Stock Market Participation, Portfolio Choice and Pensions over the Life-Cycle

Steffan G. Ball

Stock Market Participation, Portfolio Choice and Pensions over the Life-Cycle*

Steffan Ball[†]

Federal Reserve Board

12 November 2008

Abstract

The empirical evidence on stock market participation and portfolio choice defies the predictions of standard life-cycle theory. In this paper we develop and estimate a model of portfolio choice that can account for the limited stock market participation and substantial portfolio diversification seen in the data. We present three realistic extensions to the basic framework: per period fixed costs, public pension provision, and a small chance of a disastrous event in the stock market. The estimated model is able to explain observed patterns at reasonable wealth levels, while keeping to a fairly simple framework. We demonstrate that it is no longer necessary to assume counterfactual asset holdings, heterogeneity in preferences, or implausible parameter values, in order to match key financial statistics.

Keywords: rare disasters, precautionary saving, portfolio choice, stock market participation and uninsurable labor income risk

JEL Classification: G11, H31

1 Introduction

This paper presents a life-cycle portfolio choice model, with realistically calibrated stochastic labor income and reasonable risk aversion, explaining important stylized facts of household asset allocations. Empirical studies consistently find that approximately fifty percent of US households do not invest in the stock market (whether directly or indirectly), and those that do participate hold only a fraction of their wealth in risky assets. There is substantial turnover in the set of participants, with widespread movements both in and out of the stock market. Despite recent developments in financial markets leading to greater levels of participation and higher shares of risky assets in household portfolios, the empirical evidence still presents a significant challenge to the life-cycle model. This paper demonstrates that the addition of a per period fixed cost to stock market participation, public pension provision, and a small possibility of a disastrous event in the stock market enables us to explain the observed limited participation and substantial portfolio diversification, while matching assets.

A wide literature presents models of intertemporal choice incorporating precautionary and retirement motives for saving; many of the empirical patterns of wealth accumulation and consumption have been accounted for within this framework (Hubbard, Skinner and Zeldes, 1995; Carroll, 1997; Attanasio *et al*, 1999; Gourinchas and Parker, 2002). More recently, a literature has emerged which allows for the simultaneous determination of consumption and portfolio allocation within a life-cycle framework. However, it has proved difficult to explain asset allocations without assuming unrealistic wealth accumulation, extreme parameter values or heterogeneity in preferences. Cocco, Gomes and Maenhout (2005) present a thorough analysis of the standard household portfolio choice model without any fixed cost considerations. They are able to force portfolio shares down to reasonable levels, in addition to obtaining significant age effects. However, this is achieved by accepting unrealistically high levels of saving.[1] Further, their model predicts one hundred

---

[1] A combination of using a very high coefficient of risk aversion ($\gamma = 10$) and assuming a small probability

percent participation at all ages, which we do not see in the data.

It has been argued that some form of information cost is required to move away from the prediction that all households should participate in the stock market at all times. This is corroborated by the empirical work of Paiella (2001) and Vissing-Jorgensen (2002), both of whom have shown that fixed costs to stock market participation can empirically rationalize the observed limited participation. These fixed costs can be one-off entry costs or per period costs, and the estimates of these costs are extremely low.[2] There has been significant progress with incorporating entry costs into household portfolio choice models. Alan (2006) calibrates the level of this cost to match moments from the PSID, and with this she is able to match the participation rate fairly precisely. But Alan's framework does not match the wealth data, and the resulting low levels of saving means that her model cannot address the issue of portfolio diversification.[3] Gomes and Michaelides (2005) set an exogenous entry cost, and attain reasonable age profiles of both participation and shares. However, they only achieve this by assuming preference heterogeneity, Epstein Zin utility functions, and very high levels of savings. While these entry costs are a convincing way to explain lower participation early in life, they cannot be the causal factor behind the large degree of turnover in the stock market or the low levels of participation for older households. Despite the empirical evidence strongly in favor of per period fixed costs, there has been little progress incorporating such considerations into a household portfolio choice model.

The striking result from research to date is that the standard life-cycle model appears able to explain only one key statistic at a time: matching wealth, or participation, or shares; but never more than one concurrently. The contribution of this paper is to illustrate how this limitation can be overcome with some relatively simple extensions to the standard framework. We increase the model's realism with three novel features, and demonstrate

---

of a zero income event, which triggers the precautionary response too much.

[2] Paiella (2001) finds per period fixed costs of US$ 95-175, while Vissing-Jorgensen (2002) finds that costs of US$ 260 can explain the behavior of 75% of nonparticipants.

[3] Alan's model predicts complete specialization in stocks for all participants.

how this enables us to calibrate a life-cycle model to match average moments of all three statistics, using data on US households' wealth and asset allocations.

We introduce these three innovations sequentially in order to separate out their effects. First, we account for the time costs involved with determining and managing an optimal portfolio by incorporating per period fixed costs to stock market participation. We find that introducing such cost considerations into the standard household model has a large effect on our simulated statistics, and in particular on the age profile of stock market participation. With this extension we can account for low participation at both early and late stages of the life-cycle at actual wealth levels, and we are able to explain some of the observed turnover in stock market participants. Second, we introduce a stylized pay-as-you-go public pension scheme, funded by a proportional tax on labor income. The retirement motive for saving is of great importance for wealth accumulation, and tends to be ignored or treated rather imprecisely in the literature on household portfolios.[4] Third, we modify the distribution for risky returns to allow for a rare disaster in the stock market. This simple extension enables the risky return to encompass large departures from the mean,[5] which can be thought of as representing episodes such as the crash of 1987, the collapse of Enron in 2001, or the financial crisis of 2008. This is an application of the "rare events hypothesis" first put forward by Rietz (1998), which has recently received growing attention in the context of the equity premium puzzle (Barro, 2006; Weitzman, 2007; Gabaix, 2007; Gourio, 2008; Veronesi, 2004). By allowing for such large negative movements we are able to match the household model to data with reasonable parameter values.

This is the first study to simultaneously match participation, shares and asset holdings with observables. In addition, we obtain reasonable life-cycle profiles for each of these variables, despite our model not calibrating on any age characteristics. We achieve all this

---

[4]For example, Alan (2006) only considers asset allocations for households up to the age of sixty, completely disregarding participation and portfolio share decisions during retirement.

[5]Notably absent under the usual assumption of i.i.d. normal returns, where all the action takes place in close proximity to the mean.

with a per period fixed cost of less than one per cent of the permanent component of labor income, a proportional tax rate of thirty percent and a small probability of a stock market crash.

In the next section we provide some data on asset holdings, participation and portfolio shares for US households. Section 3 develops our household model of portfolio choice. Section 4 provides simulated life-cycle profiles of asset accumulation, stock market participation rates and risky asset shares; showing the effect of introducing per period fixed costs, varying the public pension provision, changing the level of risk aversion, and introducing the possibility of stock market crashes. Section 5 concludes.

2   Data on Household Assets

This section details some stylized facts on asset accumulation and allocations which are addressed by our life-cycle model. Asset accumulation profiles are well documented in the consumption literature: mean asset holdings are seen to be low at the start of working life, to rise sharply when households are in their 30s and 40s, and then diminish gradually during retirement.[6] Low (2005), for example, details such a triangular age profile for asset accumulation; mean asset holdings are shown to grow until retirement, where they peak at a magnitude of around seven times greater than mean income. Low (2005) calculates that the US median asset-to-income ratio across working life is 1.84,[7] and it is to this statistic that we calibrate our simulated asset holdings.

When addressing household investment allocations it is important to distinguish between two distinct decisions: participation and portfolio shares. First, participation represents the binary choice of whether to participate in the stock market or not. Stock market participation in the US is currently just above fifty percent: Bertaut and Starr-McCluer

---

[6]This decline is not always present in median asset holdings, see Bernheim (1987).

[7]He uses the 1995 PSID wealth supplement for respondents aged 22 to 65, defining wealth to include housing wealth and using PSID weights.
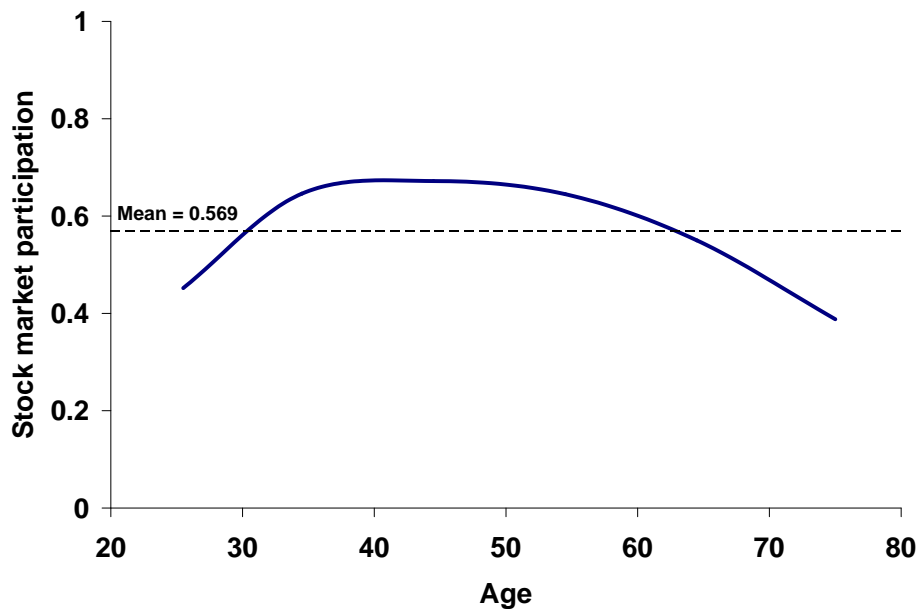
(a)
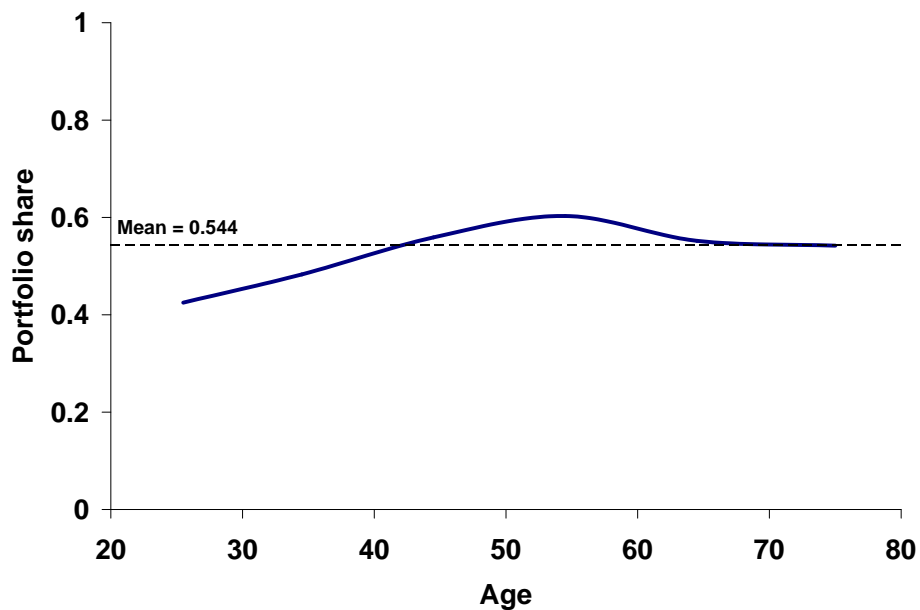
(b)

Figure 1: Asset allocation profiles in the Survey of Consumer Finances. Graph (a) shows the share of households holding risky assets. Graph (b) shows the portfolio shares conditional on stock market participation. The results for both graphs are taken from Bertaut and Starr-McCluer (2002), using weighted data from 1998 SCF. Risky assets include directly held stock; stock held through mutual funds, retirement accounts, trusts and other managed assets; corporate, foreign and mortgage-backed bonds; business equity; and investment real estate.

(2002) report that 56.9% of American households hold risky assets.[8] There is a clear age effect on participation, and a hump shape in the age profile of people who hold risky assets has been much discussed (Alan, 2006; Ameriks and Zeldes, 2001; Poterba and Samwick, 1997; Gomes and Michaelides, 2005). The top panel in figure 1 uses data from the US Survey of Consumer Finances on stock market participation.[9] We see limited participation at early and late stages of life, with a peak around age forty-five. This clearly details the hump shape discussed in the literature. Participation is also increasing in wealth, with richer households much more likely to hold risky assets than poorer households. Using multivariate probit techniques, Bertaut and Starr-McCluer (2002) find that both age and wealth effects are significant for the participation decision in the US. Interestingly, there is substantial turnover in this set of participants, with many people moving both in and out of the stock market. Vissing-Jorgensen (2002) estimates an extremely high turnover rate of around 28% for households with positive wealth, with exit accounting for a substantial portion of those changing participation status.[10]

The second important asset allocation statistic is the portfolio share. This represents the share of risky assets to total assets, conditional on participating in the stock market. These portfolio shares are documented to be well below 100% for almost all individuals (Ameriks and Zeldes, 2001; Bertaut and Starr-McCluer, 2002; Gomes and Michaelides, 2005; and others). Bertaut and Starr-McCluer (2002) find that the average portfolio share of stock market participants in the US is 54.4%.[11] The bottom panel of figure 1 shows a fairly flat age profile, indicating that portfolio shares do not vary substantially over the life-cycle. Some decline can be seen after retirement, which is consistent with

---

[8]Using 1998 SCF weighted data. Risky assets are defined as including: directly held stock; stock held through mutual funds, retirement accounts, trusts and other managed assets; corporate, foreign and mortgage-backed bonds; business equity; and investment real estate.

[9]Note that cohort effects may be important in asset allocation decisions, but our framework focuses on age effects and does not model differences between cohorts. See Alan and Ball (2007) for an example where asset allocations are modelled for different cohorts.

[10]Comparing households' participation status in 1989 with that in 1994, using data from the PSID wealth supplements.

[11]Gomes and Michaelides (2005) find similar results, with an average share of 54.8%.

financial advisors recommending clients to shift their portfolios away from risky assets as they age.[12] However, econometric studies tend to find that age effects are not significant (see, for example, Bertaut and Starr-McCluer, 2002).

The model in this paper is calibrated to match these three average statistics (asset holdings, participation and shares), explaining the limited participation and strong portfolio diversification seen in the data. As a robustness check for our model's success, we also seek to account for the age profiles of these statistics (particularly that of participation given its documented empirical significance) and ask whether we can explain some of the turnover in stock market participants.

3 Life-Cycle Model

Individuals are assumed to maximize the sum of expected discounted lifetime utility in light of uncertain and uninsurable labor income and rate of return risk. Utility is defined over a single nondurable consumption good, and is additively separable across time. Let there be T periods in the life-cycle, of which a given K periods are spent in retirement and $T - K$ give the working life. This standard set-up gives the following constrained maximization problem:

$$\max_{\omega_s, C_s} E_t \left[\sum_{s=t}^{T} \beta^{s-t} u(C_s) \right] \tag{1}$$

subject to,

$$A_{s+1} = (A_s + Y_s^n - C_s - \mathbb{I}\{partic.\}FP_s)[\omega_s(1 + r_s^{eq}) + (1 - \omega_s)(1 + r)] \tag{2}$$

during working life ($s \leq T - K$), and

$$A_{s+1} = (A_s + \Gamma - C_s - \mathbb{I}\{partic.\}FP_{T-K})[\omega_s(1 + r_s^{eq}) + (1 - \omega_s)(1 + r)] \tag{3}$$

---

[12]So called "life-cycle" funds do this automatically for investors.

during retirement $(s > T - K)$.

We impose the terminal condition $A_{T+1} = 0$[13]. In each period β is the discount factor; C_t is consumption; A_t is the amount of assets held at the start of period t; ω_t is the portfolio share; and Y_t^n is exogenous real disposable labor income net of taxes. During working life, labor income is uncertain and non-diversifiable and is taxed at a proportional rate of τ:[14]

$$Y_{t+1}^n = (1 - \tau)Y_{t+1} \tag{4}$$

Following MaCurdy (1982), Abowd and Card (1989), and Moffitt and Gottschalk (2002), we assume that the labor income process can be broken down into a deterministic component and two stochastic components, a permanent and a transitory shock:[15]

$$Y_{t+1} = P_{t+1}v_{t+1} \tag{5}$$

where v_t is transitory income and P_t is the permanent component of labor income. The log of permanent income is assumed to follow a random walk with drift, where G_t represents the deterministic growth trend:

$$P_{t+1} = G_{t+1}P_t\varepsilon_{t+1} \tag{6}$$

Both transitory and permanent shocks are assumed to be independently and identically distributed lognormally such that $\ln(v_t) \sim N(-0.5\sigma_v^2, \sigma_v^2)$ and $\ln(\varepsilon_t) \sim N(-0.5\sigma_\varepsilon^2, \sigma_\varepsilon^2)$. The income process is truncated such that a zero level of income cannot be realized.

For the final K periods of life, individuals are retired, during which time they receive

---

[13]Allowing for bequests in the model would be a relatively simple extension. Including such an additional savings motive will simply increase the size of the fixed cost required to keep stock market participation at the correct level.

[14]Following Gomes and Michaelides (2005) Y_t represents income net of housing expenditures. This accounts for the life-cycle pattern of housing expenditures, but implicitly assumes that housing decisions are separable from consumption decisions and independent of wealth. While this is a strong assumption, modelling explicit housing decisions is beyond the scope of this paper.

[15]Although this approach is commonly used in the literature, recent studies have pointed towards the weakness of empirical evidence for such restricted income profiles (see Guvenen, 2007).

a constant and certain pension income, Γ. We construct a pay-as-you-go pension scheme, somewhat akin to the arrangements in the US, with pension payments being funded out of contemporaneous tax revenues.[16] Allowing for the tax collection requirement of public pension provision is new to the literature on household portfolios.[17] Let us assume that an economy is made up of H individuals. At any particular point in time, we assume that there exists an equal mass of individuals at each age; with working population share being $\frac{T-K}{T}$, and the fraction of households in retirement given by $\frac{K}{T}$. Tax revenue collected from individual i at age a is given by τY_a^i. Summing pension contributions over all individuals, H, and all working ages, $T - K$, and normalizing by the number of retirees, $H.K$, gives the following constant annual pension income for each retired household:[18]

$$\Gamma = \sum_{a=22}^{T-K} \frac{\sum_{i=1}^{H} \tau Y_a^i}{H.K} \tag{7}$$

We assume the portfolio investment is composed of two distinct saving tools, a riskless and a risky asset, with ω_t representing the share of assets held in risky forms. The riskless asset has a constant real return of r and the risky asset has a stochastic real return of r_t^{eq}, with shocks to the excess return assumed to be independently and identically distributed:

$$
\begin{aligned}
r_{t+1}^{eq} - r &= \mu + \eta_{t+1} \\
\eta_{t+1} &\sim N(0, \sigma_{eq}^2)
\end{aligned}
\tag{8}
$$

We do not impose any correlation between stocks and labor income as the empirical evidence on the size and magnitude is mixed. Heaton and Lucas (2000) find insignificant estimates for all but the highest educational group, and Davis and Willen (2000), surpris-

---

[16]While the actual US pension system has some earnings-related benefit, we assume a fairly simple flat rate payment.

[17]Previous work either imposed exogenous retirement pension income, or ignored asset allocation decisions during retirement (see Alan (2006)).

[18]We invoke the law of large numbers to assure that the pay-as-you-go pension scheme is always balanced.

ingly, find negative correlations for low educational groups. Further, the effect of such a correlation has been shown to make very little difference to participation and portfolio shares in practice (Gomes and Michaelides, 2005), unless unrealistically high correlation coefficients are chosen.

In order to invest in risky assets, in any given period, we assume an individual has to pay some fixed cost. It is clear that gathering and processing financial information plays an important role in investment decisions: acquiring data on financial markets, monitoring brokers, and keeping up to date with trends, all take considerable time. These financial decisions must be made at each and every point when an individual invests in stocks, implying that any relevant information costs must be considered on a per period basis. In addition to these information costs, holding stocks leads to substantially more complicated annual tax returns. In the US, households investing in risky assets must fill out schedules D and D1, which take considerable time to complete. Econometric studies by both Paiella (2001) and Vissing-Jorgensen (2002) have asserted the empirical significance of such fixed costs paid in each and every period of stock market participation. We denote these fixed time costs by F, and we include them in the asset evolution equations ([2] and [3]) scaled by the permanent component of labor income, FP_t.[19] This scaling follows from the motivation of these costs representing opportunity costs of time: richer households have to forego greater earnings while spending time filling out their schedule D and D1 than poor households. The fixed costs are interacted with an indicator function, $\mathbb{I}\{\cdot\}$, which is equal to one for every period of participation and zero otherwise.

Some limited borrowing is allowed, up to the discounted sum of the minimum income individuals will receive in each remaining period, with no borrowing permitted against pension income. A short sales constraint is imposed on equity such that the portfolio shares must always lie between zero and one (inclusive). Therefore, agents can only borrow at the risk free rate and can only invest in risky assets if they hold positive balances.[20]

---

[19] Assuming F is constant over the life-cycle abstracts away from issues of learning.

[20] While it is true that we do see households simultaneously borrowing and holding stocks in the data,

This set-up gives us a system with three state variables (A_t, Y_t, P_t) and two control variables (C_t, ω_t). The first order condition of the value function with respect to consumption is given by:

$$u^{'}(C_t) = \beta E_t \left\{ [\omega_t(1 + r_t^{eq}) + (1 - \omega_t)(1 + r)]u^{'}(C_{t+1}) \right\} \tag{9}$$

and the first order condition with respect to portfolio share is given by:

$$0 = E_t \left[(r - r_t^{eq})u^{'}(C_{t+1}) \right] \tag{10}$$

These two conditions can be solved for the following policy functions in each period:

$$C_t(A_t, Y_t, P_t) \tag{11}$$

$$\omega_t(A_t, Y_t, P_t) \tag{12}$$

These two optimal decision rules are solved recursively from the final period for a discrete number of points in the state space. The details of the solution method are given in Appendix A.

4 Simulation Results

Once the optimal policy functions have been determined, we can simulate the model to imitate the behavior of a population of households and report average allocations to show the effect of introducing fixed costs and changing the public pension scheme. Initially, we detail the functional forms, parameter values and calibration statistics used. Then we show the results of simulating the model for 10,000 ex-ante identical households who differ in the realization of shocks and their initial wealth.

this restriction is imposed on the model for simplification.

We put forward the results in four subsections: first, we analyze the effect of introducing per period fixed costs to our baseline life-cycle model, and show that this modification changes the participation results substantially. Second, we illustrate the effects of altering the generosity of public pensions, and explain the link between pension generosity and the calibrated level of fixed costs. Third, we calibrate the coefficient of risk aversion, and demonstrate how our model is able to match both participation and shares. Finally, we detail the effect of including a small possibility of a stock market crash, and argue that this enables us to explain all key financial statistics. This structure permits us to disentangle the different consequences of varying fixed costs, public insurance and return risk in order to better fit the model to the data.

Parameters and Calibration   The utility function is assumed to take the typical constant relative risk aversion (CRRA) form:[21]

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma} \tag{13}$$

with a baseline value of 2 for the coefficient of risk aversion. This is in line with estimates based on consumption data: Attanasio and Weber (1995) estimate the coefficient of risk aversion to be around 1.5, Gourinchas and Parker (2002) estimate risk aversion to range from 0.28 to 2.29, and Alan and Browning (2003) find ranges from 1.2 to 1.95.

The riskless rate of return is set at a constant value of 1.6%, which represents the average annualized real rate of US 3-month treasury bills from 1960-2000. The risky return is drawn from a distribution with a mean of 5.6%, corresponding to a 4% equity premium,[22] and a standard deviation of 0.2.[23] Following Haliassos and Michaelides (2003)

---

[21] Our choice of CRRA utility stems from our objective of matching asset allocations within a standard life-cycle framework. Alternative utility specifications may prove helpful in explaining household portfolios; we leave such analysis for future work.

[22] A moderate equity premium is common in this literature; see Gomes and Michaelides (2005), Haliassos and Michaelides (2002), Yao and Zhang (2005), Cocco (2000), Campbell et al. (2001). Alternatively this can be thought of a 7% return on equity net of 20% tax.

[23] Taken from Alan (2006). Gomes and Michaelides (2005) and Haliassos and Michaelides (2002) use a similar value of 0.18

Parameter	Value
CRRA (γ)	2
Riskless rate (r)	1.6
Mean risky return (r^{eq})	5.6
Variance of risky return	0.04
Variance of permanent shock (σ_n^2)	0.01
Variance of transitory shock (σ_u^2)	0.0225
Growth of permanent component of income (G)	0.03
Proportional tax (τ)	0.3

Table 1: Exogenous parameters

we set the deterministic trend of income, G, to 3%. Standard magnitudes for the variance of permanent and transitory shocks to income are used. We take the values directly from Gomes and Michaelides (2005), which are very similar to those used by Gourinchas and Parker (2002) and Carroll (1997). All exogenous parameters are shown in table 1.

The baseline tax rate of 30% gives a retirement replacement rate of 45% of final working period labor income. In reality, the pension tax system in the US is more complicated than the set-up assumed here, but this proportionate rate is chosen as a reasonable proxy. The implied replacement ratio is in line with those used in the literature, for example, Low (2005) used a 55% replacement rate.

The values for the fixed cost, discount rate and (in the final three scenarios) risk aversion are chosen by calibration.[24] The value of the fixed cost is chosen to match simulated mean participation with the observed mean in the data (0.57). We choose the discount rate such that simulated ratio of median assets to median incomes matches the data for the working households (1.84).[25] Additionally, in the final three scenarios, we choose the risk aversion

---

[24]While we choose to calibrate parameters to match our model to data, an alternative would be to undertake a full structural estimation procedure (such as simulated minimum distance). In practice estimation is very similar to calibration, and we take the calibration route to avoid having to estimate the variance-covariance weighting matrix. We do not expect the results to be sensitive to this choice.

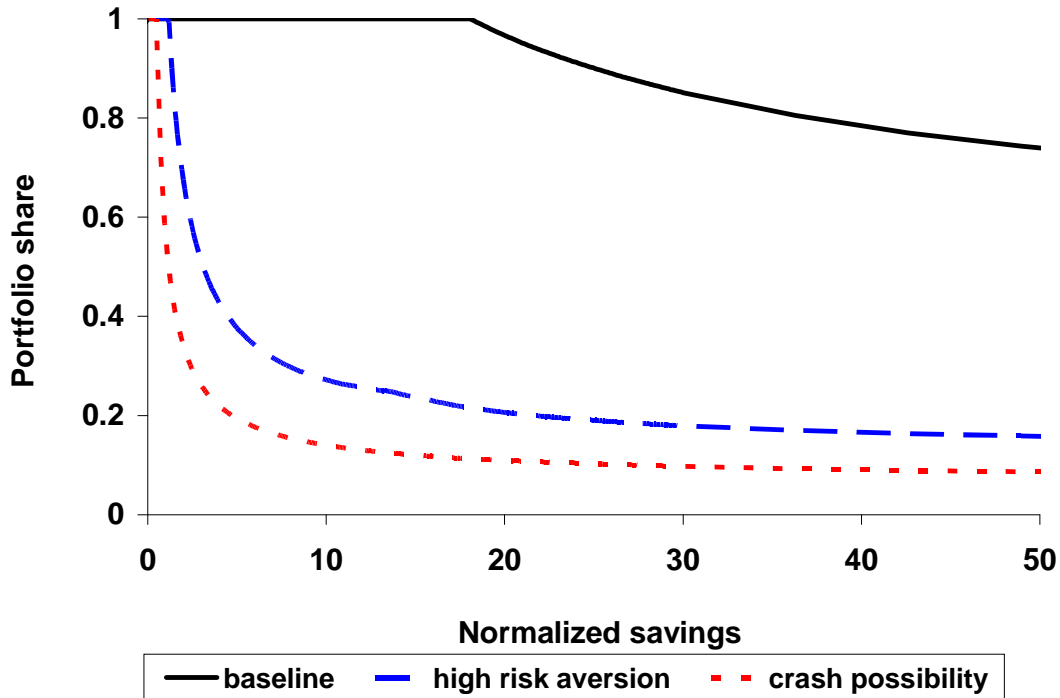[25]Except for the scenarios where we explicitly fix the discount rate.

Figure 2: Portfolio share policy functions for age 27, under different model specifications.

parameter to match the mean portfolio share statistic of 0.54. Calibration parameters and statistics are shown in table 2.

At this point, it is important to point out that one of the main contributions of this paper is to emphasize the importance of matching asset holdings. In these household models both participation and portfolio shares are highly sensitive to the level of assets. Many of the studies claiming to explain certain features of asset allocations implicitly rely on counterfactual asset holdings for their results. To understand the importance of asset levels the solid line in figure 2 plots the portfolio share policy function for an individual of age 27 under our baseline parameters.[26]

Studies such as Gomes and Michaelides (2005) are able to obtain reasonable portfolio shares by allowing asset holdings to be unrealistically high. This pushes households onto

---

[26]The choice of age 27 is completely arbitrary, analogous policy functions would be attained for any other age (see Appendix B). Portfolio shares are shown as a function of savings normalized by the permanent component of labour income.

the non-linear sections of their portfolio share policy functions (positioned on the solid line towards the far right in figure 2), enabling movements away from complete specialization in risky assets. In contrast, studies claiming to be able to match participation (Alan, 2006) achieve this by forcing assets down to levels where the borrowing constraint often binds. In fact, in these models non-participants are implicitly assumed to be those hitting their borrowing constraints. With such low amounts of saving, agents are operating on the horizontal section of their portfolio share policy functions (positioned on the solid line close to the y-axis in figure 2), hence these authors find very high shares, typically unity. In this paper we ensure a more genuine matching of asset allocations by explicitly calibrating our model to the actual level of assets observed in the US.

4.1   Introducing Per Period Fixed Costs

Standard Model without Fixed Costs In our initial parameterization we set the per period fixed costs to zero and calibrate the model to match the PSID median asset to income ratio over working life (1.84). The dashed-line graphs in figure 3 show the cross-sectional mean profiles of normalized asset holdings,[27] stock market participation and conditional portfolio shares.

Under the baseline parameters the discount rate (δ) necessary to match the level of assets is 0.5%, giving a discount factor $\beta = 0.995$. This high calibrated discount factor results from the high effective impatience introduced with our public pension scheme. The high tax, and subsequent high pension payments, increases the steepness of disposable income growth requiring only a modest δ to keep savings at the required level. The dashed line in the top panel of figure 3 shows agents accumulating assets up until retirement, and then running their balances down during the last 15 years of non-working life, reaching levels that are in the same order as those seen in the data (Low, 2005). Average assets are always positive, but some individuals do borrow, up to the endogenous borrowing limit, in

---

[27]This is asset holdings divided by the permanent component of annual labor income.

|	Simulated statistics			Parameters		
Scenario	$\frac{\text{Median}[A_{i,t}]}{\text{Median}[y_{i,t}]}$	Mean participation	Mean shares	$1+\delta$	Fixed costs	CRRA
US Data	**1.84**	**0.57**	**0.54**	-	-	
No FC tax = 0.3	1.84	0.95	0.99	1.005	-	2.0
Calibrated FC tax = 0.3	1.84	0.57	0.98	1.003	0.046	2.0
Keeping tax = 0.3	1.84	0.57	0.98	1.003	0.046	2.0
discount rate tax = 0.2	2.98	0.57	0.95	1.003	0.071	2.0
constant tax = 0.0	4.89	0.57	0.90	1.003	0.107	2.0
Calibrating tax = 0.3	4.11	0.57	0.54	1.003	0.025	9.5
CRRA tax = 0.3	1.84	0.57	0.54	1.262	0.012	12.1
Crash possibility tax = 0.3	1.84	0.57	0.54	1.090	0.008	8.8

Table 2: Calibrated parameters and statistics

the early stages of life.

The dashed line in figure 3(b) shows that participation in the stock market is 100% for all but a few years in early and late life, and only those individuals with negative savings do not participate. This is in conflict with the data, clearly illustrating the famous "portfolio participation puzzle", discussed at length in the household finance literature. Moreover, in the absence of transactions costs there is no explicit participation decision; nonparticipants are simply those agents in debt, who are constrained not to hold risky assets by the bounds on ω_t. Turnover is extremely low at 0.5%,[28] and this is dominated by households entering the stock market rather than exiting. Figure 3(c) details portfolio shares of stock market participants, and these are high across the life-cycle. The mean portfolio share is 0.99 which is well above the actual value of 0.54. These results exemplify the inability of the standard framework to explain household asset allocations at reasonable wealth levels: neither participation nor shares are close to being matched.
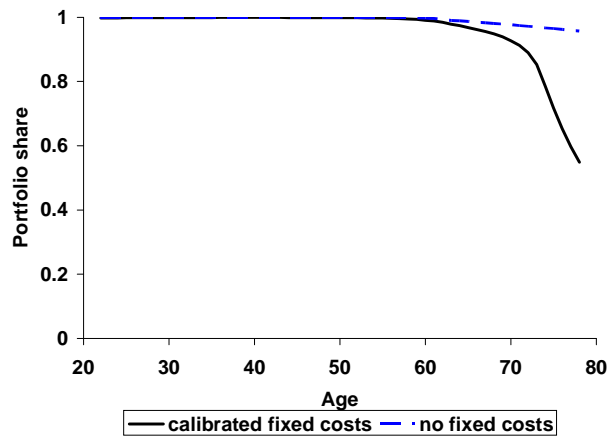
---

[28]Measured as the proportion of periods where households have positive wealth and participation status changes.

(a)

(b)

——calibrated fixed costs — ·no fixed costs
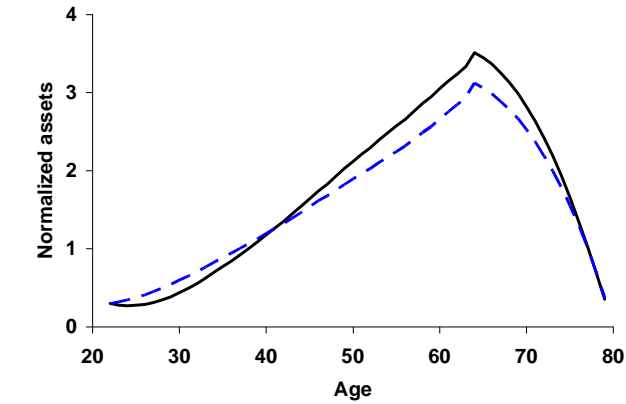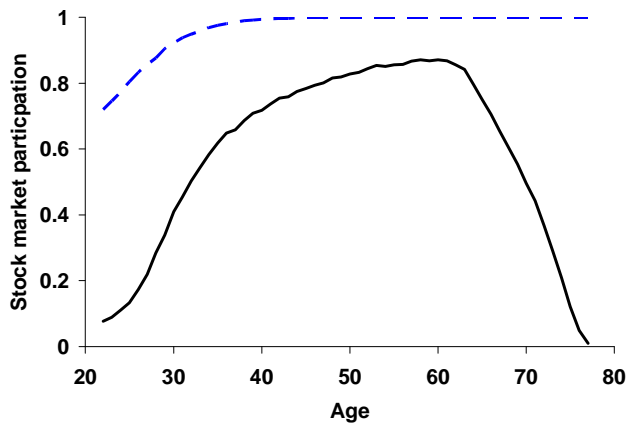
(c)

Figure 3: The effect of introducing per period fixed costs. Graph (a) shows assets normalized by permanent income. Graph (b) shows stock market participation by age. Graph (c) shows conditional portfolio shares at each age. The discount rate is calibrated separately in each scenario such that median assets match the data. For the scenario with fixed costs (represented by the solid line in the graphs), the costs are calibrated to match mean participation.

Including Per Period Fixed Costs Introducing market frictions, and calibrating the model to match the asset to income ratio of 1.84, leads to a slightly lower discount rate (from 0.5% to 0.3%).[29] These per period fixed costs increase the cost of saving, and so the discount rate must fall to keep the level of savings constant. The level of the per period fixed cost is calibrated to match the average participation rate of 0.57, giving a value of 4.6% of the permanent component of labor income. The age profiles are detailed in the solid lines in figure 3. We can see that the increased cost of saving depresses asset accumulation early in life. This reduced precautionary saving is compensated by more retirement saving in mid-life, in order to keep the median asset level calibrated to 1.84.

The introduction of per period fixed costs has a marked effect on the stock market participation profile, with a clear hump shape now evident in figure 3(b). Few individuals accumulate sufficient wealth in the first ten years of working life to make it viable to pay the fixed costs of participation, and they hold only riskless assets in their portfolio. As individuals age they accumulate assets, risky investments become profitable and participation increases - peaking at age 60 with 87% participation. As assets are drawn down later in life, the number of individuals participating is seen to fall gradually back to very low levels. It is important to note that we are not calibrating the model to match the participation age-profile, but we obtain a clear hump shape in all simulations incorporating per period fixed costs.

We see much greater turnover in participation, with the rate increasing twelvefold to 5.5% when per period fixed costs are included. As these costs must be paid during each period of participation we are able to explain exit from the stock market, as well as entry. A bad realization of labor income may drive down the wealth of stock market participants to levels where it is no longer worth paying the fixed cost, and these agents will exit the stock market. The addition of per period fixed costs enables the life-cycle model to explain exit from the stock market for agents with positive wealth. In previous work, households only

---

[29]The results from keeping the discount rate the same and only calibrating the fixed cost are not significantly different.

left the stock market when wealth hit their borrowing constraint, which is unconvincing given the empirical evidence of many households exiting with positive wealth holdings. However, our turnover rates do not reach the high levels documented by Vissing-Jorgensen (2002), though this may be partly due to differing definitions for turnover.[30]

These fixed costs ensure that stock market participants endogenously hold greater levels of savings, generating a more pronounced age profile in the conditional portfolio shares - as evident in figure 3(c). However, the model continues to predict complete specialization in risky assets early in life with a mean portfolio share of 0.98,[31] which is considerably higher than observed in the data.

4.2   Changing the Generosity of the Public Pension Scheme

In this section we analyze the effect of changing the public pension generosity in a model where fixed costs are calibrated to match the US average participation in the stock market. The public pension scheme we have introduced is a form of forced-saving, and varying the generosity of provision will lead to different levels of voluntary saving and stock market participation. Lowering the tax rate leads to less pronounced income growth over the working life and smaller state pension payments during retirement. At a given set of preference parameter values, this lower effective impatience will lead to greater amounts of precautionary and retirement savings.[32]   In this scenario we hold the discount rate constant, at the level calibrated in section 4.1, while varying the tax rate. We calibrate the per period fixed costs to match average participation.

We find that reducing the generosity of the public pensions scheme has a significant effect on the required level of fixed costs: increasing from 4.6% (when the tax rate is 30%),

---

[30]In the model turnover is not linear with age: we see a lot of entry and exit at early and late stages of life when wealth is low, with very little movement in participation during mid-life. Thus, the relatively low turnover rate partly reflects averaging over the life-cycle which is not captured by Vissing-Jorgensen's data.

[31]This is marginally smaller than the mean share predicted by the model without fixed costs, as savings behavior has only been altered modestly.

[32]To keep the amount of voluntary savings constant the discount rate would have to increase as pension provision is reduced.
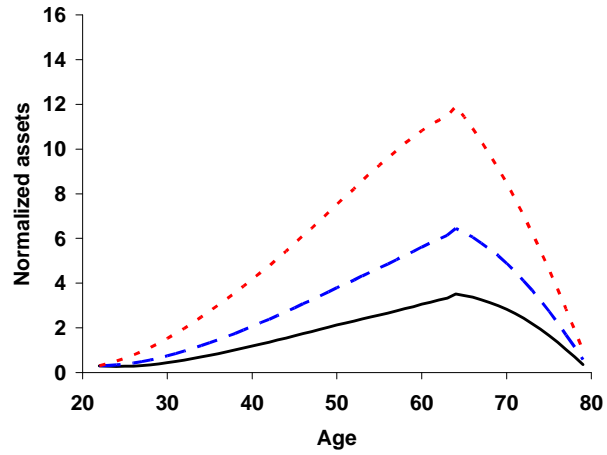
(a)
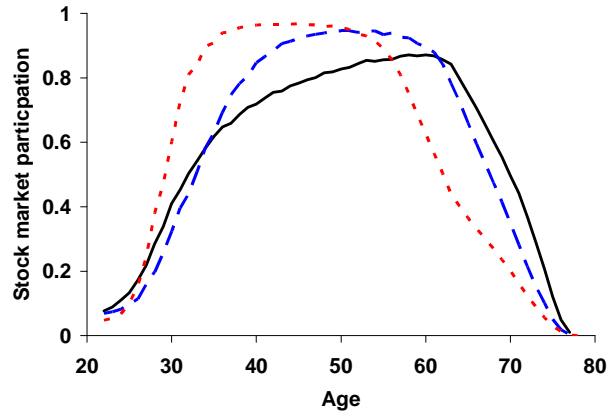
(b)
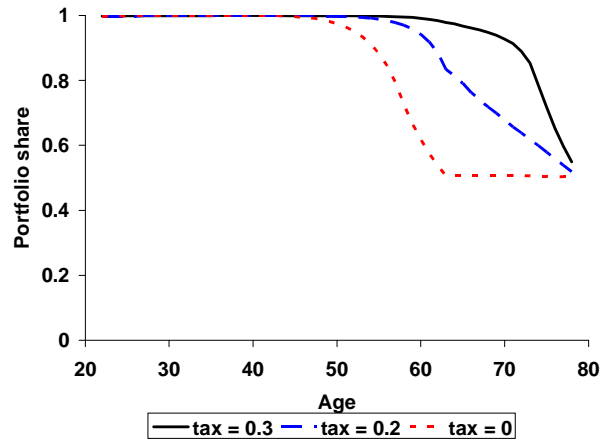
(c)

Figure 4: The effect of varying the generosity of public pension provision, keeping the discount rate constant (calibrated to the 30% tax rate). Graph (a) shows normalized assets. Graph (b) shows stock market participation by age. Graph (c) shows conditional portfolio shares at each age. The fixed costs are calibrated separately to match mean participation.

to 7.1% (when the tax rate is 10%), to 10.7% (when the tax rate is zero). In figure 4(a) the crowding out effect of our forced-saving public pensions is clearly evident. Reduced public pension provision is compensated for by greater private retirement savings, leading to very rapid wealth accumulation over the working life. This increased saving allows earlier initial entry to the stock market and higher participation rates during mid-life. Smaller tax rates lead to more diversification of assets; however, shares remain unrealistically high at early stages of life, and mean shares are substantially above their true value (see table 2).

This scenario demonstrates a number of important issues. First, households accumulate assets too rapidly in the absence of public pensions. Second, the hump shape in the age profile for participation is too pronounced for low levels of pension provision. Finally, fixed costs are negatively related to the intensity of public insurance. This final point implies that the small empirical estimates of these fixed costs should not be surprising.[33] Given the well developed social security system prevalent in the US, households hold modest retirement savings and only a small fixed cost is required to explain the observed limited participation.

4.3 Calibrating Risk Aversion

The above scenarios demonstrate how restricting ourselves to plausible participation and asset holdings lead to excessively high portfolio shares. This trade off between participation and shares has been well documented in the household finance literature. With the standard model specification, the structure of power utility dictates that large diversification can only take place if savings are substantial (see the policy function depicted in the solid line in figure 2). However, when savings are large, households are highly motivated to participate; making it impossible to reconcile limited participation with small portfolio shares. Our per period fixed cost extension enables us to break this trade off, and match both features within one model.

---

[33]See Paiella (2001) and Vissing-Jorgensen (2002).

The shapes of the portfolio share policy functions are determined to a large extent by the degree of risk aversion.[34] Furthermore, the absolute lower bound of these functions is governed by this parameter through the Samuelson rule.[35] The difficulty in matching shares stems from the fact that low risk aversion, as advocated by econometric studies on consumption, leads to a very high value for this lower bound. Under our baseline parameters, this lower bound for portfolio shares is 0.5. Given such diversification is only reached as savings go to infinity, we cannot hope to obtain shares of 0.54 at plausible levels of savings.

By increasing the coefficient of risk aversion, we can reduce this lower bound, shifting the policy functions down and inwards towards the origin. This can be seen by comparing the solid and dashed lines in figure 2.[36] This identifies a link between risk aversion and shares: for any given level of savings we obtain smaller shares with higher risk aversion. Our model set-up allows us to take advantage of this relationship, and calibrate the risk aversion parameter to match simulated average shares with the data. In this section we detail the results of this calibration in two different scenarios: firstly, keeping the discount rate at our baseline parameterization; and secondly, re-calibrating the discount rate to match actual savings.

Matching Participation and Shares In this scenario we fix the discount rate to the value calibrated in section 4.1. Rather than exogenously set the coefficient of risk aversion, we choose it by calibration such that simulated mean shares match the 0.54 value seen in the data. We calibrate fixed costs to match mean participation as before. The calibrated parameters and simulated statistics are shown in table 2, and the simulated profiles are represented by the dashed line in figure 5.

---

[34]Other parameters also determine the shape of the policy functions; see Alan and Ball (2007) for a discussion of the effect of changing the transitory income variance and moments of the expected return.
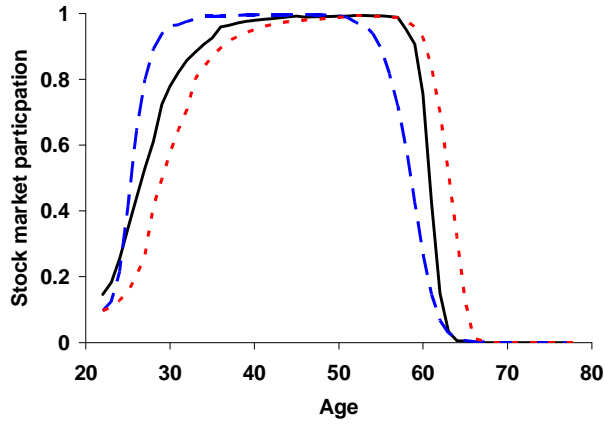
[35]As savings go to infinity, labor income becomes of negligible importance, and we approach the complete markets outcome (see footnote in section 4.2).

[36]The dashed line in figure 2 shows the policy function for a risk aversion coefficient of 8.8 with all other parameters set to the baseline levels.

(a)

(b)

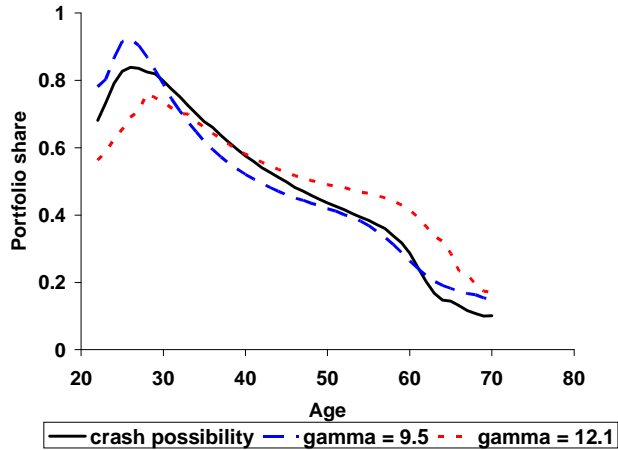| crash possibility | — · gamma = 9.5 | – – gamma = 12.1 |
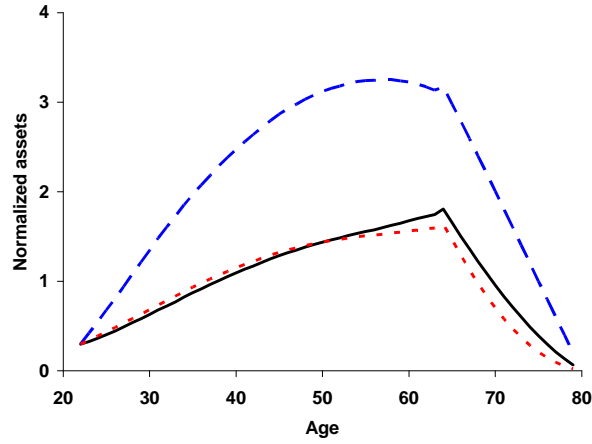
(c)

Figure 5: Calibrating risk aversion. Graph (a) shows normalized assets. Graph (b) shows stock market participation by age. Graph (c) shows conditional portfolio shares at each age. The fixed costs, discount rates and risk aversion are calibrated separately to match mean participation, asset holdings and portfolio shares respectively.

We obtain a high risk aversion parameter of 9.5 and a per period fixed cost of 2.5% of the permanent component of labor income. Figure 5(b) shows the age profile for participation displaying the expected hump shape. The portfolio share profile in figure 5(c) is now much more realistic, with values well below unity for most of life and declining as households age.[37] Thus, we have moved away from the prediction of complete specialization in stocks at young ages. This scenario demonstrates that we are able to match both mean participation and mean shares, within the standard life-cycle framework. However, at this discount rate our asset holdings are very high. This result is comparable to that of Gomes and Michaelides (2005), without requiring preference heterogeneity or non-power utility.

Matching Participation, Shares and Wealth In this scenario our model is shown to match all three key financial statistics. We calibrate the discount rate to match median asset holdings (1.84), risk aversion to match mean portfolio shares (0.54) and per period fixed costs to match mean participation (0.57). The results are detailed in the dotted lines in figure 5.

The age profiles are all reasonable, in particular we have portfolio diversification at all ages - see figure 5(c). Calibrated fixed costs now fall to 1.2% of the permanent component of labor income, which is close to the empirical findings discussed earlier. The calibrated coefficient of risk aversion is high at just over 12. In order to counter this risk aversion, and the entailing high precautionary motive, an uncomfortably large discount rate of 26.2% is required to match asset holdings. Even the high tax rate of 30% cannot moderate impatience at these parameter values.

---

[37]The slight initial increase in portfolio shares is a result of the sharp rise in participation at these ages. At first only those who begin life with large asset holdings can participate, and these individuals have highly diversified portfolios; as more individuals participate, average wealth of the participants will fall leading to higher average shares.

4.4 Stock Market Crash Possibility

A potential explanation for the model's failure to explain the stylized facts at reasonable preference parameters is risk misspecification. Solving such a dynamic optimization problem involves specifying a subjective probability distribution for the stochastic processes. The standard approach for this involves applying the law of large numbers, and assuming we have sufficient time series of data to substitute past sample moments for future population moments. If data are drawn from a sample which is unrepresentatively tranquil, the law of large numbers may not hold and *ex post* realized returns will differ from *ex ante* expected returns. We posit that extreme movements in stock returns are more likely than assumed by the Gaussian distributions and finite sample moments that we tend to use.

Normality is adopted mainly for analytical convenience, and there is now growing debate over the adequacy of such thin-tailed distributions for approximating financial stochastic variables, most vocally by Mandelbrot (1963, 1967, 1997, 2001). Stock returns are subject to many random leaps and jumps, pushing the realized value far from the mean. Examples include the stock market crash of 1987, the collapse of Enron in 2001, the financial crisis of 2008, and to an even greater extent the Great Depression. Solving dynamic programs with discretized normal distributions does not allow for such large departures from the mean, and this is likely to lead to an underestimation of the true risk of holding stocks.

We propose a simple extension to the standard Gaussian framework, whereby we insert a discrete i.i.d. rare disastrous state in the stock market. We assume that such an outcome occurs with a known proportional reduction in returns, and with a known probability. This extension is analogous to the Carroll (1992, 1997) zero income shock possibility. We set the probability of a stock market crash to 1%, and assume that during the disastrous event households lose all wealth held in risky assets, that is $1 + r_t^{eq} = 0$.[38] Our assumed 1%

---

[38] We assume that disasters do not lead to default on government bills, and so the risk free rate is unaffected by these rare events. While it could be argued that a disastrous event in the stock market is likely to affect the risk free rate (especially if inflation becomes high), historically it seems that government bills perform better than stocks in most economic crises not caused by war and outright default is rare in major developed economies (Barro, 2006).

probability is similar to the probabilities associated with disasters in the related literature on extreme events and the equity premium puzzle (Rietz, 1988; Barro, 2006). However, we are assuming a potential loss of much greater magnitude than any realized during the sample period. Individuals are assumed to worry about the possibility of the stock market being completely wiped out by a crash, even though no such extreme event has been witnessed: this is an application of the Peso Problem Hypothesis.

This small probability of a stock market crash has a significant effect on the portfolio share policy functions. The dotted line in figure 2 shows how this extension shifts the policy function further down and inwards towards the origin.[39] This enables significant portfolio diversification at low, realistic levels of savings. Under these assumptions, we again calibrate the discount rate, fixed cost and risk aversion parameters to match wealth, participation and shares respectively (as in the previous scenario). The results are detailed in the final row in table 2 and simulated profiles are shown in the solid line in figure 5.

The life-cycle profiles are all reasonable: with assets being accumulated at the correct pace, a hump shape in participation, and portfolio shares being diversified across all ages. We have a 5% turnover rate in stock market participants. The most striking result is on the calibrated parameters, we are now able to match the data at plausible parameter values. A risk aversion of 8.8 is required to enable modest share holdings. This estimate is somewhat higher than suggested by empirical studies using consumption data, however, it remains within the bounds considered reasonable by Mehra and Prescott (1985) and is consistent with the macro finance literature.[40] The per period fixed cost is now calibrated to less than one percent of annual labor income, which is closely in line with the empirical estimates of Paiella (2001) and Vissing-Jorgensen (2002). The calibrated discount rate is 9%, which is well within the range estimated by Alan and Browning (2003).

---

[39] The coefficient of risk aversion is 8.8 and all other parameters are at baseline values.

[40] Haliassos and Michaelides (2003) use a risk aversion parameter of 8 and Cocco, Gomes and Maenhout (2005) use 10.

5 Conclusions

In this paper we analyze household saving decisions in the presence of two distinct saving tools and public pension insurance. We are able to match the life-cycle averages of wealth and asset allocations, while keeping to a fairly simple framework with realistic parameter values. It is important to stress that the literature on household portfolios has tended to ignore the matching of wealth, and has focused on matching one asset allocation statistic at a time (participation or shares, but not both). Household models of portfolio choice are very sensitive to the level of wealth, and if we wish to have a coherent explanation of asset allocations we must ensure realistic wealth accumulation.

We demonstrate that a standard portfolio choice life-cycle model, calibrated to match the level of wealth of US households, gives extreme rates of stock market participation and very high portfolio shares. This is at odds with empirical findings, and is the basis for many of the puzzles highlighted in the literature. Introducing per period fixed costs enables us to calibrate the model to match the observed limited participation at actual wealth levels. These costs generate the desired hump-shaped age profile for participation, and help resolve the issue of turnover. We present a stylized pay-as-you-go pension scheme, and show that reasonable pension provision is required to keep the asset accumulation and participation age profiles in line with data.

Our framework allows us to calibrate the model to match both participation and shares, in contrast to previous work which asserted that there must exist a trade-off between matching these two statistics. However, this result comes at some cost, requiring us to accept either high wealth accumulation or high parameter values. Adapting the subjective probability distribution of returns to allow for a disastrous event in the stock market has a considerable effect on our results. We are now able to match all three statistics at reasonable parameter values, without resorting to heterogeneous preferences or moving away from power utility functions. Through our three extensions, we are able to resolve both the "portfolio participation puzzle" and the "portfolio composition puzzle" with a

per period fixed cost of 0.8% of annual labor income, risk aversion of 8.8 and discount rate of 9%.

Appendix A: Solution and simulation methods

The results presented in section 4 use standard techniques to solve a life-cycle model by backwards induction. We start from a terminal condition on assets, and obtain the optimal policy functions for each age, which map the state variables into the controls. Using these functions the model is simulated forward from $t = 1$ with 10,000 ex-ante homogenous individuals who differ in shock realizations. Following Alan (2006), initial assets are drawn from a lognormal distribution, whose underlying normal distribution has a mean of -3.15 and standard deviation of 1.96. Mean life-cycle profiles of accumulated wealth and asset allocations are then computed.

Solving the Euler equations corresponds to the determination of a fixed point within an infinite dimension state space, involving expectations over a non-linear marginal utility function, where the unknown is a function over a continuum of points. Such complexity means that the model cannot be solved analytically, which entails the implementation of numerical techniques. The state space is discretized into a finite number of nodes and interpolated using local approximation methods.[41]

The grids are defined so as to avoid the need for extrapolation outside the grid.[42] The concavity of the consumption function leads us to choose a non-uniform spacing of the asset nodes. Extra points are positioned close to the lower bound, where the consumption policy function displays a significant amount of curvature. The nodes are more spread out at high asset levels, at which point the functions become approximately linear. The savings grid is also non-uniformly spaced as the portfolio share policy function is non-

---

[41]Four hundred points are used for both the asset and savings grids. Linear splines are used for interpolation.

[42]Extrapolation is much less reliable that interpolation, especially where the policy functions are non-linear.

linear. More points are positioned around the kink in the policy function, where the short sales constraint ceases to bind, and fewer nodes at high levels of savings, where the policy function becomes almost horizontal approaching the complete markets outcome. The solution is found using NAG routines,[43] except for when these methods fail to converge, in which case the non-linear system is solved using a bisection method.[44]

We perform all numerical integration using Gaussian quadrature to approximate the distributions of labor income and the risky asset. The income shocks are discretized into six values and the risky return uses three point quadrature. In the simulations, the permanent shock to labor income is approximated as a continuous random variable. Each time period is taken as a year of life. T is taken to be 58 years and K as 15 years, giving a working life of 43 years (from age 22 to 65) and life coming to an end at 80.

A check on the accuracy of the solution method is undertaken by computing the realized values of the Euler equations. When averaged across individuals, these results do not deviate substantially from their expected values.

Appendix B: Portfolio share policy functions

Figure B1 shows the portfolio share policy functions for three different ages, under parameters outlined in table 1. At each age the policy function follows a highly nonlinear pattern. At low levels of saving, the short sales constraint is binding and agents hold all of their wealth in risky assets; represented by the horizontal section of the function. This derives from non-tradable labor (or pension) income acting as a substitute for the risk free asset.[45] At low levels of saving the agent is highly endowed with this implicit risk free asset, driving portfolio shares to 100%. As savings are increased, income becomes a relatively less important fraction of wealth, giving low implicit risk free asset holdings and

---

[43]Fortran code is available from the author on request.

[44]Bisection is an iterative procedure that computes the root of a one-dimensional function on a bounded interval of the real line. It is one of the most robust procedures but it converges slowly, hence, it is only used when the NAG routine fails.

[45]Income is risky but bounded below, and hence reduces effective risk tolerance.
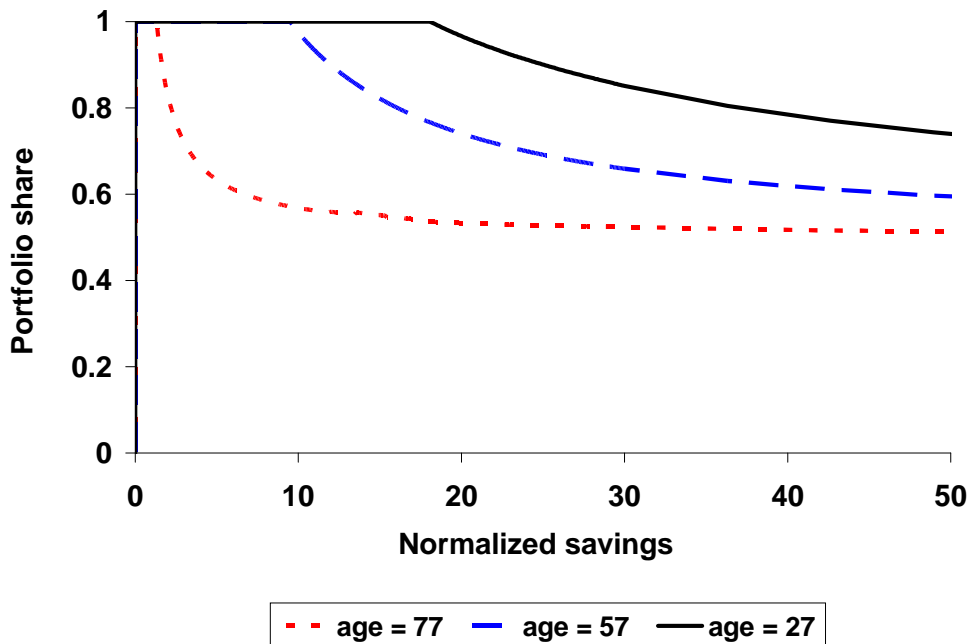
Figure B1: Portfolio share policy functions for three different ages, under baseline parameters.

allowing more diversification. At very high levels of saving, income becomes an insignificant determinant of wealth, and portfolio shares tend to the complete market solution, given by $\omega = \frac{\mu}{\gamma\sigma^2}$.[46] This depicts what has been termed the "portfolio specialization puzzle", whereby only wealthy agents are able to diversify away from complete specialization in risky assets.

As agents age, the policy functions in figure B1 shift inwards, enabling more portfolio diversification for a given level of saving. The young have a high present discounted value of lifetime income, and this represents a large implicit holding of riskless assets, resulting in high portfolio shares. As they grow older, the proportion of total lifetime wealth accounted for by the present discount value of income declines, leading to lower implicit holdings of riskless assets and a tilting of portfolios away from equities.
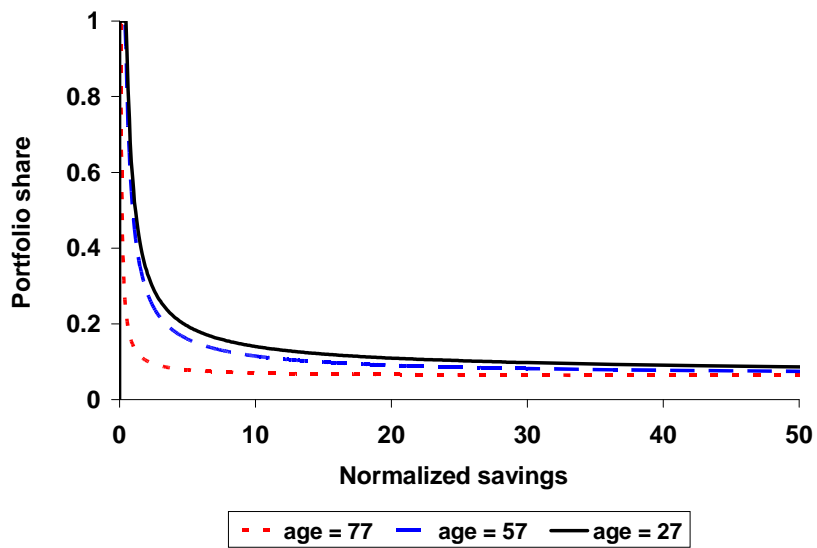
The policy functions are sensitive to the degree of risk aversion. In the low risk aversion case in figure B1, very high levels of saving are required before individuals operate on the

---

[46]Under the baseline parameter values the Samuelson rule gives a complete markets portfolio share of 0.5.
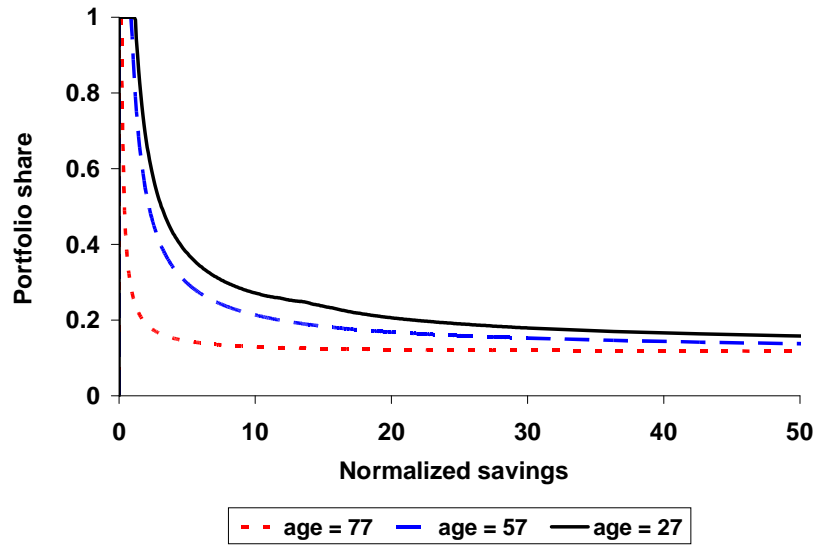
Figure B2: The effect of varying risk aversion and the risk of stocks on the portfolio share policy functions. Panel (a) shows the policy functions with high risk aversion. Panel (b) shows the policy functions with high risk aversion and a small probability of a stock market crash.

nonlinear sections of their policy functions, and move away from complete specialization in stocks. At high levels of risk aversion the policy functions shift in towards the origin, as shown in panel (a) of figure B2, which plots the policy functions for a risk aversion parameter of 8.8. As agents become more risk averse, the range over which the short-sales constraint binds becomes smaller, leading to less specialization in stocks. In addition, the increase in risk aversion reduces the lower asymptote of the policy functions, as the complete markets share falls with γ, enabling more diversification of portfolios. Increasing the risk of investing in stocks, through introducing a small probability of a stock market crash, shifts the policy functions further in towards the origin. This can be seen in panel (b) of figure B2, which depicts the policy function under the same parameters are panel (a), but incorporating a zero return shock.

References

Abowd J, and Card D, 1989, "On the covariance structure of earnings and hours changes," *Econometrica*, 57, 411-445

Alan S, 2006, "Entry costs and stock market participation over the life cycle," *Review of Economic Dynamics*, 588-611

Alan S, and Ball S, 2007, "Household portfolio choice and expected stock market returns," *mimeo*, University of Cambridge

Alan S, and Browning M, 2003, "Estimating intertemporal allocation parameters using simulated residual estimation," *CAM Working Paper* 2003-03

Ameriks J, and Zeldes S, 2001, "How do household portfolio shares vary with age?" *Working Paper, Columbia Business School*

Attanasio O, Banks J, Meghir C, and Weber G, 1999, "Humps and bumps in consumption," *Journal of Business and Economic Statistics*, 17, 22-35

Attanasio O, Banks J, and Tanner A, 2002, "Asset holdings and consumption volatility," *Journal of Political Economy*, 110, 771–792

Attanasio O, and Weber G, 1995, "Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey," *Journal of Political Economy*, 103, 1121-1157

Banks J, and Tanner S, 2002, "Household portfolios in the UK," in Guiso L, Haliassos M and Jappelli T, *Household Portfolios*, MIT Press

Barro R J, 2006, "Rare disasters and asset markets in the twentieth century," *Quarterly Journal of Economics*, 121(2), 823-66

Bernheim B, 1987, "Dissaving after retirement: testing the pure life-cycle hypothesis," in Ashenfelter O, and Card D, *Issues in Pension Economics*, North-Holland

Bertaut C, 1997, "Stockholding behavior of US households: evidence from the 1983-1989 survey of consumer finances," *The review of Economics and Statistics*, 80(2), 263-275

Bertaut C, and Haliassos M, 1997, "Precautionary portfolio behavior from a life-cycle perspective," *Journal of Economic Dynamics and Control*, 21, 1511-42

Bertaut C, and Starr-McCluer M, 2002, "Household portfolios in the United States," in Guiso L, Haliassos M and Jappelli T, *Household Portfolios*, MIT Press

Blundell R, Browning M, and Meghir C, 1994, "Consumer demand and the life-cycle allocation of household expenditure," *Review of Economic Studies*, 61, 57-80

Cagetti M, 2003, "Wealth accumulation over the life cycle and precautionary savings," *Journal of Business Economics and Statistics*, 21 (3), 339–353

Campbell J, Cocco J, Gomes F, and Maenhout P, 2001, "Investing retirement wealth: a life cycle model," in Campbell J, and Feldstein M, *Risk aspects of social security reform*, University of Chicago Press

Carroll C D, 2006, "Lecture notes on solution methods for microeconomic dynamic stochastic optimization problems," *mimeo*, Johns Hopkins University

Carroll C D, 1997, "Buffer-stock saving and the life-cycle/permanent income hypothesis," *Quarterly Journal of Economics*, 112, 1-56

Carroll C D, 1992, "The buffer-stock theory of savings: some macroeconomic evidence," *Brooking Papers on Economic Activity*, 2, 61-156

Cocco J, 2000, "Portfolio choice in the presence of housing," *Working Paper, London Business School*

Cocco J, Gomes, F, and Maenhout P, 2005, "Consumption and portfolio choice over the life cycle," *Review of Financial Studies*, 18(2), 491-533

Davis S, and Willen P, 2000, "Occupational level income shocks and asset returns: their covariance and implications for portfolio choice," *NBER Working Papers 7905*, National Bureau of Economic Research

Deaton A, 1991, "Saving and liquidity constraints," *Econometrica*, 59, 1121-42

Dynan K, Skinner J, and Zeldes S, 2002, "The importance of bequests and life-cycle saving in capital accumulation: a new answer," *American Economic Review*, 92, 274-278

Gabaix C, 2007, "A simple, unified, exactly solved framework for ten puzzles in macro finance," *MIT Working Paper*

Gakidis H, 1998, "Stocks for the old? Earnings uncertainty and life cycle portfolio choice," *PhD dissertation*, MIT

Gomes F, and Michaelides A, 2005, "Optimal life-cycle asset allocation: understanding the empirical evidence," *Journal of Finance*, 60(2), 869-904

Gomes F, and Michaelides A, 2003, "Portfolio choice with internal habit formation: a life-cycle model with uninsurable income risk," *Review of Economic Dynamics*, 6, 729-766

Gourinchas P-O, and Parker J, 2002, "Consumption over the life-cycle," *Econometrica*, 70, 47-89

Gourio F, 2008, "Disasters and recoveries," *American Economic Review* (forthcoming)

Guiso L, Haliassos M and, Jappelli T, 2002, *Household Portfolios*, MIT Press

Guvenen F, 2007,"Learning your earning: are labor income shocks really very persistent" *American Economic Review,* 97(3), 687-712

Haliassos M, and Michaelides A, 2003, "Portfolio choice and liquidity constraints," *International Economic Review*, 44(1), 143-177

Haliassos M, and Michaelides A, 2002, "Calibration and computation of household portfolio models," in Guiso L, Haliassos M and Jappelli T, *Household Portfolios*, MIT Press

Heaton J, and Lucas D, 2000, "Portfolio choice in the presence of background risk," *Economic Journal*, 109, 1–26

Heaton J, and Lucas D, 1996, "Market frictions, savings behavior, and portfolio choice," *Macroeconomic Dynamics*, 104, 443-487

Hubbard R, Skinner J, and Zeldes S P, 1995, "Precautionary saving and social insurance," *Journal of Political Economy*, 103(2), 360-399

Julliard C, and Ghosh A, 2008, "Can rare events explain the equity premium puzzle," *mimeo*, London School of Economics

Low H W, 2005, "Self-insurance in a life-cycle model of labor supply and savings," *Review of Economic Dynamics*, 8, 945-975

MaCurdy T, 1982, "The use of time series processes to model the error structure of earnings in a longitudinal data analysis," *Journal of Econometrics*, 18, 83-114

Mandelbrot B B, 2001, "Scaling in financial prices: Tails and dependence," *Quantitative Finance*, 1, 113–124

Mandelbrot B B, 1997, *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*, Selecta Volume E, Springer-Verlag

Mandelbrot B B, 1967, "The variation of some other speculative prices," *Journal of Business*, 40, 393-413

Mandelbrot B B, 1963, "The variation of certain speculative prices," *Journal of Business*, 36, 394-419

Meghir C, and Pistaferri L, 2004, "Income variance dynamics and heterogeneity," *Econometrica*, 72, 1-32

Mehra R, and Prescot E C, 1985, "The equity premium: a puzzle," *Journal of Monetary Economics*, 15(2), 145-61

Moffitt R, and Gottschalk P, 2002, "Trends in the transitory variance of earnings in the United States," *Economic Journal*, 112(478), C68-C73

Paiella M, 2001, "Transaction costs and limited stock market participation to reconcile asset prices and consumption choices," *Institute for Fiscal Studies Working Paper*, WP 01/06

Poterba J, and Samwick A, 1997, "Household portfolio allocations over the life-cycle," *NBER Working Papers 6185*, National Bureau of Economic Research

Rietz T A, 1998, "The equity risk premium - a solution," *Journal of Monetary Economics*, 22(1), 117-31

Rotemberg J, and Woodford J, 1998, "An optimization-based econometric framework for the evaluation of monetary policy: expanded version," *NBER Technical Working Papers 233*, National Bureau of Economic Research

Samuelson P A, 1969, "Lifetime portfolio selection by dynamic stochastic programing," *Review of Economics and Statistics*, 51, 239-46

Veronesi P, 2004, "The peso problem hypothesis and stock market returns," *Journal of Economic Dynamics and Control*, 28(4), 707-25

Vissing-Jorgensen A, 2004, "Perspectives on Behavioral finance: does 'irrationality' disappear with wealth? Evidence from expectations and action," *NBER Macro Annual 2003*, Cambridge MA, MIT Press

Vissing-Jorgensen A, 2002, "Towards an explanation of household portfolio choice heterogeneity: nonfinancial income and participation cost structures," *NBER Working Papers 8884*, National Bureau of Economic Research

Weitzman M, 2007, "Subjective expectations and asset-return puzzles," *American Economic Review*, 97(4), 1102-1130

Yao R, and Zhang H, 2005, "Optimal consumption and portfolio choices with risky housing and stochastic labor income," *Review of Financial Studies*, 18, 197-239