; Weber et al. 1993; Hertwig et al. 2004).

We examine empirically whether individuals’ risk attitudes in financial decisions differ depending on the macroeconomic history they experienced over the course of their lives. In particular, we test whether individuals who experienced low stock-market returns are less willing to invest in stocks and express more risk aversion, and whether individuals who lived through high-inflation periods are wary of investing in long-term bonds. We also ask whether the impact of experience depends on the time lag, e.g. whether more recent experiences have a stronger impact.

Our analysis does not attempt to disentangle the channel through which risk attitudes are affected (e.g., preferences versus beliefs), nor to contrast cognitive and friction-based explanations. Rather, we aim to improve our model of risk-taking by exploring the predictive power of life-time experiences and distinguishing them from effects of demographics and outcome variables such as wealth.

To test our hypothesis, we use repeated cross-section data on household asset allocation from the Survey of Consumer Finances (SCF) from 1964-2004. We utilize the triennial SCF starting in 1983 for

data on portfolio allocations and elicited risk aversion. For stock-market participation, we are able to extend the sample back to 1964, using survey waves from the precursor of the present-day SCF.

A key implication of the experience hypothesis is that old and young people differ in their risk attitudes in the direction of their life-time experiences. A simple scatter-plot of differences in stock-market participation against differences in experienced stock market returns illustrates the hypothesis.

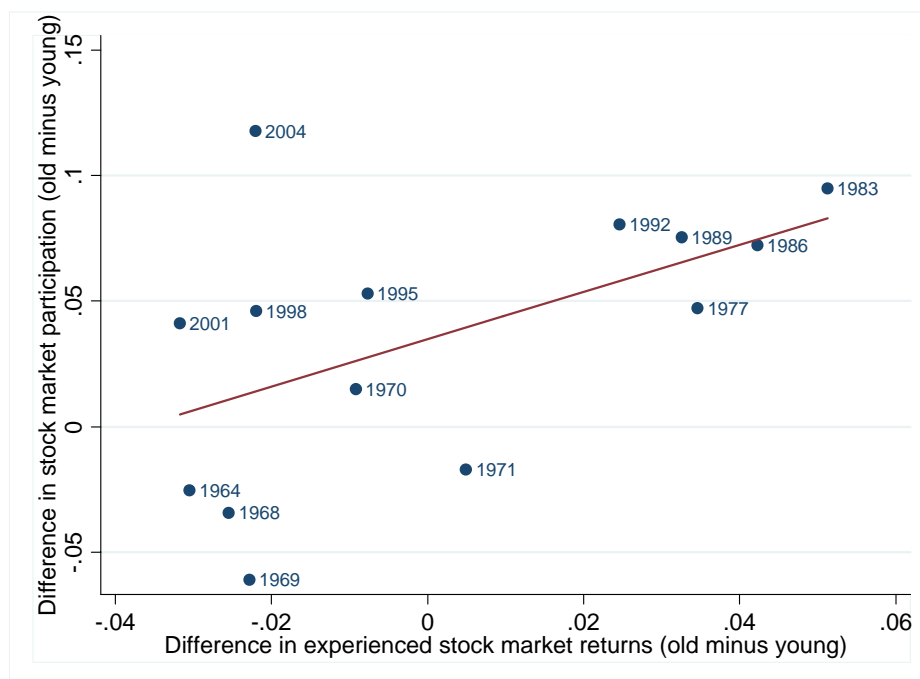


Figure 1: Differences in stock-market participation rates of old and young individuals plotted against differences in lifetime average stock-market returns. Stock market participation rates are the fraction of households who invest in stocks or mutual funds. The y-axis shows the participation rate of 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 SCF survey waves.

Figure 1 shows that in years when old people had a more *positive* stock-market experiences than young people (e.g., after the depression of the 1970s), their stock-market participation rates are higher, relative to those of young people, than in years when old people had more *negative* stock-market experiences than young people (e.g., in the 1960s when older individuals at the time still had the memory of the Great Depression). In this paper, we test whether these differences between individuals with different macroeconomic histories persist when we use a broader range of risk-attitude proxies, allow for different

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

We employ four different measures of risk-taking. The first measure is based on responses to a survey question about individuals' willingness to take financial risk. The second measure is stock-market participation. A third measure, applicable only for households that participate in the stock market, is the proportion of liquid assets (including bonds, cash, and cash equivalents) invested in stocks or mutual funds. The fourth measure is the proportion of liquid assets other than stocks that are invested in bonds, i.e., in assets subject to inflation risk. All four measures are likely to reflect a mixture between risk aversion and beliefs about future payoffs on risky investments.

We relate these measures of risk-taking to households' histories of stock returns and inflation. We calculate, for each household, at each SCF survey date, the annual real U.S. stock-market returns and inflation since the birth year of the household head. The extent to which individuals have "experienced" past returns and inflation differs, of course, depending on previous investments, interest in economic matters, and other personal circumstances that we cannot observe. The lack of such controls introduces noise in our explanatory variable. It biases our estimate only if these idiosyncratic factors are correlated with aggregate return or inflation, a concern we address below. In our estimation, we allow recent returns and returns early in life to carry different weights (if any) in influencing current risk-taking. In other words, we let the data speak, simultaneously, on how households weight past observations on returns and inflation and how strongly their risk-taking is correlated with the resulting weighted averages.

We find that households' risk taking is strongly related to their stock-market and inflation experiences. Households with higher life-time weighted average stock-market returns have lower elicited risk aversion, higher rates of stock-market participation, and a higher allocation to stocks. Households with higher life-time weighted average inflation invest fewer non-stock liquid assets in bonds. The estimated weights are remarkably similar for all four risk-taking measures. More recent returns and inflation receive higher weights, and thus have a stronger influence on risk-taking than those early in life, but even returns and inflation experienced decades earlier still have some impact for older households.

All of our estimations include a vector of year effects, which control for time trends or any aggregate effects. For example, a mechanical positive relationship between recent stock returns and households' allocation to stocks arises from market clearing: holding the supply of stocks fixed, the average portfolio share invested in stocks must increase whenever stock market prices increase in the aggregate (for example, because of investors' risk aversion declined). The inclusion of year effects removes this effect. The identification comes from cross-sectional *differences* in risk-taking and in macroeconomic histories, and from changes of those cross-sectional differences over time, not from *common* variation over time. For example, our data show that young households in the early 1980s, having experienced the dismal stock returns of the 1970s, had lower rates of stock-market participation, lower allocation to stocks, and reported higher risk aversion than older households. For older households, the experience of the low 1970s stock-market returns was moderated by having experienced the high returns of the 1950s and 1960s. Following the boom years of the 1990s, this pattern flipped. Now young households had higher life-time average returns and, consistent with the experience-effect interpretation, also higher rates of stock-market participation, higher allocation to stocks, and lower reported risk aversion than older households. It is these correlated changes in the age profile of life-time weighted average returns and risk-taking that our identification comes from.

Our estimation also accounts for age effects. As consumers grow older, they may reduce their risky-asset share or even abstain from stock-market participation (see Hurd, 1990), though it is not clear whether such behavior is optimal – a question discussed at least since Samuelson (1969). All regressions include a third-order polynomial in age or, alternatively, a full set of age dummies, in addition to our dummies for retirement, ruling out that any time-invariant life-cycle effect explains our findings.

Another potential confound are wealth effects. If life-time average returns are correlated with current wealth and if risk aversion is wealth-dependent, variation in wealth could explain the relation between current risk taking and life-time average return. We address this concern in two ways. First, all estimations include wealth and income controls. Second, to the extent that unobserved differences in wealth remain, we argue that wealth is unlikely to provide a common explanation for all of our risk-taking

measures. The experience of real stock returns could be positively correlated with real wealth, but the same is unlikely to be true for inflation. Moreover, 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 and elicited risk aversion (Brunnermeier and Nagel, *forthcoming*; Sahm 2007).

A last potential confound are unspecified cohort effects. Previous work, which has looked at cross-cohort differences in risk-taking with cohort dummy variable regressions (see, e.g., Ameriks and Zeldes, 2004), faced the problem that cohort effects cannot be separated from age and time effects, due to the collinearity of age, time, and cohort (see, e.g., Heckman and Robb 1985, and the discussion in Campbell, 2001). An advantage of our analysis is that our hypothesis predicts a *specific*, signed relationship between macroeconomic experiences and risk-taking. Since our identification strategy does not rely on estimating cohort effects, we can control for age and year effects simultaneously. Moreover, since life-time weighted average return and inflation vary not only across, but also within cohorts over time, we can even include an almost full set of cohort dummies and therefore control for any omitted variable that has cohort-level variation.

In summary, our findings suggest that individual investors' willingness to bear financial risk depends on personal experiences of macroeconomic history. This behavior could be explained either with endogenous preferences, where risk aversion depends on the risky asset payoffs experienced in the past, or with learning, where current beliefs depend on the realizations experienced in the past. In the latter case, learning from personal experience would lead to beliefs that do not converge across overlapping generations, even in the long-run. Such belief heterogeneity is a departure from standard learning models, in which all agents at a given point in time have access to and make use of the same history of past data.

Our paper connects to several strands of literature. While there is no prior literature, to the best of our knowledge, documenting the effect of long-term macroeconomic experiences on economic decisions,

several papers in macroeconomics and public finance analyze the impact of age and demographic composition. Most closely related is the work by Poterba (2001), who studies the effect of age on individual investment decisions, controlling for cohort fixed effects but (to avoid collinearity) not controlling for time effects. Other work links demographic changes to the aggregate demand for stocks and bonds (Goyal, *forthcoming*; Ang and Maddaloni, 2005; Geneakoplos, Magill, and Quinzii, 2004), and evaluates the effect of cohort size on a wide range of economic outcomes, including family choices (Easterlin 1987), social security (Auerbach and Lee, 2001; Gruber and Wise, 1999), college graduation (Card and Lemieux, 2000; Bound and Turner, 2003), research and development (Acemoglu and Lin, 2004), industry returns (DellaVigna and Pollet, 2007), and a range of macro variables (Fair and Dominguez, 1991). None of the above papers consider cohort experiences beyond those induced by size.

Another strand of the economics literature discusses endogenous preference formation more broadly. Bowles (1998) argues that market institutions influence the evolution of taste, e.g., by determining the types of activities and situations experienced by individuals and hence the social norms towards these activities. For example, the Great Depression may have affected stock-market participation by changing the attitudes of the society towards investing in the stock market. Prior to the Depression, with robber barons accumulating large gains in the 1800s, stock market investment had become more and more widespread. After the losses, social norms changed and people were less willing to participate in the stock market. Closely related to the emphasis in our paper, Palacios-Huerta and Santos (2004) apply endogenous preference formation to risk attitudes, formed as function of exposure to market risk, market incompleteness and uncertainty.

The role of personal experience relative to “learned” or “observed” information is discussed in several strands of the experimental and behavioral economics literature. The literature on reinforcement learning posits that subjects’ choice of actions strongly depends on the payoffs they obtain from the same actions in the past, even if circumstances (beliefs about other players’ behavior and hence predicted payoffs) have changed. The experimental tests of the “experience-weighted attraction” model in Camerer and Ho (1999), which links reinforcement and belief learning, show that the actual payoffs obtained from

past behavior have a large impact on subsequent choices. Relatedly, Schlag's (1999a and 1999b) models and experimental tests of social learning suggest that individuals tend to imitate behavior that has worked well in the past. Most relatedly, Simonsohn, Karlsson, Loewenstein, and Ariely (2007) show, in a series of repeated weak-link and prisoner's-dilemma games, that players' behavior is more strongly affected by the behavior of players they personally interacted with in previous rounds (and the resulting payoffs to their own actions) than by the behavior of players they did not interact with. This behavior is observed even though players are randomly and anonymously rematched in each round and observe the behavior of all players after each round. Another related experimental finding regards the role of advice: As reported in Schotter (2003), subjects adjust their actions in response to advice of previous "generations" of players more than in response to historical data about the behavior and outcomes in the games of those previous generations.

In the context of financial decision making, Greenwood and Nagel (2007) show that young mutual fund managers had more exposure to technology stocks in the late 1990s than older managers, particularly after quarters with high technology stock returns, 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 following the stock-market boom in the late 1990s young retail investors with little investment experience had the highest stock-market return expectations. While these two papers focus on effects of recent returns on young investors in the late 1990s, our paper uses a long-term sample and a broad range of risk-taking measures to estimate the long-run effect of stock-market returns on risk-taking and controls for age effects.

A couple of papers, which focus on different topics, include circumstantial evidence consistent with the view that personal experience matters. Piazzesi and Schneider (2006) report that in the late 1970s old households expected lower inflation than young households. Young households apparently had a stronger tendency to extrapolate from their recent personal experience of high inflation at the time. Graham and Narasimhan (2004) find that corporate managers that have lived through the Great Depression in the 1930s choose a more conservative capital structure with less leverage. Finally, Cogley

and Sargent (2005) build a model that explains the equity premium based on the assumption that the Great Depression had a long-lasting effect on investors' model uncertainty about the 'true stochastic model' determining consumption growth and hence investment behavior, along the lines suggested by Friedman and Schwartz (1963). If individuals learn from personal experiences of economic events and asset payoffs, as our evidence suggests, a big disaster like the Great Depression would indeed have these kinds of effects.

II. Data and Methodology

A. Survey of Consumer Finances

We use data from 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 2004, obtained from the Board of Governors of the Federal Reserve System and available every three years. The second source is the SCF precursor, 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. We found that the data prior to 1964 is not usable for our purposes since information on stock holdings is either missing or very crude, and since the sampling unit is the "spending unit" rather than the "family unit" that is used from 1964 on. To ensure comparability across years we start in 1964 and use all survey waves that offer stock-market participation information, i.e., the 1964, 1968, 1969, 1970, 1971, and 1977 surveys.

The 1983-2004 waves oversample high-income households with substantial asset holdings. As we will see, the oversampling of high-income households is helpful for our analysis of relative asset allocation decisions, but may induce selection bias. In our estimation, we include controls for income and wealth and weight with the sampling weights provided in the SCF² to account for the potential bias.

² 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 (2000), we normalize the sample weights each year so that the sum of the weights in each year is the same.

All of the asset holdings variables that we define below exclude assets in retirement accounts since the survey waves prior to 1989 do not provide information on the composition of assets in retirement accounts (e.g., IRA, Keogh, and 401(k) plans). Even from 1989 on (but prior to 2004), the SCF offers only very coarse information on the allocation of retirement assets (mostly stocks, mostly interest bearing, or split), precluding any meaningful calculation of stock holdings. We do, however, conduct robustness checks with data that includes retirement account holdings.

The key variables for our analysis are past stock-market returns that occurred during the lifetime of the household head, and several measures of risk-taking. For each household we calculate the annual real returns on the S&P500 index from the time of the household head's birth up to the end of the year preceding the survey date. For example, for a household head that is 50 years old in 1983, we take the real returns on the S&P500 index from 1933 to 1982. We use the same approach for annual inflation, based on the Consumer Price Index (CPI). Both indices are from Shiller (2005) and go back to 1871.³

The first measure of risk taking is the risk aversion elicited in the SCF waves in 1983 and 1989-2004. The SCF asks whether the interviewee is willing to (1) take substantial financial risks expecting to earn substantial returns; (2) take above average financial risks expecting to earn above average returns; (3) take average financial risks expecting to earn average returns; (4) not willing to take any financial risk. We code the answer as an ordinal variable with values from 1 to 4.

The survey answer is an imperfect measure of risk aversion for several reasons. First, individuals may differ in their interpretation of, say, "substantial" or "above average" risks and returns. For this reason, we cannot interpret the measure in a cardinal sense. Second, the answers are affected by differences in beliefs about the future payoffs of risky assets. An individual who believes that expected equity premium is high (expecting to earn a "substantial return") would, presumably, be willing to put a large proportion of her portfolio into stocks ("take substantial risks"). Thus, the measure represents, at best, the combined effect of Arrow-Pratt risk aversion and beliefs. Despite these shortcomings, prior

³ 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.

literature documents that the measure predicts individual willingness to take risks. For example, Faig and Shum (2006) find that households that report higher risk aversion in response to this question have a lower allocation to risky assets. Shaw (1996) shows that the measure helps explain differences in the willingness to make risky human capital investments and in wage growth. In our analysis, using both the survey question and the direct measures of asset allocation described below ameliorates concerns about alternative interpretations. At the same time, we do not claim that the survey answer reflects only risk aversion. We refer to the measure as “elicited risk aversion” for ease of reference.

Our second measure is a binary variable for stock-market participation, available from 1964-2004. We define stock holdings as the sum of directly held stocks (including stock held through investment clubs) and the total amount held in stock and bond mutual funds. We include bond funds, because stock and bond funds are not reported separately in the survey waves prior to 1989 (only money market mutual fund holdings are reported separately, and we do not include those in our stock holdings variable). The amount invested in bond funds is relatively small (in the 2004 SCF, for example, bond funds account for only about 20% of total non-money market mutual fund holdings) and so the treatment of bond funds is unlikely to have much impact on our results. Nevertheless, as a robustness check, we re-run our tests on the 1989-2004 sample including in the calculation of stock holdings only those mutual funds that invest mostly in stocks.

Our third measure of risk taking is the fraction of liquid assets invested in stocks (directly held stocks plus mutual funds), available from 1983-2004. Liquid assets are defined as stock holdings plus bonds plus cash and cash equivalents (checking accounts, savings accounts, money market mutual funds, certificates of deposit) plus the cash value of life insurance plus other liquid assets.

Our fourth measure of risk taking is the fraction of liquid assets other than stocks that are invested in bonds (as opposed to cash, savings accounts, and short-term money market investments, for example). While we do not know the maturity structure of households’ bond positions reported in the SCF, it is reasonable to assume that a significant portion has maturities of several years or more. These bond holdings are risky in real terms, because of future unexpected inflation. Our life-time experience

hypothesis suggests that the extent to which investors hold long-term bonds is influenced by past experience of inflation. Cohorts that have lived through high-inflation periods should be wary of investing in long-term bonds and prefer short-term instruments; cohorts that have lived only during periods of low inflation should be more willing to invest their non-stock liquid assets in bonds.

As a control variable for income we use total family income. All income, wealth, and asset holdings variables are deflated into September 2004 dollars using the consumer price index. When we use the liquid wealth variables from the 1964 survey in our regressions, we always interact them with a 1964 dummy, because the definition of the wealth variables in that year differs from the other survey years. (In 1964, the liquid-wealth variable includes, for example, some real-estate assets.)

We remove observations that are likely to be miscoded and households for which the asset allocation issue does not apply because they do not have any liquid asset holdings, following previous SCF literature.⁴ Specifically, we require that households have at least \$100 of liquid assets and annual family income greater than \$1,000. 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.

For our summary statistics and graphical descriptive analyses, we weight the data using SCF sample weights. The weighted statistics are representative of the U.S. population. In our subsequent econometric estimation we start with unweighted estimates, since weighting is, in principle, inefficient use of the data (see, e.g., Deaton 1997, p. 70). Instead, we employ control variables for wealth and income. For robustness, we also present results from weighted estimation.

B. Methodology

Our aim is to investigate the relationship between risk-taking and prior stock-market returns and inflation experienced by the household head since birth. We also want to allow for the possibility that experiences in the distant past have a different influence than more recent experiences. For example, the

⁴ 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.

memory of past stock-market returns might fade away as time progresses. Or, experiences at young age might be particularly formative and have a relatively strong influence on individuals' decisions today. Both hypotheses are not mutually exclusive: the impact of past returns may generally decay, but perhaps with a lower decay rate for the "first experience." We aim to allow for both possibilities. Our goal is to estimate the partial effect of each of the yearly returns and inflation on risk-taking.

A flexible estimation of the effect of all past returns on current risk-taking faces two hurdles. First, it would be problematic to run regressions with an exceedingly large number of explanatory return variables, say, with 50 annual returns for a 50-year old household head, and to leave the coefficients on each return variable unconstrained. The standard errors would be too large to allow any meaningful inference. Second, since we would like to consider returns and inflation back until birth, the number of explanatory variables differs across households depending on their age. To solve both problems, we use a weighted average of the household head's experienced returns and inflation since birth. This is equivalent to imposing constraints on the coefficients of each of the yearly return or inflation measures since birth. 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 the time lag since birth. In other words, we 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 stock returns:

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

where R_{t-k} is the real stock-market return in year $t-k$. The weights w_{it} depend on the age of the household head and a parameter λ that controls the shape of the weighting function. We estimate λ from the data. If $\lambda < 0$, then the weighting function is 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 stock-market returns since birth. With $\lambda > 0$ weights are decreasing in the lag k (concave for $\lambda < 1$, linear for $\lambda = 1$, and convex for $\lambda > 1$). We apply the exact same methodology to calculate the weighted average of past inflation.

Figure 2 provides an example of the weighting functions for three values of λ for a household of age 50. As the figure shows, the weighting function is quite flexible in accommodating different weighing schemes. It cannot accommodate “humps”, however. The weights are either monotonically increasing, decreasing, or flat. It is possible that the true weighting function is more complex than what our specification accommodates. In this case, however, our restriction simply biases the estimation against finding any significant effect of the resulting weighted average returns or inflation on risk-taking.

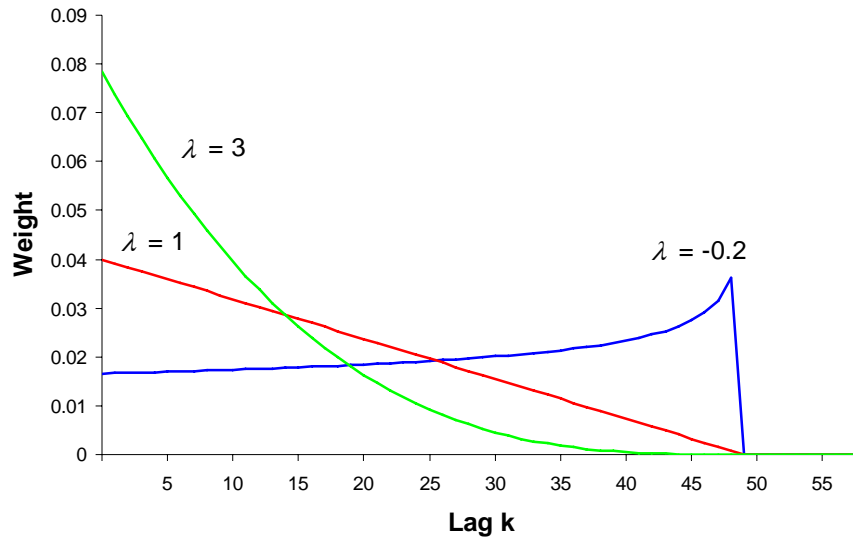


Figure 2: Three examples for the life-time stock-market returns weighting function for a household with a 50-year old household head.

As an example for how we estimate the weights and the sensitivity of risk-taking to the life-time average returns calculated with those weights, consider the following generic regression model, with y_{it} as

the dependent variable and weighted-average 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 . Note that $A_{it}(\lambda)$ is a non-linear function of the weighting parameter λ , and hence non-linear estimation is required. For Probit models, we choose β and λ to maximize the likelihood; for regression models, we choose them to minimize the sum of squared errors.

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 (and normalizing by the sum of weights, $\sum_{k=1}^{age_{it}-1} w_{it}(k, \lambda)$) yields, for a household of that age, the partial effect of a return (or inflation) experienced k years ago on the dependent variable. As an example, if $\lambda = 0$, then all returns (or inflation) in the household head's history since birth are weighted equally, and so their partial effects are all equal to their weight (one divided by age) times β .

C. Summary Statistics

Table I provides some summary statistics on our sample. Panel A (1964 – 2004) and B (1983 – 2004) include all households that satisfy our sample requirements. Panel C (1983 – 2004) further restricts the sample to stock-market participants, i.e., households that have at least \$1 in stocks or mutual funds, and Panel D restricts the sample to bond-market participants, i.e., households that have at least \$1 directly invested in bonds. Comparing Panels B and C, it is apparent that stock-market participants tend to be wealthier. For example, the median holding of liquid assets is \$13,245 in the full 1983-2004 sample, but \$65,200 in the sample of stock-market participants. Panel D shows that bond-market participants are also wealthier, with median liquid assets of \$30,399, though less than stock-market participants. The pattern is the same for median income. In the full 1983-2004 sample, median income is \$48,674, in the sample of stock-market participants, it is \$75,654, and in the sample of bond-market participants, it is \$65,748. The

median and the lower half of the income distribution are very similar in the full 1964-2004 and the (full) 1983-2004 samples. The upper half of income, however, widens in 1983-2004.

As Panel B shows, 28.5% of households participate on average in the stock market in the 1983-2004 period. This number is strikingly similar to the 28.6% participation rate in the full 1964-2004 period shown in Panel A. As described above, these rates represent the U.S. population (not the SCF sample) since we apply the SCF sample weights.⁵ This finding is somewhat surprising. It is sometimes argued that stock-market participation rates have been trending upward since the 1980s because of lower participation costs due to improved communications technology and reduced transaction costs (Choi, Laibson, Metrick, 2002). However, the early SCF data shows that participation rates were quite high in the 1960s, too, suggesting that the technological improvements story may not be the sole explanation for the recent surge. Our hypothesis that past returns experienced by investors over their lifetime play a role in generating variation in stock-market participation over time and across individuals may help explain this pattern.

The three other risk-aversion measures, available only for the 1984-2004 sample, also show considerable dispersion across households. The proportion of liquid assets invested in stocks in Panel C has 10th and 90th percentiles of 4.4% and 87.8%. The proportion of non-stock liquid assets invested in bonds in Panel D has 10th and 90th percentiles of 0.5% and 63.2%. The 10th and 90th percentiles for elicited risk aversion in Panel B are 2.0 and 4.0, respectively. It is noteworthy that mean elicited risk aversion is *lower* for the stock-market participants in Panel C (2.792) than for the full sample in Panel B (3.126) and lies in the middle for bond market participants in Panel D (3.003), which suggests that the elicited risk-aversion measure is indeed correlated with households' actual attitudes towards financial risk-taking as revealed by their participation choice.

Our main question of interest is whether the variation in risk-taking measures across households is related to the life-time average stock return and the life-time average inflation experienced by the household head's birth-cohort. To get a sense of the variation in life-time average stock returns for the

⁵ 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 C is higher than 28.5% of the number of observations in Panel B.

households in our sample, we calculate the weighted average of stock returns, $A_{it}(\lambda)$, from Eq. (1), 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 real life-time average stock return are 5.9% and 11.0% in the 1964-2004 sample. The 10th and 90th percentile for the real life-time average inflation with $\lambda = 1.00$ are 2.3% and 5.5%. Hence, there are considerable differences in the life-time average returns and inflation experienced by different cohorts. The amount of variation in the life-time average stock return is similar for a range of values around $\lambda = 1.25$. For example, with $\lambda = 0.75$ and $\lambda = 1.75$, values that are roughly the boundaries of the interval that contains the point estimates we obtain subsequently, we get differences between the 10th and 90th percentile of 3.8% and 5.5% for returns, respectively. The same is true for life-time average inflation. If we set $\lambda = 0.50$ we get a difference between the 10th and 90th percentile of 2.8%; if we set $\lambda = 1.50$, we get a quite similar difference of 3.4%.

III. Results

A. Elicited Risk Aversion

We start by relating life-time average returns to elicited risk aversion. 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, α_j denote the cutoff points that must be estimated ($\alpha_1 = 0 < \alpha_2 < \alpha_3 < \alpha_4 = \infty$), $A_{it}(\lambda)$ is the weighted life-time average return, x_{it} is a vector of control variables. Unlike the standard ordered probit estimation, $\Phi(\cdot)$ does not map a linear function of explanatory variables into the response probability P. Instead, $A_{it}(\lambda)$ 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 . The coefficient vector β does not have a direct economic interpretation. To interpret the results

of the ordered probit estimation, we focus on the partial effects of the life-time average return $A_{it}(\lambda)$ on the probabilities for being in one of the four risk-aversion categories, i.e., $\partial P(y_{it} = j | x_{it}, A_{it}(\lambda)) / \partial A_{it}(\lambda)$. We evaluate the partial effects at each sample observation, given the estimated parameters and observations on x_{it} and $A_{it}(\lambda)$ and calculate the average partial effect across sample observations.

The vector x_{it} includes income controls (log income, log income squared), demographics controls (a third-order polynomial in age to allow for a non-linear age profile, a second-order polynomial in the number of children, dummies for retirement, completed high school education, completed college education, marital status, and race) and year dummies. We also control for the level of liquid assets held by the household (log liquid assets and log liquid assets squared).

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 year (Heckman and Robb 1985). Instead, the experience hypothesis predicts that a specific variable (life-time average 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. This may prove to be a concern if variations in stock returns over time are caused at least partially by variations in aggregate risk aversion of investors. For example, stock prices might rise as investors become less risk averse, and we get a negative correlation between past stock-market returns and investors' current risk aversion, but with causality running from current risk aversion to past stock returns. This concern is addressed by the inclusion of year dummies, which absorb all aggregate time effects including variation in average risk aversion. The effect of life-time average stock returns is therefore identified from cross-sectional differences in risk taking, not from aggregate time-variation. For the other risk-taking measures that we consider below, the year dummies also absorb all other unobserved aggregate factors that might lead to

changes in stock prices and, hence, simultaneously change past stock returns and investors' allocation to stocks (through market clearing).

Table II presents the results of the ordered probit model, with the parameter estimates at the top and the average partial effects at the bottom. Each average partial effect shows how a partial change in $A_{it}(\lambda)$ affects the probability of being in the respective risk-aversion category, $P(y_{it} = j | x_{it}, A_{it}(\lambda))$. Column (i), estimated on the 1983-2004 sample, shows that higher life-time average returns increase the probability that risk aversion is in the low categories (1 and 2), have little effect on the probability of being in category 3, and decrease the probability that the reported risk aversion is in the highest category (category 4). Thus, stock-market returns experienced in the past have a significant and positive effect on risk attitudes. Recall from Table I that the difference between the 10th and 90th percentile of life-time average stock returns is about 5.1%. Applied to the average partial effects in Table II, Column (i), this means that a change from the 10th to the 90th percentile implies about $-1.364 \times 5.1\% = -7.0\%$ decrease in the probability of being in the highest risk-aversion category.

The estimate of 1.546 (s.e. 0.355) for the weighting parameter λ implies that even returns experienced many years in the past still affect households' level of risk aversion. For values of λ around 1.0 the weighting function has approximately linearly declining weights (recall 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 ruled out. Nevertheless, the estimates imply non-negligible weights of returns early in life. Apparently, the memory of these early experiences fades away only very slowly.

As Table II shows, adding the liquid asset controls in Column (ii), or applying SCF sample weights in Columns (iii) and (iv) does not lead to any substantial change in the results.

B. Stock-market Participation

For our second estimation, the effect of life-time average returns on stock-market participation, we can use the long 1964-2004 sample. 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 $A_{it}(\lambda)$ on the probability of stock-market participation and focus on the partial effect $\partial P(y_{it} = 1 | x_{it}, A_{it}(\lambda)) / \partial A_{it}(\lambda)$. Given the estimated β and λ , we evaluate this partial effect at every sample observation and average across all observations to obtain the average partial effect.

The vector x_{it} includes the same income and demographics controls as in (3). Controlling for the level of liquid assets is particularly important in this context since a fixed participation-cost explanation predicts that stock-market participation is positively related to the level of liquid assets. Given that past stock returns are likely to be positively correlated with current liquid assets, a positive relation between stock-market participation and past stock-market returns could arise just from omitting of liquid assets from the model.

Table III reports the estimates from our probit model. We show the estimates of the parameters of interest (β and λ), and the average partial effects for the life-time average returns variable.⁶ As we can see from Column (i), the life-time average returns have a positive and highly significant effect on stock-market participation. The average partial effect of 1.929 (s.e. 0.322) means that a change from the 10th to the 90th percentile of life-time average returns (5.1%, taken from Table I) leads to an increase of about $1.929 \times 5.1\% \approx 9.8\%$ 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.

As in the previous Subsection, the estimate of 1.290 (s.e. 0.212) 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 obtained in the risk-aversion model in Table II, even though the first

⁶ The unreported coefficients of the control variables have the sign and magnitude that one would expect given the prior literature. Education, income, and liquid assets all have a strong positive effect on stock-market participation; race matters, too.

measure is based on risk aversion reported by the interviewee and, thus, very different from risk-taking measures based on asset holdings. Yet, a significant part of the variation in both of them can be traced to between-cohort variation in experienced stock-market returns with roughly similar weights on the history of past returns.

In Column (ii), we add the liquid assets controls. The estimated average partial effect of life-time average returns (1.719; s.e. 0.303) is slightly lower than in Column (i). The point estimate for λ is 0.994 (s.e. 0.184), which suggests somewhat higher weights on returns in the distant past compared with Column (i).

Columns (iii) and (iv) repeat the analysis of Column (ii) with the sample split into the old (1964-1977) and the new (1983-2004) SCF sample. The results are remarkably similar. In particular, in both subsamples the estimated average partial effect is close to the value in Column (ii), suggesting that the relationship we are estimating is stable over time. The standard error in the old SCF subsample is considerably larger, though, reflecting the lower number of observations.

Finally, Columns (v) and (vi) redo the estimation for the 1983-2004 sample with observations weighted with SCF sample weights. These weights undo the oversampling of high-income households in the 1983-2004 SCF. As the table shows, this has little effect on the results.

C. Proportion Invested in Risky Assets

Table IV shows the estimated effect of life-time average stock returns on the risky asset share, i.e., the proportion of liquid assets that households invest in stocks and mutual funds. This measure allows us to control for fixed costs of stock-market participation, which are likely to affect stock-market participation but not the risky asset share, conditional on participating.

We use a nonlinear regression model to estimate the effect of life-time average 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 risky assets. The model is nonlinear, because the life-time average 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 .

The control variables are the same as in Tables II and III, and the sample period is again restricted to 1983-2004, because we do not have quantitative information on asset holding in the early SCF sample. As Column (i) shows, the life-time average return has a positive and large effect on the proportion of liquid assets invested in risky assets. The point estimate of 1.139 (s.e. 0.485) implies that a change from the 10th to the 90th percentile of life-time average returns (5.1%) leads to an increase of about $1.139 \times 5.1\% \approx 5.8\%$ in the proportion allocated to risky assets.

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 are significantly correlated with the portfolio risky asset share (see Curcuru, Heaton, Lucas, and Moore (2004), and Brunnermeier and Nagel (*forthcoming*) for recent evidence). In light of this evidence, life-course experience of 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 close to 1.0, which suggests weights that are approximately linearly declining from the prior year going back to zero weight in the birth year. The estimate for λ is in the ballpark of the estimates for λ in the elicited risk-aversion model in Table III and the stock-market participation model in Table IV, even though stock-market participation and choice of the risky asset share conditional on participation are possibly quite distinct decisions. That the returns in the distant past carry roughly similar weights is reassuring for our interpretation that the three measures capture a common attitude to financial risks and are subject to a common influence. The similarity in the estimates for elicited risk-aversion and for the risky asset share is particularly remarkable since the two models use very different approaches (survey question versus investment choice).

Adding the liquid asset controls in Column (iii) has little effect on the estimates. Weighting observations with SCF sample weights also does not change the results much: In Column (iv) the point estimate for β is almost identical to Column (ii), only the weighting parameter λ is estimated to be a bit higher (1.428, s.e. 0.080).

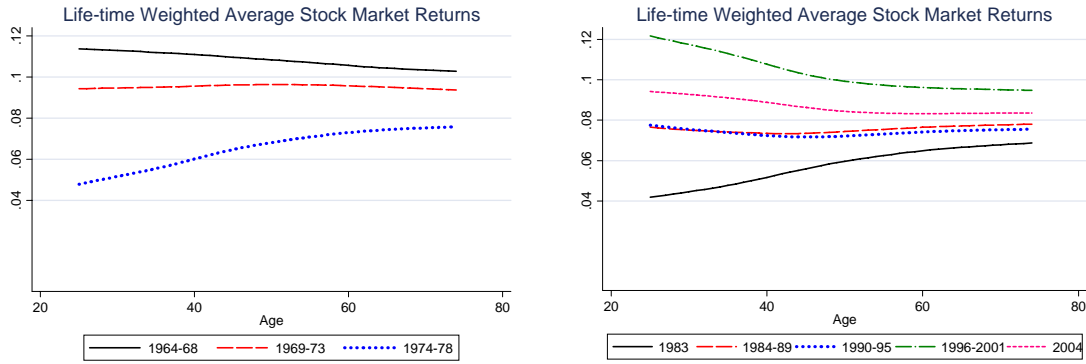
D. Graphical Summary

Figure 3 provides a graphical summary of the results for the first three measures of risk taking. We split the 1964-2004 sample into five-year subperiods, so that each contains one or more survey waves. Within each subperiod, we plot our key explanatory variable, life-time average returns, and our three outcomes variables, stock-market participation rates, risky asset shares, and elicited risk aversion, as a function of age.

The two graphs at the top show life-time average returns as a function of age. For each five-year subperiod, we calculate the average of life-time average returns (using $\lambda = 1.25$) across households with the same age, weighted by the SCF sample weights. We then employ a kernel regression to smooth the age-profile. The left plots shows how the low stock-market returns of the 1970s shifted down average returns but also increased the slope of the age profile. In the late 1970s, the life-time average return of young households is dominated by the low returns of the depression years. In contrast, in the late 1990s shown in the right plot, young households' experienced return histories are dominated by the boom years of the 1980s and 1990s. For older households, the differences are less extreme due to the longer history over which returns are averaged, and so the positive slope flattens from the early 1980s to the late 1990s. The post-2000 down-market finally resulted again in a somewhat steeper curve in 2004.

In the next row, in the middle row, we plot "residual risk aversion," defined as the difference between actual and predicted elicited risk aversion, where predicted risk aversion is estimated in the Ordered Probit model of Column (ii) in Table II, but without including the life-time average return

Panel A: Return Experiences



Panel B: Measures of Risk-taking

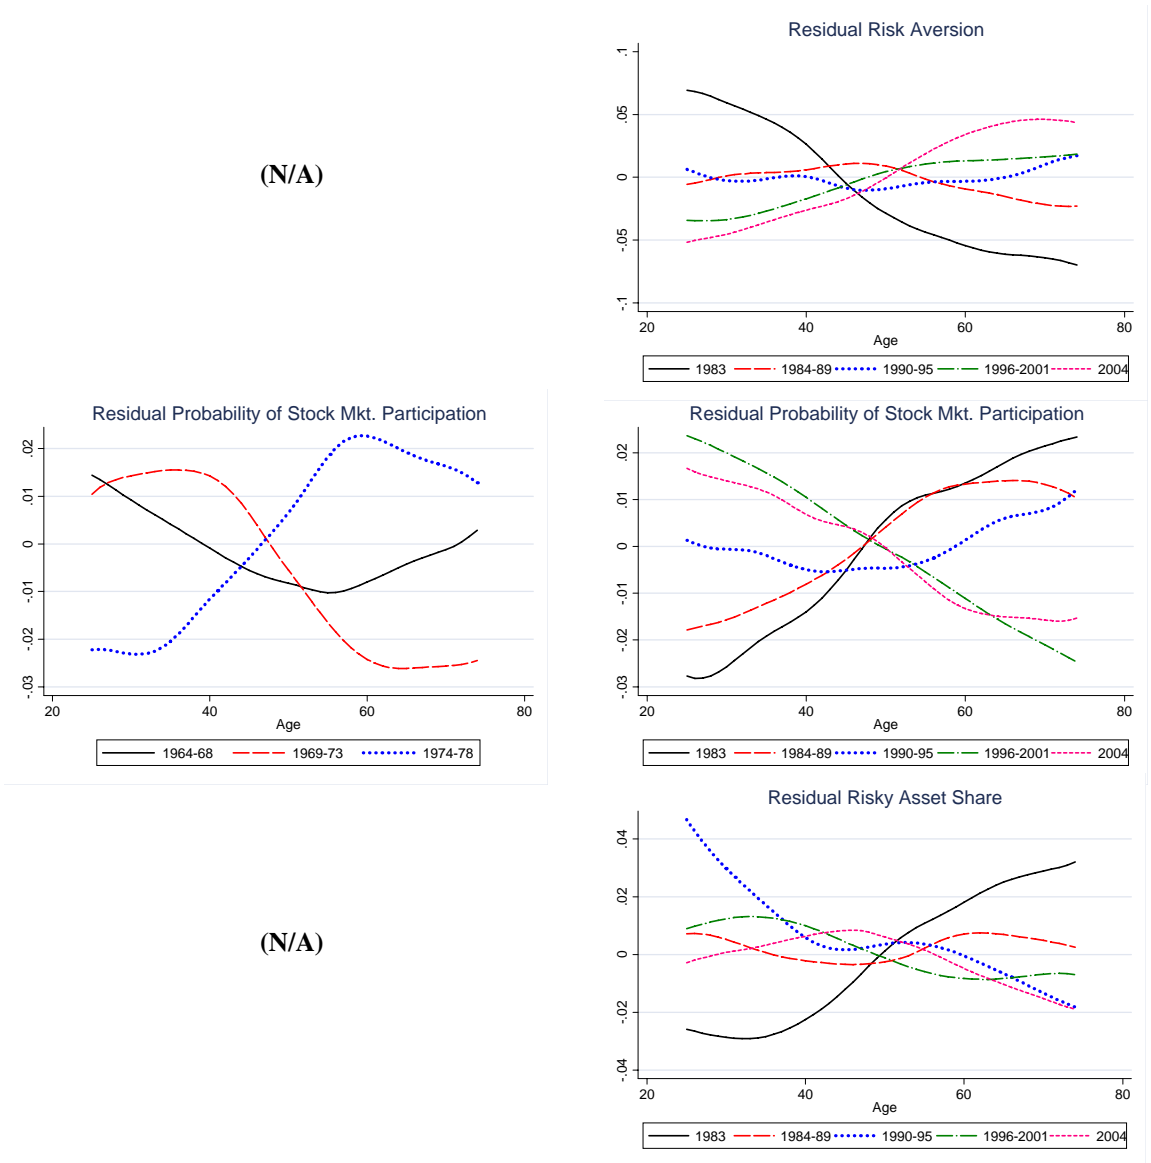


Figure 3: Life-time weighted average stock-market returns and residual risk-taking measures for as a function of age. We set $\lambda = 1.25$ in the calculation of life-time weighted average returns. The plots are smoothed with kernel regressions (simple local averaging with tri-cube weights, bandwidth 0.8).

variable. The residual risk aversion is then averaged by age within each five-year subperiod, and smoothed with a kernel regression. Comparing the risk aversion graph to the corresponding return graph in the first row, we see an almost monotone inverse relation between the slope of the age profile of residual risk aversion and the slopes of life-time average returns. Thus, when younger people have higher life-time average returns than older people, they tend to have lower risk aversion than older people, consistent with the experience hypothesis.

We perform similar exercises for our other two risk-taking measures in the bottom row of Figure 3. The two graphs in the middle row plots residual stock market participation probabilities, based on the Probit model residuals from Column (ii) in Table III, but excluding the life-time average returns from the model. It is apparent that the slopes of the average return profiles are positively related to the slopes of the stock market participation profiles. Going from the early 1960s to the late 1970s, the residual participation probability of young households dropped substantially relative to older households; during the 1980s and 1990s, we see a reversion towards a downward sloping age profile. This parallels the initial steepening and subsequent flattening of the life-time average returns age profile.

For the residual risky asset share in the bottom row, the pattern is roughly similar to that of residual stock-market participation: an upward sloping age profile in the early 1980s, and subsequent reversal to a downward sloping profile in the 1990s, but the patterns are not as clean as for stock-market participation.

Overall, the plots illustrate how financial risk-taking among young and old households changes over time, and that relative differences between young and old are closely related to changes in the relative life-time average returns experienced by young and old.

E. Bond holdings and inflation

We turn to our fourth risk-taking measure, the proportion of non-stock liquid assets invested in bonds, and relate it to life-time inflation. We estimate the non-linear least squares model specified in Section III.C, substituting the risky-asset share of liquid assets with the bond share of non-stock liquid

assets and substituting life-time average returns with life-time average inflation. Our hypothesis is that past experiences of high inflation should reduce the willingness to hold bonds.

As Table V shows, this hypothesis is borne out in the data. The coefficient on life-time average inflation in Column (i) is negative, $\beta = -2.727$ (s.e. 1.145). Variations in the set of control variables, or the weighting of observations has little effect on the results. The point estimates in Columns (ii), (iii), and (iv) of Table V are very similar.

Interestingly, the point estimate of the weighting parameter in Column (i), $\lambda = 1.008$ (s.e. 0.061), is similar to those that we estimated when we related the other three risk-taking measures to life-time average stock returns. For both inflation and stock returns, the implied weights are roughly linearly declining. Hence, the life-time experience effect appears to play out in a similar way for different aspects of economic risk-taking and it may be possible to treat them all in a common framework of experience-based beliefs or preferences.

The confirmation of the experience effect in the inflation-bond context is also important since it addresses a number of alternative explanations. One potential explanation for the positive relationship between past stock returns and risk-taking is an unobserved wealth effect on risk-taking that is not captured by our controls. While stock returns are clearly positively associated with wealth changes, it is not clear that such a link exists for inflation, or if there is one, in which direction it would be. For example, the owner of a long-term nominal bond would suffer a loss in real terms if there is unexpected inflation, while someone borrowing on a fixed-rate mortgage may gain from unexpected inflation. Moreover, real stock-market returns and inflation have little correlation.⁷ The fact that the effect of life-time experience on risk-taking is so similar for stock returns and inflation implies that wealth effects cannot fully explain our results.

In terms of economic magnitudes, the effect of life-time average inflation is sizeable. The difference between the 10th and the 90th percentile of life-time average inflation from Table I, Panel B, is

⁷ In our data, the correlation of annual real stock returns and inflation rates since 1871 is -0.14, and not statistically significantly different from zero.

1.7% during the 1983-2004 period. The estimated β implies a variation in the share of bonds of $-2.727 \times 1.7\% \approx -4.6\%$, which is sizeable relative to the mean bonds share of 20.8% (Table I, Panel D). In periods of higher inflation volatility, such as the full period 1964-2004, for which the difference between the 10th and 90th percentile of life-time average inflation is 3.2% (Table I, Panel A), the effect could be even bigger. Unfortunately, detailed data on bond holdings is not available prior to 1983.

IV. Methodological Variations and Robustness Checks

We check the robustness of our results to several variations in methodology, reported in Table VI. We focus on the stock-market participation model for which we can draw on the biggest sample going back to 1964, because standard errors tend to increase for many of these variations when we include additional variables in the model. In addition, we also fix the weighting parameter at $\lambda = 1.00$, approximately equal to our earlier estimates in Table III. Fixing the weighting parameter greatly reduces the computational challenges in estimating the Probit model and makes it feasible to include age dummies or cohort dummies, for example. For brevity, we report only the estimates of the coefficient β ; the parameter estimates are in general similar to the earlier ones in Table III, which implies that the average partial effects are similar, too. We also estimate all specifications in Table VI with observations weighted with SCF sample weights, with similar results.

Age dummies. In our main tests, we controlled for life-cycle effects using a third-order polynomial in age. To check whether this approach misses some aspects of life-cycle patterns, Column (i) in Table VI reports estimates where we use a full set of age dummies instead of the third-order polynomial. Compared with the earlier estimate of 7.053 for the coefficient β in Table III, the estimate of 6.927 (s.e. 1.226) in Table VI is almost identical. Thus, the third-order polynomial appears to control appropriately for any age effects.

Cohort dummies. Our main explanatory variable, the life-time (weighted) average return, not only varies across, but also within cohorts. This variation within cohorts is neither fully captured by age effects nor by time effects. This means that we can, in principle, identify the effect of life-time average return on risk-taking just from within-cohort variation. Therefore, in Column (ii) of Table VI we add cohort dummies to the model. As discussed in the introduction, we cannot add a complete set of cohort dummies, because we would have linear dependence between age, time, and cohort dummies. However, we can add cohort dummies just up to the point that there is no exact linear dependence, which means that we have to drop one of the cohort dummies. This is innocuous since we are not interested in estimating and interpreting the age, time, and cohort effects. Thus, we can use the maximum possible explanatory power of the age, time, and cohort controls. As Column (ii) in Table VI shows, including the cohort dummies has little effect on the point estimate for β (7.911; s.e 1.971), but standard errors are quite a bit higher, reflecting the fact that we are using only within-cohort variation to identify the effect. The results show that our findings cannot be explained by some unobserved correlated cohort fixed effects. This also highlights the advantage of our approach of looking at a specific hypothesis of how risk-taking should vary across and within cohorts, as opposed to attempting to infer cross-cohort variation in risk-taking from estimated cohort fixed effects. If we try to introduce dummies in our risky-asset share or elicited risk-aversion regressions, where we have a much shorter sample, and hence less within-cohort variation in life-time average returns, standard errors are too large to obtain meaningful results.

Controlling for risk. So far we focused on life-time experience of (weighted) average returns. It is conceivable that differences in life-time experiences of return volatility also lead to differences in risk-taking. The case for volatility experiences to matter for risk-taking is somewhat weaker, though, at least to the extent that the channel is beliefs. Since unconditional mean of returns is much harder to estimate than the second moment (Merton, 1980), there is presumably more scope for investors to disagree and be influenced by life-time experiences of mean returns rather than volatility. Our empirical estimation confirms this prior. Table VI reports results on a Probit model where we included the life-time volatility of returns, measured by the standard deviation of returns since birth, with observations weighted in the

same way (with $\lambda = 1.00$) as for the life-time average return. Column (iii) shows that the point estimate of the coefficient on life-time volatility is insignificantly negative and small. Thus, the (weighted) average return appears to be the more relevant summary measure of life-time experience. This does not rule out the possibility that experience of extreme events, in particular extreme downside events, may have some effect on risk-taking of households, nor that households consider the volatility of stock returns when making investment decisions. Our results say that differences in experienced volatility does not seem to explain *differences* in risk-taking between individuals.

More precise split of mutual fund holdings. As mentioned in Section II, the post-1989 data allow us to split non-money market mutual fund holdings into mutual funds that hold mostly stocks and those that do not (“combination” mutual funds are assumed to be split half-half between stocks and bonds). Our main estimation uses the full sample (allowing us to control for age and to exploit variation in life-time average returns unrelated to age) and, hence, cannot differentiate between different types of mutual funds. In order to check whether the results are robust, we re-run all of our tests on the short 1989-2004 sample, using the improved definition of risky asset shares, with mutual fund holdings split between stock funds and bond funds. We obtain similar results, albeit with less precision. For example, Column (iv) of Table VI shows a point estimate of $\beta = 8.630$ (s.e. 2.325) for the stock-market participation model, with λ set to 1.00. Hence, the more precise split of mutual fund holdings has little effect on the results.

Including retirement assets. Starting in 1989, the SCF offers information on the allocation of retirement assets (IRA, Keogh, 401(k), etc.) between stocks and other assets. The allocation information is based on very coarse categories of “most,” “some,” or “no” assets being in stocks. Nevertheless, we check whether an alternative definition of stockholdings that includes retirement assets invested in stocks makes a difference in the stock-market participation regressions. We find that the results are similar to before. As shown in Column (v) of Table VI, we obtain the coefficient estimate $\beta = 7.899$ (s.e. 2.243), very close to the estimate in Column (iii), and the estimates in Table III.

Restricting the sample to older investors. One potential concern is that the results are driven by the behavior of young, inexperienced investors and that the experience effects might be minimal among older investors. To check, we re-estimate our models including only investors older than 49 years. For the stock-market participation probit model of Table III, Column (ii), we obtain $\beta = 9.254$ (s.e. 5.441) and $\lambda = 2.950$ (s.e. 1.471). Hence, the same life-time experience effects as in the full sample appear in the sample of older investors, though due to the smaller sample the standard error for β is considerably larger than for the full sample. The point estimate for λ is higher than in Table III, suggesting that for older investors weights might decline somewhat faster, although the difference between the estimates is less than two standard errors. For the risky asset share regressions, we obtain $\beta = 1.446$ (s.e. 1.730) and $\lambda = 1.075$ (s.e. 0.052). Again, the point estimates are similar to the full sample regressions, although the standard error for β is too high in the restricted sample to assess the magnitude of the effect with much confidence. The risky asset share regressions do not suggest that weights of older investors decline faster.

Flexible starting point of return history. In our estimation so far, we assumed that the starting point for our weighting function is the birth year. Returns before birth therefore receive zero weight, and returns realized after birth receive weights greater than zero. This assumption is not crucial, because our flexible weighting function can accommodate high as well as low weights in the early years. If investors are not influenced much by the returns from, say, the first two decades of their life, then our estimated weighting function have relatively quickly decaying weights. Moreover, if we misspecified the starting point, this would make it less likely that we would find a significant effect of life-time weighted average returns on risk-taking. By specifying a better starting point, we can only do better. Nevertheless, to get a better sense of the weighting of experiences early in life, we now consider a flexible starting point. We treat it as a parameter to be estimated and we allow it to be before or after birth.

We find that if we have both λ and the starting point as free parameters, it is difficult to achieve convergence of the ML estimator. The reason is that the pattern of weights produced with, say, $\lambda = 1.0$ and starting point set at birth is quite similar to the pattern of weights one obtains with $\lambda = 1.5$ and

starting point about 10 years before birth. In other words, increasing λ and pushing the starting point earlier have roughly offsetting effects on the life-time weighted average return. The ML algorithm then has difficulty finding a unique combination of λ and the starting point that maximizes the likelihood. This simply reflects the fact that where exactly the starting point is set is not very important, at least in our case, where the estimated weights decline towards birth. Choosing an earlier starting point would lead to a higher estimate for λ , choosing a later one would lead to a lower estimate of λ .

If we fix λ at 1.0, i.e. roughly the point estimate from the models in Table III, and repeat the regressions of Column (ii) in Table III, but now with a flexible starting point, our estimated starting point is at an age of 3 years (s.e. 0.258), not very different from our assumption that the birth date represents the starting point. With 5.926 (s.e. 1.043), the estimate of β is also only slightly different from that in Table II. If we fix λ at a higher value of 1.29 (the maximum estimate across the different specifications in Table II), the estimate for the starting point is 13 years before birth (s.e. 1.73). This shows that, indeed, λ and the starting point have offsetting effects. The data can be fit quite well with both combinations of parameters.

It is clear, though, that it is difficult to measure precisely the weights that apply to experiences from the early years in life. Our estimates suggest that they are relatively low (for a 30-year old, and $\lambda = 1.0$, the total weight on the first 10 years is about 10%), and there is considerable uncertainty about λ and the starting point of personal experience. Our results suggest with a high level of confidence that life-time experiences have a strong effect on risk-taking, but we have less confidence about how strongly the early years in life matter. It is possible that they do not matter much, which may seem plausible, given that children and teenagers are unlikely to follow the stock market or news about inflation. It is also possible, though, that the early years do matter to some extent, perhaps because risk attitudes are partly acquired through social channels from parents, older siblings, or other close individuals.

V. An Aggregate Perspective

Our microdata estimates suggest that investors' willingness to take risks depends on the personally experienced history of stock-market returns and inflation. These experience-based changes in risky asset demand could in turn influence the dynamics of stock prices. To provide some perspective on this, we perform a simple aggregation exercise: We aggregate the life-time average returns in each SCF survey year and relate it to the aggregate demand for risky assets. Since life-time average return is correlated with risky asset demand at the micro-level, the aggregated life-time average return should be correlated with aggregate risky asset demand.

We calculate the life-time average returns for each household in each wave of the SCF, using a weighting parameter of $\lambda = 1.25$ (which is roughly the average of the weighting parameter estimates across all our specifications). Then we compute the weighted average across households, where the weight of each household is proportional to the liquid assets of the household (higher financial wealth means proportionally higher impact on aggregate demand for stocks) multiplied with the SCF survey weight. The result is shown in Figure 4. Each bar represents the aggregated life-time average returns of U.S. investors in each one of the SCF survey years that we use in our study.

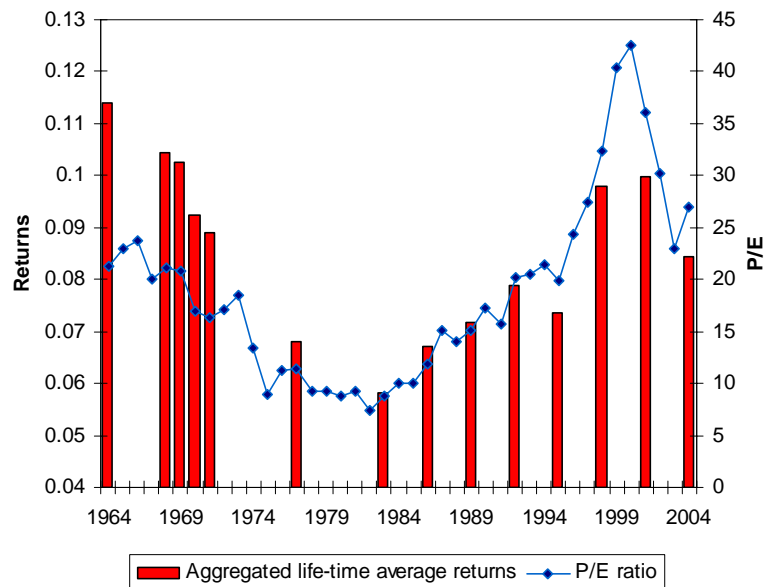


Figure 4: Aggregated life-time average returns ($\lambda = 1.25$) and equity market valuation.

For comparison, Figure 4 also plots the annual price-to-earnings (P/E) ratio from Shiller (2005), which uses a ten-year moving average of earnings in the denominator and which is known to be negatively related to future stock-market returns. The two series are highly positively correlated. Periods of high equity market valuations (the 1960s and 1990s) coincide with periods when investors have high life-time average returns, and periods of low valuation coincide with investors' having low life-time average returns (late 1970s and early 1980s).

To interpret this correlation, it is noteworthy that the distribution of liquid asset across age groups did not change much over our sample period. As plotted in Figure 5, the only exception is the year 1964, because the wealth definitions in the SCF in that year differ from those of other years. (Recall that this is also the reason why we interacted the 1964 wealth variables with a 1964 dummy in all of our regressions). We conclude that the movements in the aggregated life-time average returns over time, as shown in Figure 4, come from changes in investors' life-time return histories rather than compositional changes of age-specific wealth in the population.

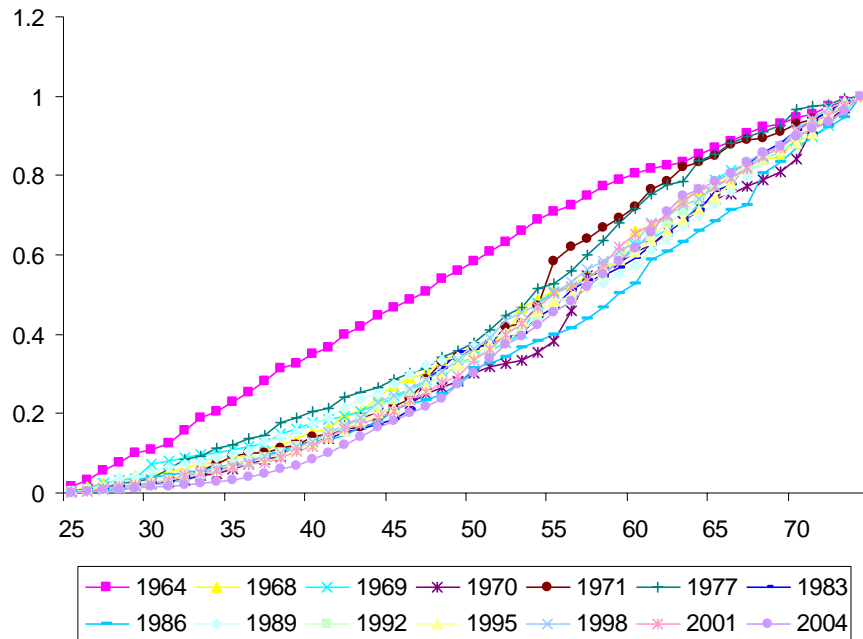


Figure 5: Distribution of liquid assets across age groups in each survey year. For each survey year, we calculate the proportion of total aggregate liquid assets held by each age group, weighted by SCF sample weights. The graph shows the cumulated proportions from age 25 to 74.

This relative stability of the liquid asset distribution suggests that we could extend our calculation of the aggregate life-time average returns before 1964, if we are willing to make the assumption that the liquid asset distribution was stable before 1964, too. Once we have the liquid asset distribution across age groups, we can aggregate the life-time average returns. Figure 6 presents the results from this exercise, assuming that the liquid asset distribution in all years from 1946 to 2004 is equal to the average of the liquid asset distributions of the years 1968-2004 in Figure 5.⁸ We then calculate the aggregate life-time average returns and compare them with the P/E ratio in each year. As one can see from Figure 6, we obtain the same pattern as with the more limited sample in Figure 4: there is a strong positive correlation (0.54) between the two series.

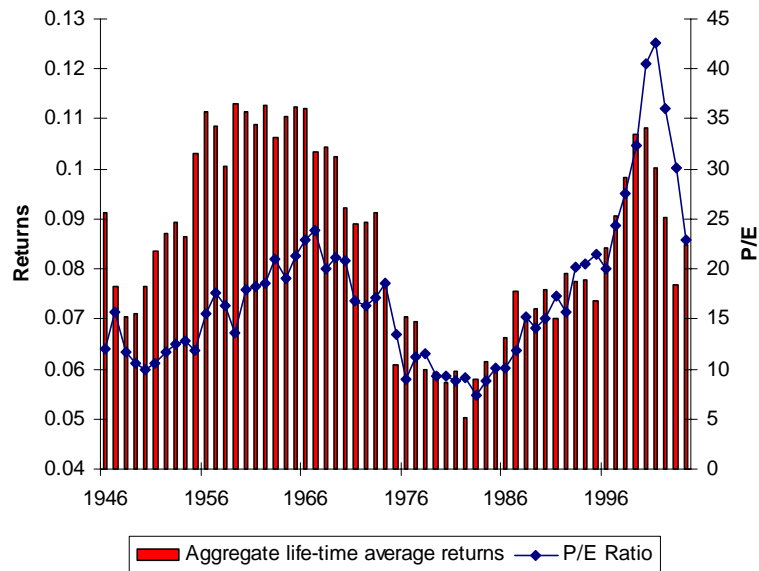


Figure 6: Aggregated life-time average returns ($\lambda = 1.25$) and equity market valuation with extended data series 1946-2004.

Note that this correlation does not simply reflect the well-known correlation between P/E ratios and past returns. We estimate the weighting parameter λ from cross-sectional *microdata* on investors' risk-taking decisions. It is not chosen to match movements in the P/E ratio over time. For example, the

⁸ We cannot go back further than 1946 because the oldest investors in our sample are 74 years old and the stock return data starts in 1871.

weighting parameter estimated from the microdata could, in principle, 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 (recall Figure 2). In that case, the aggregate life-time average return would have been uncorrelated with recent stock-market returns and the time pattern of the bars in Figures 4 and 6 would look very different.

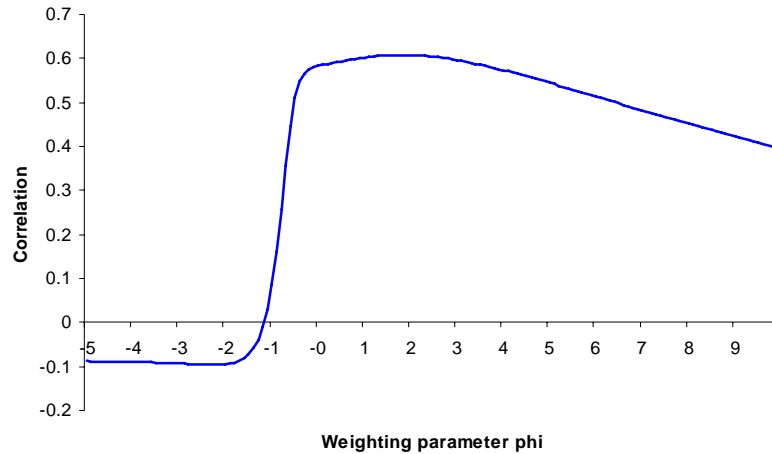


Figure 7: Correlation between aggregated life-time average returns and P/E ratio for different choices of the weighting parameter λ .

This point is underscored in Figure 7. The figure shows the correlation between aggregated life-time average 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 if λ had turned out differently. The value of $\lambda = 1.25$ is actually close to the maximum in Figure 7. And the range of point estimates between 1.0 and 1.5 that we obtained in most of our estimated models all yield a high correlation around 0.6.

The high correlation between aggregate life-time average 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 the possibility that personally experienced stock-market returns affect equity valuation 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.

VI. Conclusion

Our results show that stock returns and inflation experienced over the course of an individual's life have a significant effect on the willingness to take financial risks. Individuals that have experienced high stock-market returns report lower aversion to financial risks, are more likely to participate in the stock market, and allocate a higher proportion of their liquid asset portfolio to risky assets. Individuals that have experienced high inflation invest a lower proportion of their non-stock liquid assets in bonds. Differences between cohorts in their life-course stock return and inflation experiences appear to strongly predict heterogeneity in willingness to bear financial risks at a given point in time and controlling for age, wealth, income, and other demographics. While individuals put more weight on recent stock-market returns and inflation than on more distant realizations, the impact fades only slowly with time. According to our estimates, even experiences several decades ago still have some impact on current risk-taking of older households.

Our results are consistent with the view that economic events experienced over the course of one's life have a more significant impact on individuals' beliefs or preferences than historical facts learned from summary information in books and other sources. If all investors at a given point in time were influenced by the same set of historical data, and all placed the same weight on past return observations, then that effect would be absorbed by the time dummies in our regressions. It is the differential weighting of returns in the past by investors of different age that our life-time average return and inflation variables pick up in our regressions. Such dependence on "experienced data"—as opposed to "available data" in standard rational and boundedly rational learning models—could have important implications for both explaining heterogeneity between economic agents at the micro-level and the dynamics of asset prices at the macro level.

References

- Acemoglu, Daron and Joshua Linn, 2004. Market Size in Innovation: Theory and Evidence From the Pharmaceutical Industry, *Quarterly Journal of Economics*, 119, 1049-1090.
- Ameriks, John, and Stephen Zeldes, 2004. How Do Household Portfolio Shares Vary With Age? *Columbia University Working Paper*.
- Ang, Andrew and Angela Maddaloni, 2005. Do Demographic Changes Affect Risk Premiums? Evidence from International Data, *Journal of Business*, 78, 341-380.
- Auerbach, Alan and Ronald Lee (eds), 2001. *Demographic Change and Fiscal Policy*. Cambridge University Press, Cambridge.
- Bound, John and Sarah Turner, 2003. Cohort Crowding: How Resources Affect Collegiate Attainment? *University of Michigan Working Paper*.
- Bowles, Samuel, 1998. Endogenous Preferences: The Cultural Consequences of Markets and other Economic Institutions, *Journal of Economic Literature* 36, 75-111.
- Brunnermeier, Markus, and Stefan Nagel, *forthcoming*. Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation, *American Economic Review*.
- 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.
- Card, David, and Lemieux, Thomas, 2000. Dropout and Enrollment Trends in the Post-War Period: What Went Wrong in the 1970s? In: J. Gruber (ed.), *An Economic Analysis of Risky Behavior Among Youth*. University of Chicago Press, Chicago, 439-482.
- 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, 2005. The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?, *New York University Working Paper*.
- Choi, James; Laibson, David, and Andrew Metrick, 2002. Does the Internet Increase Trading? Evidence from Investor Behavior in 401k Plans, *Journal of Financial Economics* 64, 397-421.
- Curcuro, Stephanie, John Heaton, Deborah Lucas, and Damien Moore, 2004. Heterogeneity and Portfolio Choice: Theory and Evidence, *Handbook of Financial Econometrics*, forthcoming.
- Deaton, Angus, 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Johns Hopkins University Press, Baltimore MD.
- DellaVigna, Stefano and Joshua Pollet, 2007. Demographics and Industry Returns, *American Economic Review*, 97, 1667-1702.

- 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.
- Easterlin, Richard A., 1987. *Birth and Fortune: The Impact of Numbers on Personal Welfare*. University of Chicago Press, Chicago.
- Faig, Miquel, and Pauline Shum, 2006. What Explains Household Stock Holdings?, *Journal of Banking and Finance* 30, 2579-2597.
- Fair, Ray C. and Kathryn M. Dominguez, 1991. Effects of the Changing U.S. Age Distribution on Macroeconomic Equations, *American Economic Review* 81, 1276-1294.
- Friedman, Milton, and Anna J. Schwartz, 1963. *A Monetary History of the United States, 1867-1960*, Princeton University Press, Princeton, NJ.
- Geneakoplos, John, Michael J.P. Magill and Martine Quinzii, 2004. Demography and the Long-Run Predictability of the Stock Market, *Brookings Papers on Economic Activity*, 1, 241-307.
- Goyal, Amit, *forthcoming*. Demographics, Stock Market Flows, and Stock Returns, *Journal of Financial and Quantitative Analysis*.
- 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, 2007. Inexperienced Investors and Bubbles, *Stanford University Working Paper*.
- Gruber, Jonathan and David Wise, 1999. Social Security Programs and Retirement Around the World, *Research on Labor Economics*, 18, 1-40.
- Hertwig, Ralph, Greg Barron, Elke U. Weber, and Ido Erev, 2004. Decisions from Experience and the Effect of Rare Events, *Psychological Science* 15, 534-539.
- Hurd, Michael D., 1990. Research on the Elderly: Economic Status, Retirement, and Consumption and Saving, *Journal of Economic Literature*, 28, 565-637.
- Merton, Robert C., 1980. On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics*, 8, 323-361.
- Nisbett, Richard E., and Lee D. Ross, 1980. *Human Inference: Strategies and Shortcomings of Social Judgment*, Englewood Cliffs, N.J.: Prentice Hall.
- Palacios-Huerta, Ignacio and Tano J. Santos, 2004. A Theory of Markets, Institutions, and Endogenous Preferences, *Journal of Public Economics* 88, 601– 627.
- Piazzesi, Monika, and Martin Schneider, 2006. Inflation and the Price of Real Assets, *University of Chicago Working Paper*.

- Poterba, James M., 2001. Demographic Structure and Asset Returns, *Review of Economics and Statistics* 83, 565-584.
- 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.
- Sahm, Claudia R., 2007. Stability of Risk Preference, *Federal Reserve Board of Governors Working Paper*.
- Samuelson, Paul A., 1969. Lifetime Portfolio Selection by Dynamic Stochastic Programming, *Review of Economics and Statistics* 51, 239-243.
- Schlag, Karl, 1999a. Which One Should I Imitate?, *Journal of Mathematical Economics* 31, 493-522.
- Schlag, Karl, 1999b. Why Imitate, and If So, How?, *Journal of Economic Theory* 78, 130-156.
- Schotter, Andrew, 2003. Decision Making with Naïve Advice, *American Economic Review, Papers & Proceedings* 93, 196-201.
- 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.
- Simonsohn, Uri; Karlsson, Niklas; Loewenstein, George; and Dan Ariely, 2007. The tree of Experience in the Forest of Information: Overweighing Experienced Relative to Observed Information, *Working Paper*.
- 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, *NBER Macroeconomics Annual*.

Table I: Summary Statistics

	10 th ptile	Median	90 th ptile	Mean	Stddev	N
<i>Panel A: All households 1964 – 2004</i>						
Liquid assets	800	12,372	216,060	122,909	726,749	33,955
Income	16,430	48,475	109,705	65,457	177,594	33,955
Life-time avg. stock return ($\lambda = 1.25$)	0.059	0.087	0.110	0.086	0.021	33,955
Life-time avg. inflation ($\lambda = 1.00$)	0.023	0.042	0.055	0.039	0.012	33,955
Stock mkt. participation	0	0	1	0.286	0.452	33,955
<i>Panel B: All households 1983 – 2004</i>						
Liquid assets	745	13,245	167,187	107,953	895,198	24,914
Income	16,422	48,674	121,526	71,833	226,353	24,914
Life-time avg. stock return ($\lambda = 1.25$)	0.054	0.076	0.103	0.079	0.020	24,914
Life-time avg. inflation ($\lambda = 1.00$)	0.041	0.048	0.058	0.048	0.006	24,914
Stock mkt. participation	0	0	1	0.285	0.452	24,914
% Liquid assets in stocks	0	0	0.551	0.120	0.250	24,914
% Non-stock liquid assets in bonds	0	0	0.176	0.056	0.161	24,914
Risk aversion	2	3	4	3.126	0.834	22,537
<i>Panel C: Stock-market participants 1983 – 2004</i>						
Liquid assets	7,780	65,200	539,606	297,979	1,624,036	10,481
Income	29,391	75,654	202,645	121,553	401,129	10,481
% Liquid assets in stocks	0.044	0.378	0.878	0.421	0.303	10,481
% Non-stock liquid assets in bonds	0	0	0.347	0.090	0.199	10,481
Risk aversion	2	3	4	2.792	0.776	9,702
<i>Panel D: Bond market participants 1983 – 2004</i>						
Liquid assets	2,939	30,399	359,625	208,943	1,354,162	8,208
Income	27,152	65,748	159,166	98,919	315,101	8,208
% Liquid assets in stocks	0	0	0.626	0.164	0.262	8,208
% Non-stock liquid assets in bonds	0.005	0.090	0.632	0.208	0.254	8,208
Risk aversion	2	3	4	3.003	0.801	7,443

Note: Sample period runs from 1964 to 2004. All wealth and income variables are deflated by the CPI into September 2004 dollars. Observations are weighted by SCF sample weights.

Table II: Elicited Risk Aversion and Life-time Average Stock Returns, Ordered Probit Model

	(i) 1983-2004	(ii) 1983-2004	(iii) 1983-2004 w/ SCF sample weights	(iv) 1983-2004 w/ SCF sample weights
<i>Ordered Probit coefficient estimates:</i>				
Life-time average stock-market return coefficient β	-4.551 (1.015)	-4.387 (1.017)	-4.990 (1.197)	-4.734 (1.203)
Weighting parameter λ	1.546 (0.355)	1.498 (0.358)	1.815 (0.423)	1.841 (0.443)
Income controls	Yes	Yes	Yes	Yes
Liquid assets controls	-	Yes	-	Yes
Demographics controls	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
<i>Average partial effect of life-time average stock-market return on category probability</i>				
Risk Aversion = 1	0.496 (0.111)	0.476 (0.110)	0.660 (0.158)	0.629 (0.160)
Risk Aversion = 2	0.827 (0.184)	0.783 (0.181)	0.740 (0.178)	0.708 (0.180)
Risk Aversion = 3	0.041 (0.009)	0.037 (0.009)	0.102 (0.025)	0.067 (0.017)
Risk Aversion = 4	-1.364 (0.304)	-1.296 (0.300)	-1.503 (0.360)	-1.405 (0.357)
#Obs.	22,537	22,537	22,537	22,537
Pseudo R ²	0.09	0.10	0.06	0.08

Notes: Ordered probit model estimated with maximum likelihood. Average partial effects are the sample averages of partial effects on each category probability (given the estimated β and λ) evaluated at each sample observation. Demographic controls (coefficients not reported in the table) include a third-order polynomial in age, number of children and number of children squared, as well as dummies for marital status, retirement, race, and education. Standard errors are shown in parentheses.

Table III: Stock-market Participation and Life-time Average Stock Returns, Probit Model

	(i) 1964- 2004	(ii) 1964- 2004	(iii) 1964- 1977	(iv) 1983- 2004	(v) 1983- 2004 w/ SCF sample weights	(vi) 1983- 2004 w/ SCF sample weights
<i>Probit coefficient estimates:</i>						
Life-time average stock-market return coefficient β	6.743 (1.124)	7.053 (1.244)	6.988 (3.078)	7.053 (1.380)	6.053 (1.495)	6.229 (1.639)
Weighting parameter λ	1.290 (0.212)	0.994 (0.184)	1.076 (0.708)	1.190 (0.280)	1.049 (0.373)	0.992 (0.363)
Income controls	Yes	Yes	Yes	Yes	Yes	Yes
Liquid assets controls	-	Yes	Yes	Yes	-	Yes
Demographics controls	Yes	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
<i>Average partial effect of life-time average stock-market return on participation probability</i>						
	1.929 (0.322)	1.719 (0.303)	1.773 (0.781)	1.682 (0.329)	1.579 (0.390)	1.401 (0.369)
#Obs.	33,955	33,955	9,041	24,914	24,914	24,914
Pseudo R ²	0.24	0.35	0.26	0.38	0.17	0.29

Notes: Probit model estimated with maximum likelihood. Average partial effects are the sample averages of partial effects evaluated at each sample observation (given the estimated β and λ). Demographic controls (coefficients not reported in the table) include a third-order polynomial in age, number of children and number of children squared, as well as dummies for marital status, retirement, race, and education. Standard errors are shown in parentheses.

Table IV: Percentage of Liquid Assets Invested in Stocks and Life-time Average Stock Returns

	(i) 1983-2004	(ii) 1983-2004	(iii) 1983-2004 w/ SCF sample weights	(iv) 1983-2004 w/ SCF sample weights
Life-time average stock-market return coefficient β	1.139 (0.485)	1.288 (0.485)	1.062 (0.398)	1.243 (0.389)
Weighting parameter λ	0.967 (0.058)	0.934 (0.061)	0.843 (0.055)	1.428 (0.080)
Income controls	Yes	Yes	Yes	Yes
Liquid assets controls	-	Yes	-	Yes
Demographics controls	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
#Obs.	10,481	10,481	10,481	10,481
R ²	0.07	0.08	0.09	0.10

Notes: Model estimated with nonlinear least squares. Demographic controls (coefficients not reported in the table) include a third-order polynomial in age, number of children and number of children squared, as well as dummies for marital status, retirement, race, and education. Standard errors are shown in parentheses.

Table V: Experienced Inflation and the Percentage of Non-Stock Liquid Assets Invested in Bonds

	(i) 1983-2004	(ii) 1983-2004	(iii) 1983-2004 w/ SCF sample weights	(iv) 1983-2004 w/ SCF sample weights
Life-time average inflation coefficient β	-2.727 (1.145)	-3.453 (1.099)	-3.866 (1.045)	-3.874 (0.936)
Weighting parameter λ	1.008 (0.061)	1.051 (0.055)	0.790 (0.034)	0.965 (0.036)
Income controls	Yes	Yes	Yes	Yes
Liquid assets controls	-	Yes	-	Yes
Demographics controls	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
#Obs.	8,208	8,208	8,208	8,208
R ²	0.12	0.16	0.03	0.10

Notes: Model estimated with nonlinear least squares. Demographic controls (coefficients not reported in the table) include a third-order polynomial in age, number of children and number of children squared, as well as dummies for marital status, retirement, race, and education. Standard errors are shown in parentheses.

Table VI: Methodological variations for Stock-market Participation, Probit Model

	(i) Age dummies instead of age polynomial	(ii) Age dummies and Cohort dummies	(iii) Volatility	(iv) Mutual funds split into bond and stock funds	(v) Mutual funds split and defined contribution accounts included
Sample period	1964-2004	1964-2004	1964-2004	1989-2004	1989-2004
Weighting parameter λ fixed at	1.00	1.00	1.00	1.00	1.00
Life-time average return coefficient β	6.927 (1.226)	7.911 (1.971)	7.000 (1.234)	8.630 (2.325)	7.899 (2.243)
Life-time volatility coefficient			-0.729 (1.455)		
Income controls	Yes	Yes	Yes	Yes	Yes
Liquid assets controls	Yes	Yes	Yes	Yes	Yes
Demographics controls	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes
Age dummies	Yes	Yes	Yes	Yes	Yes
Cohort dummies ^a	-	Yes	-	-	-
#Obs.	33,955	33,955	33,955	19,499	19,704
R ²	0.34	0.34	0.34	0.38	0.38

^aOne cohort dummy dropped to prevent collinearity of cohort, age, and year dummies.

Notes: Probit model estimated with maximum likelihood. Demographic controls (coefficients not reported in the table) include the number of children and number of children squared, as well as dummies for marital status, retirement, race, and education. Life-time volatility is the standard deviation of returns, estimated using the same weights and the same time periods as for the life-time weighted average return. Standard errors are shown in parentheses.