THE JOURNAL OF FINANCE • VOL. LX, NO. 2 • APRIL 2005

Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence

FRANCISCO GOMES and ALEXANDER MICHAELIDES*

ABSTRACT

We show that a life-cycle model with realistically calibrated uninsurable labor income risk and moderate risk aversion can simultaneously match stock market participation rates and asset allocation decisions conditional on participation. The key ingredients of the model are Epstein–Zin preferences, a fixed stock market entry cost, and moderate heterogeneity in risk aversion. Households with low risk aversion smooth earnings shocks with a small buffer stock of assets, and consequently most of them (optimally) never invest in equities. Therefore, the marginal stockholders are (endogenously) more risk averse, and as a result they do not invest their portfolios fully in stocks.

IN THIS PAPER, WE PRESENT A LIFE-CYCLE ASSET allocation model with intermediate consumption and stochastic uninsurable labor income that provides an explanation for two very important empirical observations: low stock market participation rates in the population as a whole, and moderate equity holdings for stock market participants.

Our life-cycle model integrates three main motives that have been identified as quantitatively important in explaining individual and aggregate wealth accumulation. First, a precautionary savings motive driven by the presence of undiversifiable labor income risk (Deaton (1991) and Carroll (1992, 1997)). Second, pension income is lower than mean working-life labor income, implying that saving for retirement becomes important at some point in the life cycle. The combination of precautionary and retirement saving motives has recently been shown to generate realistic wealth accumulation profiles over the life cycle.[1] Third, we explicitly incorporate a bequest motive that has recently been

*Francisco Gomes is from the London Business School; Alexander Michaelides is from the London School of Economics and CEPR. We thank Viral Acharya, Orazio Attanasio, Ravi Bansal, Michael Brennan, Joao Cocco, Pierre Collin-Dufresne, Steve Davis, Karen Dynan, Bill Dupor, Lorenzo Forni, Joao Gomes, Rick Green, Luigi Guiso, Michael Haliassos, Burton Hollifield, Urban Jermann, Deborah Lucas, Pascal Maenhout, Monica Paiella, Valery Polkovnichenko, Ken Singleton, Nicholas Souleles, Harald Uhlig, Raman Uppal, Annette Vissing-Jørgensen, Tan Wang, Paul Willen, Amir Yaron, Steve Zeldes, Harold Zhang, two anonymous referees, and seminar participants at Carnegie Mellon, Columbia, the Ente-Einaudi Center, the Federal Reserve Board, H.E.C., Leicester, LBS, LSE, Tilburg, UCL, UNC Chapel Hill, Wharton, the 2002 NBER Summer Institute, MFS and SED meetings for helpful comments. Previous versions of this paper have circulated with the title: "Life-Cycle Asset Allocation: A Model with Borrowing Constraints, Uninsurable Labor Income Risk and Stock Market Participation Costs." We are responsible for any remaining errors.

[1] See, for instance, Hubbard, Skinner, and Zeldes (1995), Carroll (1997), Gourinchas and Parker (2002), Dynan, Skinner, and Zeldes (2002), and Cagetti (2003).

shown to be important in matching the skewness of the wealth distribution (de Nardi (forthcoming) and Laitner (2002)).

More recently, life-cycle models incorporating some (or all) of these motives have been extended to include an asset allocation decision, both in an infinite horizon[2] and in a finite horizon, life-cycle setting.[3] However, several important predictions of these models are still at odds with empirical regularities. First, low stock market participation in the population (Mankiw and Zeldes (1991)) persists. The latest Survey of Consumer Finances (2001) reports that only 52% of U.S. households hold stocks either directly or indirectly (through pension funds, for instance), while these models predict that, given the equity premium, all households should participate in the stock market as soon as saving takes place. Second, households in the model invest almost all of their wealth in stocks, in contrast to both casual empirical observation and to formal empirical evidence (see Poterba and Samwick (1999) or Ameriks and Zeldes (2001), for instance).

We develop a life-cycle asset allocation model that tries to address these two puzzles. We argue that it is possible to simultaneously match stock market participation rates and asset allocation conditional on participation, with moderate values of risk aversion (between one and five), and without extreme assumptions about the level of background risk. Our model has three key features. First, we include a fixed entry cost for households that want to invest in risky assets for the first time. A large literature has concluded that some level of fixed costs seems to be necessary to improve the empirical performance of asset pricing models.[4] Since the excessive demand for equities predicted by asset allocation models is merely the portfolio-demand manifestation of the equity premium puzzle, introducing a fixed cost in the model seems to be a natural extension. Moreover, recent empirical work suggests that small entry costs can be consistent with the observed low stock market participation rates (see Paiella (2001), Degeorge et al. (2002), and Vissing-Jørgensen (2002b)).

The other two key features are motivated by the (perhaps surprising) implication of the model that participation rates are an increasing function of risk aversion, at least over a wide range of parameter values. Specifically, changing risk aversion generates two opposing forces for determining the participation decision. On the one hand, more risk-averse households optimally prefer to invest a smaller fraction of their wealth in stocks. On the other hand, risk aversion determines prudence and more prudent consumers accumulate significantly

[2] See, for example, Telmer (1993), Lucas (1994), Koo (1998), Heaton and Lucas (1996, 1997, 2000), Polkovnichenko (2000), Viceira (2001), and Haliassos and Michaelides (2003).

[3] See, for instance, Cocco, Gomes, and Maenhout (1999), Cocco (2000), Campbell et al. (2001), Hu (2001), Storesletten, Telmer, and Yaron (2001), Davis, Kubler, and Willen (2002), Dammon, Spatt, and Zhang (2001, forthcoming), Polkovnichenko (2002), Yao and Zhang (forthcoming), and Gomes and Michaelides (2003). Bertaut and Haliassos (1997) and Constantinides, Donaldson, and Mehra (2002) analyze three-period models where each period amounts to 20 years.

[4] See, among others, Constantinides (1986), Aiyagari and Gertler (1991), He and Modest (1995), Saito (1995), Heaton and Lucas (1996), Luttmer (1996, 1999), Basak and Cuoco (1998), and Vayanos (1998).

more wealth over the life cycle. We show that the higher wealth accumulation motive dominates for moderate coefficients of relative risk aversion (RRA) (i.e., not greater than five). As a result, the less risk-averse investors have a weaker incentive to pay the fixed cost. This explains why previous attempts to match participation rates in the context of a life-cycle model were fairly unsuccessful. If we try to match asset allocation decisions by assuming high values of risk aversion, the implied participation rates are counterfactually high (e.g., Campbell et al. (2001)). Motivated by this result, we allow for preference heterogeneity in the population, the second key feature of the model. As argued before, since the less risk-averse investors accumulate less wealth over the life cycle, the majority optimally chooses not to pay the fixed cost. Therefore, *endogenously* stock market participants tend to be the more risk-averse investors, and consequently, even after paying the fixed cost, they do not invest their portfolios fully in equities.

The final important feature of the model is the assumption of Epstein–Zin preferences, which allows us to separate risk aversion from the elasticity of intertemporal substitution (EIS). In the context of a life-cycle model with labor income, wealth accumulation is a crucial determinant of both the stock market participation and the asset allocation decision. Within the power utility framework, households with low risk aversion also have a high EIS. Given that the expected return from investing in the stock market is higher than the discount rate, a higher EIS increases savings. As a result, even though the less risk-averse agents would not save much for precautionary reasons, they would have a strong incentive to save for retirement (and for a potential bequest motive). Thus, breaking the link between risk aversion and the EIS is crucial for delivering predictions that are consistent with the observed empirical evidence.

Therefore, in our model, households with very low risk aversion and low (moderate) EIS smooth idiosyncratic earnings shocks with a small buffer stock of assets, and most of them never invest in equities (thus behaving as in the Deaton (1991) infinite-horizon model).[5] This seems to describe adequately the behavior of a large fraction of the U.S. population that retires without significant financial assets (and does not participate in the stock market). Within the low EIS and low risk aversion group, only a small fraction owns stocks, and they do so only as they get close to retirement. On the other hand, investors with high prudence and high EIS are the ones who participate in the stock market from early on, since they accumulate more wealth and therefore have a stronger incentive to pay the fixed cost. Therefore, the marginal stockholders are (endogenously) more risk averse and as a result they do not invest their portfolios fully in stocks.

The heterogeneous agent model can simultaneously match the stock market participation rate and the average equity allocation conditional on participation, from the Survey of Consumer Finances (SCF). The life-cycle profile of the

[5] It is important to point out that we do not need heterogeneity in the EIS to obtain our results. As we will show, the less risk-averse investors can have the same EIS as the more risk-averse, just as long as this value is not too high (hence the need for Epstein–Zin preferences).

participation rate is also very close to the one observed in the data. On the negative side, the model still counterfactually predicts that young households that have already paid the participation cost will invest most of their portfolio in equities.[6] Finally, the degree of heterogeneity in the wealth distribution is quite comparable to the one observed in the data.

The rest of the paper is organized as follows. Section I summarizes results from the existing empirical literature on life-cycle asset allocation, while Section II outlines the model and calibration. In Sections III and IV, we discuss the results in the absence and presence of the fixed entry cost, respectively. Finally, Section V concludes.

I. Empirical Evidence on Life-Cycle Asset Allocation and Stock Market Participation

In most industrialized countries, stock market participation rates have increased substantially during the last decade. Nevertheless, a large percentage of the population still does not own any stocks (either directly or indirectly through pension funds). Moreover, even those households that do own stocks still invest a significant fraction of their portfolios in alternative assets.

Figures 1A and B summarize evidence reported in Ameriks and Zeldes (2001).[7] The results are sensitive to the identifying assumptions regarding time versus cohort effects. Time effects can arise, for example, from changes in market structure (e.g., transaction costs or information) or because investors use past returns to forecast future expected returns. Cohort effects can be due to differences in lifetime earnings potential, or different institutional settings (e.g., the social security system). Since age (a), time (t), and cohort (c, birth year) are linearly dependent ($a \equiv t - c$), when constructing age profiles, it is impossible to simultaneously identify time and cohort effects.

Figure 1A plots the average life-cycle equity holdings for stock market participants (as a share of total financial wealth), based on the 1989, 1992, 1995, and 1998 samples of the SCF. Although the life-cycle profiles are very sensitive to the inclusion of time dummies versus the inclusion of cohort dummies, the average stock holdings are significantly below 100% in both cases.[8] Figure 1B plots the corresponding stock market participation rate, obtained by running a probit regression on the same data. These results are less sensitive to the choice of time versus cohort dummies. As expected, a very large fraction of the population does not own equities. In both cases the participation rate gradually

[6] Hu (2001) and Yao and Zhang (forthcoming) are able to reduce the equity demand of young households by considering models with an explicit housing allocation decision.

[7] Guiso, Haliassos, and Japelli (2002) obtain similar conclusions using cross-sectional information for five different countries (United States, United Kingdom, The Netherlands, Germany, and Italy).

[8] The OLS regression with cohort effects predicts a share of financial wealth invested in stocks above 100% for the oldest age groups. This is just the result of imposing the same cohort effects on the full sample as in fact, in every individual cross-section, these age groups never invest more than 60% of their wealth in equities.

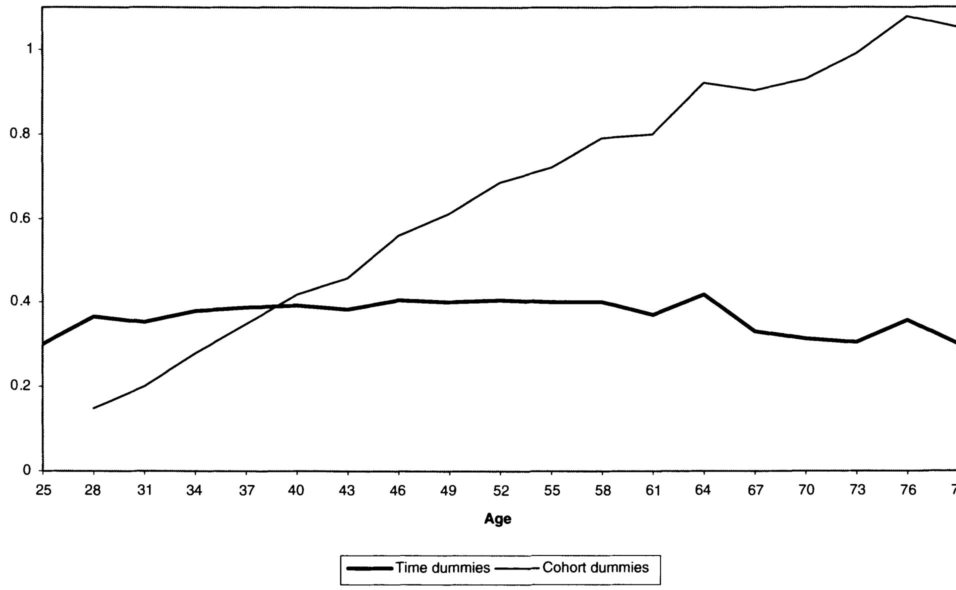

Figure 1A. Equity holdings as a fraction of total financial wealth for stock market participants. The results are taken from Ameriks and Zeldes (2001), and they are obtained from OLS regressions with age dummies and either time or cohort dummies. The data includes the 1989, 1992, 1995, and 1998 samples of the SCF.

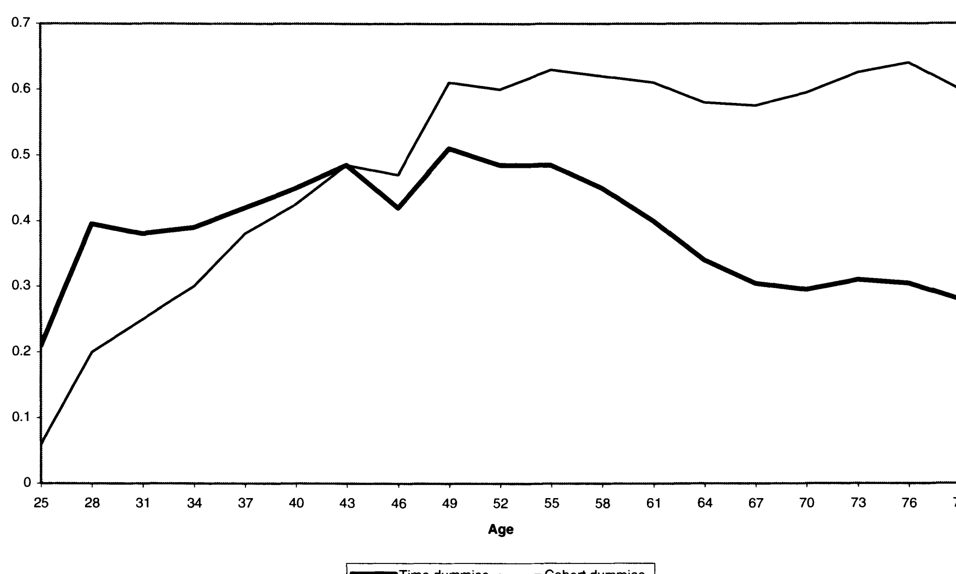

Figure 1B. Stock market participation rate. The results are taken from Ameriks and Zeldes (2001), and they are obtained from probit regressions with age dummies and either time or cohort dummies. The data includes the 1989, 1992, 1995, and 1998 samples of the SCF.

increases until approximately age 50. When including cohort dummies, the profile is flat after age 50, while with time dummies it is decreasing. Ameriks and Zeldes (2001) obtain the same results after redoing the analysis using TIAA-CREF data from 1987 to 1996, and so do Poterba and Samwick (1999), using SCF data.

We can summarize the existing evidence as follows.[9] First, the stock market participation rate in the U.S. population is close to 50%. Using the latest numbers from the SCF, we compute it as 51.9% (details given in Appendix C). Second, participation rates increase during working life and there is some evidence suggesting that they might decrease during retirement, although this might also be due to cohort effects. Third, conditional on stock market participation, households invest a large fraction of their financial wealth in alternative assets. According to the latest numbers from the SCF, the average equity holdings as a share of financial wealth for stock market participants is 54.8%. Fourth, there is no clear pattern of equity holdings over the life cycle.

II. The Model

A. Preferences

Time is discrete and t denotes adult age, which following the typical convention in this literature, corresponds to effective age minus 19. Each period corresponds to 1 year and agents live for a maximum of 81 (T) periods (age 100). The probability that a consumer/investor is alive at time $(t+1)$ conditional on being alive at time t is denoted by p_t (p_0 equal to 1).

Households have Epstein–Zin utility functions (Epstein and Zin (1989)) defined over one single nondurable consumption good. Let C_t and X_t denote respectively consumption level and wealth (cash on hand) at time t. Then, the household's preferences are defined by

$$V_t = \left\{ (1-\beta p_t)C_t^{1-1/\psi} + \beta E_t \left[p_t\left[V_{t+1}^{1-\rho}\right] + (1-p_t)b\frac{(X_{t+1}/b)^{1-\rho}}{1-\rho} \right]^{\frac{1-1/\psi}{1-\rho}} \right\}^{\frac{1}{1-1/\psi}}, \quad (1)$$

where ρ is the coefficient of RRA, ψ is the EIS, β is the discount factor, and b determines the strength of the bequest motive.[10] Given the presence of a bequest motive, the terminal condition for the recursive equation (1) is:

$$V_{T+1} \equiv b\frac{(X_{T+1}/b)^{1-\rho}}{1-\rho}. \quad (2)$$

[9] We must point out that several papers have contributed to this research. See, for example, Guiso, Jappelli, and Terlizzese (1996) (who focus mostly on the impact of background risk on asset allocation), King and Leape (1998), Heaton and Lucas (2000) and the papers in the volume edited by Guiso, Haliassos, and Japelli (2002).

[10] For more motivation and details on the modeling of bequest motives in life-cycle models see Laitner (2002), or de Nardi (forthcoming).

B. Labor Income Process

Following the standard specification in the literature, the labor income process before retirement is given by

$$Y_{it} = P_{it} U_{it}, \tag{3}$$

$$P_{it} = \exp(f(t, Z_{it})) P_{it-1} N_{it}, \tag{4}$$

where $f(t, Z_{it})$ is a deterministic function of age and household characteristics Z_{it}, P_{it} is a permanent component with innovation N_{it}, and U_{it} is a transitory component. We assume that $\ln U_{it}$ and $\ln N_{it}$ are independent and identically distributed with mean $\{-0.5 * \sigma_u^2, -0.5 * \sigma_n^2\}$, and variances σ_u^2 and σ_n^2, respectively. The log of P_{it} evolves as a random walk with a deterministic drift, $f(t, Z_{it})$.

For simplicity, retirement is assumed to be exogenous and deterministic, with all households retiring in time period K, corresponding to age 65 ($K = 46$). Earnings in retirement ($t > K$) are given by $Y_{it} = \lambda P_{iK}$, where λ is the replacement ratio (a scalar between zero and one). This specification, also standard in this literature, considerably facilitates the solution of the model, as it does not require the introduction of an additional state variable (see Section II.E).

Durable goods, and in particular housing, can provide an incentive for higher spending early in life. Modeling these decisions directly is beyond the scope of the paper, but nevertheless we take into account these potential patterns in life-cycle expenditures. Using the Panel Study of Income Dynamics, for each age (t) we estimate the percentage of household income that is dedicated to housing expenditures (h_t) and subtract it from the measure of disposable income.[11] More details on this estimation are given below, when we discuss the calibration of the model.

C. Financial Assets

The investment opportunity set is constant and there are two financial assets, one riskless (Treasury bills or cash) and one risky (stocks). The riskless asset yields a constant gross return, R^f, while the return on the risky asset (denoted by R_t^S) is given by

$$R_{t+1}^S - R^f = \mu + \varepsilon_{t+1}, \tag{5}$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$.

We allow for positive correlation between stock returns and earnings shocks. We let $\phi_N (\phi_U)$ denote the correlation coefficient between stock returns and permanent (transitory) income shocks.

Before investing in stocks for the first time, the investor must pay a fixed lump sum cost, $F * P_{it}$. This entry fee represents both the explicit transaction cost from opening a brokerage account and the (opportunity) cost of acquiring information about the stock market. The fixed cost (F) is scaled by the level of

[11] A similar approach is taken by Flavin and Yamashita (2002) in a model without labor income.

the permanent component of labor income (P_{it}), as this significantly simplifies the solution of the model. However, this specification is also motivated by the interpretation of the entry fee as the opportunity cost of time.

D. Wealth Accumulation

We denote cash on hand as the liquid resources available for consumption and saving. We define a dummy variable I_P that is equal to one when the fixed entry cost is incurred for the first time and is zero otherwise. The household's next period cash on hand ($X_{i,t+1}$) is given by

$$X_{i,t+1} = S_{it} R^S_{t+1} + B_{it} R^f + (1 - h_t) Y_{i,t+1} - FI_P P_{i,t+1}, \tag{6}$$

where S_{it} and B_{it} denote respectively stock holdings, and riskless asset holdings (cash) at time t, and h_t is the fraction of income dedicated to housing-related expenditures. Since the household must allocate cash on hand (X_{it}) between consumption expenditures (C_{it}) and savings we also have

$$X_{it} = C_{it} + S_{it} + B_{it}. \tag{7}$$

Finally, we prevent households from borrowing against their future labor income. More specifically we impose the following restrictions:

$$B_{it} \geq 0, \tag{8}$$

$$S_{it} \geq 0. \tag{9}$$

E. The Optimization Problem and Solution Method

The complete optimization problem is then

$$\max_{\{S_{it}, B_{it}\}_{t=1}^T} E(V_0), \tag{10}$$

where V_0 is given by equations (1) and (2) and is subject to the constraints given by equations (5) to (9), and to the stochastic labor income process given by (3) and (4) if $t \leqslant K$, and $Y_{it} = \lambda P_{iK}$ if $t > K$.

Analytical solutions to this problem do not exist. We therefore use a numerical solution method based on the maximization of the value function to derive the optimal decision rules. The details are given in Appendix A, and here we just present the main idea. We first simplify the solution by exploiting the scale-independence of the maximization problem and rewriting all variables as ratios to the permanent component of labor income (P_{it}). The laws of motion and the value function can then be rewritten in terms of the normalized variables, and we use lowercase letters to denote them (for instance, $x_{it} \equiv \frac{X_{it}}{P_{it}}$). This allows us to reduce the number of state variables to three: age (t), normalized cash on hand (x_{it}) and participation status (whether the fixed cost has already been paid or not). In the last period, the policy functions are determined by the bequest motive and the value function corresponds to the bequest function. We can now use this value function to compute the policy rules for the previous period, and given these, obtain the corresponding value function. This procedure is then iterated backwards.

F. Computing Transition Distributions

After solving for the optimal policy functions, we can simulate the model to replicate the behavior of a large number of households and compute, for example, the corresponding average allocations. Here we propose an alternative method of computing various statistics that is based on the explicit calculation of the transition distribution of cash on hand from one age to the next. The computational details are given in Appendix B, but the intuitive idea is straightforward. Once we have solved for the policy functions, we can substitute those in the budget constraint to obtain the distribution of x_{t+1} as a function of x_t. Doing this for every possible x_t, we are effectively computing the full transition matrix.[12]

Once we have these distributions, the unconditional mean consumption for age t can then be computed as[13]

$$\bar{c}_t = \theta_t \left\{ \sum_{j=1}^{J} \pi_{t,j}^{\mathrm{I}} * c^{\mathrm{I}}(x_j, t) \right\} + (1-\theta_t) \left\{ \sum_{j=1}^{J} \pi_{t,j}^{\mathrm{O}} * c^{\mathrm{O}}(x_j, t) \right\}, \tag{11}$$

where J is the number of grid points used in the discretization of normalized cash on hand, and $\pi_{t,j}^{\mathrm{I}}$ and $\pi_{t,j}^{\mathrm{O}}$ are the probability masses associated with each grid point at time t, for stockholders and nonstockholders, respectively. The participation rate at age $t\,(\theta_t)$ is given by

$$\theta_t = \theta_{t-1} + (1-\theta_{t-1}) * \sum_{x_j > x^*} \pi_{t,j}^{\mathrm{O}}, \tag{12}$$

where x^* is the trigger point that causes participation, which is determined endogenously through the participation decision rule.

Finally, if we use α_t to denote the share of liquid wealth invested in the stock market at age t, then the unconditional portfolio allocation is computed as:

$$\bar{\alpha}_t = \frac{\theta_t * \left\{ \sum_{j=1}^{J} \pi_{t,j}^{\mathrm{I}} * \alpha(x_j, t) * \left(x_j - c^{\mathrm{I}}(x_j, t)\right) \right\}}{\theta_t * \sum_{j=1}^{J} \left[\pi_{t,j}^{\mathrm{I}} * \left(x_j - c^{\mathrm{I}}(x_j, t)\right)\right] + (1-\theta_t) * \sum_{j=1}^{J} \left[\pi_{t,j}^{\mathrm{O}} * \left(x_j - c^{\mathrm{O}}(x_j, t)\right)\right]}. \tag{13}$$

[12] The results in the paper were computed both from the transition distributions and using Monte Carlo simulations. The results were found to be identical, as long as the number of simulations is not too small (2,000 or more).

[13] Superscript I denotes households participating in the stock market, while superscript O denotes households out of the stock market.

G. Parameter Calibration

G.1. Preference Parameters

We start by presenting results for a relatively standard choice, (risk aversion) $\rho = 5$, (EIS) $\psi = 0.2$, and (discount factor) $\beta = 0.96$. However, later on we report results for several different values of both the coefficient of RRA (ρ) and the EIS (ψ), as these parameters have very important implications for our results. We use the mortality tables of the National Center for Health Statistics to parameterize the conditional survival probabilities.

The importance of the bequest motive (b) is set at 2.5. As we discuss below, this parameter choice is motivated by the desire to match the wealth accumulation profiles observed in the data, but we present some sensitivity analysis with respect to this parameter.

G.2. Labor Income Process

The deterministic labor income profile ($f(t, Z_{it})$ reflects the hump shape of earnings over the life cycle, and the corresponding parameter values, just like the retirement transfers (λ), are taken from Cocco, Gomes, and Maenhout (1999). With respect to standard deviations of the idiosyncratic shocks, the estimates range from 0.35 for σ_u and 0.12 for σ_n (Cocco, Gomes, and Maenhout) to 0.1 for σ_u and 0.08 for σ_n (Carroll (1992)). We use numbers similar to the ones in Gourinchas and Parker (2002): $\sigma_u = 0.15$ and $\sigma_n = 0.1$. It is common practice to estimate different labor income profiles for different education groups (college graduates, high school graduates, households without a high school degree). In our paper, we only report the results obtained with the parameters estimated from the subsample of high school graduates, as the results for the other two groups are very similar.

G.3. Asset Returns, Correlation and Fixed Cost

The constant net real interest rate ($R^f - 1$) is set at 2%, while for the stock return process we consider a mean equity premium (μ) equal to 4% and a standard deviation (σ_ε) of 18%. Considering an equity premium of 4% (as opposed to the historical 6%) is a fairly common choice in this literature (e.g., Yao and Zhang (forthcoming), Cocco (2001) or Campbell et al. (2001)). Even after having paid the fixed entry cost, the average retail investor still faces nontrivial transaction costs, mostly in the form of mutual fund fees. This adjustment is a shortcut representation for those costs, since the dimensionality of the problem prevents us from modeling them explicitly (as in Heaton and Lucas (1996), for example).

The evidence on the magnitude of the correlation between stock returns and permanent labor income shocks is mixed.[14] Davis and Willen (2001) and Heaton

[14] Moreover, it has been argued that these estimations suffer from a small sample bias, since the time-series dimension is too short in micro data, and estimations using macro data usually yield larger and more significant correlations (see, e.g., Jermann (1999)).

and Lucas (2000) do not distinguish between the two components of labor income (permanent and transitory) when computing the correlation coefficients. For the purposes of calibrating our model, we need to know the magnitude of the correlation coefficient for these two shocks separately. Campbell et al. (2001) estimate the correlation between the permanent component of labor income shocks and stock returns, and obtain a correlation coefficient of 0.15.[15] They do not estimate a correlation between transitory shocks and stock returns and just assume it to be equal to zero. We use these numbers ($\phi_N = 0.15$ and $\phi_U = 0.0$) for our benchmark calibration, and perform sensitivity analysis around these values.

With respect to the fixed cost of participation we consider two limit cases: one where the cost is zero, and one where it equals 0.025 (2.5% of the household's expected annual income). This parameter reflects both the monetary cost associated with the initial investment in the stock market, and the opportunity cost associated with obtaining the necessary information for making such investment.[16]

G.4. Housing Expenditures

We measure housing expenditures using data from the Panel Study of Income Dynamics from 1976 until 1993.[17] For each household, in each year, we compute the ratio of annual mortgage payments and rent payments (housing-related expenditures—H) relative to annual labor income (Y):

$$h_{it} \equiv \frac{H_{it}}{Y_{it}}. \tag{14}$$

We combine mortgage payments and rent together, since we are not modeling the housing decision explicitly. We identify the age effects by running the following regression on the full panel:

$$h_{it} = A + B_1 * age + B_2 * age^2 + B_3 * age^3 + time\ dummies + \zeta_{it}, \tag{15}$$

[15] It is important to realize that in their tables, Campbell et al. (2001) actually report the correlation of the *aggregate* component of permanent labor income shocks with stock returns. This explains their high estimates: 45.6%. To obtain the correlation with the "total permanent shock," we need to adjust for the standard deviation of the aggregate component relative to the total, which gives the 15% number.

[16] Consider the average household that has an annual labor income of $35,000. If the time cost were zero, then this value of F would imply a monetary cost of $875. If instead the monetary cost were zero, then this would imply a time cost of 9.1 days (6.3 working days). More generally, any convex combination of these two costs is acceptable, for example, a time cost of 1 (2) day(s) and a monetary cost of $779 ($683). Paiella (2001) and Vissing-Jørgensen (2002b) used Euler equation estimation methods to obtain implied participation costs from observed consumption choices. They find values in the $75–200 range, but these are per-period costs, so our number is quite reasonable when compared to their estimates.

[17] Before 1976 there is no information on mortgage expenditures, and 1993 is the last year available on final release from the PSID.

Table I

Regression of the Ratio of Housing Expenditures to Labor Income (he_{it}), on Age Polynomials, and Time Dummies

The data are taken from the Panel Study of Income Dynamics from 1976 until 1993. For each household, in each year, we compute the ratio of annual mortgage payments plus rent payments relative to annual labor income, and regress this ratio against a constant, a cubic polynomial of age (where age is defined as the age of the head of the household), and time dummies. We eliminate all observations with age greater than 75.

	Coefficient	T-Stat
Constant	0.703998	5.47
Age	−0.0352276	−3.70
Age^2	0.0007205	3.17
Age^3	−0.0000049	−2.84
Adj. R^2	0.025	

where *age* is defined as the age of the head of the household. We eliminate all observations with $age > 75$.[18] The estimation results are reported in Table I.

In the model we use

$$h_t = \max(A + B_1 * age + B_2 * age^2 + B_3 * age^3, 0), \tag{16}$$

which, given our parameter estimates, truncates h_t at zero for $age \geqslant 80$.

III. Results without the Fixed Participation Cost

A. Consumption and Wealth Accumulation

Figure 2A plots mean normalized consumption ($\bar{c}_t$), mean normalized wealth ($\bar{w}_t$), and mean normalized income net of housing expenditures ($(1 - h_t) * \bar{y}_t$). The preference parameters are $\rho = 5$ and $\psi = 0.2$, and the importance of the bequest motive (b) is set at 2.5. Early in life, the household is liquidity constrained and saves only a small buffer stock of wealth. From approximately ages 30–35 onwards, she starts saving for retirement and bequests, and wealth accumulation increases significantly. During retirement, consumption decreases as a result of the very high effective discount rate (high mortality risk). Wealth does not fall towards zero due to the presence of the bequest motive.[19]

Table II shows the mean consumption to wealth ratio for different values of the preference parameters. We report results for values of risk aversion between one and five and for values of the EIS between 0.2 and 0.8, since this is the

[18] There are several reasons for eliminating these households. First, there are very few observations within each age group beyond age 75. Second, for most of these households, the values of h_{it} are equal to zero. Third, this is consistent with the estimation procedure used for the labor income process.

[19] Net income increases during the first years of retirement because the housing expenditures (h_t) are still positive and decreasing toward zero.

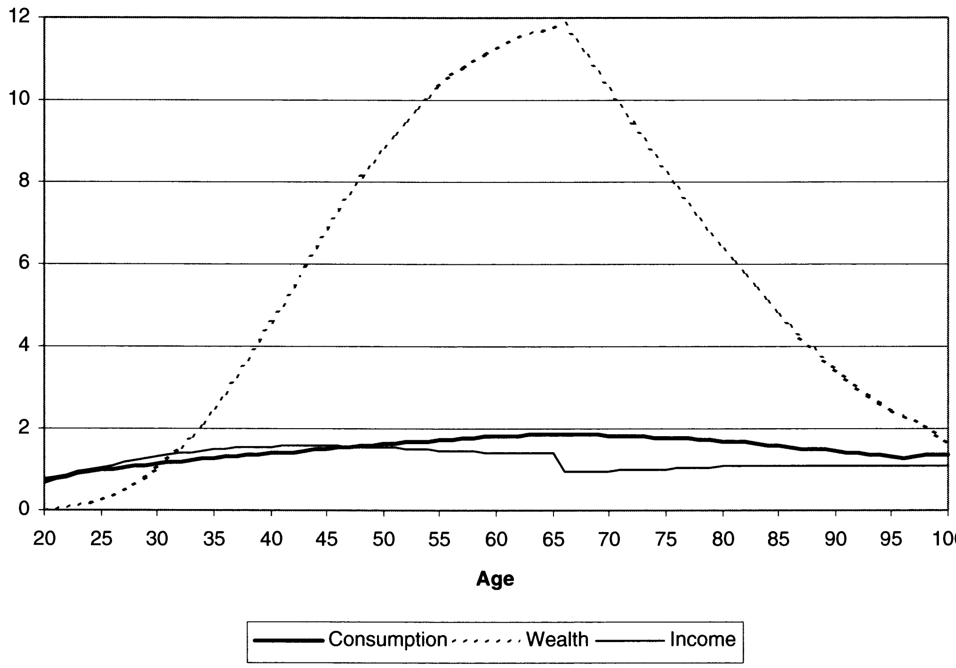

Figure 2A. Life-cycle profiles of consumption, income, and wealth for the baseline preference parameters: coefficient of RRA = 5, elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

Table II

Average Consumption–Wealth Ratio (*C/X*) Implied by the No-Fixed Cost Model for Different Values of Both the Coefficient of Risk Aversion (ρ) and the EIS (ψ), and for Different Age Groups

	Age 20 until Age 35			Age 36 until Age 65			Age 66 until Age 100		
ρ	$\psi = 0.8$	$\psi = 0.5$	$\psi = 0.2$	$\psi = 0.8$	$\psi = 0.5$	$\psi = 0.2$	$\psi = 0.8$	$\psi = 0.5$	$\psi = 0.2$
1	98%	99%	99%	98%	99%	99%	100%	100%	100%
1.2	87%	92%	93%	43%	88%	94%	88%	100%	100%
2	76%	86%	90%	18%	35%	67%	25%	71%	97%
4	61%	67%	75%	14%	18%	27%	23%	29%	59%
5	55%	60%	66%	13%	16%	19%	25%	26%	47%

range that we consider in the remaining part of the paper, and it is consistent with existing empirical evidence (see the discussion in Section IV.C.1). The top panel considers the first adult years (20–35) during which wealth accumulation is mostly driven by the precautionary savings motive. As a result, the optimal consumption to wealth ratio is significantly more affected by prudence than by the EIS. Since the more risk-averse investors are also the more prudent ones, the consumption to wealth ratio is a decreasing function of risk aversion. For very low values of risk aversion (close to 1), C/X converges to the 100% limit imposed by the borrowing constraint.

The second panel of Table II summarizes the remaining preretirement period (36–65), during which savings are now determined by the preference for low-frequency consumption smoothing, while the bottom panel reports the results for the retirement period (66–100). The results are qualitatively identical in both cases.[20] The optimal consumption to wealth ratio is driven by the trade-off between the (endogenous) expected return on invested wealth and the discount rate, combined with the household's sensitivity to these incentives (the EIS). The less risk-averse households invest a larger fraction of their portfolio in stocks, and therefore the expected return on their invested wealth is higher. Thus, since for a given $\psi < 1$ the income effect dominates the substitution effect, they have higher consumption to wealth ratios. The same intuition explains why, for a given ρ, C/X is a decreasing function of the EIS. However, for both the highest and lowest values of ρ that we consider, this pattern becomes weaker. In the first case, the expected return on invested wealth is very close to the discount rate and the consumption to wealth ratio is almost independent of ψ.[21] In the second case, C/X is close to the 100% limit given by the borrowing constraint.

From the results in Table II, we can conclude that for the range of values that we consider, the consumption to wealth ratio is a decreasing function of both ρ and ψ at every stage of the life cycle.

B. Asset Allocation

Figure 2B graphs the unconditional mean asset allocation in equities ($\bar{\alpha}_t$) for the same preference parameters as in Figure 2A ($\rho = 5$, $\psi = 0.2$, and $b = 2.5$). Even though earnings risk is uninsurable, cash is a closer substitute for future labor income than are stocks (see Heaton and Lucas (1997)). Young households are overinvested in their human capital and view this nontradeable asset as an implicit riskless asset in their portfolio. Given that the holdings of this relatively riskless asset are larger in the early part of the life cycle, all young households participate in the stock market and they allocate most of their financial wealth to stocks.[22] As retirement approaches, and financial wealth increases relative to the present value of future labor income, agents start investing in cash. When retirement savings is at its peak, more than 50% of total wealth is now being invested in the riskless asset.

During retirement, both future labor income (the present value of the pension transfers) and financial wealth are falling, so that the optimal asset allocation

[20] The results are qualitatively similar to the ones obtained by Campbell and Viceira (1999) in the context of an infinite-horizon portfolio choice model without labor income and liquidity constraints.

[21] As risk aversion increases further (not reported), the return on the portfolio falls below the discount rate and the consumption to wealth ratio becomes an increasing function of the EIS.

[22] During the very first years of adult life households hold a small fraction of their wealth in cash, since the present value of future labor income is actually still increasing.

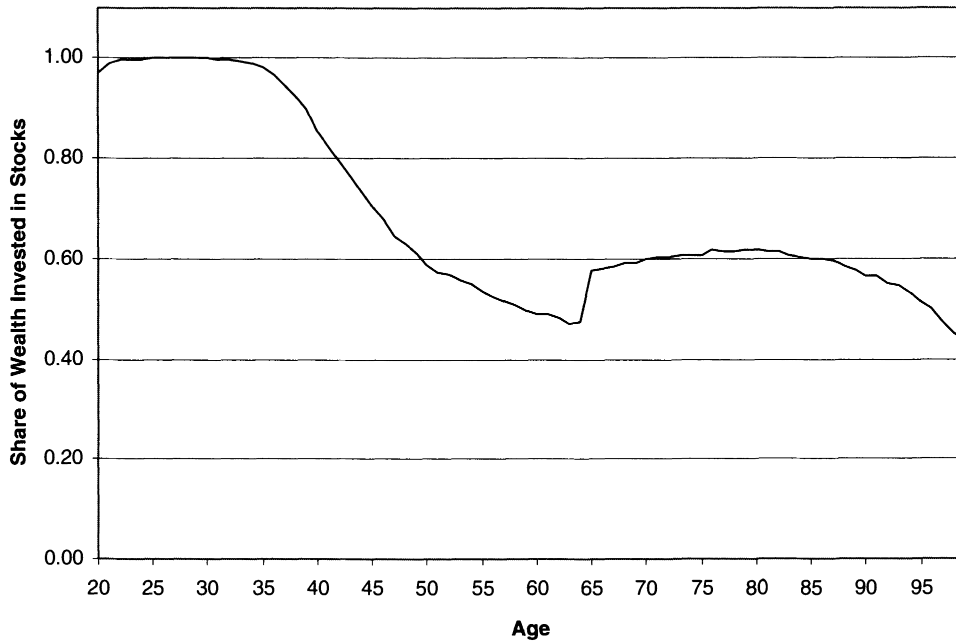

Figure 2B. Life-cycle asset allocation for the baseline preference parameters: coefficient of RRA $= 5$, elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

is determined by the relative speed at which these two decrease. Naturally, this depends both on the discount rate (adjusted for the survival probabilities) and the strength of the bequest motive. Given our parameter values, during most of the retirement period, future labor income and wealth decay at similar rates, and as a result, the share of wealth invested in stocks remains approximately constant.[23]

IV. Results with the Fixed Participation Cost

A. Baseline Case

We start by reporting the results for the baseline preference parameters ($\rho = 5$, $\psi = 0.2$, and $b = 2.5$). In the next sections, we consider different values.

A.1. Participation Decision and Asset Allocation

The participation decision is determined by four factors. First, it is an increasing function of wealth accumulation. Intuitively, households that accumulate more wealth over the life cycle have a stronger incentive to enter the stock market. Second, for the same level of wealth accumulation, participation is a

[23] This is true, except during the last years, when most households have very little financial wealth left.

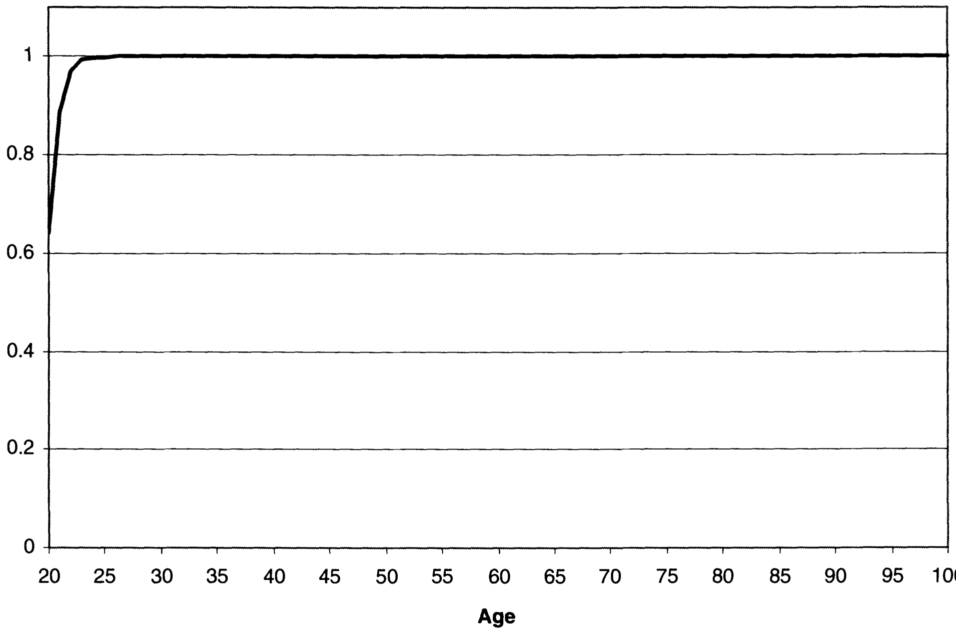

Figure 3A. Stock market participation over the life cycle for the baseline preference parameters: coefficient of RRA = 5, elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

positive function of the optimal share of wealth invested in stocks. Third, since F is one-time cost, participation is also a positive function of the investment horizon. Fourth, since the cost must be paid at the time of entry, the likelihood of participating in the stock market is a negative function of current marginal utility.

Figure 3A shows the stock market participation rate for the baseline preference parameters. Since young households are liquidity constrained, their marginal utility is extremely high, and as a result they do not participate in the stock market until sufficient wealth has been accumulated. As we can see from Figure 3A, given these preference parameters, this happens very fast and by age 25 the participation rate is almost 100%. As a result, the average life-cycle profiles of wealth accumulation, consumption, and equity shares are almost identical to the ones obtained without the fixed cost (reported in Figures 2A and B), and are therefore omitted here.

A.2. Wealth Distributions

Figure 3B plots the evolution of the distributions of cash on hand for the two types of agents: stock market participants and nonparticipants at age 30, still with $\rho = 5$ and $\psi = 0.2$. There is a pronounced spike around the normalized cash-on-hand level of 0.75; beyond that level, stock market participation

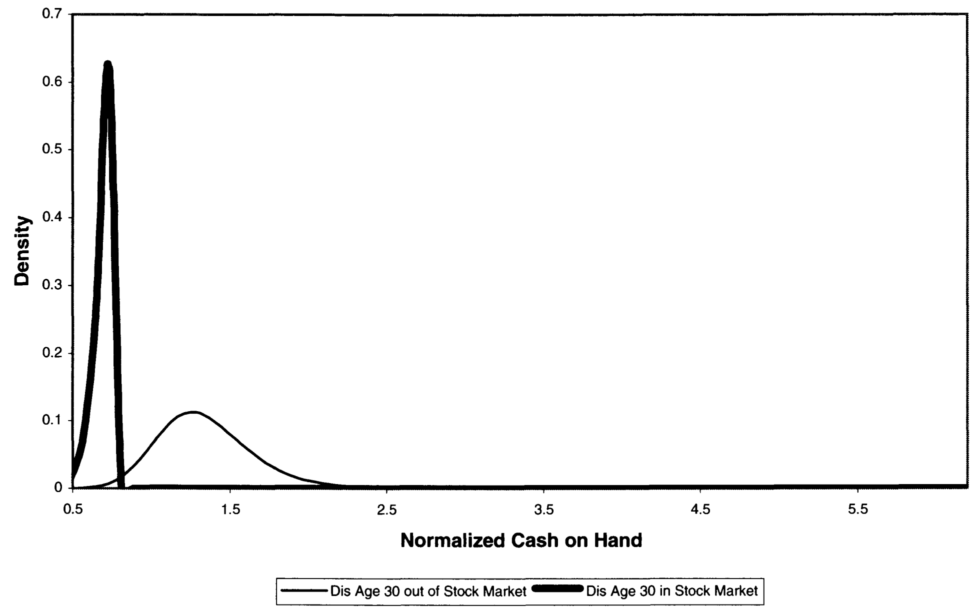

Figure 3B. Distributions for normalized cash on hand, at age 30, for the baseline preference parameters: coefficient of RRA = 5, elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

becomes optimal and the two distributions overlap for a small interval, mostly representing the incurrence of the fixed entry cost. Figure 3C plots the distributions of cash on hand for ages 50 for both types of agents. Conditional on age, the distribution of cash on hand for stockholders has a much higher variance than the wealth distribution for the households that have not participated in the stock market.

A.3. Sensitivity Analysis

Next, we perform some sensitivity analysis with respect to the importance of the bequest motive. Figure 3D plots wealth accumulation for different values of the parameter b, while Figure 3E plots the corresponding conditional asset allocations. A stronger bequest motive increases wealth accumulation at every stage of the life cycle, and the effect is strongest during retirement. The increase in wealth accumulation leads to a modest reduction in the share of equity investment during working life. Since both of these effects have a fairly modest impact until retirement, the implied participation rates are not significantly affected, and therefore we do not report them. During retirement, an increase in the bequest motive decreases the speed at which wealth is being drawn down, and leads to a higher ratio of financial wealth to labor income. As a result, for a given age, a stronger (weaker) bequest motive decreases (increases) the optimal equity share.

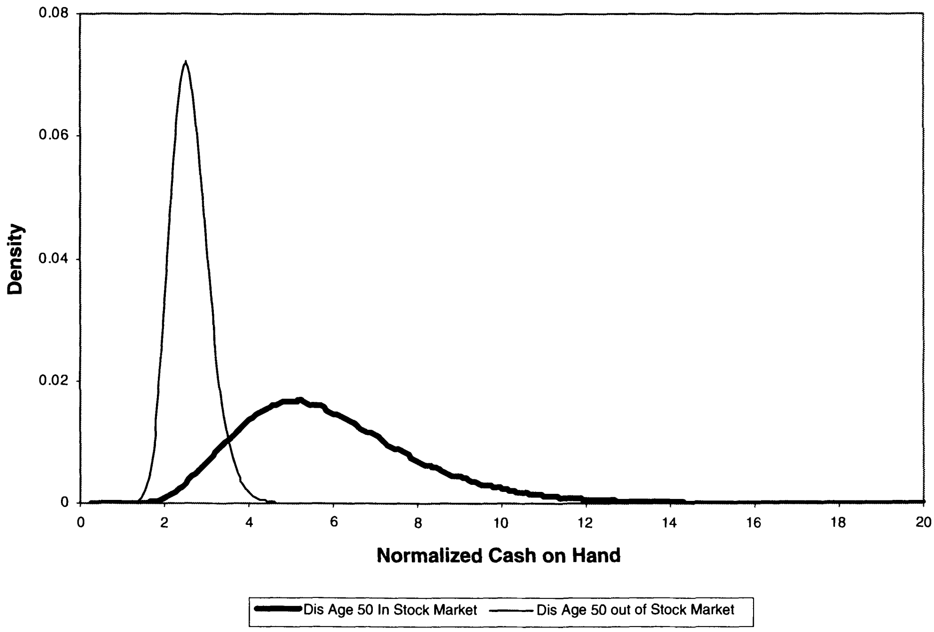

Figure 3C. Distributions for normalized cash on hand, at age 50, for the baseline preference parameters: coefficient of RRA = 5, elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

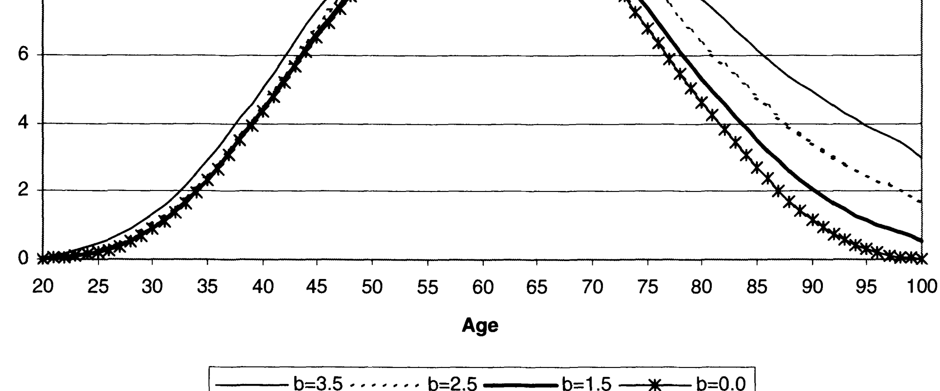

Figure 3D. Wealth accumulation for different values of the bequest parameter (*b*), and for the baseline preferences parameters: coefficient of RRA = 5 and elasticity of intertemporal substitution equal to 0.2.

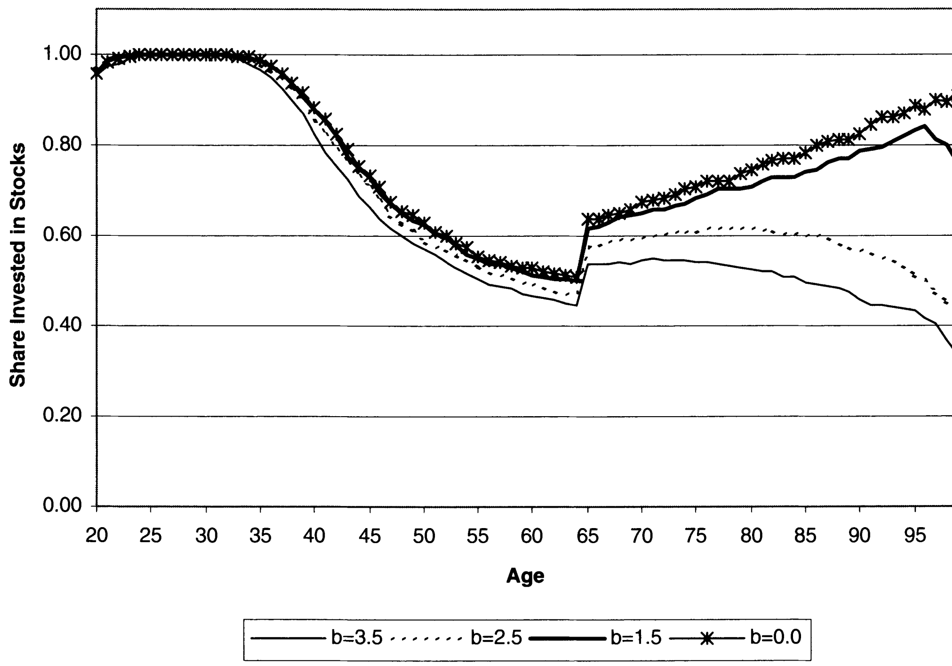

Figure 3E. Asset allocation for stock market participants, for different values of the bequest parameter (*b*), and for the baseline preference parameters: coefficient of RRA = 5, elasticity of intertemporal substitution equal to 0.2.

As mentioned before, the empirical evidence is mixed on the magnitude of the correlation coefficients between stock returns and the different labor income shocks (transitory and permanent). In our baseline calibration, we have assumed $\phi_N = 0.15$ and $\phi_U = 0.0$, following the estimation of Campbell et al. (2001). In Figure 3F, we now check whether the results are sensitive to these values. We only report the asset allocation decisions, since the participation decision is almost identical in all cases.

Campbell et al. (2001) do not actually estimate $\phi_U = 0.0$; they just assume it. Therefore, in our first experiment, we allow for a positive correlation between stock returns and the transitory labor income shocks, in particular we consider $\phi_N = 0.15$ and $\phi_U = 0.1$. The results are very similar to the ones obtained for the baseline case. In the second experiment, we now set $\phi_N = 0.0$ and assume that the correlation is instead driven exclusively by the transitory shocks, thus setting $\phi_U = 0.15$. We again obtain results that are extremely close to our baseline case. The share invested in equities is higher, but only marginally so. Finally, we consider the case in which there is no correlation between stock returns and labor income shocks. Only in this case do we find a visible difference relative to the benchmark calibration, as investors now allocate a higher fraction of their wealth to equities.

So far we have assumed that all households start at age 20 with zero initial wealth. Table III reports summary statistics for the wealth distribution for

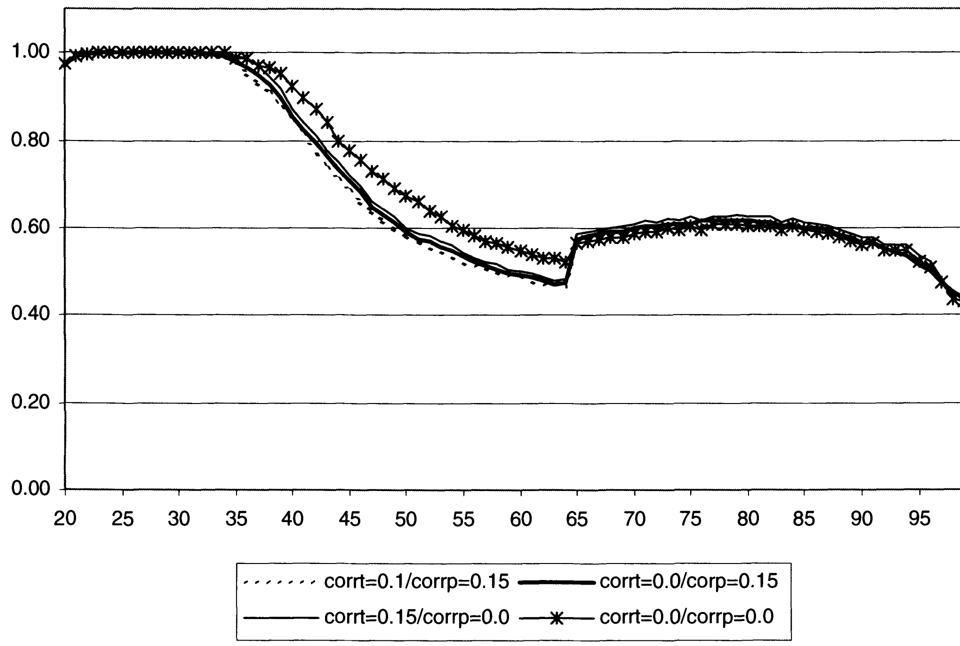

Figure 3F. Asset allocation for stock market participants, for different values of the correlation between stock returns and transitory (permanent) labor income shocks, denoted by corrt (corrp), with the baseline preference parameters: coefficient of RRA = 5, elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

Table III

The Wealth Distribution (Wealth-to-Income Ratios) for Households with Head Aged 20 or Less

The data are taken from the 2001 SCF (details in Appendix C). The variable X defines liquid wealth and Y denotes labor or pension income.

Decile (%)	10	20	30	40	50	60	70	80	90
X/Y	0.000	0.015	0.043	0.113	0.167	0.236	0.267	0.406	0.863

households of age 20 (or lower) from the SCF (details given in Appendix C). When we use this distribution as the initial condition in our model, with the exception of the first few years, the wealth profiles are virtually indistinguishable and therefore we do not report them.[24] This result occurs because young households are liquidity constrained and they therefore prefer to consume all of this additional wealth rather than save it.

In our baseline calibration, we assume that housing expenditures constitute a fixed proportion of labor income. We now allow for a stochastic component

[24] In the implementation, we truncate the distribution from the SCF at its 90 percentile.

in this ratio. More precisely, disposable income is now given by $(1 - \tilde{h}_t)Y_{i,t+1}$, where

$$\tilde{h}_t = h_t * \exp\left(\varepsilon_t^h\right), \tag{17}$$

and ε_t^h follows a normal distribution with zero mean and variance $\sigma_{\varepsilon^h}^2$.

In this experiment we set $\sigma_{\varepsilon^h} = 0.25$. This effectively corresponds to an increase in the level of background risk, and it is equivalent to an increase in the variance of the transitory labor income shocks. We find that the results are quite similar to the baseline case and therefore we do not report them. The increase in background risk reduces the willingness to invest in stocks, but since these are transitory shocks, the effect is not very large. For the same reason, the wealth accumulation and the participation rate are almost unaffected.

B. Changing Risk Aversion and the Impact of Background Risk

The stock market participation rate implied by the baseline parameters is counterfactually high. In this section, we explore the model's ability to produce more realistic results by considering different preference parameter values. As mentioned before, the participation decision is an increasing function of both wealth (X) and the optimal share of wealth invested in risky assets (α). Decreasing risk aversion increases the optimal share invested in stocks, but as shown in Section III.A, it also decreases wealth accumulation at every stage of the life cycle. Therefore, the impact on the participation decision resulting from changes in risk aversion depends on which of these two effects dominates.

B.1. Wealth Accumulation

We start by decreasing ρ from 5 to 2, while maintaining the power utility assumption, thus increasing the EIS (ψ) to 0.5. In Figure 4A, we plot the wealth accumulation for this case and for the baseline parameter values ($\rho = 5$, $\psi = 0.2$, and $b = 2.5$). As expected, wealth accumulation is significantly reduced at every stage of the life cycle. As previously shown in Section III.A (see Table II), the average consumption to wealth ratio is now 86% for the age group 20–35, and 35% for the age group 36–65, as opposed to 66% and 19% respectively.

However, from the results in Table II, we know that if we depart from power utility and decrease both risk aversion (ρ) and the EIS (ψ) simultaneously, this significantly reduces wealth accumulation. Consider then decreasing ρ to 2, but now keeping ψ at 0.2. The consumption to wealth ratio for the first age group is not significantly affected (90% instead of 86%) because, at this stage of the life cycle, savings are essentially driven by prudence (which remains constant). However, for the second age group, wealth accumulation is determined mostly by the EIS. As a result, the average consumption to wealth ratio is now almost doubled, increasing from 35% to 67%. As shown in Figure 4A, this leads to a very significant reduction in life-cycle wealth accumulation.

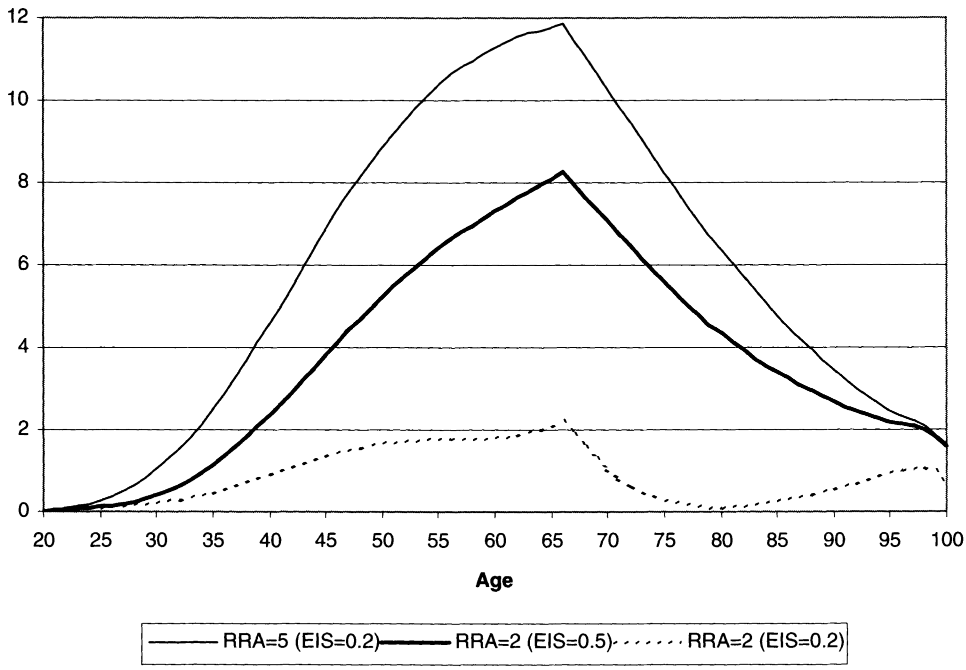

Figure 4A. Wealth accumulation over the life cycle, for different values of the preference parameters (coefficient of RRA and elasticity of intertemporal substitution) with bequest motive equal to 2.5.

B.2. Stock Market Participation Rates and Asset Allocation

Figure 4B plots the participation rates for different values of risk aversion, with the EIS equal to 0.2. Given the large differences in wealth accumulation, it is not surprising that the wealth effect dominates with respect to the participation decision. The less prudent households save less, and as a result, their participation rate is smaller. While almost all high prudent households have already paid the fixed cost by age 25, only 75% of the households with $\rho = 2$ have done so, although by age 35, even all of these investors have already paid the fixed cost as well. However, from the less risk-averse households (i.e., $\rho = 1.2$), only a very small fraction (<20%) will ever invest in stocks. On the other hand, as shown in Figure 4C, the reduction in risk aversion generates counterfactually high equity holdings for those investors that have paid the fixed cost.[25]

B.3. The Impact of Background Risk

The previous results illustrate one important trade-off generated by the level of background risk. When faced with more background risk (e.g., due to more

[25] For $\rho = 1.2$, the conditional equity share is always 100%, and therefore it is not included in the figure.

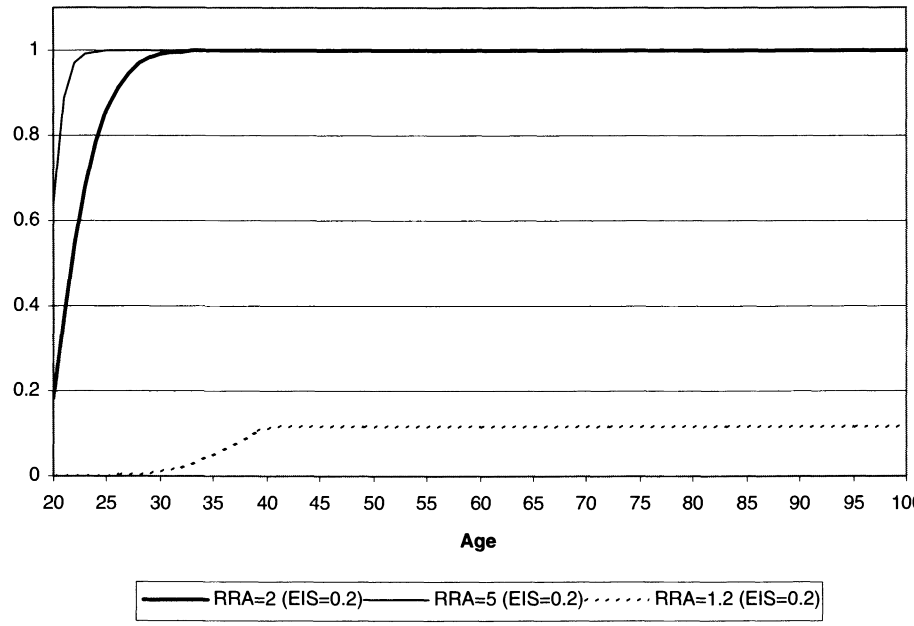

Figure 4B. Stock market participation rate for different values of the coefficient of RRA, with elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

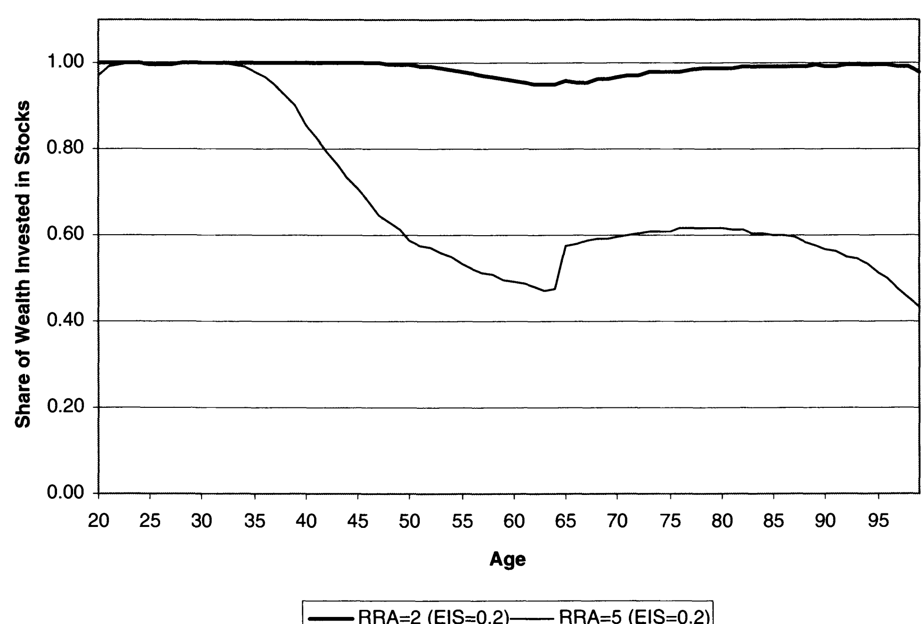

Figure 4C. Asset allocation for stock market participants, for different values of the coefficient of RRA, with elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

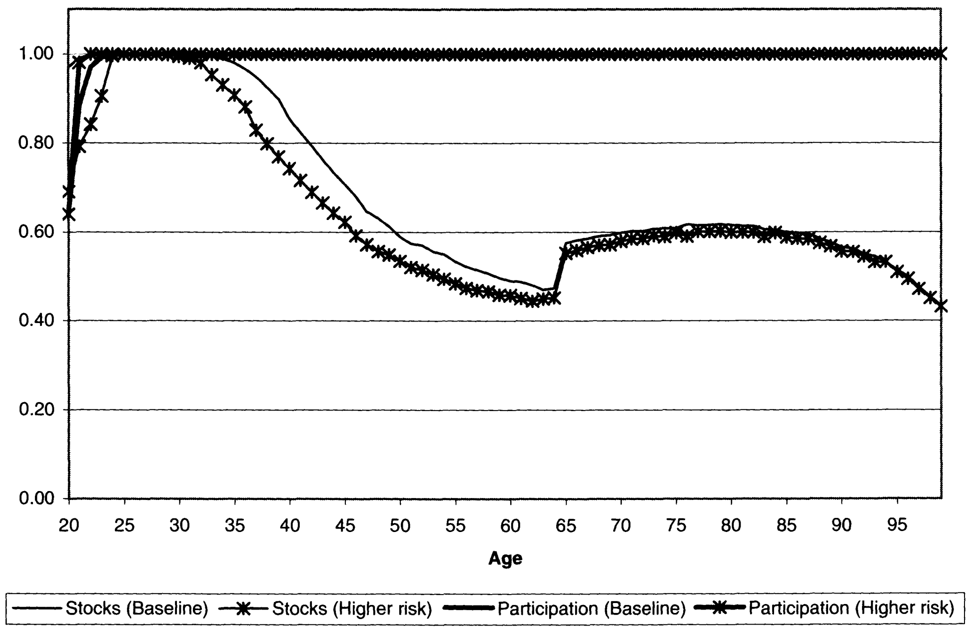

Figure 4D. Stock market participation rate and share of wealth invested in stocks for different levels of background risk, with the baseline preferences: coefficient of RRA = 5, elasticity of intertemporal substitution equal to 0.2, and bequest motive equal to 2.5.

labor income risk, consumption risk, or housing/mortgage risk) agents invest a smaller fraction of their financial wealth in risky assets. However, they also accumulate a larger buffer stock of wealth, and thus have a stronger incentive to enter the stock market. We have considered three different experiments in which we have increased the investor's background risk. In the first two, we have assumed a higher variance of respectively transitory and permanent labor income shocks, and in the third, we have included a positive probability of a disastrous labor income shock. Figure 4D shows the results for the case of the first experiment, where the variance of the transitory labor income shocks was increased by a factor of 3.[26] As expected, background risk crowds out stock holdings and households invest a smaller fraction of their portfolios in equities. However, they also increase their buffer stock of wealth, and as a result the stock market participation rate is higher than before.

C. Asset Allocation and Participation Rates with Preference Heterogeneity

C.1. Matching Participation Rates and Conditional Asset Allocations

Given our previous results, we can simultaneously match stock market participation rates and asset allocation conditional on participation with moderate

[26] The results for the other two cases are qualitatively identical, and they are available upon request.

degrees of risk aversion, if we allow for preference heterogeneity. Households with very low risk aversion and low EIS smooth idiosyncratic earnings shocks with a small buffer stock of assets, and most of them never invest in equities (thus behaving as in the Deaton (1991) infinite-horizon model). This seems to describe adequately the behavior of a large fraction of the U.S. population that retires without significant financial assets (and does not participate in the stock market). Within the low EIS and low risk-aversion group, only a small fraction of them owns stocks, and they do so only as they get close to retirement. On the other hand, investors with high prudence and high EIS are the ones who participate in the stock market from early on, since they accumulate more wealth and therefore have a stronger incentive to pay the fixed cost. Therefore, the marginal stockholders are (endogenously) more risk averse, and as a result they do not invest their portfolios fully in stocks.

In this final section, we try to evaluate how much heterogeneity we need to match the data. In other words, can the model consistently explain the two facts for a plausible distribution of preference parameters across the population? Table IV reports participation rates and average equity shares for stock market participants, for different distributions, and compares them with the empirical evidence from the SCF. We first consider a 50% split between investors with both low risk aversion and low EIS ($\rho = 1.2$ and $\psi = 0.2$), and investors with moderate risk aversion and moderate EIS ($\rho = 5$ and $\psi = 0.5$). The model gives a participation rate of 52.1% and an equity share of 54.5% for stock market participants (Case 1 in Table IV), which matches fairly well with the empirical evidence reported in Section I (and summarized in the first row of Table IV).

It is important to mention that this form of heterogeneity is consistent with the existing empirical evidence. Attanasio, Banks, and Tanner (2002) show that the CRRA coefficient is much higher (thus much lower EIS) for nonstockholders

Table IV

The Average Stock Market Participation Rate ($\bar{P}$) and Average Stock Holdings for Stock Market Participants ($\bar{\alpha}_P$)

The first row reports data from the 2001 SCF (details in Appendix C), while the other four panels report the results from the model for different distributions of investors. Case 1 assumes two groups of agents, ($\rho = 1.2$ and $\psi = 0.2$) and ($\rho = 5$ and $\psi = 0.5$), with 50% weight each. Case 2 also assumes two groups of agents, but now ($\rho = 1.1$ and $\psi = 0.2$) and ($\rho = 5$ and $\psi = 0.5$), with 50% weight each, and with the initial wealth distribution calibrated from the SCF. In Case 3 we have again two groups, ($\rho = 1.07$ and $\psi = 0.5$) and ($\rho = 5$ and $\psi = 0.5$), with 50% weight each. Finally, Case 4 considers three groups of agents, ($\rho = 1$ and $\psi = 0.2$) and ($\rho = 3$ and $\psi = 0.5$) and ($\rho = 5$ and $\psi = 0.5$), with weights 40%, 30%, and 30%, respectively.

	$\bar{P}$ (%)	$\bar{\alpha}_P$ (%)
Data	51.94	54.76
Model (Case 1)	52.14	54.48
Model (Case 2)	50.36	53.32
Model (Case 3)	54.42	56.24
Model (Case 4)	56.98	56.56

than for stockholders. Vissing-Jørgensen (2002a) focuses on this distinction and argues that "accounting for limited asset market participation is crucial for obtaining consistent estimates of the EIS" (p. 827). Vissing-Jørgensen then obtains estimates of the EIS greater than 0.3 for risky asset holders, while for the remaining households the EIS estimates are small and insignificantly different from zero. Vissing-Jørgensen and Attanasio (2003) further stress that loosening the link between risk aversion and intertemporal substitution can generate implications about the covariance of stock returns and individual consumption growth for stockholders that are not rejected in the data. They offer risk aversion estimates for stockholders at around 5–10 and EIS estimates around 1. Overall, the existing estimates of EIS and risk aversion are consistent with the values that we use in this paper.

C.2. Life-Cycle Profiles

We now report the life-cycle profiles of stock market participation and asset allocation implied by the model. As argued in Section I, in the data these profiles are not very robust to specific assumptions about cohort or time effects (see Figures 1A and B). As a result, in this paper, we have focused mostly on life-cycle averages.

Figure 5A plots the stock market participation rate implied by the model for different age groups, together with the corresponding numbers from the SCF (see Appendix C for details), while Figure 5B does the same, but now for

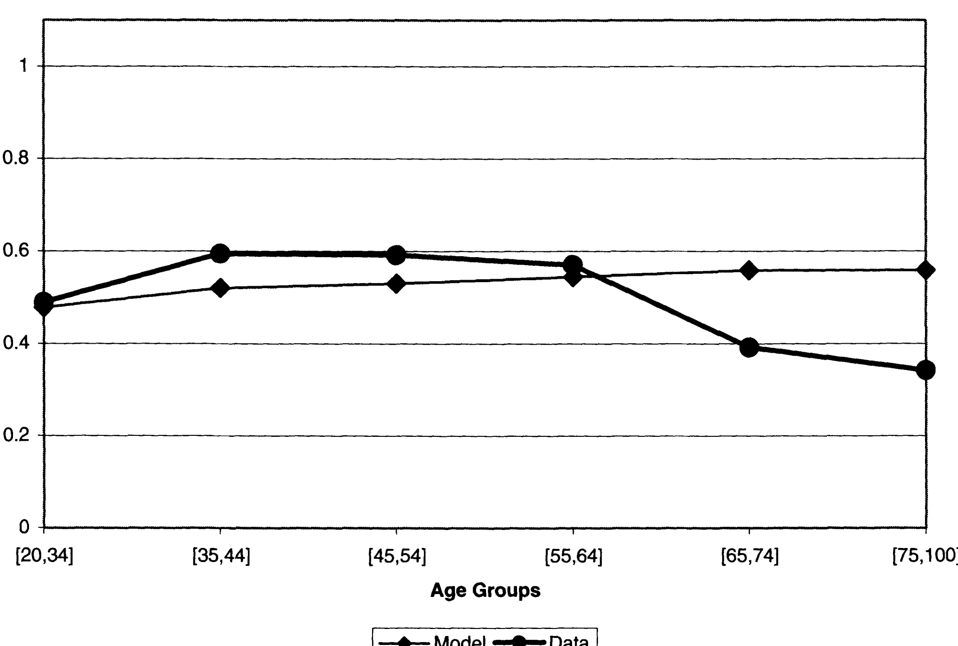

Figure 5A. Stock market participation rate implied by the heterogeneous agent model and stock market participation rate from the 2001 sample of the SCF.

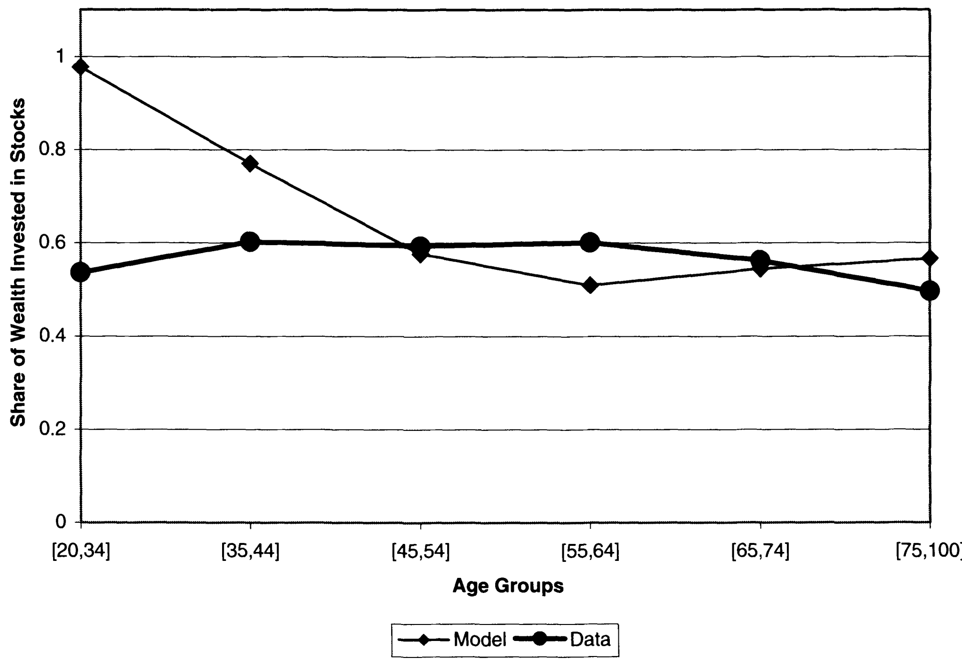

Figure 5B. Asset allocation for stock market participants rate implied by the heterogeneous agent model and asset allocation for stock market participants from the 2001 sample of the SCF.

the average asset allocation of stock market participants. The participation rates are extremely similar, with the largest difference occurring at retirement when the participation rate in the data declines, while it remains constant in the model. However, as shown in Figure 1B, this is exactly one of the results that is not robust to the assumption of cohort dummies versus time dummies. With respect to the asset allocation decisions, we do observe a more significant difference, in this case for young households. In the model, these agents invest a significant fraction of their portfolio in equities, while in the data, regardless of the controls, this does not happen.

C.3. Sensitivity Analysis and Robustness

So far we have assumed that households start at age 20 with zero initial wealth, since we have seen in Section IV.A.3 that, for $\rho = 5$, using the initial wealth distribution calibrated from the SCF does not affect the results. However, this might not be the case for investors with low risk aversion and low EIS, since they save very little. By giving these households some positive initial wealth, we are likely to see an increase in stock market participation rates. In the third row of Table IV (Case 2) we show that this effect is not too large, and we can replicate the previous results by considering a slightly lower value of risk aversion (1.1).[27]

[27] Alternatively, we could consider a lower value of the EIS. With $\psi = 0.1$ we would again obtain very similar results.

It is important to point out that we do not need to assume a very low value of the EIS to generate large nonparticipation, since we can compensate for a higher ψ by decreasing risk aversion even further. This is shown in the fourth row of Table IV (Case 3), where we fix the EIS coefficient equal to 0.5 for both types of investors. To reproduce the results in the first panel, we find that we need to decrease ρ to 1.07 for the less risk-averse group.

Given our previous discussion, we know that households with risk aversion between 1.5 and 4 tend to participate in the stock market from early on, and invest almost all of their wealth in stocks. Naturally, it is not reasonable to assume that the distribution of coefficients of risk aversion mysteriously collapses to the two extremes that we have previously considered (1.2 and 5). In the fifth row of Table IV (we now consider a smoother distribution, with ρ ranging from 1 to 5 (Case 4)). It is important to point out that this is not a uniform distribution, as there is a slightly higher fraction of less risk-averse households. If we want to match both facts simultaneously, with a (relatively) smooth distribution, we need it to exhibit some negative skewness. As predicted, both the participation rate and the equity share are now higher than before, but not significantly so. The equity share is now 57%, while the participation rate is also equal to 57%, numbers that are still extremely close to the empirical evidence (row 1).

D. Wealth Distribution

In this section, we compare the wealth accumulation predicted by the model with the empirical evidence in the SCF. Given the (exogenous) differences in the preference parameters and the (endogenous) differences in the participation decision, our model generates a large degree of heterogeneity in wealth accumulation. To illustrate, we compare both median wealth accumulation and the extremes of the distribution (percentiles 10 and 90) to see if the model generates the degree of heterogeneity observed in the SCF. We divide households into the three usual age groups: buffer stock savers (20–35), retirement savers (36–65), and retirees (over 66).

The results are shown in Table V. The model can replicate the low wealth accumulation patterns of the poorer households in the data. Households with the lowest income realizations tend not to participate in the stock market and accumulate very little wealth over the life cycle. This is consistent with the results in Hubbard, Skinner, and Zeldes (1995), who illustrate in a similar model how the presence of social insurance (pensions) can crowd out private saving over the life cycle for the poorest quintile of the wealth distribution. Nevertheless, in the SCF these households still accumulate some nonnegligible wealth during retirement, something that does not happen in the model. For the median household, the model does quite well early in life; it overshoots for the second age group; and it undershoots at retirement. Finally, at the high end of the distribution we can generate extremely large wealth accumulation, although not quite as high as in the data. This difference is most significant during the retirement period and early in life. Overall the degree of heterogeneity in the wealth distribution is comparable to the one observed in the data. The model

Table V

The Distribution of Wealth-to-Labor Income Ratios from the 2001 SCF for Different Age Groups

Appendix C provides more details. Results are provided for two different versions of the model: with zero initial wealth and with the initial wealth distribution calibrated from the SCF. Results are shown for different age groups.

	Age Groups		
	20–35	36–65	$\geqslant 65$
	Data		
10[th] percentile	0.002	0.071	0.371
Median	0.287	2.170	7.931
90[th] percentile	2.702	10.648	33.363
	Model (zero initial wealth)		
10[th] percentile	0.006	0.005	0.006
Median	0.261	3.115	4.838
90[th] percentile	0.748	8.184	17.539
	Model (initial wealth from the SCF)		
10[th] percentile	0.006	0.005	0.006
Median	0.263	3.116	4.839
90[th] percentile	0.886	8.371	17.865

consistently generates low wealth accumulation at retirement, which would suggest the presence of a stronger bequest motive, but as shown in Section IV.A.3, a stronger bequest motive also increases wealth accumulation at mid life (Figure 3D).

In the last panel of Table V, we simulate the model with the initial distribution of wealth calibrated from the SCF. The results are almost identical to the previous ones, with a minor increase in wealth accumulation at the 90 percentile.

V. Conclusion

In this paper, we present a life-cycle asset allocation model with realistically calibrated uninsurable labor income risk that provides an explanation for two very important empirical observations: low stock market participation rates in the population as a whole, and moderate equity holdings for stock market participants. We do not rely on high values of risk aversion, or on extreme assumptions about background risk.

In our model, households with very low risk aversion and low EIS accumulate very little wealth and as a result, most of them never invest in equities. On the other hand, the more prudent investors are the ones who participate in the stock market from early on, as they accumulate more wealth and therefore have a stronger incentive to pay the fixed entry cost. Therefore, the marginal stockholders are (endogenously) more risk averse and as a result they do not invest their portfolios fully in stocks. On the negative side, the model still counterfactually predicts that young households that have already paid the participation cost invest most of their portfolio in equities.

Appendix A: Numerical Solution Method

We first simplify the solution by exploiting the scale-independence of the maximization problem and rewriting all variables as ratios to the permanent component of labor income (P_{it}). The laws of motion and the value function can then be rewritten in terms of these normalized variables, and we use lowercase letters to denote them (for instance, $x_{it} \equiv \frac{X_{it}}{P_{it}}$). This allows us to reduce the number of state variables to three: one continuous state variable (cash on hand, X_{it}) and two discrete state variables (age, t, and participation status, whether the fixed cost has been paid or not). We discretize the state space along the cash on hand dimension (the only continuous state variable), so that the relevant policy functions can now be represented on a numerical grid.

We solve the model using backward induction. For every age t prior to T, and for each point in the state space, we optimize using grid search. So we need to compute the value associated with each level of consumption, the decision to pay the fixed cost, and the share of liquid wealth invested in stocks. From the Bellman equation, these values are given as current utility plus the discounted expected continuation value ($E_t V_{t+1}(\cdot, \cdot)$), which we can compute once we have obtained V_{t+1}. In the last period, the policy functions are determined by the bequest motive and the value function corresponds to the bequest function, regardless of whether the fixed cost has been paid before or not. This gives us the terminal condition for our backward induction procedure. We perform all numerical integrations using Gaussian quadrature to approximate the distributions of the innovations to the labor income process and the risky asset returns. For points that do not lie on state space grid, we evaluate the value function using a cubic spline interpolation.

Once we have computed the value of all the alternatives, we just pick the maximum, thus obtaining the policy rules for the current period (S_t and B_t). At each point of the state space, the participation decision is computed by comparing the value function conditional on having paid the fixed cost (adjusting for the payment of the cost itself) with the value function conditional on nonpayment. Substituting these decision rules in the Bellman equation, we obtain this period's value function ($V_t(\cdot, \cdot)$), which is then used to solve the previous period's maximization problem. This process is iterated until $t = 1$.

Appendix B: Computing the Transition Distributions

To find the distribution of cash on hand, we first compute the relevant optimal policy rules: bond and stock policy functions for stock market participants and nonparticipants and the $\{0, 1\}$ participation rule as a function of cash on hand. Let $b^{\mathrm{I}}(x)$ and $s^{\mathrm{I}}(x)$ denote respectively the bonds and stock policy rules for individuals participating in the stock market, and let $b^{\mathrm{O}}(x)$ be the savings decision for the individual out of the stock market. We assume that households start their working lives with zero liquid assets. During working lives, for the households that have not paid the fixed cost, the evolution of normalized cash on hand is given by[28]

$$x_{t+1} = \left[b^{\mathrm{O}}(x_t)R_f\right]\left\{\frac{P_t}{P_{t+1}}\frac{\exp(f(t, Z_t))}{\exp(f(t+1, Z_{t+1}))}\right\} + (1-h_{t+1})U_{t+1}$$
$$= w\left(x_t \,\middle|\, \frac{P_t}{P_{t+1}}, \frac{\exp(f(t, Z_t))}{\exp(f(t+1, Z_{t+1}))}\right) + (1-h_{t+1})U_{t+1}, \tag{B1}$$

where $w(x)$ is defined by the last equality and is conditional on $\{\frac{P_t}{P_{t+1}}\}$ and the deterministically evolving $\frac{\exp(f(t, Z_t))}{\exp(f(t+1, Z_{t+1}))}$. Denote the transition matrix of moving from x_j to x_k,[29] conditional on not having paid the fixed cost, as T^{O}_{kj}. Let Δ denote the distance between the equally spaced discrete points of cash on hand. The random permanent shock $\frac{P_t}{P_{t+1}}$ is discretized using Gaussian quadrature with H points: $\frac{P_t}{P_{t+1}} = \{N_m\}_{m=1}^{m=H}$. $T^{\mathrm{O}}_{kj} = \Pr(x_{t+1=k} \mid x_{t=j})$ is found using[30]

$$\sum_{m=1}^{m=H} \Pr\left(x_{t+1} \mid x_t, \frac{P_t}{P_{t+1}} = N_m\right) * \Pr\left(\frac{P_t}{P_{t+1}} = N_m\right). \tag{B2}$$

Numerically, this probability is calculated using

$$T^{\mathrm{O}}_{kjm} = \Pr\left(x_k + \frac{\Delta}{2} \geqslant x_{t+1} \geqslant x_k - \frac{\Delta}{2} \,\middle|\, x_t = x_j, \frac{P_t}{P_{t+1}} = N_m\right). \tag{B3}$$

Making use of the approximation that for small values of σ_u^2, $U \sim N(\exp(\mu_u + 0.5 * \sigma_u^2), (\exp(2 * \mu_u + (\sigma_u^2)) * (\exp(\sigma_u^2) - 1))$, and denoting the mean of $(1-h_t)U$ by $\bar{U}$ and its standard deviation by σ, the transition probability conditional on N_m equals

$$T^{\mathrm{O}}_{kjm} = \Phi\left(\frac{x_k + \frac{\Delta}{2} - w(x_t \mid N_m) - \bar{U}}{\sigma}\right) - \Phi\left(\frac{x_k - \frac{\Delta}{2} - w(x_t \mid N_m) - \bar{U}}{\sigma}\right), \tag{B4}$$

where Φ is the cumulative distribution function for the standard normal. The unconditional probability from x_j to x_k is then given by

$$T^{\mathrm{O}}_{kj} = \sum_{m=1}^{m=H} T^{\mathrm{O}}_{kjm} \Pr(N_m). \tag{B5}$$

[28] To avoid cumbersome notation, the subscript i that denotes a particular individual is omitted in what follows.

[29] The normalized grid is discretized between $(x_{\min}, x_{\max})$ where $x_{\min}$ denotes the minimum point on the equally spaced grid and $x_{\max}$ denotes the maximum point.

[30] The dependence on the deterministically evolving $\frac{\exp(f(t,Z_t))}{\exp(f(t+1,Z_{t+1}))}$ is implied and is omitted from what follows for expositional clarity.

Given the transition matrix $\mathbf{T}^{\mathrm{O}}$ (letting the number of cash-on-hand grid points equal to J, this is a J by J matrix; T_{kj}^{O} represents the $\{k^{\mathrm{th}}, j^{\mathrm{th}}\}$ element), the next period probabilities of each of the cash-on-hand states can be found using

$$\pi_{kt}^{\mathrm{O}} = \sum_{j} T_{kj}^{\mathrm{O}} * \pi_{jt-1}^{\mathrm{O}}. \tag{B6}$$

We next use the vector Π_t^{O} (this is a J by 1 vector representing the mass of the population out of the stock market at each grid point; π_{kt}^{O} represents the $\{k^{\mathrm{th}}\}$ element at time t) and the participation policy rule to determine the percentage of households that optimally choose to incur the fixed cost and invest in risky assets. This is found by computing the sum of the probabilities in Π_t^{O} for which $x > x^*, x^*$ being the trigger point that causes participation (x^* is determined endogenously through the participation decision rule). These probabilities are then deleted from Π_t^{O} and are added to Π_t^{I}, appropriately renormalizing both $\{\Pi_t^{\mathrm{O}}, \Pi_t^{\mathrm{I}}\}$ to sum to one. The participation rate (θ_t) can be computed at this stage as

$$\theta_t = \theta_{t-1} + (1 - \theta_{t-1}) * \sum_{x_j > x^*} \pi_{t,j}^{\mathrm{O}}. \tag{B7}$$

The same methodology (but with more algebra and computations) can then be used to derive the transition distribution for cash on hand conditional on having paid the fixed cost, $\mathbf{T}_t^{\mathrm{I}}$. The corresponding normalized cash-on-hand evolution equation is

$$\begin{aligned} x_{t+1} &= \left[b(x_t)R^f + s(x_t)R_{t+1}^S\right]\left\{\frac{P_t}{P_{t+1}}\frac{\exp(f(t, Z_t))}{\exp(f(t+1, Z_{t+1}))}\right\} + (1 - h_{t+1})U_{t+1} \\ &= w\left(x_t \mid R_{t+1}^S, \frac{P_t}{P_{t+1}}\right) + (1 - h_{t+1})U_{t+1}, \end{aligned} \tag{B8}$$

where $w(x)$ is now conditional on $\{R_{t+1}^S, \frac{P_t}{P_{t+1}}\}$.[31] The random processes R_{t+1}^S and $\frac{P_t}{P_{t+1}}$ are discretized using Gaussian quadrature with H points: $R_{t+1}^S = \{R_l^S\}_{l=1}^{l=H}$ and $\frac{P_t}{P_{t+1}} = \{N_n\}_{n=1}^{n=H}$. $T_{kj}^{\mathrm{I}} = \Pr(x_{t+1=k} \mid x_{t=j})$ is obtained from

$$\sum_{l=1}^{l=H}\sum_{n=1}^{n=H} \Pr\left(x_{t+1} \mid x_t, R_{t+1}^S = R_l^S, \frac{P_t}{P_{t+1}} = N_n\right) * \Pr\left(R_l^S\right) * \Pr(N_n), \tag{B9}$$

where $\Pr(R_l^S)$ and $\Pr(N_n)$ stand respectively for $\Pr(R_{t+1}^S = R_l^S)$ and $\Pr(\frac{P_t}{P_{t+1}} = N_n)$, and where the independence between $\frac{P_t}{P_{t+1}}$ and R_{t+1}^S was used.[32] Numerically, this probability is calculated using

[31] The dependence on the nonrandom earnings component is omitted to simplify notation.

[32] The methodology can be applied for an arbitrary correlation structure using the Choleski decomposition of the variance-covariance matrix of the innovations.

$$T^{\mathrm{I}}_{kjln} = \Pr\left(x_k + \frac{\Delta}{2} \geqslant x_{t+1} \geqslant x_k - \frac{\Delta}{2} \,\middle|\, x_t = x_j, \frac{P_{it}}{P_{it+1}} = N_n, R^S_{t+1} = R^S_l\right). \quad \text{(B10)}$$

The transition probability conditional on N_n, R^S_l, and R^B_m equals

$$T^{\mathrm{I}}_{kjln} = \Phi\left(\frac{x_k + \frac{\Delta}{2} - w\big(x_t \mid N_n, R^S_l\big) - \bar{U}}{\sigma}\right) - \Phi\left(\frac{x_k - \frac{\Delta}{2} - w\big(x_t \mid N_n, R^S_l\big) - \bar{U}}{\sigma}\right). \quad \text{(B11)}$$

The unconditional probability from x_j to x_k is then given by

$$T^{\mathrm{I}}_{kj} = \sum_{l=1}^{l=H} \sum_{n=1}^{n=H} T^{\mathrm{I}}_{kjln} \Pr\big(R^S_l\big) \Pr(N_n). \quad \text{(B12)}$$

Given the matrix $\mathbf{T}^{\mathrm{I}}$, the probabilities of each of the states are updated by

$$\pi^{\mathrm{I}}_{kt+1} = \sum_j T^{\mathrm{I}}_{kj} * \pi^{\mathrm{I}}_{jt}. \quad \text{(B13)}$$

Appendix C: SCF Data

The SCF is probably the most comprehensive source of data on U.S. household assets. The SCF uses a two-part sampling strategy to obtain a sufficiently large and unbiased sample of wealthier households (the rich sample is chosen randomly using tax reports). To enhance the reliability of the data, the SCF makes weighting adjustments for survey nonrespondents; these weights were used in computing the values reported in the tables. The specific names in the codebook for the variables used are given below.

We construct a measure of labor income that matches as closely as possible the process for Y_{it} (earnings) in the text. We therefore define labor income as the sum of wages and salaries (*X5702*), unemployment or worker's compensation (*X5716*) and Social Security or other pensions, annuities, or other disability or retirement programs (*X5722*). Liquid wealth is variable *FIN* in the publicly available SCF data set, to which home equity was added. Variable *FIN* is made up of *LIQ* (all types of transaction accounts—checking, saving, money market, and call accounts), *CDS* (certificates of deposit), total directly held mutual funds, stocks, bonds, total quasiliquid financial assets (the sum of IRAs, thrift accounts, and future pensions), savings bonds, the cash value of whole life insurance, other managed assets (trusts, annuities, and managed investment accounts in which the household has equity interest), and other financial assets: This includes loans from the household to someone else, future proceeds, royalties, futures, nonpublic stock, and deferred compensation. We define home equity as the value of the home less the amount still owed on the first and second/third mortgages and the amount owed on home equity lines of credit. This definition of wealth is consistent with both the definitions in Hubbard, Skinner, and Zeldes (1995) and in Heaton and Lucas (2000).

Financial assets invested in the risky asset can be directly held stock, stock mutual funds, or amounts of stock in retirement accounts. We follow the procedures the SCF uses to construct this number for each household (variable *EQUITY*). Specifically, this is done by computing the full value of stocks, adding the full value if an asset is described as a stock mutual fund, and half the value if the asset refers to a combination of mutual funds. For this purpose, IRAs/Keoghs invested in stock are computed by adding the full value if mostly invested in stock, half the value if split between stocks/bonds or stocks/money market, and one-third of the value if split between stocks/bonds/money market. We also add other managed assets with equity interest (annuities, trusts, and MIAs) by adding the full value if mostly invested in stock, half the value if split between stocks/MFs & bonds/CDs, or "mixed/diversified," and one-third of the value if "other." We also add thrift-type retirement accounts invested in stock: the full value if mostly invested in stock and half the value if split between stocks and interest earning assets. Stock market participation is then determined by checking whether the full value of stocks (*EQUITY*) is greater than zero (variable *HEQUITY*).

We construct the share of wealth in stocks conditional on *HEQUITY* being positive as (*EQUITY*)/(*FIN*) where all the variables have been defined above.

REFERENCES

Aiyagari, Rao, and Mark Gertler, 1991, Asset returns with transactions costs and uninsured individual risk, *Journal of Monetary Economics* 27, 311–331.

Ameriks, John, and Stephen Zeldes, 2001, How do household portfolio shares vary with age? Working paper, Columbia Business School.

Attanasio, Orazio, James Banks, and Sarah Tanner, 2002, Asset holding and consumption volatility, *Journal of Political Economy* 110, 771–792.

Basak, Suleyman, and Domenico Cuoco, 1998, An equilibrium model with restricted stock market participation, *Review of Financial Studies* 11, 309–341.

Bertaut, Carol, and Michael Haliassos, 1997, Precautionary portfolio behavior from a life cycle perspective, *Journal of Economic Dynamics and Control* 21, 1511–1542.

Cagetti, Marco, 2003, Wealth accumulation over the life cycle and precautionary savings, *Journal of Business and Economic Statistics* 21, 339–353.

Campbell, John, Joao Cocco, Francisco Gomes, and Pascal Maenhout, 2001, Investing retirement wealth: A life cycle model, in John Campbell and Martin Feldstein, eds:. *Risk Aspects of Social Security Reform* (University of Chicago Press, Chicago).

Campbell, John, and Luis Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics* 114, 433–495.

Carroll, Christopher, 1992, The buffer-stock theory of saving: some macroeconomic evidence, *Brookings Papers on Economic Activity* 2, 61–156.

Carroll, Christopher, 1997, Buffer stock saving and the life cycle/permanent income hypothesis, *Quarterly Journal of Economics* 112(1), 3–55.

Cocco, Joao, 2000, Portfolio choice in the presence of housing, Working paper, London Business School.

Cocco, Joao, Francisco Gomes, and Pascal Maenhout, 1999, Portfolio choice over the life cycle, Working paper, Harvard University.

Constantinides, George, 1986, Capital market equilibrium with transaction costs, *Journal of Political Economy* 94, 842–862.

Constantinides, George, John Donaldson, and Rajnish Mehra, 2002, Junior can't borrow: A new perspective on the equity premium puzzle, *Quarterly Journal of Economics* 117, 269–296.

Dammon, Robert, Chester Spatt, and Harold Zhang, 2001, Optimal consumption and investment with capital gains taxes, *Review of Financial Studies* 14, 583–616.

Dammon, Robert, Chester Spatt, and Harold Zhang, Optimal asset location and allocation with taxable and tax-deferred investment, *Journal of Finance* 59, 999–1037.

Davis, Steven, Felix Kubler, and Paul Willen, 2002, Borrowing costs and the demand for equity over the life cycle, Working paper, University of Chicago and Stanford University.

Davis, Steven, and Paul Willen, 2001, Occupation-level income shocks and asset returns: Their covariance and implications for portfolio choice, Working paper, University of Chicago.

Deaton, Angus, 1991, Saving and liquidity constraints, *Econometrica* 59, 1221–1248.

de Nardi, Maria Cristina, Wealth inequality and intergenerational links, *Review of Economic Studies* (forthcoming).

Degeorge, Francois, Dirk Jenter, Alberto Moel, and Peter Tufano, 2002, Selling company shares to reluctant employees: France Telecom's experience, Working paper, HEC and Harvard University.

Dynan, Karen, Jonathan Skinner, and Stephen Zeldes, 2002, The importance of bequests and life cycle saving in capital accumulation: A new answer, *American Economic Review* 92, 274–278.

Epstein, Lawrence, and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Flavin, Marjorie, and Takashi Yamashita, 2002, Owner-occupied housing and the composition of the household portfolio, *American Economic Review* 92, 345–362.

Gomes, Francisco, and Alexander Michaelides, 2003, Portfolio choice with internal habit formation: A life cycle model with uninsurable labor income risk, *Review of Economic Dynamics* 6, 729–766.

Gourinchas, Pierre-Olivier, and Jonathan Parker, 2002, Consumption over the life cycle, *Econometrica* 70, 47–89.

Guiso, Luigi, Michael Haliassos, and Tullio Japelli, 2002, Household portfolios: An international comparison, in Luigi Guiso, Michael Haliassos, and Tullio Japelli, eds:. *Household Portfolios* (MIT Press, Cambridge, MA).

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, 1996, Income risk, borrowing constraints and portfolio choice, *American Economic Review* 86, 158–172.

Haliassos, Michael, and Alexander Michaelides, 2003, Portfolio choice and liquidity constraints, *International Economic Review* 44, 144–177.

He, Hua, and David Modest, 1995, Market frictions and consumption-based asset pricing, *Journal of Political Economy* 103, 94–117.

Heaton, John, and Deborah Lucas, 1996, Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.

Heaton, John, and Deborah Lucas, 1997, Market frictions, savings behavior, and portfolio choice, *Macroeconomic Dynamics* 1, 76–101.

Heaton, John, and Deborah Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.

Hu, Xiaoqing, 2001, Portfolio choices for home owners, Working paper, University of Illinois at Chicago.

Hubbard, Glenn, Jonathan Skinner, and Stephen Zeldes, 1995, Precautionary saving and social insurance, *Journal of Political Economy* 103, 360–399.

Jermann, Urban, 1999, Social security and institutions for intergenerational, intragenerational and international risk-sharing: A comment, *Carnegie-Rochester Conference Series on Public Policy* 50, 205–212.

King, Mervyn, and Jonathan Leape, 1998, Wealth and portfolio composition: Theory and evidence, *Journal of Public Economics* 69, 155–193.

Koo, Hyeng, 1998, Consumption and portfolio selection with labor income: A continuous time approach, *Mathematical Finance* 8, 49–65.

Laitner, John, 2002, Wealth inequality and altruistic bequests, *American Economic Review* 92, 270–273.

Lucas, Deborah, 1994, Asset pricing with undiversifiable risk and short sales constraints: Deepening the equity premium puzzle, *Journal of Monetary Economics* 34, 325–341.

Luttmer, Erzo, 1996, Asset pricing in economies with frictions, *Econometrica* 64, 1439–1467.

Luttmer, Erzo, 1999, What level of fixed costs can reconcile consumption and stock returns? *Journal of Political Economy* 107, 969–997.

Mankiw, N. Gregory, and Stephen Zeldes, 1991, The consumption of stockholders and nonstockholders, *Journal of Financial Economics* 29, 97–112.

Paiella, Monica, 2001, Transaction costs and limited stock market participation to reconcile asset prices and consumption choices, IFS Working paper 01/06.

Polkovnichenko, Valery, 2000, Heterogeneous labor income and preferences: Implications for stock market participation, Working paper, University of Minnesota.

Polkovnichenko, Valery, 2002, Life cycle consumption and portfolio choice with additive habit formation preferences and uninsurable labor income risk, Working paper, University of Minnesota.

Poterba, James, and Andrew Samwick, 1999, Household portfolio allocation over the life cycle, NBER Working paper no. 6185.

Saito, Makoto, 1995, Limited market participation and asset pricing, Working paper, University of British Columbia.

Storesletten, Kjetil, Chris Telmer, and Amir Yaron, 2001, Asset pricing with idiosyncratic risk and overlapping generations, Working paper, Carnegie Mellon University.

Telmer, Chris, 1993, Asset-pricing puzzles and incomplete markets, *Journal of Finance* 48, 1803–1832.

Vayanos, Dimitri, 1998, Transaction costs and asset prices: A dynamic equilibrium model, *Review of Financial Studies* 11, 1–58.

Viceira, Luis, 2001, Optimal portfolio choice for long-horizon investors with nontradable labor income, *Journal of Finance* 55, 1163–1198.

Vissing-Jørgensen, Annette, 2002a, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825–853.

Vissing-Jørgensen, Annette, 2002b, Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures, Working paper, Northwestern University.

Vissing-Jørgensen, Annette, and Orazio Attanasio, 2003, Stock-market participation, intertemporal substitution and risk aversion, *American Economic Review* 93, 383–391.

Yao, Rui, and Harold Zhang, Optimal consumption and portfolio choice with risky housing and borrowing constraints, *Review of Financial Studies* (forthcoming).