Stock Market Participation: The Role of Human Capital*

Kartik Athreya † | Felicia Ionescu ‡ | Urvi Neelakantan §

FRB Richmond | Federal Reserve Board | FRB Richmond

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 invariant across households, while returns to human capital vary. We demonstrate in this paper that once human capital investment is allowed for and, critically, disciplined to match observed dispersion in earnings, a 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 alters the role of borrowing costs and short sales constraints in limiting stock market participation.

JEL Codes: E21; G11; J24;

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

---

*We are grateful to Marco Cagetti, Thomas Crossley, John Bailey Jones, Marios Karabarbounis, Stephen Zeldes, anonymous referees, and seminar and conference participants at the Allied Social Science Association Meeting, Bureau of Labor Statistics, Central Bank of Hungary, University of Central Florida, 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, George Mason University, Iowa State University, Midwest Macro meetings, Society for Economic Dynamics, FRB-St.Louis-Tsinghua University Conference, Stony Brook University, University of Connecticut, and Virginia Commonwealth University for helpful comments and suggestions. We thank especially Michael Haliassos and Yi Wen for their detailed input. We thank Nika Lazaryan for excellent research assistance. 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 the high returns stocks offer. By contrast, human capital investment is widespread early in life. The expected returns to stocks are invariant across investors and do not change with the amount invested. By contrast, the payoffs to human capital investment (earnings) vary across individuals and change with the amount (of time) invested. The objective of this paper is to evaluate the role of human capital investment for the path of stock market participation. We demonstrate that once human capital investment is allowed for and, importantly, quantitatively disciplined to match empirical measures of *heterogeneity* in earnings, an entirely standard model of portfolio choice explains the observed life-cycle path of stock market participation rates, both in the aggregate and across income and wealth groups. To our knowledge, our work is the first to demonstrate that the ability of households to accumulate human capital—and especially the effect of variation across individuals in this ability—provides a quantitatively plausible account of observed stock market participation over the life cycle.

Why might human capital investment decisions matter for life-cycle stock market participation? The answer is rooted in the fact noted above: the expected returns to human capital investment vary across individuals and over time while the expected returns to stocks do not. As we will show, for many individuals, the return to investing in human capital when young is extremely high. These individuals optimally choose to give up some earnings to spend time learning and anticipate rapid growth in their earnings over time. Intertemporal smoothing motives then lead them to borrow to finance consumption when young and keep the m away from stocks. For other young individuals, the expected return on human capital investment is lower, leading them to prefer long positions in stocks even as they invest in human capital early in life. And for all individuals, the marginal return to further investment in human capital will decrease and the opportunity cost will increase as they age and accumulate human capital. As a result, stocks grow relatively more attractive later in life, resulting in more widespread participation.

While intuitively appealing, the variation in human capital investment returns does not a priori guarantee a *quantitatively* plausible account of observed stock market participation behavior. The specific contribution of our paper is to demonstrate that it in fact does. We will show that the workhorse human capital model of Ben-Porath (1967)—with heterogeneity in learning ability and initial human capital disciplined to match heterogeneity in life-cycle earnings—can account well for stock market participation when embedded in a standard portfolio choice model. Heterogeneity is critical to our findings because it makes returns to human capital, and hence the comparison to financial investment returns, individual-specific. We calibrate this heterogeneity solely to match earnings and do not rely in any way on empirical information on financial investment choices. It is therefore noteworthy that allowing for human capital accumulation alone enables an otherwise essentially off-the-shelf model to produce variation in stock market participation decisions consistent with the data. Moreover, our model's implications for household financial wealth levels—both total wealth and the amounts invested in risky and risk-free assets—are in line with the data. These successes of the model along nontargeted dimensions suggest that human capital investment likely plays an important role in driving household financial investment over the life cycle.

The human capital mechanism we emphasize also helps clarify the role of credit constraints in stock market participation. Specifically, our approach helps explain why households may not borrow to invest in the stock market even when borrowing costs are low or borrowing limits are lax. Consider a young low-wealth investor facing marginal returns to investment in human capital that are high enough to dominate those available on stocks. All else equal, this individual will not find the strategy of borrowing to purchase stocks useful. They will, however, still find borrowing useful because the proceeds can be used to finance current consumption and thereby ease the hardship associated with spending time investing in human capital rather than earning. In other words, for such individuals, the "first dollars" of any borrowing will finance consumption, not purchases of risky financial assets. Moreover, the leverage associated with this strategy creates risk for the borrower: future consumption grows more uncertain with leverage, as debt repayment obligations loom while the payoff to human capital, like stocks, is risky. As a result, for this type of investor, leveraged risk-taking via stock market investment is unattractive. In fact, when borrowing costs are high, this investor would want to *short* stocks if they could. In this respect, our work builds on a classic argument of Friedman (1962) that ideally, individuals facing a risky payoff from human capital accumulation would like, if allowed, to issue equity claims against their future earnings. Our paper shows that individuals operating in a quantitatively plausible setting would indeed prefer to issue (via short sales) risky equity in order to finance human capital investment and, critically, that the strength of this incentive to short-sell is consistent with observed non-participation in the stock market. In other words, if individuals cannot short-sell human or financial wealth, they will proceed sequentially by accumulating human capital first, perhaps borrowing along the way, and only later accumulate financial assets.

Endogenous human capital investment is central to the preceding logic: in settings where agents are 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 those settings, the agent must only decide whether borrowing to invest in stocks makes sense at the margin, which restores the power of borrowing costs to prevent investors from holding long positions in stocks when young. Our work therefore sheds light on the question of whether households are deprived of access to lucrative financial assets by credit constraints or if they simply choose not to invest in them because they are instead engaged in human capital accumulation.

2 Related Literature

The principal result in this paper is that giving households the option to invest in human capital changes their financial portfolio allocation decisions in a manner that yields outcomes consistent with the data. Our work therefore builds on the insights of a large body of work, as we discuss below.

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).[1] These papers, building on the 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.[2] 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 the risk-free savings rate is capable of generating limited stock market participation. By contrast, we emphasize the role played by the availability of

[1] 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.

[2] 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.

an additional high-return investment option in limiting participation, even in the absence of the wedge.[3]

Though we are not directly concerned with providing a resolution to the equity premium puzzle, our model shares many features with models in the asset pricing/equity premium literature, including the presence of both uninsurable idiosyncratic labor income risk and borrowing and short sales constraints (see, for example, Lucas, 1994; Heaton and Lucas, 1996; Gomes and Michaelides, 2008). We allow households to borrow using the risk-free asset up to a limit, but we do not allow households to short stocks. Note, however, that unlike some work in this literature, we abstract from stock market participation costs and assume no correlation between earnings and stock market returns. This enables us to focus on the role played by human capital investment in stock market (non)participation and ensures that we do not deliver limited participation through other channels.

We now briefly highlight the role that the assumptions we share with this literature play in our results. Our accommodation of uninsurable idiosyncratic risk allows us to capture the substantial empirical heterogeneity across individuals of any given age. As we will demonstrate, such heterogeneity, when endogenized in an empirically disciplined manner, is precisely what generates a plausible account of variation human capital investment returns—and hence in stock market participation—across individuals of a given age.[4] As in the equity premium literature, our work also provides insight into the role played by borrowing and short sales constraints on stock market participation. For example, Constantinides, Donaldson, and Mehra (2002) demonstrated in an endowment economy that borrowing constraints provide sufficient quantitative

[3]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"), we also document the model's implications for shares in Appendix A.3. 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).

[4]We provide an example that illustrates that the returns to human capital can far exceed equity market returns for some individuals, and recent work of Huggett and Kaplan (2011) finds that, early in life, mean human capital returns exceed those of stocks.

bite to strongly limit stock market investment—especially among the young—thus resolving the puzzle.[5] Our work complements theirs by demonstrating that when households have access to the investment opportunity presented by human capital, there is once again a binding constraint that helps reconcile high equity returns with nonparticipation in stocks, especially among the young. But this time, as we show, that constraint is no longer the limit on borrowing the risk-free asset, but rather the limit on the ability of individuals to short sell stocks.[6]

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 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

[5]The crux of their explanation lies in differentiating the relative riskiness posed by risky equity to the consumption of agents of different ages: the young value stocks as diversification, while the middle-aged do not. Given binding borrowing constraints on the young, equity is effectively priced by the most risk-averse agents in the economy. We follow their structure and allow both for a life cycle and for the diversification-related benefits to the young from stock market equity by assuming zero correlation between wage and stock returns, but we show that once human capital is allowed for, there is a set of individuals for whom these benefits are overwhelmed by the returns available on human capital.

[6]Storesletten, Telmer, and Yaron (2007) is a paper in this literature that allows for short sales. In their setting, earnings and stock market returns are perfectly correlated, and households with a negative position in the risk-free asset would want to short stocks to reduce their exposure to risk. In our setting, earnings and stock market returns are uncorrelated, but young households for whom the returns to human capital dominate returns to stocks would still want to short stocks if they could, especially when borrowing costs on the risk-free asset are relatively high.

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 that invest in human capital early in life use borrowing to smooth consumption, which leads them to not want to hold long positions in stocks 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.

3 Data

Our empirical analysis proceeds as follows. First, we detail the procedure by which we derive life-cycle profiles of stock market participation from the Survey of Consumer Finances (SCF). Following that, we describe salient properties of the earnings data that we obtain from the Current Population Survey (CPS): these are key because they serve as targets that we use to discipline our model.

3.1 Household Portfolios

We obtain salient facts about household financial portfolios from the 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 stock market participation under the identifying assumption that time effects are zero. 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 dataset. 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 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 define a household as participating in the stock market if they have a positive amount of financial assets invested in equity. The SCF reports both directly held equity as well as the amount of equity 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 the part of their life cycle that we observe in the data. 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 over this age range.

The fact that households of different cohorts participate at different rates at the same age suggests that cohort effects could be important. We control for these effects using a standard probit model of the decision to invest in stocks:

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

Here $S_i = 1$ 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

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

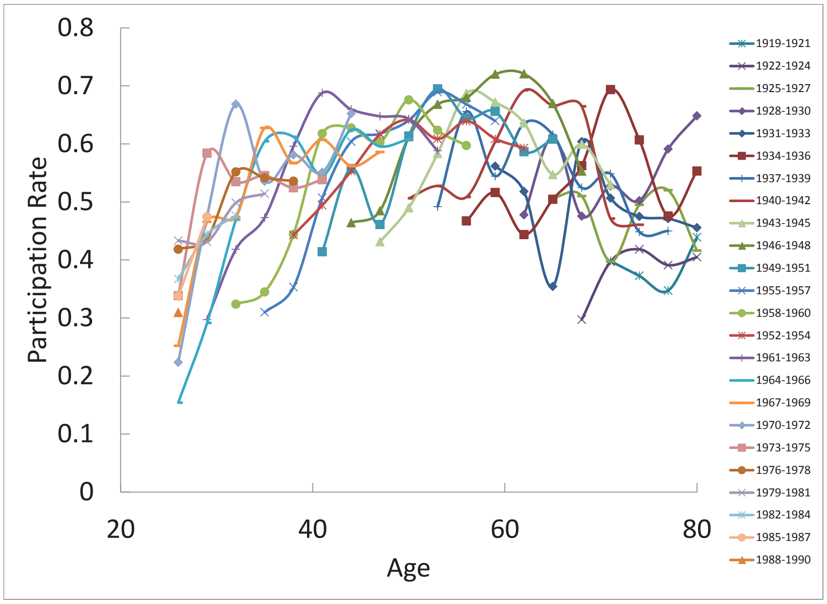

to obtain estimates that are representative of the U.S. population. As in Poterba and Samwick (1997), we estimate Equation (1) using year-specific sample weights normalized such that the sum of the weights (which equals the population represented) remains constant over time. The results of the estimation are reported in Table 2 in Appendix A.1.[7]

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.

[7]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)

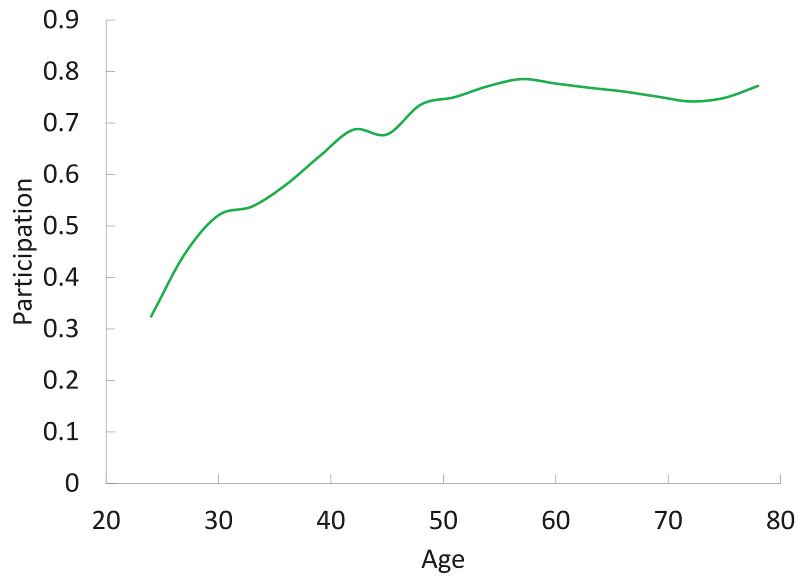

3.1.2 Time Effects

We recognize that making different identifying assumptions can generate different life-cycle estimates (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 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 results of the estimation are reported in Table 3 in Appendix A.1.

The resulting life-cycle profile is shown in Figure 3. Under the assumption that time effects matter and that cohort effects are zero, we obtain a hump-shaped rather than an increasing profile for participation. Our findings are consistent with those previously reported by Ameriks and Zeldes (2004).

Figure 3: Estimated Participation Rate over the Life Cycle (SCF, 2013 base year)

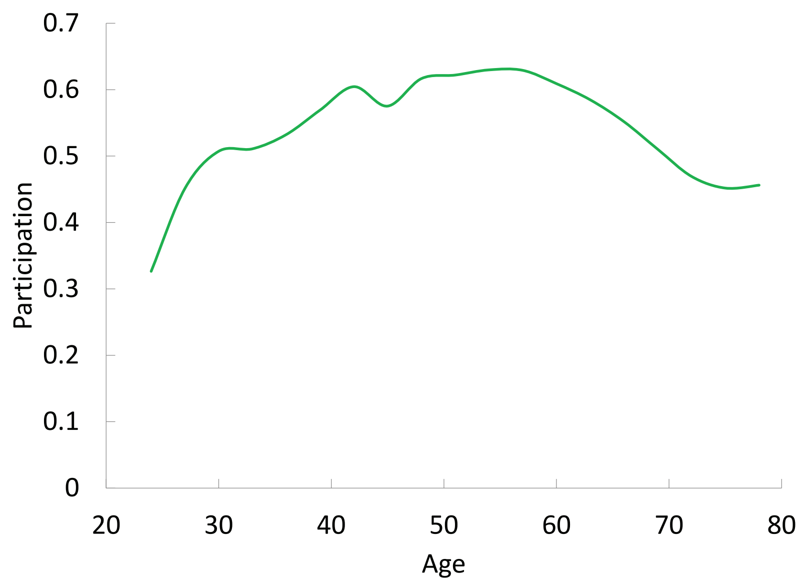

Since the two different identifying assumptions do indeed lead to different estimates for the life-cycle profiles for participation, we compare our model results to both estimates.

3.2 Earnings

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 Appendix A.2.[8]

4 Model

We turn now to the quantitative assessment of the role played by human capital investment in explaining observed stock market participation. 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). Specifically, 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 risk-free assets.

The economy is populated by a continuum of agents who value consumption throughout a finite life. Age is discrete and indexed by $t = 0, ..., T$. Agents start life in the model as high school graduates and retire at age $t = J$. 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 may be positive or negative. Borrowing is subject to a limit: $b_t \geq -\underline{b}$, with $\underline{b} > 0$.

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.,

[8] 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.

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,

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

$\Pi(\Psi_0)$ 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.[9]

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

[9] 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).

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.[10]

Risky assets

Stocks yield their owners a stochastic gross real return in period $t+1$, $R_{s,t+1}$:

$$R_{s,t+1} - R_f = \mu + \eta_{t+1}. \tag{2}$$

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 $N(0, \sigma_\eta^2)$.

Importantly, and as is standard in the models we follow (see, e.g., Cocco, 2005; Davis, Kubler, and Willen, 2006), we do not allow households to take short positions in stocks: $s \geq 0$.

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

[10] 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.

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

$$h_{t+1} = h_t(1 - \delta) + H(a, h_t, l_t), \tag{3}$$

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

$$log(y_{it}) = G(w_t, h_t, l_t) + z_{it}, \tag{4}$$

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.[11]

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}:

$$\begin{aligned} z_{it} &= u_{it} + \epsilon_{it} \\ \text{where } u_{it} &= \rho u_{i,t-1} + \nu_{it} \end{aligned}$$

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.

income

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 are a vast array of social-insurance programs that, if effective, bind 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 that 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 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

$$\tau_t(t, y_t, x_t) = \max\{0, \underline{\tau} - (\max(0, x_t) + y_t)\}, \tag{5}$$

[11] The growth rates for wages are estimated from data, as described later.

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 the risky and risk-free asset (subject to the short sales and borrowing constraints), and, in what is novel in our paper, the allocation of time between 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, it 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

$$V^R(t,a,b,s,y_J) = \sup_{b',s'}\{\frac{c_t^{1-\sigma}}{1-\sigma} + \beta E_{R_s'} V^R(t+1,a,b',s',y_J)\}, \tag{6}$$

$$\begin{aligned} \text{where } c + b' + s' &\leq \psi y_J + R_i b + R_s s \\ b &\geq \underline{b} \\ s &\geq 0. \end{aligned}$$

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, $\tau(t,y,x)$.

$$V(t,a,h,b,s,u,\epsilon) = \sup_{l,h',b',s'}\{\frac{c_t^{1-\sigma}}{1-\sigma} + \beta E_{u'|u,\ R_s'} V(t+1,a,h',b',s',u',\epsilon')\}, \tag{7}$$

where

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

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	$87.08, 35.11, 0.57$

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. In Appendix A.5, we report the effects of decreasing risk aversion to $\sigma = 3$ as well as of increasing it to $\sigma = 10$, the upper bound of values considered reasonable by Mehra and Prescott (1985). The discount factor chosen ($\beta = 0.96$) 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, $\sigma_\eta = 0.157$.[12] We assume that innovations to excess returns are uncorrelated with innovations to the aggregate component of permanent labor income.[13]

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). 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

[12]In Appendix A.5, we also study the effect on participation of raising or lowering the risk of stocks.

[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)

limit to be too loose.

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 working life 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) and range from 0.5 to 0.9. In Appendix A.4, we report the effects of different values of α on stock market participation.

To parameterize 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]

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

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

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.

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.[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] We obtain

[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

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

Figure 4 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. The model performs well given riskiness of assets and stochastic earnings in the current paper.

Figure 4: Life-cycle earnings

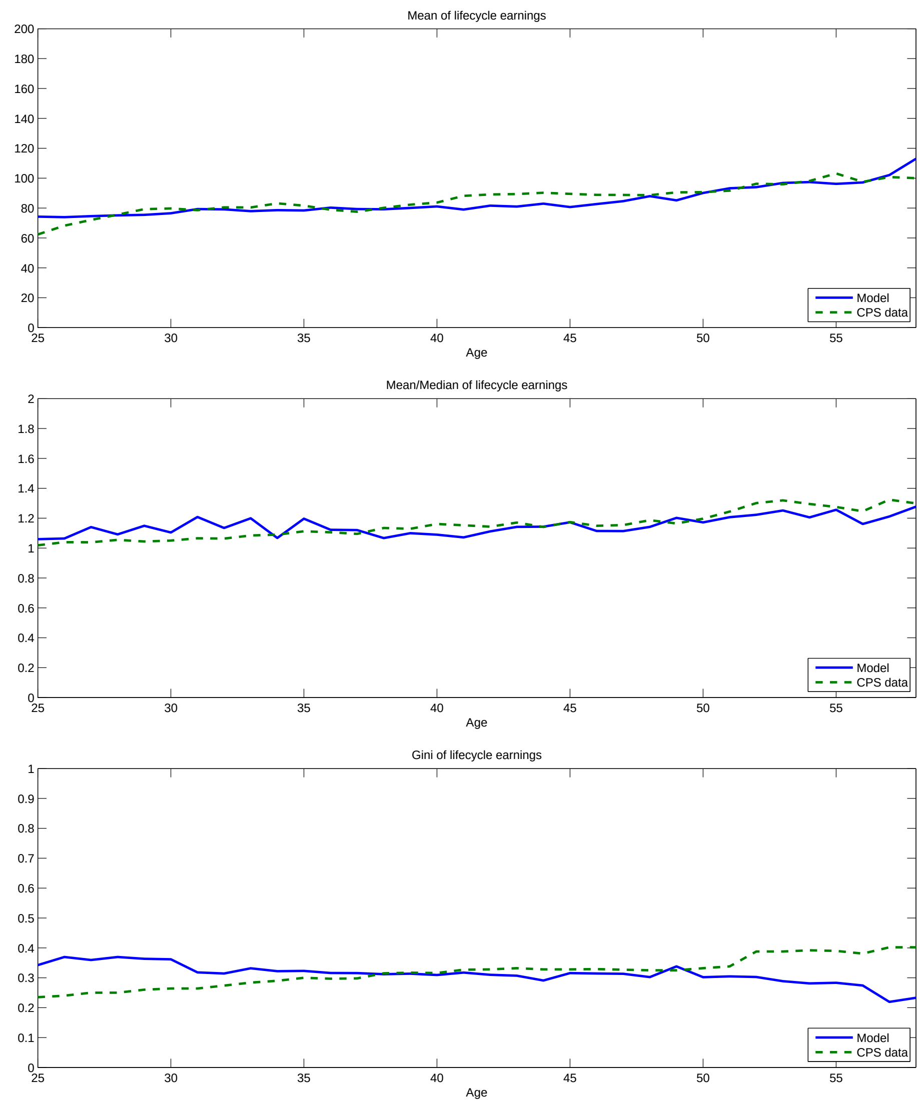

Ionescu (2009).

6 Results

We now turn to our central result: stock market participation over the life cycle can be well understood in an entirely standard model of portfolio choice when it includes the option to invest in human capital. Figure 5 compares our model predictions for stock market participation with our two empirical estimates (considering time effects and cohort effects, respectively) from SCF data. As seen, our model—with human capital investment disciplined to match only earnings—yields stock market participation rates that are broadly consistent with the data.

Figure 5: Life-Cycle Stock Market Participation

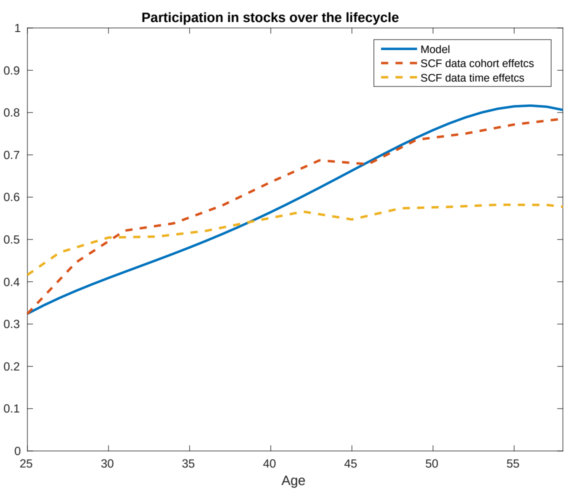

In the remainder of this section we will (i) demonstrate that the model's cross-sectional implications are also consistent with the data and (ii) elaborate on the mechanisms driving our main result.

6.1 Model vs Data: Cross-Sectional Implications

We now show that our model also delivers accurate predictions along several other salient dimensions. We remind the reader that our calibration targeted only earnings and not financial wealth or its allocation. Hence, the empirical evaluation of the model that follows is based solely on comparisons to data that our calibration did not target.

Figure 6: Life-Cycle Wealth Accumulation

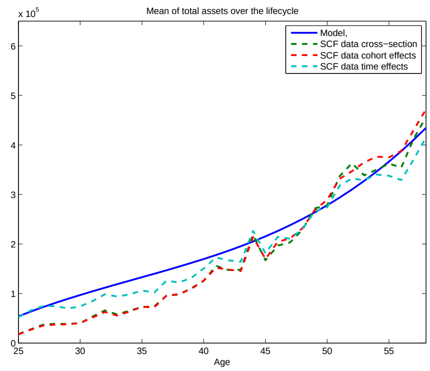

(a) Total Assets

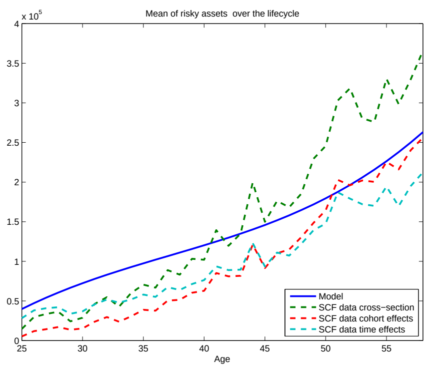

(b) Risky Assets

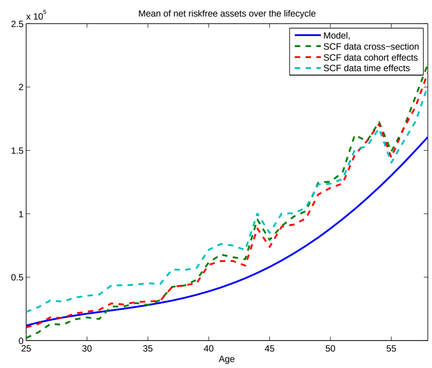

(c) Net Risk-free Assets

What are the implications of our model for life-cycle wealth accumulation? The answer is given in Figure 6: 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.[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, 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 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.

To what extent does our model generate accurate predictions for who participates in the stock market? Specifically, what are the model's implications for stock market participation across earnings and wealth groups? We first examine the model's implications across earnings groups. We report the results here and explain the mechanisms underlying them in Section 6.2.3. In Figure 7, we divide the data into three groups based on household earnings at each age: the top quartile, the bottom quartile, and the middle two quartiles taken together. For each group, we calculate stock market participation rates over the life cycle. Panels 7a and 7b of the figure represent the data after controlling respectively for cohort and time effects. The data reveal that earnings and participation are positively related: top earners participate at higher rates than the bottom two groups at every age. The model, for its part, captures this ordering of participation rates over the life cycle. For the top two groups, the model underpredicts participation at younger ages and overpredicts participation later in life, while for the bottom group, the reverse is true.[18]

Figure 7: Participation by Cross Section of Earnings: Model vs. Data

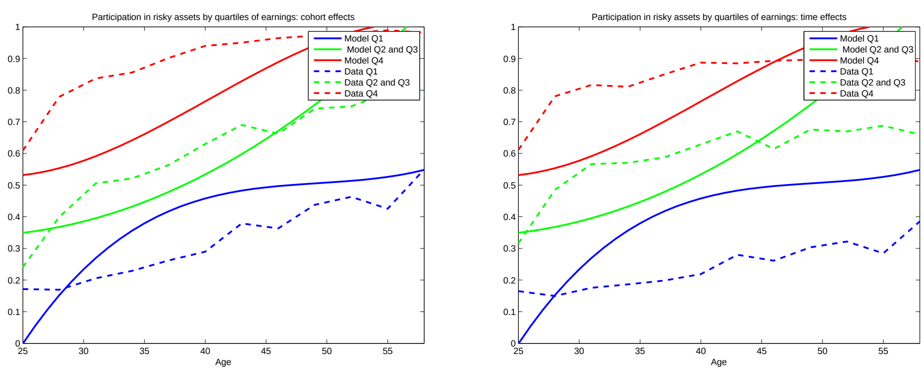

(a) Cohort Effects

(b) Time Effects

We next examine the implications of the model for stock market participa-

[18]Note that some of this discrepancy may be attributable to the differences in the way in which we construct the three groups in the model versus the data. In the model, we order households by income at each age, and divide them into the bottom 25 percent, the middle 50 percent and the top 25 percent. Because we weight the SCF data, we do not attempt to divide groups by size, but rather calculate cutoffs for the weighted data.

tion across wealth levels. We divide the population in the model and the data into three groups using the same methodology that we employed for earnings. As seen in Figure 8, the model's predictions are broadly borne out by the data. The model captures the very high and sustained rates of participation among the wealthiest households and the radically lower participation of the wealth-poorest over the entire life cycle.

Figure 8: Participation by Cross Section of Wealth: Model vs. Data

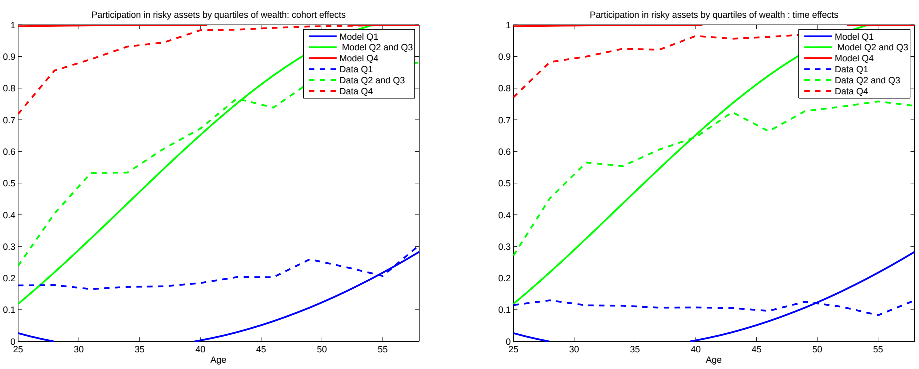

(a) Cohort Effects

(b) Time Effects

Finally, while stock market participation is the principal focus of this paper, readers may also be interested in our model's predictions for the share of stocks invested in risky assets. These are reported in Appendix A.3. Because our model does not target shares in any way, these results serve as further evidence of its empirical plausibility.

6.2 Model Mechanism

In this section, we detail the forces in the model that drive our results. We begin by describing how agents allocate time between learning and earning over the life cycle. Time allocation, in turn, is influenced by expected returns to human capital accumulation relative to stocks. We illustrate the role of this trade-off—and particularly the role of heterogeneity in human capital returns across agents—using a stylized example. The heterogeneity in returns stems from underlying differences in ability and human capital, as we discuss at the end of this section.

6.2.1 Time Spent on Human Capital Accumulation

Figure 9: Time Allocated to Human Capital over the Life Cycle

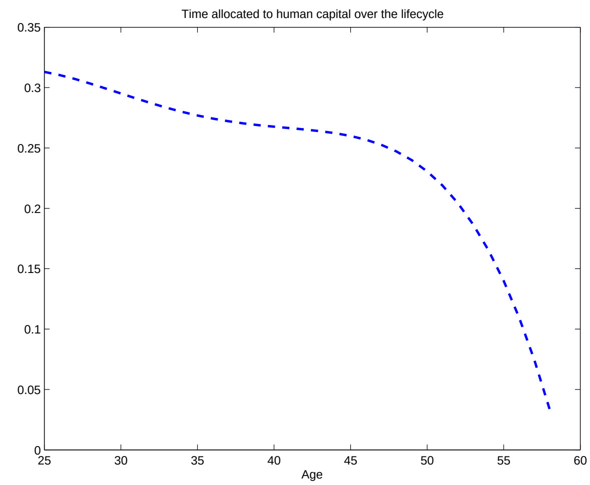

To understand our model's predictions for the path of stock market participation, a first step is to study its predictions for the trajectory of time invested in human capital accumulation (Figure 9). The key payoff from using the Ben-Porath framework is that it allows us to use observed earnings profiles over the life cycle to draw inferences about how much of the agent's time is spent accumulating skills that pay off in the future.[19] This interpretation of the data suggests that at age 25, households on average spend about a third of their time on human capital accumulation. Figure 5 shows that, at this age, only around 30 percent of all households participate in the stock market. As agents age, diminishing returns to human capital investment and a shorter horizon to recoup these returns lead them to spend less time on human capital investment. 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

[19]The abstract nature of human capital in the Ben-Porath model is precisely what allows the data to be decomposed into time spent on activities with immediate payoffs and time spent on activities that increase earnings only in the future, without limiting the interpretation of learning to only those activities that are observably associated with human capital accumulation (such as schooling).

increases with age, reaching around 80 percent at retirement.

Note that our model's predictions for time spent on human capital accumulation makes explicit an insight of Constantinides, Donaldson, and Mehra (2002). In their work, agents are endowed with a time path of income that is meant to be interpreted precisely as reflecting human capital acquisition early in life (see, e.g., p.272). Indeed, because a central aspect of our work is to model human capital acquisition explicitly, and discipline it to match the path of life-cycle earnings, our work can be seen as simply taking this point very seriously.

What underlies the trade-off between time allocated to human capital investment and stock market participation? The answer lies in the relative rates of return to each investment option. While the return to investing in the stock market is the same for all agents, the return to investing in human capital varies with each agent's endowment of ability and initial human capital as well as with their age. Critically, for some types of individuals, human capital will dominate stock market investment early in life. As we will show, these individuals would short stocks in the absence of a short sales constraint. Other types of individuals for whom the returns to human capital investment are not as high will choose to diversify by holding long positions in stocks while investing in human capital. The overall (non)participation rate at any age thus depends on the proportion of households for whom human capital provides the dominant investment payoff.

The preceding logic implies that to generate an empirically plausible prediction for stock market participation, it is critical to construct an empirically accurate representation of heterogeneity across individuals with respect to their ability and initial human capital. We achieve this by setting the Ben-Porath parameters to match earnings. We illustrate this logic below in a simplified setting to show how differences in initial human capital and ability influence the return to human capital and hence the relative payoff to financial assets.

6.2.2 Human Capital and Rate of Return Dominance: A Stylized Example

To illustrate that human capital returns may far exceed returns to financial assets, especially for individuals with relatively high ability but low initial human capital, we now employ a stylized model. The model is constructed to isolate the essential trade-off between the agent-specific, and diminishing, return to human capital investment and the common, and constant, return to financial assets. Specifically, the model features only one financial asset (whose return is constructed to be representative of a portfolio composed of both risky and risk-free financial assets) and abstracts from borrowing constraints and risks to returns on both human capital and financial assets.

Consider the problem of an agent who lives for T periods and chooses how much to consume, c_t, how much to invest in financial assets, x_t, and how to divide their endowment of a unit of time each period between human capital accumulation (learning), l_t, and working in the labor market, $1-l_t$. The agent is endowed with immutable learning ability, a, initial human capital, h_0, and initial assets, x_0. The agent solves

$$\max_{c_t, x_{t+1}, l_t} E_0 \sum_{t=0}^{T} \beta^t u(c_t)$$

subject to

$$c_t + x_{t+1} = R_t x_t + w_t h_t (1 - l_t) \ \forall t \tag{8}$$

$$h_{t+1} = (1-\delta) h_t + a(h_t l_t)^\alpha \ \forall t. \tag{9}$$

Equation (9) is the agent's budget constraint each period and Equation (9) is the law of motion for human capital. R_t is the rate of return on the financial asset, w_t is the rental rate on human capital, and δ is the rate at which human capital depreciates. The production function for human capital is as in Ben-Porath (1967), with α governing the productivity of this technology.

Denote the marginal productivity of learning by

$$MPL_t = \alpha a \frac{(h_t l_t)^\alpha}{l_t}.$$

In this setting, the price of human capital at time t is

$$P_t^h = \frac{w_t l_t}{MPL},$$

and the solution to the agent's problem yields

$$P_t^h u'(c_t) = \beta E_t[u'(c_{t+1})(w_{t+1} + P_{t+1}^h(1-\delta))].$$

The rate of return to human capital is thus

$$\frac{w_{t+1} + (1-\delta)P_{t+1}^h}{P_t^h}.$$

Figure 10 displays the rates of return in this stylized model for agents with different endowments of ability and initial human capital. We have chosen agents with ability and human capital levels from two places in the joint distribution arising from the baseline model: one with the mean level of ability and initial human capital and the other with the highest ability level but relatively low initial human capital for agents of his or her type. We see clearly from this setting that human capital investment offers rates of return that exceed those on the financial asset early in life. The difference is particularly large for individuals with high ability but relatively low human capital: Figure 10b shows that returns to human capital are as high as 30 percent for such agents early in life. Note that in our baseline setting, the agent of this type does not participate in the stock market early in life, but does so later in life. This confirms that returns to human capital can be high enough for some agents to defer participation in the stock market while they accumulate human capital. The agent with mean levels of ability and initial human capital earns a 7 percent return on human capital investment in our stylized setting. In the baseline model, this agent invests in human capital and participates in the stock market even early in life.

Figure 10: Rates of Return to Human Capital Investment by Ability and Initial Human Capital

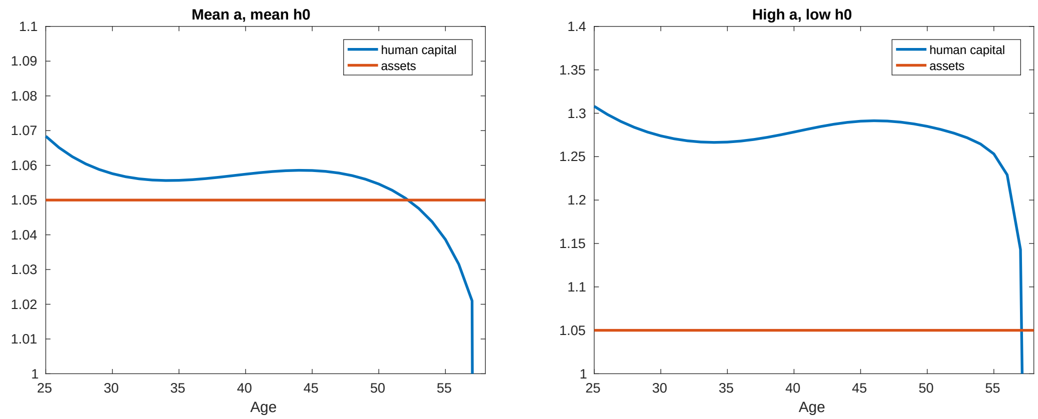

(a) Mean a, mean h_0 (b) High a, low h_0

6.2.3 The Role of Ability and Human Capital Levels

Recall that our model predicts, correctly, that those with low human capital (as manifested in low earnings) participate in the stock market at lower rates than those with high human capital (see Figure 7). We now take a closer look at the relationship between human capital (and ability) and stock market participation by comparing the distributions of these two attributes across participants and nonparticipants at various ages.

The panels of Figure 11 display the distributions of human capital and ability for participants and nonparticipants at age 25 and age 45. At age 25, the distributions of both human capital and ability for nonparticipants are to the right of the distributions for participants (Figures 11a and 11c). Why does this occur? The answer emerges from observing that, all else equal, the marginal product of time spent learning,

$$MPL_t = \alpha a \frac{(h_t l_t)^\alpha}{l_t},$$

is higher for those with greater learning ability and current human capital. Further, notice that both human capital and time spent learning are subject to diminishing returns. For any given level of ability, those with high human capital will not only obtain low marginal returns from additional human capital investment (in terms of the gains to their future earnings), but will also face high opportunity costs of doing so in the form of forgone current earnings.

Figure 11: Ability and Human Capital Distributions of Participants and Non-participants by Age

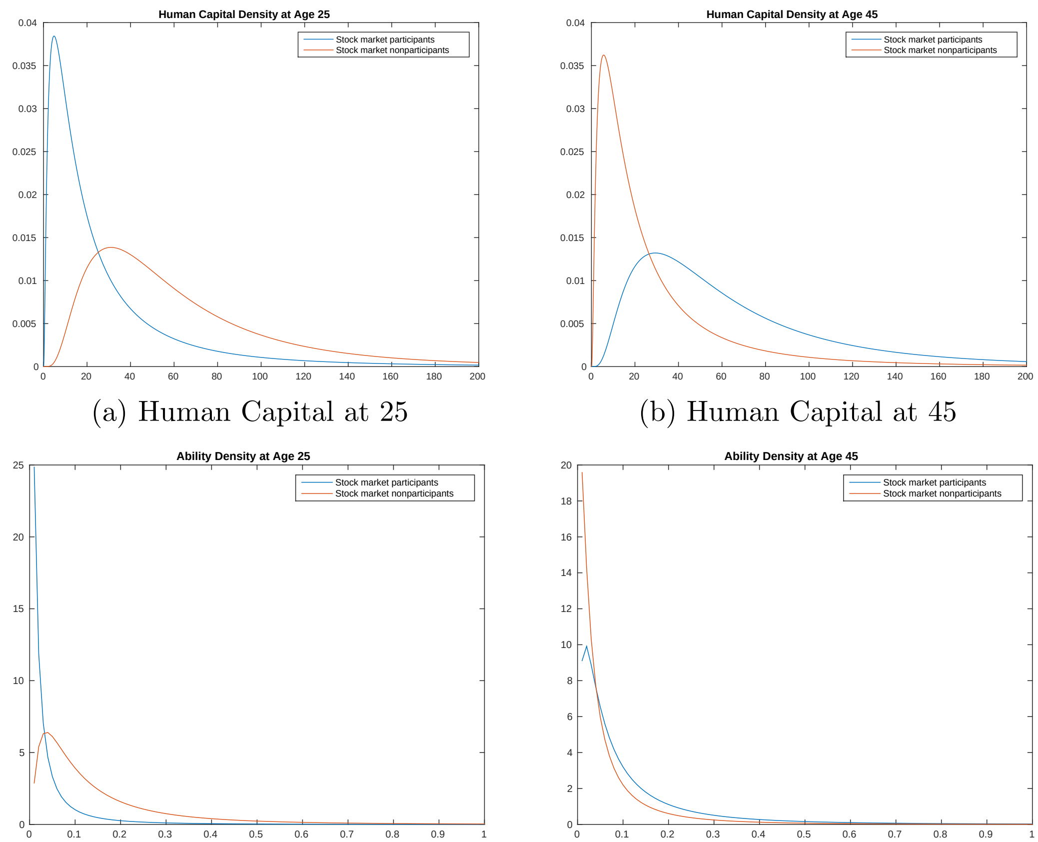

(a) Human Capital at 25

(b) Human Capital at 45

(c) Ability at 25

(d) Ability at 45

The preceding logic implies that those who begin with high ability will, all else equal, make decisions that leave them with relatively high human capital, and hence high earnings, by middle age. These individuals will also be more likely to have begun life-cycle savings and, in turn, investing in the stock market. Indeed, as seen in Figures 11b and 11d, older participants in the stock market are disproportionately those with high human capital (and hence high earnings) and high ability, while the reverse is true for nonparticipants. As a quantitative matter, these dynamics resolve themselves in favor of the strongly positive relationship between earnings, wealth, and stock market participation seen earlier in Figures 7 and 8.

To further illustrate the mechanism we describe above, we look at households with high initial wealth, defined here as being in the top 10 percent of the wealth distribution at age 25. (Note that, since these are young households, they are not likely to be rich relative to the overall population.) It turns out that young participants and nonparticipants within this group have similar ability distributions (Figure 12a). However, when it comes to human capital, participants and nonparticipants differ substantially, with the distribution of human capital among the latter being to the left of the former (Figure 12b).

The similarity in ability across participants and nonparticipants within this group allows us to ascribe differences in participation to differences in initial human capital. Given the similarity of ability between these two groups, those with relatively low human capital face higher returns to human capital investment, high enough to defer participation in stocks. By contrast, those with relatively high human capital face lower marginal returns and hence elect to also invest in stocks.

Figure 12: Distribution of Ability and Human Capital across Participants and Nonparticipants (Wealthy Households at Age 25)

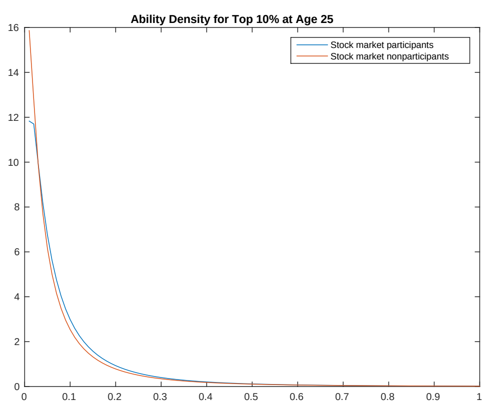

(a) Ability

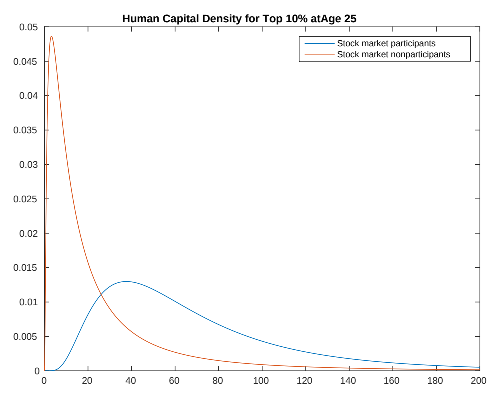

(b) Human Capital

6.3 Heterogeneity and Aggregate Stock Market Participation

We have described how an individual's endowments of ability and human capital influence their returns to human capital investment and hence their stock market participation decision. It follows that an accurate calibration of heterogeneity in these endowments is important for generating quantitatively plausible stock market participation rates in our economy. We demonstrate this by considering 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 13a.

Figure 13: Participation and Heterogeneity

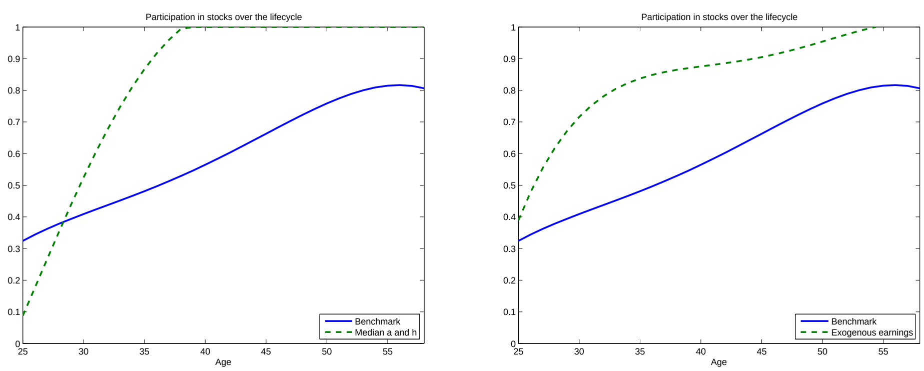

(a) Human Capital Investment with No Heterogeneity

(b) Heterogeneity with No Human Capital Investment

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 do not invest in stocks. However, by the time they reach their mid-30s, it becomes optimal for them to spend less time learning, accumulate savings, and enter the stock market. This results in participation rates rising rapidly to 100 percent at this time.

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 from Figure 13b that in a model where earnings are empirically accurate but assigned without requiring explicit human capital investment, participation is higher than in 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 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 more steeply than in the benchmark economy. We remind the reader that despite the fact that earnings are exogenous, this experiment presented in Figure 13b embodies a great deal of earnings heterogeneity. This is because we carry earnings over from our benchmark in which earnings heterogeneity results not just from idiosyncratic risk, but also from a rich specification (a continuum) of types defined by ability and human capital. In particular, this richness in types is in stark contrast to the experiment presented in Figure 13a in which, conditional on earnings shocks, there is only one expected path of earnings among all individuals of any given age. As a result, even our exogenous earnings setting allows for much greater variation in the benefits and costs of participation across agents.

We will show next in Section 6.4 that another important benefit to endogenizing human capital is that it clarifies the role played by borrowing costs in observed stock market participation.

6.4 Short Sales Constraints, Borrowing Constraints, and Aggregate Stock Market Participation

We showed in section 6.2.2 that, for many young households, the rate of return to human capital investment can far exceed the rate of return to financial assets. These young households find it optimal to borrow to finance consumption while spending time investing in human capital and would in fact short stocks if they could. Indeed, as we will show in this section, in our baseline environment it is short sales constraints on stocks and not borrowing constraints that are material in creating a set of households with zero holdings of stocks.

We illustrate the role of the short sales constraint by asking how agents in our benchmark model would behave if this constraint were relaxed. Figure 14a shows the fraction of households that would hold short and long positions over the life cycle and Figure 14b compares the fraction holding long positions in the absence of a short sales constraint to the participation rate in the benchmark.

Figure 14: Short Sales Constraint and Participation

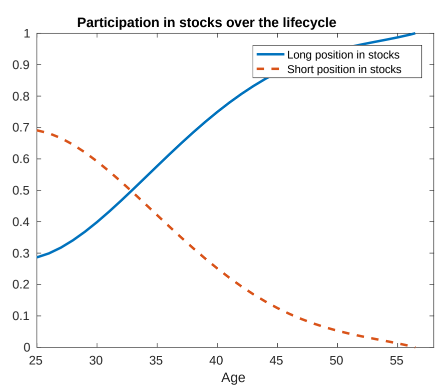

(a) No Short Sales Constraint

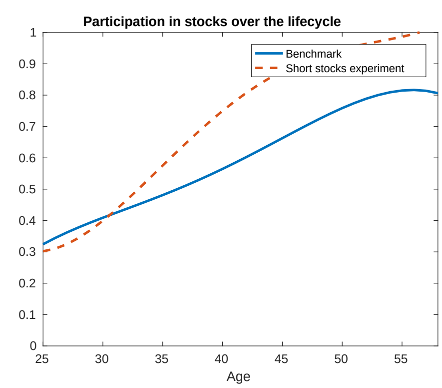

(b) Long Positions: No Short Sales Constraint vs. Benchmark

Observe from Figure 14a that early in life, nearly 70 percent of households would choose to hold short positions in stocks. This is almost exactly the fraction of nonparticipants in our benchmark case. For these households, the return to investing in human capital is high enough for it to be optimal for them to hold short positions in stocks to finance additional consumption while spending even more time on human capital investment than in the benchmark. Figure 14b shows that, in the absence of a short sales constraint, the fraction of households holding long positions in stocks eventually exceeds the participation rate in the benchmark. Without the short sales constraint, households have an additional channel by which to finance consumption while accumulating human capital. This leads to faster accumulation, higher earnings, and thereby to a higher fraction of households eventually holding long positions than in the benchmark. Put another way, in our baseline model, a binding short sales constraint is what leads a nontrivial share of households to maintain no exposure to the stock market early in life.

It is natural to ask whether nonparticipation in our model is also driven by the presence of an interest-rate wedge that makes borrowing costly. We now demonstrate that—while critical to obtaining nonparticipation in an exogenous-earnings setting—borrowing costs have little effect on stock market participation in our setting. To illustrate this, we consider a case in which there is no wedge at all between the interest rate on savings and borrowing. The blue and dashed red lines in Figure 15 compare, respectively, participation in our benchmark to participation in a model that is identical to our benchmark, except that there is no wedge between the borrowing and saving rate. Observe that households do not significantly change their stock market participation despite having access to cheaper credit. In contrast, the yellow and dashed purple lines in Figure 15 show that, in a version of our setting where earnings are exogenous, there is a large difference in stock market participation rates with and without the wedge. In particular, participation reaches nearly 100 percent early in the life cycle in the absence of a wedge (as in Davis, Kubler, and Willen, 2006).

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 or not it makes sense to borrow to invest in stocks. As long as investing in stocks yields a higher payoff than the cost of borrowing, households will choose to

Figure 15: The Role of the Borrowing Wedge in Stock Market Participation

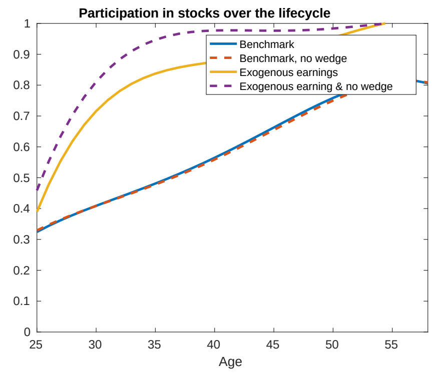

do so. Thus participation in that setting would be high in the absence of a sufficiently high interest rate on borrowing.

In contrast, when households can invest in human capital, as in our model, but borrowing is costly, many young households would short stocks if they could, instead of borrowing to hold long positions in stocks. However, when the wedge on borrowing is removed in this setting, shorting stocks is no longer attractive because borrowing is now both less costly and less risky. However, most households still do not borrow to hold long positions in stocks. This is because human capital is the dominant investment for these households, and the lower borrowing costs encourage them to borrow more and spend additional time on human capital accumulation. Borrowing to finance current consumption makes future consumption risky, especially in the face of uninsurable earnings risk. In sum, households that have incurred debts to ease risky human capital accumulation do not want to engage in further leveraged risk-taking to hold a long position in the stock market.[20]

[20]The experiment with no wedge gives going long on stocks the "best chance" in the sense that households are now faced with an environment with not only no participation costs and no positive correlation between earnings and stock market returns, but also a cheap and risk-free way to borrow. We do obtain a small fraction of households that borrow using the risk-free asset and hold long positions in stocks, but the fact that the majority of households still elect to stay away from stocks reiterates that endogenous human capital investment is indeed the driving force behind our results. However, as is common in the literature, compared to the data, our model generates a much smaller fraction of households with positive net worth and no stockholdings. Resolving this particular aspect of portfolio choice likely requires additional impediments to trade or departures from fully rational

We have discussed in detail the implications of the cost of borrowing. It is also of interest to understand how the limit on borrowing itself alters outcomes. Specifically, is limited participation in the stock market early in life driven by borrowing constraints that bind for a large fraction of individuals? The answer is no, for the same reasons as discussed above. Consider an experiment in which agents in our model face a more generous borrowing limit, twice the limit in the benchmark. Some agents will now choose to borrow more than they did before to smooth consumption while investing in human capital. Notice that the higher leverage *increases* risk for the borrower: their debt obligations are larger while the rewards from the greater investment in human capital remain risky. There are few households for whom it will be optimal to pursue this strategy. Moreover, for these individuals, borrowing only increases their incentive to reduce exposure to risk, i.e., to stay away from stocks. The forces are reflected in Figure 16, which shows that a more generous borrowing constraint does not have a large impact on overall stock market participation.

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

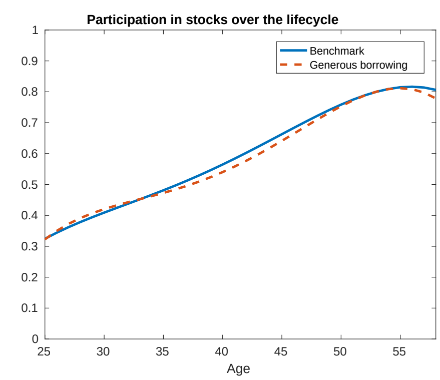

7 Conclusion

The contribution of this paper is to demonstrate that stock market participation—both over the life cycle and across earnings and wealth cross-sections—can be well accounted for once human capital investment is allowed

behavior, and hence remains a task for future work.

and disciplined to match labor earnings. Our results suggest, further, that the power of financial frictions such as borrowing constraints to limit overall equity-market participation depends strongly on the availability of human capital as an investment option. Indeed, we show that most young investors would, if allowed, short the stock market to finance human capital acquisition.

Our approach further underscores the importance of the broader mechanism identified in Constantinides, Donaldson, and Mehra (2002): that the solution to the equity premium likely lies in the constraints limiting the ability of the young to bring resources forward from the future. In our setting, this constraint is on short-selling equity. Note, interestingly, that the incorporation of human capital investment also changes the role played by the constraint on debt: our results suggest that relaxing the debt constraint would, all else equal, lower the equity premium. This occurs because the young would use debt to finance consumption while accumulating human capital instead of attempting to short stocks for the same purpose. Quantifying the strength of this force requires a general equilibrium approach, of course, and so is left for future work.

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?," 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 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," *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 Short Sales Restrictions," *Journal of Finance*, 33(2), 589–601.

Browning, M., L. P. Hansen, and J. J. Heckman (1999): "Micro data and general equilibrium models," *Handbook of macroeconomics*, 1, 543–633.

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*, pp. 439–482. University of Chicago Press.

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*, pp. 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, 1–53.

Deaton, A. (1985): "Panel Data from Time Series of Cross-sections," *Journal of Econometrics*, 30(1–2), 109–126.

Friedman, M. (1962): in *Capitalism and Freedom*University of Chicago Press.

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.

——— (2008): "Asset Pricing with Limited Risk Sharing and Heterogeneous Agents," *Review of Financial Studies*, 21(1), 415–448.

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. J. Lucas (1996): "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104(3), 443–487.

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

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

——— (2011): "Sources of Lifetime Inequality," *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*, 121(2), 427–447.

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

Lucas, D. J. (1994): "Asset pricing with undiversifiable income risk and short sales constraints: Deepening the equity premium puzzle," *Journal of Monetary Economics*, 34(3), 325–341.

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.

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," Working Paper 6185, National Bureau of Economic Research.

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

Storesletten, K., C. I. Telmer, and A. Yaron (2007): "Asset pricing with idiosyncratic risk and overlapping generations," *Review of Economic Dynamics*, 10(4), 519–548.

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

A Appendix

A.1 Estimation of Participation over the Life Cycle

We estimate life-cycle profiles of stock market participation under two alternative identifying assumptions: i) time effects are zero (cohort effects matter) and ii) cohort effects are zero (time effects matter). Table 2 shows the results of the estimation under the first assumption and Table 3 reports the results under the second assumption.

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: 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		

A.2 Earnings

We use a synthetic cohort approach to compute life-cycle earnings profiles from the CPS following Ionescu (2009). We track the earnings of household heads belonging to the cohort aged 23–27 in 1969. We use this age range because it represents the beginning of that portion of life in which households make nontrivial investments in financial assets. 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. Life-cycle profiles for the mean, inverse skewness, and Gini of earnings are shown in Figure 17.[21]

[21] 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.

Figure 17: Earnings Statistics (CPS)

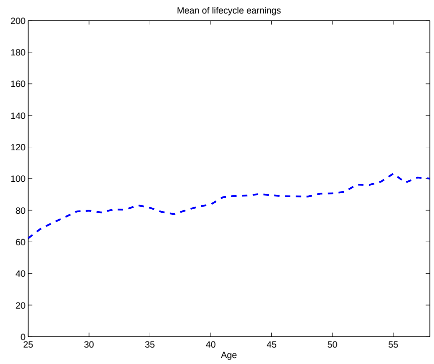

(a) Mean

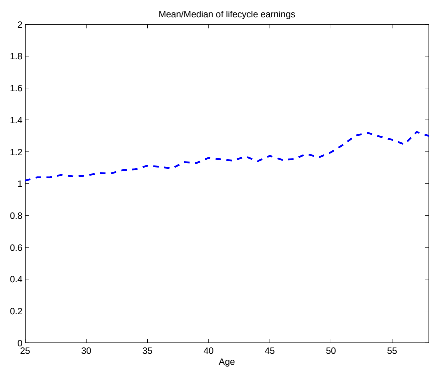

(b) Skewness

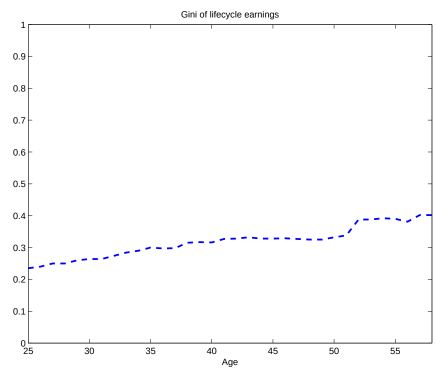

(c) Gini

A.3 Implications for the Share of Stocks in Household Portfolios

We compare our model's implications for the share of wealth invested in stocks over the life cycle to our estimates of shares from the SCF data. As with participation, we report two estimates, one adjusting for time effects and the other for cohort effects.[22] The results are reported in Figure 18.

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

[22]The details of the estimation are available upon request.

Figure 18: Fraction of Stocks in Household Portfolio

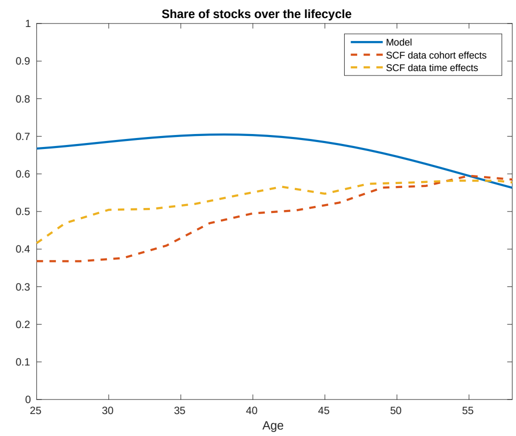

the presence of human capital remains far below 100 percent. Importantly, this occurs despite the fact that households in our model retain the ability to increase their labor supply to undo poor stock market returns. Third, the hump-shaped profile for shares generated by our model is empirically more 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.

Further, an interesting observation that follows from our model's results is that the forces that determine participation are separable from those determining shares. While endogenous human capital investment and heterogeneity in its returns are key to explaining participation, these forces matter little for shares, as we show next. On the other hand, risk and attitudes toward risk do not greatly influence participation behavior, as we will show in Appendix A.5, but are important determinants of shares.

We begin by studying the effect of earnings heterogeneity—which was key to explaining participation—on the share of household wealth invested in stocks. 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 19, 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 19: Life-Cycle Stock Market Shares with No Heterogeneity in Ability and Initial Human Capital

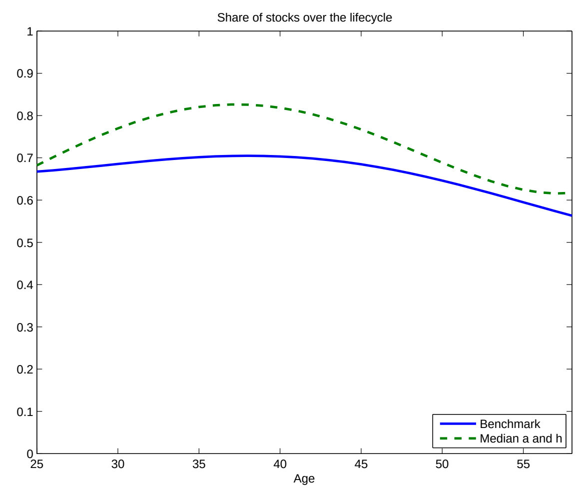

Having shown earlier that endogenous human capital dramatically limited the role of borrowing costs for stock market participation, it is of interest to see if this applies to the intensive margin as well. The answer is no. The reason is this: given participation, the question for a household is the 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. This is driven home by the fact that in our model, borrowing costs have almost no effect on the risk exposure that households choose (Figure 20).

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

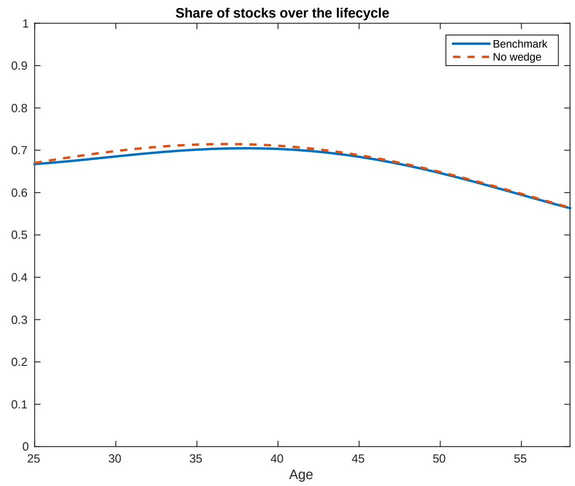

Having asserted that risk considerations are critical for explaining shares of wealth held in stocks, we can be more explicit. In Figure 21, we see that when stocks are risky, households that engage in the stock market reduce their holdings at all ages. In the case of higher-than-baseline riskiness of 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.

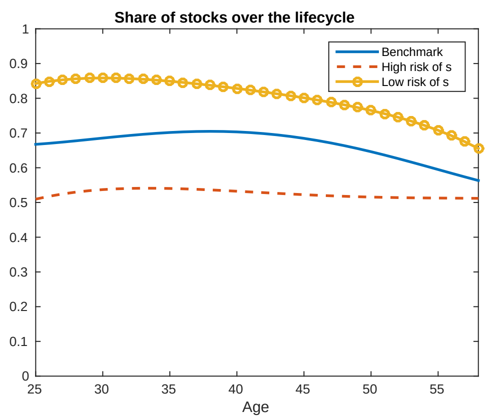

Figure 21: Shares under Higher or Lower Risk of Stocks

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 22, this is exactly what happens.

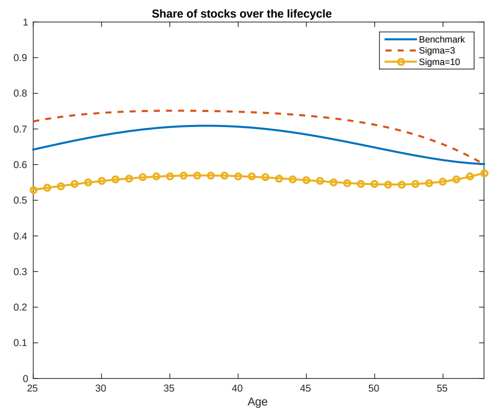

Figure 22: Shares under High and Low Risk Aversion

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. 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.

A.4 The Role of the Elasticity of Human Capital Production

The rate of return to human capital is affected not merely by ability and initial human capital, but also by the elasticity of human capital accumulation with respect to time spent learning, α. How would changes in the returns to human capital investment affect stock market participation in the baseline economy? To answer this, we consider the effect of making human capital investment less productive (α=0.5) or more productive (α=0.9), holding fixed all other parameters. Figure 23 provides the answer: the relative attractiveness of the two investment options plays a decisive role for stock market participation. In other words, if agents in our economy were to be faced with a more(less) productive human capital technology, their stock market participation would be lower(higher) than in the baseline.

Figure 23: The Effect of the Elasticity of Human Capital Production in the Baseline Model

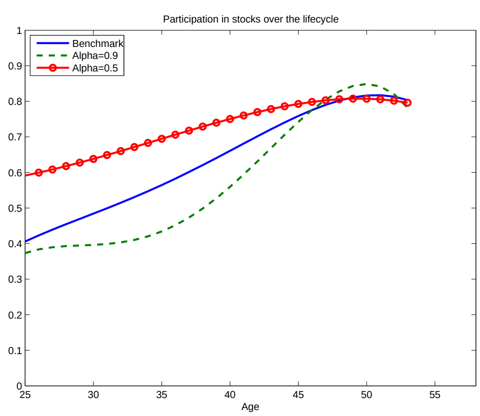

A clear implication of the logic of our model is that the better the technology for learning, the less attractive stock market investment will be, all else equal. After all, if the earnings we observe in the data were generated by a more productive human capital technology than in the benchmark, then 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 bench-

[23]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.

mark. The marginal densities for ability and initial human capital obtained from the recalibration are to the left of those in the benchmark.

Figure 24: The Effect of the Elasticity of Human Capital Production in a Recalibrated Model

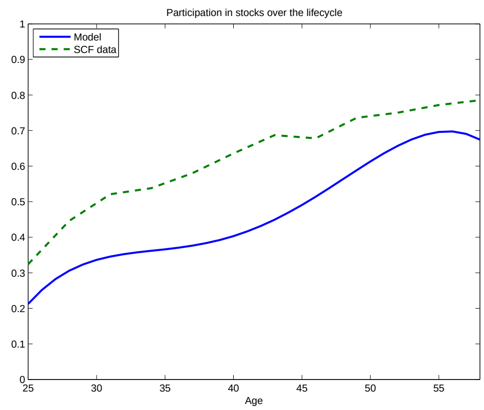

The main results are reported in Figure 24. Participation in the stock market is indeed much lower than in the benchmark. This is for two reasons. First, as we have stressed throughout, this is consistent with the idea that human capital competes with financial assets as an investment option. With a 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.

A.5 The Role of Risk

Our baseline model 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 study both the effect of changing the risk of stocks and the effect of changing agents' risk aversion in our setting.

A.5.1 Stock Market Risk

Figure 25: The Role of Risk in Stock Market Participation

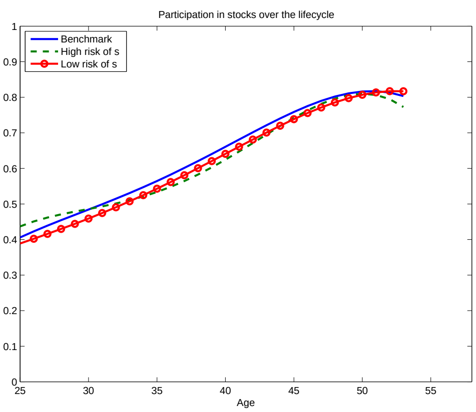

(a) Stock Market Risk

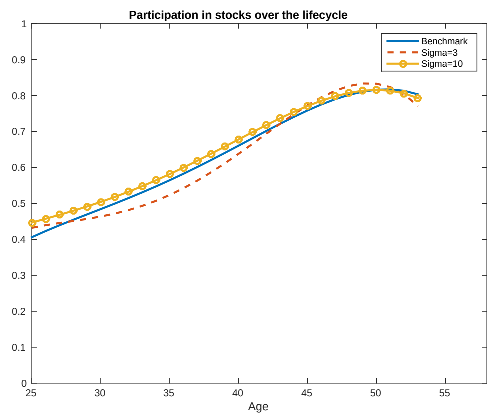

(b) Risk Aversion

The stock market, while clearly offering 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 25a, 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.

A.5.2 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 25b.

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 as being critical determinants of the participation decision.