Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Life-Cycle Portfolio Choices and Heterogeneous Stock Market Expectations

Mateo Velásquez-Giraldo

2024-097

Please cite this paper as:
Velásquez-Giraldo, Mateo (2024). "Life-Cycle Portfolio Choices and Heterogeneous Stock Market Expectations," Finance and Economics Discussion Series 2024-097. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.097.

Life-Cycle Portfolio Choices and Heterogeneous Stock Market Expectations*

Mateo Velásquez-Giraldo†

Board of Governors of the Federal Reserve System‡

December 12, 2024

Abstract

Survey measurements of households' expectations about U.S. equity returns show substantial heterogeneity and large departures from the historical distribution of actual returns. The average household perceives a lower probability of positive returns and a greater probability of extreme returns than history has exhibited. I build a life-cycle model of saving and portfolio choices that incorporates beliefs estimated to match these survey measurements of expectations. This modification enables the model to greatly reduce a tension in the literature in which models that have aimed to match risky portfolio investment choices by age have required much higher estimates of the coefficient of relative risk aversion than models that have aimed to match age profiles of wealth. The tension is reduced because beliefs that are more pessimistic than the historical experience reduce people's willingness to invest in stocks.

JEL Codes: G11, G40, G51, G53, E21, D15

*This paper has been supported by the Alfred P. Sloan Foundation Pre-doctoral Fellowship in Behavioral Macroeconomics, awarded through the NBER. I thank Christopher D. Carroll, Nicholas Papageorge, and Francesco Bianchi for their guidance and support in the development of this paper. For their helpful comments, I thank Daniel Barth, Michael Keane, Eva F. Janssens, Bence Bardóczy, Kevin Thom, and Stelios Fourakis. The paper has benefited from discussions in the Economics Department at Johns Hopkins and the Board of Governors of the Federal Reserve. Finally, I thank participants at the 2021 Summer School in Dynamic Structural Econometrics, the 2022 Cherry Blossom Financial Education Institute, the 4th Behavioral Macroeconomics Workshop, the 2022 LACEA-LAMES annual meetings, and the 2023 IAAE annual conference (which also supported this paper with a travel grant).

†`mateo.velasquezgiraldo@frb.gov`.

‡Views and conclusions stated in this paper are the responsibility of the author alone and do not necessarily reflect the views of the Federal Reserve Board or other members of its staff.

1. Introduction

When a canonical model of portfolio choice over the life cycle is calibrated to reproduce the fact that most households do not invest most of their wealth in equities, coefficients of relative risk aversion exceeding 10 are usually required.[1] In contrast, the literature on consumption and saving over the life cycle has found that coefficients of relative risk aversion of 2 or less are able to fit the data comfortably when labor income uncertainty is calibrated to match facts from widely available sources of micro data.[2] This discrepancy generates difficulties for studies that attempt to simultaneously reproduce both groups of facts using life cycle models. Virtually all of such studies calibrate household beliefs about equity returns to match the statistical properties of actual realized returns over long periods of history.[3] The thesis of this paper is that the ability of these models to simultaneously replicate portfolio allocations and savings dramatically improves, and the tension in their parameter estimates is greatly reduced, if they are calibrated using survey measurements of consumers' actual expectations instead of historical data.

To evaluate this proposition, I use a model with two main blocks. The first is a measurement system based on Ameriks, Kézdi, et al. (2020) that I use to infer the distribution of beliefs about equity returns across U.S. households from survey measurements in the Health and Retirement Study (HRS). The second is a life cycle model of saving and portfolio choices that builds on the workhorse model by Cocco, Gomes, and Maenhout (2005). This model adds a monetary cost of entering the stock market, which is common in the literature, and a proportional tax on stock sales that represents early-withdrawal penalties from retirement plans. It also incorporates a bequest motive and age-varying medical

[1] The canonical life-cycle model of portfolio choice is due to Cocco, Gomes, and Maenhout (2005). While the purpose of that paper is not to reproduce empirical portfolio shares, it shows that, even with a coefficient of relative risk aversion of 10 and a conservative equity premium, the model-implied portfolio shares are higher than those empirically observed. The high estimated coefficients of relative risk aversion can be seen, for instance, in Fagereng, Gottlieb, and Guiso (2017), Catherine (2021) and this paper, all of which compare their full models with baseline specifications that build on Cocco, Gomes, and Maenhout (2005).

[2] See, for instance, Carroll (1997), Attanasio et al. (1999), and Gourinchas and Parker (2002).

[3] As measured, for example, by Mehra and Prescott (1985) and Shiller (1990) since the late XIX century.

expense risks to replicate savings late in life.[4] I estimate the life cycle model targeting the age profiles of savings and portfolios in the U.S. and compare the results using beliefs from survey measurements with those obtained using the standard beliefs based on historical data. The fit of the model calibrated with survey measurements of expectations is dramatically better and its parameter estimates are much more plausible. For college graduates, using the beliefs from survey measurements reduces the distance between model-implied and empirical moments from the Survey of Consumer Finances (SCF) by 46 percent[5] and reduces the estimated coefficient of relative risk aversion from 11.4 to 5.1. The high risk aversion required in the baseline calibration produces extremely high precautionary savings which, to match observed wealth, must be offset with an implausibly low time-discount factor of 0.6—suggesting that consumers discount future utility by 40 percent per year. With the beliefs from survey measurements, the discount factor increases to 0.9. Finally, the estimated monetary cost of entering the stock market falls from 1.0 percent of annual income to $ 0.[6]

Several features of the beliefs estimated from survey measurements contribute to the improved performance of the life cycle model. The estimates imply that only 60 percent of high-school graduates and 72 percent of college graduates think there is any equity premium at all. This helps to rationalize the limited rates of stock ownership and their education gaps, and lowers the estimated monetary costs of entry. Among stock owners, the expected risk-adjusted returns are on average 13% and 18% lower than those implied by historical calibrations for high-school and college graduates respectively, allowing the model to match moderate portfolio shares with lower risk aversion.[7] The consequent weakening of the precautionary motive lets the model replicate savings with more patient

[4] See. e.g., De Nardi, French, and Jones (2010) and Ameriks, Briggs, et al. (2020).

[5] Measured by the objective function of the method of simulated moments minimized in estimation. See Section 4 for details.

[6] For high school graduates, the distance to empirical moments falls by 75 percent, the coefficient of relative risk aversion falls from 8.6 to 4.2, the annual preference time-discount factor raises from 0.3 to 0.8, and the monetary cost of entering the stock market falls from 3.1 to 2.5 percent of annual income.

[7] These figures refer to the average Sharpe ratio implied by the estimated beliefs of those who think there is an equity premium, compared to the historical Sharpe ratio of the S&P 500 index.

discounting of future utility. Finally, heterogeneity in the beliefs of stock owners generates considerable dispersion in their portfolio shares, which is an empirical fact particularly difficult to reproduce using historical calibrations of beliefs.

Among the possible ways to reduce the difficulties in modeling savings and portfolios, beliefs about returns have the virtue of being susceptible of estimation from the individual-level measurements of expectations that a growing number of household surveys now include. Manski (2018), Caplin (2021), and Almås, Attanasio, and Jervis (2023) recommend using this type of measurements to resolve the challenge of separately identifying preferences and beliefs from observed choices, which traditional portfolio-choice models circumvent by assuming that beliefs match historical data. The measurements also produce additional empirical facts against which models can be tested. Of particular importance among these facts is that measured expectations predict portfolio allocations, a reality demonstrated by a vast literature and corroborated by this paper.

Because the estimated beliefs differ from the historical experience, individuals with those beliefs would suffer welfare shortfalls if equities continued to perform as they have in the past. These individual welfare shortfalls would be large, quantified as a share of permanent income. They follow a hump shape across the life cycle, starting at less than 3.5 percent at age 24 and peaking at the age of retirement at averages of 8.1 percent for high-school graduates and 14.2 percent for college graduates. Those who do not own stocks due to their pessimistic beliefs suffer the greatest welfare shortfalls. I analyze the variation of these shortfalls across individuals and age groups and relate the model's predictions to findings in the financial literacy literature (Lusardi and Mitchell 2023).

Beyond its effect on the fit of wealth and portfolios, the choice of how to specify beliefs has stark implications for the welfare and counterfactual questions that life cycle models answer. For example, the different preferences that these models need to fit the data under historical and survey-based beliefs translate to different valuations of social programs. Using changes to unemployment insurance and Social Security benefits as illustrative examples, I show that estimates of the welfare effects of policy changes can easily differ by factors of two to ten across models with different specifications of beliefs.

Related literature and contributions

This paper relates and contributes to various groups of studies in household finance and behavioral macroeconomics.

The first group of related studies explores the reasons for the discrepancies between the actual portfolio choices made by households throughout their lives and the predictions of theoretical models like Merton (1969), Samuelson (1969), Viceira (2001), and Cocco, Gomes, and Maenhout (2005). Since the detection of these discrepancies, numerous studies have attempted to address them by adding various features to their models including: more flexible specifications of households' preferences (Gomes and Michaelides 2003, 2005; Wachter and Yogo 2010; Calvet et al. 2021), richer models of labor income and its risks (Chang, Hong, and Karabarbounis 2018; Catherine 2021), and the addition of different costs that could be associated with stock ownership (Khorunzhina 2013; Campanale, Fugazza, and Gomes 2015; Fagereng, Gottlieb, and Guiso 2017). My paper contributes to this literature showing that if beliefs are aligned with survey measurements, the predictions of the model come closer to the actual choices of households. This force can complement the mechanisms identified in this group of papers; for example, my estimates show that entry costs and rebalancing frictions become more powerful when households expect the lower risk-adjusted returns that their responses imply.

The second group of related papers analyzes survey measurements of expectations about future stock market returns.[8] In this literature, a large body of work has demonstrated that expectations are heterogeneous across people and that differences in expectations are predictive of portfolio choices.[9] These facts have been corroborated in multiple

[8] See Hurd (2009) and Manski (2018) for reviews on the measurement of economic expectations in surveys.

[9] See, e.g.:Dominitz and Manski (2007), Hurd, Van Rooij, and Winter (2011), Amromin and Sharpe (2014), Drerup, Enke, and von Gaudecker (2017), Ameriks, Kézdi, et al. (2020), Giglio et al. (2021), and Calvo-Pardo,

surveys, samples, and countries, as well as by using different ways of eliciting expectations. In addition to their predictive power, other important features of measured expectations about stock returns, such as pessimism (Dominitz and Manski 2007; Hurd, Van Rooij, and Winter 2011), socioeconomic gradients (Das, Kuhnen, and Nagel 2020), and rounding (Manski and Molinari 2010; Giustinelli, Manski, and Molinari 2022) have been established. A critical characteristic for modeling these measured expectations is that most of their variation comes from cross sectional differences that persist over time—individual fixed effects. Only a small part of the variation of fixed effects across individuals is explained by sociodemographic characteristics (Giglio et al. 2021). My contribution to this literature is a model that I use to estimate an interpretable representation of the persistent component of beliefs (structural analogues to individual fixed effects) that can be incorporated into life cycle portfolio-choice models. The model, which builds on Kézdi and Willis (2011), Ameriks, Kézdi, et al. (2020), and Giustinelli, Manski, and Molinari (2022), accounts for persistent and heterogeneous rounding patterns, and its estimates capture many of the empirical features of beliefs that have been highlighted in the literature.

The formation and dynamics of expectations are important areas that this paper does not address. It is a well established fact in this domain that experiences—recent and distant, personal and vicarious—have an effect on expectations (see, e.g., Malmendier and Nagel 2011, 2016; Amromin and Sharpe 2014; Greenwood and Shleifer 2014; Coibion and Gorodnichenko 2015; Bailey et al. 2018; Bordalo, Gennaioli, Porta, and Shleifer 2019). However, despite the robustness of this fact and its macroeconomic significance, experiences and common belief revisions (time fixed effects) capture only a small share of the micro-level variation in measured expectations about stock returns (Giglio et al. 2021). Therefore, with the goal of modeling households' individual choices, this paper focuses instead on the persistent components of individual expectations, which capture around half of their variation.[10]

Oliver, and Arrondel (2022).

[10]Campanale (2011), Peijnenburg (2018), and Foltyn (2020) have analyzed portfolio-choice models in

This paper also relates to a growing literature that uses measurements of individual expectations in the estimation of structural economic models.[11] Studies such as Guiso, Jappelli, and Terlizzese (1992), Dominitz and Manski (1997), Lusardi (1997, 1998), and Caplin et al. (2023) examine measurements of households' expectations about their income dynamics, showing that the expectations differ from standard estimates that use administrative data, and using the measured expectations in models of saving decisions and job-transitions. Similarly, measurements of beliefs are increasingly used in models of other economic decisions, such as educational and occupational choices (Arcidiacono et al. 2020; Wiswall and Zafar 2021), parental investments (Almås, Attanasio, and Jervis 2023), and purchase decisions (Erdem et al. 2005). Measured expectations have also been shown to improve the performance of macroeconomic and asset pricing models (Nagel and Xu 2022; Bordalo, Gennaioli, and Shleifer 2022; Bordalo, Gennaioli, Porta, and Shleifer 2024; Bordalo, Gennaioli, Porta, OBrien, et al. 2024; Bianchi, Ilut, and Saijo 2024). However, despite the well documented differences between households' measured expectations and the standard historically-based calibrations, this paper is the first to use survey measurements of expectations about equity returns in a life cycle model to explain the savings and portfolio choices of U.S. households to the best of my knowledge.

Finally, this paper contributes to the literature on wealth differences between socio-demographic groups. Studies in this literature have identified cross-group differences that are difficult to explain using life cycle or permanent income models of consumption and saving. Precautionary savings, differences in time-preference rates and the differential effects of social programs on the incentive to save have been explored as explanations for these difficulties (Carroll 1994; Hubbard, Skinner, and Zeldes 1995; Cagetti 2003). More recently, Lusardi, Michaud, and Mitchell (2017) show that cross-group differences in

which individuals learn about the distribution of risky returns form their experiences. While learning, coupled with participation costs, helps to replicate participation patterns, it does not amend the basic model's prediction about conditional portfolio shares. Therefore, some of the studies rely on additional mechanisms like ambiguity aversion.

[11] See Koşar and O'Dea (2023) for an excellent review of this literature and Manski (2018), Caplin (2021), and Almås, Attanasio, and Jervis (2023) for arguments in favor of this approach.

financial proficiency have the potential to explain a large part of the empirical relationship between savings as a fraction of income and educational attainment, and account for a large share of wealth inequality between groups with different levels of education. I add to this literature by demonstrating that, indeed, when a life cycle model accounts for measurable differences in expectations about asset returns, it can replicate educational differences in wealth and portfolios with much smaller cross-group differences in preferences.

The rest of this paper is organized as follows. Section 2 presents basic empirical facts about U.S. households' wealth, stock holdings, and beliefs about future stock returns. Section 3 presents the model of beliefs and life cycle saving and portfolio choices. Section 4 discusses estimation. Section 5 presents estimates and discusses their implications. Section 6 quantifies the welfare losses that individuals may suffer from misspecified beliefs about future stock returns. Section 7 estimates the welfare effects of counterfactual policy changes under different specifications of beliefs. Section 8 concludes.

2. The Portfolios and Expectations of U.S. Households

This section reviews various empirical facts about stockholding that challenge the predictions of portfolio-choice models that calibrate households' beliefs to match historical returns. Then, using 16 years of measured expectations, I show that U.S. households' subjective distributions of stock returns appear to deviate from the historical distribution of actual stock returns. The differences between measured expectations and the historical distribution of returns have various features that could explain some of the discrepancies between the predictions offered by traditional models and U.S. households' actual stockholding behaviors.

2.1 Aggregate Patterns of Stockholding

This section examines the stock market participation and portfolio choices of U.S. households. I highlight patterns that deviate from the predictions of standard life cycle models of portfolio choice. Deviations include low participation rates and shares of wealth in stocks, and a relatively flat share of wealth in stocks across the life cycle.

To study aggregate patterns in U.S. households' savings and stock holdings, I use the triennial *Survey of Consumer Finances* (SCF). I restrict my analysis to the survey waves between 1995 and 2019.[12] The SCF provides a comprehensive picture of American households' balance sheets, including "summary files" with useful aggregates such as the total financial assets and stock holdings of each surveyed economic unit. These measures include stocks owned both directly and indirectly through, e.g., mutual funds and retirement accounts. Indirect stock holdings are based on respondents' descriptions of the types of assets that a given fund or account invests in. Although account level data can offer more precise measurements of stock holdings (as noted by Parker et al. 2022), the advantage of the SCF lies in its comprehensive coverage of a nationally representative sample of economic units and all their financial accounts.

Table 1 presents summary statistics for the main variables of interest, as well as demographic variables that describe the sample. The statistics are shown for the full set of observations and split by the respondent's highest level of education. The average respondent is 51 years old. Out of respondents, 12% do not have a high school degree, 56% have a high school degree but no college degree, and 32% have obtained a college degree. Due to changes in educational access, those without a high school degree were born, on average, ten years earlier (1948) than those with high school or college degrees (1957 and 1958).

I focus on high school and college graduates only. There are three main reasons why I

[12] The waves before 1995 do not have the question of income "in normal times" that I use in my analysis. I stop at 2019 because it is the latest pre-COVID wave.

Table 1: Summary statistics: main variables of interest

	All		Less than H.S.		High School		College	
Variable	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Age and education								
Birth Year	1,956.00	18.37	1,947.84	20.68	1,956.77	18.16	1,957.73	16.99
Age	51.00	17.16	58.23	18.55	50.06	17.16	49.92	15.94
Less Than H.S.	0.12	0.32	—	—	—	—	—	—
High-School	0.56	0.50	—	—	—	—	—	—
College	0.32	0.47	—	—	—	—	—	—
Income, wealth, and stock ownership								
Income (1000s)	67.42	145.88	30.82	48.65	51.07	61.79	109.60	236.75
Fin. Assets (1000s)	243.02	1,645.72	42.33	324.56	109.67	775.27	550.55	2,686.12
Owns stocks?	0.53	0.50	0.21	0.41	0.47	0.50	0.75	0.43
Stocks/Fin. Assets	0.25	0.31	0.09	0.23	0.21	0.30	0.37	0.32
Cond. stock share	0.46	0.29	0.44	0.31	0.44	0.29	0.49	0.28

The summary statistics in this table come from pooling the observations from the 1995 to 2019 SCF waves. The sample is restricted to respondents above the age of 21 with non-negative financial assets and a stock-share of financial wealth between 0 and 100%. All calculations use survey weights, which are pooled and re-scaled so that the total weight of each wave is the same. The unit of analysis in SCF it is the "primary economic unit". Wealth and income are expressed in 2010 U.S. dollars and were adjusted using the using the CPI index. I refer to individuals that do not posses a high-school diploma or GED as "Less than H.S. " to those with a high-school diploma or GED but no college degree as "High school" and to those with a college degree as "College."

exclude those without a high school degree. First, their number of available observations is much lower than that of high school and college graduates. When grouped into age-bins, as my analysis requires, the number of observations per group becomes too low to produce sufficiently precise estimates of the moments of interest. Second, a large fraction of respondents without a high school degree answer "does not know/refuse" to the probabilistic questions in the *Health and Retirement Study* that I use in my analysis. Third, as Table 1 shows, households in which the respondent does not have a high school degree have low incomes and levels of wealth; as argued by Hubbard, Skinner, and Zeldes (1995) the saving decisions of these households are severely impacted by social programs that I do not model in this paper.

Table 1 reproduces the low rates of stockholding that constitute the "stockholding puzzle" (Haliassos and Bertaut 1995). The table shows that, out of all observed economic units, only 53% own stocks directly or indirectly. This fact is at odds with the prescription that all households should own stocks, which is a feature of frictionless models in which everyone expects the stock market to perform in the future as it has historically. There is a positive correlation between level of education and stock market participation, with the lowest participation rate (21%) observed among those without a high school degree and the highest participation rate (75%) among those with a college degree. Multiple factors could contribute to this correlation. One explanation could be that individuals with higher education are more likely to believe in the existence of an equity premium. Another possibility could be the existence of barriers like monetary participation costs that disproportionately affect those with lower education and income.

The shares of wealth in stocks among those who participate in the stock market shows a weaker relationship with education than the rate of participation, but they are also lower than what baseline models predict. Given the low rates of participation, the mean share of wealth in stocks is also low at 25%. A more meaningful measure is the share of wealth invested in stocks among those who participate, commonly referred to as the conditional share of wealth in stocks. College graduates have a slightly higher average conditional share of 49% compared to that of high school graduates, which is 44%. These levels are low compared to the prescriptions of life cycle portfolio-choice models like that of Cocco, Gomes, and Maenhout (2005) which, even with a relative risk aversion coefficient of 10 and a moderate equity premium, finds an optimal share of wealth in stocks ranging from 60 to 100 percent depending on the age of the household.

The age patterns of stock market participation and conditional wealth in stocks are similar for high school and college graduates. Figure 1 shows stock market participation rates and quartiles of conditional shares calculated on five-year age bins. The figure shows that, despite differences in levels, the stock market participation rates of both high school and college graduates follow a similar inverted "U" shape as they age. For both groups, the highest participation rates occur between the ages of 56 and 60, with 55% for high school graduates and 81% for college graduates. In contrast, the conditional share of wealth in stocks remains relatively stable across different age bins, showing little variation in the

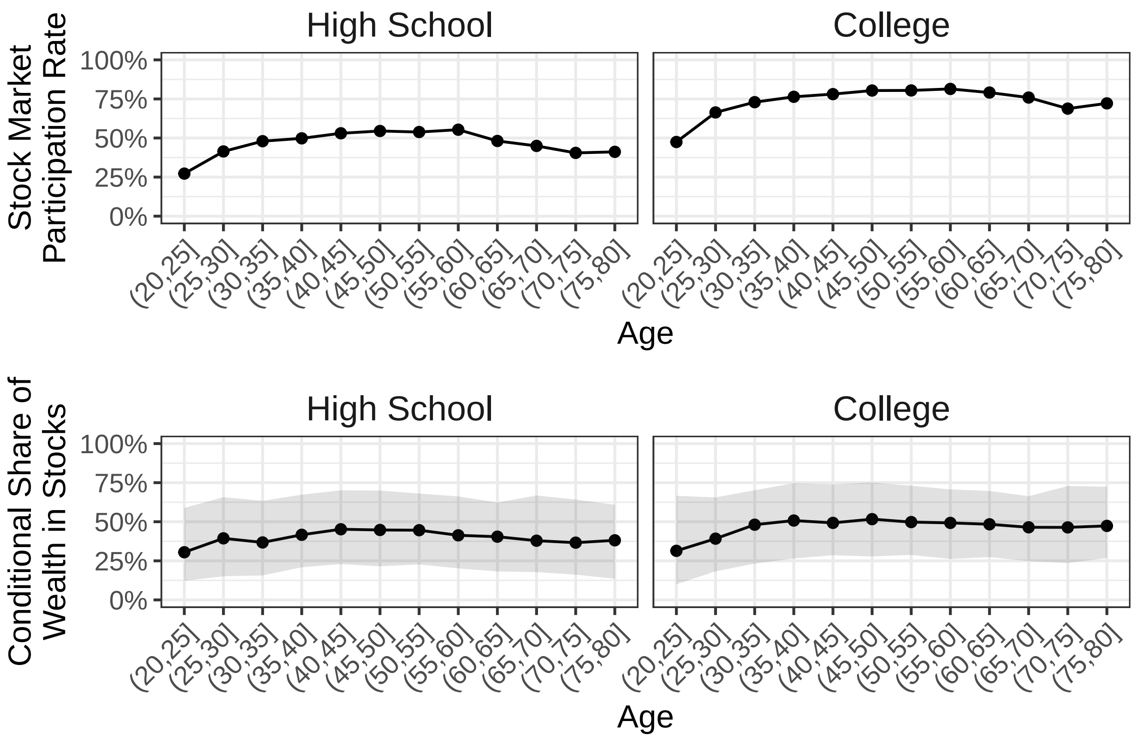

The summary statistics in this table come from pooling the observations from the 1989 to 2019 SCF waves and grouping them into 5-year age bins. For the participation rate, the solid line and points displays the fraction of stock market participants. For the conditional share of wealth in stocks, the solid line and points display the median, and the shaded areas span from the 25th to the 75th percentile. The sample is restricted to respondents with non-negative financial wealth and a stock-share of financial wealth between 0 and 100%. All calculations use pooled survey weights. Wealth is expressed in 2010 U.S. dollars and was adjusted using the using the CPI index.

Figure 1: Stockholding over the life cycle

25th, 50th, and 75th percentiles of its distributions. In every age group, the distributions of the conditional share of wealth in stocks for high school and college graduates are similar to each other.

The stability of conditional shares of wealth in stocks over different age bins in the SCF is inconsistent with traditional portfolio-choice models, which prescribe that these shares must decline with age. This prescription comes from the assumption that a person's future lifetime earnings—their "*human wealth*"—act as a hedge against stock market fluctuations. Therefore, it is optimal for a young person with high human wealth to allocate most of his investable wealth to stocks and to reduce his exposure as he ages and his human wealth decreases. Indeed, in Cocco, Gomes, and Maenhout's (2005) benchmark calibration, young agents invest 100% of their wealth in stocks and gradually lower this share as they age until around 60 percent. Parker et al. (2022) show that the increasing popularity of target-date funds has brought the conditional shares of recent cohorts more in line with the declining patterns prescribed by life cycle models.

2.2 Households' Expectations About Stock Returns

Survey measurements of people's expectations about the future performance of the stock market vary substantially across individuals and deviate from historical benchmarks. Compared to the historical experience, the average person underestimates the probability of positive returns and overestimates the probability of extreme returns (positive and negative). The magnitude of these deviations decreases with education. A large fraction of the variation in these survey measurements corresponds to persistent heterogeneity in people's expectations and this heterogeneity robustly associates with differences in their portfolio choices.

To characterize U.S. households' perceptions about the future performance of the stock market, I use the *Health and Retirement Study* (HRS). The HRS is a biennial longitudinal survey of U.S. adults over the age of 50 that gathers detailed information on respondents' health, financial status, employment, and expectations. I restrict my sample to the "financial respondent" of each household: the person that answers most questions about the income and assets of the household.[13]

Since 2002, the expectations module of the HRS has included questions about the future performance of the stock market. I use the following questions:

- [$P^{\geq 0}$] *"By next year at this time, what is the percent chance that mutual fund shares invested in blue chip stocks like those in the Dow Jones Industrial Average will* ***be worth more*** *than they are today?"*

- [$P^{\geq 20}$] *"By next year at this time, what is the percent chance that mutual fund shares invested in blue-chip stocks like those in the Dow Jones Industrial Average will* ***have gained in value by more than 20 percent*** *compared to what they are worth today?"*

- [$P^{\leq -20}$] *"By next year at this time, what is the percent chance that mutual fund shares invested in blue-chip stocks like those in the Dow Jones Industrial Average will have* ***fallen in value by more than 20 percent*** *compared to what they are worth today?"*

I use $P^{\geq 0}$, $P^{\geq 20}$, and $P^{\leq -20}$ to denote these measurements. The HRS first measured $P^{\geq 0}$ in 2002, with $P^{\geq 20}$ and $P^{\leq -20}$ following in 2008.[14]

Responses to the questions about future stock returns are disperse and their averages deviate considerably from the historical performance of the stock market. Figure 2 depicts the distribution of $P^{\geq 0}$, $P^{\geq 20}$, and $P^{\leq -20}$ across all survey waves and shows that, far from concentrating around estimated answers from a common subjective distribution of returns, the responses span wide ranges without signs of agreement. I compare the responses to annual returns of the S&P 500 index from 1881 to 2018, as reported by Shiller (1990). During this period, the S&P 500 saw positive returns on 72% of the years, returns

[13] I also restrict my sample to those above the age of 50. There is a small number of younger respondents that are interviewed because they are married to someone above the age of 50.

[14] The 2008 wave asked various different combinations of "gain/fall in value by X%" to different individuals. The "gain/fall in value by 20%" versions of the question were incorporated in 2010. I use the individuals who drew $X = 20$ in 2008 to construct $P^{\geq 20}$ and $P^{\leq -20}$ for that wave.

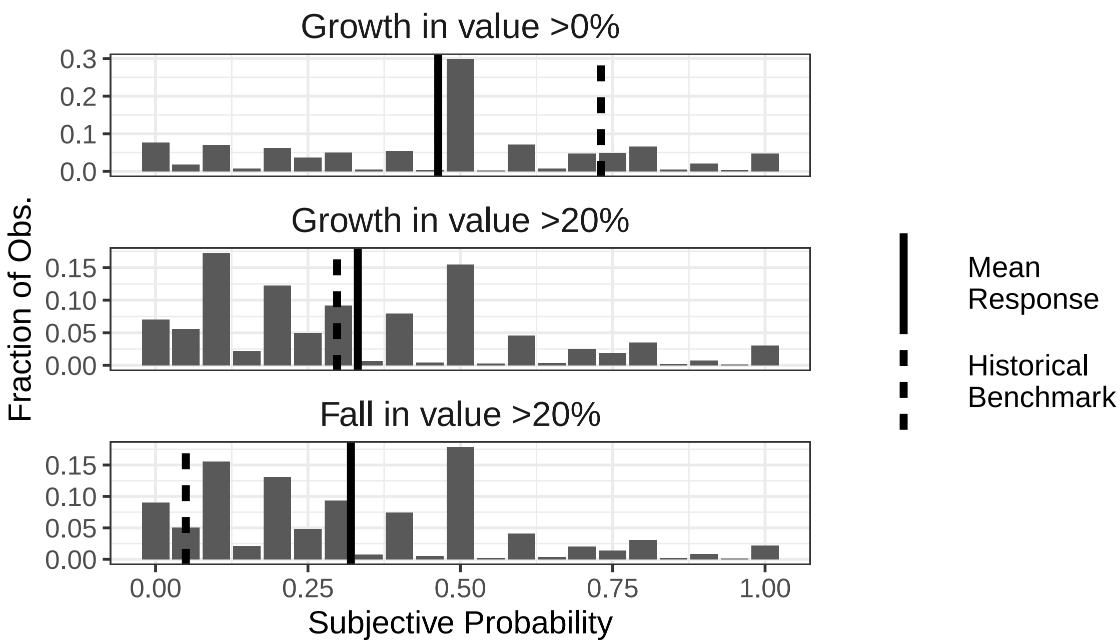

Responses are rounded to the nearest multiple of 5% and each bar reports the fraction of observations corresponding to each multiple. The sample consists of individuals above the age of 50 who report being the financial respondent of the household. The question about positive growth in value (first row) was added in 2002 and the other two were added in 2008; therefore, the samples of both columns do not exactly match. The "historical benchmark" lines correspond to the fraction of years between 1881 and 2018 that the S&P500 index had returns higher that 0%, higher than 20%, and lower than -20%; these calculations are based on the accompanying data file to Chapter 26 of Shiller (1990).

Figure 2: Probabilistic assessments about stock returns

Table 2: Probabilistic assessments about stock returns and education

Question	Mean	St. Dev.	N. Obs	Fract. DK/RF
Less than High School				
$P^{\geq 0}$	0.40	0.30	19,175	0.29
$P^{\geq 20}$	0.38	0.30	5,136	0.05
$P^{\leq -20}$	0.30	0.27	5,401	0.04
High School				
$P^{\geq 0}$	0.45	0.26	59,273	0.14
$P^{\geq 20}$	0.34	0.25	21,341	0.03
$P^{\leq -20}$	0.33	0.25	21,380	0.02
College				
$P^{\geq 0}$	0.53	0.25	23,257	0.06
$P^{\geq 20}$	0.30	0.23	10,594	0.01
$P^{\leq -20}$	0.31	0.21	10,391	0.01

The sample consists of individuals above the age of 50 who report being the financial respondent of the household. The question about positive growth in value (first row) was added in 2002 and the other two were added in 2008; therefore, the samples for the questions do not exactly match. Furthermore, each wave the questions $P^{\geq 20}$ and $P^{\leq -20}$ are asked only to participants who do not answer "does not know/refused" to $P^{\geq 0}$.

greater than 20% on 29% of the years, and returns below −20% on 5% of the years. The average response for the probability of positive nominal returns, which is 46%, is 26 percentage points below its historical benchmark. Conversely, the average responses for the probabilities of extreme returns, which were 33% for $P^{\geq 20}$ and 32% for $P^{\leq -20}$, exceed their respective historical benchmarks by 4 and 27 percentage points. The deviations of average responses from historical benchmarks and households' pessimism about the chances of positive returns in particular are well known facts that have found support across multiple surveys, conducted in the U.S. and abroad (see Hurd 2009; Manski 2018, for reviews).

Expectations about future stock returns show a systematic relationship with educational attainment. Yet, significant variability remains among individuals with the same level of education. Table 2 presents summary statistics of $P^{\geq 0}$, $P^{\geq 20}$, and $P^{\leq -20}$ for respondents with different levels of education. While all groups are pessimistic about the probability of positive returns ($P^{\geq 0}$), the average response increases steeply with education, from 40% for those without a high-school degree to 53% for college graduates. This pattern is consistent with the findings of past studies, which have shown that more educated individuals tend to have a more optimistic outlook on stock returns (Dominitz and Manski 2011; Hurd, Van Rooij, and Winter 2011; Das, Kuhnen, and Nagel 2020). The degree to which the average respondent overestimates the probability of extreme returns also varies with educational attainment, with more educated households generally giving lower responses.[15] However, within-group variability is greater than cross-group differences, as within-group standard deviations are higher than 20 percentage points for all questions.

Table 2 also shows the fraction of participants who refused to answer each of the questions or answered with "do not know." The refusal/unsure rate is much higher for $P^{\geq 0}$ than for $P^{\geq 20}$ and $P^{\leq -20}$ because participants who refuse to answer $P^{\geq 0}$ or answer this question with "do not know" are not asked $P^{\geq 20}$ or $P^{\leq -20}$. While the refusal/unsure rates for high school and college graduates are moderate, it is 29% for $P^{\geq 0}$ for those without a high-school degree. Such a high fraction of refusal/unsure answers casts doubt on the representativeness of respondents without a high school degree who do answer the probabilistic questions. For this reason, in addition to those presented in Section 2.1, I limit the modeling exercises in this paper to high school and college graduates.

Rounding is pervasive in probabilistic assessments about future stock returns. Figure 2 shows that there are large masses of answers in focal points like 0%, 50%, and 100%, and that all multiples of 10% occur more frequently than their neighboring multiples of 5%. Table 3 presents the fraction of responses that belong to different groups of frequent answers. The groups are a similar partition to the one proposed by Giustinelli, Manski, and Molinari (2022). For each question, more than 8% of the answers are multiples of 100% (0 or 100%), more that 80% are multiples of 10%, and more than 97% are multiples of 5%. Using the full expectations module of the HRS, Giustinelli, Manski, and Molinari

[15]The only exception to this pattern is the average $P^{\leq -20}$ of those without a high-school degree, which is the lowest of the three groups.

Table 3: Fractions of "rounded" probabilistic responses

	Fraction of Answers in Group					
Question	$\{0\%, 100\%\}$	50%	$\{25\%, 75\%\}$	Other×10%	Other×5%	Other
$P^{\geq 0}$	0.12	0.30	0.09	0.44	0.05	0.01
$P^{\geq 20}$	0.09	0.15	0.07	0.57	0.09	0.02
$P^{\leq -20}$	0.10	0.18	0.06	0.55	0.08	0.02

The sample consists of individuals above the age of 50 who report being the financial respondent of the household. The question about positive growth in value (first column) was added in 2002 and the other two were added in 2008; therefore, the samples of the three columns do not exactly match. Other×10% = $\{10, 20, 30, 40, 60, 70, 80, 90\}$%. Other×5% = $\{5, 15, 35, 45, 55, 65, 85, 95\}$%. "Other" represents answers that do not fall into any of the other groups.

Table 4: Sources of variation in probabilistic assessments

	Models' R^2		
Question	Time F.E.	Indiv. F.E.	Two-Way F.E
$P^{\geq 0}$	0.01	0.46	0.47
$P^{\geq 20}$	0.01	0.63	0.64
$P^{\leq -20}$	0.01	0.60	0.60

The table reports R^2 statistics from regressing each measurement on individual fixed-effects, month-of-interview fixed-effects, and both. The sample consists of individuals above the age of 50 who report being the financial respondent of the household. The question about positive growth in value (first row) was added in 2002 and the other two were added in 2008, therefore the samples of the three rows do not exactly match.

(2022) show that individuals tend to round questions in the same domain (e.g., health or finances) to consistent levels of coarseness, even though the level of rounding coarseness varies between individuals and across domains. Based on these findings, the model of beliefs that I use in this paper accounts for rounding practices that are stable over time but heterogeneous across individuals.[16]

A large fraction of the variation in probabilistic assessments about stock returns comes from cross-sectional differences between individuals that persist over time. In their analysis of the macroeconomic beliefs of Vanguard account holders, Giglio et al. (2021) show that, for all the measurements of expectations in their analysis, persistent differences in

[16]Survey responses that reflect rounding and the use of heuristics are pervasive in other related contexts such as the hypothetical choice of retirement wealth allocations (Bateman et al. 2017).

expectations across individuals (individual fixed effects) explain a vastly larger fraction of variation than common movements in expectations over time (time fixed effects). I replicate their analysis with the HRS sample, which includes both investors and non-investors and spans a longer period of time. Table 4 presents the R^2 statistics of regressions of the form:

$$\begin{aligned} \text{Indiv. F.E.,} \quad P^x_{i,t} &= a_i + \epsilon_{1,i,t} \\ \text{Time F.E.,} \quad P^x_{i,t} &= b_t + \epsilon_{2,i,t} \\ \text{Two-way F.E.,} \quad P^x_{i,t} &= a_i + b_t + \epsilon_{3,i,t}, \end{aligned}$$

where P^x is one of the probabilistic assessments ($P^{\geq 0}$, $P^{\geq 20}$, or $P^{\leq -20}$), a_i are individual fixed effects, b_t are month-of-interview fixed effects, and $\epsilon_{\cdot,i,t}$ are time-specific idiosyncratic errors. The results are consistent with the findings of Giglio et al. (2021): individual fixed effects explain a much larger fraction of the variance in responses than time fixed effects. For the probability of positive returns $P^{\geq 0}$, individual fixed effects capture 46% of the variance. For the probabilities of extreme returns $P^{\geq 20}$ and $P^{\leq -20}$, they capture 63% and 60% of the variance, respectively. Time fixed effects, in contrast, capture no more than 1% of the variance in any of the questions. As Giglio et al. (2021) point out, the persistent cross-sectional heterogeneity in beliefs manifest in these measurements stands at odds with many of the models used in macroeconomics and finance.

To study the relationship between expectations and portfolios, I restrict the HRS sample further to keep only respondents below the age of 65. The financial literacy literature has found that the age around retirement is when people are best able to respond financial questions (see, e.g., Agarwal et al. 2009; Lusardi and Mitchell 2023); I impose this restriction to limit concerns about cognitive decline. My measure of equity holdings includes both direct and indirect investments.[17] For financial wealth, I use the value of stocks, mutual

[17] For direct holdings, I use the RAND HRS aggregate `astck`. For indirect holdings, I use the total amounts and shares invested in stocks of the households' retirement accounts.

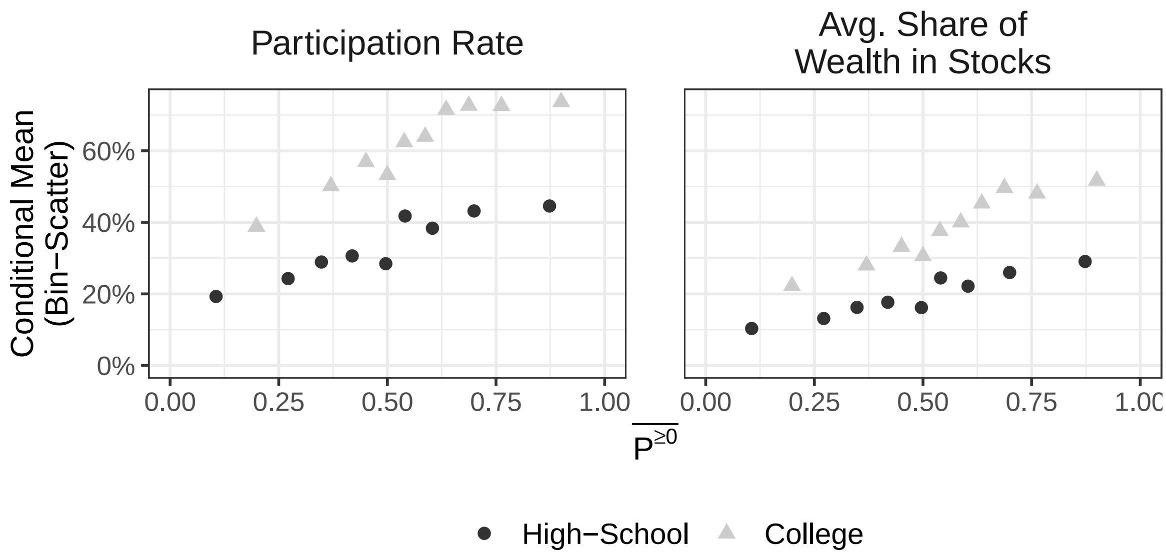

The sample consists of individuals between the ages of 50 and 65 who report being the financial respondent of the household. Binary participation, stock-shares of wealth, and the expectation response are averaged over time at the individual level. The figure displays the conditional mean "bin-scatter" estimates on the resulting dataset of one averaged observation per household.

Figure 3: Stockholding and the subjective probability of positive returns

funds, investment trusts; checking, savings, and money market accounts; fixed-income assets; and IRA and Keogh accounts.[18]

An extensive literature has consistently shown that cross sectional differences in individuals' beliefs about stock returns predict their stock-holding behavior.[19] Giglio et al. (2021) show that the persistent component of expectations—the individual fixed effect—is the main driver of this predictive relationship; belief revisions have a much weaker link to portfolios. To illustrate this relationship in my HRS sample, Figure 4 displays "bin-scatter" estimators for the conditional mean of stock market participation and the share of wealth in stocks as a function of the $P^{\geq 0}$ expectation, aggregated at the individual level by taking their average over time. The left panel of the figure shows a steep relationship between stock market participation and the reported probability of positive returns: those with the most optimistic answers participate at roughly twice the rate of the least optimistic.

[18]This is the sum of RAND HRS aggregates `astck`, `achck`, `acd`, `abond`, and `aira`.

[19]Dominitz and Manski (2007), Hurd, Van Rooij, and Winter (2011), Amromin and Sharpe (2014), Drerup, Enke, and von Gaudecker (2017), Ameriks, Kézdi, et al. (2020), Giglio et al. (2021), and Calvo-Pardo, Oliver, and Arrondel (2022) are some examples. See Hurd (2009) and Manski (2018) for reviews.

The right panel shows that the relationship also holds for the share of financial wealth invested in stocks. These relationships hold for both levels of education and they have similar magnitudes. A college degree shifts participation rates and portfolio shares up at every level of optimism.

This section reviewed several challenges faced by traditional life cycle models in explaining U.S. households' stockholding patterns. Stockholding rates are low and correlated with education, those who own stocks do not allocate as much wealth to them as these models prescribe, and the optimal relationship between portfolio shares and age predicted by these models does not match its empirical counterpart. These conflicting predictions arise in models that calibrate agents' expectations about stocks using the historical data on actual returns, which is at odds with survey measurements of these expectations. Several properties of measured expectations could help bridge the gap between portfolio choice models and U.S. households' behavior. For instance, heterogeneous beliefs correlated with education could help explain limited participation rates and pessimism could rationalize low shares of wealth in stocks. In the next sections, I evaluate this possibility, testing whether a life cycle model that matches expectations from survey measurements can fit savings and portfolio choices better than its historically-calibrated counterpart.

3. A Model of Beliefs and Stock-Holding

The model that I propose consists of two parts. The first is a measurement system that I use to interpret the probabilistic responses, $P^{\geq 20}$, $P^{\geq 20}$, and $P^{\leq -20}$, and to estimate the distribution of beliefs about stock returns across the population. The second part is a life cycle model of saving and portfolio choices in which agents' beliefs about asset returns are fixed and exogenous. I discuss each part in turn.

3.1 Representing the Beliefs of U.S. Households

To represent the beliefs about stock returns of U.S. households, I construct and estimate a model that maps their probabilistic assessments to heterogeneous subjective distributions of stock returns. The model represents beliefs in a way that can be estimated directly from survey measurements and then plugged into life cycle models. I estimate the model using almost twenty years of longitudinal measurements of U.S. households' expectations about stock returns. The estimates suggest that there are permanent differences in beliefs about stock returns across people, that the average person is more pessimistic about stocks than the historical experience would suggest, and that more educated people are more optimistic about stocks. All these features are consistent with previous findings in the literature.

3.1.1 *A model of beliefs and probabilistic assessments*

The model is an adaptation of the ones used by Kézdi and Willis (2011) and Ameriks, Kézdi, et al. (2020). Every person believes that stock returns follow a distribution that can change from one person to the next but does not change over time. People use their subjective distributions to produce probabilistic assessments but their answers are also perturbed by time-varying shocks that represent survey errors and short term fluctuations in their beliefs. People also round their answers to different but personally-stable degrees: some round all their answers to the nearest multiple of 5%, others to the nearest multiple of 10%, 25%, 50%, or 100%. I identify and estimate the distribution of the persistent part of beliefs across the population using the longitudinal nature of the data and the fact that multiple questions about stock returns are asked in various survey waves.

In the model, people believe that nominal stock returns follow a log-normal distribution[20] with individual-specific parameters μ_i and σ_i:

$$\ln \tilde{R}_{t+1} \overset{i}{\sim} \mathcal{N}(\mu_i, \sigma_i).$$

The individual-specific parameters of people's beliefs are fixed over time and follow a distribution Ω across the population, $(\mu_i, \sigma_i) \sim \Omega$. This assumption is a parsimonious way of modeling the fact that most of the panel variation in probabilistic assessments about stock returns comes from persistent differences across individuals, as shown in Section 2.2 and in Giglio et al. (2021). The literature has suggested various mechanisms that could generate this persistent heterogeneity, offering differences in lived experiences (Malmendier and Nagel 2011) or in costs of and returns to learning about stocks (Kézdi and Willis 2011) as two examples. I take belief heterogeneity as given and model it using the distribution Ω, which I estimate.

People make probabilistic assessments about stock returns using their log-normal beliefs, but their responses are subject to time- and question-specific disturbances and rounded to different degrees. Manski and Molinari (2010) and Giustinelli, Manski, and Molinari (2022) demonstrate that rounding is prevalent in the answers to probabilistic questions in the HRS, and that the degree or "coarseness" of rounding varies across respondents but is stable over time. These studies show that ignoring the rounding patterns present in the data and taking probabilistic assessments at face value can alter the estimates of econometric models and their precision. I account for rounding assuming that each person i has a "rounding type" (or rounding behavior) $\mathcal{R}_i \in \{5, 10, 25, 50, 100\}$. An individual of rounding type $\mathcal{R}_i = x$ rounds all of his answers to probabilistic questions about stock returns to the nearest multiple of x% at every point in time. Rounding types $\mathcal{R}_i$ are independent of (μ_i, σ_i) and I use $\vec{\wp} = \{\wp_5, \wp_{10}, \wp_{25}, \wp_{50}, \wp_{100}\}$ to denote their

[20] The main qualitative results of Section 5 hold under different distributional assumptions. For example, if the return factor of stocks is normal (with censoring at zero) instead of log-normal, it is still the case that the life cycle model with beliefs estimated from surveys fits the data better than its historical counterpart for both high school and college graduates, with smaller estimates of risk aversion and entry costs, and larger estimates of time-discount factors.

frequencies across the population.

In every survey wave, a person might be asked to estimate the chances of positive returns, returns greater than 20%, or returns lower than −20%. In the model, person i's responses to these questions at time t are:

$$P^{\geq 0}_{i,t} = \left[\Phi\left(\frac{\mu_i}{\sigma_i} + \varepsilon^{\geq 0}_{i,t}\right)\right]_{\mathcal{R}_i},$$
$$P^{\geq 20}_{i,t} = \left[\Phi\left(\frac{\mu_i - \ln 1.20}{\sigma_i} + \varepsilon^{\geq 20}_{i,t}\right)\right]_{\mathcal{R}_i}, \qquad P^{\leq -20}_{i,t} = \left[\Phi\left(\frac{\ln 0.8 - \mu_i}{\sigma_i} + \varepsilon^{\leq -20}_{i,t}\right)\right]_{\mathcal{R}_i}, \tag{1}$$

where the operator $[\,\cdot\,]_x$ rounds its argument to the nearest multiple of x%, $\Phi(\cdot)$ is the univariate standard normal CDF, and the random disturbances $\left\{\varepsilon^{\geq 0}_{i,t}, \varepsilon^{\geq 20}_{i,t}, \varepsilon^{\leq -20}_{i,t}\right\}'$ follow the joint normal distribution $\mathcal{N}(\vec{0}, \Sigma)$. I assume that Σ is diagonal, so that the shocks are independent. Without rounding or disturbances, Equation 1 would imply that people perfectly calculate and report the queried moments of their subjective log-normal distribution every wave—their answers would not change over time. The time-specific random disturbances represent survey errors and the effects of short term information that might shift people's responses but not their long term beliefs about stock returns.

3.1.2 *The estimated distribution of beliefs*

I represent the distribution of beliefs across the population Ω using grids of (μ, σ) pairs. The grids are discretizations of bivariate normal distributions that condition on the event that subjective standard deviations must be positive:

$$\begin{bmatrix}\mu_i \\ \sigma_i\end{bmatrix} \overset{\texttt{discretized}}{\sim} \mathcal{N}\left(\begin{bmatrix}\nu_\mu \\ \nu_\sigma\end{bmatrix}, \begin{bmatrix}\Psi_{1,1} & \Psi_{1,2} \\ \Psi_{2,1} & \Psi_{2,2}\end{bmatrix}\right) \mid \sigma_i > 0. \tag{2}$$

For any set of parameters $\{\nu_\mu, \nu_\sigma, \Psi\}$, I produce a set of 25 equiprobable (μ, σ) pairs that approximate the conditioned normal distribution in Equation 2; Appendix A discusses the discretization procedure in detail. There are two main advantages of this representation. First, it is flexible enough to accommodate distributions where beliefs have different averages, levels of dispersion, and correlations between subjective means and standard deviations. Second, I can incorporate the resulting discrete set of estimated (μ, σ) into a life cycle model as a set of possible "belief types."

I estimate the model by maximum likelihood, using responses to the three probabilistic assessments in the nine biennial waves of the HRS between 2002 and 2018. The full set of parameters to estimate in my representation of beliefs comprises those in the distribution of μ and σ in Equation 2, the covariance matrix of random disturbances Σ, and the prevalence of different rounding types $\vec{\wp}$: $\vartheta^B \equiv \{\nu_\mu, \nu_\sigma, \Psi, \Sigma, \vec{\wp}\}$. Appendix A describes the likelihood function. My estimation sample consists of all person-year observations in which: a) the person being interviewed is the financial respondent of the household, b) the person being interviewed is between 50 and 65 years old, and c) the respondent did not refuse to answer nor answer "do not know" to any of the three questions regarding future stock returns.[21] These restrictions, in addition to considering only respondents with at least a high school degree, yield a sample of 35, 211 individual-wave observations from 12, 025 unique individuals.

Because of the relationship between education and probabilistic assessments documented in Section 2.2, I estimate the model separately for those with and without a college degree. This lets people with different levels of educational attainment have different average beliefs, levels of disagreement, rounding patterns, and covariance structures in the random disturbances of their responses. Previous studies have found evidence of these differences (see, e.g., Kézdi and Willis 2011; Das, Kuhnen, and Nagel 2020; Ameriks, Kézdi, et al. 2020). A relationship between beliefs about stock returns and education could explain part of the educational gradients in stockholding documented in Section 2.

The parameter estimates, which are presented in Table 11 of Appendix A, suggest that

[21] Out of the person-year observations that satisfy a) and b), the fraction of observations that I drop for not satisfying c) are 11.8% for those with a high-school degree and 5.2% for college graduates.

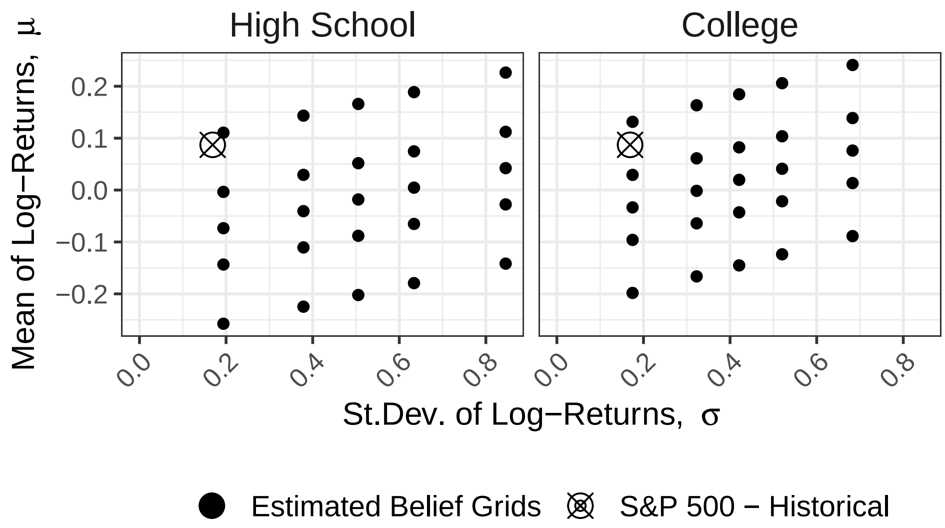

The belief grids are equiprobable discretizations of the joint distribution in Equation 2, see Appendix A for details. The S&P 500 point depicts the average and standard deviation of annual log-returns to that index between 1881 and 2018. The data for this calculation comes from the accompanying data file to Chapter 26 of Shiller (1990).

Figure 4: Estimated belief grids

different degrees of rounding are prevalent in the data and that, as found by Giustinelli, Manski, and Molinari (2022), more educated individuals tend to round their answers to finer levels. The fraction of individuals belonging to the finest rounding type (multiples of 5%) is 40% for high school graduates and 51% for college graduates. Coarse rounding is non-negligible: 17% of high school graduates and 9% of college graduates round their answers to levels coarser than 10% (25%, 50% or 100%).

The estimated models assign a large role to persistent differences in people's beliefs for explaining their probabilistic assessments.[22] Figure 4 depicts the estimated grids of possible (μ, σ) pairs for every level of education, showing that they span large ranges of the mean-variance space. The breadth of the grids is a consequence of the fact that, as discussed in Section 2.2, most of the variation in this data comes from persistent cross-sectional differences, which in this model correspond to (μ_i, σ_i).

[22]While the functional forms that I impose assume that (μ_i, σ_i) are fixed individually, they do not limit the scale of their distribution across the population. In principle, the model could adjudicate the differences in responses to the time-specific shocks $\left\{\varepsilon_{i,t}^{\geq 0}, \varepsilon_{i,t}^{\geq 20}, \varepsilon_{i,t}^{\leq -20}\right\}'$ and find narrow distributions for (μ_i, σ_i).

Table 5: Summary statistics of estimated beliefs about stock returns

	$\mu_i = E_i[\ln \tilde{R}]$		$\sigma_i = \sqrt{V_i(\ln \tilde{R})}$		$E_i[\tilde{R}] - R$	Sharpe Ratio*	
	Mean	S.D.	Mean	S.D.	Fract > 0	Mean	S.D.
Estimated Beliefs							
High School	-0.017	0.131	0.512	0.221	0.60	0.283	0.111
College	0.020	0.118	0.424	0.172	0.72	0.266	0.169
Historical Realizations							
S&P 500 (1881-2018)	0.087	-	0.169	-	-	0.342	-

The Sharpe ratio and its summary statistics are computed only for those beliefs for which it is positive. All summary statistics are taken over the points in the estimated beliefs grids for every level of education, depicted in Figure 4. The Sharpe ratios are computed with the nominal "risk-free" return factor as a benchmark. I take the average yearly risk-free return factor between 1881 and 2018 from the accompanying data file to Chapter 26 of Shiller (1990), which is 1.044.

For every level of educational attainment, the estimated distributions of persistent beliefs imply that most people are more pessimistic about future stock log-returns than standard calibrations based on the historical experience. Figure 4 compares the estimated distribution of beliefs with the historical mean and standard deviation of log-returns to the S&P500 index. For both levels of educational attainment, the majority of points lie below and to the right of the S&P500. This means that most people believe log-returns to stock investments are lower on average and more volatile than those historically experienced by the S&P500. Moreover, many of the points fall below the $\mu = 0$ line, suggesting that a considerable fraction of the population believes that average log-returns are in fact negative. As suggested by Dominitz and Manski (2007), this "pessimism" about stock returns could help explain why a fraction of U.S. households do not own stocks at all, despite having substantial wealth.[23]

The estimated distributions of beliefs imply a steep relationship between educational attainment and anticipated rewards from investing in stocks. Table 5 presents summary

[23] In other countries, studies have also shown that stock-market participation is insensitive to wealth windfalls (Andersen and Nielsen 2011; Briggs et al. 2021). These facts are difficult to accommodate for models that try to explain non-participation using monetary costs. For instance, Catherine (2021) assumes a fraction of the population—which he estimates to be 46%-47%—exogenously avoids the stock market. Pessimistic beliefs could be behind this persistent non-participation.

statistics of the distribution of persistent beliefs for every level of educational attainment. Expected log-returns (μ) vary widely within educational attainment groups, with within-group standard deviations of 1,300 and 1,200 basis points for high-school and college graduates, respectively. Average expected log-returns are higher for college graduates (200 basis points) than for high school graduates (-170 basis points). Subjective assessments of volatility (σ) vary considerably within educational attainment groups but also show a relationship with education: both their mean and standard deviation across individuals are lower for college graduates. The average subjective standard deviation of log-returns is 4,200 basis points for college graduates and 5,100 basis points for high school graduates. Both values are much greater than the historical standard deviation of the S&P500's log-returns, which has been around 1,700 basis points.

The estimated distributions of beliefs also imply differences in people's expected rewards from the risks associated with stock market participation. For high school and college graduates, these differences align qualitatively with their differing investment patterns. The fourth column of Table 5 shows that not all individuals believe that the expected returns to stocks are greater than those of a safe bond: only 60% of high school graduates and 72% of college graduates do. For those who expect a premium from stocks, the fifth and sixth column calculate the Sharpe ratio, which measures the expected excess returns per unit of risk that they believe stocks offer. The average Sharpe ratios of high school and college graduates who believe there is an equity premium are similar (0.28 and 0.27 respectively) but lower than the historical Sharpe ratio of the S&P500 index which, based on 1881 to 2018 data from Shiller (1990), has been around 0.34.

The model can produce individual-level estimates of the persistent component of beliefs that are interpretable, account for measurement error and rounding, and aggregate the information of the multiple probabilistic questions. Similarly to Ameriks, Kézdi, et al. (2020), the estimates of beliefs that I construct are the individual level expectation of

Table 6: Estimated beliefs and portfolio choices

	Stock Market Participation		Stock Share of Fin. Wealth	
	Probit Marg.	Probit Marg.	Tobit	Tobit
$\widehat{\mu}_i$	1.66	0.85	2.08	1.14
	[1.31; 2.01]	[0.63; 1.07]	[1.68; 2.49]	[0.88; 1.38]
$\widehat{\sigma}_i$	−1.08	−0.52	−1.24	−0.59
	[−1.32; −0.88]	[−0.66; −0.38]	[−1.47; −1.02]	[−0.76; −0.46]
Outcome Mean	0.45	0.45	0.27	0.27
Std. Dev. ($\widehat{\mu}_i$)	0.08	0.08	0.08	0.08
Std. Dev. ($\widehat{\sigma}_i$)	0.11	0.11	0.11	0.11
Pseudo R^2	0.13	0.27	0.11	0.24
Controls		✓		✓
Num. obs.	5447	5447	5447	5447

The sample are HRS financial respondents between the ages of 50 and 65. I aggregate the data at the individual level by taking averages over the time dimension. For multinomial variables like the willingness to take risk, I take the mode. The control variables are: college degree, age, sex, log-income, general willingness to take risk (10 possible levels), and financial planning horizon (5 possible levels).The first two columns report average marginal effects from Probit regressions. The second two columns report coefficients from Tobit regressions that limit the outcome to the $[0, 1]$ interval. Reported confidence intervals are 95% coverage sets from 150 bootstrap replications. To account for the fact that $\widehat{\mu}_i$ and $\widehat{\sigma}_i$ are generated regressors, each bootstrap replication re-samples households, re-estimates the full beliefs model, re-calculates $\widehat{\mu}_i$ and $\widehat{\sigma}_i$ for each sampled household, and re-estimates the regressions reported in this table.

(μ_i, σ_i) given the full set of responses of the person:

$$\widehat{\mu}_i = \mathbb{E}\left[\mu_i \,\middle|\, \left\{P^{\geq 0}_{i,t}, P^{\geq 20}_{i,t}, P^{\leq -20}_{i,t}\right\}_{t\in\mathcal{T}(i)}\right] \quad \text{and} \quad \widehat{\sigma}_i = \mathbb{E}\left[\sigma_i \,\middle|\, \left\{P^{\geq 0}_{i,t}, P^{\geq 20}_{i,t}, P^{\leq -20}_{i,t}\right\}_{t\in\mathcal{T}(i)}\right], \quad (3)$$

where $\mathcal{T}(i)$ are the waves in which i appears satisfying the sample criteria.

The individual-level estimates of beliefs about stock returns are predictive of portfolio choices. Table 6 reports regressions that relate these estimates to the stock ownership and portfolio shares of financial respondents in the HRS sample. Individuals with higher subjective means of log-returns participate at higher rates and have higher shares of their wealth invested in stocks; the converse is true for those with higher expected volatilities. These estimated relationships are precise and robust to the inclusion of socio-demographic controls and measures of patience and risk preferences. The estimates imply that an individual with a subjective mean $\widehat{\mu}_i$ that is higher by 100 basis points is 0.85 percentage points more likely to own stocks, and has an (unconditional) share of wealth in stocks that is 1.14 percentage points higher. An individual with a subjective volatility $\widehat{\sigma}_i$ that is higher by 100 basis points is 0.52 percentage points less likely to own stocks, and has a share of wealth in stocks that is 0.59 percentage points lower. These associations are empirically weaker than the ones implied by frictionless models—Ameriks, Kézdi, et al. (2020) call this fact "the attenuation puzzle"— however, their magnitudes are material given the wide variation in $\widehat{\mu}_i$ and $\widehat{\sigma}_i$, which have standard deviations of 800 and 1,100 basis points respectively.[24]

Qualitatively, the fact that not all households believe that there is an equity premium could explain why some of them do not invest in stocks. The low perceived Sharpe ratios could explain why those who do own stocks invest a limited share of their wealth in them. In addition, the educational differences in beliefs about returns are also consistent with the different stockholding patterns of high school and college graduates. Differences in the share of respondents who believe there is an equity premium could explain the relationship between education and participation rates. The similarity in the subjective Sharpe ratios of those who believe there is a premium could explain why the conditional shares of wealth in stocks are similar for high school and college graduates. To evaluate these possibilities quantitatively, I now present a life cycle model of saving and portfolio choices.

3.2 Life Cycle Model of Saving and Portfolio Choices

The life cycle model has several features in common with portfolio choice models in the literature (see, e.g., Cocco, Gomes, and Maenhout 2005; Gomes and Michaelides 2005; Campanale, Fugazza, and Gomes 2015; Fagereng, Gottlieb, and Guiso 2017; Catherine 2021). Households save to smooth their consumption against fluctuations in their income,

[24] Enke et al. (2024) show that portfolio choice is only one of many instances in which people's decisions are empirically less responsive to economic fundamentals than we would expect based on theoretical models. They argue that this "behavioral attenuation" is due to information processing constraints.

which come from deterministic changes as they age and random shocks, both permanent and transitory. My model features a bequest motive and age-varying health expenditure shocks as additional reasons for saving; both are important motives in explaining post-retirement wealth (De Nardi, French, and Jones 2010; Ameriks, Briggs, et al. 2020). Agents face two different financial frictions when deciding how to allocate their savings between assets. First, they must pay a monetary cost before owning stocks for the first time. Second, they face a 10% early withdrawal penalty when liquidating stocks before retirement.

I now discuss the main components of the model and leave its full mathematical description and treatment for Appendix C.

3.2.1 *Lifespan, utility, and mortality*

Time periods in the model represent a year. Agents enter the model at age 24 and can live up to a maximum age of 100. At the end of every year, they face an exogenous risk of death that becomes certain at the maximum age. The probability of surviving from age t to $t+1$ is represented by δ_{t+1} and the probability of not surviving is $\not\delta_{t+1} \equiv 1 - \delta_{t+1}$.

Agents derive utility from consumption. Their utility function follows a constant relative risk aversion specification,

$$u(C) = \frac{C^{1-\rho}}{1-\rho}, \tag{4}$$

where ρ is the coefficient of relative risk-aversion.

If an agent dies at the end of a year, after making all his choices, he derives warm-glow utility from bequeathing his total wealth. The utility derived from bequeathing wealth x is:

$$\mathbb{B}(x) = \mathbb{b} \times \frac{(x/\mathbb{b})^{1-\rho}}{1-\rho} = \mathbb{b}^{\rho} \times \frac{x^{1-\rho}}{1-\rho} = \mathbb{b}^{\rho} \times u(x),$$

where $\mathbb{b} \geq 0$ is a parameter that controls the intensity of the bequest motive. This is the same specification used by, e.g., Gomes and Michaelides (2005).

3.2.2 *Income process*

Agents supply labor inelastically and retire exogenously at the end of the year in which they turn 65. Their labor earnings, denoted by $Y_{i,t}$, are a product of two factors: a permanent component represented by $P_{i,t}$ and a transitory stochastic component represented by $\theta_{i,t}$. Labor earnings and their permanent component follow:

$$\ln Y_{i,t} = \ln P_{i,t} + \ln \theta_{i,t} \quad \text{and} \quad \ln P_{i,t} = \ln P_{i,t-1} + \ln \Gamma_{i,t} + \ln \psi_{i,t},$$

where Γ_t is a deterministic growth factor that captures life cycle patterns in earnings, and $\ln \psi_{i,t} \sim \mathcal{N}(-\sigma_\psi^2/2, \sigma_\psi^2)$ is a multiplicative shock to permanent income.[25]

The transitory component of earnings $\theta_{i,t}$ is a mixture that represents unemployment and temporal fluctuations in income that occur while employed:

$$\ln \theta_{i,t} = \begin{cases} \ln \mathcal{U}, & \text{With probability } \mho \\ \ln \tilde{\theta}_{i,t} \sim \mathcal{N}(-\sigma_\theta^2/2, \sigma_\theta^2), & \text{With probability } 1 - \mho. \end{cases}$$

$\mho$ denotes the probability of unemployment, and $\mathcal{U}$ denotes the replacement factor of unemployment benefits.

The sequences of growth factors $\{\Gamma_t\}_{t=25}^{100}$ differ for high school and college graduates. I take their values from Cagetti (2003). Figure 5 displays the income paths that individuals with different education levels would experience in the absence of shocks.[26] The decline after age 65 corresponds to retirement. I also take the volatilities of transitory and permanent income shocks (σ_ψ and σ_θ) used by Cagetti (2003), which come from Carroll and Samwick's (1997) estimates.

After retirement, individuals are no longer subject to transitory and permanent shocks to their earnings. Instead, they face out-of-pocket medical expenditure shocks. As people

[25] The mean of the shock is set so that $E[\psi_{i,t}] = 1$.

[26] This is $P_{i,24} \times \prod_{j=25}^{t} \Gamma_{i,t}$.

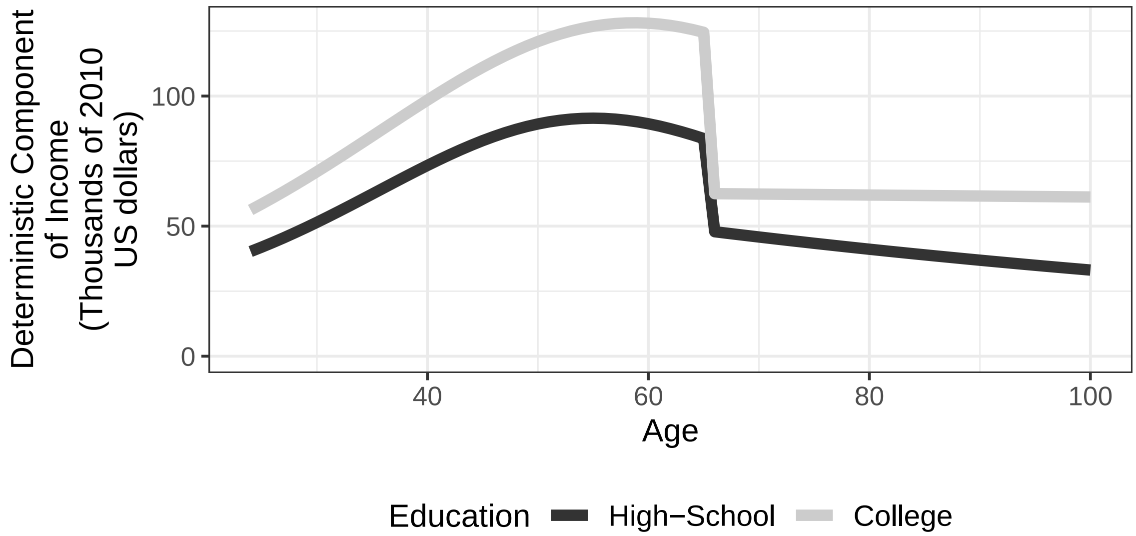

These estimated trajectories for the deterministic component of permanent income come from Cagetti (2003). I thank the author for sharing his exact estimates.

Figure 5: Deterministic component of income by education level

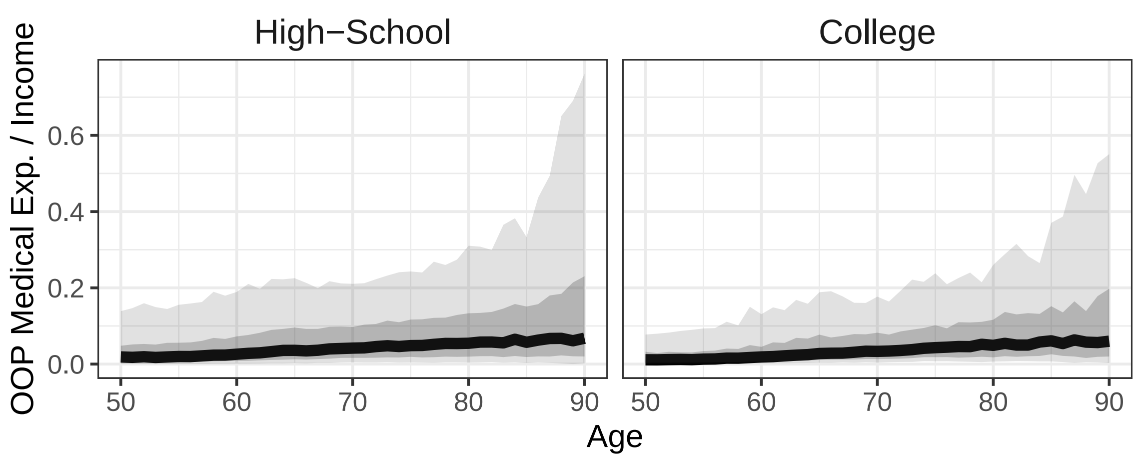

The figure depicts estimates of the distribution of people's out-of-pocket medical expenditures expressed as a ratio of their annual income. At every age, the solid line represents the median, the dark shaded area represents the 25th and 75th percentiles, and the light-shaded area represents the 10th and 90th percentiles.

Figure 6: Out-of-pocket medical expenditures over the life cycle

age, medical expenditures increase rapidly (see Figure 6). The anticipation of these rising expenditures has been shown to be one of the reasons why the elderly do not spend their wealth as quickly as a basic life cycle model would predict (De Nardi, French, and Jones 2010; Ameriks, Briggs, et al. 2020). The literature that specializes in the study of these expenditures has identified several important features, such as their dependence on persistent health states (Kopecky and Koreshkova 2014; Ameriks, Briggs, et al. 2020), their relationship with permanent income (De Nardi, French, and Jones 2010), and the prevalence of "catastrophic" shocks (French and Jones 2004). To incorporate these shocks into my model, I adopt a parsimonious representation that matches the distribution of expenditures across the population and matches the fact that they increase with age and permanent income. Every year, agents draw a shock $\texttt{oop}_{i,t}$ that represents the fraction of their earnings used up by out-of-pocket medical expenses. I assume that government programs cover any health expense above an agent's income, so that income net of medical expenses cannot be negative. The process for earnings net of costs and permanent income becomes:

$$
\begin{aligned}
Y_{i,t} &= P_{i,t} \times \max\{0.0, 1 - \texttt{oop}_{i,t}\} \\
P_{i,t} &= P_{i,t-1}\Gamma_i.
\end{aligned}
$$

The shocks $\texttt{oop}_{i,t}$ are independent across time and follow age- and education-specific distributions that approximate the patterns in Figure 6. I calibrate these distributions using the RAND HRS longitudinal file; I describe the process in Appendix B.

3.2.3 *Financial assets and frictions*

Agents try to smooth their consumption by saving and they have two assets available for this purpose. The first asset is a risk-free liquid account with a constant per-period return factor R. The second asset is a stock fund with a stochastic return factor $\tilde{R}$ that agents view as log-normally distributed and independent across time. I denote the dollar amounts available to agent i at the start of period t in the risk-free account and the stock fund with $M_{i,t}$ and $N_{i,t}$, respectively. The flows between the two assets are one of the agents' control variables and denoted with $D_{i,t}$, with $D_{i,t} > 0$, representing a movement of funds from the risk-free to the risky account.

The model has four different financial frictions. First, agents cannot short-sell any of the assets or borrow against their future income. Second, agents enter the model not having access to the stock fund and must pay a one-time financial cost to access it. As in Gomes and Michaelides (2005), the cost represents the money and time spent opening a brokerage account and getting familiarized with the stock market. The cost is proportional to agents' permanent income, $P_{i,t} \times F$, where F is a parameter to be estimated. Third, withdrawals from the stock fund are taxed at a constant rate $\tau = 0.1$ before agents retire.[27] This friction represents early retirement fund withdrawal penalties and the costs associated with liquidating stock positions. Finally, agents must pay for their consumption using funds from their risk-free accounts only.

As demonstrated by Campanale, Fugazza, and Gomes (2015), the combination of rebalancing penalties and the fact that consumption must be paid for using risk-free funds generates a reason for young people to not allocate all of their wealth to stocks, as is predicted by standard life cycle portfolio models (see Cocco, Gomes, and Maenhout 2005). An agent who anticipates the possibility of consuming part of his savings in the next period—because of an unemployment spell, for instance—might keep a buffer of risk-free funds to avoid having to pay the stock withdrawal tax if this is the case. The size of the desired buffer will depend, among other factors, on the agent's beliefs about the equity premium and its volatility, the volatility of his income, and the magnitude of the withdrawal tax.

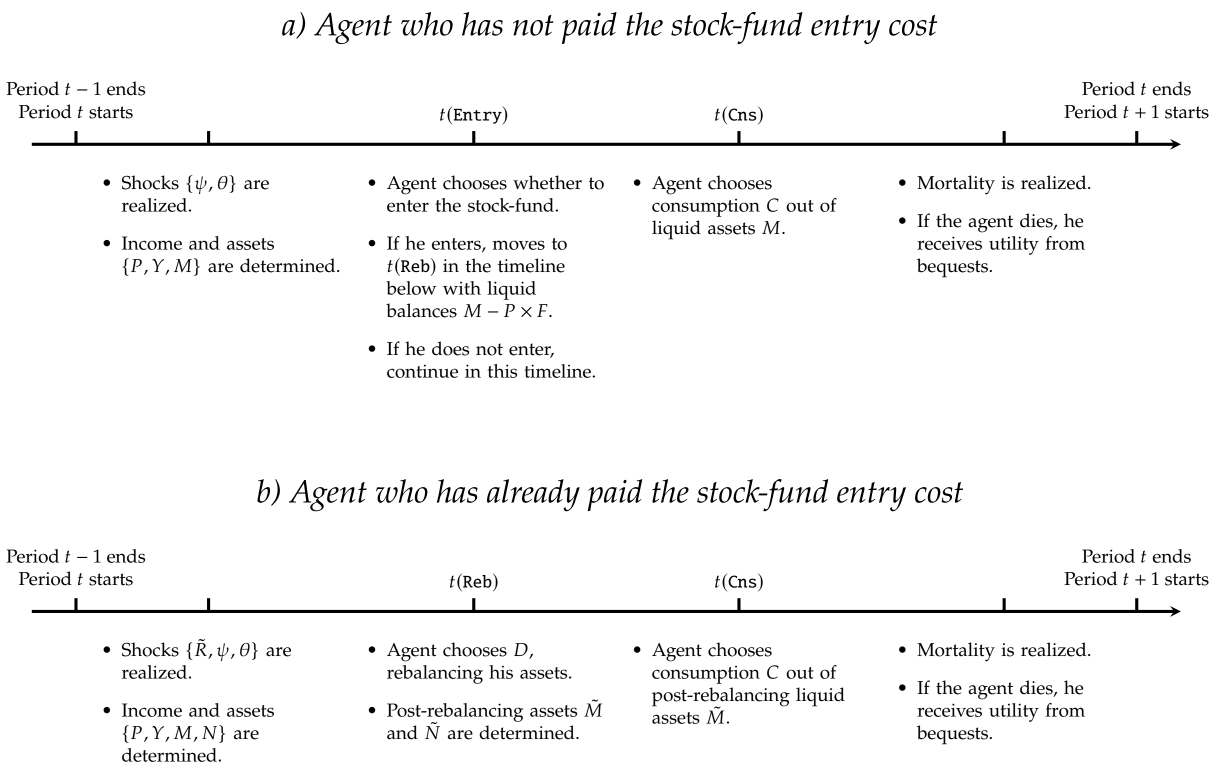

Figure 7: Summary of timing in the life cycle model

3.2.4 *Timing and recursive representation*

Figure 7 summarizes the timing of stochastic shocks and optimizing decisions that occur within a period of the life cycle model. Agents enter the model in timeline a), not having paid the stock fund entry cost $P_{i,t} \times F$. They are presented with the option to pay the cost and enter the fund every year. Once they pay the cost, they move to the portfolio-rebalancing stage $t(\texttt{Reb})$ of timeline b) and remain on timeline b) for the rest of their lives.

To illustrate the choices and constraints faced by agents succinctly, Equation 5 presents the recursive value function of an agent who has paid the financial participation cost and therefore has access to the stock-fund.[28] I present the value function of the agent who has not paid the financial participation cost in Appendix C, which also discusses various alternative representations of the model that I use in its solution.

[27]The withdrawal tax rate becomes $\tau = 0$ after agents retire.

[28]Individual subindices i are dropped for simplicity.

$$
\begin{aligned}
V_t^{\text{In}}(M_t, N_t, P_t) = \max_{C_t, D_t} u(C_t) + \beta \delta_{t+1} \mathbb{E}_t \left[V_{t+1}^{\text{In}}(M_{t+1}, N_{t+1}, P_{t+1}) \right] \\
+ \not{\delta}_{t+1} \mathbb{B}(A_t + \tilde{N}_t)
\end{aligned}
$$

Subject to:

$$
-N_t \leq D_t \leq M_t, \quad 0 \leq C_t \leq \tilde{M}_t
$$

$$
\begin{aligned}
\tilde{M}_t &= M_t - D_t \left(1 - 1_{[D_t \leq 0]} \tau\right) \\
\tilde{N}_t &= N_t + D_t \\
A_t &= \tilde{M}_t - C_t \\
M_{t+1} &= R A_t + Y_{t+1} \\
N_{t+1} &= \tilde{R}_{t+1} \tilde{N}_t \\
P_{t+1} &= \Gamma_{t+1} \psi_{t+1} P_t \\
Y_{t+1} &= \theta_{t+1} P_{t+1}
\end{aligned}
\quad . \tag{5}
$$

4. Estimation

To determine whether survey measurements of beliefs can improve the life cycle model's capacity to fit U.S. households' savings and portfolios, I now estimate the model under alternative specifications of beliefs about stock returns. In the first specification, all the simulated agents believe that future stock fund returns will follow a distribution that approximates their historical behavior—this is the commonly used *full-information rational-expectations* (F.I.R.E.) specification. In the second specification, the agents' beliefs about future stock returns are heterogeneous and distributed across the population following the specifications that I estimated from survey measurements in Section 3.1.2. For each level of education and each specification of beliefs, I estimate the model's unobservable parameters that govern agents' preferences and barriers to stock market participation. The estimation strategy searches for the parameters that best replicate the life cycle profiles of U.S. households' savings, stock market participation rates, and shares of financial wealth in stocks.

4.1 Data, Sample Restrictions, and Targeted Variables

The estimation exercise targets the life cycle profiles of wealth-to-income ratios, stock market participation rates, and the shares of wealth in stocks conditional on participation of U.S. households. I construct these targeted variables using the nine waves of the *Survey of Consumer Finances* (SCF) between 1995 and 2019. My variables of interest rely on definitions and calculations in the SCF's *summary files*, which produce standard measures of households' wealth and income from the raw survey.[29]

The *wealth-to-income ratio* measures the financial savings of an individual relative to what would be their "usual" income—their permanent income in the life cycle model. For my measure of wealth, I take the economic unit's total financial assets (`fin`). For my measure of income, I take the sum of wage and salary income (`wageinc`) and social security and pension income (`ssretinc`).

To measure stock holdings, I use the SCF's estimated total value of financial assets invested in stocks (`equity`), which includes direct and indirect investments. I define *stock market participation* as a binary variable that takes the value of one if the given economic unit has stock holdings greater than zero, and zero otherwise. Finally, I calculate the *share of wealth in stocks* as a unit's stock holdings divided by their wealth; if a unit's wealth is zero, I set the share to zero.

To generate a sample that more closely matches the type of household represented in the model and to ensure that the variables of interest are well defined, I apply various filters to the data before using it to compute the targeted moments. First, I keep only

[29] Variables in `teletype font` denote calculations that are readily available in the summary files.

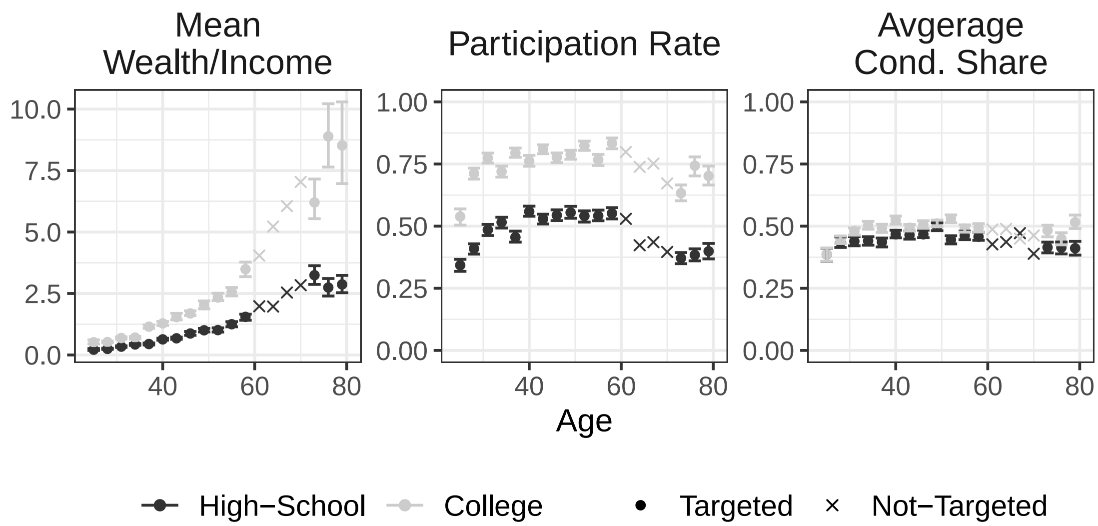

Each point represents the relevant statistic computed over the group of SCF respondents with the given level of educational attainment and falling in a 3-year age bin. See the main text for definitions of the sample and relevant variables. Error bands correspond to 95% point-wise confidence intervals calculated using 500 bootstrapped samples for each level of education.

Figure 8: Targeted moments

economic units whose respondent was born between 1920 and 1995. To make the wealth-to-income ratio comparable to a model analogue that uses permanent income, I only keep units who report that their income was "normal" in the given year.[30] To ensure that the ratio is defined, I only keep units with positive incomes. Finally, I exclude business owners from the sample.[31]

I group observations by the respondent's level of education and age. For education, I split the sample into those without a high school degree (which I do not analyze), those with a high school but not a college degree, and those with a college degree. For age, I form three-year bins starting at age 24 and up to a maximum age of 80, for a total of 19 groups: $\{24, 25, 26\}$, $\{27, 28, 90\}$, ..., $\{75, 76, 77\}$, $\{78, 79, 80\}$. The moments that my estimation routine targets are summary statistics of the variables of interest calculated over the education-by-age groups.

[30] The question of whether income was unusual was added in the 1995 wave. For this reason, I exclude previous waves from this part of the analysis.

[31] I define business owners as those with bus > 0.

For each group of observations, the moments that I target in estimation are:

- **Average wealth-to-income ratio**: the average of the wealth-to-income ratio. The ratio can take on extreme values for agents with a low measured wage income. To limit the influence of these extreme observations, I winsorize the wealth-to-income ratios at their within-group 95th percentile before taking their average.

- **Stock market participation rate**: the average of the binary stock market participation variable.

- **Average conditional share of wealth in stocks**: the average share of wealth in stocks of those who participate in the stock market.[32]

I use survey weights for the calculation of these moments and rescale the weights of different waves so that each of them has an equal representation in the moments' calculations.

Figure 8 displays each of the moments for every education-by-age group. The figure also presents 95% confidence intervals for each targeted moment. The confidence intervals come from calculating the targeted moments on 500 bootstrapped samples for each level of education. The error bars correspond to the 2.5th and 97.5th percentiles of the bootstrapped values of each moment. The bars demonstrate that the only moments with considerable sampling variation are the post-retirement wealth ratios of college graduates.

4.2 Objective Function and Optimization

For every level of education and specification of beliefs, I estimate the preferences and participation costs that minimize the distance between the targeted moments in the SCF and their model-implied counterparts.

[32]The model counterpart of the average conditional share can behave poorly in regions of the parameter space where participation is very low (for instance, if entry costs are very high). To improve this behavior and ensure that the SMM loss function is defined for all parameter values, I target a weighted average of the unconditional share of wealth in stocks that gives very little weight (instead of a 0 weight) to non-participants. Using $\text{Share}_{i,t}$ to denote the unconditional share, I target $\sum_i \omega_{i,t}\text{Share}_{i,t}/\sum_i \omega_{i,t}$, where $\omega_{i,t} = 1$ for participants ($\text{Share}_{i,t} > 0$) and $\omega_{i,t} = 10^{-3}$ for non-participants ($\text{Share}_{i,t} = 0$). This moment is practically

Table 7: Non-estimated parameter values

Symbol	Interpretation	Value	Source
$\mathbb{b}$	Bequest Intensity	2.5	Gomes and Michaelides (2005)
$\{\delta_t\}_{t=0}^{T}$	Survival probabilities	-	S.S.A. Actuarial tables 2010[33]
$\{\Gamma_t\}_{t=0}^{T}$	Permanent income drift (educ.)	-	Cagetti (2003)
σ_ψ, σ_θ	Volatility of income shocks (educ.)	-	Carroll and Samwick (1997)
$\mho$	Probability of unemployment	0.050	-
$\mathcal{U}$	Unemp. benefits replacement factor	0.500	National median, Ganong, Noel, and Vavra (2020)
π	Log-inflation rate	0.024	Mean CPI Log-Inflation.
r	Log risk-free rate (nominal)	0.043	Mean 1-year U.S. bond log-returns.
R	Risk-free return factor (real)	1.019	$\exp\{r-\pi\}$
μ^{SP500}	Mean stock log-return	0.085	S&P500 Index (nominal).
σ^{SP500}	St. Dev. stock log-returns	0.170	S&P500 Index (nominal).

Parameters that depend on educational attainment are marked with "(educ.)." Averages and standard deviations of financial variables are all taken over the 1881-2018 period. The data on financial assets and the CPI index come from the 'Chapter 26' file in Robert Shiller's website: http://www.econ.yale.edu/ shiller/data.htm.

The set of parameters that I estimate structurally consists of the coefficient of relative risk aversion (ρ), the time-discount factor (β), and the size of the cost of accessing the stock fund for the first time (F). I denote this set of parameters with $\vartheta \equiv \{\rho, \beta, F\}$. I set other parameters related to the income process and mortality to historical estimates or values from the literature and they remain fixed throughout the estimation process; I summarize their values or sources in Table 7. The parameters that govern the returns to different assets and agents' expectations about them are discussed in detail in the next section.

For any given set of parameters, I solve the life cycle model and simulate populations of agents that I use to find model-implied counterparts to the targeted moments. I solve the model by backward induction using a combination of the techniques outlined in Carroll (2006), Iskhakov et al. (2017), and Druedahl (2021); I describe the process in detail in Appendix E. I use the resulting policy functions to simulate populations of agents on which I calculate model-counterparts to the targeted empirical moments. The model is not well suited to accommodate the transitional dynamics of households' savings and portfolios as they move into retirement; it assumes that all agents retire exogenously at

identical to the average conditional share at the participation levels observed in the data and targeted by the model.

age 65. Therefore, I exclude the bins spanning ages 60 to 71, leaving a total of 15 targeted age-bins and 45 moments for each level of education.[34] For a level of education e, I use m_0^e to denote a vector of the 45 targeted empirical moments and $\hat{m}^e(\vartheta)$ to denote its model-implied counterpart under parameters ϑ.

The loss function that I minimize is:

$$L^e(\vartheta) = \left(m_0^e - \hat{m}^e(\vartheta)\right)' \mathbb{W}^e \left(m_0^e - \hat{m}^e(\vartheta)\right), \tag{6}$$

where $\mathbb{W}^e$ is a diagonal weighting matrix. Since the moments have different scales, I set $\mathbb{W}^e$ so that moment deviations are expressed as fractions of the average relevant statistic across age groups.[35] I obtain estimates as:

$$\hat{\vartheta}^e = \arg\min_{\vartheta} L^e(\vartheta). \tag{7}$$

To solve the minimization problem, I use the `TikTak` algorithm (Arnoud, Guvenen, and Kleineberg 2019) as implemented in the `estimagic` toolbox (Gabler 2022). I use 2,500 initial "exploration points" and allow for 10 full local-optimization runs using the `DFO-LS` algorithm (Cartis et al. 2019), which takes advantage of the least-squares structure of the optimization problem.

4.3 Agents' Expectations and Returns to Financial Assets

In the model, returns of the stock fund follow a distribution that approximates the historical behavior of the S&P500 index. In the first specification of agents' beliefs, everyone's expectations are consistent with this data generating process. In the second specification, agents' expectations match survey measurements instead.

[34] 15 age-bins times three moments of interest (median wealth-to-income ratio, participation rate, and average conditional share of wealth in stocks).

[35] For example, for individuals with a college degree, the diagonal positions of $\mathbb{W}^{\text{Coll.}}$ that multiply errors in the stock market participation rate are set to $1/(\bar{\text{Part}})^2$ where $\bar{\text{Part}}$ is the average of the 15 stock market participation rates in $m_0^{\text{Coll.}}$.

The simulated log-returns of the stock fund follow a normal distribution. Their mean and variance match those of the S&P500's log-returns between 1881 and 2018, and I make a constant adjustment for inflation π, which I take to be its average over the same period:

$$\ln \tilde{R} \sim \mathcal{N}(\mu^{SP500} - \pi, \sigma^{SP500}). \tag{8}$$

Table 7 presents the values of μ^{SP500}, σ^{SP500}, and π.

The first specification of beliefs that I use is *full-information rational-expectations* (F.I.R.E.). Under this specification, agents have correct beliefs about the stock-fund's returns (Equation 8) when solving their dynamic optimization problem. For each level of education, I simulate populations of 500 agents that run for 1,000 years and use them to compute the model's counterparts to targeted moments, $\hat{m}^e(\vartheta)$. The second specification of beliefs that I use is *estimated beliefs*. Under this specification, I replace every agent of the F.I.R.E. simulation with 25 agents whose beliefs about the stock fund's returns come from the distributions I estimated in Section 3.1.2.[36] The jth agent believes that $\ln \tilde{R} \sim \mathcal{N}(\hat{\mu}_j - \pi, \hat{\sigma}_j)$, where $\{\hat{\mu}_j, \hat{\sigma}_j\}_{j=1}^{25}$ is the grid of estimated beliefs for the given level of education, displayed in Figure 4. I use the resulting populations of 12,500 agents to compute model-implied moments.

5. Life Cycle Model Estimates

The results from structurally estimating the life cycle model confirm that incorporating survey measurements of beliefs improves the model's capacity to explain the savings and portfolios of U.S. households. For both high school and college graduates, replacing model-consistent (F.I.R.E.) beliefs with a specification that fits survey measurements reduces the distance between the model's predictions and the targeted moments of the

[36] The 25 estimated-beliefs agents share the same shock realizations of the F.I.R.E. agent they are replacing. They differ only in their beliefs about stock-fund returns.

Table 8: S.M.M. estimated parameters under different belief models

	College		High-School	
	F.I.R.E.	Est. Beliefs	F.I.R.E.	Est. Beliefs
CRRA (ρ)	11.396	5.114	8.607	4.231
	[11.344; 11.571]	[5.058; 5.150]	[8.494; 8.612]	[4.191; 4.252]
Disc. Fact. (β)	0.634	0.886	0.331	0.761
	[0.621; 0.642]	[0.884; 0.890]	[0.311; 0.355]	[0.751; 0.770]
Entry Cost ($F \times 100$)	1.041	0.000	3.116	2.576
	[0.485; 1.825]	[0.000; 0.000]	[2.848; 3.513]	[2.292; 2.673]
SMM Loss, $L^e(\hat{\vartheta})$	5.264	2.857	15.984	3.998

The brackets under each point estimate are 95% confidence intervals that come from estimating a surrogate model on bootstrapped moments, see Appendix F for details. "F.I.R.E" stands for full-information rational-expectations and "Est. Beliefs" corresponds to the heterogeneous beliefs specification, both described in Section 4. The "SMM Loss" row displays the value of the Simulated Method of Moments loss function (Equation 6) attained by the given parameter values and belief specifications.

data. In both cases, the improvement comes from the capacity to fit low participation rates with low participation costs and low portfolio shares with moderate levels of risk aversion. This is possible because in the estimated distribution of beliefs not everyone thinks that there is an equity premium, and those who do underestimate the risk-return compensation that the stock market offers.

For both high school and college graduates, the life cycle model fits the targeted moments more closely when it uses estimated beliefs instead of F.I.R.E. beliefs. Table 8 displays the estimated parameters for each level of educational attainment and specification of beliefs, along with the loss function (Equation 7) evaluated at the estimates. The loss function aggregates squared differences between the model-implied and empirical moments; thus, it serves as an index of how well each model fits the age profiles of wealth-to-income ratios, stock market participation rates, and conditional shares of wealth in stocks. Specifications that use the beliefs estimated from survey measurements have lower losses than their F.I.R.E. counterparts. The reductions are substantial: 75% for high school graduates (15.984 to 3.998) and 46% for college graduates (5.264 to 2.857).

Figure 9 compares the predicted moments of different model specifications with their

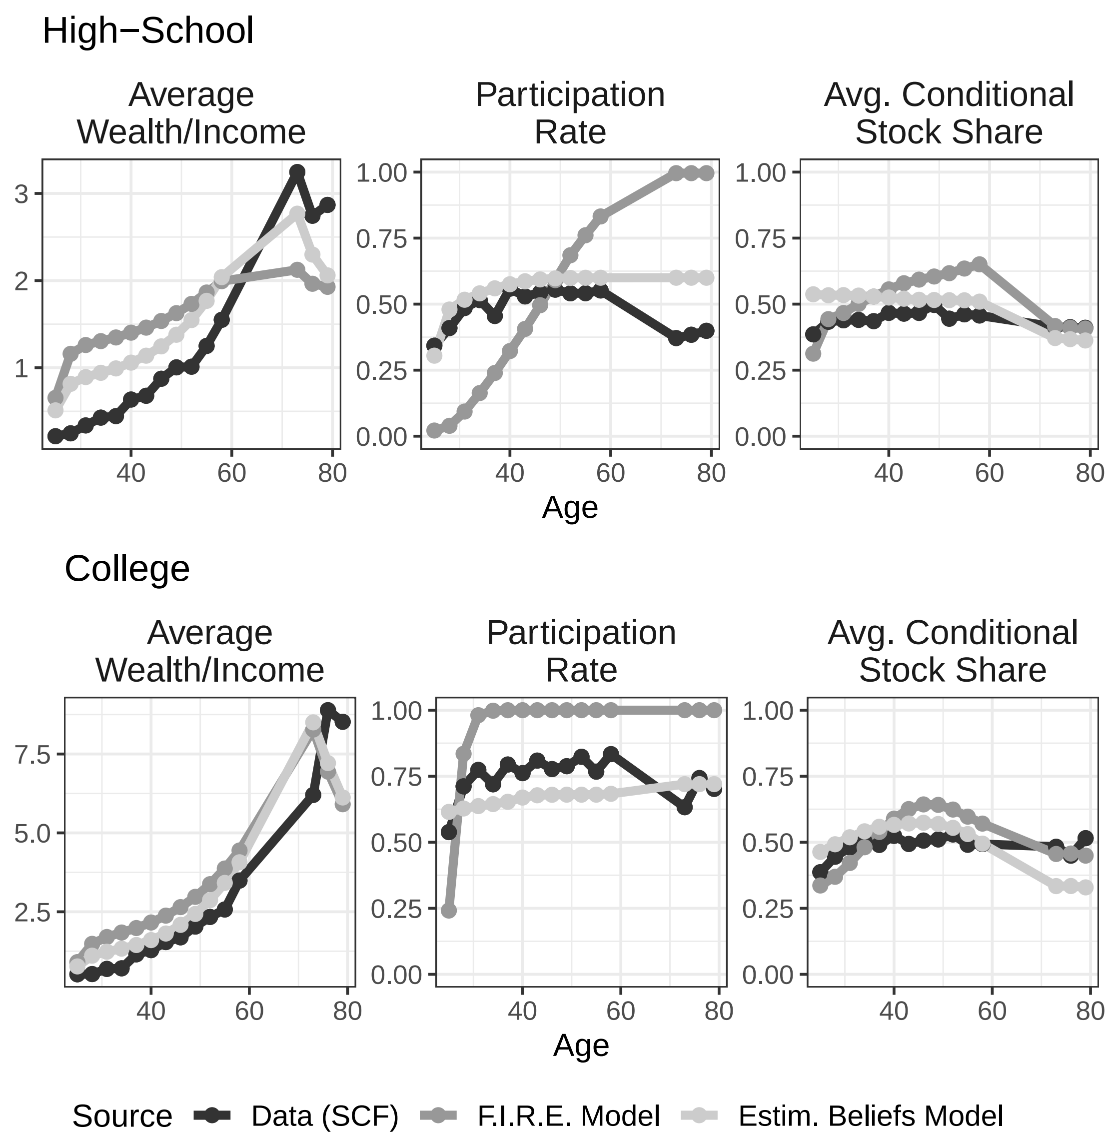

Each dot in the figure represents the relevant statistic calculated over a three-year age-group. See the main text for precise definitions of the age groups and statistics. Ages 60 to 71 are omitted because of transitional dynamics into retirement that the model does not account for.

Figure 9: Model fit of targeted moments

empirical counterparts, revealing the sources of the improvement in their measured fit.

The F.I.R.E. model does not replicate the humped-shaped participation rates of high school graduates, which peak at less than 60%. Instead, it produces participation rates that start at 2% and increase with age until they reach 100%. The reason is that under this specification of beliefs every agent who saves even modest amounts wants to participate and the only way to prevent them from doing so is to impose high entry costs. The costs can only generate participation rates that increase with age because once an agent has paid them he can participate for the rest of his life. Therefore, to generate a participation profile with a low peak of less than 60%, the F.I.R.E. model would need high entry costs that would make participation at younger ages counterfactually low. Instead, the best-fitting entry cost of 3.1% of annual income generates participation rates that do not match the shape of the true age profile but whose average across age bins (51%) is close to that in the data (48%).

In contrast, the model that uses the estimated beliefs specification accurately reproduces the participation rates of high school graduates before retirement. Because not all agents believe that there is an equity premium under this specification, the model generates participation rates that are moderate. This would be the case even in the absence of entry costs. The model fits increasing participation rates between ages 24 and 40 with a cost of 2.6% of permanent income. After age 40, the model uses the fraction of the population who does not believe in an equity premium to match the plateauing of participation rates until age 60. Without these agents, participation would continue to grow. Neither specification of beliefs can replicate the decline in participation rates after retirement that occurs in the data. This is due to the structure of costs: since there are no per-period costs associated with owning stocks, agents who already participate have little incentive to completely exit the stock market.

For college graduates, neither specification of beliefs replicates stock market participation rates perfectly: the F.I.R.E. model overestimates them and the estimated beliefs model underestimates them. As was the case with high school graduates, the F.I.R.E. model is constrained by the fact that all agents with sufficient savings want to own stocks. For college graduates, it uses a cost of 1% of annual permanent income that reduces early participation, but all agents overcome this cost by age 40, leading to a 100% participation rate for most of the life cycle. The opposite problem occurs in the estimated beliefs model: the fraction of agents that think there is an equity premium is lower than the actual participation rates of college graduates aged 40 to 60. Therefore, despite its null estimated participation cost, this model underestimates participation rates for most of the life cycle.

The differences between the wealth-to-income ratios and conditional stock shares generated by the two specifications of beliefs are subtler. For high school graduates, both models overestimate the wealth-to-income ratios of young agents and underestimate those of retirees; these errors are greater for the F.I.R.E. model. Similar issues appear in the wealth-to-income ratio of college graduates but to a lesser degree, with both specifications of beliefs tracking the empirical age profile more closely. The average conditional shares of wealth in stocks produced by the models are close to the empirical age profiles for both levels of educational attainment and both specifications of beliefs. The profiles are flatter than those predicted by baseline frictionless calibrations and this brings them closer to the empirical profiles. The main noticeable discrepancy in the model with estimated beliefs is the reduction of the average conditional share of wealth in stocks of college graduates after retirement. This reduction comes from a change in the composition of participants when the rebalancing tax is removed: some agents with pessimistic beliefs (but who still think that there is an equity premium) enter the market, and they drive down the average conditional share.

While the two specifications of beliefs produce qualitatively similar wealth-to-income ratios and conditional shares of wealth in stocks, they rely on different mechanisms to generate them. The main difference is how the two specifications reduce the conditional share of wealth in stocks to the moderate levels observed in the data. Models with F.I.R.E.

beliefs rely on high relative risk aversion coefficients of $\rho = 8.6$ for high school graduates and $\rho = 11.4$ for college graduates. Since all agents believe in a large equity premium under this specification, the only way to dissuade participants from allocating large shares of their savings to stocks is to make them extremely risk averse. On the other hand, under the estimated beliefs specification, Table 5 shows how even agents who believe in an equity premium think (on average) that the risk-return trade-off offered by stocks is not as attractive as historical benchmarks suggest. This feature enables the specification with estimated beliefs to match the empirical age-profiles of the conditional shares of wealth in stocks with lower relative risk aversion coefficients of $\rho = 4.2$ for high school graduates and $\rho = 5.1$ for college graduates.

The different relative risk aversion coefficients required by the F.I.R.E. and estimated-beliefs specifications produce differences in how the models fit wealth-to-income ratios. High relative risk aversion coefficients increase agents' precautionary saving, preventing the models from producing agents with low wealth. This is evident in Figure 9, where both belief specifications struggle to match the savings of younger agents, especially high school graduates. To counteract the effect of precautionary saving on the wealth of young agents, the models use lower time-discount factors (β) than those typically found in the macroeconomics and labor-economics literature.[37] This effect is stronger for F.I.R.E. models because of their higher relative risk aversion coefficients; it results in $\beta = 0.33$ for high-school graduates and $\beta = 0.63$ for college graduates, implying that agents discount their future utility at rates of 67% and 37% per year respectively. The models that use the estimated beliefs instead can afford higher time-discount factors of $\beta = 0.76$ for high-school graduates and $\beta = 0.89$ for college graduates because of the lower pressure of conditional portfolio shares on the relative risk aversion coefficients.

[37]Most studies that use consumption-saving models to match wealth find annual discount factors $\beta > 0.9$, see e.g., Gourinchas and Parker (2002), Cagetti (2003), De Nardi, French, and Jones (2010), and Carroll, Slacalek, et al. (2017). However, because of the higher relative risk aversion coefficients needed to match portfolio shares, lower estimated discount factors are frequently found in the household finance literature, see e.g., Fagereng, Gottlieb, and Guiso (2017) and Catherine (2021).

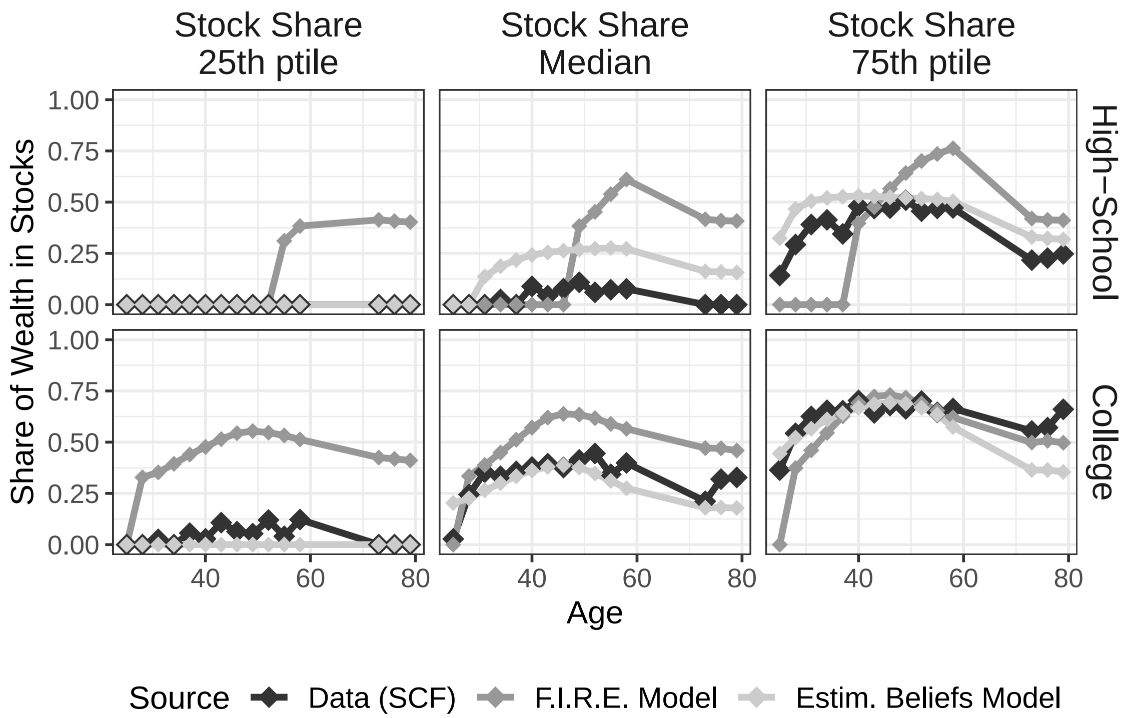

Each marker in the figure represents the relevant statistic calculated over a three-year age-group. See the main text for precise definitions of the age groups. None of the moments in this figure were targeted in estimation. Ages 60 to 71 are omitted because of transitional dynamics into retirement that the model does not account for.

Figure 10: Model fit of non-targeted moments

In sum, replacing F.I.R.E. beliefs with the estimated specification of beliefs produces a better fit of targeted moments, lower estimates of relative risk aversion coefficients, higher estimates of discount factors, and lower estimates of entry costs. These conclusions are robust to the sampling variation of the targeted moments shown in Figure 8. In Appendix F, I follow Chen, Didisheim, and Scheidegger 2021 and Catherine et al. 2022 in approximating the structural life cycle models with deep neural networks. I use the approximate models to demonstrate that these qualitative conclusions would hold if the estimation exercise was repeated for 500 different bootstrapped values of the targeted parameters. The confidence intervals in Table 8 come from the distributions of these bootstrapped estimates.

Figure 10 presents age profiles for the quartiles of the unconditional share of wealth in

Table 9: Average share of wealth in stocks by subjective mean quintiles

	Quintile of $\widehat{\mu}_i$				
Dataset	1	2	3	4	5
HRS	17.2%	24.5%	25.1%	32.8%	37.4%
Simulated HRS	9.3%	26.3%	34.4%	38.6%	40.9%

For each dataset (true and simulated HRS), I calculate the individual-level estimate of the mean of the subjective distribution of log-returns $\widehat{\mu}_i$ (see Equation 3. I aggregate the datasets at the individual level by taking the average of the share of wealth in stocks across waves. Then, I split the datasets by quintiles of $\widehat{\mu}_i$ and report the average share of wealth in stocks between those quintiles. The "Simulated HRS" matches the sample composition and response patterns of the true HRS, but uses my estimated structural model to produce probabilistic responses and portfolio choices.

stocks, which were not directly targeted in estimation. The estimated beliefs specification models fit these moments better than their F.I.R.E. counterparts. The F.I.R.E. models struggle to produce agents that participate in the stock market but invest a low share of their wealth in stocks. Because they rely on the participation cost, the age profiles that they imply for unconditional shares tend to start at 0% and then jump to higher values when agents decide to enter. Additionally, they do not generate much variation in the share of wealth in stocks of those who participate: the different percentiles of the unconditional share are close at most ages. In contrast, the models with estimated beliefs generate a distribution for the unconditional share whose different percentiles follow qualitatively different trajectories across the life cycle—the levels and shapes of the 25th and 75th percentiles are different. The different percentiles implied by the models with estimated beliefs track their empirical counterparts closely for both high school and college graduates.

A core prediction of the model that is not targeted by the estimation process is the relationship between people's beliefs and their portfolio choices. To study this relationship, I use the estimated models of beliefs and life cycle portfolio choices to simulate a synthetic dataset that resembles the HRS. For each respondent in the HRS sample[38] I create a simulated counterpart with the same level of education. Then, I draw the beliefs

[38]I use the same sample from Section 3.1.2: financial respondents aged 50 to 65 that had at least a high school degree and answered a probabilistic question about stocks.

(μ_i, σ_i) and level of rounding $\mathcal{R}_i$ of every simulated agent from the estimated measurement system (Section 3.1.2). I generate the simulated responses to $\{P_{i,t}^{\geq 0}, P_{i,t}^{\leq -20}, P_{i,t}^{\geq 20}\}$ using the estimated distributions for the $\{\varepsilon_{i,t}^{\geq 0}, \varepsilon_{i,t}^{\geq 20}, \varepsilon_{i,t}^{\leq -20}\}$ shocks. Given their beliefs, I simulate the savings and portfolios of each agent using the life cycle model. Finally, I drop observations from the simulated dataset to exactly reproduce non-response patterns of the HRS sample: I observe the simulated counterpart of every respondent exactly the same number of times, answering the same questions, at the same ages.

The model does well at reproducing the relationship between the portfolios of HRS respondents and their expected log-returns. I use the noisy and rounded probabilistic responses in the simulated dataset to calculate individual estimates of the mean log-return $\widehat{\mu}_i$ as I did with the true HRS (see Equation 3). Using these estimates instead of the simulated μ_i approximates the role of noise and rounding in attenuating the relationship between beliefs and portfolio choices in the true HRS. Table 9 presents the average unconditional share of wealth in stocks for different quintiles of $\widehat{\mu}_i$, comparing the true and simulated HRS samples. The table shows that the true and simulated relationship between beliefs and portfolios are similar. This is notable given that this relationship was not targeted: the model of beliefs and the portfolio choice model were estimated independently.

Overall, using the specification of beliefs from survey measurements improves the life cycle model's fit of U.S. households' portfolios with moderate levels of risk aversion, lower financial costs of entry, and higher time-discount factors than its F.I.R.E. counterparts. The model also reproduces non-targeted features of the distribution of portfolio shares and its relationship with expectations.

6. The Welfare Costs of Misspecified Beliefs

This section considers a scenario in which stocks continue to perform as they have historically. In this scenario, the beliefs estimated from survey measurements would be misspecified because they differ from the historical distribution of returns. An objective observer expects lower welfare for an agent with misspecified beliefs than for a counterpart with accurate beliefs because the decisions of the former are based on an incorrect model of the world. Using the estimated life cycle model, I quantify these welfare shortfalls for agents with the beliefs estimated from survey measurements. The metric that I use is the fraction of permanent income that the objective benevolent observer would take from each agent in exchange for correcting their beliefs.[39] Average welfare shortfalls start out at less than 3.5% of permanent income for young agents but follow a hump shape that peaks around the age of retirement at 8% to 14% of permanent income, depending on their level of education.

Since beliefs can differ from the data generating process for risky returns, the discounted welfare that agents expect can differ from the discounted welfare that an *objective observer*—one who knew the true data generating processes and their decision rules—would expect them to receive. Let $\mathfrak{V}_t(P_t, M_t, N_t, \texttt{Paid}_t; \mu, \sigma)$ denote the discounted welfare that we objectively expect an agent to receive starting from age t and state $(P_t, M_t, N_t, \texttt{Paid}_t)$[40] if his beliefs about returns are (μ, σ). The function $\mathfrak{V}_t(\cdot)$, which I define formally in Appendix G, uses the agent's own preferences to discount the future and assumes that he will follow the policy functions that solve his dynamic problem according to his beliefs. The difference between $\mathfrak{V}_t(\cdot)$ and the agent's subjective value function $V_t(\cdot)$ is that $\mathfrak{V}_t(\cdot)$ uses the true data-generating parameters $(\mu^{SP500}, \sigma^{SP500})$ in its expectations. Thus, if $(\mu, \sigma) \neq (\mu^{SP500}, \sigma^{SP500})$, then $\mathfrak{V}_t(\cdot; \mu, \sigma) \neq V_t(\cdot)$ in general.

To measure the expected welfare shortfalls that misspecified beliefs cause, I use a compensating variation. The metric is the proportional reduction in permanent income (present and future) that would need to accompany the correction of an agent's beliefs to leave him at his current level of objectively-expected welfare. Formally, for an agent with

[39]The objective benevolent observer is necessary because agents are not aware that their beliefs are misspecified. They would not pay to change their beliefs because they think that they are correct.

[40]$\texttt{Paid}_t$ indicates whether the agent has paid the fixed stock-market entry cost or not.

state variables $(P_t, M_t, N_t, \texttt{Paid}_t)$ and beliefs (μ, σ), I find the λ that satisfies:

$$\mathfrak{V}_t\left(P_t, M_t, N_t, \texttt{Paid}_t; \mu, \sigma\right) = \mathfrak{V}_t\left((1-\lambda)\times P_t, M_t, N_t, \texttt{Paid}_t; \mu^{SP500}, \sigma^{SP500}\right). \quad (9)$$

The metric λ can be interpreted as the maximum fraction of permanent income that an altruistic and objective planner would be willing to take from the agent in exchange for amending his beliefs about returns. Equation 9 cannot be solved for λ analytically, but it shows that an agent's expected welfare shortfall is contingent on his age, permanent income, assets, whether or not he has paid the entry cost, and his beliefs.[41]

I calculate the expected welfare shortfalls from the estimated heterogeneous beliefs for individuals with the preferences and assets implied by the estimated life cycle model. For each level of education, I find the expected welfare functions $\{\mathfrak{V}_t(\cdot)\}_{t=24}^{100}$ using the preference and cost parameters from Table 8 for the estimated beliefs specifications. Then, I simulate the lives of agents with beliefs about stock returns drawn from the education-specific estimated distributions. These simulations use the same population sizes and shock realizations as in the estimation process (see Section 4). Finally, I evaluate the estimated welfare shortfall λ of every agent at every year.

The welfare shortfalls from misspecified beliefs follow a hump shape across the life cycle, starting low in agents' youth and peaking at the age of retirement. Figure 11 presents the expected welfare shortfalls λ of high school and college graduates at every age, evaluated at their simulated assets and incomes. For both levels of education, the average shortfall is lowest at the youngest age, with values of 0.9% and 3.4% of permanent income for high school and college graduates respectively. Accurate beliefs yield their benefits later in life, when agents have accumulated more savings and they rely on them for their consumption. The impatient discounting of these future benefits results in the low initial welfare shortfalls. The shortfalls increase as agents accumulate wealth and

[41] Because the model is homothetic in permanent income, λ can be expressed as a function of assets normalized by permanent income, and the other states.

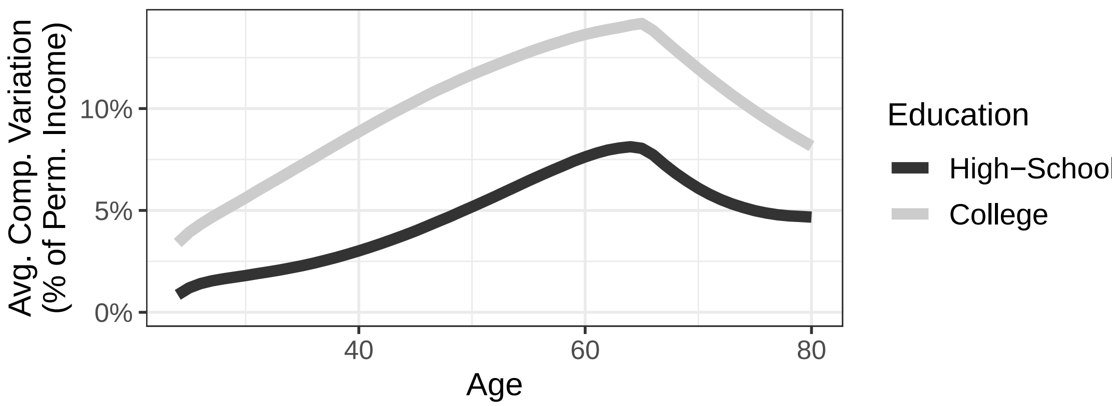

The figure presents average welfare losses from misspecified beliefs about risky-asset returns at every age. I simulate populations of agents that behave according to the beliefs and preferences estimated in Sections 3.1.2 and 5. Then, for every agent-period observation, I compute the expected welfare shortfall λ defined in Equation 9. I report the average of this measure for every age-education combination.

Figure 11: Expected welfare shortfalls across the life cycle

approach retirement, reaching their peak at age 64 for high school graduates with a value of 8.1% of permanent income and at age 65 for college graduates with a value of 14.2% of permanent income. The shortfalls decline progressively thereafter, as agents deplete their wealth and life expectancy. Despite college graduates having beliefs closer on average to the historical benchmark (see Table 5), their average welfare shortfalls are greater than those of high school graduates at all ages. This is due to their higher estimated discount factor, higher levels of savings, and lower replacement rates of retirement income.

The magnitude, age patterns, and educational differences in the welfare shortfalls from distorted beliefs are consistent with several findings from the financial literacy literature. First, Figure 11 indicates that agents would derive the greatest (discounted) benefits from correcting their misconceptions about the risky asset at ages close to retirement. Empirically, different measures of financial literacy follow a similar hump shape across the life cycle of U.S. respondents and peak close to this age range (Lusardi and Mitchell 2023). Additionally, models that allow for endogenous accumulation of financial knowledge prescribe that it must peak around the age of retirement and at higher levels for those with more education (Lusardi, Michaud, and Mitchell 2017, 2020). While these models

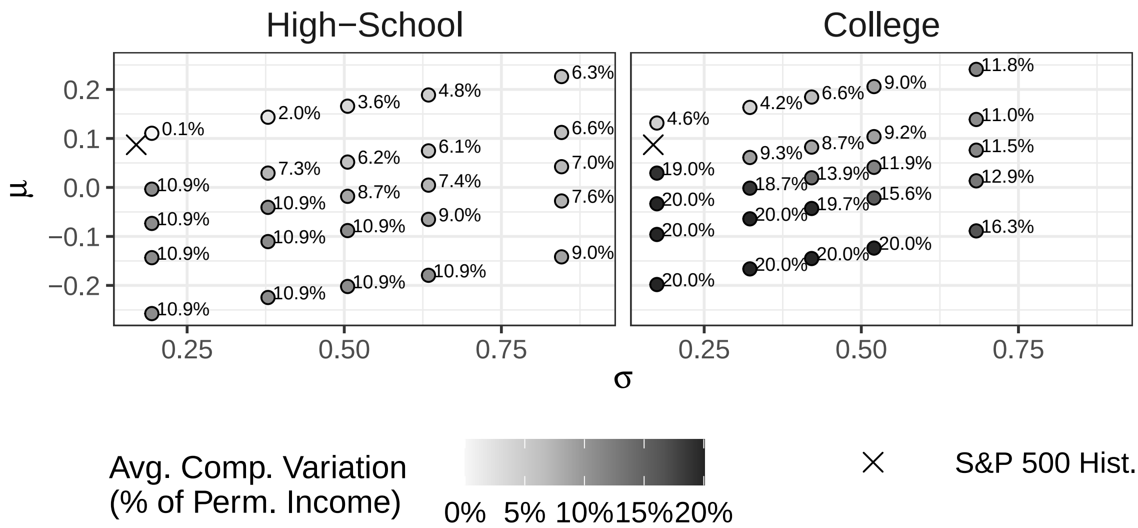

The figure presents average welfare losses from misspecified beliefs about risky-asset returns for every set of beliefs considered in the life-cycle model estimation. I simulate populations of agents that behave according to the beliefs and preferences estimated in Sections 3.1.2 and 5. Then, for every agent-period observation, I compute the expected welfare shortfall λ defined in Equation 9. I report the average of this measure for every belief-education combination at age 65.

Figure 12: Expected welfare shortfalls from different beliefs at age 65

find start-of-life welfare shortfalls of similar magnitudes to those that I calculate,[42] my findings show that, as a percentage of income, the shortfalls can be far greater around the age of retirement. These patterns highlight the years preceding retirement as a potential "teachable moment" for financial knowledge interventions since people have accumulated enough wealth to put their knowledge to use and anticipate that they will become more reliant on this knowledge for their support. Identifying these "teachable moments" has been highlighted as a crucial determinant of the success of these interventions in changing downstream behaviors (Kaiser and Menkhoff 2017). Workplace interventions, for instance, are a modality of financial knowledge program that has gathered increasing interest (see Clark 2023; Lusardi and Mitchell 2023) and which these results favor over earlier interventions.

The shortfalls in welfare vary substantially with individuals' beliefs about stock returns,

[42] For example, in their model of costly investments in financial knowledge, Lusardi, Michaud, and Mitchell (2017) find that endowing 25-year-olds with perfect financial knowledge would increase their welfare by magnitudes comparable to 2-3% increases in their lifetime consumption.

and the largest impacts fall on those with beliefs that discourage them from ever owning stocks. Figure 12 displays the average expected welfare shortfall for simulated individuals with different beliefs and levels of education at age 65, their last working year. For both levels of education, the largest welfare shortfalls occur among individuals with low subjective means and volatilities of log-returns. These individuals do not perceive an equity premium and, therefore, they refrain from ever participating in the stock market. The average welfare shortfall for these groups of non-participants is 10.9% of permanent income for high school graduates and 20.0% for college graduates. Welfare shortfalls are lower for the group of agents who perceive an equity premium, progressively increasing as beliefs move further away from the model-consistent benchmark. The average welfare shortfall among those who believe there is an equity premium is 6.1% of permanent income for high-school graduates and 11.9% for college graduates.

The results in this section highlight considerations for the design and application of interventions aimed at changing financial knowledge and behaviors. Lusardi and Mitchell (2023) stress the importance of pinpointing the groups for which these interventions have the greatest effects and suggest that they might not be cost-effective for groups that will not use their financial knowledge in a timely manner. The foregoing analysis suggests that interventions aimed at improving the financial knowledge of young and disadvantaged individuals with low savings can have smaller effects on welfare than those targeted at wealthier and more educated individuals. This conclusion is consistent with the findings of, e.g., Lusardi, Michaud, and Mitchell (2020). This socioeconomic gradient notwithstanding, the potential welfare effects of information treatments on the most disadvantaged are large. The question of how to design interventions effective at changing both financial knowledge and behaviors remains a crucial area of ongoing research (see Kaiser and Menkhoff 2017; Kaiser, Lusardi, et al. 2022; Clark 2023).

7. Misspecified Beliefs and the Value of Public Policies

The preceding sections demonstrate that different assumptions about beliefs can produce different estimates of preferences in saving and portfolio choice models. Do these differences change the answers that these models offer to counterfactual questions about the welfare impacts of policy changes? This section demonstrates that they do, by calculating the welfare impacts of two simple policy changes under the different models estimated in Section 5.

The first policy change that I evaluate is the elimination of unemployment insurance. I model this change by making the unemployment benefits replacement factor zero, $\mathcal{U}^{\text{C.F.}} = 0$; this makes the worst transitory income shock that working agents can draw much worse. The second policy change that I evaluate is a reduction in Social Security benefits that makes the retirement income of all agents 20% lower. I model this change by shrinking the permanent income growth factor that is applied at retirement, $\Gamma_{i,66}^{\text{C.F.}} = 0.8 \times \Gamma_{i,66}$. The goal of these exercises is to clearly illustrate the effects of the estimated preferences, not to represent realistic policy proposals. The change in unemployment benefits highlights differences in the implied value of insurance and the change to retirement benefits highlights differences in the value of future income.

The metric that I use to represent the welfare effects of policy changes is the one-time cash transfer that, if bundled with the policy change, would leave each agent at their original level of objectively-expected welfare. I express this quantity as a fraction of permanent income. Using the objective expected welfare function $\mathfrak{V}(\cdot)$ introduced in Section 6 and denoting its counterpart under the counterfactual changes with $\mathfrak{V}^{\text{C.F.}}(\cdot)$, the welfare metric CV solves

$$\mathfrak{V}_t\left(P_t, M_t, ...\right) = \mathfrak{V}_t^{\text{C.F.}}\left(P_t, M_t + CV \times P_t, ...\right). \tag{10}$$

Table 10: Average Welfare Losses from Policy Changes Under Different Models

	High School		College	
Model	Age 30	Age 60	Age 30	Age 60
Welfare Loss from Removal of U.I. (as % of a year's wages)				
F.I.R.E.	46%	25%	120%	10%
Est. Beliefs	41%	7%	50%	4%
Welfare Loss from Reduction of S.S. (as % of a year's wages)				
F.I.R.E.	0%	6%	1%	98%
Est. Beliefs	0%	60%	8%	123%

See the main text for the definition of the policy changes. The welfare losses are calculated at the simulated states and ages of the populations used in the estimation of the life-cycle models. Each row indicates which preferences and (and corresponding policy functions) from Table 8 were used to calculate welfare losses. Losses are expressed as a fraction of agents' permanent income; see Equation 10.

I evaluate this metric for each agent and period in the simulated populations used to estimate the life cycle model (see Section 4). This gives the metric the interpretation of the losses incurred by agents at different ages and states if the policies are introduced as a surprise.

The preference estimates and policy functions generated by different specifications of beliefs produce very different estimates of the welfare effects of the policy changes. Table 10 presents the average welfare costs of the policy changes for individuals of different ages and levels of education, comparing their implied values under the F.I.R.E. models and the models that use estimated beliefs. As the table shows, the average welfare losses estimated under the F.I.R.E. models can be multiples of those estimated under the model with heterogeneous beliefs and vice-versa. These differences come mainly from the different preference estimates that the models imply (see Table 8). The removal of unemployment benefits produces much greater welfare losses under the F.I.R.E. model, because its higher estimates of risk aversion raise the value of insurance. With a relative-risk aversion (ρ) of 11, 30-year-old college graduates in the F.I.R.E. model value unemployment benefits at 120% of a year's wages; this falls to less than half with the risk aversion of 5 in the model that uses estimated beliefs. For retirement benefits that occur in the future, valuations are higher under the model that uses estimated beliefs due to its higher estimates of the time-discount factor (β). For instance, at age 60, the difference in discount factors for high-school graduates make the welfare losses from the reduction in social security benefits implied by the model with estimated beliefs 10 times larger than those implied by the F.I.R.E. model.

Different specifications of beliefs could simply be considered alternative strategies to explain the same set of empirical facts. With beliefs given, other features of our models—preferences in this case—can change to accommodate the targeted moments. However, as the exercises in this section demonstrate, the alternative explanations of facts derived from different specifications of beliefs can have starkly different implications for the type of counterfactual and welfare questions that the structure of our models affords.

8. Concluding Remarks

Dominitz and Manski (2007) note that "many households (...) are not as convinced as economists are about the existence of an equity premium." Since the their pioneering work, the HRS has expanded the range of available measurements of household expectations in both time and variety, having collected nearly two decades of measurements and now including three different questions regarding equity returns. This paper uses this expanded set of measurements to show that many households remain unconvinced of the existence of an equity premium and that even those who seem to believe in its existence deem it smaller in risk-adjusted terms than what economists usually assume.

These facts about measured expectations provide a qualitative explanation for why most households do not invest most of their wealth in equities. The exercises carried out in this paper quantitatively evaluate the plausibility of this explanation. They demonstrate that the explanation has several attractive features: it substantially enhances the capacity of the considered model to reproduce both portfolio choices and savings, and their relationship with age and education; it brings the estimates of unobserved preference parameters to ranges more consistent with alternative sources of evidence; and it is backed by a robust body of measurements and empirical results.

Important challenges and questions remain. The model put forward in this paper faces difficulties in matching the low savings of groups like young households and those without a college degree. Allowing households to borrow and modeling the social programs that low-wealth households use to smooth their consumption are possible ways to reduce these difficulties. Additionally, the model proposed in this paper assumes that households do not change their beliefs about equity returns or that they do not react to short-term fluctuations in their opinions. I make this assumption to replicate features of belief measurements like the "dominance of individual fixed effects" (Giglio et al. 2021) and the fact that households' portfolios have weak responses to changes in their elicited expectations. Questions such as why households do not learn in ways that eliminate the persistent heterogeneity in their measured beliefs or why the association between changes in their elicited expectations and changes in their portfolios is weak are left unresolved. The increasing availability of individual-level measurements of expectations and portfolios can help address these questions.

Author Affiliation

Mateo Velásquez-Giraldo,
Board of Governors of the Federal Reserve System,
Division of Research and Statistics,
Macroeconomic and Quantitative Studies Section.

Appendix

A. Estimating the Model of Beliefs

As discussed in Section 2.2 and in previous studies like Giustinelli, Manski, and Molinari (2022), people round their answers to probabilistic questions. For each of the probabilistic questions about stock returns, Table 3 shows the fraction of all answers that are multiples of 5%, 10%, 25%, 50% and 100%. The Table shows that, for each question, less than 2.5% of the answers are not multiples of 5%. Based on this fact, 5% is the finest level of rounding in my model and I round the few answers that are not multiples of 5% to the nearest 5% multiple.

A.1 The Likelihood Function

The likelihood function and its derivation are similar to those in Kézdi and Willis (2011) and Ameriks, Kézdi, et al. (2020).

Denote the set of parameters of the beliefs model with

$$\vartheta^{\mathrm{B}} \equiv \{\nu_\mu, \nu_\sigma, \Psi, \Sigma, \vec{\wp}\},$$

and let the data consist of triplets of responses to probabilistic questions

$$\left\{\left\{P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20}\right\}_{t\in\mathcal{T}(i)}\right\}_{i\in\mathcal{I}}$$

where $\mathcal{I}$ is the set of respondents and $\mathcal{T}(i)$ denotes the set of time periods in which individual i answered the probabilistic questions.

The first step in evaluating the likelihood function for ϑ^{B} is to find the equiprobable grid for (μ, σ) that is associated with $\vartheta^{\mathrm{B.\,Grid}} \equiv \{\nu_\mu, \nu_\sigma, \Psi\}$. I construct an equiprobable grid with n^2 points that approximates distribution 2 using the following steps.

1. Find an equiprobable n-point grid for $x \sim \mathcal{N}(\nu_\sigma, \Psi_{2,2}) | x > 0$. Denote the grid with $\sigma^\#$.

2. For every σ in $\sigma^\#$, find an equiprobable n-point grid for the distribution of μ_i conditional on $\sigma_i = \sigma$, which given Equation 2 is

$$\mathcal{N}\left(\nu_\mu + \frac{\Psi_{1,2}}{\Psi_{2,2}}(\sigma - \nu_\sigma)\,,\, \Psi_{1,1} - \frac{\Psi_{1,2}^2}{\Psi_{2,2}}\right).$$

   Denote that grid with $\mu^\#(\sigma)$.

3. The joint n^2-point grid will be

$$(\mu, \sigma)^\# \equiv \{(\mu, \sigma) : \mu \in \mu^\#(\sigma), \sigma \in \sigma^\#\}.$$

The likelihood of an agent's response depends both on the parameters of their subjective distribution of returns (μ_i, σ_i) and on the degree to which they round their answers. I denote agent i's level of rounding with $\mathcal{R}_i$ and consider 5%, 10%, 25%, 50%, and 100% as the possible levels to which agents round their answers. Therefore, $\forall i \;\; \mathcal{R}_i \in \{5, 10, 25, 50, 100\}$.

The likelihood of an answer that has been rounded is that of the interval of all the real numbers that round to that answer. To facilitate the representation of these intervals, define the following two sets of functions:

- $\underline{u}_{\mathcal{R}}(x)$ gives the lowest number in $[0, 1]$ that rounds to x when the level of rounding is $\mathcal{R}$. For instance, $\underline{u}_5(0.15) = 0.125$, $\underline{u}_{10}(0.30) = 0.25$, $\underline{u}_{25}(0.5) = 0.375$, $\underline{u}_{50}(100) = 0.75$, and $\underline{u}_{100}(0.0) = 0.0$.

- $\bar{u}_{\mathcal{R}}(x)$ gives the highest number in $[0, 1]$ that rounds to x when the level of rounding is $\mathcal{R}$. For instance, $\bar{u}_5(0.15) = 0.175$, $\bar{u}_{10}(0.30) = 0.35$, $\bar{u}_{25}(0.5) = 0.625$, $\bar{u}_{50}(100) = 1.0$, and $\bar{u}_{100}(0.0) = 0.5$.

With these functions and the response model from Equation 1, we can say that if agent i rounds their answers to the $\mathcal{R}_i$-level and has subjective-distribution parameters (μ_i, σ_i), then

$$\begin{aligned} P_{i,t}^{\geq 0} = x \leftrightarrow & \underline{u}_{\mathcal{R}_i}(x) \leq \Phi\left(\frac{\mu_i}{\sigma_i} + \varepsilon_{i,t}^{\geq 0}\right) \leq \bar{u}_{\mathcal{R}_i}(x) \\ \leftrightarrow & \Phi^{-1}(\underline{u}_{\mathcal{R}_i}(x)) - \frac{\mu_i}{\sigma_i} \leq \varepsilon_{i,t}^{\geq 0} \leq \Phi^{-1}(\bar{u}_{\mathcal{R}_i}(x)) - \frac{\mu_i}{\sigma_i}, \end{aligned} \tag{11}$$

$$\begin{aligned} P_{i,t}^{\geq 20} = y \leftrightarrow & \underline{u}_{\mathcal{R}_i}(y) \leq \Phi\left(\frac{\mu_i - \ln 1.20}{\sigma_i} + \varepsilon_{i,t}^{\geq 20}\right) \leq \bar{u}_{\mathcal{R}_i}(y) \\ \leftrightarrow & \Phi^{-1}(\underline{u}_{\mathcal{R}_i}(y)) - \frac{\mu_i - \ln 1.20}{\sigma_i} \leq \varepsilon_{i,t}^{\geq 20} \leq \Phi^{-1}(\bar{u}_{\mathcal{R}_i}(y)) - \frac{\mu_i - \ln 1.20}{\sigma_i}, \end{aligned} \tag{12}$$

and

$$\begin{aligned} P_{i,t}^{\leq -20} = z \leftrightarrow & \underline{u}_{\mathcal{R}_i}(z) \leq \Phi\left(\frac{\ln 0.8 - \mu_i}{\sigma_i} + \varepsilon_{i,t}^{\leq -20}\right) \leq \bar{u}_{\mathcal{R}_i}(z) \\ \leftrightarrow & \Phi^{-1}(\underline{u}_{\mathcal{R}_i}(z)) - \frac{\ln 0.8 - \mu_i}{\sigma_i} \leq \varepsilon_{i,t}^{\leq -20} \leq \Phi^{-1}(\bar{u}_{\mathcal{R}_i}(z)) - \frac{\ln 0.8 - \mu_i}{\sigma_i}. \end{aligned} \tag{13}$$

Equations 11-13 and the assumption that $(\varepsilon^{\geq 0}, \varepsilon^{\geq 20}, \varepsilon^{\leq -20}) \sim \mathcal{N}(0, \Sigma)$ allow me to compute

$$\mathbf{P}\left(P_{i,t}^{\geq 0} = x, P_{i,t}^{\geq 20} = y, P_{i,t}^{\leq -20} = z | \left(\mu_i, \sigma_i\right), \mathcal{R}_i\right) \tag{14}$$

as the integral of a normal density over a cube. Since I do not use observations in which the answer to any of the questions is “do not know/refuse,” observations where responses for at least one of the question is missing correspond to instances where not all questions were asked. For instance, in all observations before 2008, only $P^{\geq 0}$ was asked. For these observations, the likelihood of the given answers omits the questions that were not asked and it becomes an integral over a real interval (if only one question is asked) or a rectangle (if two questions were asked). With this clarification, I use the same notation in Equation 14 for complete and incomplete sets of answers.

Now, I can write the likelihood of observing an individual i with responses

$$\{(P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20})\}_{t \in \mathcal{T}(i)}$$

conditional on his rounding type as

$$\ell \left(\left\{ (P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20}) \right\}_{t \in \mathcal{T}(i)} | \mathcal{R}_i \right) = \frac{1}{n^2} \sum_{(\mu_i, \sigma_i) \in (\mu, \sigma)^{\#}} \left(\prod_{t \in \mathcal{T}(i)} \mathbf{P} \left(P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20} | \left(\mu_i, \sigma_i \right), \mathcal{R}_i \right) \right),$$

where I have integrated over the n^2 equiprobable (μ, σ) grid-points. The unconditional likelihood follows from integrating over the rounding types using the prior $\vec{\wp}$,

$$\ell \left(\left\{ (P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20}) \right\}_{t \in \mathcal{T}(i)} \right) = \sum_{\mathcal{R}_i \in \{5,10,25,50,100\}} \wp_{\mathcal{R}_i} \times \ell \left(\left\{ (P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20}) \right\}_{t \in \mathcal{T}(i)} | \mathcal{R}_i \right)$$

Finally, the log-likelihood function comes from aggregating over individuals

$$\ln \mathcal{L} \left(\left\{ \left\{ P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20} \right\}_{t \in \mathcal{T}(i)} \right\}_{i \in \mathcal{I}} | \vartheta^{\mathrm{B}} \right) = \sum_{i \in \mathcal{I}} \ln \ell \left(\left\{ (P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20}) \right\}_{t \in \mathcal{T}(i)} \right).$$

A.2 Parameter Estimates

I estimate the beliefs model by maximum likelihood for every level of education,

$$\vartheta_E^{\mathrm{B}} = \arg\max_{\vartheta} \ln \mathcal{L} \left(\left\{ \left\{ P_{i,t}^{\geq 0}, P_{i,t}^{\geq 20}, P_{i,t}^{\leq -20} \right\}_{t \in \mathcal{T}(i)} \right\}_{i \in \mathcal{I}_E} | \vartheta \right)$$

where E indexes educational attainment levels—high school and college graduates.

Table 11 presents the parameter estimates for every level of education. Readers should

Table 11: Maximum-likelihood estimates of the beliefs model

	High School	College
ν_μ	-0.108	-0.071
	(0.014)	(0.009)
ν_σ	0.499	0.418
	(0.009)	(0.008)
$\Psi_{1,1}$	0.019	0.016
	(0.001)	(0.001)
$\Psi_{2,1}$	0.011	0.008
	(0.002)	(0.001)
$\Psi_{2,2}$	0.060	0.035
	(0.005)	(0.003)
$\Sigma_{1,1}$	0.607	0.453
	(0.007)	(0.008)
$\Sigma_{2,2}$	0.574	0.460
	(0.010)	(0.011)
$\Sigma_{3,3}$	0.642	0.397
	(0.011)	(0.009)
$\wp 5$	0.403	0.512
	(0.006)	(0.009)
$\wp 10$	0.425	0.400
	(0.006)	(0.009)
$\wp 25$	0.043	0.029
	(0.003)	(0.004)
$\wp 50$	0.116	0.052
	(0.004)	(0.005)
$\wp 100$	0.013	0.007
	(0.002)	(0.002)
Log-Likeligood	-121667.124	-59763.868
N. Obs	24027	11184
N. Individuals	8463	3562
N. Excluded DK/RF Obs.	3216	613

Standard errors come from the inverse of the negative of the hessian of the log-likelihood function, evaluated exactly at the parameter estimates using automatic differentiation tools.

note, however, that the estimates of Equation 2 do not have the traditional mean-covariance interpretation of multivariate normal parameters. The reason is that the actual distributions from which (μ, σ) are drawn condition on the event $\sigma > 0$. These estimates are presented for completeness; for interpreting and comparing the belief distributions that go into the structural model, I refer readers to the depictions of the discretized belief distributions (Figure 4 in the main text).

B. Calibration of the Medical Expenditures Process

The shocks $\texttt{oop}_{i,t}$ represent the ratio of an agent's income that is used up by out-of pocket medical expenditures in a year. To approximate their distribution at different ages and for agents with different levels of education, I use the RAND HRS longitudinal file, which constructs various variables of interest in a manner that is consistent across HRS waves.

I start by defining a measure of what would be a retiree's household's income. The measure that I use includes the household's earnings, income from pensions and annuities, income from Social Security Disability and Supplemental Security Income, and income from Social Security retirement. This corresponds to the sum of the RAND HRS variables `IEARN`, `IPENA`, `ISSDI`, and `ISRET`. It also corresponds to their measure of "total income" minus government transfers, capital income, and income from "other sources."

Starting with the third wave of the HRS, The RAND HRS longitudinal file includes a measure of out of pocket medical expenses over the previous two years at the time of the interview, `OOPMD`.[43] I divide this measure by two to obtain an estimate of medical expenses over a year at the household level. The ratio between this measure and the previously defined income of the respondent's household—for households with strictly positive income—is what I take as the ratio of out-of-pocket medical expenditures to income over a year; I denote it with `oop`. Figure 6 depicts the distribution of `oop` for individuals of different ages and levels of education.

To construct discrete distributions that approximate the variability of `oop`, I start by grouping observations according to their level of education and age. I use 5-year age bins $[66, 70]$, $[71, 75]$,...,$[85, 90]$ and a final $[91, 100]$ bin. For each combination of age-group and education, I construct discrete equiprobable distributions using quantiles of the empirical distribution of `oop`. First, I split the $[0, 1]$ interval in n intervals of the same length, where n is the number of points of the discrete approximation—in my case, $n = 7$. Then, I

[43] The measure is not constructed for the first wave. In the second wave, the question that is used to build the measure had a different time horizon and therefore I exclude it.

Table 12: Discrete approximations of medical expenditures/income ratios

Age Group	Equiprobable Points						
High-School							
[50,55]	0.000	0.004	0.010	0.019	0.033	0.064	0.205
(55,60]	0.000	0.005	0.013	0.023	0.040	0.077	0.245
(60,65]	0.000	0.008	0.018	0.032	0.055	0.104	0.290
(65,70]	0.001	0.011	0.023	0.038	0.064	0.111	0.264
(70,75]	0.002	0.014	0.028	0.046	0.074	0.126	0.293
(75,80]	0.001	0.015	0.031	0.053	0.084	0.143	0.346
(80,85]	0.001	0.016	0.033	0.059	0.096	0.168	0.433
(85,90]	0.000	0.016	0.036	0.066	0.110	0.229	0.849
(90,100]	0.000	0.011	0.034	0.069	0.131	0.301	1.479
College							
[50,55]	0.000	0.003	0.007	0.012	0.021	0.039	0.121
(55,60]	0.001	0.005	0.010	0.016	0.027	0.049	0.163
(60,65]	0.002	0.007	0.014	0.023	0.040	0.078	0.227
(65,70]	0.003	0.010	0.019	0.031	0.050	0.089	0.227
(70,75]	0.004	0.013	0.024	0.039	0.060	0.103	0.262
(75,80]	0.004	0.015	0.028	0.047	0.074	0.123	0.294
(80,85]	0.004	0.017	0.033	0.054	0.089	0.155	0.410
(85,90]	0.002	0.015	0.033	0.057	0.100	0.191	0.719
(90,100]	0.000	0.015	0.039	0.075	0.160	0.389	1.485

The table presents approximations to the distribution of health expenditure shocks as a fraction of income. I approximate de distribution of these shocks for each age group and level of education with an equiprobable discrete distribution. Each row displays the seven points used to approximate the distribution for each age group and level of education. Each point has a probability of 1/7. See the text for a description of how I obtain the points.

take the midpoint of each interval and denote with Q the set of midpoints—for $n = 7$, $Q = \{0.071, 0.214, ..., 0.786, 0.929\}$. Finally, for every $q \in Q$, I obtain the q quantile of the empirical distribution of oop for the given group. My approximation of the distribution of oop is a discrete random variable where the possible draws are the n previously obtained quantiles and each of them occurs with probability $1/n$. Table 12 displays the points that I use in the model for every age group and level of education.

C. Recursive Formulation and Normalization of the Model

Individual subscripts are dropped for simplicity throughout this section.

An agent starts his life not having paid the stock-market entry cost. In periods when the cost has not been paid, the agent observes his risk-free resources and his permanent income, and then decides whether to enter the stock market or not. His value function is

$$V_t^{\mathtt{Out}}(M_t, P_t) = \max\{V_t^{\mathtt{Stay}}(M_t, P_t),\ V_t^{\mathtt{In}}(M_t - F \times P_t, 0, P_t)\},$$

where $V_t^{\mathtt{Stay}}(\cdot)$ is the value function of an agent who stays out of the stock market and $V^{\mathtt{In}}(\cdot)$ is the value function of an agent who has already paid the stock-market entry cost.

An agent who has just decided not to pay the stock-market entry cost decides how much to consume out of his assets, knowing that in the next time period he will have the opportunity to enter the stock market again. His value function is

$$V_t^{\mathtt{Stay}}(M_t, P_t) = \max_{C_t} u(C_t) + \beta\delta_{t+1}\mathbb{E}_t\left[V_{t+1}^{\mathtt{Out}}(M_{t+1}, P_{t+1})\right] + \not{\delta}_{t+1}\mathbb{B}(A_t)$$

Subject to:

$$0 \leq C_t \leq M_t$$

$$.$$

$$\begin{aligned} A_t &= M_t - C_t \\ M_{t+1} &= RA_t + Y_{t+1} \\ P_{t+1} &= \Gamma_{t+1}\psi_{t+1}P_t \\ Y_{t+1} &= \theta_{t+1}P_{t+1} \end{aligned}$$

Finally, an agent who has already paid the stock-market entry cost observes his balances in both the risky and risk-free assets and his permanent income, and then decides how to reallocate his balances and how much to consume. He forms expectations about the future knowing that he will not need to pay the entry cost again. His value function is

$$V_t^{\texttt{In}}(M_t, N_t, P_t) = \max_{C_t, D_t} u(C_t) + \beta\delta_{t+1}\mathbb{E}_t\left[V_{t+1}^{\texttt{In}}(M_{t+1}, N_{t+1}, P_{t+1})\right]$$

$$+\not\delta_{t+1}\mathbb{B}(A_t + \tilde{N}_t)$$

Subject to:

$$-N_t \leq D_t \leq M_t, \quad 0 \leq C_t \leq \tilde{M}_t$$

$$\begin{aligned}
\tilde{M}_t =& M_t - D_t\left(1 - 1_{[D_t \leq 0]}\tau\right) \\
\tilde{N}_t =& N_t + D_t \\
A_t =& \tilde{M}_t - C_t \\
M_{t+1} =& RA_t + Y_{t+1} \\
N_{t+1} =& \tilde{R}_{t+1}\tilde{N}_t \\
P_{t+1} =& \Gamma_{t+1}\psi_{t+1}P_t \\
Y_{t+1} =& \theta_{t+1}P_{t+1}
\end{aligned}.$$

I assume that the utility function $u(\cdot)$ and the bequest function $\mathbb{B}(\cdot)$ are homothetic of the same degree $(1-\rho)$. With this assumption, the problem can be normalized by permanent income, following Carroll (2022). Using lower case variables to denote their upper-case counterparts normalized by permanent income ($x_t = X_t/P_t$) and defining $\tilde{\Gamma}_t = \Gamma_t\psi_t$, we can write normalized versions of the previous value functions as

$$v_t^{\texttt{Out}}(m_t) = \max\{v_t^{\texttt{Stay}}(m_t),\, v_t^{\texttt{In}}(m_t - F, 0)\}, \tag{15}$$

$$v_t^{\texttt{Stay}}(m_t) = \max_{c_t} u(c_t)+\beta\delta_{t+1}\mathbb{E}_t\left[\tilde{\Gamma}_{t+1}^{1-\rho}v_{t+1}^{\texttt{Out}}(m_{t+1})\right] + \not\delta_{t+1}\mathbb{B}(a_t)$$

Subject to:

$$0 \leq c_t \leq m_t \qquad , \tag{16}$$

$$a_t = m_t - c_t$$

$$m_{t+1} = \frac{R}{\tilde{\Gamma}_{t+1}}a_t + \theta_{t+1}$$

and

$$v_t^{\texttt{In}}(m_t, n_t) = \max_{c_t, d_t} u(c_t)+\beta\delta_{t+1}\mathbb{E}_t\left[\tilde{\Gamma}_{t+1}^{1-\rho}v_{t+1}^{\texttt{In}}(m_{t+1}, n_{t+1})\right] + \not\delta_{t+1}\mathbb{B}(a_t + \tilde{n}_t)$$

Subject to:

$$-n_t \leq d_t \leq m_t, \quad 0 \leq c_t \leq \tilde{m}_t$$

$$\tilde{m}_t = m_t - d_t\left(1 - 1_{[d_t \leq 0]}\tau\right) \qquad .$$

$$\tilde{n}_t = n_t + d_t$$

$$a_t = \tilde{m}_t - c_t$$

$$m_{t+1} = \frac{R}{\tilde{\Gamma}_{t+1}}a_t + \theta_{t+1}$$

$$n_{t+1} = \frac{\tilde{R}_{t+1}}{\tilde{\Gamma}_{t+1}}\tilde{n}_t$$

It can be shown that

$$V_t^{\texttt{Out}}(M_t, P_t) = P_t^{1-\rho}v_t^{\texttt{Out}}(m_t),$$

$$V_t^{\texttt{Stay}}(M_t, P_t) = P_t^{1-\rho}v_t^{\texttt{Stay}}(m_t),$$

$$V_t^{\texttt{In}}(M_t, N_t, P_t) = P_t^{1-\rho}v_t^{\texttt{In}}(m_t, n_t)$$

and that the policy functions that solve each of the problems are related through

$$\begin{aligned} C_t^{\texttt{Stay}}(M_t, P_t) &= P_t c_t^{\texttt{Stay}}(m_t) \\ C_t^{\texttt{In}}(M_t, N_t, P_t) &= P_t c_t^{\texttt{In}}(m_t, n_t) \\ D_t(M_t, N_t, P_t) &= P_t d_t(m_t, n_t). \end{aligned}$$

Therefore, I solve the normalized problem and re-scale its solutions to obtain the original problem's solutions.

C.1 Partition Into Stages

An additional insight that facilitates solving the dynamic problem of the agent who has paid the stock-market entry cost is that the two decisions that he takes in a period (rebalancing his assets and consuming) can be seen as happening sequentially. This is convenient because the sequential sub-problems are easier to solve than the multi-choice full problem.

To re-express the problem, I take the order of the decisions to be: first rebalance assets, then consume. I denote the stages at which these decisions are taken with `Reb` and `Cns`. I will use $v^{\texttt{Reb}}(\cdot)$ and $v^{\texttt{Cns}}(\cdot)$ to represent the respective *stage value functions*.

I now present each stage in detail, working backwards in time.

C.1.1 *Consumption stage,* `Cns`

The important fact to realize at this stage is that the first thing that the agent will do in period $t + 1$ is make his asset-rebalancing decision. Therefore, that is the value function about which the agent forms expectations.

The consumption stage problem is

$$v_t^{\texttt{Cns}}(\tilde{m}_t, \tilde{n}_t) = \max_{c_t} u(c_t) + \beta \delta_{t+1} \mathbb{E}_t \left[\tilde{\Gamma}_{t+1}^{1-\rho} v_{t+1}^{\texttt{Reb}}(m_{t+1}, n_{t+1}) \right] + \not{\delta}_{t+1} \mathbb{B}(a_t + \tilde{n}_t)$$

Subject to:

$$\begin{aligned} 0 \leq & c_t \leq \tilde{m}_t \\ a_t = & \tilde{m}_t - c_t \\ m_{t+1} = & \frac{R}{\tilde{\Gamma}_{t+1}} a_t + \theta_{t+1} \\ n_{t+1} = & \frac{\tilde{R}_{t+1}}{\tilde{\Gamma}_{t+1}} \tilde{n}_t \end{aligned} \tag{17}$$

C.1.2 *Rebalancing stage,* `Reb`

The first decision that an agent takes is how to reallocate his assets. His payoff is given by the subsequent consumption problem's value function, evaluated at his post-rebalancing assets.

$$v_t^{\texttt{Reb}}(m_t, n_t) = \max_{d_t} v_t^{\texttt{Cns}}(\tilde{m}_t, \tilde{n}_t)$$

Subject to:

$$\begin{aligned} -n_t \leq & d_t \leq m_t \\ \tilde{m}_t = & m_t - d_t \left(1 - 1_{[d_t \leq 0]} \tau\right) \\ \tilde{n}_t = & n_t + d_t \end{aligned} \tag{18}$$

D. First Order Conditions and Value Function Derivatives

The computational solution of the model uses the first order conditions of the optimization problems and the derivatives of the value functions defined above. This appendix writes the first order conditions and value function derivatives explicitly.

D.0.1 *Agent who is staying out of the stock market,* `Stay`

The first order condition of the maximization problem in Equation 16 is

$$u'(c_t) = \beta R \delta_{t+1} \, \mathbb{E}_t \left[\tilde{\Gamma}_{t+1}^{-\rho} \frac{\partial v_{t+1}^{\texttt{Out}}}{\partial m_{t+1}} \right] + \not{\mathcal{B}}_{t+1} \mathbb{B}'(a_t). \tag{19}$$

The condition, while necessary, is not sufficient because $v^{\texttt{Out}}(\cdot)$ is not concave. Therefore, I use the DC-EGM method (Iskhakov et al. 2017) to solve this sub-problem.

D.0.2 *Consumption stage,* `Cns`

The first order condition for an interior solution ($c < \tilde{m}$) of the consumption stage problem (Equation 17) is

$$u'(c_t) = \beta R \delta_{t+1} \mathbb{E}_t \left[\tilde{\Gamma}_{t+1}^{-\rho} \frac{\partial v_{t+1}^{\texttt{Reb}}}{\partial m_{t+1}} \right] + \not{\mathcal{B}}_{t+1} \mathbb{B}'(a_t + \tilde{n}_t) \tag{20}$$

The derivatives of the stage value function are

$$\frac{\partial v_t^{\texttt{Cns}}(\tilde{m}_t, \tilde{n}_t)}{\partial \tilde{m}_t} = u'(c_t) \tag{21}$$

$$\frac{\partial v_t^{\texttt{Cns}}(\tilde{m}_t, \tilde{n}_t)}{\partial \tilde{n}_t} = \beta \delta_{t+1} \mathbb{E}_t \left[\tilde{R}_{t+1} \tilde{\Gamma}_{t+1}^{-\rho} \frac{\partial v_{t+1}^{\texttt{Reb}}}{\partial n_{t+1}} \right] + \not{\mathcal{B}}_{t+1} \mathbb{B}'(a_t + \tilde{n}_t) \tag{22}$$

D.0.3 *Rebalancing stage,* Reb

The first order condition for a solution of the type $d \in [(-n, 0) \cup (0, m)]$ in the rebalancing stage problem (Equation 18) is

$$\left(1 - 1_{[d_t \leq 0]}\tau\right) \frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{m}_t} = \frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{n}_t}, \tag{23}$$

and a necessary condition for a solution of the type $d = 0$ is

$$(1 - \tau) \frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{m}_t} \leq \frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{n}_t} \leq \frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{m}_t} \tag{24}$$

The derivatives of the stage value function are

$$\frac{\partial v_t^{\mathtt{Reb}}(m_t, n_t)}{\partial m_t} = \max\left\{\frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{m}}, \frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{n}}\right\} \tag{25}$$

$$\frac{\partial v_t^{\mathtt{Reb}}(m_t, n_t)}{\partial n_t} = \max\left\{(1 - \tau)\frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{m}_t}, \frac{\partial v_t^{\mathtt{Cns}}}{\partial \tilde{n}_t}\right\} \tag{26}$$

Table 13: Grids and Discretizations

	Symbol	# Points	Type of grid/Discretization
Grids for $a, m, n, \tilde{m}, \tilde{n}$	$a^\#$, $m^\#$, $n^\#$, $\tilde{m}^\#$, $\tilde{n}^\#$	101	Equispaced in logs between $1e-6$ and $5e3$, with 0 added.
Perm. Inc. Shock	ψ	5	Equiprobable.
Trans. Inc. Shock	$\tilde{\theta}$	5	Equiprobable.
Risky return	$\tilde{R}$	5	Equiprobable.

E. Numerical Solution of the Life-Cycle Model.

E.1 Grids and discretizations

The solution of the model uses various discrete grids over state variables and discretizations of stochastic variables. Table 13 summarizes the grids and discretization schemes that I use for every variable and shock. The only shock discretization not addressed in Table 13 is the out-of-pocket medical expenditure shock, which I discuss in detail in Appendix B.

E.2 Transformed-space interpolation

In my solution, I treat the continuous choice variables c and d as continuous, instead of discretizing them. Because of this decision and the multiple shocks in the model, I must evaluate value functions and their derivatives on values of the state vector that are not on my grids. In these instances, I interpolate (and extrapolate) using on-grid values.

To improve my approximation of properties of the value and marginal-value functions, such as their curvature, and the fact that they approach $-\infty$ or ∞ as wealth approaches zero, I perform my interpolations and extrapolations in a "transformed" space. The trick, discussed in Carroll (2022), consists in finding a transformation $T : \mathbb{R} \to \mathbb{R}$ such that

$T(f(\cdot))$ behaves more like an affine function than $f(\cdot)$, the function that we are trying to approximate. Then, with our chosen transformation, we create an interpolator $\hat{g}(\cdot)$ for $T(f(\cdot))$. When asked to approximate $f(x)$ for some off-grid x we return $T^{-1}(\hat{g}(x))$, where $T^{-1}(\cdot)$ is the inverse of $T(\cdot)$.

I apply this trick when constructing interpolators for value functions and marginal value functions. For value functions, I use

$$T(x) = u^{-1}(x) = \left((1-\rho) \times x\right)^{\frac{1}{1-\rho}}, \qquad T^{-1}(x) = u(x) = \frac{x^{1-\rho}}{1-\rho}.$$

For marginal value functions, I use

$$T(x) = u'^{-1}(x) = x^{-\frac{1}{\rho}}, \qquad T^{-1}(x) = u'(x) = x^{-\rho}.$$

E.3 Solving the Consumption Stage, Cns

For this stage, I use the method of endogenous gridpoints (Carroll 2006) over risk-free resources at different fixed levels of risky resources.

First, for every $\tilde{n}$ in the risky-asset balances grid $\tilde{n}^{\#}$,

- Apply the endogenous gridpoint method using Equation 20 over the grid $a^{\#}$ for end-of-period risk-free assets. The result is a set of optimal consumption points on an endogenous grid of post-rebalancing risk-free assets $\tilde{m}_t^{\#-\text{endog}}(\tilde{n})$,

$$c_t^*(\tilde{m}, \tilde{n}) \text{ for } \tilde{m} \in \tilde{m}_t^{\#-\text{endog}}(\tilde{n}).$$

- Denote the endogenous $\tilde{m}$ associated with $a_t = 0$ by $\tilde{m}_0(\tilde{n})$. This is the point where the liquidity constraint stops binding. If $\tilde{m}_0(\tilde{n}) > 0$, then add $\tilde{m} = 0$, $c_t^*(0, \tilde{n}) = 0$ to the set of endogenous risk-free assets and optimal consumption points.

- Use the optimal consumption points to find v_t^{Cns}, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{m}}$, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{n}}$ at $(\tilde{m}, \tilde{n})$, for $\tilde{m} \in$

$\tilde{m}_t^{\#-\text{endog}}(\tilde{n})$ using Equations 17, 21, and 22.

The result is a set of $(\tilde{m}, \tilde{n})$ points on which we know the consumption, value, and marginal value functions. This set of points is not a rectangular grid because the values of $\tilde{m}$ are different for every $\tilde{n}$. The next step is to use the current points to obtain an approximation of the functions over a rectangular grid.

I start with an exogenous grid for post-rebalancing risk-free assets $\tilde{m}^{\#}$ which I augment by adding the points where the liquidity constraint stops binding, $\{\tilde{m}_0(\tilde{n}) : \tilde{n} \in \tilde{n}^{\#}\}$. Denote the augmented grid with $\tilde{m}^{\#+}$. For every $\tilde{n}$ in the risky-asset balances grid $\tilde{n}^{\#}$,

- Use the values of c_t^*, v_t^{Cns}, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{m}}$, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{n}}$ calculated at $\{(\tilde{m}, \tilde{n}) : \tilde{m} \in \tilde{m}_t^{\#-\text{endog}}(\tilde{n})\}$ to approximate the value of c^*, v_t^{Cns}, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{m}}$, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{n}}$ at $\{(\tilde{m}, \tilde{n}) : \tilde{m} \in \tilde{m}^{\#+}\}$ using transformed-space linear interpolation (and extrapolation).

This process yields approximations of the functions of interest on a rectangular grid, $\{(\tilde{m}, \tilde{n}) : \tilde{m} \in \tilde{m}^{\#+} \text{ and } \tilde{n} \in \tilde{n}^{\#}\}$. I use these approximations to construct bilinear transformed-space interpolators for c_t^*, v_t^{Cns}, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{m}}$, $\frac{\partial v_t^{\text{Cns}}}{\partial \tilde{n}}$.

E.4 Solving the Rebalancing Stage, Reb

In this stage, I look for the optimal deposit/withdrawal function $d_t^*(m, n)$.

I start by defining the following convenient transformation of the optimal deposit/withdrawal

$$\mathfrak{d}_t(m,n) = \begin{cases} d_t^*(m,n)/m & d_t^*(m,n) > 0 \\ d_t^*(m,n)/n & d_t^*(m,n) < 0 \\ 0 & d_t^*(m,n) = 0. \end{cases}$$

The transformation simply re-scales the deposits or withdrawals by the balance of the fund that they are coming from, so that $\mathfrak{d}_t = 1$ corresponds to moving all the risk-free balances to the risky stocks fund, and $\mathfrak{d}_t = -1$ corresponds to withdrawing all balances from the stocks fund.

I search for the optimal $\mathfrak{d}_t$ in a rectangular exogenous grid $\{(m,n): \ m \in m^{\#} \text{ and } n \in n^{\#}\}$. The search uses the first order conditions in Equations 23 and 24 and proceeds as follows.

For every (m,n) in the rectangular grid,

- Evaluate $\frac{\partial v_t^{\text{Reb}}(m,n)}{\partial \tilde{m}}$ and $\frac{\partial v_t^{\text{Reb}}(m,n)}{\partial \tilde{n}}$.
- If $(1-\tau)\frac{\partial v_t^{\text{Cns}}(m,n)}{\partial \tilde{m}} \leq \frac{\partial v_t^{\text{Cns}}(m,n)}{\partial \tilde{n}} \leq \frac{\partial v_t^{\text{Cns}}(m,n)}{\partial \tilde{m}}$, then $\mathfrak{d}_t(m,n) = 0$.
- If $\frac{\partial v_t^{\text{Cns}}(m,n)}{\partial \tilde{n}} < (1-\tau)\frac{\partial v_t^{\text{Cns}}(m,n)}{\partial \tilde{m}}$
  - We know that the solution involves withdrawing funds, $\mathfrak{d}_t(m,n) < 0$. We have to check the corner solution $\mathfrak{d}_t(m,n) = -1$.
  - If
  $$\frac{\partial v_t^{\text{Cns}}(m+(1-\tau)n, 0)}{\partial \tilde{n}} < (1-\tau)\frac{\partial v_t^{\text{Cns}}(m+(1-\tau)n, 0)}{\partial \tilde{m}},$$
  then set $\mathfrak{d}_t(m,n) = -1$.
  - Otherwise, use bisection search to find the $d^* \in (-1,0)$ that solves
  $$\frac{\partial v_t^{\text{Cns}}(m-(1-\tau)d^*, n+d^*)}{\partial \tilde{n}} = (1-\tau)\frac{\partial v_t^{\text{Cns}}(m-(1-\tau)d^*, n+d^*)}{\partial \tilde{m}}$$
  and set $\mathfrak{d}_t(m,n) = d^*/n$.
- If $\frac{\partial v_t^{\text{Cns}}(m,n)}{\partial \tilde{m}} < \frac{\partial v_t^{\text{Cns}}(m,n)}{\partial \tilde{n}}$
  - We know that the solution involves depositing funds, $\mathfrak{d}_t(m,n) > 0$. We also know that the corner solution $\mathfrak{d}_t(m,n) = 1$ is not optimal because it leaves the agent without funds to consume.
  - Use bisection search to find the $d^* \in (0,1)$ that solves
  $$\frac{\partial v_t^{\text{Cns}}(m-d^*, n+d^*)}{\partial \tilde{m}} = \frac{\partial v_t^{\text{Cns}}(m-d^*, n+d^*)}{\partial \tilde{n}}$$

– Set $\mathfrak{d}_t(m,n) = d^*/m$.

The result of this process is a rectangular grid of asset-combinations and their associated optimal rebalancing solutions. I use these points to construct a bilinear interpolator for $\mathfrak{d}_t(\cdot,\cdot)$. Then, I use the fact that (from Equation 18),

$$
\begin{aligned}
v_t^{\texttt{Reb}}(m_t, n_t) =& v_t^{\texttt{Cns}}(\tilde{m}_t, \tilde{n}_t)\\
&\text{Where:}\\
d_t =& \begin{cases} \mathfrak{d}_t(m,n) \times m_t, & \text{If } \mathfrak{d}_t(m,n) \geq 0 \\ \mathfrak{d}_t(m,n) \times n_t, & \text{If } \mathfrak{d}_t(m,n) < 0 \end{cases}\\
\tilde{m}_t =& m_t - d_t \left(1 - 1_{[d_t \leq 0]}\tau\right)\\
\tilde{n}_t =& n_t + d_t
\end{aligned}
$$

to calculate $v_t^{\texttt{Reb}}(\cdot,\cdot)$ and its derivatives whenever they are needed.

E.5 Solving the Problem of the Agent Staying Out, `Stay`

Agents who have not paid the one-time entry cost must decide whether to pay it at the start of every period. They do this by comparing the value of staying out and not paying, versus entering and paying, as shown in Equation 15. An agent who enters passes onto the asset-rebalancing stage, `Reb`. An agent who does not enter must choose his consumption knowing that next period he will start outside of the stocks-fund again. I use the "DC-EGM" method (Iskhakov et al. 2017) to solve this problem.

I start with an exogenous grid for end-of-period risk-free assets, $a^\#$. I apply the endogenous-gridpoint method inversion over $a^\#$ using the first-order condition in Equation 19. The result is a set of candidate endogenous consumption and beginning-of-period-assets points, associated with the exogenous end-of-period-assets points,

$$\left\{ \left(c_t^e\left(a\right), m_t^e\left(a\right) \right) : \ a \in a^{\#} \right\}.$$

As argued by Iskhakov et al. (2017), these points will not necessarily be optimal. The future discrete decision of whether to pay the cost or not makes the value function not-concave and therefore points that satisfy the first order condition are not necessarily optimal. I calculate the discounted utility associated with the points from the endogenous-gridpoint inversion,

$$v_t^e(a) = u(c_t^e(a)) + \beta \delta_{t+1} \mathbb{E}_t \left[\tilde{\Gamma}_{t+1}^{1-\rho} v_{t+1}^{\texttt{Out}}(a) \right] + \not{\delta}_{t+1} \mathbb{B}(a_t)$$
$$m_{t+1} = \frac{R}{\tilde{\Gamma}_{t+1}} a + \theta_{t+1}.$$

Then, I apply the upper-envelope algorithm in Iskhakov et al. (2017) to the candidate points $\left\{ \left(c_t^e\left(a\right), m_t^e\left(a\right), v_t^e\left(a\right) \right) : \ a \in a^{\#} \right\}$ to eliminate non-optimal points. I add the "kink points" of the value function to the grid. The result is a set of optimal consumption and value points over a refined endogenous grid for start-of-period assets $m_t^{*\#}$, $\left\{ \left(c_t^*\left(m\right), v_t^*\left(m\right) \right) : \ m \in m^{*\#} \right\}$. I use these points to create linear transformed-space interpolators for $v_t^{\texttt{Stay}}$, $c_t^{\texttt{Stay}}$, and $\frac{\partial v_t^{\texttt{Stay}}}{\partial m}$.

F. Surrogate Bootstrap

The estimated parameters of the structural life cycle models are functions of the targeted moments. Uncertainty about these moments generates uncertainty about the best-fitting values of the parameters. I quantify this uncertainty by calculating the targeted moments on bootstrapped samples, estimating an approximate—"surrogate"—model on each set of moments and presenting various summary statistics of the resulting distribution of estimated parameters.

The bootstrapped targeted moments come from education-specific re-samplings of the SCF analytical sample defined in Section 4.1. I divide the sample into high-school and college graduates, and draw 500 bootstrapped samples for each level of education using the re-scaled survey weights. Then, I calculate the targeted moments on each of the bootstrapped samples. This results in 500 vectors of 45 targeted moments each for both levels of education, $\{m_{b,k}^{\text{HS}}\}_{k=0}^{500}$ and $\{m_{b,k}^{\text{College}}\}_{k=0}^{500}$.

To find the estimated parameters that would be result from each vector of moments, I use accurate surrogate models that approximate the relationship between parameters and moments embedded in the true structural models. Estimating the structural models 500 times for each level of education would come at a high computational arising mainly from their solution and simulation at each candidate vector of parameters. Recent studies like Chen, Didisheim, and Scheidegger (2021) and Catherine et al. (2022) show that these costly evaluations can be avoided using accurate approximations of the structural model that can be constructed using known parameter-moments pairs. These approximations are know as "surrogate models." Denoting with $\mathbb{\Theta}$ the space of admissible parameter values and with $\mathbb{M}$ the set of possible values for targeted moments, a model is a function $f : \mathbb{\Theta} \rightarrow \mathbb{M}$ and a surrogate model is a different function $\hat{f} : \mathbb{\Theta} \rightarrow \mathbb{M}$ that approximates the true model f, but which is ideally much faster to evaluate.

The surrogate models that I use to approximate the true structural models are deep

Table 14: Root-Mean-Squared-Errors of Surrogate Models over Targeted Moments

	High-School		College	
Sample	F.I.R.E.	Est. Beliefs	F.I.R.E.	Est. Beliefs
Training	5×10^{-3}	2×10^{-3}	4×10^{-3}	2×10^{-3}
Validation	1×10^{-2}	9×10^{-2}	1×10^{-2}	3×10^{-3}

neural networks. The networks have 3 inputs (the parameters $\{\rho, \beta, F\}$) and 45 outputs (the targeted moments), and 4 hidden layers with 192 neurons each. I use sigmoid-linear-unit "SiLU" activation functions for the hidden layers; for the output layer, I use "softplus" functions for positive moments (like the wealth ratio) and sigmoid functions for moments that are shares (like conditional stock-shares and participation rates). I use a different network for each combination of educational attainment and specification of beliefs (high school, college, and F.I.R.E., Est. Beliefs).

As suggested by Catherine et al. (2022), I train and validate the surrogate models using the parameter-moment points that I evaluate when estimating the true models. The optimization routine outlined in Section 4.2 evaluates each structural model at at least 2,500 points of the space of admissible parameter values Θ. The initial 2,500 points come from a Sobol sequence that covers Θ well. The local optimization runs of the `TikTak` algorithm generate additional evaluations, which concentrate around the best-fitting parameter values. I save the parameter-moments pair of every one of these evaluations for each specification of the model. Then, I randomly split the points into 90% training and 10% validation samples. I train the deep networks using the "Adam" algorithm (Kingma and Ba 2017) to minimize the root-mean-squared-error (RMSE) over the targeted moments. Table 14 presents the RMSEs of every surrogate model at the end of estimation, confirming that they do a good job of approximating the predictions of the true models both in- and out-of-sample.

For each vector of bootstrapped moments, I find the input parameters of its correspond-

Table 15: Different Percentiles of Bootstrapped Estimates

	Percentiles						
	F.I.R.E.			Est. Belifs			Difference
Parameter	P_5	P_{50}	P_{95}	P_5	P_{50}	P_{95}	Frac. > 0
High School							
CRRA (ρ)	8.51	8.56	8.61	4.20	4.22	4.25	1
Disc. Fac (β)	0.31	0.33	0.35	0.75	0.76	0.77	0
Entry Cost ($F \times 100$)	2.87	3.11	3.42	2.33	2.49	2.64	1
MSM Loss	15.45	16.01	16.69	3.69	4.04	4.44	1
College							
CRRA (ρ)	11.36	11.48	11.55	5.07	5.10	5.14	1
Disc. Fac (β)	0.62	0.63	0.64	0.88	0.89	0.89	0
Entry Cost ($F \times 100$)	0.51	0.57	1.75	0.00	0.00	0.00	1
MSM Loss	4.60	5.33	6.13	2.27	2.94	3.73	1

ing surrogate model that minimizes the SMM loss function using the same optimization routine that I use for the main estimates. The $k-th$ set of bootstrapped parameter estimates for the pair of education and beliefs-specification $(e, \mathbf{b})$ is

$$\hat{\vartheta}_k^{e,\mathbf{b}} = \arg\min_{\vartheta} \left(m_{b,k}^e - \hat{f}^{e,\mathbf{b}}(\vartheta)\right)' \mathbb{W}^e \left(m_{b,k}^e - \hat{f}^{e,\mathbf{b}}(\vartheta)\right), \tag{27}$$

where $\hat{f}^{e,\mathbf{b}}$ is the surrogate model for education level e and belief specification $\mathbb{b}$. The results are sets of bootstrapped estimates

$$\{\hat{\vartheta}_{b,k}^{\text{HS,F.I.R.E.}}\}_{k=0}^{500}, \{\hat{\vartheta}_{b,k}^{\text{College, F.I.R.E.}}\}_{k=0}^{500}, \text{ and } \{\hat{\vartheta}_{b,k}^{\text{HS, Est. Beliefs}}\}_{k=0}^{500}, \{\hat{\vartheta}_{b,k}^{\text{College, Est. Beliefs}}\}_{k=0}^{500}.$$

I additionally store and present the values of the loss functions associated with each parameter estimate.

Table 15 presents the 5th, 50th, and 95th percentiles of every parameter for every combination of educational attainment and belief specification, in addition to the minimized loss function "MSM Loss." The table shows that all the parameter estimates and attained losses are tightly distributed around the main estimates reported in Table 8. For each vector of bootstrapped moments $m^e_{b,k}$, I find the difference between the parameter estimates and attained losses under the F.I.R.E. and "Est. Beliefs" specifications. The last column in Table 15 presents the fraction of moment vectors for which this difference is positive. This column shows the robustness of the conclusions that, for both levels of education, models that use the estimated beliefs improve upon the fit of F.I.R.E. models and do so with lower levels of relative risk aversion, higher discount factors, and lower entry costs. These conclusions are true for each of the 500 bootstrapped vectors of moments for each level of education.

The 95% confidence intervals presented in Table 8 correspond to the 2.5-th and 97.5-th percentiles of the bootstrapped values of each parameter, for each model specification.

G. The expected welfare function

The welfare calculations presented in Section 6 rely on the calculation of individuals' objectively expected welfare, which can differ from their subjective expectations due to their misspecified beliefs. This section defines my measure of expected welfare and discusses how I calculate it.

For a given set of beliefs about risky returns denoted with $\mathcal{B} = (\mu, \sigma)$ and other parameters, I solve the life cycle model and its components described in Appendix C. Denote the resulting policy functions for every age t with $C_t^{\texttt{Stay}}(\cdot; \mathcal{B})$, $D_t(\cdot; \mathcal{B})$, and $C_t^{\texttt{In}}(\cdot; \mathcal{B})$. I calculate functions $\mathfrak{V}_t(\cdot)$ that allow me to find the expected lifetime welfare that an objective observer would expect an agent to derive from his remaining years of life if he behaved according to the policy functions associated with his beliefs $\mathcal{B}$. These functions are

$$\mathfrak{V}_t^{\texttt{Stay}}(M_t, P_t; \mathcal{B}) = u(C_t) + \beta\delta_{t+1}\mathbb{E}_t\left[\mathfrak{V}_{t+1}^{\texttt{Out}}(M_{t+1}, P_{t+1}; \mathcal{B})\right] + \not{\delta}_{t+1}\mathbb{B}(A_t)$$

$$\text{Where:} \tag{28}$$

$$C_t = C_t^{\texttt{Stay}}(M_t, P_t; \mathcal{B})$$

$$A_t = M_t - C_t, \quad M_{t+1} = RA_t + Y_{t+1}, \quad P_{t+1} = \Gamma_{t+1}\psi_{t+1}P_t, \quad Y_{t+1} = \theta_{t+1}P_{t+1}$$

for agents who have not paid the risky-asset entry cost and decide to not pay it in t, and

$$\mathfrak{V}_t^{\texttt{In}}(M_t, N_t, P_t; \mathcal{B}) = u(C_t) + \beta\delta_{t+1}\mathbb{E}_t\left[\mathfrak{V}_{t+1}^{\texttt{In}}(M_{t+1}, N_{t+1}, P_{t+1}; \mathcal{B})\right]$$
$$+\not\delta_{t+1}\mathbb{B}(A_t + N_t + D_t)$$

$$\text{Where:} \tag{29}$$

$$D_t = D_t(M_t, N_t, P_t; \mathcal{B}), \quad C_t = C_t^{\texttt{In}}(M_t, N_t, P_t; \mathcal{B})$$
$$A_t = M_t - D_t\left(1 - 1_{[D_t \leq 0]}\tau\right) - C_t, \quad M_{t+1} = RA_t + Y_{t+1}$$
$$N_{t+1} = \tilde{R}_{t+1} \times (N_t + D_t), \quad P_{t+1} = \Gamma_{t+1}\psi_{t+1}P_t, \quad Y_{t+1} = \theta_{t+1}P_{t+1}$$

for agents who have already paid the entry cost. Equations 28 and 29 differ from the value functions defined in Appendix C because the expectations are taken using the true distribution of risky asset returns. Numerically, I construct interpolators for these functions iterating backwards, using the solved policy functions and the same grids and discretizations described in Appendix E.

The simplified notation that I use for $\mathfrak{V}_t(\cdot)$ in the main text is corresponds to

$$\mathfrak{V}_t(P_t, M_t, N_t = 0, \texttt{Paid}_t = 0; \mathcal{B}) \equiv \max\left\{\mathfrak{V}_t^{\texttt{Stay}}(M_t, P_t; \mathcal{B}), \mathfrak{V}_t^{\texttt{In}}(M_t - F \times P_t, 0, P_t; \mathcal{B})\right\}$$

$$\mathfrak{V}_t(P_t, M_t, N_t, \texttt{Paid}_t = 1; \mathcal{B}) \equiv \mathfrak{V}_t^{\texttt{In}}(M_t, N_t, P_t; \mathcal{B}).$$

References

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson (2009). "The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation". In: *Brookings Papers on Economic Activity* 2009.2, pp. 51–117.

Almås, Ingvild, Orazio P. Attanasio, and Pamela Jervis (Jan. 2023). *Economics and Measurement: New Measures to Model Decision Making*. URL: https://www.nber.org/papers/w30839 (visited on 02/17/2023). Pre-published.

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti (June 2020). "Long-Term-Care Utility and Late-in-Life Saving". In: *Journal of Political Economy* 128.6, pp. 2375–2451.

Ameriks, John, Gábor Kézdi, Minjoon Lee, and Matthew D. Shapiro (2020). "Heterogeneity in Expectations, Risk Tolerance, and Household Stock Shares: The Attenuation Puzzle". In: *Journal of Business & Economic Statistics* 38.3, pp. 633–646.

Amromin, Gene and Steven A. Sharpe (Apr. 2014). "From the Horse's Mouth: Economic Conditions and Investor Expectations of Risk and Return". In: *Management Science* 60.4, pp. 845–866.

Andersen, Steffen and Kasper Meisner Nielsen (May 1, 2011). "Participation Constraints in the Stock Market: Evidence from Unexpected Inheritance Due to Sudden Death". In: *The Review of Financial Studies* 24.5, pp. 1667–1697.

Arcidiacono, Peter, V. Joseph Hotz, Arnaud Maurel, and Teresa Romano (Dec. 2020). "Ex Ante Returns and Occupational Choice". In: *Journal of Political Economy* 128.12, pp. 4475–4522.

Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg (Oct. 7, 2019). *Benchmarking Global Optimizers*. w26340. National Bureau of Economic Research.

Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber (1999). “Humps and Bumps in Lifetime Consumption”. In: *Journal of Business & Economic Statistics* 17.1, pp. 22–35.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel (Dec. 2018). “The Economic Effects of Social Networks: Evidence from the Housing Market”. In: *Journal of Political Economy* 126.6, pp. 2224–2276.

Bateman, Hazel, Christine Eckert, Fedor Iskhakov, Jordan Louviere, Stephen Satchell, and Susan Thorp (Feb. 1, 2017). “Default and Naive Diversification Heuristics in Annuity Choice”. In: *Australian Journal of Management* 42.1, pp. 32–57.

Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo (Jan. 1, 2024). “Diagnostic Business Cycles”. In: *The Review of Economic Studies* 91.1, pp. 129–162.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, Matthew OBrien, and Andrei Shleifer (Apr. 2024). “Long-Term Expectations and Aggregate Fluctuations”. In: *NBER Macroeconomics Annual* 38, pp. 311–347.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2019). “Diagnostic Expectations and Stock Returns”. In: *The Journal of Finance* 74.6, pp. 2839–2874.

— (May 2024). “Belief Overreaction and Stock Market Puzzles”. In: *Journal of Political Economy* 132.5, pp. 1450–1484.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (Aug. 2022). “Overreaction and Diagnostic Expectations in Macroeconomics”. In: *Journal of Economic Perspectives* 36.3, pp. 223–244.

Briggs, Joseph, David Cesarini, Erik Lindqvist, and Robert Östling (Jan. 1, 2021). “Windfall Gains and Stock Market Participation”. In: *Journal of Financial Economics* 139.1, pp. 57–83.

Cagetti, Marco (July 1, 2003). “Wealth Accumulation Over the Life Cycle and Precautionary Savings”. In: *Journal of Business & Economic Statistics* 21.3, pp. 339–353.

Calvet, Laurent E., John Y. Campbell, Francisco Gomes, and Paolo Sodini (May 2021). *The Cross-Section of Household Preferences*. URL: https://www.nber.org/papers/w28788 (visited on 10/07/2022). Pre-published.

Calvo-Pardo, Hector, Xisco Oliver, and Luc Arrondel (Jan. 2022). "Subjective Return Expectations, Perceptions, and Portfolio Choice". In: *Journal of Risk and Financial Management* 15.1 (1), p. 6.

Campanale, Claudio (Apr. 1, 2011). "Learning, Ambiguity and Life-Cycle Portfolio Allocation". In: *Review of Economic Dynamics* 14.2, pp. 339–367.

Campanale, Claudio, Carolina Fugazza, and Francisco Gomes (Apr. 1, 2015). "Life-Cycle Portfolio Choice with Liquid and Illiquid Financial Assets". In: *Journal of Monetary Economics* 71, pp. 67–83.

Caplin, Andrew (Oct. 2021). *Economic Data Engineering*. URL: https://www.nber.org/papers/w29378 (visited on 02/17/2023). Pre-published.

Caplin, Andrew, Victoria Gregory, Eungik Lee, Søren Leth-Petersen, and Johan Sæverud (Mar. 2023). *Subjective Earnings Risk*. URL: https://www.nber.org/papers/w31019 (visited on 04/05/2023). Pre-published.

Carroll, Christopher D. (1994). "How Does Future Income Affect Current Consumption?" In: *The Quarterly Journal of Economics* 109.1, pp. 111–147.

— (Feb. 1, 1997). "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*". In: *The Quarterly Journal of Economics* 112.1, pp. 1–55.

— (June 1, 2006). "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems". In: *Economics Letters* 91.3, pp. 312–320.

— (2022). "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems". In: p. 58.

Carroll, Christopher D. and Andrew A. Samwick (Sept. 1, 1997). "The Nature of Precautionary Wealth". In: *Journal of Monetary Economics* 40.1, pp. 41–71.

Carroll, Christopher D., Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White (2017). "The Distribution of Wealth and the Marginal Propensity to Consume". In: *Quantitative Economics* 8.3, pp. 977–1020.

Cartis, Coralia, Jan Fiala, Benjamin Marteau, and Lindon Roberts (Aug. 8, 2019). "Improving the Flexibility and Robustness of Model-based Derivative-free Optimization Solvers". In: *ACM Transactions on Mathematical Software* 45.3, 32:1–32:41.

Catherine, Sylvain (Dec. 24, 2021). "Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle". In: *The Review of Financial Studies*, hhab136.

Catherine, Sylvain, Mehran Ebrahimian, David Alexandre Sraer, and David Thesmar (Sept. 1, 2022). *Robustness Checks in Structural Analysis*. URL: https://papers.ssrn.com/abstract=4206149 (visited on 06/09/2023). Pre-published.

Chang, Yongsung, Jay H. Hong, and Marios Karabarbounis (Apr. 2018). "Labor Market Uncertainty and Portfolio Choice Puzzles". In: *American Economic Journal: Macroeconomics* 10.2, pp. 222–262.

Chen, Hui, Antoine Didisheim, and Simon Scheidegger (Feb. 9, 2021). *Deep Surrogates for Finance: With an Application to Option Pricing*. URL: https://papers.ssrn.com/abstract=3782722 (visited on 06/19/2023). Pre-published.

Clark, Robert L. (Apr. 28, 2023). "Effectiveness of Employer-Provided Financial Education Programs". In: *Journal of Financial Literacy and Wellbeing*, pp. 1–15.

Cocco, João F., Francisco Gomes, and Pascal J. Maenhout (July 1, 2005). "Consumption and Portfolio Choice over the Life Cycle". In: *The Review of Financial Studies* 18.2, pp. 491–533.

Coibion, Olivier and Yuriy Gorodnichenko (Aug. 2015). "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts". In: *American Economic Review* 105.8, pp. 2644–2678.

Das, Sreyoshi, Camelia M Kuhnen, and Stefan Nagel (Jan. 1, 2020). "Socioeconomic Status and Macroeconomic Expectations". In: *The Review of Financial Studies* 33.1, pp. 395–432.

De Nardi, Mariacristina, Eric French, and John B. Jones (Feb. 2010). "Why Do the Elderly Save? The Role of Medical Expenses". In: *Journal of Political Economy* 118.1, pp. 39–75.

Dominitz, Jeff and Charles F. Manski (1997). "Using Expectations Data To Study Subjective Income Expectations". In: *Journal of the American Statistical Association* 92.439, pp. 855–867.

— (May 1, 2007). "Expected Equity Returns and Portfolio Choice: Evidence from the Health and Retirement Study". In: *Journal of the European Economic Association* 5.2-3, pp. 369–379.

— (2011). "Measuring and Interpreting Expectations of Equity Returns". In: *Journal of Applied Econometrics* 26.3, pp. 352–370.

Drerup, Tilman, Benjamin Enke, and Hans-Martin von Gaudecker (Oct. 1, 2017). "The Precision of Subjective Data and the Explanatory Power of Economic Models". In: *Journal of Econometrics*. Measurement Error Models 200.2, pp. 378–389.

Druedahl, Jeppe (2021). "A Guide on Solving Non-Convex Consumption-Saving Models". In: *Computational Economics* 58.3, pp. 747–775.

Enke, Benjamin, Thomas Graeber, Ryan Oprea, and Jeffrey Yang (Sept. 2024). *Behavioral Attenuation*. URL: https://www.nber.org/papers/w32973 (visited on 11/13/2024). Pre-published.

Erdem, Tülin, Michael P. Keane, T. Sabri Öncü, and Judi Strebel (Sept. 1, 2005). "Learning About Computers: An Analysis of Information Search and Technology Choice". In: *Quantitative Marketing and Economics* 3.3, pp. 207–247.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso (2017). "Asset Market Participation and Portfolio Choice over the Life-Cycle". In: *The Journal of Finance* 72.2, pp. 705–750.

Foltyn, Richard (Feb. 24, 2020). *Experience-Based Learning, Stock Market Participation and Portfolio Choice*. URL: https://papers.ssrn.com/abstract=3543442 (visited on 06/26/2022). Pre-published.

French, Eric and John B. Jones (2004). "On the Distribution and Dynamics of Health Care Costs". In: *Journal of Applied Econometrics* 19.6, pp. 705–721.

Gabler, Janos (2022). *A Python Tool for the Estimation of Large Scale Scientific Models.*

Ganong, Peter, Pascal Noel, and Joseph Vavra (Nov. 1, 2020). "US Unemployment Insurance Replacement Rates during the Pandemic". In: *Journal of Public Economics* 191, p. 104273.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus (May 2021). "Five Facts about Beliefs and Portfolios". In: *American Economic Review* 111.5, pp. 1481–1522.

Giustinelli, Pamela, Charles F. Manski, and Francesca Molinari (Nov. 1, 2022). "Tail and Center Rounding of Probabilistic Expectations in the Health and Retirement Study". In: *Journal of Econometrics*. Annals Issue: Subjective Expectations & Probabilities in Economics. 231.1, pp. 265–281.

Gomes, Francisco and Alexander Michaelides (Oct. 1, 2003). "Portfolio Choice with Internal Habit Formation: A Life-Cycle Model with Uninsurable Labor Income Risk". In: *Review of Economic Dynamics*. Finance and the Macroeconomy 6.4, pp. 729–766.

— (2005). "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence". In: *The Journal of Finance* 60.2, pp. 869–904.

Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002). "Consumption Over the Life Cycle". In: *Econometrica* 70.1, pp. 47–89.

Greenwood, Robin and Andrei Shleifer (Mar. 1, 2014). "Expectations of Returns and Expected Returns". In: *The Review of Financial Studies* 27.3, pp. 714–746.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese (Nov. 1, 1992). "Earnings Uncertainty and Precautionary Saving". In: *Journal of Monetary Economics* 30.2, pp. 307–337.

Haliassos, Michael and Carol C. Bertaut (1995). “Why Do so Few Hold Stocks?” In: *The Economic Journal* 105.432, pp. 1110–1129.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (Apr. 1995). “Precautionary Saving and Social Insurance”. In: *Journal of Political Economy* 103.2, pp. 360–399.

Hurd, Michael (2009). “Subjective Probabilities in Household Surveys”. In: *Annual Review of Economics* 1.1, pp. 543–562.

Hurd, Michael, Maarten Van Rooij, and Joachim Winter (2011). “Stock Market Expectations of Dutch Households”. In: *Journal of Applied Econometrics* 26.3, pp. 416–436.

Iskhakov, Fedor, Thomas H. Jørgensen, John Rust, and Bertel Schjerning (2017). “The Endogenous Grid Method for Discrete-Continuous Dynamic Choice Models with (or without) Taste Shocks”. In: *Quantitative Economics* 8.2, pp. 317–365.

Kaiser, Tim, Annamaria Lusardi, Lukas Menkhoff, and Carly Urban (Aug. 1, 2022). “Financial Education Affects Financial Knowledge and Downstream Behaviors”. In: *Journal of Financial Economics* 145 (2, Part A), pp. 255–272.

Kaiser, Tim and Lukas Menkhoff (Oct. 1, 2017). “Does Financial Education Impact Financial Literacy and Financial Behavior, and If So, When?” In: *The World Bank Economic Review* 31.3, pp. 611–630.

Kézdi, Gábor and Robert J. Willis (Nov. 2011). *Household Stock Market Beliefs and Learning*. Working Paper 17614. National Bureau of Economic Research.

Khorunzhina, Natalia (Dec. 1, 2013). “Structural Estimation of Stock Market Participation Costs”. In: *Journal of Economic Dynamics and Control* 37.12, pp. 2928–2942.

Kingma, Diederik P. and Jimmy Ba (Jan. 29, 2017). *Adam: A Method for Stochastic Optimization*. URL: http://arxiv.org/abs/1412.6980 (visited on 07/03/2023). Pre-published.

Kopecky, Karen A. and Tatyana Koreshkova (July 2014). “The Impact of Medical and Nursing Home Expenses on Savings”. In: *American Economic Journal: Macroeconomics* 6.3, pp. 29–72.

Koşar, Gizem and Cormac O’Dea (Jan. 1, 2023). “Chapter 21 - Expectations Data in Structural Microeconomic Models”. In: *Handbook of Economic Expectations*. Ed. by Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw. Academic Press, pp. 647–675.

Lusardi, Annamaria (Dec. 19, 1997). “Precautionary Saving and Subjective Earnings Variance”. In: *Economics Letters* 57.3, pp. 319–326.

— (1998). “On the Importance of the Precautionary Saving Motive”. In: *The American Economic Review* 88.2, pp. 449–453.

Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell (Apr. 2017). “Optimal Financial Knowledge and Wealth Inequality”. In: *Journal of Political Economy* 125.2, pp. 431–477.

— (Oct. 1, 2020). “Assessing the Impact of Financial Education Programs: A Quantitative Model”. In: *Economics of Education Review* 78, p. 101899.

Lusardi, Annamaria and Olivia S. Mitchell (Apr. 2023). *The Importance of Financial Literacy: Opening a New Field*. URL: https://www.nber.org/papers/w31145 (visited on 04/17/2023). Pre-published.

Malmendier, Ulrike and Stefan Nagel (Feb. 1, 2011). “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?*”. In: *The Quarterly Journal of Economics* 126.1, pp. 373–416.

— (Feb. 1, 2016). “Learning from Inflation Experiences *”. In: *The Quarterly Journal of Economics* 131.1, pp. 53–87.

Manski, Charles F. (Apr. 2018). “Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise”. In: *NBER Macroeconomics Annual* 32, pp. 411–471.

Manski, Charles F. and Francesca Molinari (Apr. 1, 2010). “Rounding Probabilistic Expectations in Surveys”. In: *Journal of Business & Economic Statistics* 28.2, pp. 219–231.

Mehra, Rajnish and Edward C. Prescott (Mar. 1, 1985). “The Equity Premium: A Puzzle”. In: *Journal of Monetary Economics* 15.2, pp. 145–161.

Merton, Robert C. (1969). “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case”. In: *The Review of Economics and Statistics* 51.3, pp. 247–257.

Nagel, Stefan and Zhengyang Xu (May 1, 2022). “Asset Pricing with Fading Memory”. In: *The Review of Financial Studies* 35.5, pp. 2190–2245.

Parker, Jonathan A., Antoinette Schoar, Allison T. Cole, and Duncan Simester (Mar. 2022). *Household Portfolios and Retirement Saving over the Life Cycle*. Working Paper 29881. National Bureau of Economic Research.

Peijnenburg, Kim (Oct. 2018). “Life-Cycle Asset Allocation with Ambiguity Aversion and Learning”. In: *Journal of Financial and Quantitative Analysis* 53.5, pp. 1963–1994.

Samuelson, Paul A. (1969). “Lifetime Portfolio Selection By Dynamic Stochastic Programming”. In: *The Review of Economics and Statistics* 51.3, pp. 239–246.

Shiller, Robert J. (June 7, 1990). *Market Volatility*. Cambridge, MA, USA: MIT Press. 480 pp.

Viceira, Luis M. (2001). “Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income”. In: *The Journal of Finance* 56.2, pp. 433–470.

Wachter, Jessica A. and Motohiro Yogo (Nov. 1, 2010). “Why Do Household Portfolio Shares Rise in Wealth?” In: *The Review of Financial Studies* 23.11, pp. 3929–3965.

Wiswall, Matthew and Basit Zafar (May 2021). “Human Capital Investments and Expectations about Career and Family”. In: *Journal of Political Economy* 129.5, pp. 1361–1424.