NBER WORKING PAPER SERIES

INVESTING RETIREMENT WEALTH: A LIFE-CYCLE MODEL

John Y. Campbell
João F. Cocco
Francisco J. Gomes
Pascal J. Maenhout

Working Paper 7029
http://www.nber.org/papers/w7029

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 1999

Campbell gratefully acknowledges the financial support of the National Science Foundation, Cocco of the Banco de Portugal, Cocco and Gomes of Fundacão para a Ciencia e Tecnologia, Portugal, and Maenhout of the Fund for Scientific Research Flanders. This paper was prepared for the NBER Conference on Risk Aspects of Investment-Based Social Security Reform, Cheeca Lodge, Florida, January 15-16, 1999. We are grateful for the helpful comments of David Cutler, David Weil, Amir Yaron, our discussant, and other conference participants. The views expressed in this paper are those of the authors and do not reflect those of the National Bureau of Economic Research.

Investing Retirement Wealth: A Life-Cycle Model
John Y. Campbell, João F. Cocco,
Francisco J. Gomes, and Pascal J. Maenhout
NBER Working Paper No. 7029
March 1999

ABSTRACT

If household portfolios are constrained by borrowing and short-sales restrictions, or by fixed costs of participating in risky asset markets, then alternative retirement savings systems may affect household welfare by relaxing these constraints. This paper uses a calibrated partial-equilibrium model of optimal life-cycle portfolio choice to explore the empirical relevance of these issues. In a benchmark case, we find ex-ante welfare gains equivalent to a 3.7% increase in consumption from the investment of half of retirement wealth in the equity market. The main channel through which these gains are realized is that the higher average return on equities permits a lower Social Security tax rate on younger households, which helps households smooth their consumption over the life cycle. There is a smaller welfare gain of 0.5% of consumption when Social Security tax rates are held constant. We also find that realistic heterogeneity of risk aversion and labor income risk can strongly affect optimal portfolio choice over the life cycle, which provides one argument for a privatized Social Security system with an element of personal portfolio choice.

John Y. Campbell
Sloan School of Management
E52-456 MIT
Cambridge, MA 02142
and NBER
jyc@mit.edu

João F. Cocco
Department of Economics
Littauer Center
Harvard University
Cambridge, MA 02138

Francisco J. Gomes
Department of Economics
Littauer Center
Harvard University
Cambridge, MA 02138

Pascal J. Maenhout
Department of Economics
Littauer Center
Harvard University
Cambridge, MA 02138

1 Introduction

During the past few decades, American households have begun to display increasing financial sophistication and awareness of rates of return on alternative investments. At the same time the implicit rate of return on contributions to the Social Security system has declined as the system has matured; and this rate of return is projected to decline further in the 21st Century in response to unfavorable demographic trends (Geanakoplos, Mitchell, and Zeldes 1998). Unsurprisingly, politicians and the public have become interested in the possibility of moving to a privatized system in which retirement contributions earn market-based rates of return.

Unfortunately, it is not straightforward to compare alternative retirement systems. Three important issues affect the comparison and invalidate simple rate-of-return calculations. First, the return on the current system is low in part because of the overhang of unfunded liabilities. Past generations have received a gift that must be paid off before the economy can enjoy the steady-state benefits of any new system. Second, capital income taxation puts a wedge between pre-tax and after-tax rates of return. Welfare calculations should take account of the tax revenue generated by capital accumulation (in some systems, this tax revenue is forgiven and private retirement accounts earn higher pre-tax rates of return). Third, returns on alternative financial assets can differ if these assets have different risk characteristics. A valid comparison of rates of return must adjust for risk.

This paper focuses on the last issue, the evaluation of alternative investments with different risk characteristics. From the point of view of households, the current Social Security system represents a defined-benefit pension plan in which income realizations through life are tied to annuity payments in retirement. This is similar to a system in which households are forced to accumulate a riskless asset in a retirement savings account.

Some commentators have recently argued that households would be better off if their retirement savings were invested in risky assets such as equities that have a higher average return. This could be achieved within the present system if the Social Security trust funds were invested in equities; alternatively, within a privatized system, household retirement accounts could include equity investments.

If a household can borrow to invest in equities, however, then the accumulation of riskless assets within a Social Security account need not restrict the household's overall portfolio. The household can undo riskless Social Security accumulation by borrowing outside the retirement account; the household's overall portfolio can be made just as risky as if the retirement account itself were invested in equities. Thus any claimed benefits for a change in the Social Security system must depend on the presence of portfolio constraints that prevent this sort of asset transformation.

Two different sorts of constraints are potentially relevant. First, a household may be unable to borrow at the riskless interest rate to finance equity investments (Constantinides, Donaldson, and Mehra 1998). Second, a household may face fixed costs of equity market participation; if these fixed costs exceed the benefit of participation, the household may hold no equities (Abel 1998). These constraints may affect different households differently. The first constraint is particularly likely to bind on a household whose unconstrained optimal equity position is particularly large, while the second constraint is particularly likely to bind on a poor household with little total wealth. These different sorts of households may be differentially affected by a Social Security reform that alters portfolio constraints.

Clear understanding of these issues requires a well developed normative theory of optimal portfolio choice over the life cycle. Until very recently, however, theoretical work on this subject lagged far behind the familiar theory of single-period optimal portfolio choice. Samuelson (1969) and Merton (1969, 1971) showed that there are conditions under which long-lived investors choose the same portfolios as single-period investors, so that the investment horizon is irrelevant; unfortunately these conditions are highly restrictive in that they include power utility, returns on safe and risky investments that are independently and identically distributed over time, and, most disturbing of all, the absence of labor income.

In the last few years economists have returned to this topic and have begun to explore long-run portfolio choice when these restrictions are relaxed. Ross (1997) shows how horizon effects can emerge from more general models of preferences. Brennan and Xia (1998), Campbell and Viceira (1998), and Wachter (1998b) consider changes over time in the riskless real interest rate, while Balduzzi and Lynch (1997), Barberis (1999), Brandt (1999), Campbell and Viceira (1999), Kim and Omberg (1996), Samuelson (1991), and Wachter (1998a) consider changes over time in the equity premium, and Brennan, Schwartz, and Lagnado (1997) and Liu (1998) allow a more complex pattern of variation in both the real interest rate and the equity premium. The effect of labor income on portfolio choice has been explored by Bertaut and Haliassos

(1997), Bodie, Merton, and Samuelson (1991), Cocco, Gomes, and Maenhout (henceforth CGM, 1998), Gakidis (1997), Heaton and Lucas (1997a), Storesletten, Telmer, and Yaron (1998a), and Viceira (1997), among others.

In this paper we concentrate on the effect of labor income. The theoretical literature on this subject can be loosely summarized as follows. A household with labor income has an implicit holding of a nontradable asset, human capital, that represents a claim to the stream of future labor income. This nontradable asset can "crowd out" explicit asset holdings. If labor income is literally riskless, then riskless asset holdings are strongly crowded out and the household will tilt its portfolio strongly towards risky assets (Bodie, Merton, and Samuelson 1991). If the household is constrained from borrowing to finance risky investments, the solution may be a corner at which the portfolio is 100% risky assets. If labor income is risky but uncorrelated with risky financial assets, then riskless asset holdings are still crowded out but less strongly; the portfolio tilt towards risky assets is reduced (Viceira 1997). If labor income is positively correlated with risky financial assets, then risky assets can actually be crowded out, tilting the portfolio towards safe financial assets.

Under the assumption that income shocks are uncorrelated or only weakly correlated with stock returns, these results suggest that households who expect high future labor income—discounted at some appropriate rate and measured relative to financial wealth—should have the strongest desire to hold stocks. In a life-cycle model with a realistic age profile of income, the discounted value of expected future income increases relative to financial wealth in the very early part of adulthood, but peaks fairly early and then declines as workers approach retirement. This suggests that fairly young (but not the very youngest) households are the most likely to be affected by borrowing constraints that limit their equity positions. While empirical evidence on household portfolio allocation is fragmentary, a few recent empirical papers have found that household portfolios over the life cycle have hump-shaped equity positions, and U-shaped positions in safe assets, consistent with the message of the theoretical literature (Bertaut and Haliassos 1997, Heaton and Lucas 1997b, Poterba and Samwick 1997).

A complicating factor is that many households, particularly younger and poorer ones, appear to hold no equities at all. This is inconsistent with simple frictionless models of optimal portfolio choice, but may be explained if there is a fixed cost of participating in equity markets. Such a fixed cost would deter young households from buying equities, but later in the life cycle these households might find it worthwhile to begin participating if their wealth levels are high enough to justify paying the cost.

In this paper we explore the quantitative importance of these effects by solving a calibrated life-cycle model of consumption and portfolio choice with labor income uncertainty. The model is set in partial equilibrium and takes as given the stochastic properties of income and asset returns. We closely follow CGM (1998), but augment their model to allow us to explore alternative retirement savings systems and fixed costs of equity market participation.

We also ask whether heterogeneity across households is likely to have a large effect on optimal investment patterns. This issue is important for the debate over social security privatization. A privatized system can allow greater individual choice over the investment of retirement wealth, but opponents argue that some individuals may fail to invest optimally and that privatization may increase administrative costs. Whatever the merits of these arguments, it is important to understand the potential gains from individual choice in the absence of administrative costs and optimization failures. To explore this issue we compare the labor income risk of individuals working in different sectors of the economy, and study the sensitivity of optimal choices to differences in the rate of time preference and the coefficient of relative risk aversion.

The organization of the paper is as follows. Section 2 lays out our life-cycle model and calibrates the parameters. Section 3 presents benchmark results in graphical form. Section 4 explores heterogeneity across households. Section 5 conducts a welfare analysis, and Section 6 concludes.

2 A Life-Cycle Model of Portfolio Choice

2.1 Model Specification

2.1.1 Time parameters and preferences

We let t denote adult age. The investor is adult for a maximum of T periods, of which he works the first K. For simplicity K is assumed to be exogenous and deterministic. We allow for uncertain life-span in the manner of Hubbard, Skinner and Zeldes (1994). Let p_t denote the probability that the investor is alive at date $t+1$, conditional on being alive at date t. Then, investor i's preferences are described by the time-separable power utility function:

$$E_1 \sum_{t=1}^{T} \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{C_{it}^{1-\gamma}}{1-\gamma}, \tag{1}$$

where C_{it} is the level of date t consumption, $\gamma > 0$ is the coefficient of relative risk aversion, and $\delta < 1$ is the discount factor. We assume that the individual derives no utility from leaving a bequest.

2.1.2 The labor income process

Investor i's age t labor income, Y_{it}, is exogenously given by:

$$\log(Y_{it}) = f(t, Z_{it}) + v_{it} + \varepsilon_{it} \quad \text{for} \quad t \leq K, \tag{2}$$

where $f(t, Z_{it})$ is a deterministic function of age and other individual characteristics Z_{it}, ε_{it} is an idiosyncratic temporary shock distributed as $N(0, \sigma_\varepsilon^2)$, and v_{it} is given by

$$v_{it} = v_{i,t-1} + u_{it}, \tag{3}$$

where u_{it} is distributed as $N(0, \sigma_u^2)$ and is uncorrelated with ε_{it}. Thus log income is the sum of a deterministic component that can be calibrated to capture the hump shape of earnings over the life cycle, and two random components, one permanent and one transitory. We assume that the temporary shock ε_{it} is uncorrelated across households, but we decompose the permanent shock u_{it} into an aggregate component ξ_t and an idiosyncratic component ω_{it}, uncorrelated across households:

$$u_{it} = \xi_t + \omega_{it}. \tag{4}$$

This decomposition implies that the random component of aggregate labor income follows a random walk, an assumption made by Fama and Schwert (1977) and Jagannathan and Wang (1996). While macroeconomists such as Campbell (1996), Campbell and Mankiw (1989), and Pischke (1995) have found empirical evidence for short-term persistence in aggregate quarterly labor income growth, the simplification to a random walk should have little effect on optimal consumption and portfolio choice over the life cycle.

2.1.3 Financial assets

We assume that there are two assets in which the agent can invest: a riskless asset with gross real return $\overline{R}_f$, which we call *Treasury bills*, and a risky asset with gross real return R_t, which we call *stocks*. The excess return on the risky asset, $R_{t+1} - \overline{R}_f$, is given by:

$$R_{t+1} - \overline{R}_f = \mu + \eta_{t+1} \tag{5}$$

where η_{t+1}, the period $t+1$ innovation to excess returns, is assumed to be i.i.d. over time and distributed as $N(0, \sigma_\eta^2)$. We allow innovations to excess returns to be correlated with innovations to the aggregate component of permanent labor income, and we write the correlation coefficient as $\rho_{\xi\eta}$. We also allow for fixed costs of equity market participation: to have access to the stock market the investor must pay a one-time monetary fixed cost equal to F.

2.1.4 Retirement and liquid wealth

We model a system of mandatory saving for retirement in the following simple way. During working life the individual must save a fraction, θ, of current labor income as retirement wealth. Under this assumption disposable labor income, Y_{it}^d, is given by:

$$Y_{it}^d = (1-\theta)Y_{it} \;\; \text{for} \;\; t \leq K. \tag{6}$$

The amount θY_{it} is added to retirement wealth, denoted by W_{it}^R. During working life retirement wealth is illiquid; the individual cannot consume it or borrow against it. At age K retirement wealth is rolled into a riskless annuity, so that the individual receives in each of the retirement years the annuity value corresponding to W_{iK}^R. This assumption of riskless annuitization affects the portfolio choices of older investors. An interesting extension of our work would be to allow investors to choose between riskless and variable annuities.

We will consider several alternative systems governing the investment of retirement wealth during working life. In the first system the individual is forced to hold retirement wealth in riskless assets. This implies that $W_{it}^R = B_{it}^R$, where B_{it}^R is the dollar amount of retirement wealth investor i has in riskless assets. In alternative systems retirement wealth is partially or fully invested in risky assets, but the allocation remains constant over time and is not controlled by the investor. We interpret this as the social security administration managing the individual's retirement account. For this reason the fixed cost of investing in stocks applies only to investments outside the retirement account.

Investors also have liquid wealth outside their retirement accounts. We denote liquid wealth of investor i at date t by W_{it}^L, and liquid holdings of bills and stocks by B_{it}^L and S_{it}^L, respectively. We assume that the investor faces the following borrowing and short-sales constraints:

$$B_{it}^L \geq 0, \tag{7}$$

$$S_{it}^L \geq 0. \tag{8}$$

The borrowing constraint (7) ensures that the investor's allocation to bills in both the liquid and retirement accounts is non-negative at all dates. It prevents the investor from borrowing against future labor income or retirement wealth. The short-sales constraint (8) ensures that the investor's allocation to equities is non-negative at all dates.

2.1.5 The household's optimization problem

In each period of a household's working life ($t \leq K$) the timing of events is as follows. The investor starts the period with liquid wealth W_{it}^L and retirement wealth W_{it}^R.

Then labor income Y_{it} is realized. Following Deaton (1991) we denote *cash on hand* in period t by

$$X_{it} = W_{it}^L + (1-\theta)Y_{it}. \tag{9}$$

The investor must decide how much to consume, C_{it}, whether to pay the fixed cost of entering the stock market (if he has not done so before), and how to allocate the remaining cash on hand between stocks and bills. We denote the proportion of liquid wealth invested in stocks by α_{it}^L. The proportion of retirement wealth invested in stocks, α_{it}^R, is given exogenously by the retirement system and does not vary over time, so $\alpha_{it}^R = \alpha_i^R$ for all t. We will consider different values for α_i^R.

Next-period liquid and retirement wealth are then given by:

$$W_{i,t+1}^L = R_{p,i,t+1}^L \left[W_{it}^L + (1-\theta)Y_{it} - C_{it} - (f_{it} - f_{i,t-1})F \right], \tag{10}$$

$$W_{i,t+1}^R = R_{p,i,t+1}^R \left[W_{it}^R + \theta_t Y_{it} \right], \tag{11}$$

where f_{it} is a binary variable which equals zero until the investor pays the fixed cost of entering the stock market and equals one thereafter, and $R_{p,i,t+1}^j$ is the return on the portfolio held from period t to period $t+1$:

$$R_{p,i,t+1}^j \equiv \alpha_{it}^j R_{t+1} + (1-\alpha_{it}^j)\overline{R}_f, \qquad j = L, R. \tag{12}$$

Here α_{it}^L is freely chosen when $f_{it} = 1$, and equals zero when $f_{it} = 0$.

After retirement ($t > K$), the problem takes the same form except that retirement wealth no longer accumulates. Instead, it is annuitized and provides riskless income $A(W_{iK}^R)$. After-tax labor income $(1-\theta)Y_{it}$ in (9) and (10) is replaced by $A(W_{iK}^R)$.

The problem the investor faces is to maximize (1) subject to the working-life and retirement versions of (2) through (12), plus the constraints that consumption must be non-negative at all dates. The control variables of the problem are $\{C_t,\ \alpha_{it}^L, f_{it}\}_{t=1}^T$. The state variables are $\left\{t,\ X_{it}, W_{it}^R, v_{it}, f_{i,t-1}\right\}_{t=1}^T$. The problem is to solve for the policy rules as a function of the state variables, i.e., $C_{it}(X_{it}, W_{it}^R, v_{it}, f_{i,t-1})$, $\alpha_{it}^L(X_{it}, W_{it}^R, v_{it}, f_{i,t-1})$ and $f_{it}(X_{it}, W_{it}^R, v_{it}, f_{i,t-1})$.

2.1.6 Numerical solution

This problem cannot be solved analytically. We derive the policy functions numerically by discretizing the state-space and the variables over which the choices are made, and by using Gaussian quadrature to approximate the distributions of the innovations to the labor income process and risky asset returns (Tauchen and Hussey, 1991). The problem is then solved by using backward induction. In period T the investor consumes all his wealth and the value function coincides with the instantaneous utility. In every period t prior to T, and for each admissible combination of the state variables, we compute the value associated with each level of consumption, decision to pay the fixed cost of entering the stock market, and share of liquid wealth invested in stocks. This value is equal to current utility plus the expected discounted continuation value. To compute this continuation value for points which do not lie on the grid we use cubic spline interpolation. The combinations of the choice variables ruled out by the constraints of the problem are given a very large (negative) utility such that they are never optimal. We optimize over the different choices using grid search.

When the fixed cost of equity market participation F is equal to zero, we simplify the solution by exploiting the scale-independence of the maximization problem and rewriting all variables as ratios to the permanent component of labor income.

2.2 Calibration

2.2.1 Time parameters and preferences

Adult age starts at age 20 for households without a college degree, and at age 22 for households with a college degree. The age of retirement is set to 65 for all households. The investor dies with probability one at age 100. Prior to this age we use the mortality tables of the National Center for Health Statistics to parameterize the conditional survival probabilities, p_j for $j = 1, ...T$. We set the discount factor δ to 0.96, and the coefficient of relative risk aversion γ to 5. In variations of the benchmark case, we also consider investors who are extremely impatient with $\delta = 0.80$, comparatively risk-tolerant with $\gamma = 2$, and extremely risk-averse with $\gamma = 10$.

2.2.2 The labor income process

To estimate the labor income process we follow CGM (1998). Here we briefly describe the data and estimation method.

We use the family questionnaire of the PSID to estimate equations (2) and (3) which give labor income as a function of age and other characteristics. Families that are part of the Survey of Economic Opportunities subsample are dropped to obtain a random sample. Only households with a male head are used, as the age profile of income may differ across male- and female-headed households and relatively few observations are available for female-headed households. Retirees, non-respondents, students, homemakers, and household heads younger than 20 (22 for college-educated households) or older than 65 are also eliminated from the sample.

Like CGM and Storesletten, Telmer, and Yaron (1998a,b), we take a broad definition of labor income so as to implicitly allow for insurance mechanisms—other than asset accumulation—that households use to protect themselves against pure labor income risk. Such insurance mechanisms include welfare programs that effectively set a lower bound on the support of non-asset income, endogenous variation in the labor supply of both male and female household members, financial help from relatives and friends, and so on. Thus we define labor income as total reported labor income plus unemployment compensation, workers' compensation, social security, supplemental social security, other welfare, child support and total transfers (mainly help from relatives), all this for both head of household and if present his spouse. Observations which still report zero for this broad income category are dropped. Labor income defined this way is deflated using the Consumer Price Index, with 1992 as the base year. The sample starts in 1970, so a household appears at most 24 times in the sample. Households with fewer observations are retained in the panel.

The estimation controls for family-specific fixed effects. To control for education the sample is split into three groups: households without high school education, a second group with high school education but without a college degree, and finally college graduates. This sample split is intended to accommodate the well-established finding that age profiles differ in shape across education groups (Attanasio 1995, Hubbard, Skinner and Zeldes 1994). For each education group the function $f(t, Z_{it})$ is assumed to be additively separable in t and Z_{it}. The vector Z_{it} of personal characteristics, other than age and the fixed household effect, includes marital status and household composition. Household composition equals the additional number of family members in the household besides the head and (if present) his spouse.

Ideally one should also control for occupation. Using PSID data this is problematic because from the 1975 wave onwards the majority of the unemployed report no occupation, and are categorized together with people who are not in the labor force. But modelling unemployment as a switch in occupation is inappropriate as the possibility of unemployment through layoff is one of the main sources of labor income risk. We explore this in greater detail in Section 4 below.

To obtain age profiles suitable for the simulation model of life-cycle portfolio choice, we fit a third-order polynomial to the age dummies estimated from the PSID. The resulting income profiles are similar to those used in Gourinchas and Parker (1996), Attanasio (1995), and Carroll and Summers (1991). They are plotted in Figure 1 along with the underlying age dummies for each of the three educational groups.

To estimate the variances of permanent and temporary shocks to labor income, we follow Carroll and Samwick (1997). Defining Y_{it}^* as

$$\log(Y_{it}^*) \equiv \log(Y_{it}) - \widehat{f}(t, Z_{it}), \tag{13}$$

then

$$Var(\log(Y_{i,t+d}^*) - \log(Y_{it}^*)) = d * \sigma_u^2 + 2 * \sigma_\varepsilon^2. \tag{14}$$

We estimate σ_u^2 and σ_ε^2 by running an OLS regression of $Var(\log(Y_{i,t+d}^*) - \log(Y_{it}^*))$ on d and a constant term. We find that groups with lower education tend to have more variable transitory income shocks, but less variable permanent shocks, than groups with higher education. Table 1 reports these variances.

We use a similar procedure to estimate the correlation between labor income shocks and stock returns, $\rho_{\xi\eta}$. The change in $\log(Y_{it}^*)$ can be written as

$$\Delta \log(Y_{it}^*) = \xi_t + \omega_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}. \tag{15}$$

Averaging across individuals gives

$$\overline{\Delta \log(Y_{i,t}^*)} = \xi_t. \tag{16}$$

The correlation coefficient is then easily computed from the OLS regression of $\overline{\Delta \log(Y_t^*)}$ on demeaned excess returns:

$$\overline{\Delta \log(Y_{it}^*)} = \beta \left(R_{t+1} - \overline{R}_f - \mu \right) + \psi_t. \tag{17}$$

As an empirical measure for the excess return on our stylized risky asset, we use CRSP data on the New York Stock Exchange value-weighted stock return relative to the T-bill rate. For all education groups, the regression coefficients are strikingly low and insignificant. To allow for potential lags in the realization of labor income, we repeat the exercise with the excess stock return lagged one year. The relationship becomes much stronger: the regression coefficient now varies from 0.06 to 0.10, and the correlation coefficient from 0.32 to 0.52, as reported in Table 1. Interestingly, the correlation of labor income with the stock market is larger and more significant for households with higher education.[2]

In our portfolio choice model, allowing for lags in the relationship between innovations in stock returns and permanent shocks to labor income unfortunately requires an additional state variable. We therefore assume the correlation is contemporaneous. The model requires the variances of both ξ_t, the aggregate permanent labor income shock that is correlated with stock market risk and ω_{it}, the idiosyncratic permanent shock to labor income. The first variance is obtained immediately as the variance of $\overline{\Delta \log(Y_t^*)}$. Subtracting this variance from the total variance of u_{it} gives then the variance of ω_{it}.

2.2.3 Other parameters

The riskless real interest rate is assumed to be constant at 2%. We set the equity premium μ equal to 4%. This is well below the long-run historical average, but represents a reasonable compromise between that average and lower forward-looking estimates based on the observation that stock prices have tended to increase in recent years relative to corporate earnings (Blanchard 1993, Campbell and Shiller 1998). We set the standard deviation of innovations to the risky asset σ_η to 0.157. Recall that the classic formula for the risky asset portfolio share, under power utility with iid returns and no labor income, is $\mu/\gamma\sigma_\eta^2$. With these parameters the implied risky asset share would be about 1/3 at the benchmark risk aversion of 5; we find higher optimal shares in our model only because of the presence of labor income. We set

[2]We also examined the relationship of labor income shocks with lagged returns on long-term government bonds. These results are reported in Section 4.

the fixed cost of equity market participation to zero in the benchmark case, but we go on to consider a $10,000 fixed cost.

The proportion of labor income θ that is added to retirement wealth is equal to 10% of current labor income when retirement wealth accumulates at the riskless rate. This value implies an average replacement ratio at age 65 of 60%. When retirement wealth is also invested in stocks we either fix θ at the same 10% value, or adjust it so as to maintain the replacement ratio at 60%. We will show that the value of θ has a very important effect on our results. Table 1 summarizes the parameters used in the baseline case.

3 Benchmark Results: A Graphical Summary

The first comparison we consider is between a system with riskless retirement accumulation ($\alpha^R = 0$) and one in which at each age half of retirement wealth is invested in stocks ($\alpha^R = .5$). In the latter system, we reduce the social security tax rate from 10% to 6% so that on average the replacement ratio is the same in both systems and equal to 0.6. At retirement, the account is annuitized at the riskless interest rate, so that on average, and given survival probabilities, the system has zero balance.

To study the behavior of the variables in the model, we calculate cross-sectional averages across 10,000 households receiving different draws of income and asset returns, and plot them against age. Figure 2 plots labor income net of social security contributions, consumption, liquid wealth and retirement wealth for households with a high-school degree (the life-cycle patterns for other education groups are similar). Figure 2a illustrates the system in which retirement wealth is fully invested in the riskless asset, and Figure 2b illustrates the system in which retirement wealth is partially invested in stocks.

In both systems, the average consumer is borrowing-constrained early in life. Consumption tracks net income very closely and little savings accumulate outside the retirement account until after age 40. These limited savings early in life are driven by the precautionary savings motive; thus like Gourinchas and Parker (1996) we find that younger consumers are buffer-stock savers rather than life-cycle savers in the classic sense. Consumption rises with income early in life because of borrowing constraints, and falls later as increased mortality drives up the effective rate of time preference; thus consumption profiles are hump-shaped over life, as found in the literature on life-cycle consumption behavior.[3]

Investment of some retirement wealth in stocks has an income effect. Because the average return on stocks is higher than the average return on bills, and because younger consumers have neither the desire nor the liquid wealth to offset a shift of retirement wealth into stocks, the shift increases average lifetime resources. Since we reduce the social security tax rate to keep the average replacement ratio constant

[3]We could generate a more pronounced hump shape in consumption if we added age-specific preference shocks to the model.

across systems, the investment of retirement wealth in stocks frees up resources in the working years. These additional resources are consumed early in life, since at this stage households are borrowing-constrained.

Of course the investment of retirement wealth in stocks has a cost: It imposes additional risk on households. In midlife households react by increasing their precautionary saving, accumulating more liquid wealth and consuming less relative to income. After retirement the additional wealth is run down, since at this stage retirement wealth is rolled into a riskless annuity. These patterns show up in the paths of consumption relative to income in Figure 2.

Figure 3 plots liquid wealth and liquid holdings of equities and bills over the life cycle. In each retirement system the borrowing constraint binds for young households; they would like to take more equity risk but are unable to do so. For approximately the first 20 years of life, they hold 100% of their portfolios in the form of equity. Households in midlife hold bills, but these holdings decrease again after retirement.

Figure 4 plots the portfolio share of stocks in liquid wealth. The crucial variables for portfolio composition are liquid wealth, retirement wealth and future labor income. In the model future labor income, although risky, can be thought of as implicit holdings of a riskless asset. Innovations to labor income are positively correlated with innovations to stock returns, but this correlation is not sufficiently large for future labor income to resemble more closely stocks than bills. Since early in life the implicit holdings of the riskless asset in the form of future labor income are large, the investor wishes to invest what little liquid wealth he has fully into stocks. From age 40 onwards liquid wealth increases relative to future labor income and retirement wealth, so that implicit holdings of the riskless asset become less important. This induces a shift in the composition of liquid wealth towards bills. After retirement liquid wealth is run down more rapidly than the implicit annuitized holdings of the riskless asset. As this happens the implicit holdings of the riskless asset become relatively more important, inducing a shift in portfolio composition back towards stocks.

These considerations explain the life-cycle patterns in both Figure 4a and Figure 4b. There are however important differences in magnitudes between the two figures. In Figure 4b the midlife decrease in the share invested in stocks is much more dramatic. Since the investor already holds risky assets in the retirement account, he wishes to hold a safer liquid portfolio.

Another way to understand these results is to compare the patterns of current utility of consumption over the life cycle. Figure 5a shows the ratio of average current utility for households in the 50/50 system to the average current utility of households in the 100/0 system. We see that most of the gain from investing retirement wealth in stocks occurs early in life. The source of this gain is the higher levels of consumption that a lower social security contribution allows. Return risk in the retirement account allows some households to end up poorer so that after age 55 current utility is on average higher in the 100/0 system.

Return risk also increases the dispersion of utility. The standard deviations of current utility across households with different income and return realizations are higher in the 50/50 system than in the 100/0 system. Figure 5b reports the ratio of these standard deviations in the two systems. These ex-post differences raise important practical issues for designers of retirement systems because they may create an incentive to bail out cohorts negatively affected by lower stock return realizations.

Of course, a proper welfare analysis requires the discounting of current utility over the life cycle. We defer such an analysis to section 5 below.

One limitation of the results reported so far is that they counterfactually predict 100% stock market participation among younger investors. However we can modify this prediction, with little effect on other aspects of the model, by adding a fixed cost of stock market participation. Figure 6 reports results with a $10,000 fixed cost. The fraction of households who have paid the fixed cost and the average share of assets invested in stocks are plotted for each retirement system. Early in life the two series move almost perfectly together, showing that young investors are either entirely in or entirely out of the market; later in life all investors have paid the fixed cost, and the model behaves much as it did in the absence of the cost.

4 Heterogeneity

In the previous section we studied a representative household with a high school education but no college degree. Results are similar for representative households in the other two education groups. However households may differ along other dimensions. For example, labor income processes may differ for households who work in different industries, and for self-employed households. Also some households may be more impatient or risk-averse than others as found by Barsky, Juster, Kimball, and Shapiro (1997). These differences across households may have important effects on optimal investment strategies. In this section we consider this issue.

4.1 Measuring Heterogeneity in Labor Income

We first consider variation in the stochastic structure of the labor income process across industries, and then study differences between self-employed and non-self-employed households.

We use the 2-digit SIC classification to split households into 12 different industries. Starting in 1972, the PSID reports both the current industry of the household head if currently working and the last industry if currently unemployed. This is the information we use. Three caveats apply however. First, we ignore the industry of the spouse. This might be problematic because the spouse's labor income is added to the head's labor income, yet it might have quite different risk characteristics. Second, on average 16% of our respondents switch industries each year. Business-cycle considerations (like the anticipation of a recession) might force people out of cyclical sectors and into less volatile industries. As we do not model the switching decision, our estimates of the sensitivities of labor income shocks to financial market risks for different industries might be biased. However we did not find any significant effects when we regressed the number of industry switchers onto innovations in business cycle indicators. Third, there is a timing issue because the labor income reported in the PSID is for the previous calendar year, while the industry concerns the current job.

Table 2 reports the number of household-year observations for each of the 36 different education-industry cells. There is tremendous variation across industries, with particularly small numbers in mining, personal services, and recreation. These industries are omitted from Tables 3 through 5. As a further cut-off, we drop cells in which any PSID wave contains fewer than 20 observations; these cells are left blank in Tables 3 through 5. We do however include observations in these small cells when estimating column and row totals, that is when reporting the results for a given industry across all education levels or for a given education group across all industries.

We reestimate age profiles of labor income for the shorter PSID sample beginning in 1972, and including industry dummies in the vector Z_{it} of personal characteristics. That is, we allow industry to influence only the level and not the shape of the age profile for a household with a given amount of education. We then estimate the stochastic model of labor income separately for each education-industry cell.[4]

Table 3 reports the total variance of income, and its decomposition into permanent and transitory components, for each different education-industry cell. Agriculture has by far the highest variance of labor income shocks. Other industries subject to significant labor income shocks are construction and business services. The variance decomposition indicates that labor income shocks for construction workers without a high school degree are entirely temporary. At the other extreme, permanent income shocks are especially important for college graduates in public administration. As a general pattern, the relative importance of permanent shocks seems to increase with educational attainment. This was already documented for the column totals, but seems robust within individual industries.

The bottom of Table 3 splits the sample in a different way, by distinguishing self-employed from non-self-employed households. (We included both types of households in the industry analysis, since there are too few self-employed households to allow an industry decomposition.) Income variability is dramatically larger for self-employed households. Income shocks are entirely temporary for the self-employed without a high school degree, but are disproportionately permanent for the self-employed in the two higher education groups.

Table 4 considers heterogeneity in the sensitivity of different households' income shocks to lagged stock returns. Table 5 repeats the exercise replacing lagged stock returns with lagged returns on long-term government bonds. Unfortunately the small

[4]We also report results aggregating across industries and education groups. Note that these differ slightly from the results reported in Section 3 because here our sample starts in 1972 and we include industry dummies in the estimation of age profiles for labor income.

cell sizes mean that the results are often statistically insignificant for individual industries, but there are many interesting patterns. Stock market risk seems especially relevant for people in manufacturing, construction and public administration. Interest rate risk shows up for agriculture, professional services and finance, real-estate and insurance, in addition to the stock-market-sensitive sectors. Among college graduates, the self-employed are especially exposed to stock market risk, while interest rate risk is far more important for the non-self-employed. This finding supports the conclusion of Heaton and Lucas (1997b) that privately owned business risk is an especially important substitute for stock market risk in the portfolios of many wealthy households.

4.2 Effects of Heterogeneity on Portfolio Choice

In this section we illustrate the effects of investor heterogeneity on optimal consumption and portfolio choice. First, we consider heterogeneity of preferences, calculating optimal behavior for highly risk-averse investors with $\gamma = 10$ and impatient investors with $\delta = 0.8$. Second, we consider differences in labor income risk of the sort illustrated in section 4.1. To highlight these differences, we simulate the behavior of households whose income is particularly risky and highly correlated with asset returns: self-employed college graduates.

Table 6 shows average consumption and liquid wealth (in thousands of dollars) and the share of liquid wealth invested in stocks for different age groups. This is a more compact way for us to summarize the information presented graphically in Figures 2 through 4. The first three columns of Table 6 use the baseline parameters and consider retirement systems with $\alpha^R = 0$; $\alpha^R = 0.5$ and the same 10% tax rate as with $\alpha^R = 0$; and $\alpha^R = 0.5$ and a lower 6% tax rate that maintains the same average replacement ratio as with $\alpha^R = 0$.

The next two columns of Table 6 present results for a higher risk aversion coefficient of 10 rather than 5. Since the tax rates do not depend on γ, they are the same as in the baseline case. To understand the results for higher γ it is important to remember that with isoelastic preferences this parameter measures both risk aversion and prudence (Kimball 1990). Greater prudence increases precautionary savings and explains why highly risk-averse investors consume less, and save more, until age 65. After this age the precautionary savings motive is reduced since there is no labor income risk and retirement wealth is converted into a riskless annuity. Thus highly risk–averse investors consume more after retirement.

Table 6 also shows that, as one would expect, highly risk-averse investors have a lower portfolio share in stocks. One interesting pattern that is not visible in the table is that very early in life, these investors' equity portfolio share is increasing with age. This pattern does not show up for investors with $\gamma = 5$ because in early life these investors are constrained by their inability to borrow to finance equity investments. The reason for this increasing pattern is explained in CGM. In the presence of an increasing labor income profile the annuity value of future labor income, equivalent to implicit holdings of the riskless asset, increases with age at first as peak earnings years move closer in time. Investors respond to this increase by shifting liquid wealth towards risky financial assets. Later on the annuity value of future labor income starts to decrease as peak earnings are realized and retirement approaches; investors respond by shifting out of stocks in middle age.

The next two columns of Table 6 show optimal consumption, wealth, and portfolio allocation for impatient households with $\delta = 0.8$. These households consume more early in life (roughly up to age 50) and less later. They accumulate almost no wealth, never holding more than about \$10,000 in liquid assets. What little wealth they do accumulate they hold in stocks; their exposure to the stock market is so small that they are extremely tolerant of equity risk.

The last two columns of Table 6 report results for self-employed college-educated households. The preference parameters are the same as in the baseline case, but the results are quite different from the first three columns of Table 6 which apply to households with only a high-school education. The tax rate in the 50/50 system is set to 6.75% to maintain the average replacement ratio for self-employed college-educated households. Looking at the share of liquid wealth invested in stocks, there are two distinctive features: the share invested in risky financial assets is much lower, and it exhibits a clear hump-shape. The higher variance of labor income shocks and the large positive correlation between the latter and innovations to stock returns crowd out investment in risky financial assets. This effect is particularly strong early in life, when the investor has accumulated little liquid wealth.

5 Welfare Analysis

We have presented detailed results describing the effects of alternative retirement systems on consumption, wealth accumulation, and portfolio choice. We now turn to the welfare implications of these systems.

We evaluate each system by discounting current utilities back from the end to the beginning of adult life (20, or 22 for college-educated households). We renormalize discounted utility into consumption-equivalent units, so that a 5% increase represents the increase in utility that would be produced by a 5% increase in consumption at every date. We calculate the expectation of discounted lifetime utility at age 20 across all realizations of income and risky asset returns; in order to measure the variability of outcomes, we also calculate the standard deviation of discounted lifetime utility across these realizations. Finally, we assess welfare effects on retired households by repeating the same exercise discounting only back to age 65.

The top panel of Table 7 reports results for the benchmark case of a household with risk aversion $\gamma = 5$ and a high-school education. In the top row all retirement wealth is invested in the riskless asset; in the second row half of retirement wealth is invested in stocks but the tax rate is held constant at 10%; in the third row half of retirement wealth is invested in stocks and the tax rate is lowered to 6% to maintain the average replacement ratio at 60%; and in the fourth row all retirement wealth is invested in stocks and the tax rate is lowered to 3.75% to maintain the average replacement ratio. The third and fourth columns give the mean and standard deviation of discounted lifetime utility in consumption-equivalent units (thousands of dollars of consumption) at age 20; the fifth and sixth columns report the same moments at age 65; and the remaining columns report the percentage changes in these moments relative to the benchmark case of riskless retirement wealth.

The table shows that a shift from a riskless retirement system to one that is 50% invested in stocks, with a lower payroll tax rate, increases the welfare of the typical 20-year-old household by 3.7%. Most of the welfare gain is due to the reduction in the tax rate; the household gains only 0.5% if the payroll tax rate is held constant. The standard deviation of lifetime utilities across households with different income and asset-return realizations also increases modestly, by 3.4% if the payroll tax rate is reduced and 2.9% if it is held constant.

Households also prefer the risky retirement system at age 65; their discounted utility is on average 7.8% higher if the payroll tax rate is reduced and 20.3% higher if it is held constant (for then they have accumulated greater wealth over their working life and at age 65 are able to enjoy the benefits). A further shift in the retirement system to 100% equity investment, with a payroll tax rate that is lower again, produces even larger utility gains but also considerably greater variability of outcomes.

The next panel of Table 7 considers a highly risk-averse household with $\gamma = 10$. This household is not usually constrained in its equity holdings under the riskless retirement system, and it actually loses slightly if the retirement system is shifted into equities without any reduction in the payroll tax rate. But if the payroll tax rate is reduced, the highly risk-averse household actually gains more than the benchmark household, 4.4% rather than 3.7%. The reason is that under power utility, higher risk aversion implies a lower elasticity of intertemporal substitution—that is, a stronger desire to smooth consumption intertemporally—and the higher disposable income associated with a lower tax rate allows the household to smooth consumption more effectively over the life cycle.

The next panel of Table 7 considers a comparatively risk-tolerant household with $\gamma = 2$. The results here are the mirror image of those for the highly risk-averse household. The risk-tolerant household gains more than the benchmark household if the payroll tax rate is fixed (for it is particularly anxious to relax constraints on its equity holdings) but gains less if the payroll tax rate is reduced (since improved opportunities to smooth consumption over the life cycle are less important for this household).

The last panel of Table 7 reports results for a household headed by a self-employed college graduate. This household has risky labor income that is unusually correlated with the stock market, so its desired stock holdings are smaller than those of the benchmark household; it actually loses slightly from investment of retirement wealth in equities with a fixed tax rate. On the other hand, this household also has a particularly pronounced hump shape in labor income so it is particularly anxious to smooth consumption over the life cycle. The household gains 3.8% from a risky retirement system with a lower payroll tax rate, comparable to the results for the benchmark household.

All the results in Table 7 can be criticized on the grounds that they are derived from a model in which there is no role for a Social Security system. Since households save and invest their liquid wealth optimally, the mandatory saving and rigid asset allocation of the retirement system can only reduce their welfare. In this setting, any reform that effectively reduces the scale of Social Security will increase welfare.

As a partial response to this concern, in Table 8 we report a "paternalistic" welfare analysis in which the government uses different utility parameters than those of the household itself. In the first panel, the household is impatient, with $\delta = 0.8$; left to its own devices it will do very little saving for retirement, as illustrated in Table 6. The government, however, discounts utility using $\delta = 0.96$ as in our benchmark case. Here there is a modest welfare gain of 2.0% from a shift of retirement wealth towards equities with a fixed tax rate, but a small welfare loss of -0.6% if the payroll tax rate is reduced. The reason, of course, is that from the government's point of view people make poor use of their tax cuts, spending them early in life and failing to save enough for retirement.

The second panel of Table 8 reports results when the household is highly risk-averse, with $\gamma = 10$, but the government is risk-tolerant, with $\gamma = 2$. In this case a shift of retirement wealth towards equities forces households to take on more risk, which improves their welfare from the government's point of view. Welfare rises by 1.7% and 2.4%, respectively, in the cases with fixed and reduced payroll tax rates.

All the results we have reported so far assume that the riskless real interest rate is fixed at 2% and the equity premium at 4%. As a final exercise, in Table 9 we consider variations in these parameters. This is a valuable check on the robustness of our basic results; it also enables us to consider what might happen to welfare if the investment of retirement wealth in risky assets reduced the equilibrium equity premium.

The top panel of Table 9 repeats the first three rows of Table 7 for the benchmark case. The second and third panels consider two alternative scenarios in which the equity premium is 1 percentage point lower at 3%. In the first alternative, the equity premium falls but the riskless interest rate is unchanged, so the expected equity return falls by 1 percentage point. This is a scenario envisaged by critics of Social Security investment in equities who worry that such a reform would drive up stock prices and drive down expected stock returns. In the second alternative, the equity premium falls but the riskless interest rate rises by 1 percentage point, leaving the expected equity return unchanged. This scenario is predicted by general equilibrium models in which the return to risky capital is fixed by technology, such as Cox, Ingersoll, and

Ross (1985) or Abel (1999).

Within each of the alternative scenarios, the welfare gains produced by risky investment of retirement wealth are similar to but slightly smaller than those in the benchmark case. In the rows marked with double asterisks, Table 9 compares welfare in the alternative scenarios with risky investment of retirement wealth, to welfare in the benchmark case with riskless investment of retirement wealth. This is a crude way to capture the possibility that risky investment of retirement wealth might lower the equity premium. It turns out that the results are critically dependent on the way in which the equity premium falls. If it falls through lower stock returns, as in the first alternative, then welfare gains are reduced from 3.7% to 2.8%. If it falls through a higher riskless rate, as in the second alternative, then welfare gains are actually increased to 5.0%.

6 Conclusion

Decisions about the quantity and form of retirement saving are among the most important that a typical household takes in the course of a lifetime. Despite the importance of the issue, until very recently financial economists have had little quantitative understanding of the factors that should affect this decision. This gap in our knowledge has made it hard to give sound advice to policymakers considering reforms in retirement systems.

In this paper we have built a partial-equilibrium life-cycle model that can be used to explore the properties of alternative systems. In our benchmark case, we find a welfare gain equivalent to 3.7% of consumption from the investment of half of retirement wealth into equities, accompanied by a reduction in the Social Security tax rate to maintain the same average replacement rate of income in retirement. The main channel through which these gains are realized is that a lower Social Security tax rate helps households smooth their consumption over the life cycle. The gains from equity investment of retirement wealth are smaller, about 0.5% of consumption, when the Social Security tax rate is held constant at its initial level. Interestingly, in our model particularly risk-averse households are particularly keen to smooth consumption and thus experience even larger gains from reduced tax rates made possible by equity returns on retirement wealth. This is true despite the fact that risk-averse households are less enthusiastic equity investors.

While these results are encouraging for proponents of equity investment by the Social Security system, we note two caveats. First, lower Social Security tax rates reduce welfare by 0.6% of consumption in a model in which investors are extremely impatient, with a low time discount factor of 0.8, but the government judges their welfare using a higher time discount factor of 0.96. This calculation is a crude way to capture factors that might lead to inadequate private saving and justify a mandatory retirement saving system. In this model, however, there is a substantial gain of 2.0% from Social Security equity investment with fixed tax rates. These findings suggest that the appropriate adjustment of tax rates will depend on detailed assumptions about the behavior of households, but that under a wide range of assumptions there are welfare gains to be had by investing some retirement wealth into equities.

Second, a system with partial investment of retirement wealth in the equity market has greater variability of outcomes across cohorts. Particularly negative outcomes for some cohorts might provoke pressure for political bailouts, and the anticipation of bailouts might change the consumption behavior modelled here. This is an important issue for research on Social Security reform.

Using data from the PSID, we study heterogeneity in labor income processes across households. We find that some households—particularly self-employed college graduates—are exposed to much greater volatility in their labor income than are typical households. The labor income of these households also tends to be more highly correlated with returns on stocks and long-term bonds. This heterogeneity affects optimal investment strategies and may help to justify Social Security reform that includes an element of personal choice.

We also consider the possibility that investment of retirement wealth in equities might reduce the equity premium. We do not build a general equilibrium model to study this issue, but we do compare results under alternative assumptions about the equity premium. We find that it matters greatly whether the equity premium falls through a decline in the expected return on stocks or through a rise in the riskless interest rate. In the former case the welfare gains from investing retirement wealth in equities are reduced, while in the latter case they are actually greater than in our benchmark model.

In evaluating our results, it is important to be aware of several respects in which our model is oversimplified. First, we consider only self-financing retirement systems in which there is no net payment from any household to any other. Thus we ignore the redistributive features of the present Social Security system, and we have nothing to say about the overhang of liabilities to previous generations implied by the present system. Second, we assume that asset returns are independently and identically distributed. Thus we ignore the variation in real interest rates and equity premia that is the subject of much recent research. Third, we abstract from the existence of owner-occupied housing. This is an important omission since housing is the main component of wealth for many people. Cocco (1998) takes a first step towards the realistic incorporation of housing into a life-cycle model. Fourth, we assume that labor income shocks have constant variances. Some researchers have argued that the variance of idiosyncratic shocks to labor income is higher when the economy is weak and risky asset returns are low; this can have important effects on the demand for risky assets as shown by Mankiw (1986), Constantinides and Duffie (1996), and Storesletten, Telmer, and Yaron (1998a). Finally, we assume that labor income and the retirement age are exogenous to the household. Bodie, Merton, and Samuelson (1991) point out that households with flexible labor supply can afford to hold riskier portfolios because they can adjust to negative asset returns both by changing their consumption and by changing their labor supply. An important task for future research is to incorporate these and other realistic complications into the basic life-cycle model of portfolio choice.

7 References

Abel, Andrew B., 1998, "The Aggregate Effects of Including Equities in the Social Security Trust Fund", unpublished paper, Wharton School of the University of Pennsylvania.

Abel, Andrew B., 1999, "The Social Security Trust Fund, the Riskless Interest Rate, and Capital Accumulation", unpublished paper, Wharton School of the University of Pennsylvania.

Attanasio, Orazio, 1995, "The Intertemporal Allocation of Consumption: Theory and Evidence," *Carnegie-Rochester Conference Series on Public Policy* 39–89.

Balduzzi, Perluigi, and Anthony Lynch, 1997, "The Impact of Predictability and Transaction Costs on Portfolio Choice in a Multiperiod Setting," unpublished paper, Boston College and New York University.

Barberis, Nicholas C., 1999, "How Big Are Hedging Demands? Evidence from Long-Horizon Asset Allocation," forthcoming *Journal of Finance.*

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, 1997, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study", *Quarterly Journal of Economics* 112, 537–579.

Bertaut, Carol C. and Michael Haliassos, 1997, "Precautionary Portfolio Behavior from a Life-Cycle Perspective," *Journal of Economic Dynamics and Control* 21, 1511–1542.

Blanchard, Olivier J., 1993, "Movements in the Equity Premium", *Brookings Papers on Economic Activity* 2, 75–118.

Bodie, Zvi, Robert C. Merton and William Samuelson, 1991, "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model," *Journal of Economic Dynamics and Control*, 16, 427-49.

Brandt, Michael W., 1999, "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach", forthcoming *Journal of Finance.*

Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, "Strategic Asset Allocation," *Journal of Economic Dynamics and Control* 21, 1377–1403.

Brennan, Michael J. and Y. Xia, 1998, "Resolution of a Financial Puzzle", unpublished paper, Anderson Graduate School of Management, UCLA.

Campbell, John Y., 1996, "Understanding Risk and Return", *Journal of Political Economy* 104, 298–345.

Campbell, John Y. and N. Gregory Mankiw, 1989, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence", in Olivier Blanchard and Stanley Fischer eds. *NBER Macroeconomics Annual 1989*, Cambridge, MA: MIT Press.

Campbell, John Y. and Robert J. Shiller, 1998, "Valuation Ratios and the Long-Run Stock Market Outlook", *Journal of Portfolio Management* 24, 11–26.

Campbell, John Y. and Luis M. Viceira, 1998, "Who Should Buy Long-Term Bonds?", NBER Working Paper No. 6801.

Campbell, John Y. and Luis M. Viceira, 1999,"Consumption and Portfolio Decisions When Expected Returns are Time Varying", forthcoming *Quarterly Journal of Economics.*

Carroll, Christopher D. and Lawrence H. Summers, 1991, "Consumption Growth Parallels Income Growth: Some New Evidence," in B.D. Bernheim and John B. Shoven, eds., *National Savings and Economic Performance*, Chicago, IL: University of Chicago Press.

Cocco, João F., 1998, "Owner-Occupied Housing, Permanent Income, and Portfolio Choice", unpublished paper, Harvard University.

Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout, 1998, "Consumption and Portfolio Choice over the Life-Cycle", unpublished paper, Harvard University.

Constantinides, George M., John B. Donaldson, and Rajnish Mehra, 1998, "Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle", NBER Working Paper No. 6617.

Constantinides, George M. and Darrell Duffie, 1996, "Asset Pricing with Heterogeneous Consumers", *Journal of Political Economy* 104, 219–240.

Cox, John, Jonathan E. Ingersoll, Jr., and Stephen A. Ross, "A Theory of the Term Structure of Interest Rates", *Econometrica* 53, 385–408.

Deaton, Angus S., 1991, "Savings and Liquidity Constraints," *Econometrica* 59, 1221–1248.

Fama, Eugene F. and G. William Schwert, 1977, "Human Capital and Capital Market Equilibrium", *Journal of Financial Economics* 5, 115–146.

Gakidis, Harry, 1997, "Earnings Uncertainty and Life-Cycle Portfolio Choice," unpublished paper, MIT.

Geanakoplos, John, Olivia S. Mitchell, and Stephen P. Zeldes, 1998, "Social Security Money's Worth", forthcoming in Olivia S. Mitchell, Robert J. Myers, and Howard Young eds. *Prospects for Social Security Reform*, Pension Research Council and University of Pennsylvania Press.

Gourinchas, Pierre-Olivier and Jonathan Parker, 1996, "Consumption Over the Life Cycle," unpublished paper, MIT.

Heaton, John and Deborah J. Lucas, 1997a, "Market Frictions, Saving Behavior and Portfolio Choice," *Macroeconomic Dynamics* 1, 76–101.

Heaton, John and Deborah J. Lucas, 1997b, "Savings Behavior and Portfolio Choice: Which Risks Matter?," unpublished paper, Northwestern University.

Hubbard, Glenn, Jonathan S. Skinner and Stephen Zeldes, 1994, "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," in Allan H. Meltzer and Charles I. Plosser, eds., *Carnegie-Rochester Conference Series on Public Policy*, 40, 59–125.

Jagannathan, Ravi and Narayana R. Kocherlakota, 1996, "Why Should Older People Invest Less in Stocks Than Younger People?," *Federal Reserve Bank of Minneapolis Quarterly Review*, 11–23.

Jagannathan, Ravi and Zhenyu Wang, 1996, "The Conditional CAPM and the Cross-Section of Expected Returns", *Journal of Finance* 51, 3–53.

Kim, Tong Suk, and Edward Omberg, 1996, "Dynamic Nonmyopic Portfolio Behavior," *Review of Financial Studies* 9, 141—161.

Kimball, Miles, 1990, "Precautionary Saving in the Small and in the Large", *Econometrica* 58, 53–73.

Liu, Jun, 1998, "Portfolio Selection in Stochastic Environments", unpublished paper, Stanford University.

Mankiw, N. Gregory, 1986, "The Equity Premium and the Concentration of Aggregate Shocks", *Journal of Financial Economics* 17, 211–219.

Merton, Robert C., 1969, "Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case," *Review of Economics and Statistics* 51, 247–257.

Merton, Robert C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory* 3, 373–413.

Pischke, Jörn-Steffen, 1995, "Individual Income, Incomplete Information, and Aggregate Consumption", *Econometrica* 63, 805–840.

Poterba, James M. and Andrew A. Samwick, 1997, "Household Portfolio Allocation Over the Life Cycle", NBER Working Paper No. 6185.

Ross, Stephen A., 1997, "Samuelson's Fallacy of Large Numbers Revisited", unpublished paper, Yale University.

Samuelson, Paul A., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics* 51, 239–246.

Samuelson, Paul A., 1991, "Long-Run Risk Tolerance when Equity Returns are Mean Regressing: Pseudoparadoxes and Vindication of a 'Businessman's Risk"' in William C. Brainard, William D. Nordhaus and Harold W. Watts eds., *Money, Macroeconomics, and Economic Policy: Essays in Honor of James Tobin*, Cambridge, MA: MIT Press, 181–200.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 1998a, "Asset Pricing with Idiosyncratic Risk and Overlapping Generations", unpublished paper, Carnegie Mellon University.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 1998b, "The Risk Sharing Implications of Alternative Social Security Arrangements", unpublished paper, Carnegie Mellon University.

Tauchen, George and R. Hussey, 1991, "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models", *Econometrica* 59, 371–396.

Viceira, Luis M., 1997, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Income," unpublished paper, Harvard University.

Wachter, Jessica, 1998a, "Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets", unpublished paper, Harvard University.

Wachter, Jessica, 1998b, "Risk Aversion and Allocation to Long-Term Bonds", unpublished paper, Harvard University.

Table 1: Baseline Parameters

Description	Parameter Value
Retirement age (K)	65
Discount factor (δ)	.96
Risk aversion (γ)	5
Variance of transitory shocks (σ_ε^2)	
No high school	.1056
High school	.0738
College	.0584
Variance of permanent shocks (σ_u^2)	
No high school	.0105
High school	.0106
College	.0169
Sensitivity to stock returns (β)	
No high school	.0956
High school	.0627
College	.0733
Correlation with stock returns ($\rho_{\xi\eta}$)	
No high school	.3280
High school	.3709
College	.5155
Riskless rate ($\overline{R}_F - 1$)	.02
Mean excess return on stocks (μ)	.04
Std stock return (σ_η)	.157
Fixed cost (F)	0 or 10000
Social Security tax rate (θ)	.10

Table 2: Cell Sizes

Industry	No High School	High School	College	Total
Agriculture	765	1165	381	2311
Mining	119	360	98	577
Manufacturing	3116	7319	2430	12865
Construction	1414	2938	494	4846
Transp., comm.	828	3172	701	4701
Trade	1072	4661	1445	7178
Fin., RE, ins.	109	995	965	2069
Bus. serv.	439	1437	548	2424
Personal serv.	159	358	108	625
Recreation	39	236	110	385
Prof. serv.	356	1738	3883	5977
Public adm.	339	2357	1097	3793
Self-employed	1115	3562	2242	6919
Non self-employed	7640	23174	10018	40832
Total	8755	26736	12260	47751

Table 3: Variance Decomposition

	No High School		High School		College		All	
Industry	Total	Perm/Total	Total	Perm/Total	Total	Perm/Total	Total	Perm/Total
Agriculture	-	-	0.3094	0.1350	-	-	0.3166	0.0786
Manufacturing	0.0579	0.0973	0.0506	0.0912	0.0469	0.2483	0.0529	0.1315
Construction	0.1097	0.0000	0.1030	0.1394	-	-	0.1080	0.1451
Transp., comm.	0.1119	0.3058	0.0553	0.0833	-	-	0.0672	0.1616
Trade	0.0997	0.1707	0.0890	0.1254	0.0721	0.2586	0.0895	0.1591
Fin., RE, ins.	-	-	0.0930	0.1264	0.0826	0.2382	0.0906	0.1471
Bus. serv.	-	-	0.1275	0.1820	-	-	0.1188	0.1789
Prof. serv.	-	-	0.0702	0.1180	0.0732	0.1953	0.0773	0.1478
Public adm.	-	-	0.0470	0.1657	0.0361	0.2594	0.0450	0.1882
Self-employed	0.3695	0.0000	0.2199	0.2192	0.1985	0.1909	0.2302	0.1634
Non self-employed	0.0655	0.0762	0.0535	0.0969	0.0418	0.1768	0.0545	0.1124
All	0.1057	0.0683	0.0829	0.1306	0.0744	0.2390	0.0868	0.1413

Note: All positive variances are significantly different from zero at the 0.1% level. Two variances estimated to be negative have been set to zero.

Table 4: Regression of Permanent Aggregate Shock on Lagged Excess Stock Returns

	No High School		High School		College		All	
Industry	β	ρ	β	ρ	β	ρ	β	ρ
Agriculture	-	-	0.3333	0.3507	-	-	0.1754	0.2218
Manufacturing	0.1025	0.3381	0.0750	0.2786	0.0577	0.2057	0.0829*	0.3524
Construction	0.3197**	0.6577	0.0770	0.2423	-	-	0.1333**	0.5527
Transp., comm.	-0.1778	0.2983	0.0488	0.3234	-	-	0.0294	0.1879
Trade	0.1163	0.2128	-0.0050	0.0245	-0.0440	0.1095	0.0075	0.0283
Fin., RE, ins.	-	-	0.0360	0.0693	-0.0093	0.0224	0.0461	0.1281
Bus. serv.	-	-	-0.0156	0.0265	-	-	0.0804	0.1949
Prof. serv.	-	-	0.0320	0.0860	0.0622	0.3457	0.0479	0.2644
Public adm.	-	-	0.0634	0.2302	0.2195**	0.5094	0.1072**	0.4237
Self-employed	0.0916	0.0964	0.0753	0.1652	0.1869*	0.4164	0.1119	0.3159
Non self-employed	0.1022*	0.3970	0.0681*	0.4148	0.0462	0.2998	0.0710**	0.4638
All	0.0937	0.3170	0.0689*	0.3785	0.0721**	0.4943	0.0762**	0.4562

Note: ** Significant at the 5% level. * Significant at the 10% level.

Table 5: Regression of Permanent Aggregate Shock on Lagged Excess Bond Returns

	No High School		High School		College		All	
Industry	β	ρ	β	ρ	β	ρ	β	ρ
Agriculture	-	-	0.5750**	0.4244	-	-	0.5735**	0.5083
Manufacturing	0.2020**	0.4674	0.0750	0.2786	0.1050	0.2627	0.1221*	0.3637
Construction	0.1371	0.1977	0.1788*	0.3961	-	-	0.1378*	0.4005
Transp., comm.	-0.1723	0.2027	0.0632	0.2943	-	-	0.0485	0.2177
Trade	0.1822	0.2324	0.0835	0.2608	0.0012	0.0022	0.0754	0.1980
Fin., RE, ins.	-	-	0.1562	0.2074	0.1995	0.3283	0.1936*	0.3779
Bus. serv.	-	-	-0.2202	0.2638	-	-	-0.0941	0.1600
Prof. serv.	-	-	0.0199	0.0374	0.1319**	0.5138	0.1174**	0.4540
Public adm.	-	-	0.1539*	0.3926	0.1230	0.2000	0.1329*	0.3685
Self-employed	0.3653	0.2693	0.1912	0.2941	0.1391	0.2173	0.2124**	0.4204
Non self-employed	0.1223	0.3330	0.1021**	0.4361	0.1290**	0.5874	0.1152**	0.5279
All	0.1511*	0.3586	0.1168**	0.4501	0.1316**	0.6326	0.1300**	0.5457

Note: ** Significant at the 5% level. * Significant at the 10% level.

Table 6: Life Cycle Profiles.

Consumption									
	Baseline Case			$\gamma = 10$		$\delta = 0.8$		Self Employed	
Age	0/100	50/50*	50/50	0/100	50/50	0/100	50/50	0/100	50/50
20-35	20.22	20.26	21.07	20.13	20.88	20.53	20.61	25.09	26.10
36-50	25.48	26.11	26.15	25.12	24.86	26.50	26.47	38.39	38.28
51-65	24.61	26.06	25.46	24.23	24.54	23.94	23.78	35.23	35.35
66-80	22.43	26.64	24.19	22.65	24.46	15.95	15.73	32.67	34.02
81-100	16.98	25.61	18.27	19.04	20.60	14.27	14.27	27.26	28.70
Wealth									
	Baseline Case			$\gamma = 10$		$\delta = 0.8$		Self Employed	
Age	0/100	50/50*	50/50	0/100	50/50	0/100	50/50	0/100	50/50
20-35	5.94	5.77	6.39	8.20	7.92	3.39	2.73	12.84	13.75
36-50	29.34	23.17	35.87	39.28	50.57	7.25	5.64	65.75	81.99
51-65	75.77	40.34	96.26	100.16	126.78	10.23	7.83	173.70	195.02
66-80	77.28	26.06	92.81	105.50	128.71	5.71	4.71	159.76	169.73
81-100	13.60	4.15	18.40	30.85	39.94	0.11	0.11	46.75	52.00
Liquid Portfolio Share in Stocks									
	Baseline Case			$\gamma = 10$		$\delta = 0.8$		Self Employed	
Age	0/100	50/50*	50/50	0/100	50/50	0/100	50/50	0/100	50/50
20-35	1.00	0.93	1.00	0.97	0.97	0.99	0.99	0.57	0.53
36-50	0.99	1.00	0.98	0.95	0.72	1.00	1.00	0.91	0.68
51-65	0.88	0.81	0.61	0.61	0.08	1.00	0.90	0.57	0.14
66-80	0.90	0.96	0.85	0.57	0.49	1.00	1.00	0.54	0.51
81-100	0.92	0.58	0.90	0.68	0.57	1.00	1.00	0.61	0.53
Tax Rates									
	10.00%	10.00%	6.00%	10.00%	6.00%	10.00%	6.00%	10.50%	6.75%

* Refers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

Table 7: Welfare Analysis.

Baseline	Tax Rate	Welfare Age 20		Welfare Age 65		Gain Age 20		Gain Age 65	
		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.00%	18.67	1.75	22.69	3.42				
50/50*	10.00%	18.76	1.80	27.30	7.41	0.48%	2.86%	20.32%	116.67%
50/50	6.00%	19.36	1.81	24.46	6.22	3.70%	3.43%	7.80%	81.87%
100/0	3.50%	19.59	1.96	24.92	9.77	4.93%	12.00%	9.83%	185.67%
$\gamma = 10$	Tax Rate	Welfare Age 20		Welfare Age 65		Gain Age 20		Gain Age 65	
		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.00%	16.08	2.17	22.53	2.54				
50/50*	10.00%	16.05	2.19	27.54	6.62	-0.18%	0.95%	22.22%	160.81%
50/50	6.00%	16.78	2.26	24.42	5.68	4.35%	4.15%	8.39%	123.62%
$\gamma = 2$	Tax Rate	Welfare Age 20		Welfare Age 65		Gain Age 20		Gain Age 65	
		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.00%	20.71	1.17	22.03	3.27				
50/50*	10.00%	21.18	1.30	29.86	8.57	2.31%	11.15%	35.57%	161.92%
50/50	6.00%	21.50	1.33	23.11	6.25	3.83%	13.79%	4.91%	90.94%
Coll. S.E.	Tax Rate	Welfare Age 22		Welfare Age 65		Gain Age 22		Gain Age 65	
		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.50%	17.64	2.90	31.20	5.09				
50/50*	10.50%	17.63	2.89	37.96	10.13	-0.07%	-0.36%	21.66%	99.11%
50/50	6.75%	18.31	2.95	32.83	7.79	3.80%	1.72%	5.22%	53.05%

* Refers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

Table 8: Paternalistic Welfare Analysis.

δ Too Low	Tax Rate	Welfare Age 20		Welfare Age 65		Gain Age 20		Gain Age 65	
		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.00%	18.13	1.64	17.26	1.31				
50/50*	10.00%	18.50	1.77	25.33	8.05	2.04%	7.93%	46.76%	514.50%
50/50	6.00%	18.03	1.62	17.03	1.20	-0.55%	-1.22%	-1.33%	-8.40%
γ Too High	Tax Rate	Welfare Age 20		Welfare Age 65		Gain Age 20		Gain Age 65	
		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.00%	20.67	1.09	20.26	2.75				
50/50*	10.00%	21.02	1.22	30.59	7.34	1.69%	12.12%	50.96%	166.82%
50/50	6.00%	21.17	1.21	27.39	6.24	2.42%	11.01%	35.19%	126.91%

* Refers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

Table 9: Welfare Analysis with Alternative Mean Asset Returns.

Baseline Case	Tax Rate	Welfare Age 20		Welfare Age 65		Gain Age 20		Gain Age 65	
$\mu = 4\%$, $\overline{R}_F = 1.02$		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.00%	18.67	1.75	22.69	3.42				
50/50*	10.00%	18.76	1.80	27.30	7.41	0.48%	2.86%	20.32%	116.67%
50/50	6.00%	19.36	1.81	24.46	6.22	3.70%	3.43%	7.80%	81.87%
Alternative 1	**Tax Rate**	**Welfare Age 20**		**Welfare Age 65**		**Gain Age 20**		**Gain Age 65**	
$\mu = 3\%$, $\overline{R}_F = 1.02$		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	10.00%	18.57	1.70	22.15	3.53				
50/50*	10.00%	18.63	1.78	25.13	6.64	0.33%	4.43%	13.48%	87.89%
**						-0.22%	1.73%	10.75%	94.05%
50/50	7.25%	19.20	1.82	25.49	6.74	3.38%	6.69%	15.11%	90.74%
**						2.82%	3.93%	12.35%	97.00%
Alternative 2	**Tax Rate**	**Welfare Age 20**		**Welfare Age 65**		**Gain Age 20**		**Gain Age 65**	
$\mu = 3\%$, $\overline{R}_F = 1.03$		Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
0/100	7.00%	19.30	1.79	24.20	4.16				
50/50*	7.00%	19.32	1.84	26.76	6.88	0.12%	2.81%	10.60%	65.38%
**						3.50%	5.16%	17.96%	101.16%
50/50	4.75%	19.61	1.83	25.07	6.11	1.61%	2.23%	3.60%	46.88%
**						5.03%	4.57%	10.49%	78.65%

* Refers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

** Computes the welfare gain relative to the baseline case, i.e. under the initial assumption for asset returns of $\mu = 4\%$ and $\overline{R}_F = 1.02$.

Figure 1 - Labor Income Profiles
(age dummies and fitted polynomials)

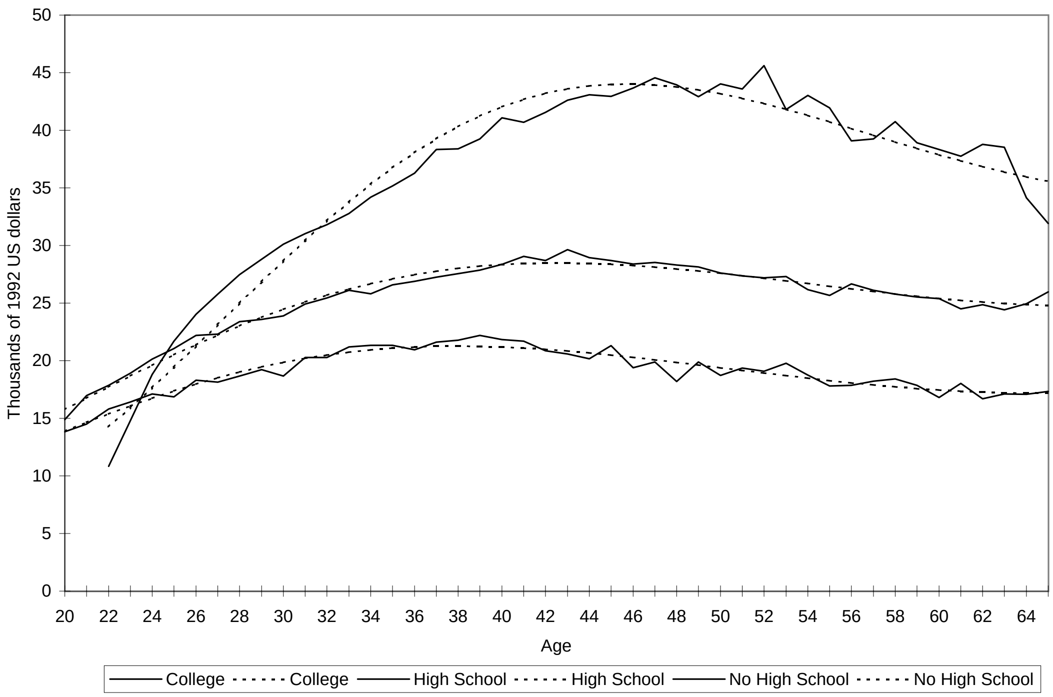

Figure 2a - Consumption, Income, Wealth and Annuity
(Retirement Wealth Fully Invested in the Safe Asset)

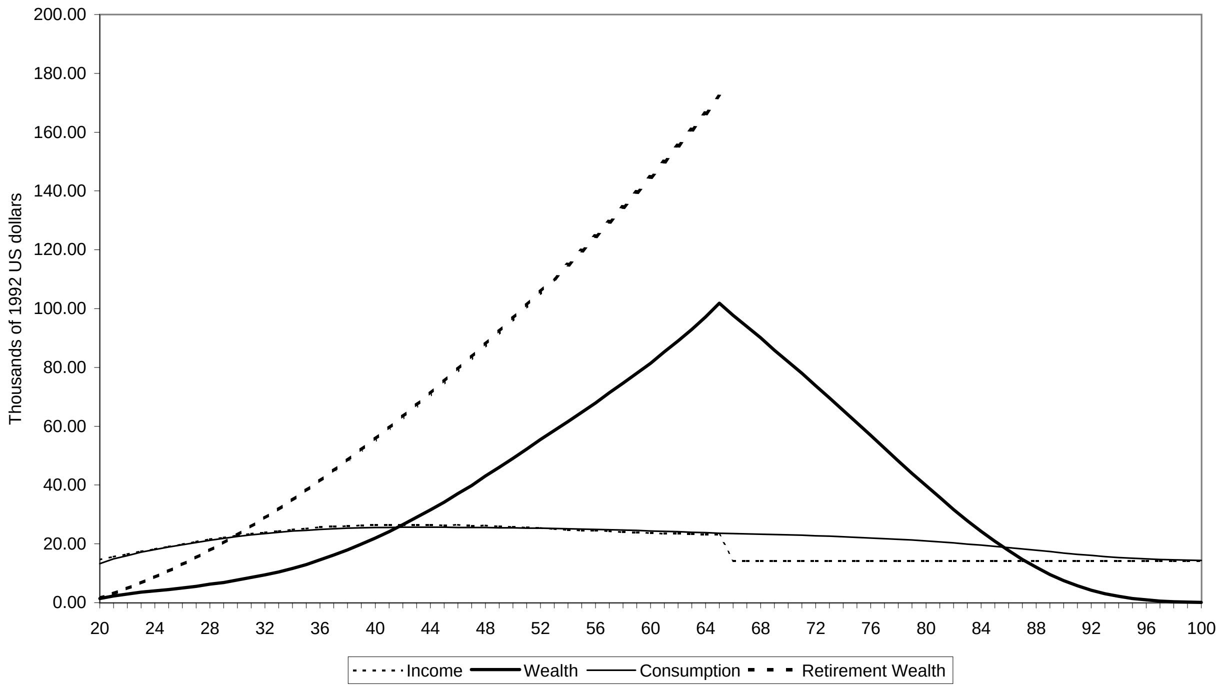

Figure 2b - Consumption, Income, Wealth and Annuity
(Retirement Wealth Invested 50/50 in Risky/Safe Asset)

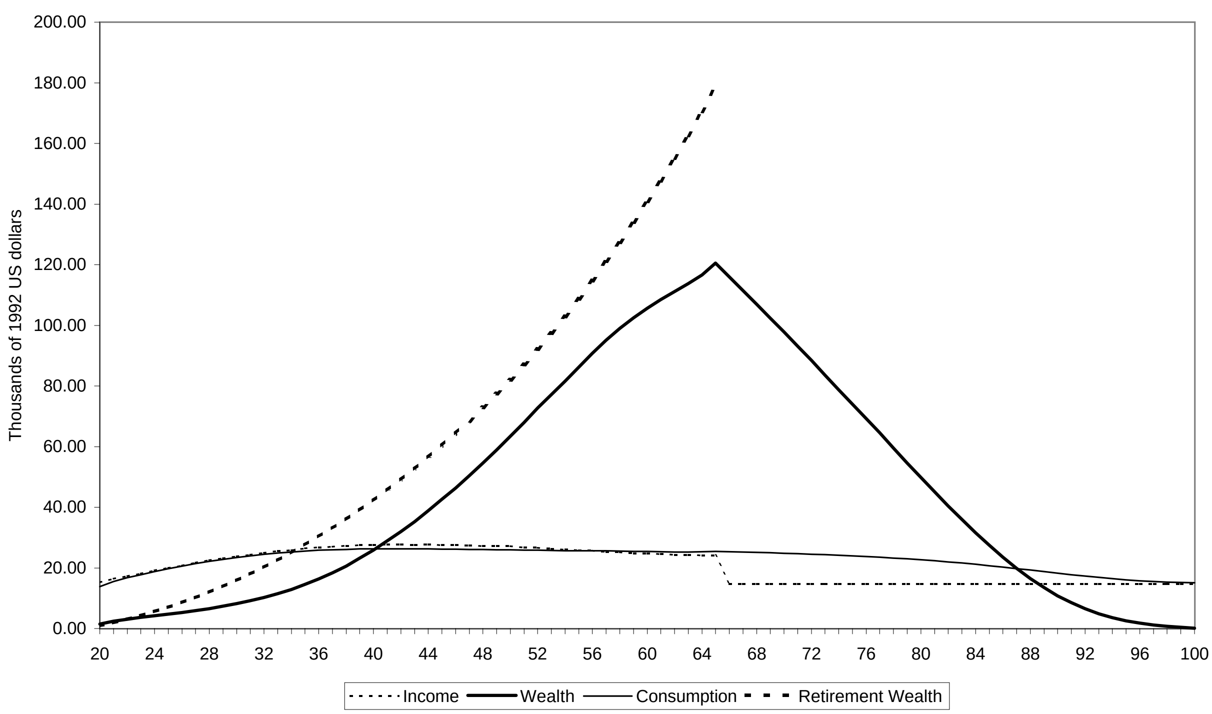

Figure 3a - Wealth, Stocks and T-Bills
(Retirement Wealth Fully Invested in the Safe Asset)

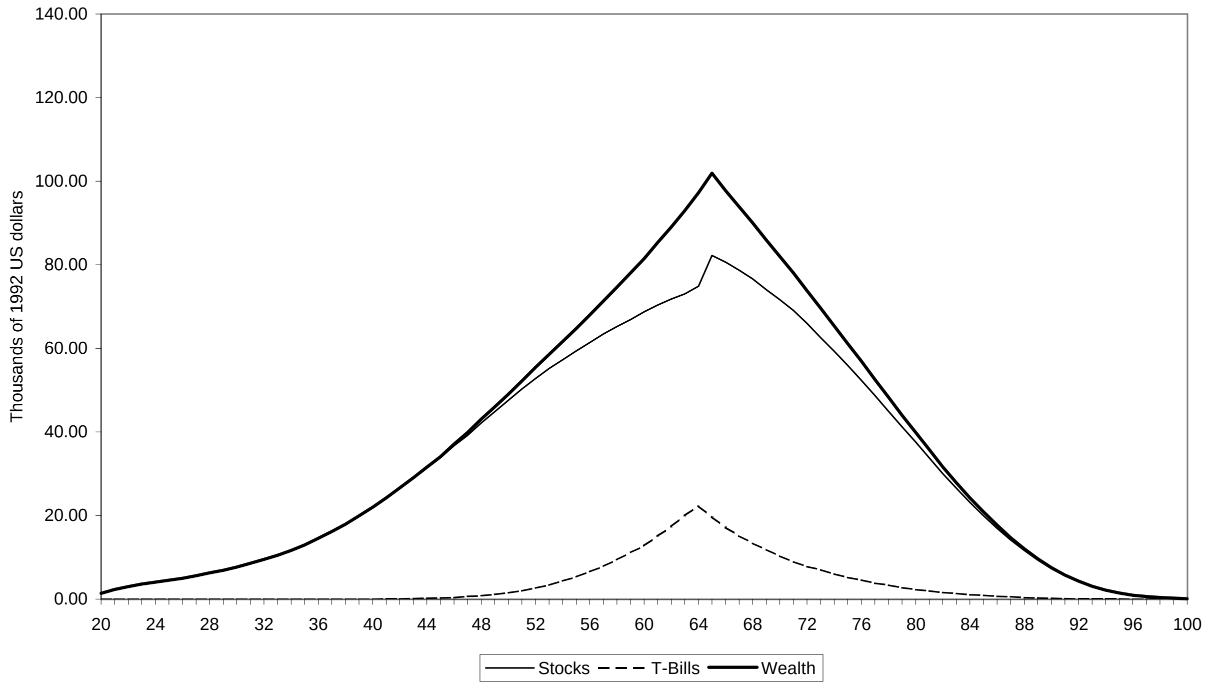

Figure 3b - Wealth, Stocks and T-Bills
(Retirement Wealth Invested 50/50 in Risky/Safe Asset)

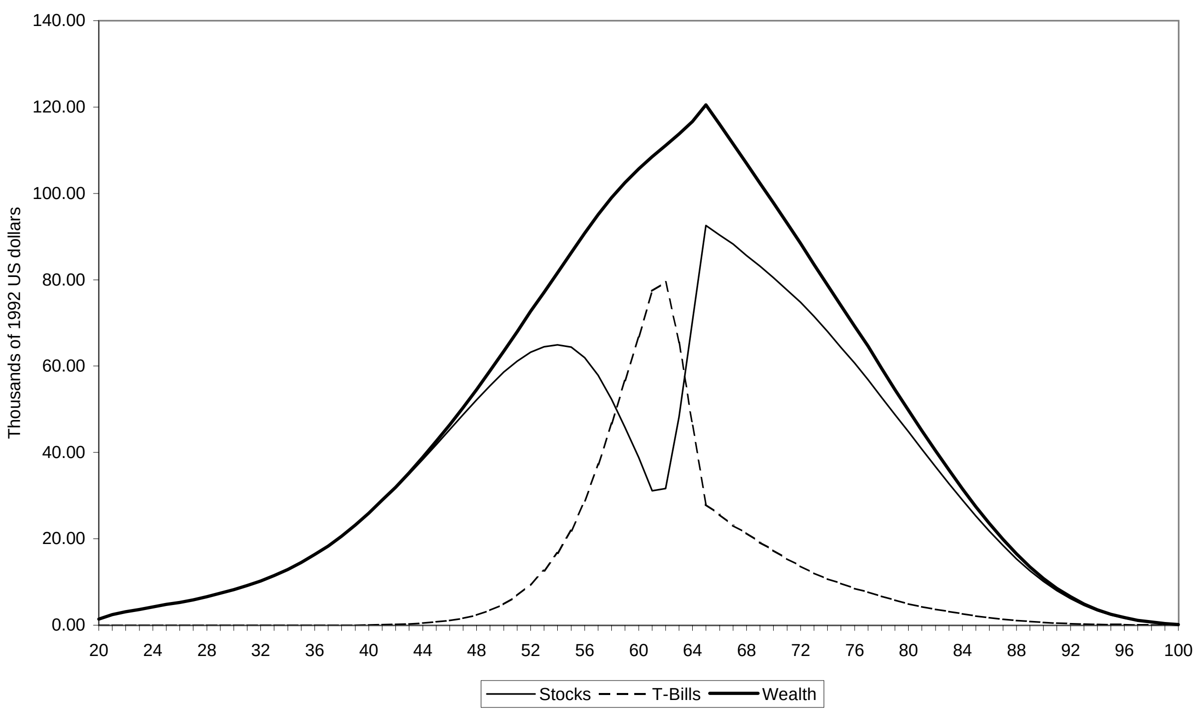

Figure 4a - Liquid Portfolio Share invested in Stocks
(Retirement Wealth Fully Invested in the Safe Asset)

Figure 4b - Liquid Portfolio Share invested in Stocks
(Retirement Wealth Invested 50/50 in Risky/Safe Asset)

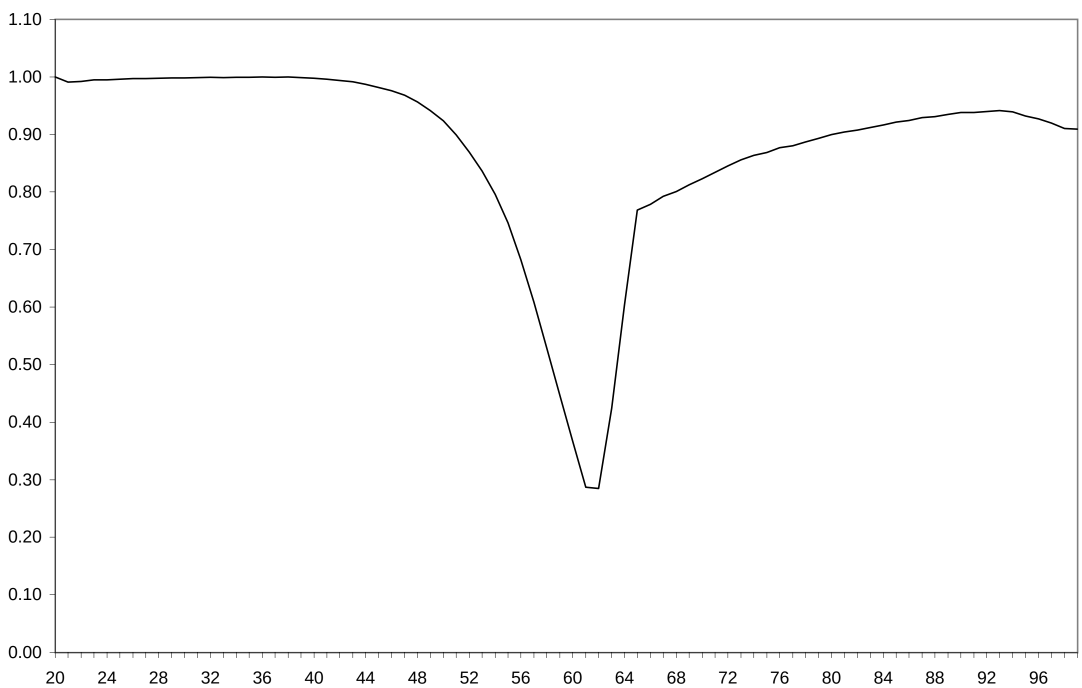

Fig 5a - Utility Gain

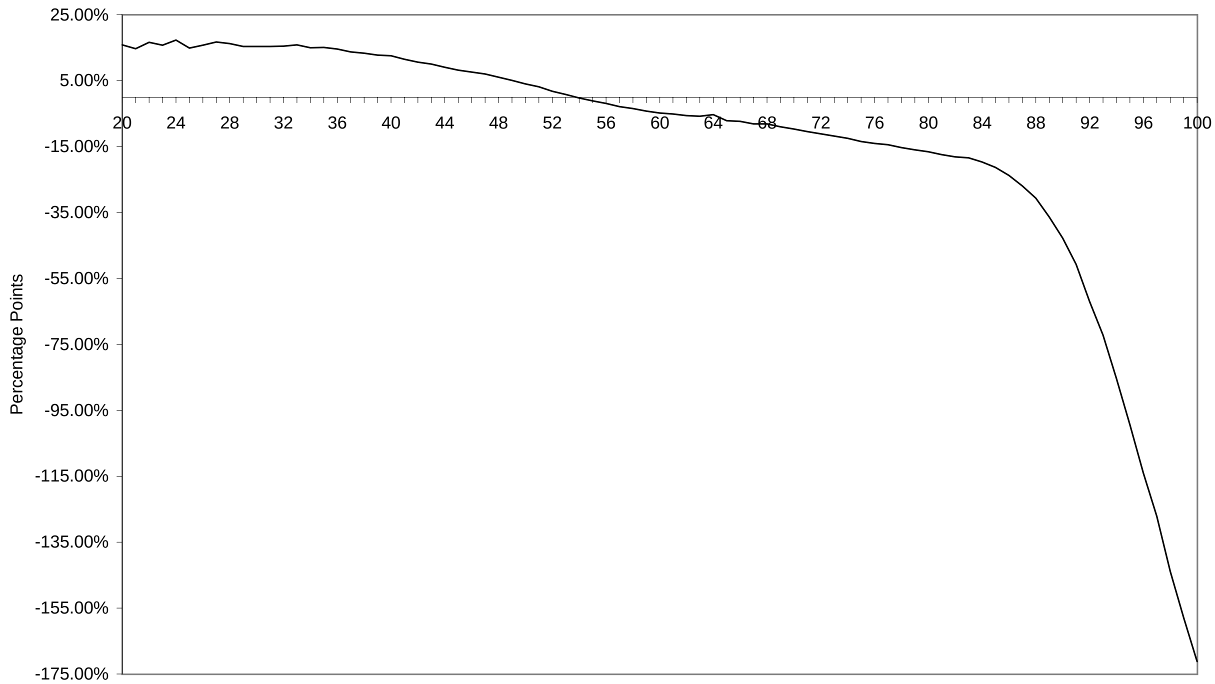

Figure 5b - Utility Dispersion
(Ratio of Standard Deviations)

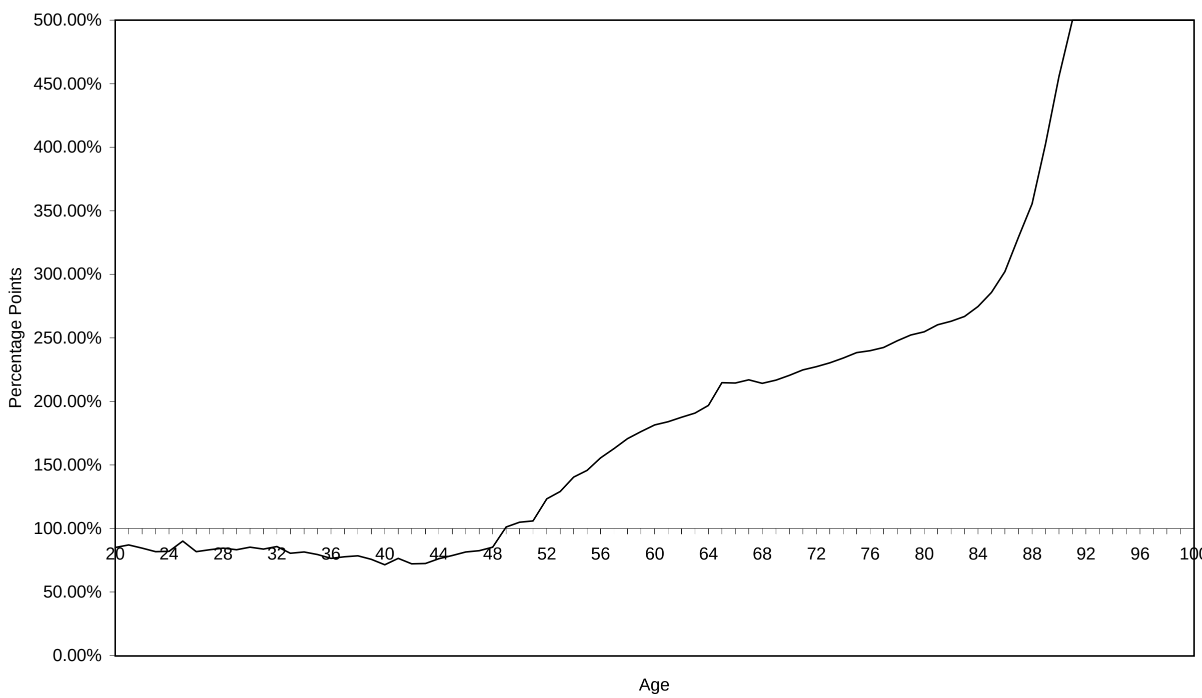

Figure 6a - Liquid Portfolio Share invested in Stocks and Participation Rate
(Retirement Wealth Fully Invested in the Safe Asset & Fixed Cost of Entering in the Stock Market)

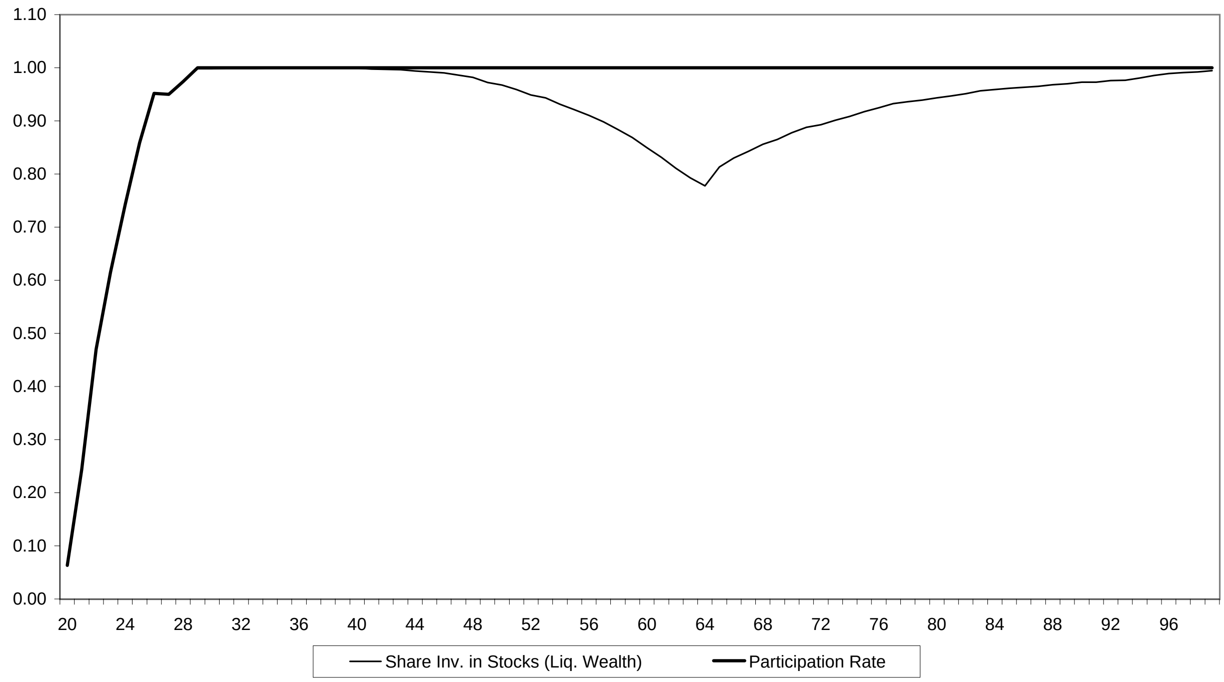

Figure 6b - Liquid Portfolio Share invested in Stocks
(Retirement Wealth Invested 50/50 in Risky/Safe Asset & Fixed Cost of Entering the Stock Market)