Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Stock Market Investment: The Role of Human Capital

Kartik Athreya, Felicia Ionescu, and Urvi Neelakantan

2015-065

Please cite this paper as:
Athreya, Kartik, Felicia Ionescu, and Urvi Neelakantan (2015). "Stock Market Investment: The Role of Human Capital," Finance and Economics Discussion Series 2015-065. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2015.065r1.

Stock Market Participation: The Role of Human Capital*

Kartik Athreya † FRB Richmond

Felicia Ionescu ‡ Federal Reserve Board

Urvi Neelakantan § FRB Richmond

June 20, 2016

Abstract

Participation in the stock market is limited, especially early in life. By contrast, human capital investment is widespread, especially early in life. Returns to equity are constant across households, while returns to human capital vary. The contribution of this paper is to demonstrate that once human capital investment is allowed for and, critically, disciplined to match observed dispersion in earnings, an entirely standard model of portfolio choice delivers stock market participation rates consistent with the data over the entire life cycle. Moreover, we show that endogenizing human capital strongly alters the role of borrowing costs in limiting stock market participation.

JEL Codes: E21; G11; J24;

Keywords: Financial Portfolios; Human Capital Investment; Life-cycle

*We thank Nika Lazaryan for excellent research assistance on this project. We are grateful to Marco Cagetti, Marios Karabarbounis, Stephen Zeldes, and seminar and conference participants at the Allied Social Science Association Meeting, The Central Bank of Hungary, Computing in Economics and Finance, Econometric Society, European Central Bank Conference on Household Finance, Federal Reserve Bank of Richmond and University of Virginia Research Jamboree, Federal Reserve Board Macro Workshop, Iowa State University, Midwest Macro meetings, Society for Economic Dynamics, the 2016 FRB-St.Louis-Tsinghua University Conference, and Stony Brook University for helpful comments and suggestions. We thank especially Michael Haliassos and Yi Wen for their detailed input. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or the Federal Reserve System. All errors are ours.

†Federal Reserve Bank of Richmond, P.O. Box 27622, Richmond, VA 23261, Kartik.Athreya@rich.frb.org, Ph:804-697-8225

‡Board of Governors of the Federal Reserve System, 2050 C St. Washington, D.C. 20551, felicia.ionescu@frb.gov, Ph: 202-452-2504

§Federal Reserve Bank of Richmond, P.O. Box 27622, Richmond, VA 23261, Urvi.Neelakantan@rich.frb.org, Ph:804-697-8913

1 Introduction

Household participation in the stock market is limited, especially early in life, despite its high rate of return. The contribution of this paper is to evaluate the role of human capital *investment* for the path of stock-market participation. We are motivated to consider this channel by the observation that these two forms of investment may compete: human capital investment is highest early in life, when stock market participation is low, and low late in life, when stock market participation is high. We demonstrate once human capital investment is allowed for and disciplined to match empirical labor income dispersion, an entirely standard model of portfolio choice closely predicts the observed life-cycle path of stock market participation rates. To our knowledge, our work is the first to examine the role played by the ability of households to accumulate human capital—and especially the effect of variation in this ability—for stock market participation over the life cycle.

Given the suggestive evidence that human capital and stock market equity *do* compete for households' resources, we can ask *why* they might. That is, why might the ability of households to make human capital investment decisions matter for the financial portfolios they choose over the life cycle? The answer lies in the manner in which human capital investment pins down the path of earnings over the life cycle, and in the costs associated with accumulating skill. Human capital is an attractive investment early in life, especially for those with high learning ability or low initial human capital: the opportunity cost of spending time learning—forgoing earnings—is relatively low, the marginal return to learning is high, and the horizon over which to recoup any payoff from learning is long. While early human capital investment thus provides the means to maximize the present value of lifetime resources, it also ensures growth of earnings over time. Forward-looking individuals who invest in human capital will therefore seek to smooth consumption in early life by either saving less than they otherwise would or by borrowing if anticipated income growth is high enough. This means that households who invest in human capital early in life will desire, absent risk, to avoid large positive net positions in financial assets and may even seek to borrow or attain negative financial net worth.[1]

While the preceding closely connects human capital investment choices to *net* financial wealth positions, it does not guarantee that those who invest in human capital

[1] In addition, we will show that the presence of risk motivates precautionary savings and leads to nonparticipation even among some households with a small buffer stock of savings in the safe asset.

will decline a positive *gross* position in stocks, i.e., that they will decide not to participate. After all, households still care about provisioning themselves for the future, and if the rate of return on one asset class (say, stocks) is high, while that on an other (say, bonds) is low, then borrowing that asset (shorting the bond) and investing the proceeds in the other (going long on stocks) may well be a sensible strategy. Absent significant borrowing costs, households with an expected upward-sloping profile of earnings will borrow (short the risk-free asset) to invest in stocks at substantial rates throughout the life cycle.

However, once we allow for the fact that human capital must be accumulated by training, the household's calculation with respect to stock participation changes. Given the decision to borrow, the household must decide whether to spend the proceeds on current consumption while investing time in human capital acquisition or to finance stock purchases. As long as the marginal return to investing in human capital (the expected increase in future earnings) is higher than the expected return from investing in stocks, the household will choose to do the former. As a result, a young investor facing high marginal returns to investment in human capital will not find the strategy of borrowing to hold stocks attractive. If they borrow, it will instead be to finance consumption. Notice that for agents for whom the marginal return to human capital exceeds the return on stocks, any wedge driving borrowing costs further above the risk-free rate has no relevance.

When agents are, instead, implicitly endowed with human capital (as is the case whenever earnings processes are modeled as exogenous), increasing future earnings through human capital investment is not an option. In this instance, the agent must only decide whether borrowing to invest in stocks makes sense at the margin. As a result, borrowing costs regain the power to prevent high rates of stock participation among the young.

Thus, once human capital is not given, but must instead be acquired, participation falls irrespective of borrowing opportunities. This is an important finding, as it sheds light on whether households are deprived by credit constraints of access to lucrative financial assets or simply choose not to invest in them because they are instead engaged in human capital accumulation.

While intuitively appealing, there is no a priori guarantee that human capital investment is capable of generating a *quantitatively* plausible account of observed stock market participation behavior. A principal contribution of our paper is to demonstrate that it is. We will show that a standard human capital model, disciplined to match *heterogeneity* in both the levels and slopes of life cycle earnings, can account well for stock market participation at all ages and, especially, limited participation early in life.

Critical to our findings is the fact that households differ in ways that make total and marginal returns to human capital agent-specific. In our model, heterogeneity in the marginal return to human capital arises from differences in agents' learning ability and endowment of initial human capital, as well as from the idiosyncratic dispersion in wages, and hence, in the opportunity cost of human capital investment. Households for whom the expected returns to human capital investment are relatively low choose to spend less time learning and more time earning; their earnings profiles are flatter, and they save and enter the stock market early. On the other hand, households with endowments that lead them to expect high returns to human capital investment spend more time learning, face a steeper earnings profile, and do not enter the stock market when young.

We proceed by embedding the classic Ben-Porath (1967) model of time allocation between working ("earning") and human-capital accumulation ("learning") into a parameterized heterogenous agent life-cycle consumption-savings model with uninsurable idiosyncratic labor income risk and financial portfolio choice. In this setting, agents differ in their learning ability as well as in their initial endowments of human capital levels and wealth.

In addition to its predictions for stock-market participation, our model's implications for household wealth levels—both total wealth and the levels invested in risky and risk-free assets—are broadly consistent with the data. We also find that those who do invest in the risky asset do not allocate all their wealth to it even when young. This makes clear that the portfolio choices we derive describe empirically relevant magnitudes for the size and division of cash flows that the household receives. These surprising successes along dimensions not targeted suggest that human capital investment likely plays an important part in driving household financial investment over the life cycle.

2 Related Literature

Our work can be seen as building on the insights of a large body of work as we discuss below. In addition, our modeling approach most closely follows four papers—Davis, Kubler, and Willen (2006), Roussanov (2010), Huggett, Ventura, and Yaron (2011), and Kim, Maurer, and Mitchell (2016).

While our quantitative evaluation of the ability to invest in human capital for households' stock market participation is new, the more general idea that labor income matters for stock-market investment is not (see, for example, the early work of Brito, 1978). In particular, our work is informed by a set of papers that study, as we do, portfolio choice in a life-cycle setting with uninsurable, idiosyncratic labor income risk. Examples include Campbell, Cocco, Gomes, and Maenhout (2001), Gomes and Michaelides (2003), Cocco, Gomes, and Maenhout (2005), Cocco (2005), Gomes and Michaelides (2005), Davis, Kubler, and Willen (2006), Polkovnichenko (2007), and Chang, Hong, and Karabarbounis (2014).[2] These papers, building on earlier work of Jagannathan and Kocherlakota (1996), argue that it is the risk properties of labor income that are likely to influence households' investment in the stock market. Importantly, however, in the preceding work, human capital is only implicitly defined by the present value of exogenously imposed labor income processes. It does not arise, as in our model, from investment choices. Another common assumption is that participation entails a cost.[3] Several of these papers assess the role of preferences such as Epstein-Zin with heterogeneity in risk preferences (Gomes and Michaelides, 2005) or habit formation (Gomes and Michaelides, 2005; Polkovnichenko, 2007) in generating empirically plausible predictions. Along these dimensions, our work is closest to that of Davis, Kubler, and Willen (2006), who assume standard Constant Relative Risk Aversion (CRRA) preferences and abstract from stock-market participation costs. These authors demonstrate that a wedge between the borrowing rate and risk-free savings rate is capable of generating limited stock market participation. By contrast, we emphasize the role played by endogenous human capital investment in limiting participation in equity markets, even in the absence of the wedge.

Many of the papers cited above focus on the share of wealth invested in stocks (the "intensive margin") and though our focus is on participation (the "extensive margin"),

[2]Chang, Hong, and Karabarbounis (2014) represents an innovation within the class of models with exogenous human capital. They focus on understanding the share of wealth held in risky assets. Their model incorporates front-loaded risk of unemployment into a model where agents must learn about the income-generating process that they are endowed with. They show that data on shares can be interpreted as optimal behavior under a particular specification of parameters, including one regulating the speed of Bayesian learning.

[3]Haliassos and Michaelides (2003) is an example of a paper that introduces a fixed cost in an infinite horizon setting. However, once this entry cost is paid, households hold their entire financial wealth in stocks. In other words, in their setting, the empirically observed coexistence of risky and risk-free asset holdings in household portfolios remains a puzzle. For an assessment of the size of stock market participation costs, though exclusively in models that abstract from human capital, see Khorunzhina (2013) and references therein.

we also will document the model's implications for shares. This connects our work to literature starting with the classic work of Merton (1969) and Samuelson (1969). In general, the studies that have examined the implications of labor income for life-cycle portfolios concur that, in spite of labor income risk, a young investor should place much of her financial wealth in the risky asset. In these models, this implication obtains because labor income shocks are assumed to be (nearly) independent from stock-market return innovations. Thus, a young investor chooses to diversify away her human capital risk by holding a high fraction of her liquid wealth in a well-diversified portfolio of stocks.[4] However, as we show, once human capital investment is (i) allowed and (ii) is disciplined to match observed earnings dispersion, the typical household's share of financial wealth held in stock-market equity is far from 100 percent. Along this dimension, our model shares with recent work the implication that shares should be hump shaped over the life cycle (see, e.g. Benzoni, Collin-Dufresne, and Goldstein, 2007, and the references therein).

Though we are not directly concerned with providing a resolution to the equity premium puzzle, it is clear that our model has implications for this. After all, our work can be viewed as asking the question: "if one gets human capital investment 'right' (i.e., ensuring that the model generates accurate earnings over the life-cycle under observed stock and bond returns), does one get equity investments right, given observed returns?" Our model says that, at least in terms of equity market participation, the answer is largely "yes." Moreover, in terms of total savings, we show that allowing human capital generates a path of total wealth over the life cycle that is remarkably close to the data, despite not being targeted in any way. Nonetheless, we do not account completely for

[4]For example, Cocco, Gomes, and Maenhout (2005) argue that as individuals age, the present value of their labor income decreases because of the decrease in the number of remaining working years. Following the logic of Jagannathan and Kocherlakota (1996), they further argue that labor income usually acts as a substitute for holding a riskless asset and, as such, should encourage households to reduce the share of stocks in their portfolio as they age. In the same spirit, Viceira (2001) shows that the fraction of savings optimally invested in stocks is larger for employed investors than for retired investors when labor income risk is uncorrelated with stock return risk. Within the class of models with exogenous human capital, recent work measures the extent to which earnings are bond-like or stock-like and studies the implications for the share of wealth held in equities (Benzoni, Collin-Dufresne, and Goldstein, 2007; Huggett and Kaplan, 2015). Others examine the role of labor supply. For example, Gomes, Kotlikoff, and Viceira (2008) endogenize the labor supply decision, thus allowing households who fare poorly on the stock market to hedge their losses by working more to increase their labor income. Chai, Horneff, Maurer, and Mitchell (2011) allow for flexibility both in work hours and in the choice of retirement age. Both papers conclude that the optimal share of stocks in the household's portfolio should be age-dependent, with the share being highest at young ages. In important early work, Heaton and Lucas (1997) find that households would want to allocate all of their savings to stocks under a variety of assumptions, including the presence of transactions costs.

the share of wealth located in stocks, which one might require of a full resolution of the puzzle.

When it comes to antecedents aimed at understanding the equity premium, our work is informed by Constantinides, Donaldson, and Mehra (2002). In their paper, as in Davis, Kubler, and Willen (2006), borrowing constraints play a key role in generating low demand for equities. Households look forward in their life-cycle planning, and would, if allowed, borrow and invest in equity. Borrowing constraints prevent this and hence lower demand for equity (and boost equity premia, both by increasing equity prices and by driving down bond prices) relative to the counterfactual with inexpensive and lax borrowing limits. However, in these settings, households are implicitly endowed with human capital that yields a flow of (stochastic) labor income over the life cycle. Critically, no investment in human capital need be undertaken, leaving high-yield risky financial assets as the sole investment option aside from risk-free assets.

We show that once human capital investment is allowed for, however, borrowing constraints do not as directly hinder investment in financial wealth. This is because the ability to borrow can simply facilitate investment in human, and not financial, wealth, primarily by allowing consumption to take place while learning early in life. Our approach acknowledges that investors in practice have not two, but three, kinds of investment opportunities among which to decide: risky equity, (risky) human capital, and riskless bonds. As a result, any pair of relative returns, such as the equity premium, depends on the investment decisions made along the third dimension, in this case human capital.[5] Recent work of Huggett and Kaplan (2011) finds that, early in life, mean human capital returns exceed those of stocks.

Despite the richness of the models employed by the work above, little work to date has studied portfolios when households may also invest in their human capital. Indeed, we are only aware of three papers that study financial portfolios in the presence of an

[5]Nonetheless, we emphasize that the same forces of differing risk appetite from stocks over the life-cycle are operative in our model as well. Investors vary systematically over the life cycle in their appetite for stock market risk. Early in life, expecting a high but uncertain future income, households welcome the hedge provided by risky equity. Thus, if allowed to borrow (cheaply), they would do so and invest in the stock market (and consume in anticipation of future earnings). The motivation to accept equity risk is heightened by the fact that equity payoffs will not matter as much for consumption (which is influenced by the uncertainty of future labor income—something large for any young person). When households reach (late) middle age, labor earnings are largely decided and uncertainty resolved. At this stage, movements in stock yields will directly impinge on consumption and make households reluctant to invest. Thus, if the young can borrow, they are likely to be the marginal investor and thereby demand only a low equity premium. But if they cannot borrow, the marginal investor will be middle aged, all else equal, and demand (and receive) a high equity premium.

option to invest in human capital. In a theoretical contribution, Lindset and Matsen (2011) provide a stylized theory of investment in financial wealth and education as "expansion options" in a complete markets infinite-horizon economy, where the rental price of human capital is perfectly correlated with the risky financial asset return. The paper provides insights into optimal portfolio weights when taking human capital into account. It is, however, abstract and not aimed at confronting empirical regularities. Roussanov (2010) is arguably the closest work to ours, as it studies portfolio choice in a setting where agents can invest in a college education once in their lifetime and cannot work until it matures, something that may take several periods. Since borrowing is disallowed in that setting, nonparticipation is driven by agents' need to save in order to finance consumption and education during the investment period. While Roussanov (2010) does not directly compare model outcomes to data, he finds that allowing human capital investment can generate reasonable implications for the share of equity in portfolios. In our model, by contrast, households may invest in human capital throughout life and may also borrow, and human capital is disciplined by the empirical distribution of earnings, both cross-sectionally and over the life cycle. We obtain nonparticipation even while allowing for borrowing because households who invest in human capital early in life use borrowing to smooth consumption rather than save in financial assets early in life. Finally, novel work of Kim, Maurer, and Mitchell (2016) examines investment management and inertia in portfolio adjustment in a model that takes into account the fact that doing so is costly in terms of forgone leisure and human capital. We follow their approach to modeling human capital accumulation, though our focus is on measuring the role of human capital accumulation, absent other costs, for life cycle stock-market participation.

Because our approach emphasizes financial investment in a setting that explicitly captures human capital and household earnings *heterogeneity* over the entire life cycle, we follow Ben-Porath (1967), Huggett, Ventura, and Yaron (2011) and Kim, Maurer, and Mitchell (2016). In particular, this work not only endogenizes human capital, but also captures both the life-cycle and cross-sectional distribution of earnings.

We now turn to a description of the data.

3 Data

3.1 Household Portfolios

We begin by describing salient facts about household financial portfolios from the Survey of Consumer Finances (SCF). The SCF is a survey of a cross section of U.S. families conducted every three years by the Federal Reserve Board. It includes information about families' finances as well as their demographic characteristics. While the SCF provides us with rich detail about household finances, it is not a panel, so it does not enable us to directly observe the evolution of finances over the life cycle.

The differences in participation rates across households may be the result of three factors: aggregate fluctuations experienced by all households living in a particular year (time effects), lifetime experiences that vary by year of birth (cohort effects), and getting older (age effects). Since we are interested in participation over the life cycle—the changes in a household's portfolio that result from that household getting older—we need to distinguish age effects from cohort and time effects. The three variables are perfectly collinear (age=year of birth–year of observation), which makes separately identifying the three effects empirically challenging. We separately consider both cohort and time effects and later, in the results section, compare our results to both sets of estimates.

3.1.1 Cohort Effects

We first estimate life cycle profiles of participation in the stock market and stockholdings by making the identifying assumption that time effects are zero. We follow a methodology similar to Poterba and Samwick (1997) to create life-cycle profiles. As Deaton (1985) describes, each successive cross-sectional survey of the population will include a random sample of a cohort if the number of observations is sufficiently large. Using summary statistics about the cohort from each cross section, a time series that describes behavior *as if* for a panel can be generated. In particular, sample cohort means will be consistent estimates of the cohort population mean.

To implement a procedure in this spirit, we begin by pooling households from all nine waves of the 1989-2013 SCF into a single data set. We assign a household to a cohort if the head of the household is born within the three-year period that defines the cohort. We have 24 cohorts in all, with the oldest consisting of households whose head was born between 1919 and 1921 and the youngest consisting of households with heads born between 1988 and 1990. We include all observations where the household head is between the ages of 23 and 79 to be consistent with assumptions we make later in our theoretical model. For the same reason, we exclude from our sample those households whose head has less than a high school diploma.

Except for the cohorts that are too young or too old to be represented in all waves of the survey, we have at least 100 observations of every cohort in each survey year. We use this data to create life-cycle profiles of cohort participation in the stock market. We will define a household as participating in the stock market if they have a positive amount of financial assets invested in equity. The variable in the SCF that measures this includes directly held stocks as well as stocks held in mutual funds, IRAs/Keoghs, thrift-type retirement accounts, and other managed assets.

In Figure 1, we plot the average participation of each of the 24 cohorts over their life cycle (defining the cohort by the mid-point of the age range of the cohort). For example, we observe the cohort born in 1943-45 from the time they are age 44–46 (in the 1989 wave of the SCF) to the time they are age 68–70 (in the 2013 SCF). Figure 1 shows that participation for this cohort increases from roughly 43 to 53 percent.

Figure 1: Household Stock Market Participation Rate by Cohort (SCF)

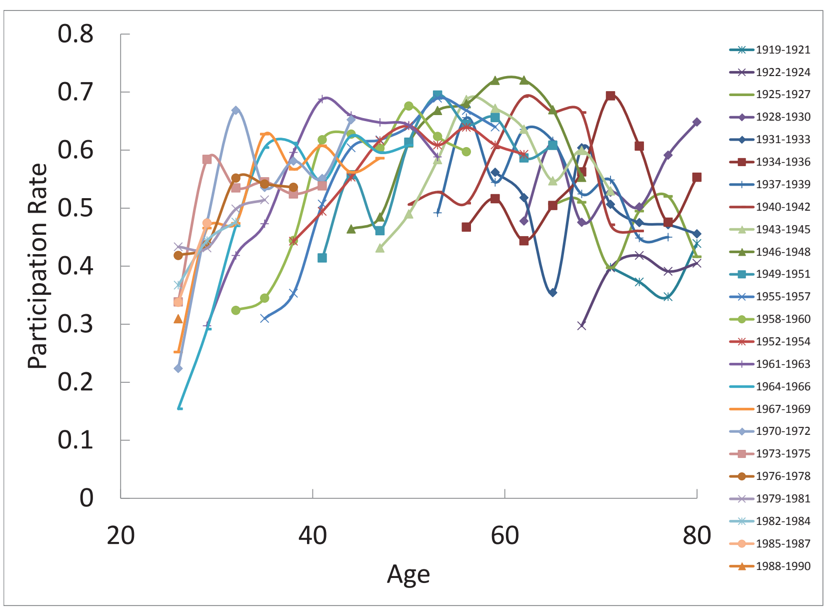

The decision to invest in stocks can be expressed using a standard probit model

$$S_i^* = \alpha + \sum_{n=2}^{21} \beta_n age_{i,n} + \sum_{m=2}^{24} \gamma_m cohort_{i,m} + \epsilon_i \tag{1}$$

where $S_i = 1$ if $S_i^* > 0$ and 0 otherwise. S_i is the discrete dependent variable that equals 1 if household i invests in stocks and zero otherwise. S_i is determined by the continuous, latent variable S_i^*, the actual amount invested in stocks. S_i^*, and thus S_i, is specified in the above as a function of $age_{i,n}$ and $cohort_{i,m}$. We include 19 dummies for age categories ranging from 23–25 to 77–79, with $age_{i,n}$ being the dummy variable that indicates whether the current age of the household head lies in one of these intervals. We include 24 cohort dummies $cohort_{i,m}$ to represent cohorts born in one of the three-year intervals in the range from 1919–21 to 1988–90.

The SCF oversamples wealthy households and therefore needs to be weighted to obtain estimates that are representative of the U.S. population. As in Poterba and Samwick (1997), we estimate Equation (1) using year-specific sample weights normalized such that the sum of the weights (which equals the population represented) remains constant over time. The results of the estimation are reported in Table 2 in the Appendix.[6] We use the coefficients to construct our estimate of the life-cycle profile of stock-market participation. Figure 2 shows the results for the cohort born in 1973–75. (Participation rates are generally lower over the life cycle for older cohorts and higher for younger cohorts.) By our estimation, participation in the stock market increases until agents reach age 60, after which it levels off.

We are also interested in portfolio allocation over the life cycle conditional on participation. In other words, we want to know how the share of assets invested in stocks evolves over the life cycle. To construct this measure, we calculate what fraction of the household's total financial *assets* is invested in equity, where equity is measured by the variable described above. To be consistent with our model, we use a measure of household financial assets from the SCF that excludes housing. Once we deduct the value of equity from this variable, what remains is a measure of the risk-free financial assets held by the households, which includes, for example, certificates of deposit and savings bonds.

Our measure of shares lies between 0 and 1 by construction, so we want our life-cycle

[6] We use all five implicates from the SCF in our estimation. While this provides accurate coefficients, the statistical significance of the results may be inflated. We only need the values of the coefficients to construct life-cycle profiles; therefore, we do not report the results of the significance tests.

Figure 2: Estimated Participation Rate over the Life Cycle (SCF, 1973–75 Birth Cohort)

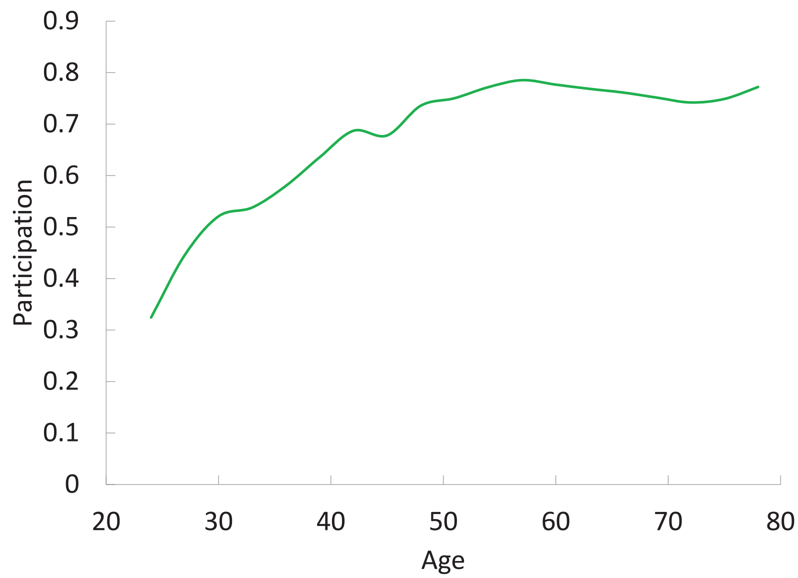

estimate of it to lie between 0 and 1 as well. To ensure this, we construct a logistic transformation to obtain the variable $Y_i = \ln \frac{\frac{s}{s+b}}{1-\frac{s}{s+b}}$. Here s denotes the equity holdings of the households and b denotes the value of risk-free assets. We run the following Ordinary Least Squares (OLS) regression on this variable.[7]

$$Y_i = \alpha + \sum_{n=2}^{21} \beta_n age_{i,n} + \sum_{m=2}^{24} \gamma_m cohort_{i,m} + \epsilon_i \tag{2}$$

The results are reported in Table 3. As we did for participation, we use the reported coefficients to estimate the life-cycle profile of portfolio allocation for the cohort born in 1973–75. Figure 3 shows the results. The estimated share of risky assets conditional on participation increases steadily after age 25.

[7]Note that, unlike Poterba and Samwick (1997), we do not use Tobit to estimate this equation. By construction, our data is not censored— values below 0 and above 1 are infeasible. Moreover, since our variable of interest is the share of risky assets in the household's portfolio conditional on participation, it will always be strictly positive. It is possible for it to exactly equal 1, but we have very few observations with this value, and in this instance we set it to 0.999999.

Figure 3: Estimated Average Fraction of Stocks in Portfolio over the Life Cycle Conditional on Participation for 1973–75 Birth Cohort (SCF)

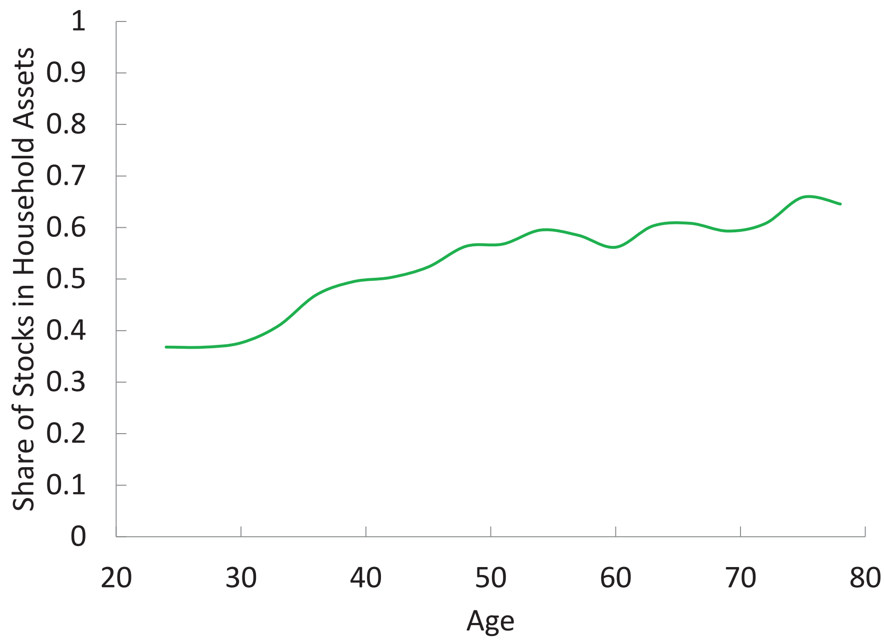

3.1.2 Time Effects

We recognize that making different identifying assumptions can generate different life-cycle estimates, particularly for shares (Ameriks and Zeldes, 2004). Moreover, because participation costs have likely fallen over the past several decades, time effects may be especially relevant for accurately measuring participation. We therefore also estimate participation and shares over the life cycle under a different identifying assumption, namely, that cohort effects are zero.

To estimate participation over the life cycle, we run a probit similar to that in Equation (1), but with time dummies for each year of the SCF instead of cohort dummies. We use 2013 as our base year for reporting the results. The resulting life-cycle profile is shown in Figure 4.[8]

Correspondingly, we run an OLS regression as in Equation (2) with time dummies

[8]The results of the estimation are reported in Table 4 in the Appendix.

Figure 4: Estimated Participation Rate over the Life Cycle (SCF, 2013 base year)

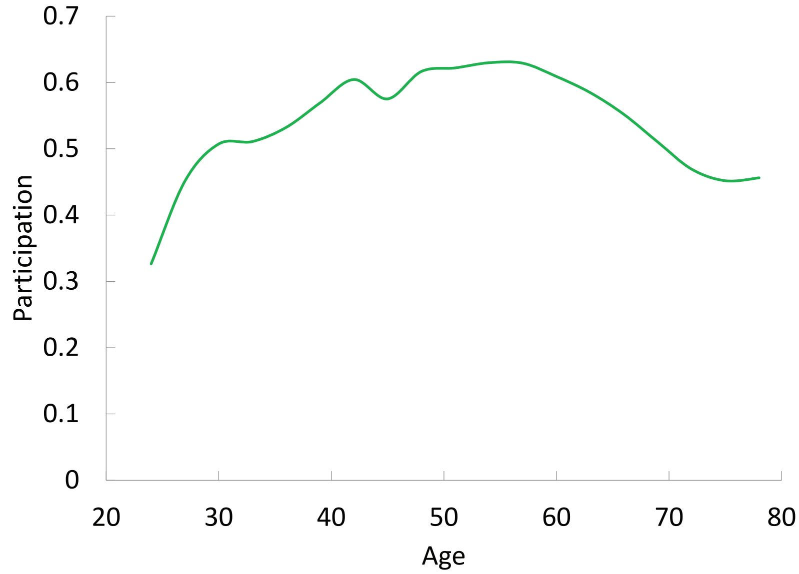

to estimate the life-cycle profile of shares. Figure 5 shows the result.[9]

Observe that different identifying assumptions do indeed lead to different estimates for the life-cycle profiles for participation and shares. In particular, under the assumption that time effects matter and that cohort effects are zero, we obtain hump-shaped rather than increasing profiles for both participation and shares. Our findings are consistent with those previously reported by Ameriks and Zeldes (2004).

3.2 Earnings

Next, we compute statistics of age-earnings profiles from the CPS for 1969-2002 using a synthetic cohort approach, following Ionescu (2009). To be precise, we use the 1969 CPS data to calculate the earnings statistics of 25-year-olds, the 1970 CPS data to compute earnings statistics of 26-year-olds, and so on. We include only those who have at least 12 years of education, to correspond with our modeling assumption that agents start life after high school. To compute the mean, inverse skewness, and Gini of earnings for households of age a in any given year, we average the earnings of household

[9]The results of the estimation are reported in Table 5 in the Appendix.

Figure 5: Estimated Average Fraction of Stocks in Portfolio over the Life Cycle Conditional on Participation (SCF, 2013 base year)

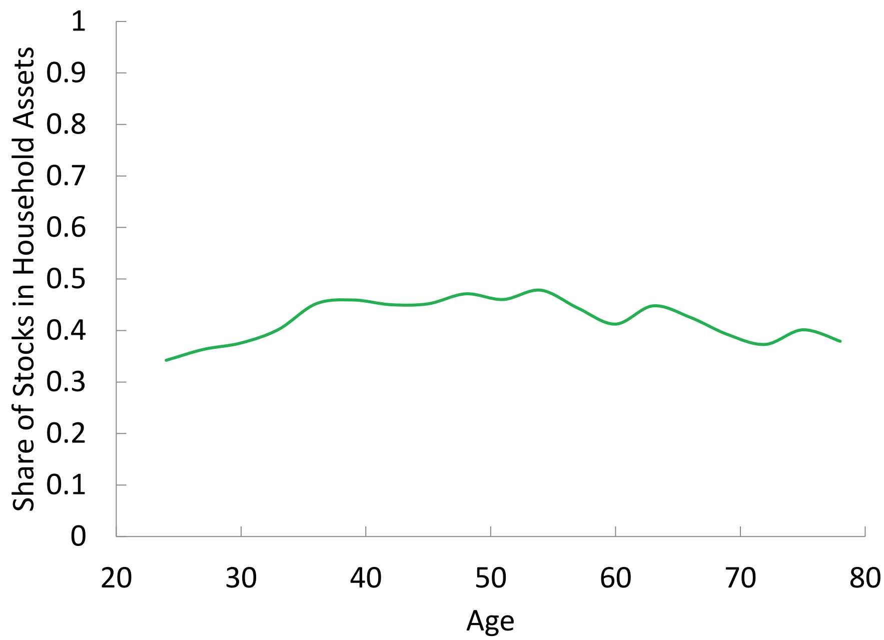

heads between the ages of $a-2$ and $a+2$ to obtain a sufficient number of observations. Life-cycle profiles for all three statistics are shown in Figure 29 in the Appendix.[10]

With these facts in hand, we turn to the description of the model.

4 Model

Our model is a standard model of life-cycle consumption and savings in the presence of uninsurable risk (e.g. Gourinchas and Parker, 2002), but it contains two enrichments. First, households choose their level of human capital, and second, households can invest in both risky and riskless assets.

The economy is populated by a continuum of agents who value consumption throughout a finite life. Age is discrete and indexed by $t = 0, ..., T$, where $t = 1$ represents

[10]We obtain real earnings in 2013 dollars using the Consumer Price Index. We convert earnings to model units such that mean earnings at the end of working life, which equal $70,800, are set to 100.

the first year after high school graduation, and $t = J$ represents the age of retirement. Agents enter the model endowed with an initial level of human capital, h_0, which varies across the population. This embodies human capital accumulated by the time agents graduate high school.

In each period, households can divide their time between work and the accumulation of human capital, as in the classic model of Ben-Porath (1967). Households consume and decide how to allocate any wealth they have in period t between a risky asset s_{t+1} and a risk-free asset b_{t+1}. Households also have the option to borrow, that is $b_t \geq -\underline{b}$, with $\underline{b} > 0$, may be positive or negative.

To capture risk and heterogeneity, we follow Huggett, Ventura, and Yaron (2011) and allow for four potential sources of heterogeneity across agents — their immutable learning ability, a; human capital stock, h; initial assets, x; and subsequent shocks to the yield on their holdings of human capital, i.e., their earnings. The set of initial characteristics are jointly drawn according to a distribution $F(a, h, x)$ on $A \times H \times X$. Lastly, households are not subject to risks once they retire, i.e., once $t > J$.

4.1 Preferences

All agents have identical preferences, with their within-period utility given by a standard CRRA function with parameter σ and with a common discount factor β. The general problem of an individual is to choose consumption over the life cycle, $\{c_t\}_{t=1}^{T}$, to maximize the expected present value of utility over the life cycle,

$$\max_{(\{c_t\} \in \Pi(\Psi_0))} E_0 \sum_{t=1}^{T} \beta^{t-1} \frac{c_t^{1-\sigma}}{1-\sigma} \tag{3}$$

$\Pi(\Psi_0)$ denotes the space of all feasible combinations $\{c_t\}_{t=1}^{T}$, given initial state $\Psi_0 \equiv \{a_0, h_0, x_0\}$. Agents do not value leisure.

4.2 Financial Markets

Our focus throughout is on the implications of human capital investment for participation in the market for risky financial assets. We therefore model the household as having access to two forms of financial assets: a risk-free asset, b_t, to be interpreted as savings (or borrowing when negative), and a risky asset, s_t, to be interpreted as stock-market equity. Of course, as an empirical matter, households have the option to accumulate real physical assets as part of their overall investment strategy, including equity in an owner-occupied home, car, and other consumer durables. However, we abstract from these additional assets for two reasons. First, while central to certain questions, the inclusion of durables is unlikely to be critical for understanding the relationship between human and financial wealth accumulation. Second, we are particularly interested in accounting for low stock-market participation early in life, a time when equity positions in durable goods (including, especially, in home equity) are typically minor for nearly all households. We acknowledge, nonetheless, that durables may exert independent influence on overall stock market participation; for a model that studies the role of housing—though in the absence of human capital investment—see Cocco (2005).

Risk-free assets

An agent can borrow or save by taking negative or positive positions, respectively, in a risk-free asset b_t. Savings ($b_t \geq 0$) will earn the risk-free interest rate, R_f. Borrowing ($b_t < 0$) resembles unsecured credit and carries an additional (proportional) cost as in Davis, Kubler, and Willen (2006), denoted by ϕ, to represent costs of intermediating credit. The borrowing rate, R_b, therefore, is higher than the savings rate and given by $R_b = R_f + \phi$. As noted above, borrowing is subject to a limit $\underline{b}$. We assume that debt is nondefaultable.[11]

Risky assets

For ease of exposition, we will refer to the risky asset as "stocks" and denote the agent's holdings of these claims between period t and $t+1$ by s_{t+1}. Stocks yield their owners a stochastic gross real return in period $t+1$, $R_{s,t+1}$ whereby the excess return on stocks is given by:

$$R_{s,t+1} - R_f = \mu + \eta_{t+1}. \tag{4}$$

The first term μ is the mean excess return to stocks. The second, η_{t+1}, represents the period $t+1$ innovation to excess returns and is assumed to be independently and identically distributed (i.i.d.) over time with distribution $N(0, \sigma_\eta^2)$.

[11] We believe that this is a reasonable assumption both because default rates on credit card debt are low in the data and because individuals close to default will likely have not accumulated resources to engage in financial market participation. Therefore the option to default on unsecured debt is not central for bond and stock market choices.

Given asset investments at age t, b_{t+1} and s_{t+1}, financial wealth at age $t+1$ is given by $x_{t+1} = R_i b_{t+1} + R_{s,t+1} s_{t+1}$, with $R_i = R_f$ if $b \geq 0$ and $R_i = R_b$ if $b < 0$.

4.3 Human Capital

The key innovation of our work is to allow for human capital investment in a model of portfolio choice. We do this by employing the workhorse model of Ben-Porath (1967), extended to allow for risks to the payoff from human capital: in each period, agents can apportion some of their time to acquiring human capital, or they may work and earn wages that depend on current human capital and shocks. At any given date, an agent's human capital stock summarizes their ability to turn their time endowment into earnings. In this sense, it reflects *earning* ability and, critically, can be accumulated over the life cycle. By contrast, *learning* ability, which governs the effectiveness of the production function that maps time to human capital investment, is fixed at birth and does not change over time. Both learning ability and initial human capital will be allowed to vary across agents and, as we will demonstrate, heterogeneity in each is implied by earnings heterogeneity in the data among the youngest cohorts and by the subsequent evolution of earnings dispersion.

Human capital investment in a given period occurs according to the human capital production function, $H(a, h_t, l_t)$, which depends on the agent's immutable learning ability, a, human capital, h_t, and the fraction of available time put into human capital production, l_t. Human capital depreciates at a rate of δ. The law of motion for human capital is given by

$$h_{t+1} = h_t(1 - \delta) + H(a, h_t, l_t) \tag{5}$$

Following Ben-Porath (1967), the human capital production function is given by $H(a, h, l) = a(hl)^\alpha$ with $\alpha \in (0, 1)$. As demonstrated by Huggett, Ventura, and Yaron (2006), the Ben-Porath model has the additional advantage of being able to match the dynamics of the U.S. earnings distribution given the appropriate joint distribution of initial ability and human capital.

4.4 Labor Income

Human capital confers a return (i.e., its rental rate, wages) in each period that is subject to stochastic shocks. Specifically, earnings are given by a product of the stochastic component, z_t, the rental rate of human capital, w_t, the agent's human capital, h_t, and the time spent in market work, $(1 - l_t)$.

Therefore, agent i's earnings in period t are given by

$$log(y_{it}) = G(w_t, h_t, l_t) + z_{it} \tag{6}$$

with $G(w_t, h_t, l_t)$ representing the deterministic component as a function of rental rate w_t, human capital stock at age t, h_t, and labor effort, $1 - l_t$, and z_t representing the stochastic component. The rental rate of human capital evolves over time according to $w_t = (1 + g)^{t-1}$ with the growth rate, g.[12]

The stochastic component, z_{it}, consists of an idiosyncratic temporary (i.i.d) shock $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ and a persistent shock u_{it}:

$$z_{it} = u_{it} + \epsilon_{it}$$

where

$$u_{it} = \rho u_{i,t-1} + \nu_{it}$$

follows an AR(1) process as in Abbott, Gallipoli, Meghir, and Violante (2013), with $\nu_{it} \sim N(0, \sigma_\nu^2)$ representing an innovation to u_{it}. The variables u_{it} and ϵ_{it} are realized at each period over the life cycle and are not correlated.

4.5 Means-Tested Transfer and Retirement Income

To accurately capture the risk-management problem of the household, it is important to make allowance for additional sources of insurance that may be present. In the United States, there is a vast array of social-insurance programs that, if effective, bound households' purchasing power away from zero. Moreover, it is well known, since at least Hubbard, Skinner, and Zeldes (1995), that such a system may be acting to greatly diminish savings among households who earn relatively little. In our model, this will consist of unlucky households, households with low learning ability, or both. To ensure that we confront households with an empirically relevant risk environment in which they choose portfolios, we specify a means-tested income transfer system, which, in addition to asset accumulation, can provide another source of insurance against labor income risk (Campbell, Cocco, Gomes, and Maenhout, 2001). Agents receive means-tested transfers from the government, τ_t, which depend on age, t, income, y_t, and net

[12] The growth rates for wages are estimated from data, as described further below.

assets, x_t. These transfers capture the fact that in the U.S. social insurance is aimed at providing a floor on consumption. Following Hubbard, Skinner, and Zeldes (1995), we specify these transfers by

$$\tau_t(t, y_t, x_t) = \max\{0, \underline{\tau} - (\max(0, x_t) + y_t)\} \tag{7}$$

Total pre-transfer resources are given by $\max(0, x_t) + y_t$ and the means-testing restriction is represented by the term $\underline{\tau} - \max((0, x_t) + y_t)$. These resources are deducted to provide a minimal income level $\underline{\tau}$. For example, if $x_t + y_t > \underline{\tau}$ and $x_t > 0$, then the agent gets no public transfer. By contrast, if $x_t + y_t < \underline{\tau}$ and $x_t > 0$, then the agent receives the difference, in which he has $\underline{\tau}$ units of the consumption good at the beginning of the period. Agents do not receive transfers to cover debts, which requires the term $\max(0, x_t)$. Lastly, transfers are required to be nonnegative, which requires the "outer" max.

After period $t = J$ when agents start retirement, they get a constant fraction ψ of their income in the last period as working adults, y_J, which they divide between risky and risk-free investments.

4.6 Agent's Problem

The agent's problem is to maximize lifetime utility by choosing asset positions in stocks and bonds (or borrowing), and, in what is novel in our paper, time allocated throughout life to market work and human capital investment.

We formulate the problem recursively. The household's feasible set for consumption and savings is determined by its age, t, ability, a, beginning-of-period human capital, h, net worth, $x(b, s)$, current-period realization of the persistent shock to earnings, u, and current-period transitory shock, ϵ.

In the last period of life, agents consume all available resources. The value function in the last period of life is therefore simply their payoff from consumption in that period. Prior to this terminal date, but following working life, agents are retired. Retired agents do not accumulate human capital and do not face human capital risk. Thus, we have $V_T^R(a, x, y_J) = \frac{c^{1-\sigma}}{1-\sigma}$, where $c = x(b, s) + \psi y_J$. Notice that, when retired, human capital is irrelevant as a state, and in what follows, is not part of the household's state. Retired households face a standard consumption-savings problem, though, as in working life, they may invest in both risk-free and risky assets. Indeed, in retirement, the only risk agents face comes from the uncertain return on stocks. Their value function for retirees is given by

$$V^R(t,a,b,s,y_J) = \sup_{b',s'}\{\frac{c_t^{1-\sigma}}{1-\sigma} + \beta E_{R_s'} V^R(t+1,a,b',s',y_J)\} \tag{8}$$

where

$$\begin{aligned} c + b' + s' &\leq \psi y_J + R_i b + R_s s \\ b &\geq \underline{b} \end{aligned}$$

In the budget constraint, we remind the reader that $R_i = R_f$ if $b \geq 0$ and $R_i = R_b$ if $b < 0$.

During working life, the agent faces uncertainty from the returns on human capital as well as from any risk assumed in the portfolio they choose. The budget constraint makes clear that current consumption c and total net financial wealth next period $(b' + s')$ must not exceed the sum of current labor earnings $w(1-l)hz$, the value of the portfolio $(R_i b + R_s s)$, and any transfers from the social safety net $\tau(t, y, x)$.

$$V(t,a,h,b,s,u,\epsilon) = \sup_{l,h',b',s'}\{\frac{c_t^{1-\sigma}}{1-\sigma} + \beta E_{u'|u,\, R_s'} V(t+1,a,h',b',s',u',\epsilon')\} \tag{9}$$

where

$$\begin{aligned} c + b' + s' &\leq w(1-l)hz + R_i b + R_s s + \tau(t,y,x) \text{ for } t = 1,..,J-1 \\ \text{s.t. } l &\in [0,1],\ h' = h(1-\delta) + a(hl)^\alpha,\ b \geq \underline{b} \end{aligned}$$

The value function $V(t,a,h,b,s,u,\epsilon)$ thus gives the maximum present value of utility at age t from states h, b, and s, when learning ability is a and the realized shocks are u and ϵ. The solution to this problem is given by optimal decision rules $l_j^*(t,a,h,b,s,u,\epsilon)$, $h^*(t,a,h,b,s,u,\epsilon)$, $b^*(t,a,h,b,s,u,\epsilon)$, and $s^*(t,a,h,b,s,u,\epsilon)$, which describe the optimal choice of the fraction of time spent in human capital production, the level of human capital, and risk-free and risky assets carried to the next period as a function of age, t, human capital, h, ability, a, and current assets, b and s, when the realized shocks are u and ϵ.

5 Mapping the model to the data

There are four sets of parameters in the model: 1) standard parameters, such as the discount factor and the coefficient of risk aversion; 2) parameters specific to asset markets; 3) parameters specific to human capital and to the earnings process; and 4) parameters for the initial distribution of characteristics. Our approach includes a combination of setting some parameters to values that are standard in the literature, calibrating some parameters directly to data, and jointly estimating those parameters that we do not directly observe in the data by matching moments for several observable implications of the model. We summarize parameter values in Table 1 and describe in detail below how we obtain them.

Table 1: Parameter Values: Benchmark Model

Parameter	Name	Value
T	Model periods (years)	53
J	Working periods	33
β	Discount factor	0.96
σ	Coeff. of risk aversion	5
R_f	Risk-free rate	1.02
R_b	Borrowing rate	1.11
μ	Mean equity premium	0.06
σ_η	Stdev. of innovations to stock returns	0.157
α	Human capital production function elasticity	0.7
g	Growth rate of rental rate of human capital	0.0013
δ	Human capital depreciation rate	0.0114
ψ	Fraction of income in retirement	0.68
$\underline{\tau}$	Minimal income level	$\$17,936$
$(\rho, \sigma_\nu^2, \sigma_\epsilon^2)$	Earnings shocks	$(0.951, 0.055, 0.017)$
μ_a, σ_a	Parameters for joint distribution of ability	$0.246, 0.418$
$\mu_h, \sigma_h, \varrho_{ah}$	and initial human capital	$87.08, 35.11, 0.57$

We follow agents from age 25 onward, as this captures the beginning of the portion of life in which households make nontrivial investments in financial assets and in learning on the job. Agents live $T = 53$ model periods, which corresponds to ages 25 to 78, and retire at age $J = 58$.

5.1 Preference and Financial Market Parameters

The per period utility function is CRRA, $u(c_t) = \frac{{c_t}^{1-\sigma}}{1-\sigma}$, with the coefficient of risk aversion $\sigma = 5$, which is consistent with values chosen in the financial literature. Risk aversion is a key parameter and so we conduct robustness checks on it, in particular we consider higher values up to the upper bound of $\sigma = 10$ considered reasonable by Mehra and Prescott (1985). We also consider lower values, such as $\sigma = 3$. The discount factor ($\beta = 0.96$) chosen is also standard in the literature.

We turn now to the parameters in the model related to financial markets. We fix the mean equity premium to $\mu = 0.06$, as is standard (e.g., Mehra and Prescott, 1985). The standard deviation of innovations to the risky asset is set to its historical value, $\sigma_\eta = 0.157$. The risk-free rate is set equal to $R_f = 1.02$, consistent with values in the literature (McGrattan and Prescott, 2000) while the wedge between the borrowing and risk-free rate is $\phi = 0.09$ to match the average borrowing rate of $R_b = 1.11$ (Board of Governors of the Federal Reserve System, 2014). Lastly, we assume that innovations to excess returns are uncorrelated with innovations to the aggregate component of permanent labor income.[13].

5.2 Human Capital and Earnings Parameters

The rental rate on human capital equals $w_t = (1+g)^{t-1}$ where g is set to 0.0013, as in Huggett, Ventura, and Yaron (2006). Given this growth rate, the depreciation rate is set to $\delta = 0.0114$, so that the model produces the rate of decrease of average real earnings at the end of the working life cycle observed in the data. The model implies that at the end of the life cycle negligible time is allocated to producing new human capital and, thus, the gross earnings growth rate approximately equals $(1+g)(1-\delta)$. We set the elasticity parameter in the human capital production function, α, to 0.7. Estimates of this parameter are surveyed by Browning, Hansen, and Heckman (1999)

[13] Evidence on this correlation is mixed, ranging from negative to strongly positive. For instance, Lustig and Van Nieuwerburgh (2008) show that innovations in current and future human wealth returns are negatively correlated with innovations in current and future financial asset returns, regardless of the elasticity of intertemporal substitution, while Benzoni, Collin-Dufresne, and Goldstein (2007) argue that the correlation in labor income flows and stock market returns is positive and large in particular at long horizons. At the same time, prior studies that have examined the relation between labor income and life-cycle financial portfolio choice assume that labor income shocks are (nearly) independent from stock market return innovations (see Cocco, Gomes, and Maenhout, 2005; Davis, Kubler, and Willen, 2006; Davis and Willen, 2013; Gomes and Michaelides, 2005; Haliassos and Michaelides, 2003; Roussanov, 2010; and Viceira, 2001)

and range from 0.5 to 0.9.

In the parametrization of the stochastic component of earnings, $z_{it} = u_{it} + \epsilon_{it}$, we follow Abbott, Gallipoli, Meghir, and Violante (2013), who use the National Longitudinal Survey of Youth (NLSY) data using CPS-type wage measures to estimate the autoregressive coefficients for the transitory and persistent shocks to wages. For the persistent shock, $u_{it} = \rho u_{i,t-1} + \nu_{it}$, with $\nu_{it} \sim N(0, \sigma_\nu^2)$ and for the idiosyncratic temporary shock, $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$, they report the following values for high school graduates: $\rho = 0.951$, $\sigma_\nu^2 = 0.055$, and $\sigma_\epsilon^2 = 0.017$. We set retirement income to be a constant fraction of labor income earned in the last year in the labor market. Following Cocco (2005) we set this fraction to 0.682, the value for high school graduates. The income floor, $\underline{\tau}$, is expressed in 2013 dollars and is consistent with the levels used in related work (e.g. Athreya, 2008).[14]

We assume a uniform credit limit across households. We obtain the value for this limit from the SCF. The SCF reports, for all individuals who hold one or more credit card, the sum total of their credit limits. We take the average of this over all individuals in our sample and obtain a value of approximately $17,000 in 2013 dollars. Note that, when we take the average, we include those who do not have any credit cards. This ensures that we are not setting the overall limit to be too loose. Lastly, in our baseline model, we assume that the returns to both risky assets (human capital and stocks) are uncorrelated.

5.3 The Distribution of Assets, Ability, and Human Capital

We turn now to parameters defining the joint distribution of initial heterogeneity in the unobserved characteristics central to human capital accumulation. There are seven parameters, and using only these, we are able to closely match the evolution, over the entire life cycle, of three functions of moments of the earnings distribution: mean earnings, the ratio of mean to median earnings, and the Gini coefficient of earnings.

To estimate the parameters of this distribution, we proceed as follows. First, for the asset distribution, we use the SCF data described in Section 3 to compute the mean and standard deviation of initial assets to be $22,568 and $24,256, respectively, in 2013 dollars. Second, we calibrate the initial distribution of ability and human capital to match the key properties of the life-cycle earnings distribution reported earlier using the CPS for 1969-2002.

[14] The results turn out to be robust to the choice of this parameter; results are available upon request.

Earnings distribution dynamics implied by the model are determined in several steps: i) we compute the optimal decision rules for human capital using the parameters described above for an initial grid of the state variable; ii) we simultaneously compute financial investment decisions and compute the life-cycle earnings for any initial pair of ability and human capital; and iii) we choose the joint initial distribution of ability and human capital to best replicate the properties of U.S. data.

Figure 6: Life-cycle earnings

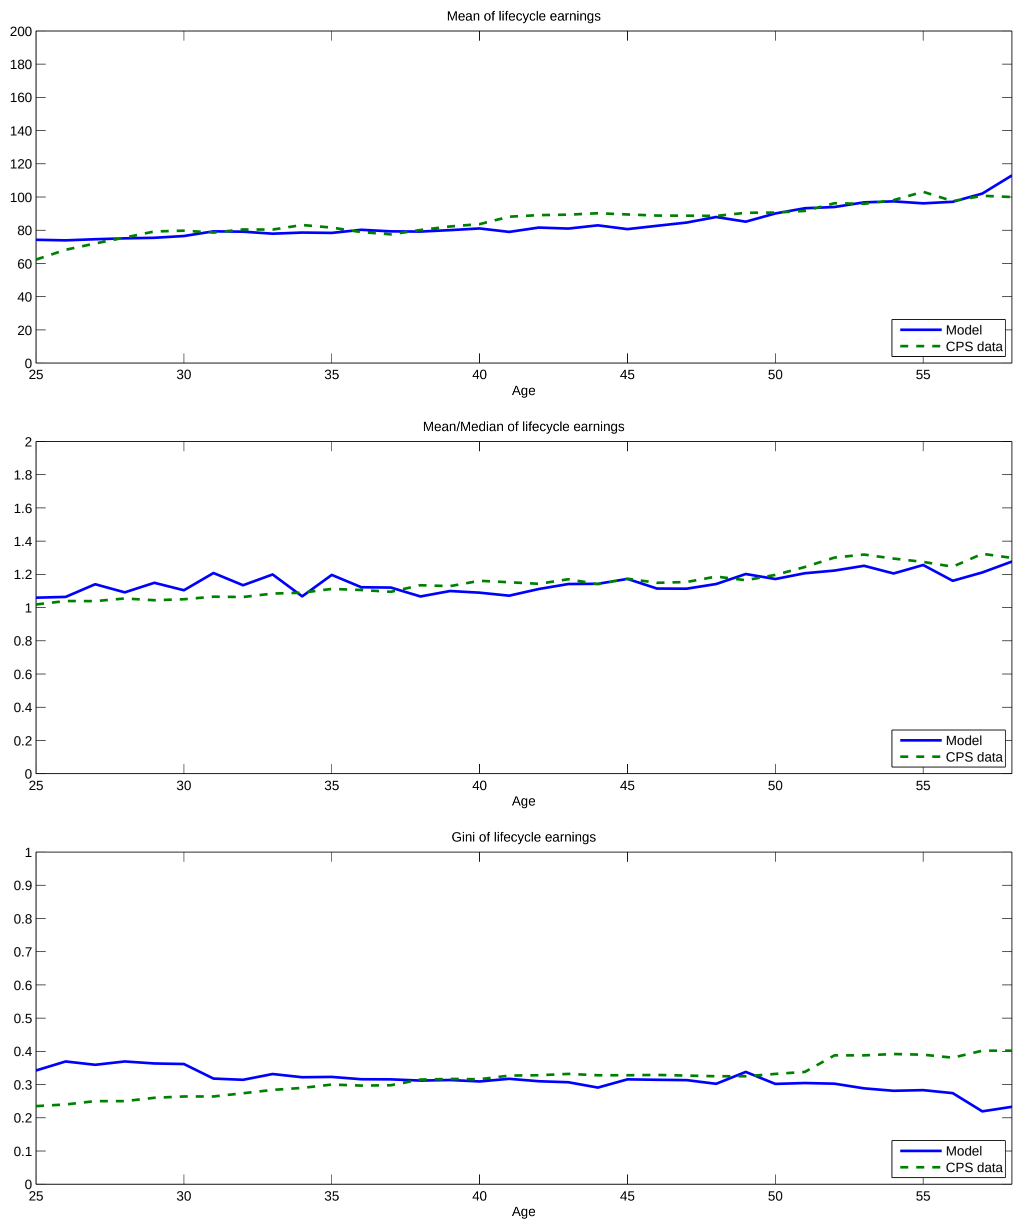

To set values for these parameters, we search over the vector of parameters that characterize the initial state distribution to minimize a distance criterion between the model and the data. We restrict the initial distribution to lie on a two-dimensional grid spelling out human capital and learning ability, and we assume that the underlying distribution is jointly log-normal. This class of distributions is characterized by five parameters.[15] We find the vector of parameters $\gamma = (\mu_a, \sigma_a, \mu_h, \sigma_h, \varrho_{ah})$ characterizing the initial distribution by solving the minimization problem

$$\min_{\gamma} \left(\sum_{j=5}^{J} |log(m_j/m_j(\gamma))|^2 + |log(d_j/d_j(\gamma))|^2 + |log(s_j/s_j(\gamma))|^2 \right)$$

where $m_j, d_j,$ and s_j are mean, dispersion, and inverse skewness statistics constructed from the CPS data on earnings, and $m_j(\gamma)$, $d_j(\gamma)$, and $s_j(\gamma)$ are the corresponding model statistics. Overall, we match 102 moments.[16] Figure 6 illustrates the earnings profiles for individuals in the model versus CPS data when the initial distribution is chosen to best fit the three statistics considered. We obtain

$$\gamma = (0.246, 0.418, 87.08, 35.11, 0.57)$$

The model performs well given riskiness of assets and stochastic earnings in the current paper.

6 Results

With the model parameterized as described above, we are now in a position to obtain a quantitative account of household financial investment—with specific attention to the extensive margin of stock-market participation over the life cycle—when household human capital investment is disciplined by earnings data.

6.1 Understanding Stock Market Participation

We begin by studying our model's predictions for the stock-market participation rate. Figure 7 compares our model results with our two empirical estimates (considering time

[15] In practice, the grid is defined by 20 points in human capital and ability.

[16] For details on the calibration algorithm see Huggett, Ventura, and Yaron (2006) and Ionescu (2009).

effects and cohort effects, respectively) from SCF data. It is clear that stock-market participation in our model is broadly consistent with the data. Importantly, we see that nonparticipation is not a pathology, but rather a direct implication of our essentially standard model.

Figure 7: Life-Cycle Stock Market Participation

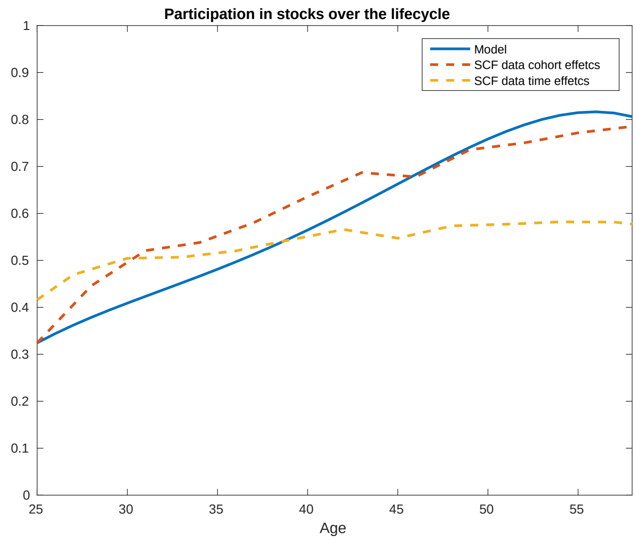

To understand the path of participation, a first step is to study the model's predictions for the trajectory of time invested in human capital accumulation (Figure 8). As is entirely standard in the Ben-Porath model, time spent on human capital accumulation is at its highest early in life, then declines somewhat, and then declines sharply. In our model, this discipline imposed by earnings data implies for instance, that at age 25, households spend about a third of their time on average on human capital accumulation. During the early part of life, we also see that only around 30 percent of all households participate in the stock market. Diminishing returns, and a shorter horizon to recoup investment, imply that human capital accumulation should fall with age, just as it does. Indeed, as retirement approaches, we see that the fraction of time allocated to human capital falls sharply, reaching below 0.05 by retirement age. Correspondingly, we see that stock-market participation steadily increases with age, reaching around 80 percent at retirement.

Figure 8: Time Allocated to Human Capital over the Life Cycle

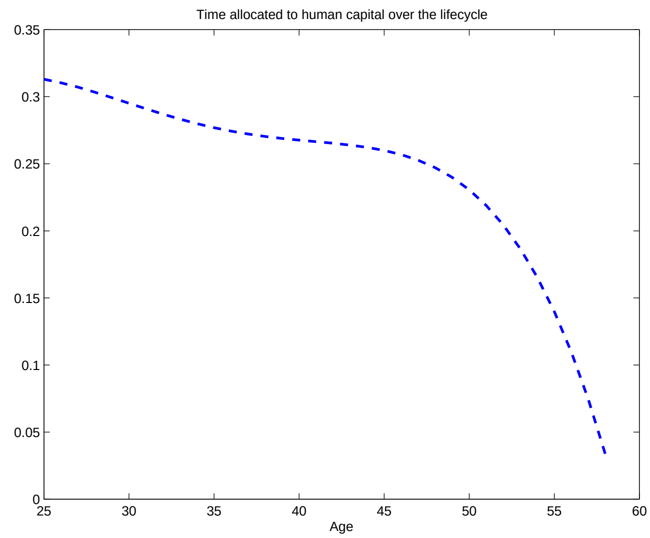

Given that households choose to allocate their time in a manner that yields a hump-shaped earnings profile—and implies low income early in life, they are unlikely to save substantial amounts in *any* asset, let alone stocks. So the immediate question is: are the model's predictions consistent with observed household asset accumulation over the life cycle? If not, the model may be identifying a force that while qualitatively possible, is quantitatively irrelevant. The answer is given in Figure 9: wealth accumulation predicted by our model—as well as the trend of each of its components (risky and risk-free assets)—is remarkably consistent with the data, despite not being targeted in any manner.[17] Thus, our findings for stock market participation arise from a model that captures the salient quantitative and qualitative features of household income and savings, and hence, of consumption as well, throughout the life cycle.

[17]As we did for participation, we report two estimates for life-cycle wealth from the SCF data, one adjusted for time effects and the other for cohort effects. In all cases, we try to make consistent comparisons with the model. The total wealth figure is reported only for those who hold nonnegative amounts in the safe asset, both in the model and in the data. However, the values reported for for the risk-free asset include those who borrow in the model, so the data comparison is with risk-free assets net of credit card debt.

Figure 9: Life-Cycle Wealth Accumulation

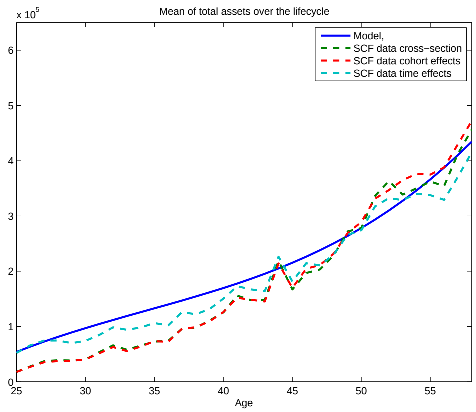

(a) Total Assets

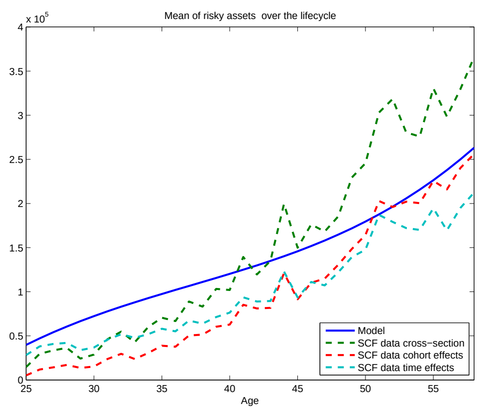

(b) Risky Assets

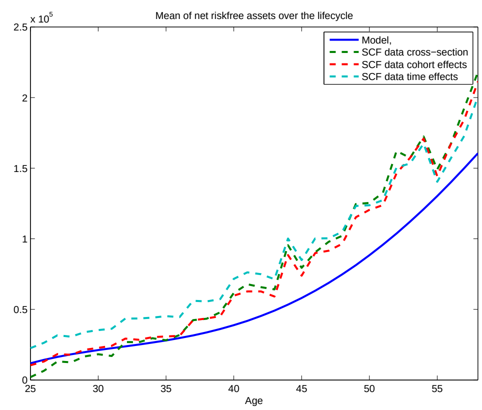

(c) Net Risk-free Assets

6.2 The Importance of Heterogeneity

While stocks earn the same return for all households, the rate of return to human capital investment varies across households. It is this heterogeneity in returns that leads households of the same age and financial wealth to make different decisions about stock-market participation. We now demonstrate that an accurate representation of earnings heterogeneity is vital for generating stock market participation rates. The easiest way to do this is to consider outcomes when important aspects of heterogeneity are shut down. Specifically, we set the values for ability and initial human capital at their respective medians. All other parameters of the model, including shocks to earnings, remain the same as in the benchmark. The results are reported in Figure 10.

Figure 10: Participation in the Absence of Heterogeneity

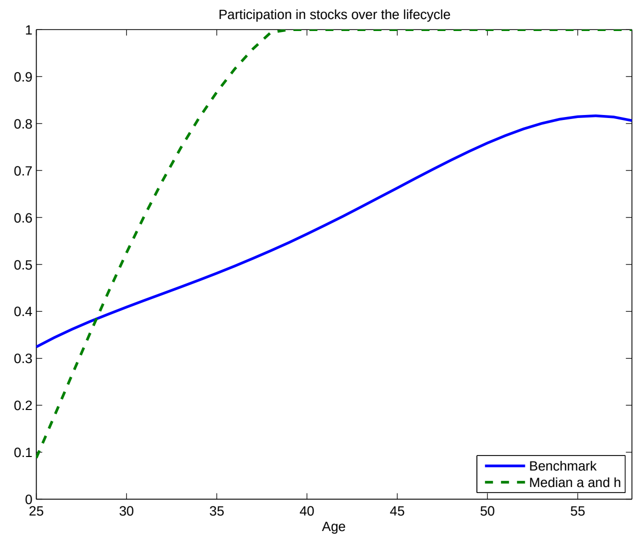

We see immediately in this case that stock market participation rises extremely rapidly and becomes universal by age 35 or so, which is a commonly-found result in the literature. Given that ability and initial human capital do not vary across households, and that the only source of variation in returns to human capital investment is earnings shocks, all households now face similar incentives to invest in human capital. As in our benchmark model, early in life, households borrow to smooth consumption while spending time accumulating human capital, and do not invest in stocks. However, by the time they reach their mid-thirties, it becomes optimal for the households to spend time earning rather than learning and to accumulate savings. At this time, they all choose to enter the stock market, which results in participation rates rising rapidly to 100%.

In the preceding exercise we simply limited heterogeneity, which means that the model's implications for earnings were inaccurate by construction. We now consider another case in which we restore heterogeneity and households face empirically accurate earnings paths. The difference between this experiment and the benchmark is that we do not allow households to invest in human capital but rather assign households the same mean earnings paths that they would have faced in the benchmark, given their ability and initial human capital (and given optimal decisions with respect to learning and earning). We see that once such a model without explicit human capital investment is calibrated to match the properties of earnings, participation once again resembles the benchmark model. This is shown in Figure 11. Note, however, that the predicted path of participation is not identical to the benchmark. The main reason for this is that in the benchmark, human capital investment is endogenous, which means households have an additional use for borrowing—to finance consumption while learning. When this channel is shut down, households participate in the stock market at a slightly greater rate than in the benchmark, and may use borrowing to do so. The absence of additional motives to borrow leads households to accumulate financial wealth at earlier ages, and as a result, leads participation to rise somewhat more steeply than in the benchmark economy. More generally, we will show in section 6.4 that an important benefit to endogenizing human capital is that it clarifies the role played by borrowing costs in observed stock market participation.

Figure 11: Participation with Heterogeneity and No Human Capital Investment

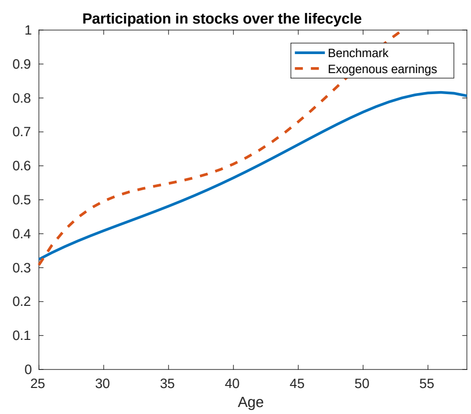

Taken as a whole, our results illustrate the importance of ensuring that the dispersion in human capital returns, and hence, the incentives to accumulate wealth and learn, are accurately represented.

6.3 The Role of Ability and Initial Human Capital

From the fact that the model is able to generate quantitatively sensible stock-market participation, it is apparent that our model captures key aspects of investor heterogeneity seen in the data. The two "initial" sources of heterogeneity in the population that we emphasize are (i) dispersion in individuals' initial human capital and (ii) dispersion in their ability to learn. We now illustrate the mechanisms at work in the model that translate differences in these characteristics into differences in equity participation.

All else equal, initial human capital and investment in human capital will go in opposite directions. This is because those with low initial human capital are faced with high marginal returns to human capital investment as well as a low opportunity cost of learning (low initial human capital implies low earnings today). As a result, we can expect that individuals with low initial human capital will earn relatively low amounts when young but relatively higher amounts when older. This upward-sloping expected earnings profile implies—by the logic of consumption smoothing—low net financial wealth accumulation when young, including stock market participation. As argued above, borrowing will, if it occurs, be channeled toward consumption, not stock market investment.

By the same argument, any individual entering adult life with substantial human capital will find additional accumulation relatively unproductive (due to diminishing marginal returns) and costly—after all, such individuals face high opportunity costs of spending time learning instead of earning, precisely because they can instead rent out their relatively large stock of human capital. As households age and retirement approaches, savings will take center stage, and stock markets will be used along with other means of saving.

Analogous intuition holds for the effect of learning ability. Higher ability implies higher proficiency in acquiring human capital, i.e., in turning time into increments to one's stock of human capital. This means, in turn, a higher current opportunity cost of failing to invest in human capital, and in equilibrium, higher expected future earnings than current earnings.

We now illustrate these forces through an experiment in which we study stock-market participation across households with different initial characteristics. Specifically, we look at participation separately by quartiles of initial human capital and ability. In the baseline model parametrization, the data are consistent with the presence of a substantial positive correlation between these two variables. As a result, each quartile of the distribution of initial human capital will be be accompanied by a different distribution of ability, and vice versa. In order to isolate the effects of initial human capital and ability separately, we set the correlation between initial human capital and ability to zero. Given the bivariate lognormal joint distribution of these variables, the implication for this experiment is that the conditional distributions of each (ability given human capital and vice versa) do not vary with the other.

6.3.1 Initial Human Capital

We turn first to initial human capital, and the implications of its dispersion across households for stock market participation.

Figure 12: Investment by Quartiles of Initial Human Capital

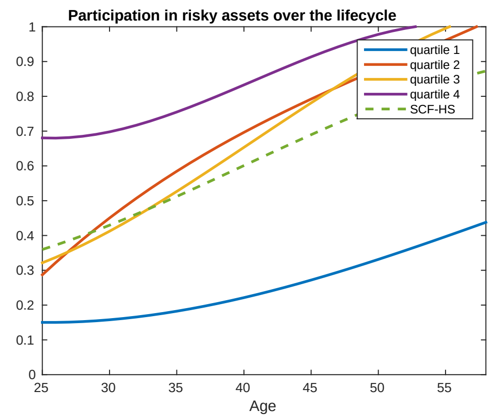

(a) Participation

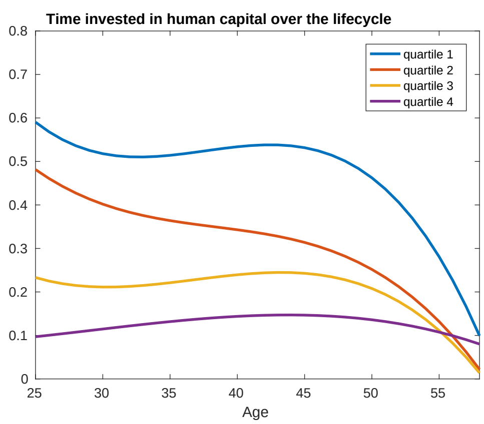

(b) Time invested in human capital

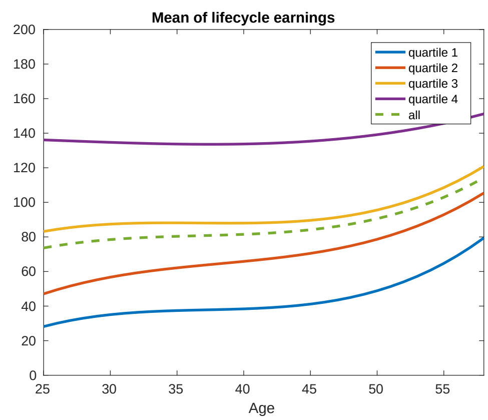

(c) Time invested in human capital

Figure 12 reports stock market participation and the time allocated to human capital investment by quartiles of initial human capital. As predicted above, and as seen in Figure 12b, time allocation as a function of initial human capital is inversely proportional to its initial level: those in quartile 1 (the lowest level of initial human capital) invest the most time, while those in the highest quartile invest the least. The intuition is natural. Those with high initial human capital face not only a high opportunity cost of additional accumulation, but also stand to reap only low marginal returns. The reverse holds for those with low initial human capital. Note that initial differences in human capital levels persist over time, although with some "catch-up" due to those with low initial human capital allocating higher amounts of time toward its accumulation. The result is the path of earnings observed in Figure 12c.

What does this imply for the accompanying investment that households make in the stock market? Those with the highest levels of initial human capital (quartile 4) participate in the stock market at the highest rates, while those with the lowest levels participate at by far the lowest rates. Specifically, participation within the top quartile is about 70 percent at age 25 and reaches 100 percent participation by age 50 (Figure 12a). Quartiles 2 and 3 participate at around a 30 percent rate early in life, and reach 100 percent participation after age 55. For the lowest quartile, participation starts at around 15 percent and remains below 50 percent throughout working life.

To recap, stock market behavior in this case is influenced by two forces. First, households with high initial human capital not only have relatively high earnings, but also do not expect earnings to rise as rapidly over the life cycle as those with low initial human capital do. As a result, their motivation to borrow early in life is limited, and the same force that leads to relatively little time spent accumulating human capital encourages stock market participation. In other words, the optimal overall portfolio for those with high initial human capital reflects the relative value of savings, even early in life, and this leads to a relatively high rate of equity market participation. By contrast, those with low human capital find it to be a far better investment than stocks and, moreover, expect future earnings to be higher than present levels. Higher expected future earnings make savings less attractive, as that would hinder the intertemporal smoothing of consumption. Indeed some of these households would value borrowing (or, at the very least, not accumulating wealth). Thus, saving via *any* financial asset, especially risky stocks, is less attractive. Additionally, individuals in the lowest quartile also earn the least of all groups, and hence face significant uninsurable risk to consumption arising from stochastic variability in the payoffs on any assets they hold, especially early in life. Thus, the riskiness of equity makes such investment unattractive for such individuals. For households in the middle quartiles of initial human capital, optimal investment behavior falls between these two extremes.

6.3.2 Learning Ability

Figure 13: Investment by Quartiles of Ability

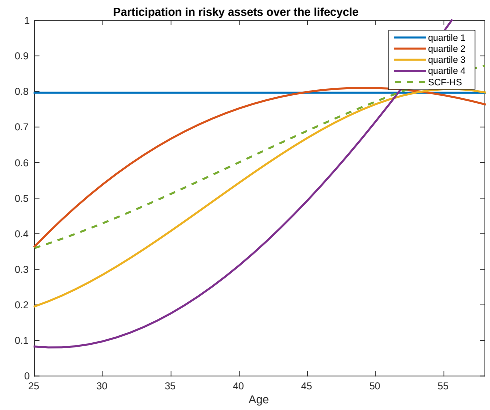

(a) Participation

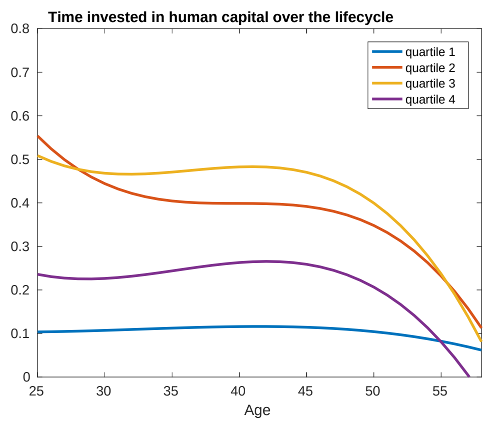

(b) Time invested in human capital

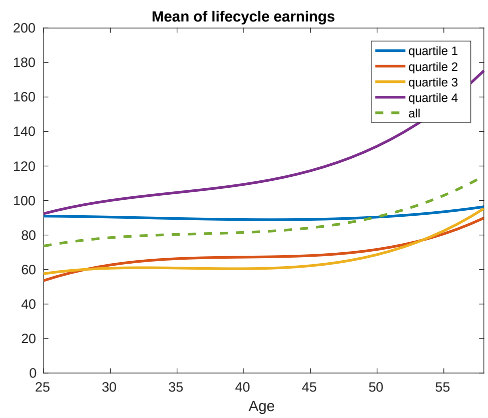

(c) Earnings

We turn next to the other dimension of "initial conditions": learning ability. Figure 13 shows participation and human capital investment behavior by quartiles of ability levels, with quartile 1 being the lowest. Precisely as predicted above, we see that agents with high ability accumulate human capital more rapidly than agents with low ability. This is driven by the fact that investing time in human capital is more productive for these agents, which increases their incentive to do so. Of course, these agents do not have to invest as much time to accumulate the same amount of human capital as those with lower ability, and as a result, will be able to enter retirement with a given wealth level with less effort by virtue of their greater earnings capacity. These two forces work in opposite directions, with the result that we observe that agents in the middle two quartiles invest the most time in human capital investment, especially early in life

(Figure 13b). Agents in the lowest quartile of ability invest the least time in human capital accumulation, and their time investment remains relatively flat over the life cycle.

When it comes to the main question of interest to us, namely, stock-market participation, we turn to Figure 13a. Recall that in the baseline model, a lower time investment in human capital is associated with a higher stock-market participation rate. This is seen in stark terms here: the lowest quartile participates at extremely high rates (80 percent). The intuition is simply that for low-ability households, the effective rate of return from human capital is much lower than from equity investment. Further, their earnings profile is relatively flat (Figure 13c), which means that their participation rate also remains flat over the life cycle. In contrast, the high initial investment in human capital, particularly for quartiles 2 and 3, and the steeper earnings profile, particularly for quartile 4, is associated with these groups exhibiting a steeply increasing stock market participation rate over the life cycle. For these households, learning, especially when young, is a better investment than earning and investing in equities.

We remind the reader that the preceding figures hold the correlation between initial human capital and learning ability at zero. In the data, participation rates presumably depend on the joint distribution of ability and initial human capital. Indeed, as clarified above, these characteristics are positively correlated in the baseline model. Thus, those who face high costs of learning—and hence wish to invest primarily in stocks—are frequently also those with low initial human capital, and who therefore wish to invest in human capital instead. The net result is that participation rates in the baseline model fall in between the levels implied by Figures 12 and 13.

Nonetheless, readers might be concerned that the model's implication that learning ability should be inversely related to equity investment is counterfactual. While it is true that, all else equal, high-ability households would participate at lower rates than low-ability households early in life, this relationship reverses later in life. Another way to look at this is to compare ability distributions across participants and nonparticipants. In Figure 14, we show the results of this comparison at various ages. Consistent with our message that the presence of a high-return alternative deters stock-market participation, we see that in the first two panels of Figure 14, when households are young, nonparticipants have substantially higher ability levels than stock-market participants. However, as agents age, as seen in Figures 14b and 14c, ability is similarly distributed across stock-market participants and non-participants, with participants now being of slightly higher ability. By middle age, marginal returns to human capital are no longer substantially higher than the returns on stocks for even those with high learning ability.

Figure 14: Ability Distribution of Participants and Nonparticipants

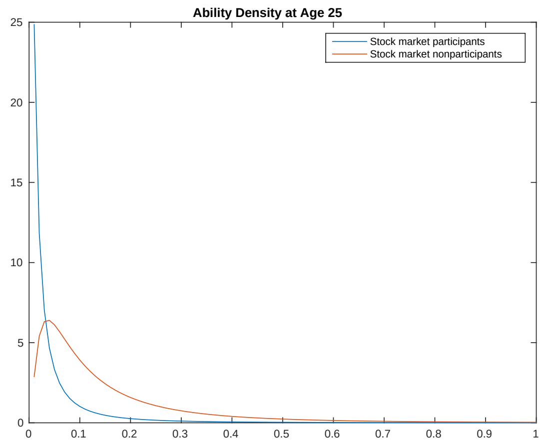

(a) Age 25

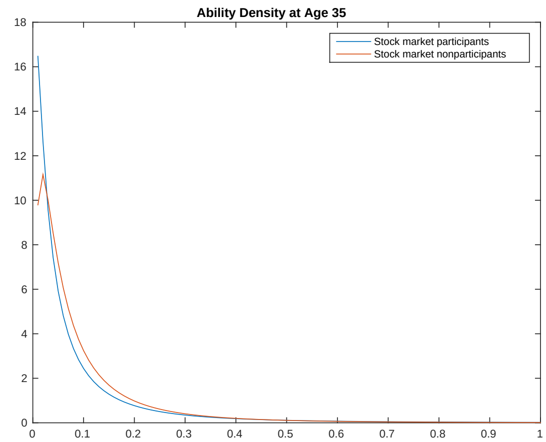

(b) Age 35

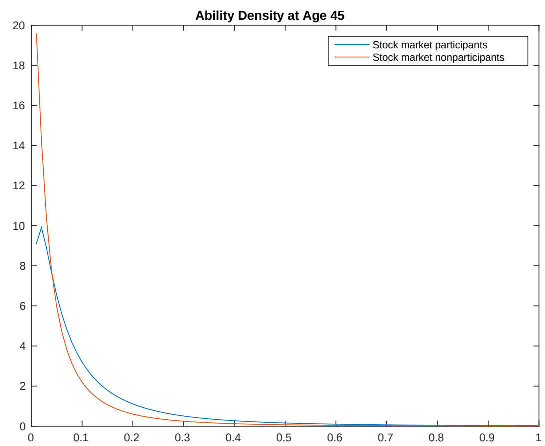

(c) Age 45

To point out that the differences in participation are not being driven by any resulting differences in wealth, we also look at households with high initial wealth, defined here as being in the top 10 percent of the wealth distribution at age 25. Figure 15 shows clearly the central mechanism that we have emphasized: within the group of households with similar ability (Figure 15a), it is precisely those with low initial human capital who elect not to participate in the stock market (Figure 15b).

Figure 15: Distribution of Ability and Human Capital across Participants and Non-Participants (Wealthy Households at Age 25)

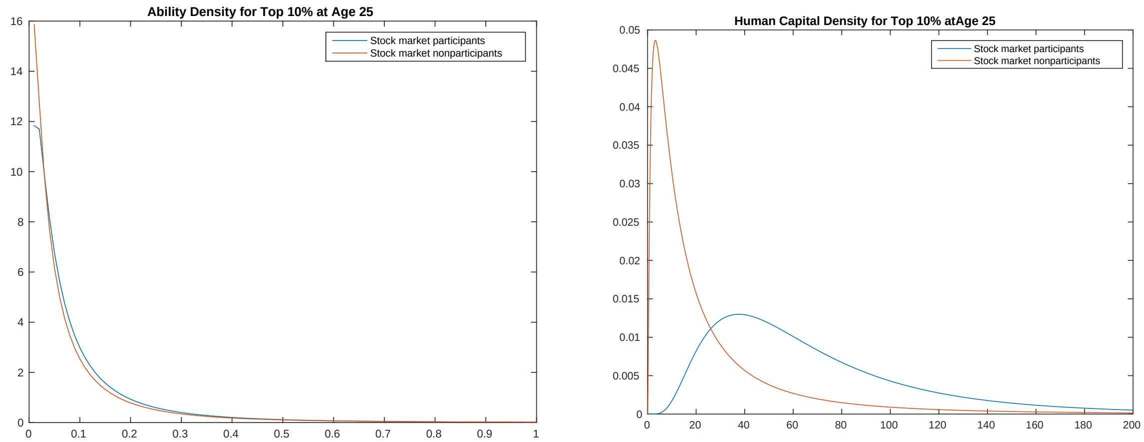

(a) Ability

(b) Human Capital

In other words, nonparticipation among those with high ability is well-understood once we look at initial human capital levels within this group. The set of nonparticipants is compsed disproportionately of those with relatively low initial human capital. The combination of high ability and low initial human capital makes human capital investment particularly attractive since both ensure high marginal returns. As a result, when young, these households forgo stock market participation in favor of human capital investment.

6.4 The Role of Borrowing Costs

It is natural to ask whether nonparticipation in our model is driven by the presence of a borrowing wedge rather than dispersion in human capital returns. We demonstrate that while borrowing costs are critical to obtaining nonparticipation in an exogenous-earnings setting, they have little effect on stock-market participation when earnings are endogenized through human capital investment. To illustrate this, we consider a case in which there is no wedge at all between the interest rate on savings and borrowing. Figure 16, which reports the results from our benchmark model with no wedge, shows that households do not significantly change their stock market participation despite having access to cheaper credit. In contrast, we know from Davis, Kubler, and Willen (2006) that, in a setting where earnings are exogenous, stock market participation reaches nearly 100 percent early in the life cycle in the absence of a wedge.

What accounts for the differential impact of borrowing costs in the two settings? The answer lies in the presence of the option to invest in human capital. When this option is not available to the household, as is the case in the exogenous earnings setting, the household must only decide whether it makes sense to borrow to invest in stocks. As long as the expected return to the investment in stocks exceeds the cost of borrowing, households will choose to do so. Thus participation will be high in the absence of a sufficiently high interest rate on borrowing.

Figure 16: The Role of the Borrowing Wedge in Stock Market Participation

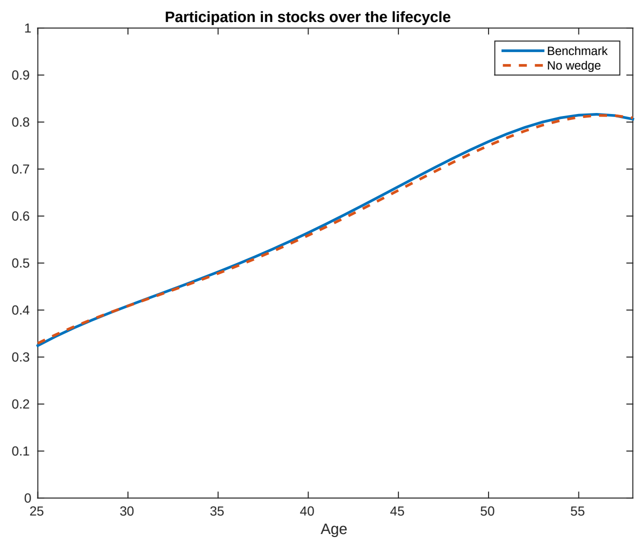

In contrast, when the household has the option to invest in human capital, it has two choices for what to do with any money it borrows—it can use the resources to buy stocks or to fund consumption while investing in human capital. The return to the former is the expected market rate of return on stocks, while to the latter is expected higher future earnings. All else equal, as long as the growth in future earnings that comes from spending additional time learning exceeds the rate of return on stocks, the household will invest borrowed money to fund consumption while spending more time learning, not to buy stocks. Thus, in this case, the relevant comparison is between the rate of return to human capital investment and the rate of return to stocks, and *not* between the rate of return to stocks and the rate of interest on borrowing. In other words, for those facing high enough marginal returns to human capital, stocks are unattractive and the borrowing wedge is essentially irrelevant to this comparison (up to the secondary indirect effect arising from the disincentive on human capital coming from the increased cost of bringing forward future earnings to the present).

Nonetheless, the wedge *is* relevant for the human capital investment decision. This is because, in the absence of the wedge, borrowing costs fall, so households who borrow spend additional time on human capital accumulation, as Figure 17b shows.[18]

Figure 17: The Effect of No Borrowing Wedge on Time Allocated to Human Capital

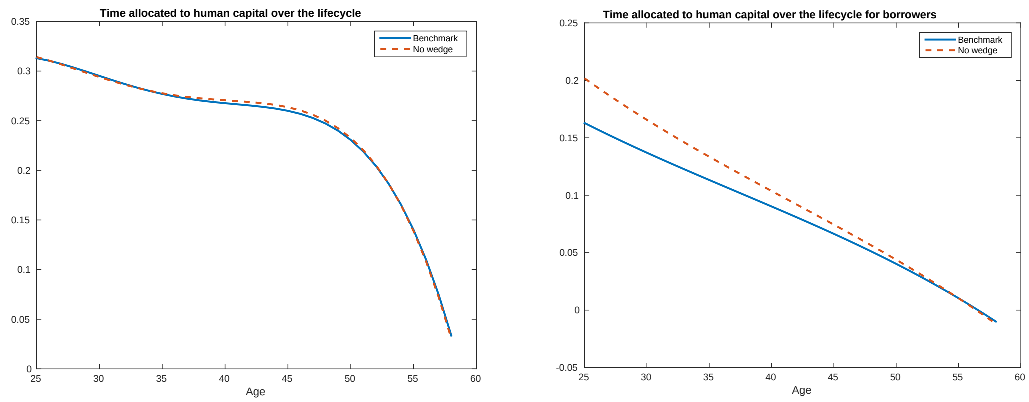

(a) All Households (b) Borrowers

6.5 The Importance of Diminishing Returns to Human Capital Investment

Figure 18: Comparison of Marginal Densities in Model with $\alpha = 0.7$ and $\alpha = 0.9$

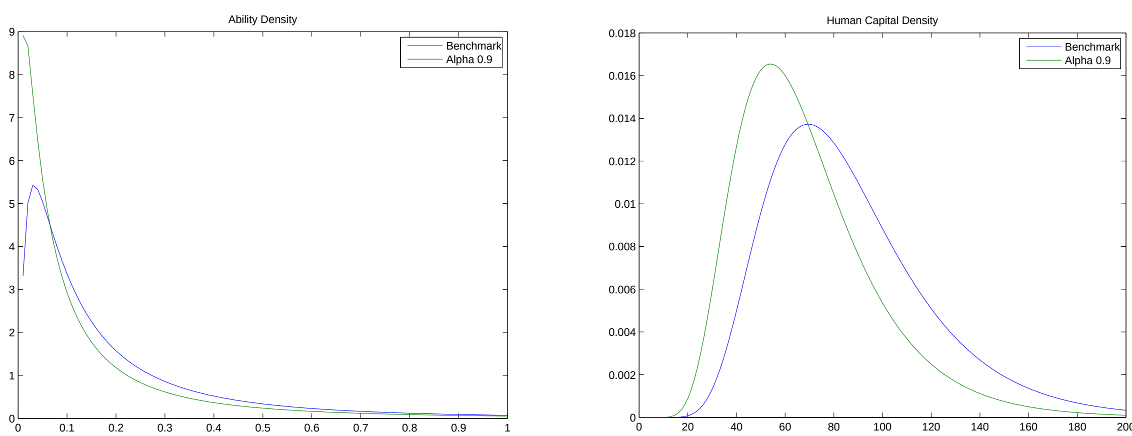

(a) Ability (b) Initial Human Capital

[18]Time spent on human capital investment does not change much relative to the benchmark because there are more borrowers when borrowing costs are lower and these individuals typically spend less time on human capital investment.

A clear implication of the logic of our model is that the better the technology for learning, the less attractive stock market investment will be, all else equal. After all, if the earnings that we observe in the data were generated by a more productive human capital technology than in the benchmark, then we should expect to see lower participation in the stock market than in the benchmark. To illustrate this, consider a case in which the human capital technology is extremely productive: $\alpha = 0.9$.[19] To preserve comparability, we recalibrate all the parameters needed to match earnings facts as in the benchmark. The marginal densities for ability and initial human capital obtained from the recalibration are to the left of those in the benchmark (Figure 18).

Figure 19: Results with $\alpha = 0.9$ in Recalibrated Model

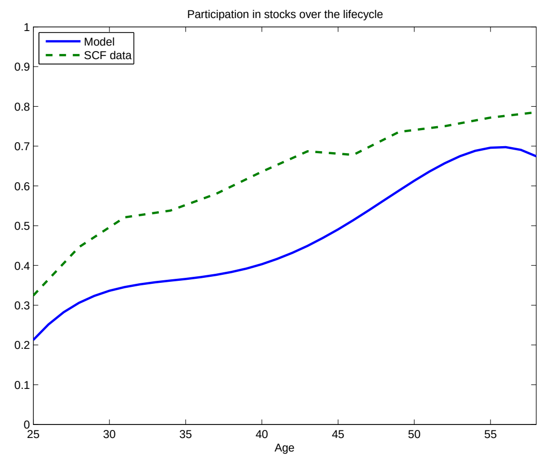

The main results are reported in Figure 19. Participation in the stock market is indeed much lower than in the benchmark. This is for two reasons. First, as we have stressed throughout, this is consistent with the idea that human capital competes with financial assets as an investment option. With a high α, human capital competes favorably for longer because households encounter marginal returns to human capital investment that diminish more slowly than in the benchmark model. Second, households in this model start life with lower initial human capital levels on average relative to the benchmark. As a result, more households choose to forgo participation in the stock market in favor of human capital accumulation.

[19]The literature provides a range of estimates for this parameter (Browning, Hansen, and Heckman, 1999). While this example reinforces one of the main mechanisms underlying our results, it is important to note that a value of $\alpha = 0.9$ is at the high end of these estimates in the literature and hence has less empirical plausibility.

Our model can also shed light on the effects of exogenous or policy-induced changes in the learning technology. How would households in our model respond if they were to be confronted with a change in the productivity of the learning technology? We address this case by considering the effect of decreasing the value of α to 0.5 and increasing it to 0.9 *without recalibrating the model*. In other words, it is as if households with initial conditions as in our model were suddenly faced with a more productive or less productive human capital investment technology. The results are reported in Figure 20.

Figure 20: The Effect of the Elasticity of Human Capital Production on Investments

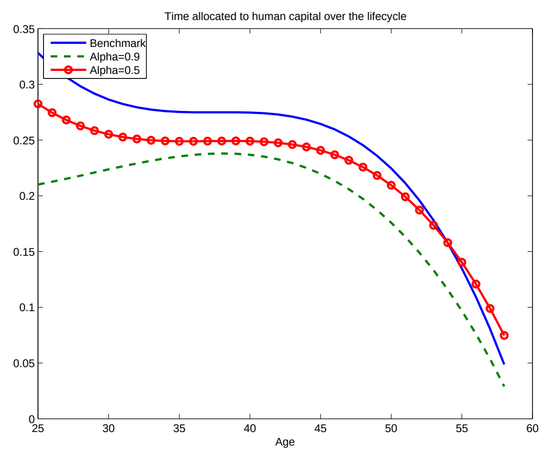

(a) Time Invested in Human Capital

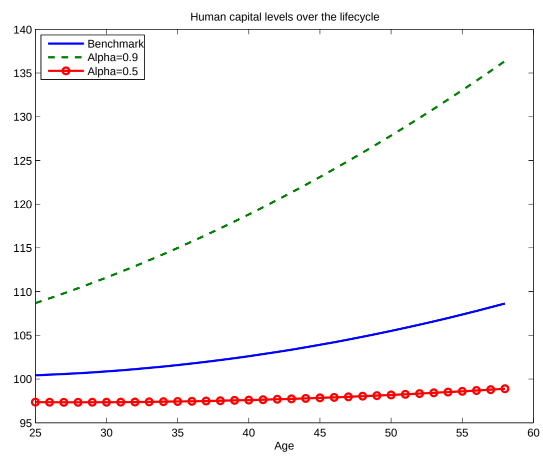

(b) Human Capital

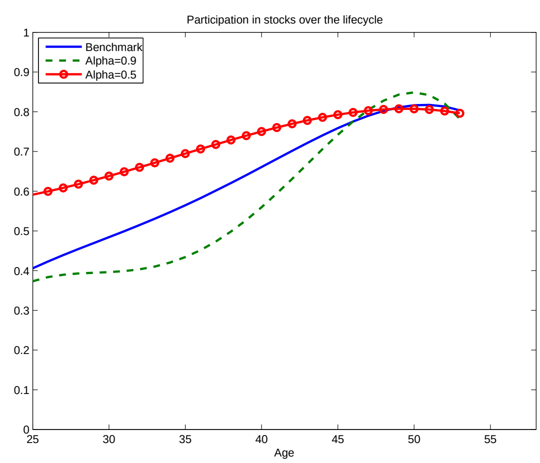

(c) Participation

First, consider the case where the human capital technology is less productive (α=0.5). Two opposing forces are at work here. On the one hand, because human capital is less productive, agents have less incentive to invest time in it. On the other, to the extent that agents do want to accumulate human capital, they need to invest more time to accumulate the same level of human capital as in the benchmark. It turns out that the first effect dominates; agents invest less time in human capital than in the benchmark, as Figure 20a shows, with the effect that their human capital levels are lower throughout working life than in the benchmark (Figure 20b). This has two effects on participation. Less time invested in human capital leads to higher stock participation early in life, while the slower growth rate of human capital over the life cycle (which translates into a flatter path for earnings) leads to a flatter profile of participation over the life cycle (Figure 20c).

In the case where the human capital technology is more productive (α=0.9), the two opposing forces described earlier also lead agents to invest less time in human capital accumulation. Despite this, their human capital levels are higher and increasing much more steeply than in the benchmark. The participation rate in the stock market is lower early in life but rises steeply to move past the rate observed in the benchmark by age 50.

This experiment reveals a more general mechanism that is at work in our model. Agents have two ways to move resources through time—using financial assets or human capital. The more human capital pays off in the future, the steeper the earnings profile and the higher the incentive to invest in human capital now. If agents can use financial assets to bring some of those future earnings into the present to smooth consumption, they will, with the result that they do not invest in stocks early in life and instead borrow to the extent possible. On the other hand, if earnings are going to be flat, or if agents don't expect high returns to human capital in the future, they will enter financial markets early. The findings are similar if we change the growth rate of the rental rate of human capital, g (results available upon request).

A common theme that emerges from the results described above is that higher human capital accumulation, if achieved through an improvement in its production technology, leads to an increase in earnings *and* stock market participation. In these instances, the agent accumulates more human capital without necessarily allocating additional time to it. On the other hand, any increase in human capital that comes from households allocating more time to human capital investment leads to lower stock-market participation.

6.6 The Role of Risk

Thus far, we have almost exclusively emphasized intertemporal tradeoffs as a key consideration in explaining both human capital investment and stock market participation. Our baseline model, however, builds in risk in both human capital and stock market returns. Our incorporation of risk was driven both by the clear consensus within the literature in favor of its presence and its essentiality in delivering observed heterogeneity in earnings and wealth. We now demonstrate that risk, while relevant for disciplining the parameters of the model, especially human capital, is not central to the question of stock-market participation. We begin by studying the effect of changing the risk of stocks and the effect of changing agents' risk aversion in our setting.

6.6.1 Stock-Market Risk

The stock market, while it clearly offers a far higher average rate of return than risk-free savings, may still not attract overwhelming participation due to the exposure that it creates for households. To study the effect of the risk properties of stock returns on participation, we examine two cases in which equity market risk is different than in the baseline model. In Figure 21, we report results under the assumptions that the standard deviation of stock market returns is low (50 percent less) or high (50 percent more) compared to our benchmark (0.078 and 0.236, respectively). Interestingly, these large differences in the risk properties of stocks have almost no effect on participation compared to the benchmark. This is because the participation decision is affected (especially when borrowing is costly) by the household's decision to hold net positive wealth in the first place. And this stance is, in turn, driven not only by the intertemporal profile of expected earnings (e.g., the steeper the earnings, the less the desire to hold positive net wealth when young), but also by precautionary saving induced by risk. Thus, even when stocks are risky, or especially if they are, households may save more at all dates, with some of those savings being channeled to stocks.

Figure 21: Stock Market Participation with Low and High Risk of Stocks

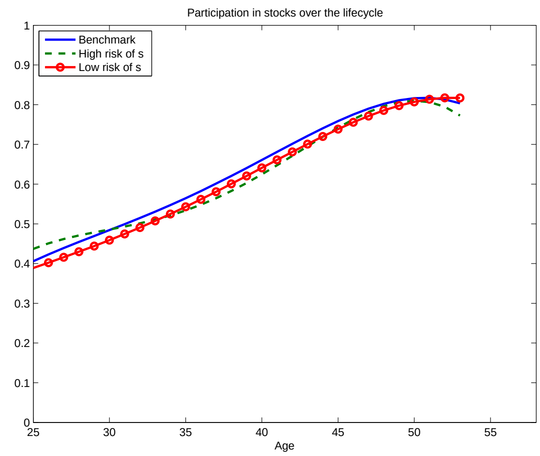

6.6.2 Agents' Risk Aversion

Having seen that risk per se is not a powerful determinant of stock market participation, one might expect that attitudes to risk do not much matter either. This intuition is borne out below. We consider two cases, $\sigma = 3$ and $\sigma = 10$. The results are shown in Figure 22.

Figure 22: Effect of Changing Risk Aversion on Stock-Market Participation

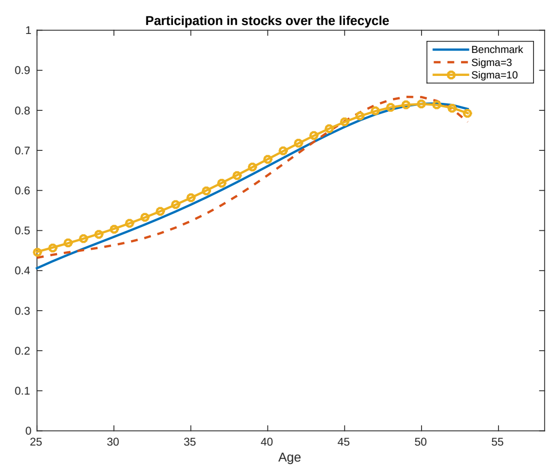

As seen clearly in the figures, the effect of changing risk aversion is qualitatively similar to changing the riskiness of stock returns, in the sense that it does not have much effect on stock market participation in the economy. One useful implication of these results is that while we have employed a risk-aversion value that is standard in the portfolio-choice literature (e.g., it is higher than the value typically assumed in macroeconomics, which ranges from 1 to 3 for example), stock-market participation is not especially sensitive to risk aversion. While primarily suggestive, as we do not recalibrate the entire model when we change risk aversion, it is consistent with the intertemporal motives we emphasize being critical determinants of the participation decision.

6.7 Participation and Savings

In gauging the model's success overall, it is useful to understand the extent to which its predictions for stock participation across savers, borrowers, and nonparticipants in the riskless asset market are borne out. As we see from Figure 23, both in the data and in the model, conditional on holding strictly positive savings levels of the risk-free asset, participation rates are high and increasing over the life cycle. Conditional on holding negative levels of the risk-free asset, participation rates are lower and hump-shaped. Relative to the data, however, the model overpredicts the participation rate of savers, underpredicts that of borrowers, and is close for the nontrivial mass of households with zero (liquid) wealth holdings. This shortcoming of the model is primarily attributable to the relatively simple structure of the liabilities and assets available to households in the model. For one, households cannot simultaneously borrow and hold positive amounts of the risk-free asset in the model while, in the data, households simultaneously hold liquid risk-free assets and credit card debt. We do observe in the data that risk-free balances are much lower for nonparticipants than for participants, and it is possible that we are classifying as "savers" those who are simply holding transactional balances.[20] Further, there is undoubtedly heterogeneity in borrowing costs across households in the data, with some able to access borrowing for lower rates than the near-credit-rates we assume are available to all following the literature. Future work may be able to improve on this aspect of the model's performance, but our conjecture is that it will require a substantially richer menu of risk-free assets and liabilities. As our focus is on understanding the role of human capital investment in overall financial investment, this is something beyond the scope of the current paper.

[20]The numbers are available upon request.

Figure 23: Participation among Borrowers and Savers: Model vs. Data

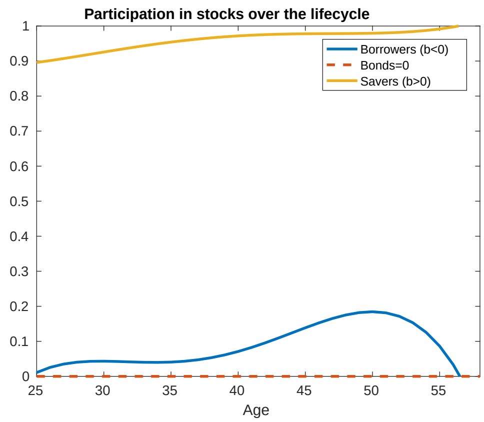

(a) Model

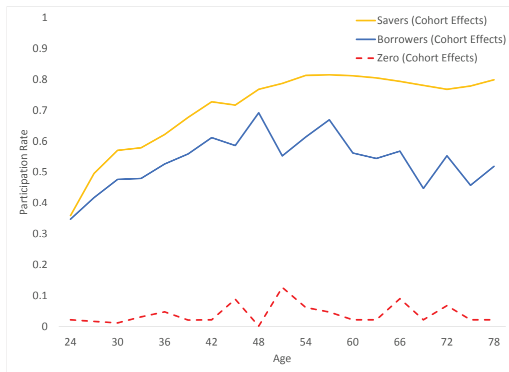

(b) Data

6.8 Understanding the Share of Stocks in Household Portfolios

Our goal from the outset has been to understand the role of human capital investment in stock-market participation. While not our principal focus, total investment in risky assets is of course of significant interest. We therefore turn now to the "intensive" margin of stock-market investment. As with participation, these data are also in no way targeted. As a result, a comparison of the model's implications to data on shares provides a useful view of the model's performance. As we will show, the fundamental lesson of the model is that the forces that determine participation are separable from those determining the proportion of wealth invested in stocks. In particular, the forces essential for explaining the behavior of shares of wealth held in stocks—risk and attitudes toward risk—have little overlap with the forces that explain participation. In one sense, this may be natural: agents in the model always have access to a safe asset to move wealth through time. Second, the investment horizon for those with significant life-cycle wealth is short (as wealth is accumulated in substantial amounts only in middle age and beyond); this means that the power of interest rates to dramatically alter the attractiveness of stocks is limited. This leaves risk as a key determinant of household decisions—especially in a setting where human capital also carries risk. While future work that better identifies the risk characteristics of equity investment (and household attitudes to risk) will allow the model to capture both participation and the intensive margin of stock-market investment, it is clear that one can approach the extensive and intensive margins of stock-market investment separately.

With respect to the share of wealth invested in stocks, we see in Figure 24 that three things are salient. First, the model implies a higher share for wealth held in equity than in our SCF data early in life, but this gap closes later in life. This is important because, in the model, as in the data, the bulk of financial wealth is accumulated late in life. As a result, our model accounts well for the share of wealth allocated to equity during the part of life in which financial wealth is largest. Second, we see that the share of wealth held in stocks in the presence of human capital remains far below 100 percent. Importantly, this occurs despite the fact that households in our model retain the ability to increase their labor supply to undo poor stock market returns. Third, the hump-shaped profile for shares generated by our model is empirically more plausible than the decreasing profile derived by much of the existing work. This is true irrespective of whether time or cohort effects are used to identify the path of shares, with model and data being closest for the case in which time effects are assumed to matter. Moreover, if we were to abstract from time and cohort effects altogether, as in Gomes and Michaelides (2005), our model's predictions for shares would be very close to the data. An interesting implication of our model is that the conventional "100 minus age" rule of thumb often prescribed in financial planning circles, and often not followed by households in the data, may not be optimal in settings where investment in human capital is an option.

Figure 24: Fraction of Stocks in Household Portfolio

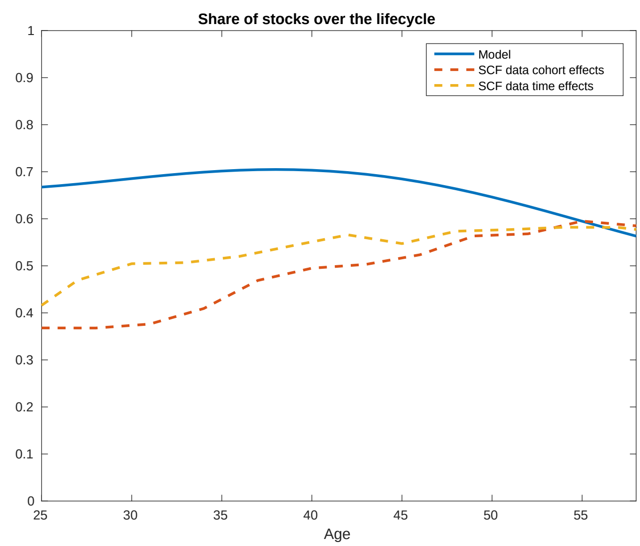

What is the role played by earnings heterogeneity for the share of wealth held in stocks by those who participate? As we did for participation earlier, we can run an experiment with no heterogeneity in ability and initial human capital to get at this. What emerges is that the proportion of wealth held in stocks is not sensitive to capturing earnings heterogeneity. As seen in Figure 25, shares in this experiment are very similar to the benchmark. This is intuitive: While participation decisions are clearly dependent on the path and marginal returns to human capital, conditional on saving, the risk-allocation problem of households does not differ in a substantive manner.

Figure 25: Life-Cycle Stock-Market Shares with No Heterogeneity in Ability and Initial Human Capital

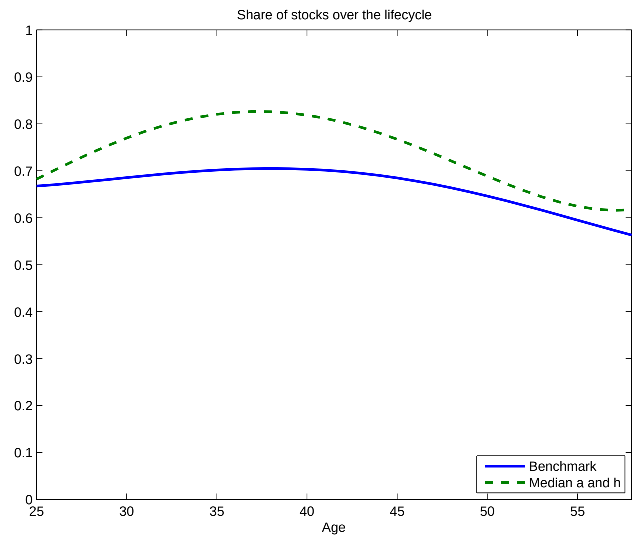

Having shown earlier that endogenous human capital dramatically limited the role of borrowing costs for stock-market participation, it is of interest to see if this applies to the intensive margin as well. The answer is no. The reason is this: Given participation, the question for a household is the extent of risk they wish to bear, and there is little reason to think that the cost of borrowing alters the willingness to bear risk in a first-order manner. After all, as we know from Davis, Kubler, and Willen (2006), borrowing costs limit demand for stocks altogether by making its mean return lower. This is driven home by the fact that in the benchmark, we see that borrowing costs have almost no effect on the risk exposure that households choose (Figure 26).

Figure 26: The Role of the Borrowing Wedge in Stock-Market Shares

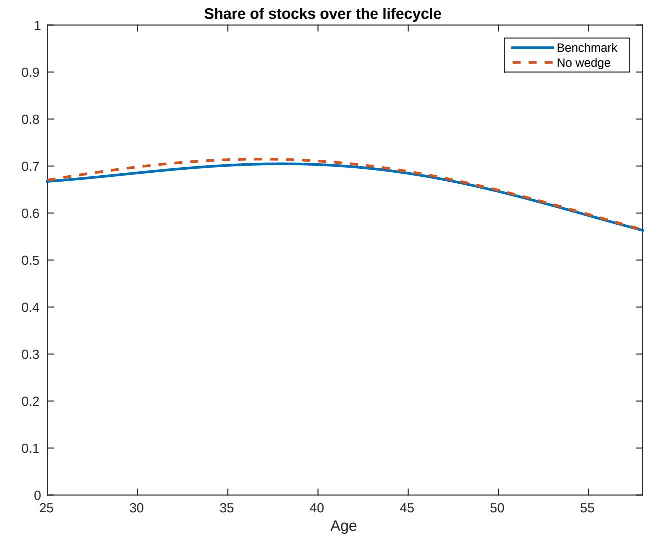

Having asserted that risk considerations are critical for explaining shares of wealth held in stocks, we can be more explicit. In Figure 27, we see that when stocks are risky, households who engage in the stock market reduce their holdings at all ages. In the case of higher-than-baseline riskiness of stock return, we find that household diversification plays a significant role and leads to much lower proportions of wealth held in stocks than in the baseline. Conversely, we observe that when stock-market risk is cut, wealth shares balloon to nearly 80 percent when averaged over the life cycle.

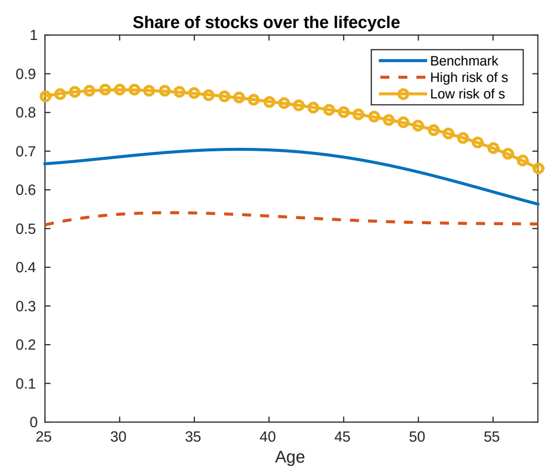

Figure 27: Shares under Higher or Lower Risk of Stocks

If risk-related considerations loom large in determining the exposure chosen by stock market participants, as seems entirely intuitive, risk aversion will matter importantly for the wealth share. As seen in Figure 28, this is exactly what happens.

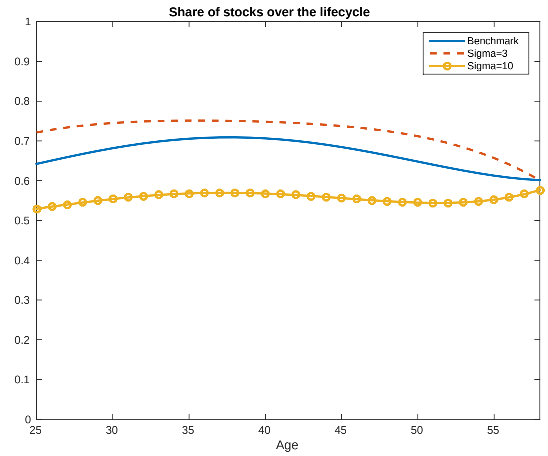

Figure 28: Shares under High and Low Risk Aversion

Thus, an interesting implication of our analysis is that while initial human capital levels and ability govern the decision to invest *at all* in the stock market, the risk of stocks is what matters for the share of wealth held in equity. Our approach has been to understand the role, in household investment, played by the availability of human capital as a special kind of investment, in a setting that is as standard as possible.

7 Conclusion

The contribution of this paper is to show that human capital investment matters for the path of stock-market participation. We have demonstrated that once human capital investment is allowed for and, critically, disciplined to match empirical labor income dispersion, an entirely standard model of portfolio choice delivers observed stock market participation rates over the entire life cycle.

Our approach is both novel and straightforward: we embed the classic human capital model of Ben-Porath (1967) into a standard life-cycle model of portfolio choice where households face uninsurable idiosyncratic shocks to productivity (e.g., Cocco, Gomes, and Maenhout, 2005). Importantly, as in Huggett, Ventura, and Yaron (2006), households in our model are heterogeneous with respect to characteristics governing initial human capital and their ability to acquire it.

Our findings flow from two simple and intuitive mechanisms: First, the returns to human capital investment are highest early in life and exceed the constant returns on financial assets for most households. As households age, this relationship reverses. At the same time, even if borrowing is allowed, households will elect to use it to offset the loss in current earnings entailed by human capital investment, not to invest in stocks. Thus, stock-market participation starts low and grows over the life cycle, just as in the data.

Our model generates empirically-accurate accumulation of both risky and risk-free assets despite not targeting these data. Our model also has reasonable, though not fully satisfactory, implications for the share of wealth held in stocks. One aspect of our findings is that the forces governing participation, which are well-captured by our allowance for human capital investment by households, are different from those governing the share of wealth held in equity. The former are predominantly intertemporal in nature, while the latter are driven by risk and attitude to risk.

Taken as a whole, our work suggests that human capital investment, and the dispersion in its returns across households, may be playing a substantial role in observed stock-market participation over the life cycle.

References

ABBOTT, B., G. GALLIPOLI, C. MEGHIR, AND G. L. VIOLANTE. 2013. "Education Policy and Intergenerational Transfers in Equilibrium." National Bureau of Economic Research Working Paper 18782.

AMERIKS, J., AND S. ZELDES. 2004. "How Do Household Portfolio Shares Vary with Age?" Columbia University Working Paper.

ATHREYA, K. B.. 2008. "Default, Insurance, and Debt over the Life Cycle." *Journal of Monetary Economics* 55(4): 752–774.

BEN-PORATH, Y.. 1967. "The Production of Human Capital and the Life-cycle of Earnings." *Journal of Political Economy* 75(4): 352–365.

BENZONI, L., P. COLLIN-DUFRESNE, AND R. GOLDSTEIN. 2007. "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated." *The Journal of Finance* 62(5): 2123–2167.

BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM. 2014. "Federal Reserve Statistical Release, G19, Consumer Credit." Board of Governors of the Federal Reserve System.

BRITO, N. O.. 1978. "Portfolio Selection in an Economy with Marketability and Short Sales Restrictions." *The Journal of Finance*, 33(2): 589–601.

BROWNING, M., L. HANSEN, AND J. HECKMAN. 1999. In "Micro data and general equilibrium models." *Handbook of Macroeconomics*, Vol. 1, Part A, ed. by John B Taylor and Michael Woodford. North Holland: Elsevier.

CAMPBELL, J. Y., J. F. COCCO, F. J. GOMES, AND P. J. MAENHOUT. 2001. "Investing Retirement Wealth: A Life-Cycle Model." In *Risk Aspects of Investment-Based Social Security Reform*, ed by John Y Campbell and Martin Feldstein, 439–482. Chicago: University of Chicago Press.

CHAI, J., W. HORNEFF, R. MAURER, AND O. S. MITCHELL. 2011. "Optimal portfolio choice over the life cycle with flexible work, endogenous retirement, and lifetime payouts." *Review of Finance* 15(4): 875–907.

Chang, Y., J. H. Hong, and M. Karabarbounis . 2014. "Life Cycle Uncertainty and Portfolio Choice Puzzles." Federal Reserve Bank of Richmond Working Paper 14-13R.

Cocco, J.. 2005. "Portfolio Choice in the Presence of Housing." *Review of Financial Studies* 18(2): 535–567.

Cocco, J., F. Gomes, and P. Maenhout. 2005. "Consumption and Portfolio Choice over the Life Cycle." *Review of Financial Studies* 18(2): 491–533.

Constantinides, G. M., J. B. Donaldson, and R. Mehra. 2002. "Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle." *Quarterly Journal of Economics* 117(1): 269–296.

Davis, S. J., F. Kubler, and P. Willen. 2006. "Borrowing Costs and the Demand for Equity over the Life Cycle." *Review of Economics and Statistics* 88(2): 348–362.

Davis, S. J., and P. Willen. 2013. "Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice." *Quarterly Journal of Finance* 03(03n04): 1350011.

Deaton, A.. 1985. "Panel Data from Time Series of Cross-sections." *Journal of Econometrics* 30(1): 109–126.

Gomes, F., and A. Michaelides. 2003. "Portfolio Choice with Internal Habit Formation: A Life-cycle Model with Uninsurable Labor Income Risk." *Review of Economic Dynamics* 6(4): 729–766.

——— . 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *The Journal of Finance* 60(2): 869–904.

Gomes, F. J., L. J. Kotlikoff, and L. M. Viceira. 2008. "Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds." *American Economic Review, Papers and Proceedings of the Annual Meeting* 98(2): 297-303. American Economic Association.

Gourinchas, P., and J. Parker. 2002. "Consumption Over the Life Cycle." *Econometrica* 70(1): 47–89.

Haliassos, M., and A. Michaelides. 2003. "Portfolio Choice and Liquidity Constraints." *International Economic Review*, 44(1): 143–177.

Heaton, J., and D. Lucas. 1997. "Market Frictions, Savings Behavior, and Portfolio Choice" *Macroeconomic Dynamics* 1: 76–101.

Hubbard, R. G., J. Skinner, and S. P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103(2): 360–399.

Huggett, M., and G. Kaplan. 2011. "Human capital values and returns: Bounds implied by earnings and asset returns data." *Journal of Economic Theory* 146(3): 897–919.

Huggett, M., and G. Kaplan. 2015. "How Large is the Stock Component of Human Capital?" Georgetown University Working Paper.

Huggett, M., G. Ventura, and A. Yaron. 2006. "Human Capital and Earnings Distribution Dynamics." *Journal of Monetary Economics* 53(2): 265–290.

Huggett, M., G. Ventura, and A. Yaron. 2011. "Sources of Lifetime Inequality." *The American Economic Review* 101(7): 2923–2954.

Ionescu, F. (2009): "The Federal Student Loan Program: Quantitative Implications for College Enrollment and Default Rates," *Review of Economic Dynamics* 12(1): 205–231.

Jagannathan, R., and N. R. Kocherlakota. 1996. "Why Should Older People Invest Less in Stocks Than Younger People?" *Federal Reserve Bank of Minneapolis Quarterly Review* 20(3): 11–23.

Khorunzhina, N.. 2013. "Structural Estimation of Stock Market Participation Costs." *Journal of Economic Dynamics and Control*, 37(12): 2928–2942.

Kim, H. H., R. Maurer, and O. S. Mitchell. 2016. "Time is money: Rational life cycle inertia and the delegation of investment management." *Journal of Financial Economics* forthcoming.

Lindset, S., and E. Matsen. 2011. "Human Capital Investment and Optimal Portfolio Choice." *The European Journal of Finance* 17(7): 539–552.

Lustig, H., and S. Van Nieuwerburgh. 2008. "The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street." *Review of Financial Studies* 21(5): 2097–2137.

McGrattan, E., and E. Prescott. 2000. "Is the Stock Market Overvalued?." *Federal Reserve Bank of Minneapolis Quarterly Review* 24(4): 20-40.

Mehra, R., and E. C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15(2): 145–161.

Merton, R.. 1969. "Lifetime Portfolio Selection under Uuncertainty: The Continuous-time Case." *The Review of Economics and Statistics*, 51(3): 247–257.

Polkovnichenko, V.. 2007."Life-Cycle Portfolio Choice with Additive Habit Formation Preferences and Uninsurable Labor Income Risk." *Review of Financial Studies* 20(1): 83–124.

Poterba, J. M., and A. A. Samwick. 1997. "Household Portfolio Allocation over the Life Cycle." National Bureau of Economic Research Working Paper 6185.

Roussanov, N.. 2010. "Human Capital Investment and Portfolio Choice over the Life Cycle." University of Pennsylvania Wharton School Working Paper.

Samuelson, P.. 1969. "Lifetime Portfolio Selection by Dynamic Stochastic Programming." *The Review of Economics and Statistics*, 51(3): 239–246.

Viceira, L. M.. 2001. "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income." *The Journal of Finance*, 56(2): 433–470.

A Regression Tables

Table 2: Probit for Stock Market Participation with Cohort Effects (SCF), N=34,008

Age	Coefficient	Cohort	Coefficient
23-25	(omitted)	1919-1921	-0.9716
26-28	0.3195	1922-1924	-1.0055
29-31	0.5079	1925-1927	-0.7505
32-34	0.5510	1928-1930	-0.6046
35-37	0.6580	1931-1933	-0.7356
38-40	0.8026	1934-1936	-0.6558
41-43	0.9430	1937-1939	-0.5859
44-46	0.9177	1940-1942	-0.5368
47-49	1.0862	1943-1945	-0.5006
50-52	1.1310	1946-1948	-0.3663
53-55	1.2002	1949-1951	-0.4259
56-58	1.2459	1952-1954	-0.3639
59-61	1.2166	1955-1957	-0.3494
62-64	1.1894	1958-1960	-0.3038
65-67	1.1660	1961-1963	-0.1609
68-70	1.1346	1964-1966	-0.1800
71-73	1.1051	1967-1969	-0.0860
74-76	1.1265	1970-1972	-0.0062
77-79	1.2015	1973-1975	(omitted)
		1976-1978	0.0339
		1979-1981	0.0143
		1982-1984	-0.0091
		1985-1987	0.0566
Constant	-1.4273	1988-1990	-0.0419

Table 3: OLS for Share of Risky Assets in Household Portfolio with Cohort Effects (SCF), N=21,778

Age	Coefficient	Cohort	Coefficient
23-25	(omitted)	1919-1921	-1.4651
26-28	-0.0010	1922-1924	-1.0181
29-31	0.0353	1925-1927	-0.9239
32-34	0.1739	1928-1930	-0.7940
35-37	0.4163	1931-1933	-0.8928
38-40	0.5209	1934-1936	-0.7637
41-43	0.5531	1937-1939	-0.6232
44-46	0.6351	1940-1942	-0.6912
47-49	0.7963	1943-1945	-0.5213
50-52	0.8147	1946-1948	-0.5880
53-55	0.9260	1949-1951	-0.4477
56-58	0.8842	1952-1954	-0.2879
59-61	0.7891	1955-1957	-0.3955
62-64	0.9596	1958-1960	-0.1467
65-67	0.9803	1961-1963	-0.1118
68-70	0.9177	1964-1966	0.0636
71-73	0.9793	1967-1969	-0.0321
74-76	1.1988	1970-1972	0.1489
77-79	1.1405	1973-1975	(omitted)
		1976-1978	0.1307
		1979-1981	-0.0045
		1982-1984	-0.0401
		1985-1987	-0.2877
Constant	-2.0059	1988-1990	-0.4191

Table 4: Probit for Stock Market Participation with Time Effects (SCF), N=34,008

Age	Coefficient	Year	Coefficient
23-25	(omitted)	1989	-0.3832
26-28	0.3273	1992	-0.2460
29-31	0.4679	1995	-0.1837
32-34	0.4772	1998	0.0593
35-37	0.5310	2001	0.1716
38-40	0.6241	2004	0.0845
41-43	0.7148	2007	0.1236
44-46	0.6395	2010	0.0138
47-49	0.7464	2013	(omitted)
50-52	0.7604		
53-55	0.7810		
56-58	0.7793		
59-61	0.7266		
62-64	0.6637		
65-67	0.5799		
68-70	0.4752		
71-73	0.3728		
74-76	0.3286		
77-79	0.3397		
Constant	-0.4498		

Table 5: OLS for Share of Risky Assets in Household Portfolio with Time Effects (SCF), N=21,778

Age	Coefficient	Year	Coefficient
23-25	(omitted)	1989	-0.8549
26-28	0.0929	1992	-0.4849
29-31	0.1463	1995	-0.0183
32-34	0.2565	1998	0.2701
35-37	0.4604	2001	0.5515
38-40	0.4902	2004	0.1675
41-43	0.4534	2007	0.1702
44-46	0.4605	2010	-0.1126
47-49	0.5385	2013	(omitted)
50-52	0.4940		
53-55	0.5668		
56-58	0.4259		
59-61	0.2997		
62-64	0.4442		
65-67	0.3530		
68-70	0.2147		
71-73	0.1343		
74-76	0.2540		
77-79	0.1602		
Constant	-0.6536		

B Figures

Figure 29: Earnings Statistics (CPS)