Depression Babies:

Do Macroeconomic Experiences Affect Risk-Taking? [*]

Ulrike Malmendier
UC Berkeley and NBER

Stefan Nagel
Stanford University and NBER

August 2007
First draft: December 2006

Abstract

We investigate whether differences in individuals' experiences of macro-economic shocks affect long-term risk attitudes, as is often suggested for the generation that experienced the Great Depression. Using data from the Survey of Consumer Finances from 1964-2004, we find that birth-cohorts that have experienced high stock market returns throughout their life 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. We also find that cohorts that have experience high inflation are less likely to hold bonds. These results are estimated controlling for age, year effects, and a broad set of household characteristics. Our estimates indicate that stock market returns and inflation early in life affect risk-taking several decades later. However, more recent returns have a stronger effect, which fades away slowly as time progresses. Thus, the experience of risky asset payoffs over the course of an individuals' life affects subsequent risk-taking. Our results 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).

"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' preferences or beliefs? For the generation of "Depression Babies" it has often been suggested that their experience of a large macro-economic shock affected their long-term risk attitude. In this paper, we ask whether people who live through different macroeconomic histories make different risky choices. Standard models in economics typically assume that individuals are endowed with stable risk preferences, unaltered by economic experiences. The standard models also assume that, when forming beliefs about risky outcomes, individuals adhere to Bayesian updating and incorporate all available historical data. The psychology literature, instead, distinguishes between knowledge obtained from "description" via statistical summary information (e.g., in books or via education) and knowledge acquired by personal experience. Data learned from personal experience has a greater influence on personal decisions than knowledge obtained from books or study, consistent with the notion of Depression Babies. The literature also suggests that, when learning from personal experience, recent events tend to get disproportionate weight. Low-probability events are under-weighted until they occur, and over-weighted once they occur (Weber et al. 1993; Hertwig et al. 2004).

In this paper, we analyze whether the distinction between personal experience and knowledge gained from historical resources matters for economic decisions. We examine empirically whether individuals' risk attitudes in financial decisions differ depending on the macroeconomic history they experienced during the course of their life. 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 more recent experiences have a stronger impact and to what extent experiences early in life are formative for long-term risk attitudes.

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 for the overweighting of experience. Rather, we aim at improving our model of risk-taking by exploring the predictive power of life-time experiences.

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 simple cross-sectional comparison provides suggestive evidence of large differences between cohorts and the role of past experiences. Figure 1 shows the stock market participation rates of individuals between the age of 36 and 45 in different birth cohorts. .
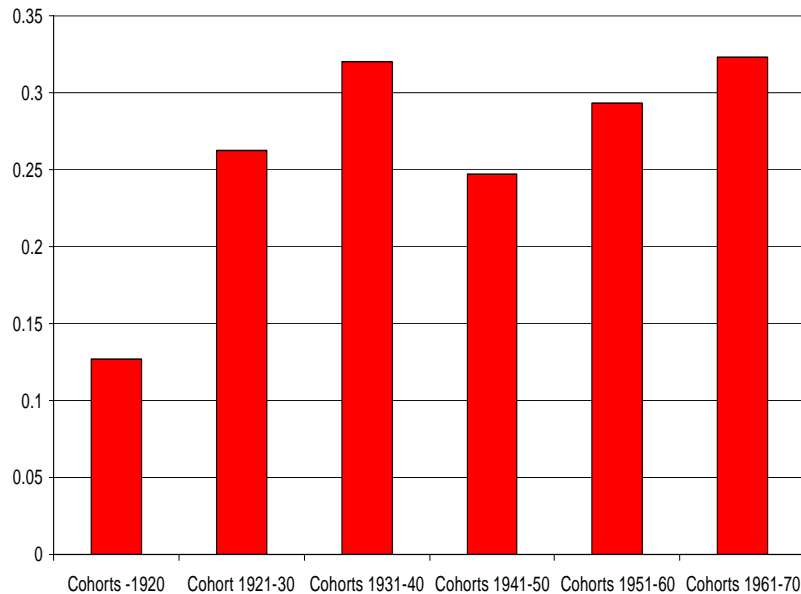
Figure 1: Average stock market participation rates at age 36-45 for different cohorts. Due to data limitations, the average for the cohorts born before 1921 contain only ages 44 and 45, the average for the cohorts 1961-70 contains only ages 36-43.

We observe a stark difference between the participation of the 1920 cohort and later cohorts. The average participation rate for the generation that experienced the 1930s Great Depression as teenagers or adults is

13%, significantly lower than the rates of all other cohorts, which range from 26 to 32%. The 1931-1940 cohort, which experienced the post-war boom years during their young adult life, has a participation rate at age 36-45 that is more than twice as high. For the 1941-50 cohort, the rate dips again, consistent with the fact that this cohort reached age 36-45 just after the depression years of the 1970s. In this paper we test whether these cohort differences persist after controlling for a wide range of demographics and, in particular, whether we can distinguish the effect of return and inflation experience from the other explanations such as year effects.

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 (as opposed to cash, savings accounts, and short-term money market investments). All four measures are likely to reflect a mixture between risk aversion of the Arrow-Pratt type and beliefs about future payoffs on risky investments, especially stocks. Our analysis does not attempt to disentangle the different interpretations. Rather, we ask whether the personal experience of stock market risk predicts the revealed willingness to take financial risk.

To measure households' experience of stock market returns, we calculate, for each birth-cohort in the SCF, the annual real U.S. stock market returns since the birth year of that cohort. To measure households' experience of bond holding, we calculate the annual inflation since the birth year. 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 of bonds have maturities of several years or more. These holdings are risky in real terms, because of future unexpected inflation. Our inflation measure is the Consumer Price Index (CPI) from Shiller (2005) that goes back to 1871.

The extent to which individuals have "experienced" past returns and inflation may, of course, differ depending on their previous investment, their interest in economic matters, and other personal

3

circumstances that we cannot observe. The lack of such controls introduces noise in our explanatory variable, which is likely to attenuate the estimated effects. It would bias our result if the some circumstances are correlated with other cohort-specific factors and stock market returns over time and yet other circumstance are correlated with cohort factors and inflation, a concern we address below.

We relate annual stock returns and inflation since birth to our measures of risk-taking, controlling for age, year fixed effects, wealth and income, and a number of socio-economic characteristics. Since we want to allow for experiences in the distant past to carry a different weight than the more recent experience, we employ a flexible, yet parsimonious weighting scheme. Our weighting function allows the partial effect of past return or inflation experiences on current risk-taking to decline, stay constant, or increase with the time lag, going as far back as the household head's birth date. The weighted averages (i.e., the "life-time average return" and the "life-time average inflation") are obtained by applying the estimated weights for each year to the history of observations since birth. We estimate these weights simultaneously with the sensitivity of households' risk-taking to the resulting weighted average of returns (for risk aversion, stock market participation, and stocks) or inflation (for bonds). In other words, we let the data speak as to how households weight past observations and how strongly their risk-taking is correlated with the weighted average.

We find that households' risk-taking decisions are strongly related to the life-time average return and inflation. Households with higher life-time average returns have lower elicited risk aversion, higher rates of stock market participation, and a higher allocation to stocks. Households with higher life-time average inflation invest fewer liquid assets in bonds. The estimated weights are remarkably similar for all four risk-taking measures. More recent returns and inflation have a somewhat 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.

An important aspect of our analysis is that we control non-parametrically for year effects. Hence, we rule out the possibility that time trends in measured risk-taking or a hard-wired link between asset allocation and stock prices in the aggregate are driving our results. Our identification comes from cross-

sectional *differences* in risk-taking and life-time average returns, and from changes of those cross-sectional differences over time, but 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 effect of the low 1970s stock market returns was moderated by the fact that their life-time experience included 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 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 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). Therefore, all regressions include a third-order polynomial in age and dummies for retirement, ruling out that any time-invariant life-cycle effect explains our findings.

The inclusion of age and year controls comes, necessarily, at the expense of not being able to include cohort dummies (see, e.g., the discussion in Campbell, 2001). Since life-time average returns do not vary within, only across cohorts, our identification strategy does not allow us to distinguish the impact of experienced stock returns and inflation from unrestricted cohort effects. Our approach is to examine *specific* hypotheses relating risk-taking to life-time average return and life-time average inflation, which provide restrictions on the cohort effects that we can test. The downside of this approach is that we cannot control for unobserved, unspecified correlated cohort effects.

Wealth effects are a leading candidate for such correlated cohort 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. While we include wealth and income controls in each regression, unobserved differences in wealth are likely to remain.

Wealth effects are, however, unlikely to explain our bond investment results, because, real wealth is unlikely to be positively correlated with life-time average inflation. Even for the stock-based risk measures, concerns about wealth effects are most relevant only for stock market participation, which is known to be positively related to wealth (see, e.g., Vissing-Jorgensen, 2002). The risky asset share for stock market participants, instead, has been shown not to be wealth-dependent (Brunnermeier and Nagel, 2006; Sahm 2007). Elicited risk attitude also helps address this issue.

Our findings suggest that individual investors' willingness to bear financial risk depends on personal 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 empirically the effect of experienced macro-economic shocks such as the Great Depression, 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 uses a subset of the SCF data employed in this paper. Poterba assumes zero time effects and evaluates the role of age, controlling for cohort effects. Other work links demographic changes and the aggregate demand for stocks and bonds (Goyal, forthcoming; Ang and Maddaloni, 2003; Geneakoplos, Magill, and Quinzii, 2002), and the effect of cohort size on 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, 2003), and a range of macro variables (Fair and Dominguez, 1991). None of these papers, however, consider cohort experiences beyond those induced by size.

Few papers analyze the effect of past cohort-specific experiences. Most closely related is the paper by Greenwood and Nagel (2006), which shows that young mutual fund managers, who had relatively more exposure to technology stocks in the late 1990s, increased their holdings particularly after quarters with high tech stock returns, consistent with our finding that young individuals' allocation to stocks is more sensitive to recent stock market returns than that of older investors. In a similar vein, Vissing-Jorgensen (2002) shows that at the end of 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 which allows us to estimate the long-run effect of stock market returns on risk-taking several decades later and to evaluate the (fading) impact over time. Moreover, our paper improves over previous work by combining results on stock investment with results on bond investment and results based on elicited risk aversion. The synopsis of those sets of results help to address alternative interpretations of previous work.

In addition, a couple of papers, which focus on different issues, include some 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 experience of high inflation at the time, consistent with a high weight being placed on life-course experience. 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' beliefs, 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 is the most comprehensive source of data on U.S. households' asset holdings. The SCF provides repeated cross-section observations on asset holdings and various household background characteristics every few years. 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, where repeated cross-sections are available every three years. We supplement this sample with data from the SCF precursor obtained from the Inter-university Consortium for Political and Social Research at the University of Michigan. The precursor surveys started in 1947, partly annually, but with some gaps in between. We found that the data prior to 1964 is not usable for our purposes since information on stock holdings is completely missing or very crude in many years, and the sampling unit prior to 1964 is the "spending unit" rather than the "family unit" that is used in later years. To ensure some 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, providing a relatively large number of observations on households with substantial asset holdings. 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 SCF[1] to account for the potential bias.

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

---

[1] 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.

returns on the S&P500 index from 1933 to 1982. We use the same approach for annual inflation, using the Consumer Price Index (CPI). Both indices are from Shiller (2005) and go back to 1871.[2]

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 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 mutual funds. The type of mutual funds held is not

---

[2] 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 more recent period). We thank Bob Shiller for providing the data on his website.

specified in the survey waves prior to 1989. (As a robustness checks, we re-run our tests on the 1989-2004 sample including only those mutual funds that invest mostly in stocks.) Our baseline measure also excludes retirement accounts since the survey waves prior to 1989 do not provide information on the composition of assets in retirement accounts (IRA, Keogh, and Thrift 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.

Our third measure of risk taking is the proportion 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. Stock holdings divided by liquid assets yields the proportion invested in stocks.

Our fourth measure of risk-taking is the proportion 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 of bond holdings are in bonds with 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 want to 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 that year the liquid wealth variable it can include some real estate assets, for example).

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.[3] 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 the 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 recently experiences. For example, the 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

---

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

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 on 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

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

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 thus $A_{it}(\lambda)$ 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 $\lambda > 1$). Figure 2 provides an example of the weighting functions for three values of λ for a household of age 50.
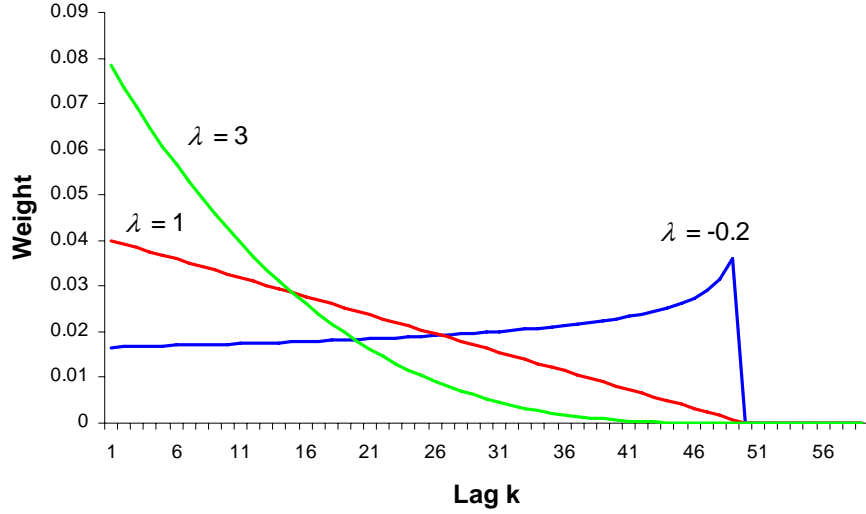
Figure 2: Three examples for the life-time stock market returns weighting function for a household with a 50-year old household head.

We apply the exact same methodology to calculate the weighted average of past inflation.

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:

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

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 those weights with β yields, for a household of that age, the partial effect of a return (or inflation) experienced k years

ago. 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, compared with $65,200 for 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 between the full 1964-2004 and the (full) 1983-2004 are very similar. 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.[4] 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.

---

[4] 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.

14

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 10^{th} and 90^{th} percentiles of 4.4% and 87.8%. The proportion of non-stock liquid assets invested in bonds in Panel D has 10^{th} and 90^{th} percentiles of 0.5% and 63.2%. The 10^{th} and 90^{th} 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 this variation in our 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 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 10^{th} and 90^{th} percentile for the real life-time average stock return are 5.9% and 11.0% in the 1964-2004 sample. The 10^{th} and 90^{th} 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 10^{th} and 90^{th} 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 10^{th} and 90^{th} 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 4 distinct categories, $y_{it} \in \{1, 2, 3, 4\}$. We model the cumulative probability of these ordinal outcomes with an ordered probit model

$$P\left(y_{it} \leq j \mid x_{it}, A_{it}(\lambda)\right) = \Phi\left(\alpha_j - \beta A_{it}(\lambda) - \gamma' x_{it}\right) \qquad j \in \{1, 2, .., 4\}, \tag{3}$$

where $\Phi(.)$ 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, and Φ is the cumulative normal distribution function. Differently from the standard ordered probit estimation, $\Phi(.)$ 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., $\partial P\left(y_{it} = j \mid x_{it}, A_{it}(\lambda)\right)/\partial A_{it}(\lambda)$. We evaluate the partial effects at each sample observation, given the estimated parameters and observations on x_{it} and $A_{it}(\lambda)$, and then we average across sample observations to get the average partial effect.

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, 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, one cannot simultaneously identify age, time, and cohort effects, without further restrictions. Since our model contains the life-time average return along with age controls and year dummies and life-time average returns do not vary within, only across birth-cohorts, we cannot distinguish our estimated effect from unrestricted cohort effects. However, our hypothesis imposes a restriction on cohort effects, which we *can* test. We predict that a specific cohort-level variable (life-time average stock returns) generates cohort-level variation in risk-taking. This breaks the linear dependence between age, time, and cohort effect and makes identification possible (Heckman and Robb 1985).

A second important identification issue is reverse causality. Our aim is to find out whether there is a positive causal effect of past stock-return experiences on current willingness to take risks. However, if variation in stock returns over time is caused, at least partly, by variations in aggregate risk aversion of investors, then the causality could be reverse. For example, as risk aversion of investors goes down, stock prices would rise, 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. Since year fixed effects absorb variation in average risk aversion, the estimated effect of life-time average stock returns reflect only cross-sectional differences in risk taking. The year dummies also absorb all other unobserved 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. The top of the table shows the parameter estimates from the ordered probit model, and towards the bottom we present the average partial effects. Each average partial effect shows how a partial change in $A_{it}(\lambda)$ affects the probability of being in the respective risk aversion category. Column (i) 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 × 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 households' risk aversion is affected by returns many years in the past. 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,

$$P\left(y_{it} = 1 \mid x_{it}, A_{it}\left(\lambda\right)\right) = \Phi\left(\alpha + \beta A_{it}\left(\lambda\right) + \gamma' x_{it}\right), \tag{4}$$

where the binary indicator y_{it} equals 1 if household i has greater than zero stock holdings at time t, and x_{it} is a vector of control variables. 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 so we focus on the partial effect $\partial P\left(y_{it} = 1 \mid x_{it}, A_{it}\left(\lambda\right)\right) / \partial A_{it}\left(\lambda\right)$. 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 (2). Control 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 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, but not the estimated coefficients for the control variables.[5] 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 10^{th} to the 90^{th} 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 personal stock market return experience of different cohorts appears to have a large effect on stock market participation.

As in 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 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).

---

[5] The control variable coefficients 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.

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. The estimates are still in the ballpark of those obtained in the other specifications.

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,

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

where, in a slight abuse of notation, y_{it} now 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.

Table IV reports the results. 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 particularly remarkably since it is a common finding 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, e.g., Vissing-Jorgensen (2002), Curcuru, Heaton, Lucas, and Moore (2004), and Brunnermeier and Nagel (2006) 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 current year going back to zero weight in the birth year. It is noteworthy that 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. Stock market participation and 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 two measures capture a common attitude to financial risks and are subject to a common influence (returns experienced in the past). The similarity in the estimates in the risk aversion model and risky asset share model is 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 life-time average returns, 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 (with $\lambda = 1.25$) across households with the same age, weighted by the SCF sample weights. We then employ a running-mean smoother and plot the smoothed 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 two graphs, in the middle row, we check whether "residual stock market participation" is correlated with life-time average returns, where the residual participation probability is the difference between the actual participation indicator and predicted participation probability and where predicted participation is estimated in the Probit model of Column (ii) in Table III, but without including the life-time average return variable. The residual participation probabilities are averaged by age within each five-year subperiod. The graphs show a monotone relation between the slope of the age profile and the slopes of life-time average returns.
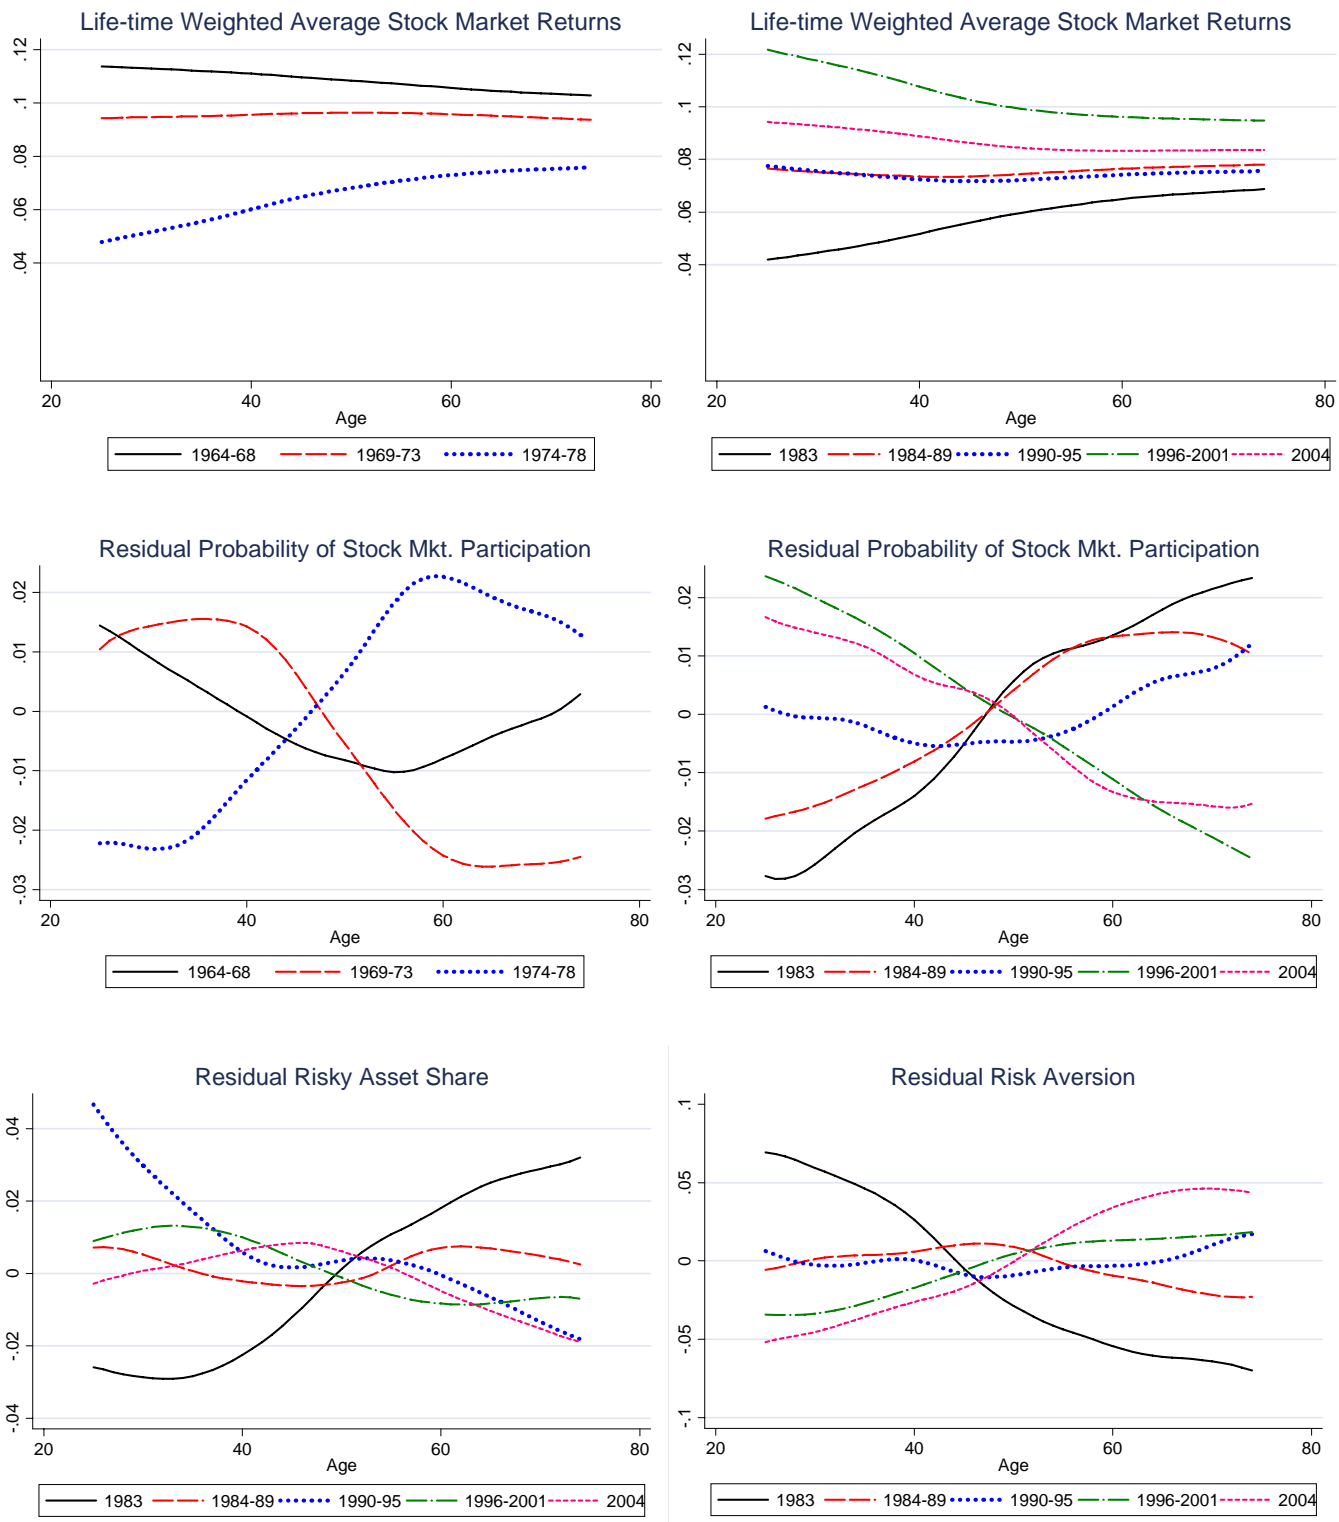
Figure 3: Life-time weighted average stock market returns (with $\lambda = 1.25$) and residual risk-taking measures for as a function of age; smoothed with tri-cube weighted running-mean smoother.

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.

We perform a similar exercise in the bottom row of Figure 3. The left-hand side graph plots residual risky asset shares, based on the residuals from Column (ii) in Table IV, but excluding the life-time average returns from the model. 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.

The right-hand side graph in the bottom row presents the residual risk aversion, obtained as the residual from the ordered probit model in Table II, Column (ii), again with life-time average returns excluded from the model. As one would expect from the estimation results in Table II, the plots show exactly the opposite pattern from the two other risk-taking measures: the age profile is downward sloping in the early 1980s and reverts to an upward sloping relationship in the 1990s. This reflects the fact that high values for elicited risk aversion imply low willingness to take risk, while high stock market participation probability and high risky asset shares imply high willingness to take risks.

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.

24

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's 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. 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 10[th] and the 90[th] percentile of life-time average inflation from Table I, Panel B, is 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 1962-2004, for which the difference between the 10[th]

and 90[th] 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.

F. Methodological variations and robustness checks

We checked robustness of our results to several variations in methodology.

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), while the earlier data do not. Our main estimation uses the full sample which considers all non-money market mutual fund holdings to be stock holdings. The long sample period is important for the estimation of life-time average return effects since it allows us to control for age, and to get variation in life-time average returns unrelated to age. But to check whether the results are robust, we re-ran all of our tests with the shorter 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, re-running the stock market participation probit model of Table III, Column (ii), yields point estimates $\beta = 8.450$ (s.e. 2.334) and $\lambda = 0.963$ (s.e. 0.276). For the risky asset share regressions from Table IV, Column (ii), we obtain point estimates of 0.934 (s.e. 0.761) for β and 1.569 (s.e. 0.127) for λ, which are quite close to those in Table IV.

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, that determine only whether most, some, or none of the assets is in stocks. Nevertheless, we check whether an alternative definition of stockholdings that includes retirement assets makes any difference in the stock market participation regressions. We find that the results are similar to before. For example, repeating the probit regressions from Table III, Column (ii), after including retirement assets and for the time period 1989-2004, we obtain the coefficient estimates $\beta =$

8.248 (s.e. 2.369) and $\lambda = 1.308$ (s.e. 0.326), which are less than a standard error away from the original estimates in Table III, and which imply effects of similar economic magnitudes.

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, but 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 that crucial, because our flexible weighting function permits high or 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 should pick that up as relatively quickly decaying weights. Nevertheless, we check the robustness of our results years other than the birth year being the start year. To do this, we treat the starting point as an additional parameter to be estimated. 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. However, if we fix λ at 1.0, i.e. roughly the point estimate from the models in Table III, we can estimate the starting point. Repeating the regressions of Column (ii) in Table III, our estimated starting point is 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. Summing up, the assumption in the main tests that the starting point of the return history is at the birth does not appear to be at odds with the data.

IV. 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 changes in risky asset demand due to past experiences 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.

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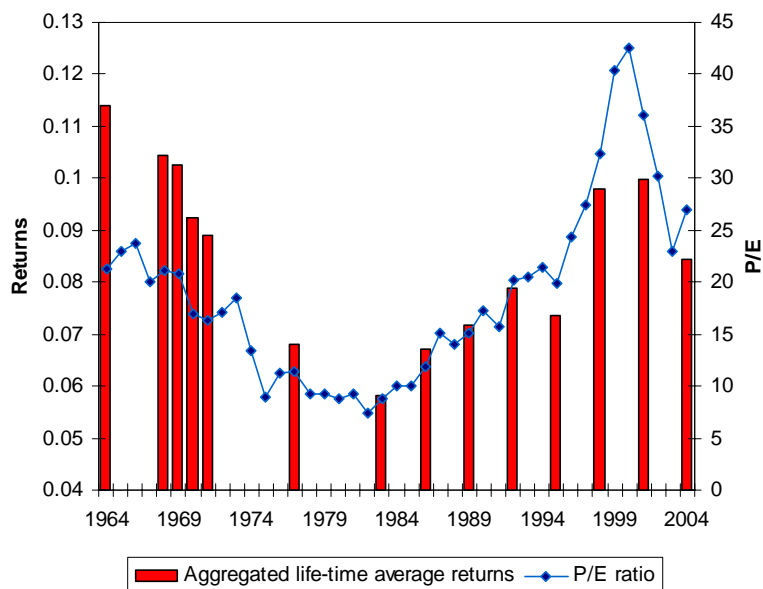
Figure 4: Aggregated life-time average returns ($\lambda = 1.25$) and equity market valuation.

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).

Movements in the aggregated life-time average returns over time come from changes in investors life-time return histories as time progresses and from compositional changes of the population (changes in the liquid asset distribution across age groups). With respect to the latter point, it turns out that the liquid asset distribution across age groups, plotted in Figure 5, did not change much over our sample period. The only exception is the year 1964, because in that year the wealth definitions in the SCF 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).
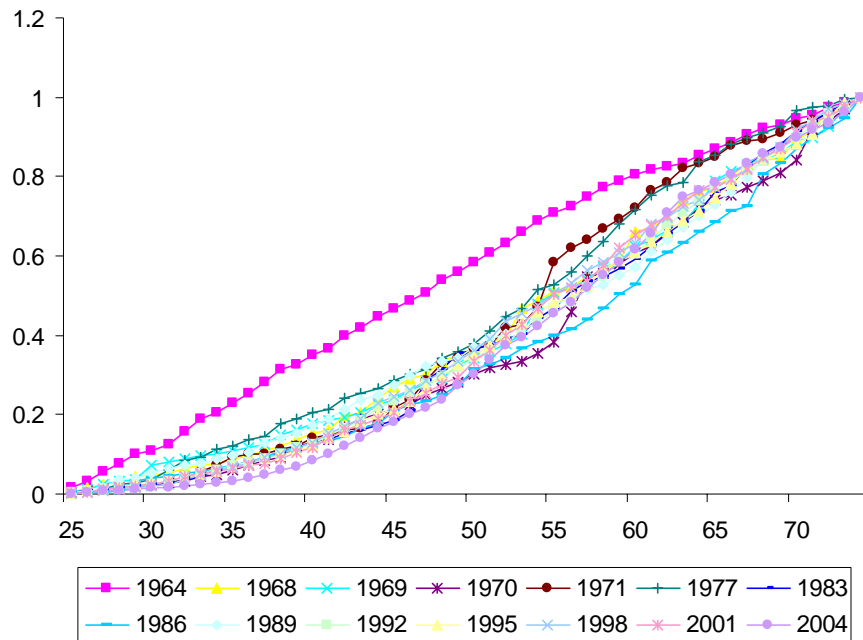
Figure 5: Distribution of liquid assets across age groups in each survey year. For each survey year, we calculate the proportion of total aggregate liquid assets held by each age group, weighted by SCF sample weights. The graph shows the cumulated proportions from age 25 to 74.

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. We assume 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.[6] 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.

---

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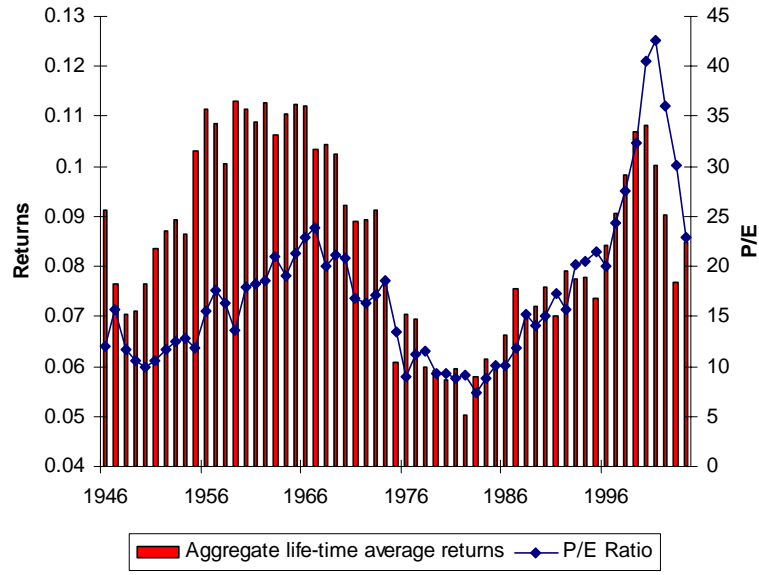
Figure 6: Aggregated life-time average returns ($\lambda = 1.25$) and equity market valuation with extended data series.

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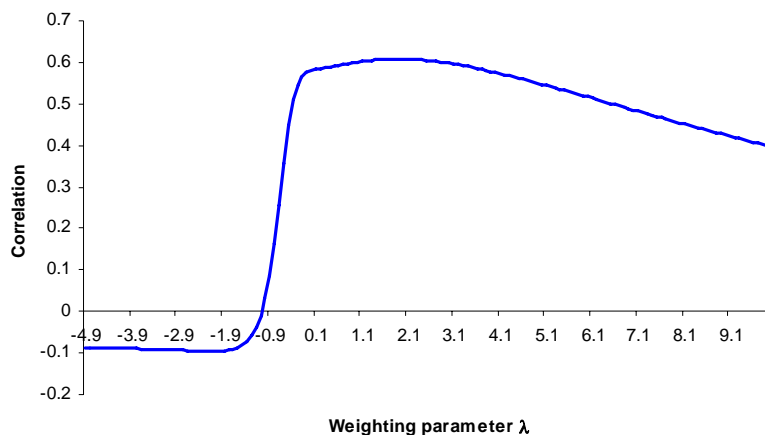
Figure 7: Correlation between aggregated life-time average returns and P/E ratio for different choices of the weighting parameter λ.

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 of λ had turned out differently. The value of λ = 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.

Overall, the high correlation between aggregate life-time average returns and stock market valuation levels adds additional credibility to our microdata estimates, as the estimates imply plausible time-variation in aggregate demand for risky assets. Our results thus raise 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.

V. Conclusion

Our results show that the stock returns and inflation experienced over the course of an individuals' 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 somewhat 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 picks 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, 2003. Market Size in Innovation: Theory and Evidence From the Pharmaceutical Industry, *NBER Working Paper* 10038.

Ang, Andrew and Angela Maddaloni, 2003. Do Demographic Changes Affect Risk Premiums? Evidence from International Data, *ECB Working Paper* 208.

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.*

Brunnermeier, Markus, and Stefan Nagel, 2006. Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation, *Stanford University Working Paper.*

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.

Curcuru, 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 Microeconometric Approach to Development Policy*, Johns Hopkins University Press, Baltimore MD.

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, 2002. Demography and the Long-Run Predictability of the Stock Market, *Yale Cowles Foundation Discussion Paper* 1380.

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, 2006. 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.

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.

Shaw, Kathryn L., 1996. An Empirical Analysis of Risk Aversion and Income Growth, *Journal of Labor Economics* 14, 626-653.

Shiller, Robert J., 2005. *Irrational Exuberance* (second edition), Princeton University Press, Princeton NJ.

Weber, Elke U., Ulf Böckenholt, Denis J. Hilton, and Brian Wallace, 1993. Determinants of Diagnostic Hypothesis Generation: Effects of Information, Base Rates, and Experience, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 1151-1164.

Vissing-Jorgensen, Annette 2002. Towards and Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures, *Northwestern University Working Paper*.

Table I: Summary Statistics

	10th pctile	Median	90th pctile	Mean	Stddev	N
Panel A: All households 1964 – 2004						
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
Panel B: All households 1983 – 2004						
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
Panel C: Stock market participants 1983 – 2004						
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
Panel D: Bond market participants 1983 – 2004						
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
Ordered Probit coefficient estimates:				
Life-time average stock market return coefficient β	-4.551	-4.387	-4.990	-4.734
	(1.015)	(1.017)	(1.197)	(1.203)
Weighting parameter λ	1.546	1.498	1.815	1.841
	(0.355)	(0.358)	(0.423)	(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
Average partial effect of life-time average stock market return on category probability				
Risk Aversion = 1	0.496	0.476	0.660	0.629
	(0.111)	(0.110)	(0.158)	(0.160)
Risk Aversion = 2	0.827	0.783	0.740	0.708
	(0.184)	(0.181)	(0.178)	(0.180)
Risk Aversion = 3	0.041	0.037	0.102	0.067
	(0.009)	(0.009)	(0.025)	(0.017)
Risk Aversion = 4	-1.364	-1.296	-1.503	-1.405
	(0.304)	(0.300)	(0.360)	(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
Probit coefficient estimates:						
Life-time average stock market return coefficient β	6.743	7.053	6.988	7.053	6.053	6.229
	(1.124)	(1.244)	(3.078)	(1.380)	(1.495)	(1.639)
Weighting parameter λ	1.290	0.994	1.076	1.190	1.049	0.992
	(0.212)	(0.184)	(0.708)	(0.280)	(0.373)	(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
Average partial effect of life-time average stock market return on participation probability	1.929	1.719	1.773	1.682	1.579	1.401
	(0.322)	(0.303)	(0.781)	(0.329)	(0.390)	(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	1.288	1.062	1.243
	(0.485)	(0.485)	(0.398)	(0.389)
Weighting parameter λ	0.967	0.934	0.843	1.428
	(0.058)	(0.061)	(0.055)	(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	-3.453	-3.866	-3.874
	(1.145)	(1.099)	(1.045)	(0.936)
Weighting parameter λ	1.008	1.051	0.790	0.965
	(0.061)	(0.055)	(0.034)	(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.