# Stock Market Investment: The Role of Human Capital[*]

Kartik Athreya [†]
FRB Richmond

Felicia Ionescu [‡]
Federal Reserve Board

Urvi Neelakantan [§]
FRB Richmond

May 15, 2016

**Abstract**

Portfolio choice models counterfactually predict (or advise) almost universal equity market participation early in life. Empirically consistent predictions have proved elusive without participation costs, informational frictions, or nonstandard preferences. We demonstrate that once investment in human capital is taken into account, standard theory predicts stock market participation much closer to that empirically observed. The mechanism at work is intuitive: Early in life, most households' initial human capital levels are low while the horizon over which they will recoup any payoff from learning is long. Early investment in human capital, especially for those with relatively high ability or low current productivity, therefore yields a sustained increase in expected future earnings, but requires forgoing current earnings. As long as they prefer smooth consumption over the life-cycle, households who invest in human capital early in life will not find saving in financial assets useful. Moreover, those who do *borrow*, will not do so to finance stock purchases, as would occur if they did not have to expend resources to acquire human capital (Davis et al. (2006)), but rather to finance education (human capital). As a quantitative matter, we show that the dispersion in human capital returns across households, when disciplined to match observed earnings paths and heterogeneity, substantially explains overall stock market participation over the life-cycle.

[*]We thank Nika Lazaryan for excellent research assistance on this project. We are grateful to Marco Cagetti, Marios Karabarbounis, Michael Haliassos, 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, Midwest Macro meetings, Society for Economic Dynamics, and Stony Brook University for helpful comments and suggestions. 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

JEL Codes: E21; G11; J24;

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

# 1 Introduction

Household participation in the stock market is generally limited. This observation has proved difficult to explain within the context of models that avoid the imposition of nonstandard preferences, stock market participation costs, or imperfect information. The contribution of this paper is to show that once human capital investment is allowed for, stock market (non) participation can be well understood within an entirely standard setting.

Our approach is premised on a simple observation: Throughout life, and especially early in it, households have human capital as an investment opportunity and in deciding how to invest over the life cycle, will take this into account. Early in life, households' initial human capital levels are relatively low while the horizon over which they will recoup any payoff from learning is long. This induces human capital investment. Early investment in human capital, in turn, yields a sustained increase in expected future earnings over a long period. To ensure that consumption remains smooth in the face of such an increase, all households who invest in human capital early in life will desire, absent risk, to *borrow*, not save in financial assets. And while the presence of risk motivates precautionary savings, such balances are trivial for much of the life-cycle: Buffer stocks, if not zero, are small (Gourinchas and Parker (2002)).

Nonetheless,the presence of an upward-sloping profile for average earnings over the life cycle is not necessarily a barrier to stock market participation. The logic is simple: A household still cares about provisioning itself for the future, and if the rate of return on one asset class (stocks) is high, while that on an other (bonds) is low, then borrowing that asset (shorting the bond) and investing the proceeds in stocks may well be a sensible strategy. Important work of Davis, Kubler, and Willen (2006) (henceforth DKW) demonstrates precisely this: They show that absent substantial 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.

So why then do we find that human capital is important in diverting households away from the stock market? The answer is that we *endogenize* human capital. Thus, to arrange a particular age-earnings profile, individuals must *invest* in human capital. And to invest in human capital requires resources. In particular, it requires time, and hence, implies foregone earnings. As a result, even given access to credit markets, a young investor facing high marginal returns to investment in human capital will not find the strategy of borrowing to hold stocks attractive. Instead, if they borrow, it will be to finance human capital accumulation–by allowing for consumption while investing time in human capital acquisition.

By contrast, When agents are instead implicitly endowed with human capital, as is the case whenever earnings processes are exogenous, borrowing costs become essential at preventing high rates of stock participation by the young. Once human capital is not given, but must instead be acquired, participation falls irrespective of borrowing opportunities. In this sense, our works builds closely on the insight of DKW.

Our mechanism suggests that financial asset positions consistent with early investment in human capital will likely involve low stock market participation when young, savings in both stocks and risk-free assets in middle age( as diversification becomes important), and dissaving in retirement (to finance consumption).[1]

To investigate this logic, we embed the classic Ben-Porath (1967) model of time allocation between working ("earning") and human-capital accumulation ("learning") into a life-cycle consumption-savings model with uninsurable idiosyncratic labor income risk and financial portfolio choice. To our knowledge, we are the first to study human capital *investment* and financial investment decisions in such a setting.

While intuitively appealing, there is no a priori guarantee that the option to invest in human capital is capable of generating a *quantitatively* plausible account of observed household wealth and portfolio behavior. A principal contribution of our paper is to demonstrate that it is. We will show that when households have access to financial assets yielding empirically accurate returns, a standard human capital investment process, disciplined to be consistent with observed earnings data alone, can account well for limited stock market participation, over the entire life cycle. Importantly, our model's predictions are broadly consistent with the observed path for household wealth levels, both total wealth and the levels invested in risky and risk-free assets. We also find that those who do invest in the risky 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 are interesting: The come from a model that makes noinvestment over the life cycle comes from a model in which there is no appeal to stock-market participation costs, behavioral assumptions, or informational imperfections. This strongly suggests the importance of human capital investment as an option over the life cycle.

An observation critical to understanding our findings is that both total and marginal returns to human capital are individual-specific and depend on the household's current holdings while, from the household's perspective, the marginal return to financial investments is invariant to any household-level characteristics—it is simply the market rate of return that the household takes

[1] Though participation is higher among savers in our mode, it is not 100%. This is because the returns to both human capital and stocks are risky. This gives households a countervailing incentive to build a buffer stock of riskless savings, leading even some of those households who save to avoid risky stocks. We will return to this point later.

as given. Heterogeneity in the marginal returns to human capital arises from differences across agents not only in their ability to learn or in their initial human capital, but also from any other source of variation in household productivity. One important source of additional heterogeneity is idiosyncratic dispersion in wages, and hence in the opportunity cost of human capital investment. To ensure that the dispersion in the marginal payoffs to human capital is of an empirically reasonable magnitude, we follow a huge body of existing work, and especially that of Huggett, Ventura, and Yaron (2011), and allow for idiosyncratic variation to the payoffs from human capital.

The idea that human capital might play a significant role in how households invest their wealth is not new (see, for example, the early work of Brito, 1978). Several papers study, as we do, portfolio choice in a life-cycle setting with uninsurable, idiosyncratic labor income risk. However, in most of these papers, human capital is only implicitly defined as the present value of exogenously imposed labor income processes. Examples include Campbell, Cocco, Gomes, and Maenhout (2001), Gomes and Michaelides (2003), 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.

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 (which we do by calibrating to 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 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 DKW, and also by Constantinides et al. (2002); the mechanisms in these two papers are very much at work in our model as well.[3]

[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]Both Constanindes, Donaldson, Mehra (2002), and Davis, Kubler, Willen (2006) demonstrate that borrowing constraints play a key role in generating low demand for equities (and hence a high equilibrium premium for stocks over bonds). In those models, households look forward in their life-cycle planning, and would, if allowed to, borrow and invest in equity. Borrowing constraints prevent this and hence lower demand for equity (and boost equity

We have a quantitatively rich structure that adds empirically-relevant within-generation variation to the between-generation variation that Constantinides et al (2002) features, and moreover, endogenize earnings itself. The latter is a critical step as it recognizes that investors in practice generates plausible average earnings paths even though it is not directly parameterized to do so. In our model, by contrast, households may invest in human capital throughout life and, in particular, after formal schooling is typically completed. They may also borrow. 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. Our approach thus emphasizes financial investment in a setting that explicitly captures human capital and household heterogeneity over the entire life cycle. This is why, in terms of specifying the mechanism for human capital accumulation, we follow Ben-Porath (1967), Huggett, Ventura, and Yaron (2011) and Kim, Maurer, and Mitchell (2013). Huggett, Ventura, and Yaron (2011), in particular, not only endogenize human capital, but also capture both the life-cycle and cross-sectional distribution of earnings. Kim, Maurer, and Mitchell (2013) examine the dynamics of portfolio adjustment in a model that takes into account the fact that doing so is costly in terms of foregone leisure and human capital. We follow their approach to modeling human capital accumulation, though our focus is on documenting the role of human capital accumulation, absent other costs, in matching life-cycle stock market participation and investment shares.

As noted above, a common assumption in many papers aimed at understanding stock-market behavior is that participation entails a cost, usually in the form of a fixed cost of entry; see, for example, Campbell, Cocco, Gomes, and Maenhout (2001) and Cocco (2005).[4] Some of the preceding papers also make assumptions on preferences, such as allowing for habit formation (Gomes and Michaelides, 2003; Polkovnichenko, 2007) or heterogeneous risk preferences (Gomes and Michaelides, 2005). Along this dimension, our work is closest to that of Davis, Kubler, and Willen (2006), who do not make additional assumptions on preferences or stock-market participation costs and obtain limited stock-market participation early in life via the presence of a wedge between the borrowing rate and risk-free savings rate. However, they do not allow for human capital investment, and, as we will show, this matters. Nevertheless, it is useful to keep in mind that our model is indeed close to theirs: In the special case of our model where human capital investment is not permitted, we find results very similar to theirs. In essence, therefore, our work can be seen as building most closely on the insights of four papers—Davis, Kubler, and Willen (2006), Roussanov (2010), Huggett, Ventura, and Yaron (2011), and Kim, Maurer, and Mitchell (2013)—to demonstrate that household financial investment behavior can be quantitatively understood with standard tools.

[4]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.

While our model's ability to closely account for participation (the "extensive margin" of stock market investment) represents a contribution to one strand in the literature, we will also document the model's implications for the share of wealth,invested in stocks (the "intensive margin"). Here again, our work is connected to Davis et al. (2006) and to Constantinides , Donaldson, and Mehra (2002). Their key insight, which is at work in our setting is that 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 observed earnings dispersion, the typical household's share of financial wealth held in stock-market equity is far from 100%. 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). The mechanism by which we obtain this result differs, however. While these authors find that shares exhibit a hump shape if labor income and stock market returns are positively correlated at long horizons, we show that positive correlation is not necessary.

Before laying out the model in detail in Section 3, we describe some facts about household portfolios and earnings in the next section. The calibration is laid out in Section 4 and results are provided in Section 5.

# 2 Data

## 2.1 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

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.

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.

### 2.1.1 Cohort Effects

We first estimate life cycle 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 s for a panel can be generated. In particular, sample cohort means will be consist75(b)-29.1249(e)-311.425..5752.8i249342t t laty o

4.579((w)4.33755(h)0.975119(ln)03908586(i(le)3.3871(o)-344.104(b)-29.1119(o)-2.83755((,)-808557(s)-326.675(s)-266.444(

filebo ctopar f participatio t thg st1190Td[(c).4254714(m)1.96045(m)1.96614(a)-2.83857(ks)4.50782(s)1.9603.92

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

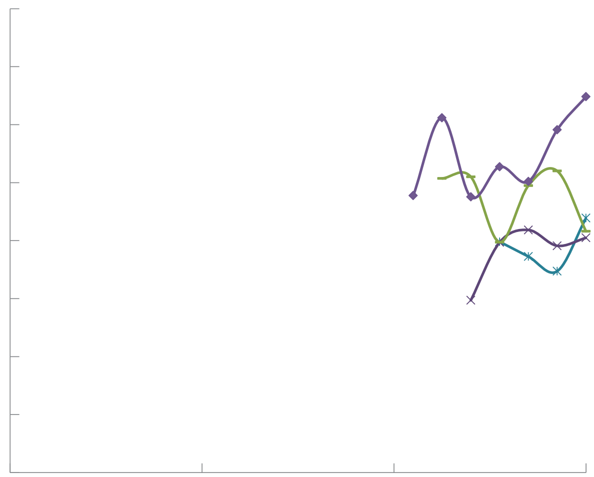

estimation are reported in Table 2 in the Appendix.[6] We use the coefficients to construct our estimate of the life-cycle profile of stock-market participation. Figure 2 shows the results for the cohort born in 1973–75.[7] By our estimation, participation in the stock market increases till agents reach age 60, after which it levels off.

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

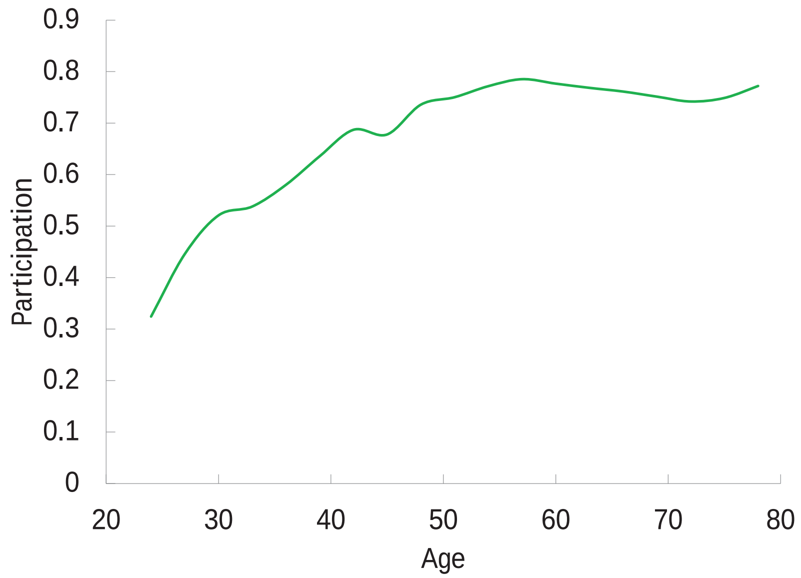


We are also interested in portfolio allocation over the life cycle conditional on participation. In other words, we want to know how the fraction of assets invested in stocks evolves over the life cycle. As we will describe later, our model will have one risk-free asset $b$ and one risky asset $s$, so the measure in which we are interested is $\frac{s}{s+b}$. As described earlier, the risky asset is the value of equity that the household holds, which includes directly held stocks and stocks in mutual funds, retirement accounts, and other managed assets. household's household's risk

We calculate the fraction $\frac{s}{s+b}$ for those who own equity. This measure lies between 0 and 1 by construction, so we want our life-cycle estimate of it tonlie between 0and 1nas6.507(w)34.4998(e)3.38586(ll.)-4
statistical significance of the results may l
e-cycle profiles; lifnle3.405.088n4(o).93142(W)87.4753(eo-5.8

we construct a logistic transformation to obtain the variable $Y_i = \ln \frac{\frac{s}{s+b}}{1-\frac{s}{s+b}}$. We run the following Ordinary Least Squares (OLS) regression on this variable.[8]

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

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.

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

### 2.1.2 Time Effects

We recognize that making different identifying assumptions can generate different life-cycle estimates, particularly for shares (Ameriks and Zeldes, 2004). 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.[9]

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

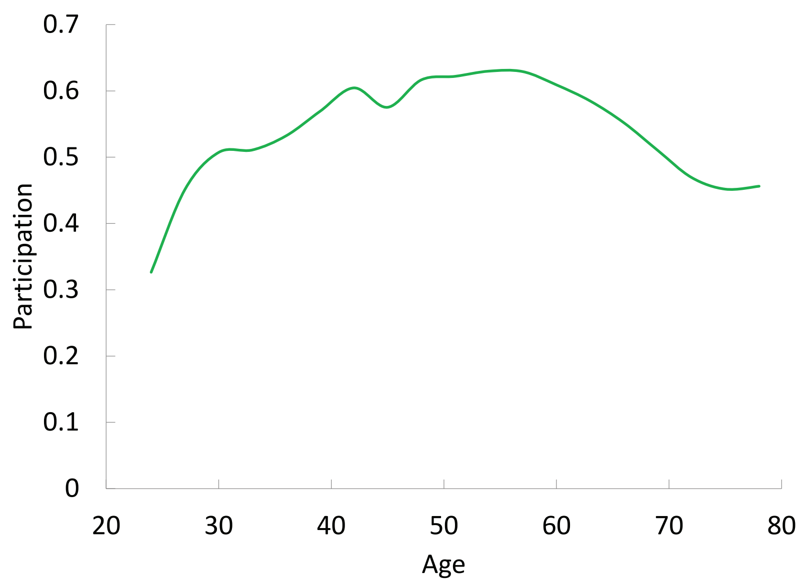


Correspondingly, we run an OLS regression as in Equation (2) with time dummies to estimate the life-cycle profile of shares. Figure 5 shows the result.[10]

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

[9]The results of the estimation are reported in Table 4 in the Appendix.
[10]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)

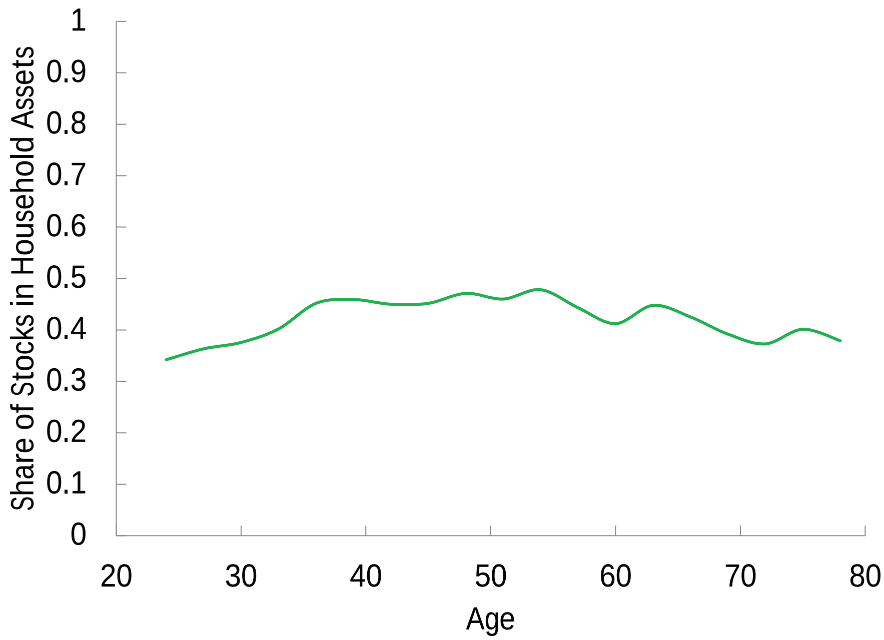


Ameriks and Zeldes (2004).

## 2.2 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 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 27 in the Appendix.[11]

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

[11] 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.

# 3 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 $t = 1, ..., T$, where $t = 1$ represents 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.

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 $t > J$.

## 3.1 Preferences

All agents have identical preferences, with their within-period utility given by a standard CRRA function with parameter $\sigma$, and with a common discount factor $\beta$. 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.

## 3.2 Financial Markets

There are two financial assets in which the agent can invest, a risk-free asset, $b_t$, and a risky asset, $s_t$, to be interpreted as stock-market equity.

### 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 ($b_t < 0$), however, carries an additional (proportional) cost as in Davis, Kubler, and Willen (2006), denoted by $\phi$, 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 $\underline{b}$.

### Risky assets

For ease of exposition, we will refer to the risky assets 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 $\mu$ is the mean excess return to stocks. The second, $\eta_{t+1}$, 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 $N(0, \sigma_\eta^2)$.

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 $b < 0$.

## 3.3 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 $\delta$. 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.

## 3.4 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

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

[12] The growth rates for wages are estimated from data. See Section 4.1 for details.

follows an AR(1) process as in Gourinchas and Parker (2002) and Hubbard, Skinner, and Zeldes (1995), with $\nu_{it} \sim N(0, \sigma_\nu^2)$ representing an innovation to $u_{it}$. The variables $u_{it}$ and $\epsilon_{it}$ are realized at each period over the life cycle and are not correlated.

## 3.5 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, $\tau_t$, which depend on age, $t$, income, $y_t$, and net 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, case 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 $\psi$ of their income in the last period as working adults, $y_J$, which they divide between risky and risk-free investments. this may

### 3.6 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. Let any period $t$ variable $j$ be denoted by $j$ and its period $t+1$ value by $j^{'}$. 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, $\epsilon$.

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 stan 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 $\epsilon$. 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 $\epsilon$.

# 4 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.

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 $J = 58$.

## 4.1 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

Table 1: Parameter Values: Benchmark Model

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

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 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, $\alpha$, to 0.7. Estimates of this parameter are surveyed by Browning, Hansen, and Heckman (1999) 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]

Borrowing limits in the model will be allowed to vary across households. We introduce heterogeneity in these limits as follows: We first group agents in the model by quartiles of initial human capital, then compute average earnings over the life cycle for each quartile. We then set the borrowing limit for all agents within a quartile to be a given percentage of the average life-cycle earnings for that quartile. We obtain the relevant percentages from the SCF by dividing the sample into income quartiles and calculating the average credit limit as a percentage of the average income within each quartile. The resulting borrowing limits as a percentage of average earnings by quartiles are: 55%, 48%, 35%, and 27%.[15] Lastly, in our baseline model, we assume that the returns to both risky assets (human capital and financial wealth) are uncorrelated. in

## 4.3 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

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

[15]We extrapolate the first percentage from the other three rather than calculating it directly because of the large numbers of zeros in the earnings data for the lowest quartile.

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 2 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.

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.

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.[16] 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 problems $\min_\gamma \left( \sum_{j=5}^{J} |log(m_j/m_j(\gamma))|^2 + |log(d_j/d_j(\gamma))|^2 + |log(s_j/s_j(\gamma))|^2 \right)$, 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.[17] 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.[18]

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

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

[18] We obtain a fit of 9.4% (0% would be a perfect fit). As a matter of perspective, we note that a close relative of this part of our model, Huggett, Ventura, and Yaron (2006), obtain a fit of 7% (for the same value of the elasticity parameter $\alpha = 0.7$). Theirs is a Ben-Porath model where the main choice is investment in human capital to maximize lifetime earnings in a framework without investments in financial assets, debt, and without earnings uncertainty. As a measure of goodness of fit, we use $\frac{1}{3J} \sum_{j=5}^{J} |log(m_j/m_j(\gamma))| + |log(d_j/d_j(\gamma))| + |log(s_j/s_j(\gamma))|$. This represents the average (percentage) deviation, in absolute terms, between the model-implied statistics and the data.

Figure 6: Life-cycle earnings

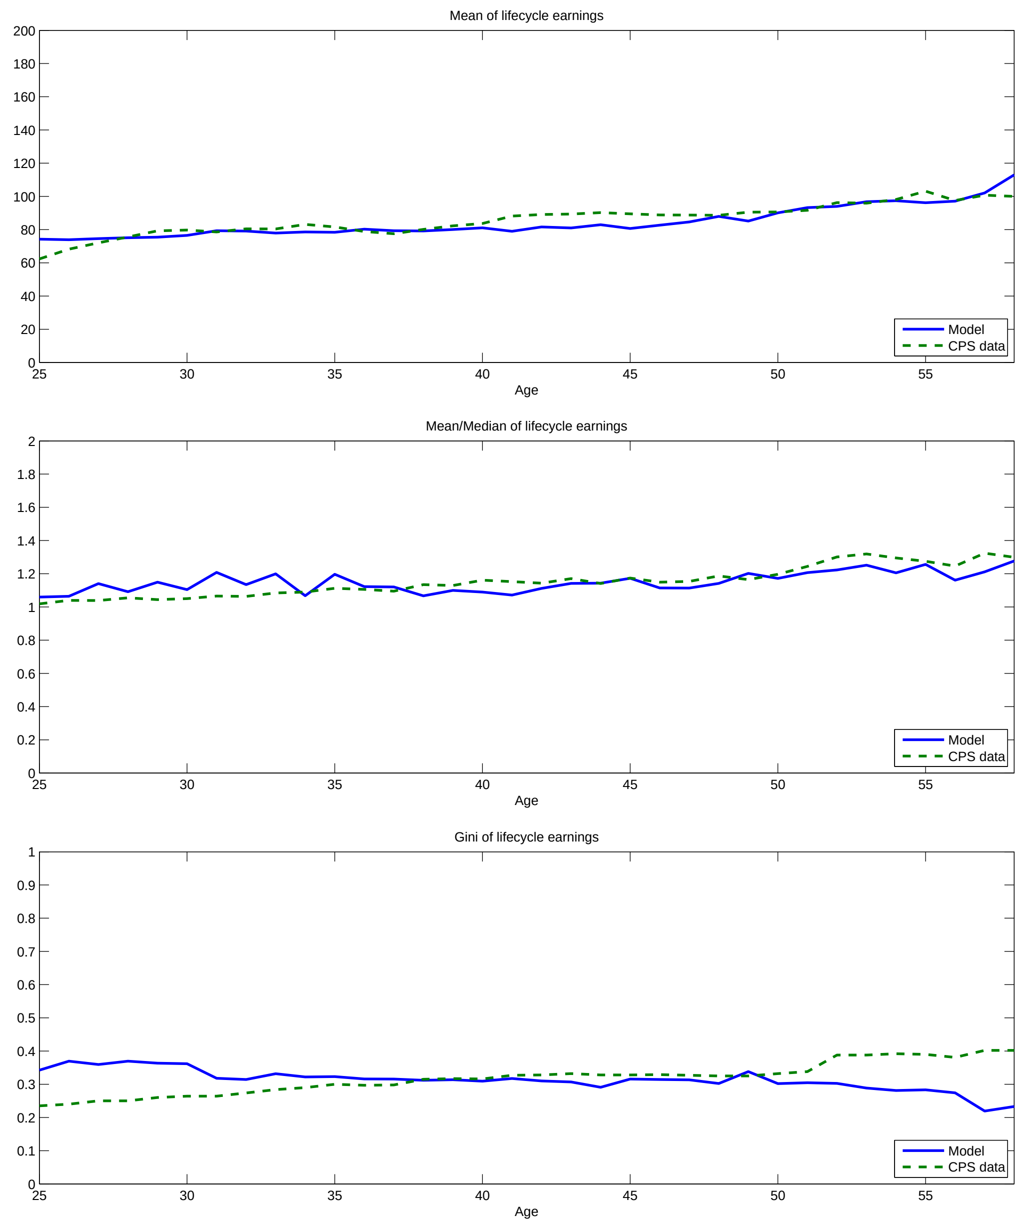


# 5 Results

Our paper provides 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 In this section we provide evidence from the model that helps explain our findings. We also study the implications for the intensive margin, i.e. the share of wealth invested in stocks. Before proceeding, however, we recall the basic mechanism at work. Early in life, forgone wages are low while marginal returns to human capital investment are high. Moreover, the horizon over which to reap the benefits of such high marginal rewards is long. The fact that acquiring human capital necessarily takes time away from working means that, unlike financial investments, it alters the time path of earnings. This implies that prior to making human capital investments, the household will consider how best to use financial assets to ensure that consumption is smooth, given that earnings will not be. Households that expect the time path of earnings to be increasing will therefore borrow to smooth consumption rather than save, and will therefore avoid all assets, including, especially, risky stocks. To illustrate this intuition more explicitly, it is useful to consider a simplified, two-period version of our model.

## 5.1 A Simple Two-Period Model

Consider a setting in which there is only one financial asset, which is risk-free and can be used for saving or borrowing. In the initial period, 0, agents choose how much to save or borrow using this asset as well as how much of their time endowment to invest in human capital. There is no uncertainty and we assume agents have standard C[redacted]A preferences, with $\sigma = 1$ (log utility). The remainder of the notation is as above, with $h_0$ denoting initial human capital, $R_b$ denoting the interest rate on risk-free assets, $\delta$ the depreciation rate of human capital, $\alpha$ the elasticity of investment in human capital, and $a$ the ability to learn. $w_1$ denotes first-period wages and $\underline{b}$ denotes the exogenous limit on risk-free borrowing. Initial-period and first-period consumption are given by $c_0$ and $c_1$, respectively. Given this, the agent's problem simplifies to:

$$\max_{l_0, b_1} \ln(c_0) + \beta \ln(c_1)$$

subject to

$$\begin{aligned} c_0 + b_1 &\leq w_0(1 - l_0)h_0 \\ c_1 &\leq w_1 h_1 + R_b b_1 \\ h_1 &= h_0(1 - \delta) + a(h_0 l_0)^{\alpha} \\ b_1 &\geq \underline{b} \end{aligned}$$

With $w_1 = w_0(1 + g)$, the optimal solution is:

$$l_0^* = \frac{1}{h_0} \left[ \frac{(1 + g)a\alpha}{R_b} \right]^{\frac{1}{(1-\alpha)}}$$

and

$$b_1^*(l_0^*) =$$

Importantly, we see that nonparticipation is not a pathology, but rather a direct implication of a standard model. As the two-period model suggests, and as we will show shortly, this result is driven primarily by the presence of human capital investment in our model. As a first step, we turn to Figure 8, which shows the trajectory of time invested in human capital over the life cycle. As is clear, time spent on human capital accumulation is at its highest early in life. For instance, at age 25, households spend about a third of their time endowment on human capital accumulation. During the early part of life, we see also that only around 30% of all households participate in the stock market. Diminishing returns, and a shorter horizon to recoup the investment, imply that human capital accumulation falls with age. As this occurs, we see that stock-market participation steadily increases, reaching around 80% at retirement age. As retirement approaches, we see that the fraction of time allocated to human capital falls sharply, reaching below 0.05 by retirement age.

Figure 7: Life-Cycle Stock Market Participation

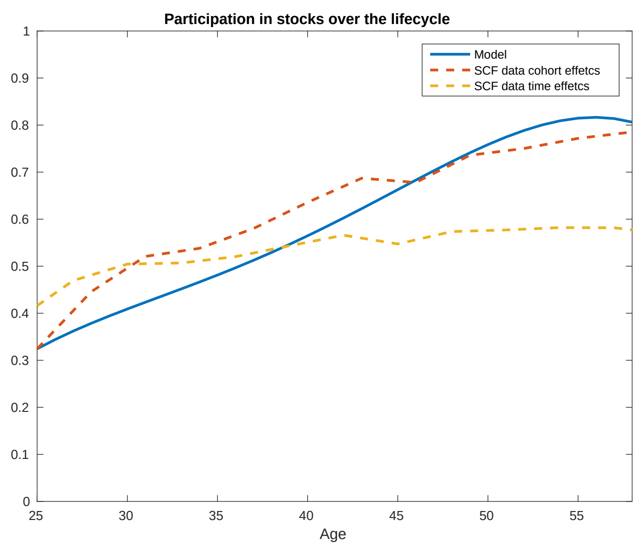

Figure 8: Time Allocated to Human Capital over the Life Cycle

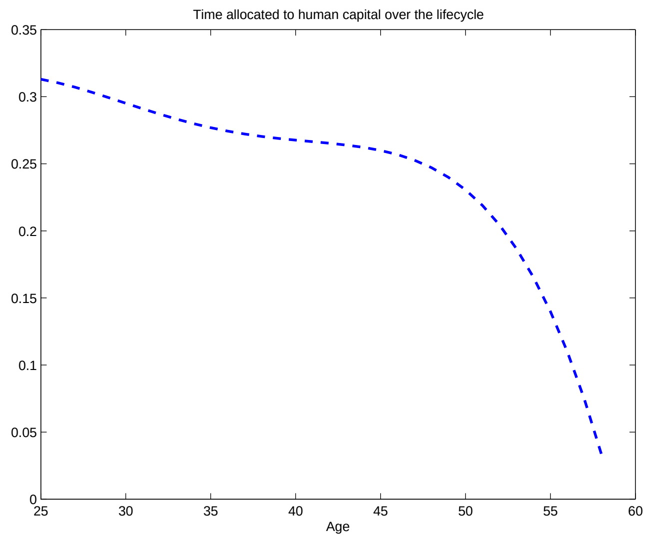


Having shown that stock-market participation can be very well accounted for by the accommodation of human capital, we turn now to the "intensive" margin of stock market investment. As seen in Figure 9, three things are salient. First, the model implies a higher share for wealth held in equity than in 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%. 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 more empirically 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. particularly difficult to square with the data (though not with advice from practitioners!). Our results suggest that perhaps a lesson is that human capital remains an investment vehicle, it is to be used.

Figure 9: Fraction of Stocks in Household Portfolio

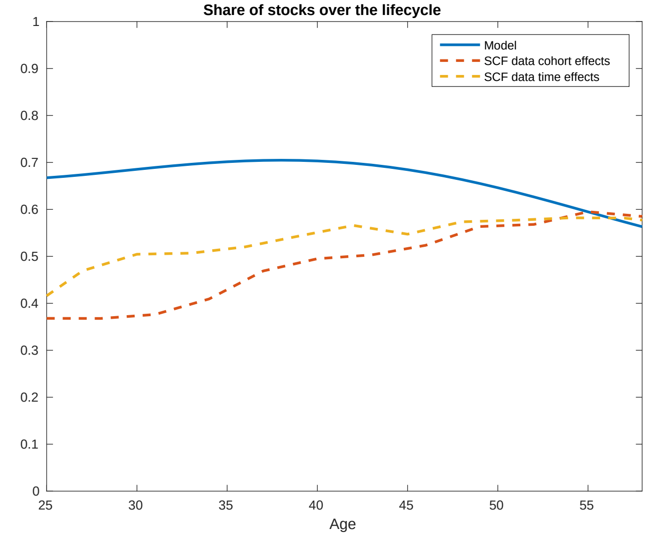

Figure 10: Life-Cycle Wealth Accumulation

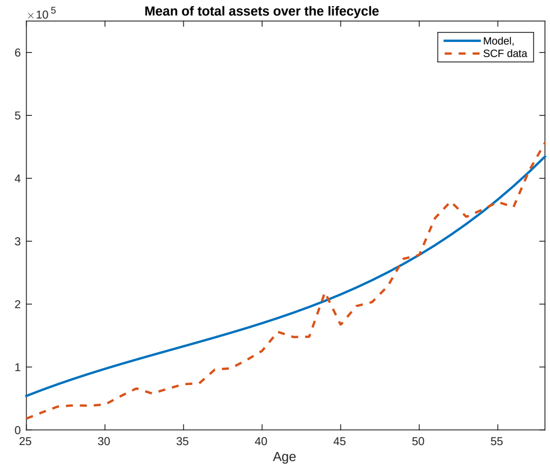


(a) Total Assets

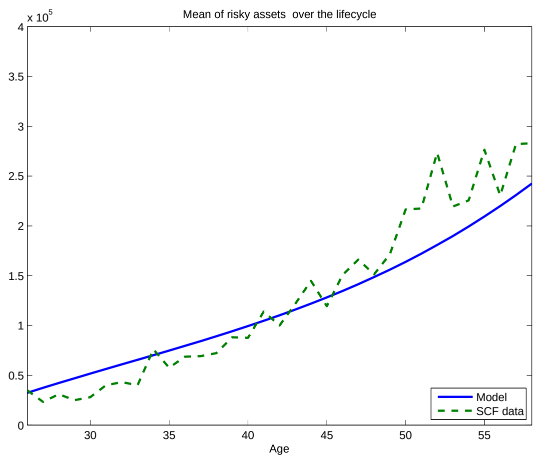


(b) Risky Assets

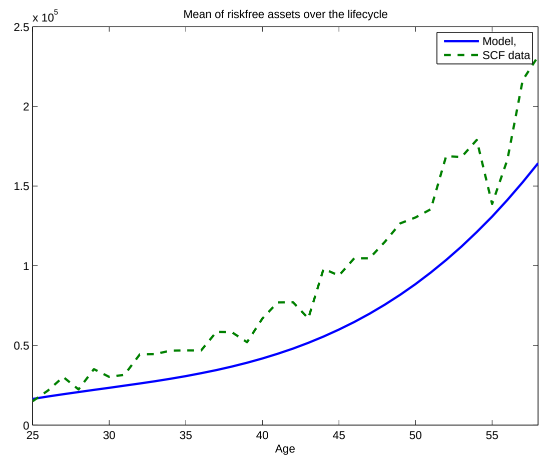


(c) Risk-free Assets

## 5.4 The Role of Endogenous Human Capital

The most direct route to seeing that that our results are driven primarily by the presence of endogenous human capital investment is to consider outcomes in which this channel is shut down. To do this, we now study a setting in which agents exogenously obtain the labor income stream generated by our [illegible] model, but do not spend any time on human capital accumulation. To study a case while retaining comparability to the benchmark model requires an additional step, however. Specifically, we "assign" earnings to agents based on their initial endowment of human capital.[20] We retain all the other features of our model, including the shocks to earnings as well as the wedge between the interest rate on borrowing and savings. As such, this setting is very close to that of Davis, Kubler, and Willen (2006). The result is a model in which the distribution of earnings unfolds as in the benchmark model, but does so without requiring human capital investment.

### 5.4.1 Exogenous Human Capital Investment and Stock-Market Investment

We first report results for participation in Figure 11.[21] Observe now that the participation rate is much higher than in the case with endogenous human capital accumulation. Indeed, we see that it reaches 100% around age 55—a result similar to what Davis, Kubler, and Willen (2006) obtain. This is an important observation because with exogenous human capital, our setting, including its quantitative implications, becomes very similar to theirs, as well as papers cited earlier (e.g., Gomes and Michaelides, 2005).[22] Indeed, we are able to recover the result in these papers that participation increases rapidly to a 100%, with the only deterrent being, just as in Davis, Kubler, and Willen (2006), the presence of the borrowing wedge. By making borrowing expensive—especially for those young households who would like to borrow—this wedge is helpful in keeping households away from stock market participation early in life. However, this result clarifies that the mechanism of high borrowing costs alone is not sufficient to explain limited stock participation later in life, when households are less likely to borrow. The relative improvement provided by our benchmark model drives home the relevance of households' ability to augment their human capital for their financial portfolio choices.

[20]Note that we still allow agents to differ in their initial endowments but assume that initial human capital and ability are uncorrelated.

[21]For ease of exposition, in this picture and those that follow, we compare model results only to our data estimates that take cohort effects into account.

[22]Note, however, that there are still quantitatively meaningful differences in our parameters and theirs, including in risk-aversion, the interest rate on borrowing, and the share of income taken into retirement.

Figure 11: Life-Cycle Stock-Market Participation Under Exogenous Earnings

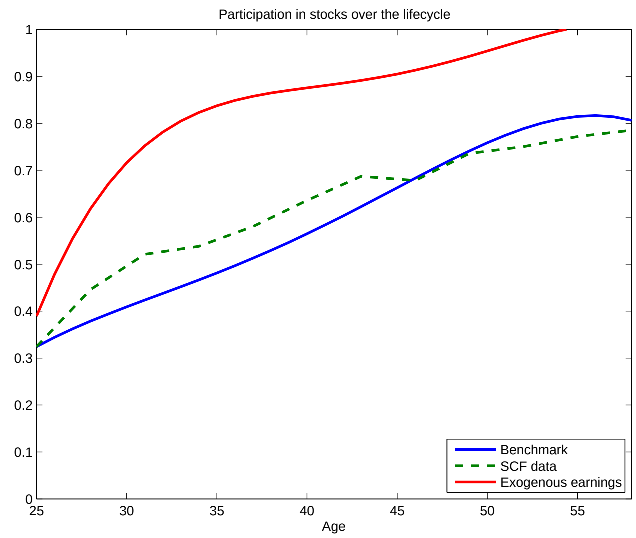


Turning next to shares, we see from Figure 12 that the exogeneity

## 5.5 Human Capital Accumulation Technology

An implication of our model is that the better the technology for learning, the less attractive stock market investment will be. In other words, 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: $\alpha = 0.9$.[23] 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 13).

Figure 13: Comparison of Marginal Densities in Model with $\alpha = 0.7$ and $\alpha = 0.9$

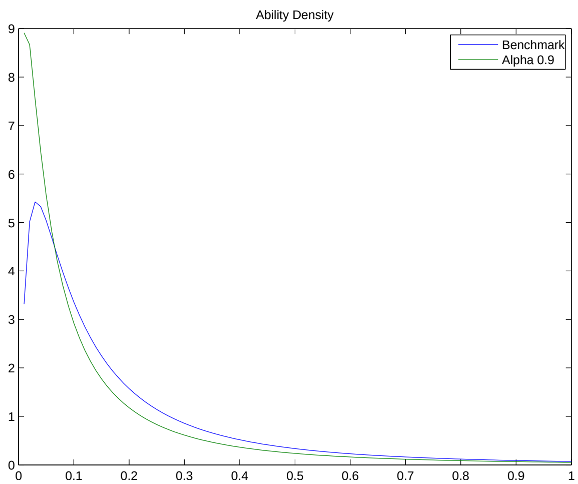


(a) Ability

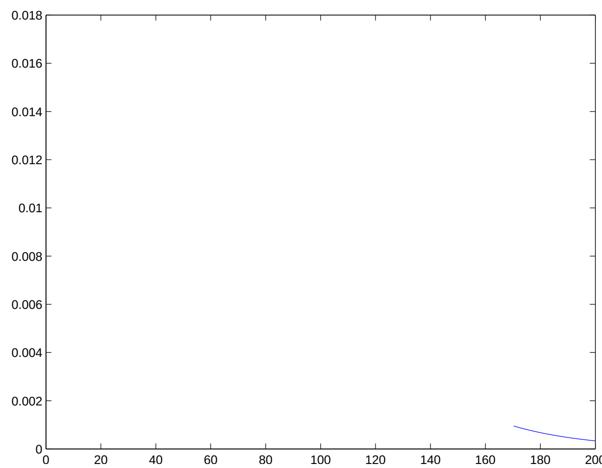

Figure 14: Results with $\alpha = 0.9$ in Recalibrated Model

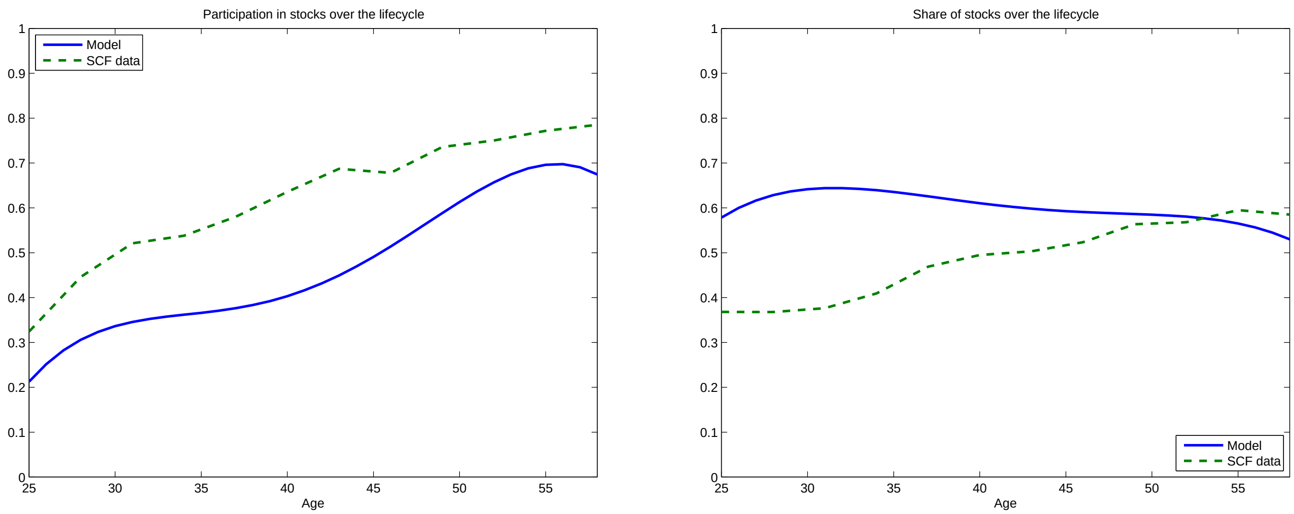


(a) Participation

(b) Shares

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 $\alpha$ to 0.5 and increasing it to 0.9. To understand the implications in this case, it is important to keep all other parameters as in the benchmark. The results are reported in Figure 15.

First consider the case where the human capital technology is less productive ($\alpha$=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 the bottom left panel shows, with the effect that their human capital levels are lower throughout working life than in the benchmark (bottom right panel). This has two effects on participation. Less time invested in human capital leads to higher 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.

In the case where the human capital technology is more productive ($\alpha$=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.

Figure 15: The Effect of the Elasticity of Human Capital Production on Investments

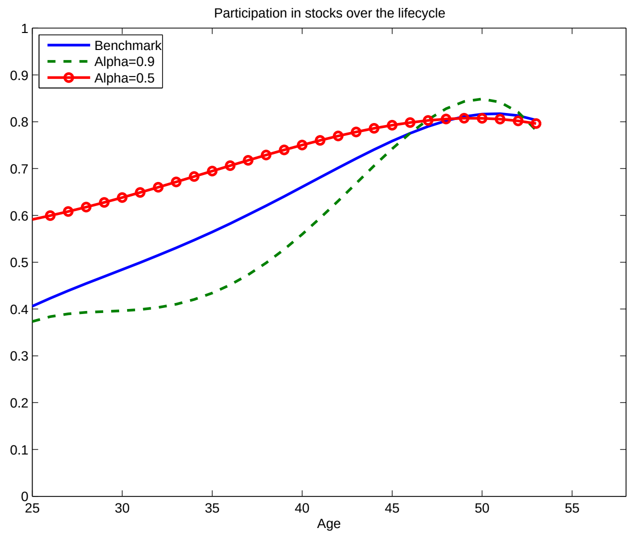

(a) Participation

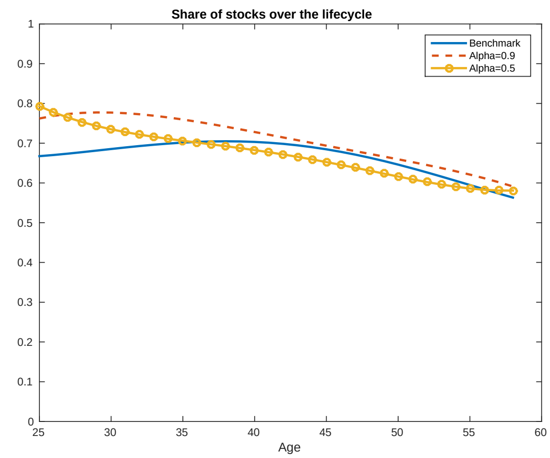

(b) Shares

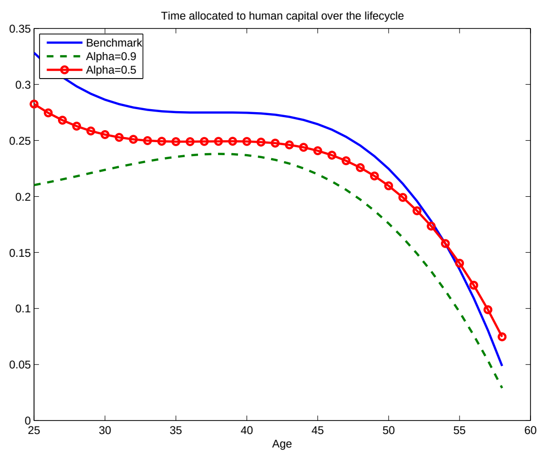

(c) Time Invested in Human Capital

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).

## 5.6 The Role of Initial Characteristics

In our setting, initial ability and initial human capital both influence the life-cycle earnings profile. Specifically, initial human capital determines the initial level of earnings, while initial ability affects the rapidity with which earnings grow from that level. We have already seen some of the effects of the earnings profile on stock-market investment in the previous experiment; here we trace these effects back to initial conditions. In our benchmark model, initial ability and human capital are positively correlated. In order to describe their effects separately, the figures below are derived from an experiment in which the conditional distributions of ability given human capital do not vary with the level of human capital, and vice versa.

Figure 16 shows participation and human capital investment behavior by quartiles of ability levels, with quartile 1 being the lowest. 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 16c). 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 one of the main questions of interest to us, namely, stock-market participation, we turn to Figure 16a. 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%). 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, 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. This analysis makes clear that once human capital investment is allowed, the model suggests that learning ability, [illegible], should be inversely related to equity investment. Of course, this fact and the fact that the model captures observed participation suggests there are other forces at work. In terms of shares, we see (Figure 16b) that quite unlike for participation, those who have chosen to invest in the stock market diversify in fairly consistent manner: Shares are quite similar across ability quartiles, especially as households age. We will see later that initial human capital as well as the risk properties of stock market returns and individual risk aversion are the main drivers of shares and none of these are in play in this experiment.

Figure 16: Investment by Quartiles of Ability

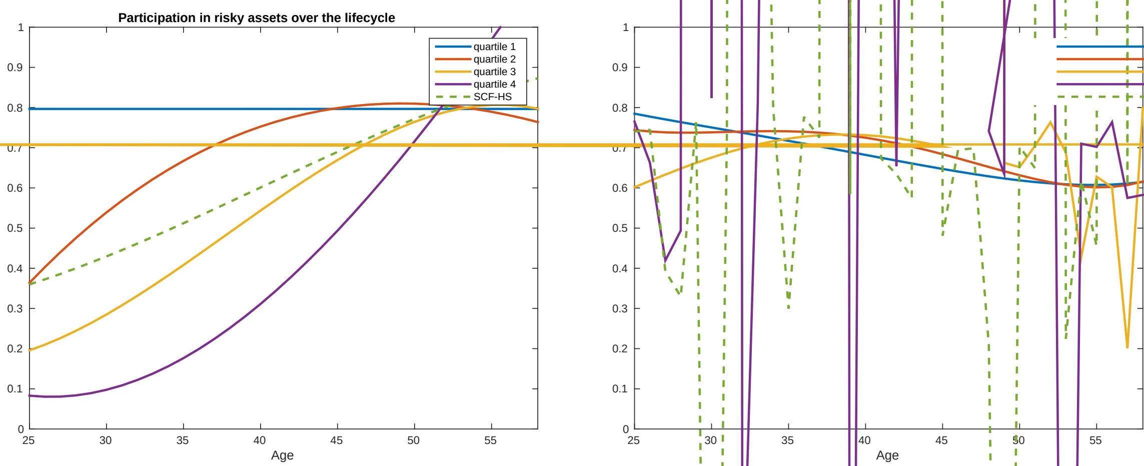


(a) Participation

well as the time allocated to human capital investment by quartiles of initial human capital. As seen in panel 17c, 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. For brevity we do not report the evolution of human capital levels but 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 towards its accumulation.

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% at age 25 and reaches 100% participation by age 50 (Figure 17a). Quartiles 2 and 3 participate at around a 30% rate early in life, and reach 100% participation after age 55. For the lowest quartile, participation starts at around 15% and remains below 50% throughout working life. When we look at the shares of financial wealth held by quartiles of initial human capital, as displayed in Figure 17b, we find that all but the lowest quartile invest a fairly similar fraction of their wealth in stocks over the life cycle.

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 low time allocation towards human capital investment 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 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. The individuals in the lowest quartile also earn the least of all groups, and hence face significant uninsurable risk, especially early in life. The riskiness of equity makes such investment unattractive. For households in the middle quartiles of initial human capital, optimal investment behavior falls between these two extremes.

While this case is instructive, it is important to note that it holds the correlation between initial human capital and learning ability at zero. Overall participation will depend, in general, on the joint distribution of ability and initial human capital. Indeed, in the baseline model, these characteristics are positively correlated. 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—who 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 16 and 17.

Figure 17: Investment by Quartiles of Initial Human Capital

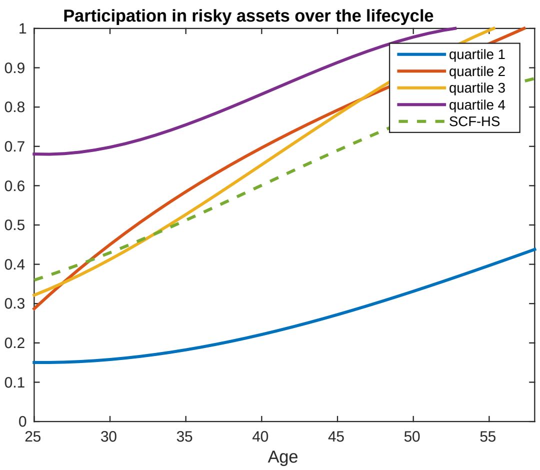


(a) Participation

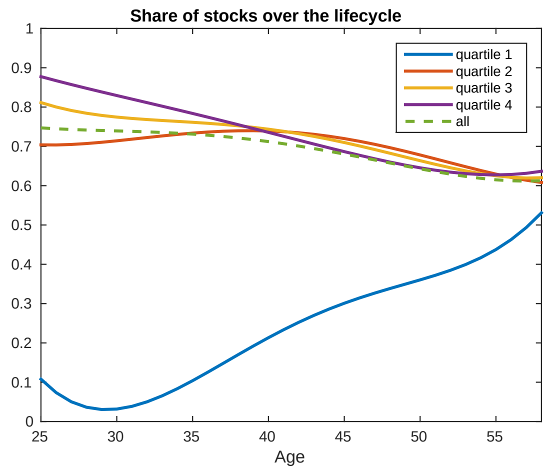


(b) Shares

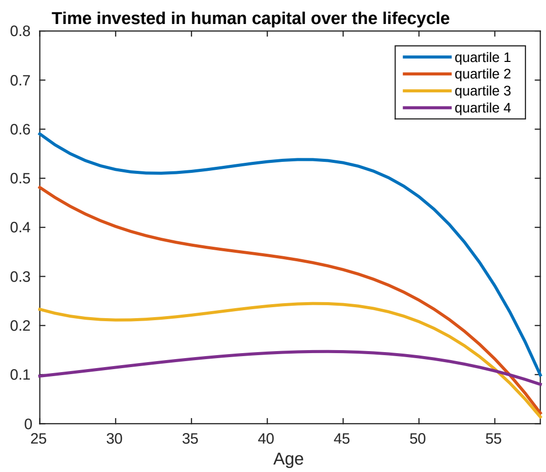


(c) Time invested in human capital

A common theme that emerges from the experiments described above is that higher human capital accumulation, if achieved through a higher initial endowment of human capital and ability or 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

Figure 18: Ability Distribution of Participants and Non-Participants

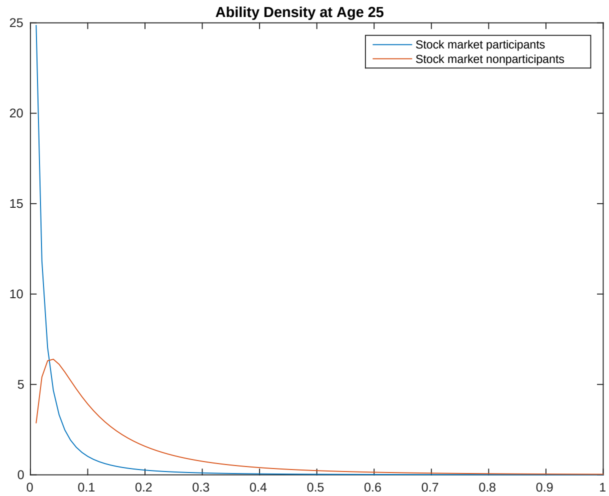


(a) Age 25

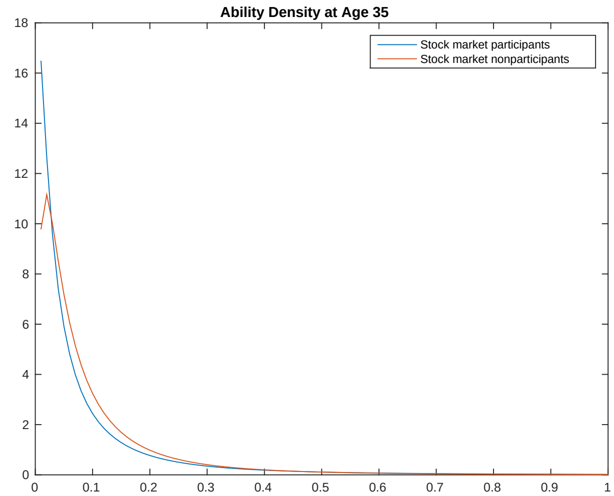


(b) Age 35

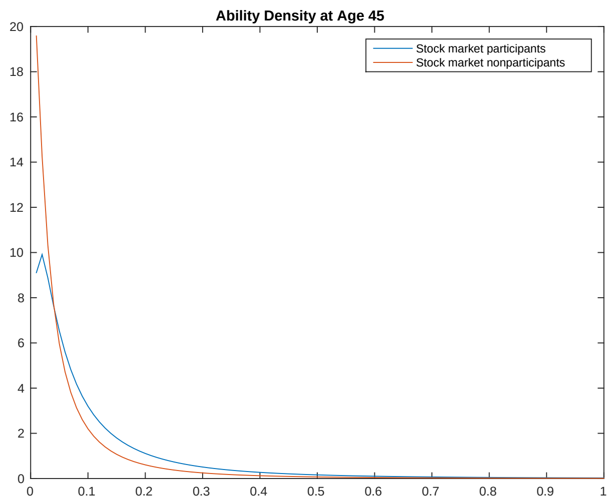


(c) Age 45

We now look at households with high initial wealth, defined here as being in the top 10 percent of the wealth distribution at age 25. Figure 19 shows clearly the central mechanism that we have emphasized: within the group of households with similar ability, it is precisely those with low initial human capital who elect not to participate in the stock market (Figure 19b).

Figure 19: Distribution of Ability and Human Capital across Participants and Non-Participants (Wealthy Households at Age 25)

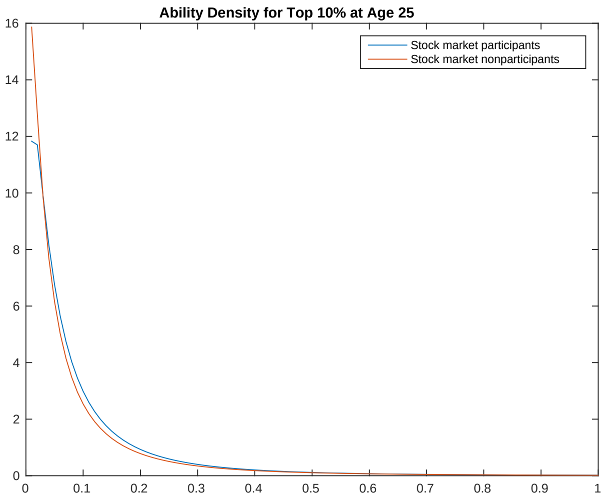


(a) Ability

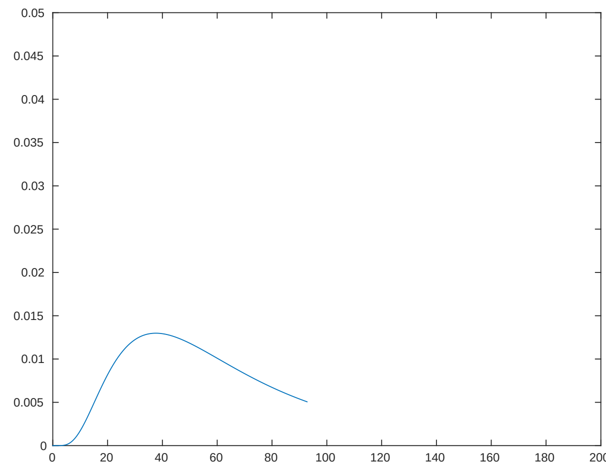

100% early in the life-cycle (Figure 22b). In this sense, our model suggests that the transactions costs used in models that abstract from human capital may be instead capturing the effect of the presence of a better alternative investment.

Figure 20: The Role of the Borrowing Wedge in Stock Market Participation

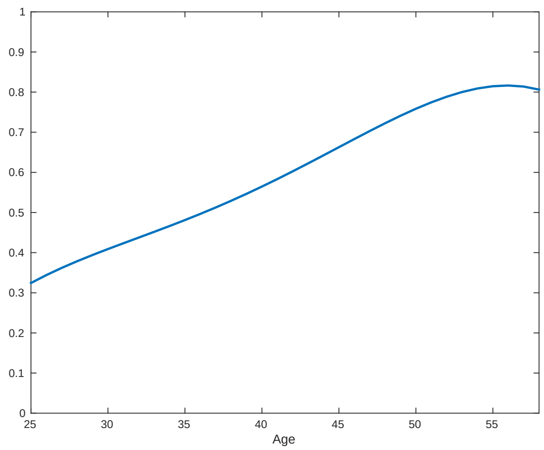

Figure 22: The Role of the Borrowing Wedge in Stock-Market Shares

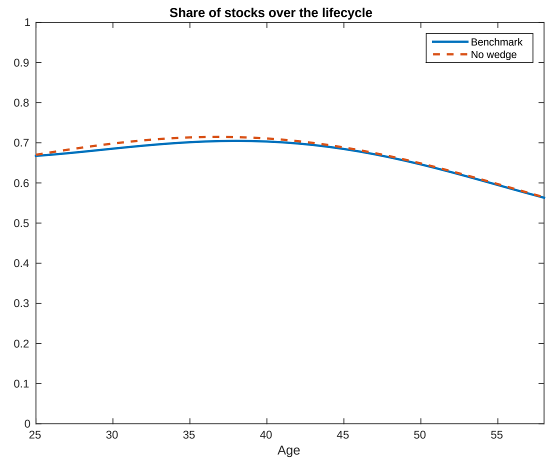


(a) Benchmark: Wedge vs. No Wedge

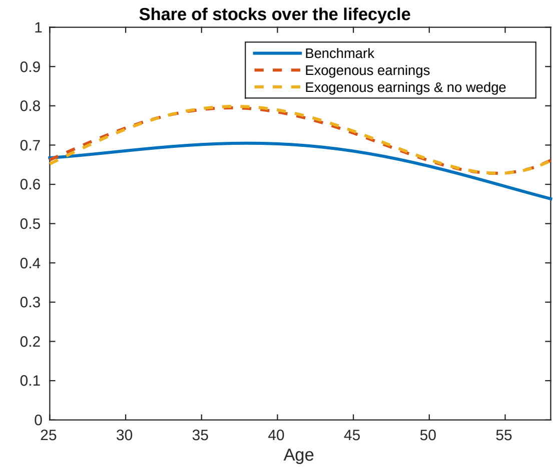


(b) Exogenous Earnings: Wedge vs. No Wedge

### 5.8.2 The Role of Expensive Credit

In our model, credit potentially serves two functions. First, it makes it possible for agents to borrow to invest in stocks. Second, it enables agents to smooth consumption while they invest time in human capital and forego current earnings. However, the benefits to using credit diminish with borrowing costs. To provide a quantitative sense of the importance of borrowing costs, we next consider a case in which the interest rate on borrowing is 22%—double the rate in the benchmark. Figures 23 and 24 show that higher borrowing costs have virtually no impact on participation and shares. This again is consistent with people borrowing primarily to smooth consumption while they invest in human capital rather than to invest in stocks. While time allocated to human capital does not change much in the aggregate, (Figure 24a), there is a marked difference if we look only at

large differences in the risk properties of stocks have almost no effect on participation compared to the benchmark. Rather, all the adjustment is on the intensive margin, and it is sizable. In the case of higher-than-baseline riskiness of stock return, we find that household diversification pays 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% when averaged over the life cycle. 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.

Figure 25: Stock Market Investment with Low and High Risk of Stocks

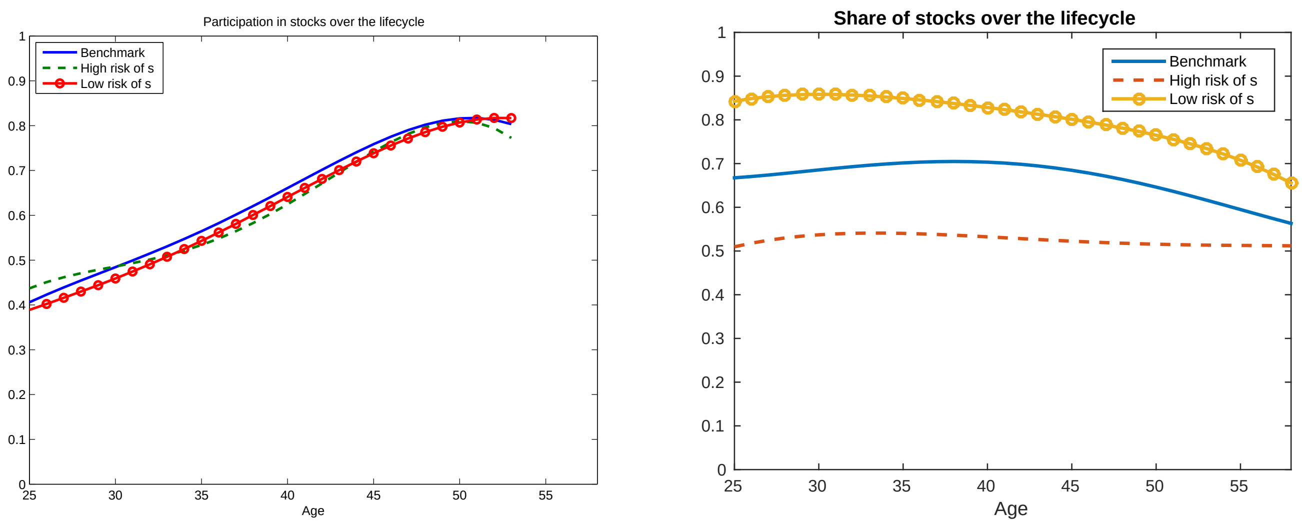


(a) Participation (b) Shares

## 5.10 Agents' Risk Aversion

We study the effect of changing agents' risk aversion in our setting. We consider two cases, $\sigma = 3$ and $\sigma = 10$. The results are shown in Figure 26.

Figure 26: Effect of Changing Risk Aversion on Household Portfolios

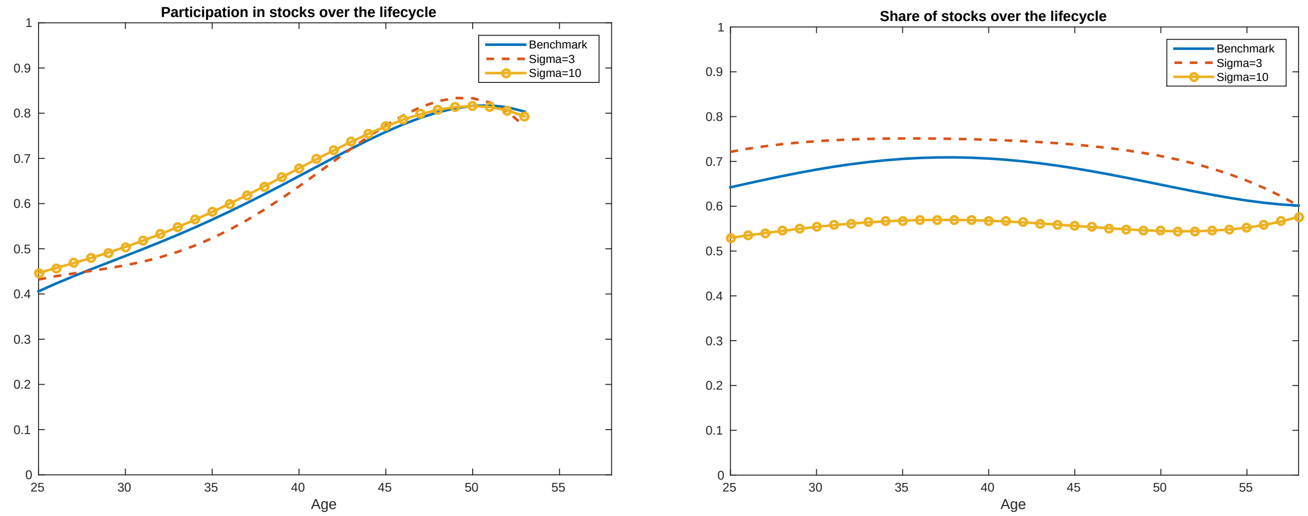


(a) Participation

(b) Shares

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. Rather, households adjust the amount of their wealth that they allocate to stocks in a completely natural manner, allocating a larger share to stocks when they are less risk averse and a smaller share when they are more risk averse. One useful implication of these results is that while we have employed a risk-aversion value that is standard in the portfolio-choice literature, it is higher than the value typically assumed in macroeconomics, which ranges from 1 to 3 for example. Therefore, it is worth noting that neither stock-market participation nor shares change substantially under lower risk aversion. This is only suggestive, however, as we do not recalibrate the entire model when we change risk aversion.

# 6 Conclusion

Research on household portfolios frequently predicts that households will almost universally participate in equity markets and allocate a high share of financial wealth to equity, especially early in life. These predictions are counterfactual and empirically consistent predictions have proved hard to obtain without stock-market participation costs, informational frictions, or departures from standard preferences. The central contribution of this paper is to demonstrate that once human capital investment is allowed for, a standard model predicts stock-market participation and equity-investment shares that are much closer to those empirically observed throughout the 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. Thus, stock-market participation starts low and grows over the life-cycle, just as in the data. As for shares, the risks to human capital limit the household's desire to hold wealth in risky financial equity. Our results suggest that the option to invest in human capital is important for understanding observed household portfolio choices over the life cycle.

# References

Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2013): "Education Policy and Intergenerational Transfers in Equilibrium," Working Paper 18782, National Bureau of Economic Research.

Ameriks, J., and S. Zeldes (2004): "How Do Household Portfolio Shares Vary with Age?," Discussion paper, working paper, Columbia University.

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 Producation 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," *Journal of Finance*, 62(5), 2123–2167.

Board of Governors of the Federal Reserve System (2014): "Federal Reserve Statistical Release, G19, Consumer Credit," Discussion paper, Board of Governors of the Federal Reserve System.

Brito, N. O. (1978): "Portfolio Selection in an Economy with Marketability and Sho

**Chang, Y., J. H. Hong, and M. Karabarbounis** (2014): "Life Cycle Uncertainty and Portfolio Choice Puzzles," Federal Reserve Bank of Richmond Working Paper.

**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.

**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.

**Gomes, F., and A. Michaelides** (2005): "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *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," in *Papers and Proceedings of the Annual Meeting*, p. 297. 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., J. Skinner, and S. Zeldes** (1994): "Expanding the life-cycle model: Precautionary saving and public policy," *American Economic Review*, 84(2), 174–179.

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

Huggett, M., and G. Kaplan (2015): "How Large is the Stock Component of Human Capital?," Discussion paper, Georgetown University.

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

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 (2013): "Time is money: Life cycle rational inertia and delegation of investment management," Discussion paper, National Bureau of Economic Research.

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 manuscript.*

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," *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," Discussion paper, National Bureau of Economic Research.

**Roussanov, N.** (2010): "Human Capital Investment and Portfolio Choice over the Life Cycle," Discussion paper, Citeseer.

**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 | 0.7--2.83857(.7)-2.83857(6)-2.8385-2.2T[632-64 | | 0676 |
| 5-67 | 0.579 | | |
| 8-70 | 0.471- | | 0.3--2.833652(8)-2.83857]TJ14.5199TLT[744-6 |
| 7-79 | 0.339 | | |

-0(4)-2.836524398 Cnan

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 27: Earnings Statistics (CPS)