Limited stock market participation, pensions and consumption over the life-cycle*

Steffan Ball†

University of Cambridge

August 2006

Abstract

In this paper we present a household model with realistically calibrated parameters, which is able to simultaneously match consumption, asset accumulation and stock market participation profiles over the life-cycle. The inclusion of fixed costs and a public pension scheme eradicates the need to assume heterogeneity in preferences or extreme parameter values in order to explain observed patterns. We find a fixed cost of just over 4 percent of permanent income can explain the limited stock market participation, with a proportional pension tax rate of twenty percent. More generous public pensions are seen to crowd out private savings and significantly influence the estimates of fixed costs.

Keywords: life-cycle models, precautionary saving, portfolio choice, stock market participation and uninsurable labour income risk

JEL Classification: G11, H31

*I would like to thank Hamish Low, Pramila Krishnan, Michael Ball, Dalibor Eterovic and participants at the Cambirdge University Macroeconomic Worshop for helpful comments and suggestions.

†Corresponding address at: Faculty of Economics, University of Cambridge, Sidgwick Avenue, Cambridge, CB3 9DD. *Email*: sgb24@econ.cam.ac.uk

1 Introduction

In this paper we present a life-cycle portfolio choice model, with realistically calibrated stochastic labour income and low risk aversion, in order to explain some of the stylised facts of household asset allocations. There is now a wide literature of household models with precautionary and retirement motives for saving, which are able to explain some of the observed patterns of wealth accumulation and consumption (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 (Cocco, Gomes and Maenhout, 1999; Cocco, 2000; Gomes and Michaelides, 2003). However, these models still have difficulties in explaining some of the main empirical evidence on household portfolios without assuming extreme parameter values or heterogeneity in preferences.

The divergence of the theoretical predictions of asset allocation models from observed facts has led to much discussion of "puzzles" within the finance literature. These puzzles arise either from the predicted participation decisions or from the optimal share of wealth held in risky assets diverging from the results of empirical studies. Given a premium on equities over risk-free saving tools, many models of optimal portfolio choice predict that all households with positive wealth holdings should participate in the stock market. However, stock market participation is limited to less than half of the population in the US, just below a third of UK households and to less than a fifth of the Italian population (Guiso, 2002), a key feature which has been difficult to reconcile with theory. With respect to portfolio composition, the theory predicts that average portfolio shares should be high and only wealthy households should diversify their portfolio away from full specialisation in risky assets. However, the data does not seem to support this with low average portfolio shares, and small savers holding mostly riskless assets. Our model aims to further the understanding of the mechanisms that underlie the limited stock market participation and relatively low average equity shares observed empirically by introducing fixed costs to stock market participation and some public pension provision.

It is clear that information plays a key role in investment decisions. Gathering and processing information is costly. These costs may be the material cost of paying for the assistance of a financial advisor, or time costs of gathering and processing the information. Information costs directly affect the cost of saving, by determining access to and utilisation of different investment technologies. This has important implications of households' ability to accumulate wealth and self-insure against income shocks. It is important to understand whether such information requirements constitute binding constraints for many households and how they affect choices over different assets. Fixed costs have been shown to be important in order to get realistic predictions out of household models of portfolio choice (Alan, 2005; Gomes and Michaelides, 2005; Haliassos and Michaelides, 2003). Empirical work by Paiella (2001) and Vissing-Jorgensen (2002) have shown that fixed costs can rationalise the low levels of stock market participation. This work has focused on incorporating fixed costs of stock market entry into a life-cycle framework but little has been advanced incorporating per period fixed costs and calibrating their size to match observed participation rates.

Many of the models used to explain asset allocation behaviour tend to impose an exogenous retirement function or an exogenous replacement rate without looking at the effect of collecting the tax revenues required to fund such pension schemes. Taxation has a direct effect on the cost of saving and this has important implications for optimal participation rates and portfolio shares.

The effect of changing public pensions on optimal portfolio allocations and on the size of the fixed cost necessary to explain the observed participation has not been analysed in the literature to date. To the best of my knowledge, this paper is the first study to quantify per period fixed costs by accounting for limited participation within a fully structural framework with public insurance. The model presented below is able to match the age profiles of consumption, asset accumulation and stock market participation in the US without resorting to unrealistically high risk aversion, preference heterogeneity or non-homotheticity. The results for portfolio shares are a little more ambiguous, hindered somewhat by the lack of empirical consensus on the age and wealth profile for shares; as they remain implausibly high for young individuals[1], but the decrease in risky asset holdings towards the latter years of life does seem to correspond with descriptive studies of portfolio choice.

In the next section we provide some data on consumption, asset holdings, participation and portfolio shares by households in the US. Section 3 develops our household model of portfolio choice. Section 4 provides simulated life-cycle profiles of consumption, asset accumulation, stock market entry and risky assets shares showing the effect of introducing per period fixed costs and public pension provision. Section 6 concludes.

2 Stylised facts on life-cycle profiles and portfolio allocations

The aim of this paper is to attempt to explain consumption and saving patterns. The model detailed in the next section can be thought of as an extension to the recent literature on buffer-stock saving (Deaton, 1991, Gourinchas and Parker,

[1]Which may be the result of background risk not being fully captured by the i.i.d. income risk.

2002) to include an asset allocation decision. This means that we do not simply want to match the consumption and asset accumulation over the life cycle but also match the stock market participation and portfolio share over the different ages of life. This section details data on some of the life-cycle patterns that the model addresses.

A hump shaped consumption profile is well documented in the literate (see Low (2005) for a recent example using US data) where consumption is seen to increase in the early years of working life, reaching a peak at around 50 years of age and subsequently declining into retirement. Average asset holdings are seen to increase gradually over the working life and then be run down during retirement (Low, 2005; Burnheim, 1987).

When addressing household investment in stocks it is important to distinguish between participation and portfolio shares. Stock market participation in the US is currently just under fifty percent[2], which is well below the prediction of many asset allocation models where participation rates of 100% are conjectured. This leads to questions as to why participation is so limited. Furthermore, there is a clear age effect on participation and there has been much discussion of a hump shape in the decision of whether or not to hold risky assets (Alan, 2005; Ameriks and Zeldes, 2001; Poterba and Samwick, 1999; Gomes and Michaelides, 2005). Figure 1 below uses data from the US Survey of Consumer Finances to show how stock market participation follows such a hump shape.

[2] Guiso et al. (2002) report a figure of 48.9%

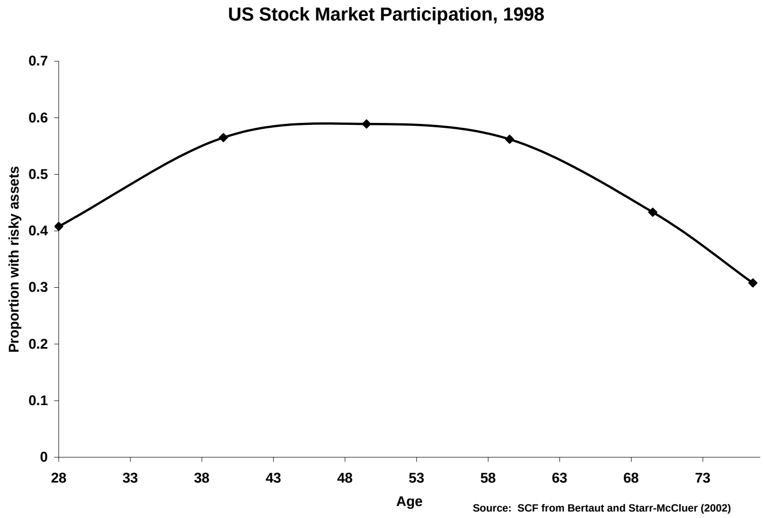

Figure 1

Participation is also increasing in wealth with richer households much more likely to hold risky assets than poorer household. Using multivariate probit techniques for the US, taking data from the Survey of Consumer Finances in 1995, Bertaut and Starr-McCluer (2002) find that both age and wealth effects are significant for the participation decision.

The second important stylised fact of asset allocations is that portfolio shares are typically well below 100% (Ameriks and Zeldes, 2001; Bertaut and Starr-McCluer, 2002; Gomes and Michaelides, 2005; and others). In the US, for example, the average portfolio share is around 60%[3]. The age and wealth effects on portfolio shares are much weaker than those on participation. It is common practice for financial advisors to tell their clients to shift their portfolios away from risky assets as they age, especially as they enter retirement, and so it seems reasonable to expect portfolio shares to fall with age in the

[3] Bertaut and Starr-McCluer (2002) report an average portfolio share of 60.5%. Gomes and Michaelides (2005) find average shares of 55%.

latter part of life. Bertaut and Starr-McCluer (2002) show some evidence to this effect in their descriptive statistics whereby portfolio shares, conditional on participation, fall from just over 60% for households aged 55-64 to around 55% for the retired. However, these results are not robust to more sophisticated econometric analysis. In their multivariate probit analysis Bertaut and Starr-McCluer (2002) find that neither wealth nor age effects were significant for portfolio share. Further, it is also commonly reported that small savers tend to hold all of their wealth in risk free assets (for example, Haliassos and Michaelides, 2002).

The prevailing evidence brings Gomes and Michaelides (2005) to conclude that the empirics strongly suggest a hump-shaped age participation profile but that the age profile of portfolio shares has no clear cut pattern. Therefore, we would like our model of household behaviour to mimic the average patterns of stock market participation as well as consumption and savings; with some degree of ambivalence towards portfolio share profiles.

3 Life-cycle model

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

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

subject to,

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

and a terminal condition, $A_{T+1} = 0$. 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; Y_t is exogenous real disposable labour income; F is the per period fixed costs of stock market participation; ω_t is the portfolio share; r and r_t^{eq} are the risk free and risky returns respectively; and $I(.)$ is an indicator function taking a value of unity when the individual participates in the stock market, and zero otherwise.

In each period t the agent has to decide how much to consume and how much to save out of cash-on-hand, which is comprised of assets held at the start of the period and realised income $(A_t + Y_t)$. Any savings, given by $(A_t + Y_t - C_t)$, are invested and give a composite portfolio return denoted by the terms in square brackets in equation (2) above. We assume the portfolio investment is comprised of two distinct saving tools, a riskless and a risky asset. The riskless asset has a constant real return of r and the risky asset has a stochastic real return of r_t^{eq}, which is assumed to be i.i.d. over time and distributed lognormally. We do not impose any correlation between stocks and labour 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 (2001) even 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. Ignoring such a correlation is likely to bias my estimates of fixed costs downwards as risky assets are more desirable when independent of labour income than if some positive correlation existed, which makes the high fixed costs that result from the model even more stark.

The portfolio share, ω_t, represents the proportion of assets held in risky forms. In order to invest in risky assets in any given period the individual has to pay a fixed cost of F. These per period fixed costs can be thought of as representing brokerage/membership fees. Alternatively, the F represents time costs of information gathering and corresponds to the opportunity costs of acquiring data on financial markets, monitoring brokers, keeping up to date with trends etc.. F is assumed to be a proportion of permanent labour income following the motivation of the fixed cost being an opportunity cost of time. Also, such a specification significantly simplifies the solution to the model as it removes the necessity to have an additional state variable.

Some limited borrowing is allowed up to 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 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.

During working life, income is uncertain with no insurance possibilities other than self-insurance. Following the standard specification in the literature,[4] we shall assume that the stochastic labour income process can be broken down into a deterministic component, which can be calibrated to match the hump shape of income over the life-cycle; and two random components, a permanent and a transitory shock:

$$Y_{t+1} = (1-\tau)G_{t+1}P_{t+1}U_{t+1} \tag{3}$$

where U_t is transitory income, P_t is permanent income, τ is a proportional

[4] See, for example, Carroll (1992), Low(2005), Alan (2005).

pension-tax and G_t is a deterministic growth trend. The log of permanent income is assumed to follow a random walk.

$$P_{t+1} = P_t N_{t+1} \tag{4}$$

Both transitory and permanent shocks are assumed to be independently and identically distributed lognormally such that $\ln(U_t) \sim N(-0.5\sigma_u^2, \sigma_u^2)$ and $\ln(N_t) \sim N(-0.5\sigma_n^2, \sigma_n^2)$. The income process is augmented such that a zero level of income cannot be realised, otherwise individuals would optimally choose never to borrow.

For the final R periods of life individuals are retired. The proportional tax rate is set exogenously and the tax revenue is used to fund retirement as in a pay-as-you-go pension scheme. For any given period there is a continuum of individuals at each age, those of working age are paying tax on their labour income which is being used to fund the retired. This gives the following certain pension income:

$$ssi = \frac{\tau . E_t \left[\sum_{t=1}^{T-R} Y_t \right]}{R} \tag{5}$$

The model detailed above provides three motives for saving; precautionary, retirement and patience. These motives are non-additive in the sense that savings accumulated for one specific purpose are not excluded for alternative uses, should there be a change in the state of nature. For example, savings held for retirement can be used to buffer a negative unexpected shock to earnings, and precautionary balances can be consumed in retirement should realisations of income prove better than expected. Such fungibility of assets lies behind Dynan, Skinner and Zeldes (2002) argument that bequest motives and precautionary motives are not distinct. For this reason we do not include any specific bequest motive in the model under the assumption that it is subsumed in the other motives for saving.

The first order condition of the value function with respect to consumption is given by:

$$u^{'}(C_t) = \beta E_t \left\{ [\omega_t(1 + r_t^{eq}) + (1 - \omega_t)(1 + r)]u^{'}(C_{t+1}) \right\} \tag{6}$$

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

$$0 = E_t \left[(r - r_t^{eq})u^{'}(C_{t+1}) \right] \tag{7}$$

The concavity of the value function ensures that the solution to these first-order conditions yields a maximum. These two conditions can be solved for the following policy functions in each period

$$C_t(A_t, Y_t, P_t) \tag{8}$$

$$\omega_t(A_t, Y_t, P_t) \tag{9}$$

This is a complicated multidimensional problem which can be simplified into a sequence of two relatively simpler optimisation problems, following Carroll (2006). First, the number of state variables can be reduced by exploiting the scale-independence of the maximisation problem and redefining all variables in terms of ratios to permanent income. Second, the asset evolution equation and the value function can be written in terms of ratios to permanent income, and lowercase letters refer to the normalised counterpart of uppercase variables. Then at each period t, a grid for savings is defined and the first order condition in (7) is solved for portfolio share policy function as a function of savings:

$$\omega_t(s_t) \tag{10}$$

This portfolio share policy function is subsequently used to solve (6) at each t for the following consumption function:

$$c_t(a_t + y_t) \tag{11}$$

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 the appendix.

4 Simulation results

Once the optimal policy functions have been determined we can simulate the model to imitate the behaviour 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 shall detail the functional forms, parameter values and calibration statistics used. Then we shall show the results of simulating the model for 10,000 ex-ante identical households who differ in the realisation of the shocks.

We then put forward the results in four subsections: first, we analyse our baseline life-cycle model with portfolio choice in order to highlight the models failure to account for the stylised facts of US households. Second, we introduce a public pension scheme and show how this affects the simulated profiles. Third, we introduce fixed costs and show this modification changes the results of the model. Finally, we show the effects of introducing fixed costs and public pensions. This structure permits us to disentangle the different effects of public insurance and fixed cost considerations in order to better fit the data.

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

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}. \tag{12}$$

with a baseline value of 2 for the coefficient of risk aversion.

The riskless rate of return is taken from Low (2005) at a constant value of 1.6%, which represents the average annualised real rate of US 3-month treasury bills from 1960-2000. The risky return is drawn from a distribution with mean of 5.6%, corresponding to a 4% equity premium[5], and a standard deviation of 0.2[6]. The deterministic trend in the income process is taken directly from Low (2005) with the parameter values emerging from a regression of log wages on a quadratic in age controlling for demographics and cohort effects. This gives the well documented concave income profile over the life-cycle, where income rises from the start of working life to a peak around an age of 50 and then declines to retirement. The values for the variance of permanent and transitory shocks to income are taken from Meghir and Pistaferri (2004)[7].

The values for the discount factor and fixed cost are chosen by calibration. As detailed in the simulations in the paper by Low (2005), we choose the discount rate such that the simulated ratio of median wealth to income matches the PSID for the working households. The value of the fixed cost, where appropriate, is chosen in order to match simulated average participation with the observed data.

[5] As is common in this literature; see Gomes and Michaelides (2005), Haliassos and Michaelides (2002), Yao and Zhang (2005), Cocco (2001), Campbell et al. (2001)

[6] See Alan (2005). Gomes and Michaelides (2005) and Haliassos and Michaelides (2002) use a similar value of 0.18

[7] Giving a value of 0.03 for the variance of permanent shocks and 0.031 for transitory shocks.

4.1 Baseline case with no public pensions or fixed costs

In our baseline model we set the per period fixed costs to zero and consider the case with no public pension provision. Figures 2-5 show the average life-cycle profiles of log consumption, asset holdings, stock market participation and portfolio shares[8].

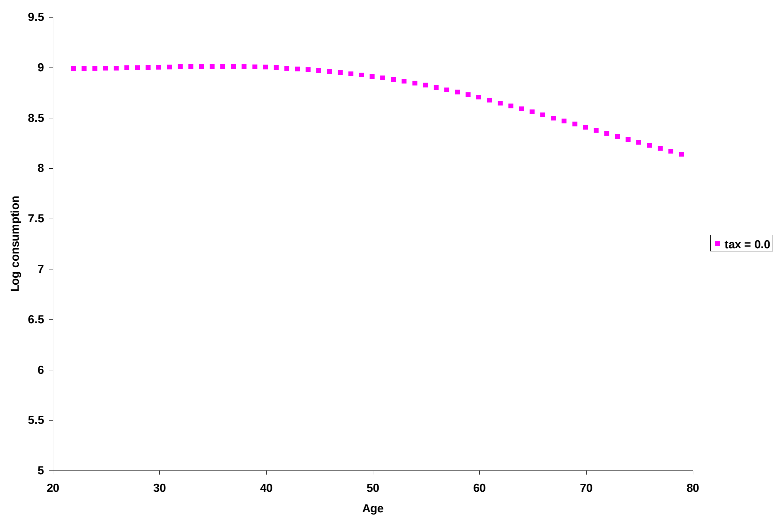

Figure 2

[8] All portfolio shares shown are conditional on participation, when participation is not 100% these conditional portfolio shares will diverge signifanctly from unconditional shares.

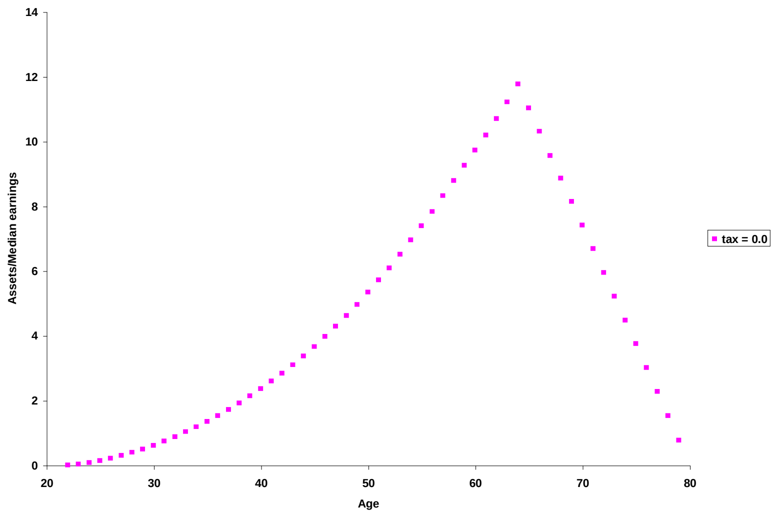

Figure 3

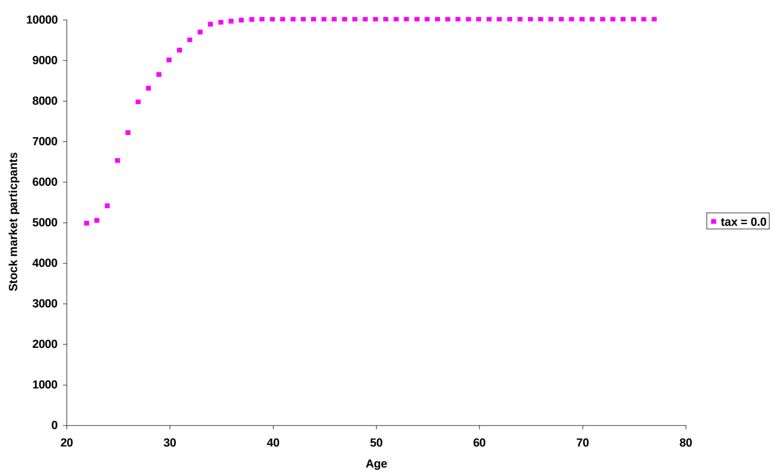

Figure 4

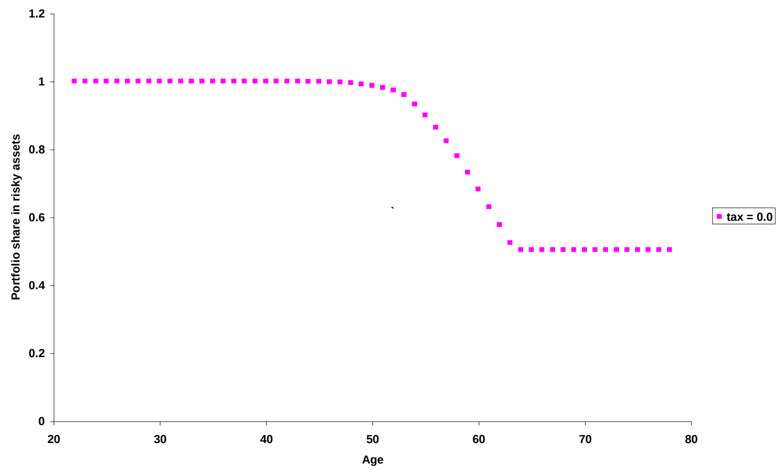

Figure 5

Under these baseline parameters the discount rate required to match the level of assets is 9.98%. Such a high level of impatience is needed to keep savings down to the empirically observed level, since the high returns to risky asset holdings make individuals very willing to save, compare with the calibrated discount rates of around 3% Low (2005) obtains from a similar framework without risky asset considerations. Figure 2 shows the log consumption profile at a comparable level to that of the US and is slightly hump shaped. However, the high level of impatience means that the initial increase is fairly modest and the decline in consumption in later life is considerable. In Figures 3, agents are seen to accumulate assets up until retirement and then run 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. Average assets are always positive, but some individuals do borrow in the early stages of life. Figure 4 shows participation in the stock market is 100% for all but a few years in early life, in fact only individuals with negative savings do not participate. This clearly shows the "portfolio participation puzzle" that the finance literature has discussed at length. This baseline model predicts that all agents with positive savings should invest in the stock market.

Figure 5 details portfolio shares conditional on participation in the stock market. Once agents optimally choose to participate, their portfolio shares are unity for the majority of working life, falling to the complete markets share during retirement. This describes what has been termed the "portfolio specialisation puzzle", whereby only very wealthy agents are able to diversify away from complete specialisation in risky assets. Households with low levels of savings invest all their assets in the risky asset as their consumption path is driven mainly by the stochastic income process rather than the return to savings. As households build up larger savings balances, their consumption paths are determined more by the return to the risky investment and once a certain threshold of wealth is reached it is no longer optimal for households to hold all of their assets in the risky saving tool. In this baseline model it takes thirty years of asset accumulation for individuals to be able to hold both risky and riskless assets in their portfolio.

Once retired there is no source of income, and hence the baseline model predicts the complete markets portfolio share[9]. The average portfolio share of those households that participate in the stock market is 81%. Although this over-predicts the actual share for the US which currently stands at around 60%, it is significantly less than the 100% often reported in models of portfolio choice with labour income risk.

The baseline model fails to match the observed consumption profile and highlights the problems that have been widely discussed when trying to model household asset allocations. In order to attempt to correct for these short-

[9] The seminal work by Samuelson (1969) analysed portfolio choice in a complete markets situation with no background risk. He showed that the portfolio share should be independent of wealth and age and be given by $\omega = \frac{\mu}{\gamma\sigma^2}$, giving a value of 0.5 using our parameter values.

comings we shall start by looking at the effect of introducing a public pensions scheme, and then determine if fixed costs to stock market participation improve the results. Finally we shall look at including both public pensions and fixed costs.

4.2 The effect of introducing a public pension scheme

Figures 6-9 shows the effect of introducing a pay-as-you-go public pension scheme, outlined in Section 3, to the baseline model. A proportional tax rate of 10% and 20% are considered giving a retirement replacement rate of 0.29 and 0.57 respectively.

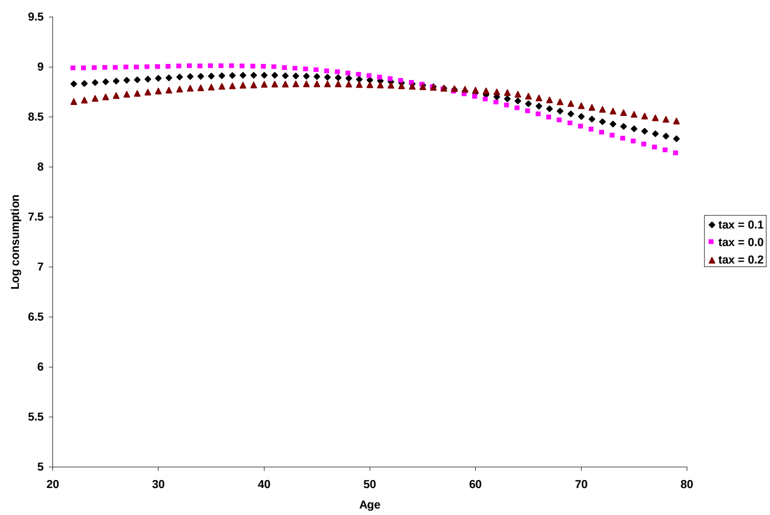

Figure 6

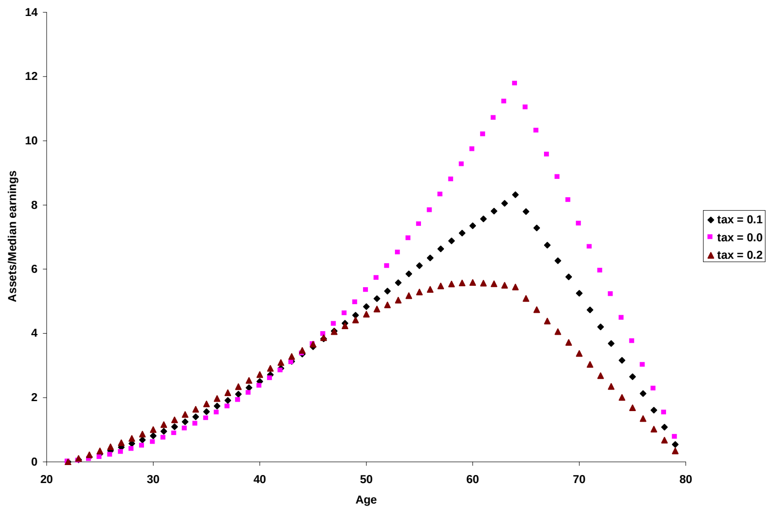

Figure 7

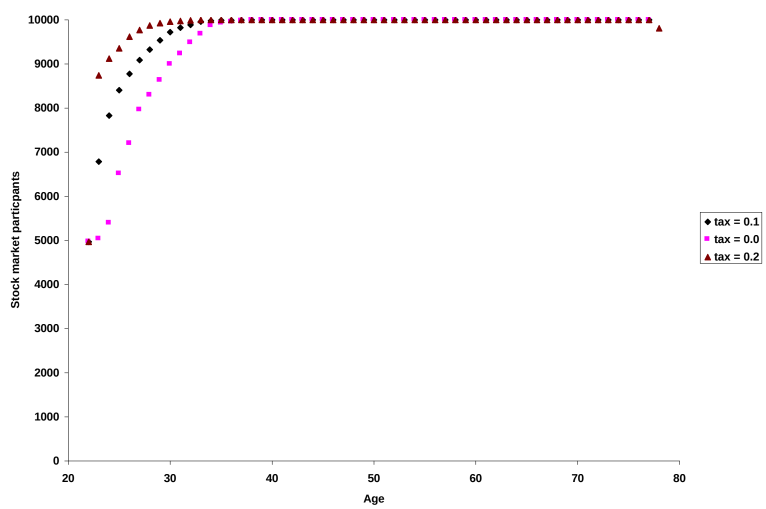

Figure 8

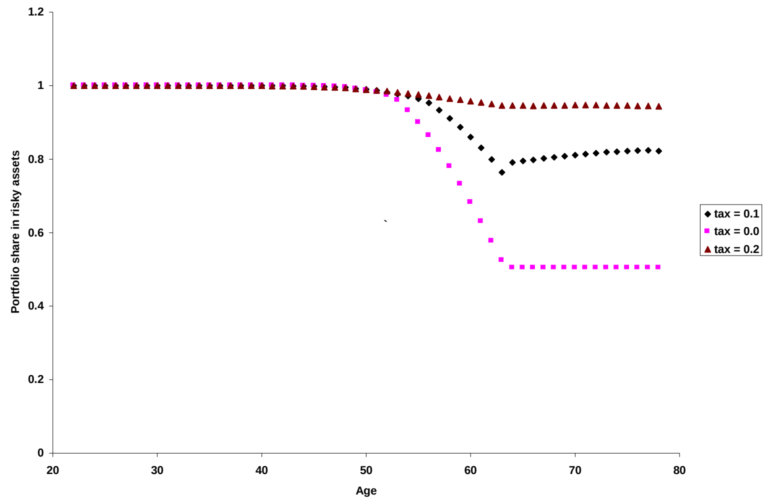

Figure 9

Changing the tax rate directly affects the growth rate of income and hence the effective patience of the individuals. The higher the tax rate the less pronounced the income growth over one's life and so the more impatient the individuals. In order to keep the amount of assets consistent with the US data the calibrated discount rate falls from 9.98% to 8.94% to 7.75% as the tax is increased from zero to 10% to 20%. If we had kept the discount rate the same across the scenarios, assets would be lower under higher tax rates due to a higher effective impatience and the crowding out of private retirement savings by the pension contributions. Ensuring that savings remain the same, leads to a redistribution of savings from mid-age to earlier years in the working life which represents a shift from retirement savings to precautionary savings, which is evident in Figure 7. The change in savings behaviour has a considerable effect on consumption, with the introduction of a state pension scheme making log consumption more hump shaped (Figure 6). Individuals consume less when young (a combination of the direct tax effect and an increase in precautionary balances) and more in retirement. This transfer away from the young in favour of older cohorts, leads to a more realistically shaped consumption profile over the life-cycle.

The shift in savings has a significant effect on portfolio allocations. Because of higher precautionary balances, fewer individuals borrow during the first 15 years of life; and this leads to even greater participation rates, given the model predicts that all agents with positive assets should be participating, as can be seen in Figure 8. Individuals hold lower retirement savings, reducing the peak level of assets, which means that they are not able to diversify away from risky assets to the same degree during mid-age (Figure 9). The introduction of the public pension scheme leads to both higher participation and higher average portfolio shares, and thus independently does not help the model explain observed asset allocation puzzles.

4.3 The effect of introducing calibrated fixed costs

Introducing fixed costs and calibrating the model to match the savings level from the PSID leads to a slightly lower discount rate (from 9.98% to 9.51%)[10]. Introducing the fixed cost effectively makes individuals more impatient and less willing to save, as savings are more costly; and so the discount rate must fall. The model is calibrated to match the average participation rate of 48.9% giving a per period fixed cost of 7.72% of permanent income and the results are detailed in Figures 10-13 below.

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

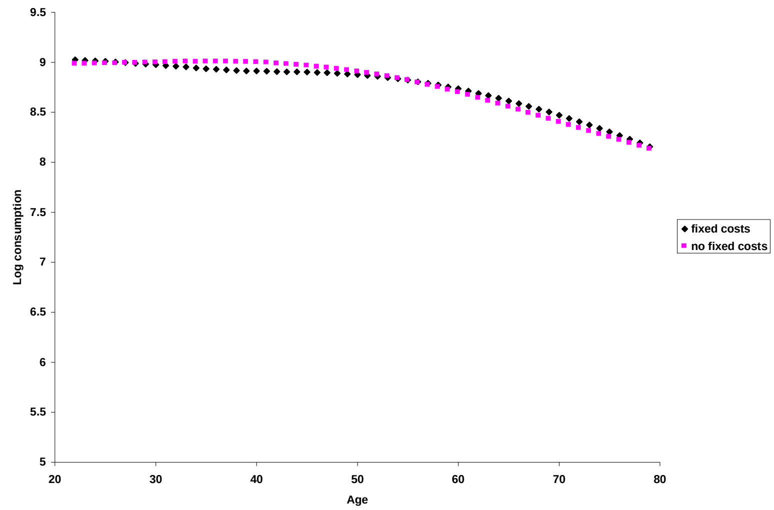

Figure 10

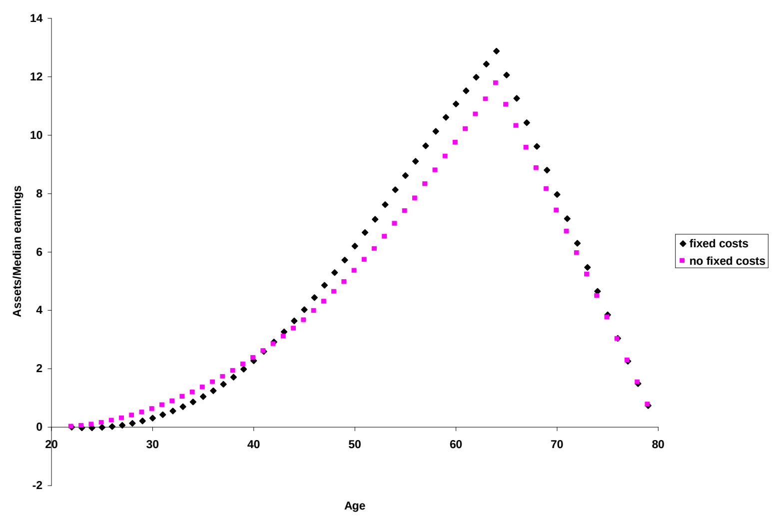

Figure 11

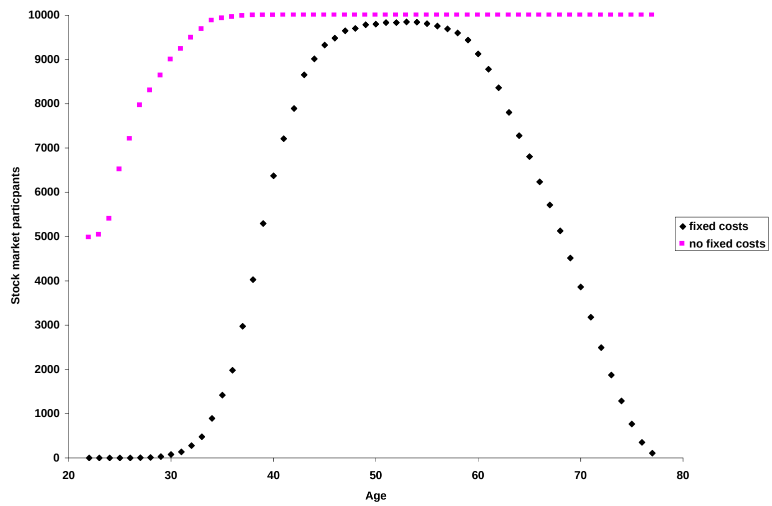

Figure 12

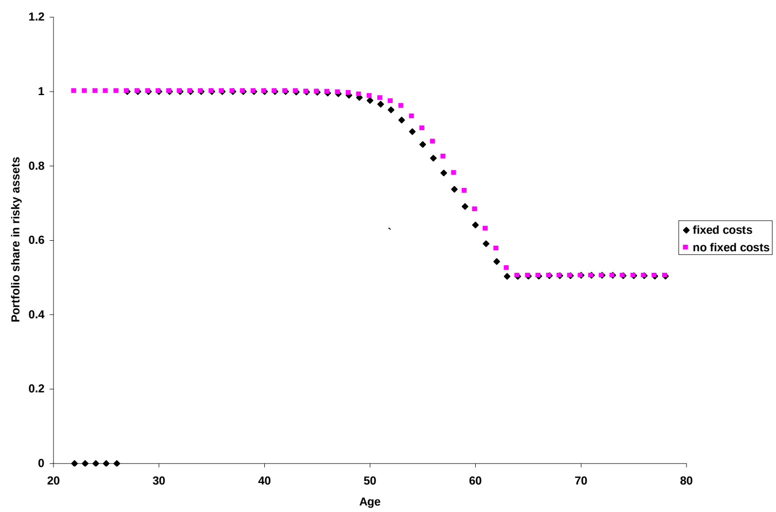

Figure 13

The increased impatience makes individuals borrow early in life rather than save, and so consumption is higher initially in the case with fixed costs. This reduced saving early in life is compensated by more saving for retirement in mid-age, see Figure 11. This leads to a consumption profile in Figure 10 that is hardly distinguishable from the case without fixed costs.

The introduction of per period fixed costs has a marked effect on the stock market participation profile, as a clear hump shape is now evident in Figure 12. Few individuals accumulate sufficient wealth in the first twenty years of working life to make it worth their while 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 to a peak in one's 50s close to 100% participation. As assets are drawn down later in life, the number of individuals participating is seen to fall gradually back to zero. This elucidates the rationale behind small savers holding entirely risk free portfolios, although the hump shape is much more pronounced than the data shown in Figure 1.

The introduction of per period fixed costs is able to explain the observed limited participation and low portfolio shares of small savers, but has little effect on the portfolio shares age profile (see Figure 13) or on average consumption[11].

4.4 The effect of introducing a public pension scheme with calibrated fixed costs

In this section we analyse the effect of changing the public pension scheme generosity in a model where fixed costs are calibrated to match the US average participation in the stock market. We consider two different scenarios; first, we shall consider the effect of varying the tax rate while keeping the level of savings constant, and allowing the discount rate to change in order to achieve this. Second, we shall consider keeping the discount rate constant, and allow savings

[11] This is because household savings behaviour is not significantly altered by the introduction of these relatively small fixed costs.

to vary. We shall increase the tax rate, and by so doing raise the generosity of the state pensions, in order to determine the effect on the calibrated level of fixed costs and on the profiles of consumption, assets accumulation and portfolio allocations.

4.4.1 Keeping savings constant

The discount rate is calibrated to match asset holdings for each of the tax rates. Moving from a zero tax rate to a 10% tax rate to a tax rate of 20% leads to the discount rate to fall from 9.55% to 8.46% to 7.40%. In order to explain the observed limited participation the per period fixed cost remains just below 8% of permanent income for each calibration. The results of these simulations are shown in Figures 14-17 below.

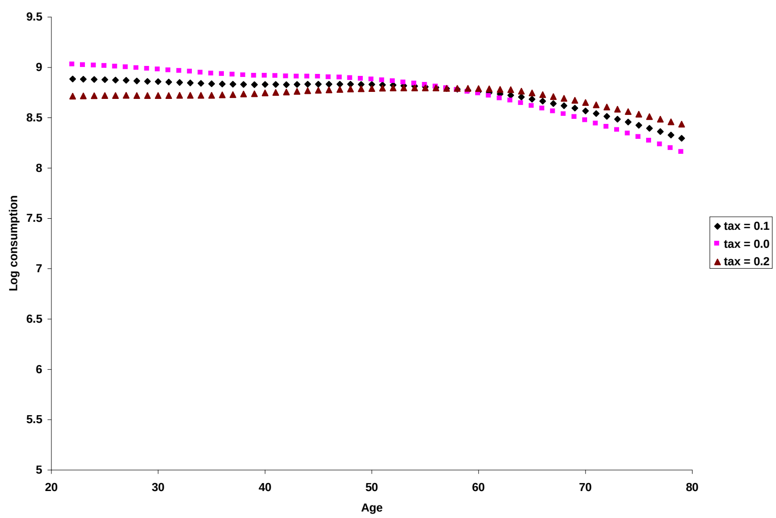

Figure 14

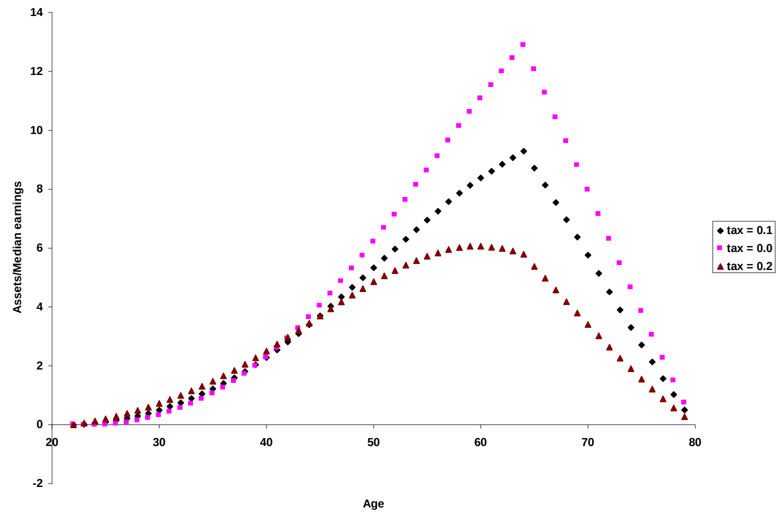

Figure 15

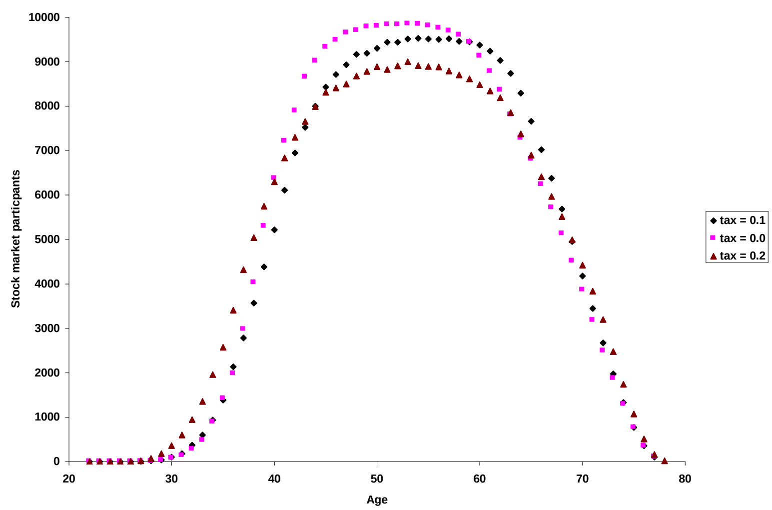

Figure 16

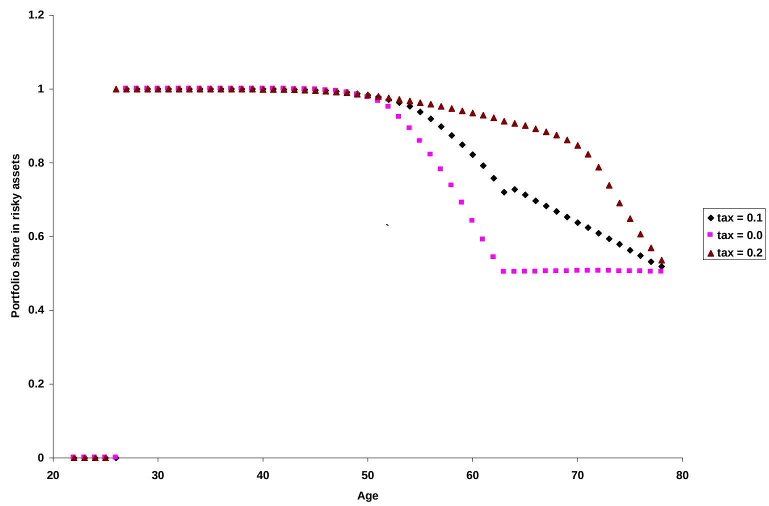

Figure 17

The introduction of a tax has a similar effect on consumption and asset profiles as in the case when fixed costs are assumed to be zero - compare Figures 6 and 7 with Figures 14 and 15. Again precautionary savings increase at the cost of retirement saving, with asset levels in Figure 15 increasing in the first years of working life in response to more generous pension provision and falling in later life. Figure 14 shows consumption being redistributed from the young to the retired and the profile taking on a more pronounced hump shape the higher the tax rate.

The introduction of the public pensions makes the hump shape in stock market participation less stark than in the case with no public provision. More young and retired individuals are seen to participate, and the peak level of participation falls to significantly below 100%. This is much closer to the actual US profile shown in Figure 1.

Figure 17 shows how the public insurance leads to higher portfolio shares, coming from the lower wealth levels during mid-life, but they are still seen to decline into retirement. Small savers still hold portfolios that consist of no risky assets but the age profile of shares does remain uncomfortably high. However, the model with fixed costs and public insurance performs well with respect to explaining the paths of consumption, assets and stock market participation over the life-cycle.

4.4.2 Keeping the discount rate constant

If we keep the discount rate constant while increasing the pensions tax rate, we effectively increase the impatience of individuals, resulting in less saving. Increasing the generosity of the public pensions scheme has a significant effect on the calibrated level of fixed costs falling from 7.73%, when the tax rate is zero; to 6.48%, when the tax rate is 10%; to 4.14%, when the tax rate is 20%. This provides a possible explanation as to why empirical estimates of fixed costs are so low, they are conditional on the generosity of the public insurance scheme. With a fairly well developed public pension scheme, individuals are well catered for in retirement and so have less incentive to save. Given savings are relatively low, only a small fixed cost is needed to explain the given limited participation in the stock market. However, in the situation where there is no public provision for old age, individuals have large asset holdings and are highly motivated to invest in risky investment tools. Thus, high fixed costs are required to keep participation at the observed level.

The profiles for the scenario where the pension scheme is changed keeping the discount rate at the level calibrated under no public insurance are shown in Figures 18-21 below.

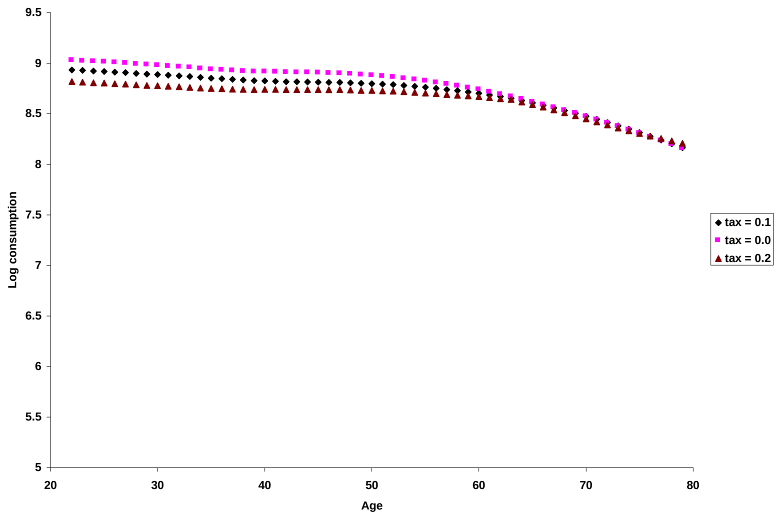

Figure 18

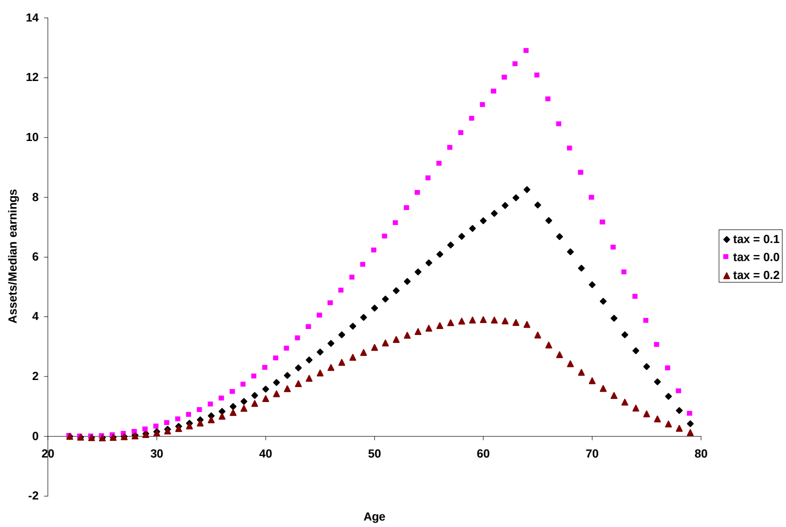

Figure 19

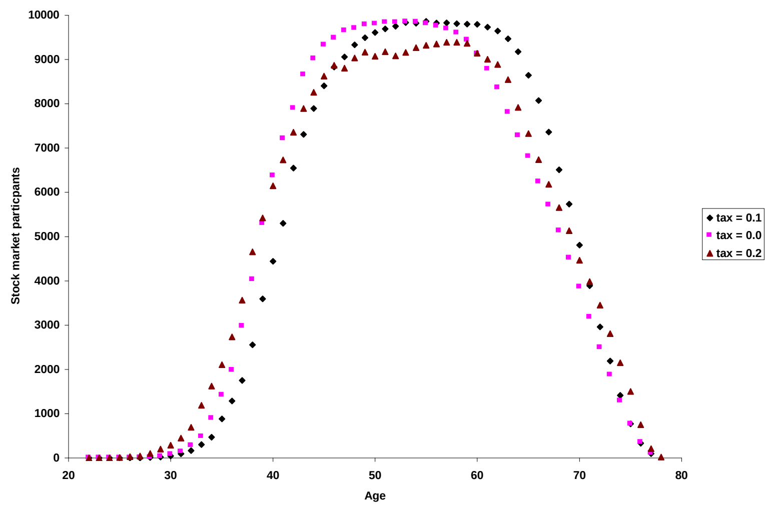

Figure 20

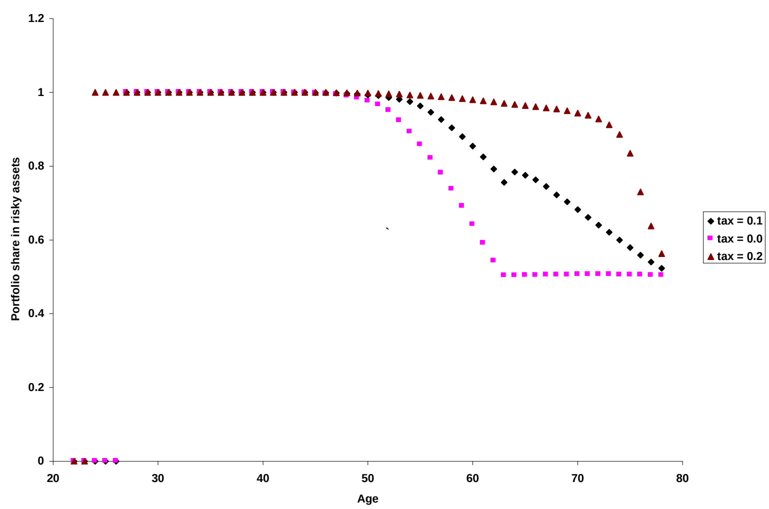

Figure 21

In Figures 18 and 19 the crowding out effect of public pensions is clearly evident. Public provision in old age substitutes for self-insurance, and individuals do not hold large retirement savings. Further, precautionary savings balances are also reduced. The tax structure reduces consumption when young but does not increase consumption for the retired, due to the crowding out, leading to a welfare loss.

5 Conclusion

In this paper we analyse household consumption and saving decisions in the presence of two distinct saving tools and public pension insurance. We are able to match the life-cycle profiles of consumption, savings and stock market participation while keeping to a fairly simple framework with realistic parameter values.

We find that a standard portfolio choice life-cycle model, calibrated to match the level of savings in the PSID, gives downward sloping consumption profiles and extreme rates of stock market participation, which is at odds with empirical findings. Introducing a pay-as-you-go public pension scheme improves the fit of the model's consumption profile but does little to aid explanation of asset allocations. Introducing per period fixed costs to the baseline model considerably improves the matching of participation rates but consumption remains close to downward sloping. Including fixed costs and public pensions enables the model to explain the observed hump shape in both consumption and stock market participation profiles. Per period fixed costs of just over 4 percent of permanent income are found to explain the limited stock market participation, with a proportional pension tax rate of twenty percent. More generous public pensions are seen to crowd out private savings and significantly influence the estimates of fixed costs.

The portfolio shares over the life-cycle remain fairly high for all of our simulations, especially in early working life. Further empirical work is needed in order to determine precise predictions of how portfolio shares change as households age, but our finding of full portfolio specialisation (albeit by only a few wealthy households) for young workers seem unrealistic. This divergence of the model's predictions from observed reality could be the result of not fully capturing effective background risk with the exogenous i.i.d. labour income process. Further work needs to be carried out relaxing this i.i.d. assumption and looking at the effects of including other risks to lifetime wealth which vary with age, for example housing.

6 Appendix

The results presented in Section 4 uses now standard techniques to solve the model by backwards induction, starting from a terminal condition, in order to obtain the optimal policy functions for each age, mapping the state variables into the controls. Using these functions the model is simulated forward from $t = 1$ with an initial asset level of zero. The model is simulated 10,000 times with ex-ante homogenous individuals who differ ex-post due to different shock realisations and the average consumption, accumulated wealth and asset allocations are 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 discretised into a finite number of nodes and interpolated using local approximation methods[12].

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

The grids are defined so as to avoid the need for extrapolation outside the grid[13]. The concavity of the consumption function leads to us choosing 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-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 horizontal as it approaches the complete markets outcome. The solution is found using NAG routines[14], except for when these methods fail to converge, in which case the non-linear system is solved using a bisection method[15].

We perform all numerical integration using Gaussian quadrature to approximate the distributions of labour income and the risky asset. The income shocks are discretised into six values and the risky return uses three point quadrature. In the simulations, the permanent shock to labour income is approximated as a continuous random variable. Each time period is taken as a year of life. T was taken to be 58 years and R 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 realised values of the Euler equations. When averaged across individuals these do not deviate significantly from their expected values.

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

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

[15] 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.

7 References

Alan S, 2005, "Entry costs and stock market participation over the life cycle," *Institute for Fiscal Studies Working Paper*, W05/01

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 optimisation? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103, 1121-1157

Bertaut C, 1997, "Stockholding behaviour 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 behaviour 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

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

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

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

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

Cocco J, Gomes, F, and Maenhout P, 1999, "Portfolio choice over the life cycle," *Working Paper*, Harvard University

Davis S, and Willen P, 2000, "Occupational level income shocks and asset returns: their covariance and implications for portfolio choice," *NBER Working Paper,* 7905

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

Gomes F, and Michaelides, 2005, "Optimal life-cycle asset allocation: understanding the empirical evidence," *CEPR Discussion Paper, 4853*

Gomes F, and Michaelides, 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

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

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, 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

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

MEGHIR C, AND PISTAFERRI L, 2004, "Income variance dynamics and heterogeneity," *Econometrica*,

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

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 Paper*, 6185

VISSING-JORGENSEN A, 2002, "Towards an explanation of household portfolio choice heterogeneity nonfinancial income and participation cost structures," *mimeo*, University of Chicago

YAO R, AND ZHANG H, 2005, "Optimal consumption and portfolio choices with risky housing and stochastic labour income," *Review of Financial Studies*, 18, 197-239