Consumption and Portfolio Choice over the Life-Cycle *

João F. Cocco

London Business School

Francisco J. Gomes

London Business School

Pascal J. Maenhout

INSEAD

*We would like to thank the Editor, John Heaton, two anonymous referees, Rui Albuquerque, Robert J. Barro, John Y. Campbell, Gary Chamberlain, Luigi Guiso, Per Krusell, David Laibson, Hanno Lustig, Alexander Michaelides, N. Gregory Mankiw, Jonathan Parker, James M. Poterba, Tony Smith and Luis Viceira for comments and discussions, and the participants of the NBER Summer Institute, EFA, ESEM and of seminars at Harvard University, Oxford University and the University of British Columbia for comments. Cocco gratefully acknowledges the financial support of the *Banco de Portugal*, Gomes of *Fundação para a Ciencia e Tecnologia*, Portugal, and Maenhout of the Fund for Scientific Research Flanders. Please send correspondence to Francisco Gomes, London Business School, Regent's Park, London NW1 4SA, United Kingdom, fgomes@london.edu. The usual disclaimer applies.

Consumption and Portfolio Choice over the Life-Cycle

Abstract: This paper solves a realistically calibrated life-cycle model of consumption and portfolio choice with non-tradable labor income and borrowing constraints. Since labor income substitutes for riskless asset holdings, the optimal share invested in equities is roughly decreasing over life. We compute a measure of the importance of human capital for investment behavior. We find that ignoring labor income generates large utility costs, while the cost of ignoring only its risk is an order of magnitude smaller, except when we allow for a disastrous labor income shock. Moreover, we study the implications of introducing endogenous borrowing constraints in this incomplete-markets setting.

The issue of portfolio choice over the life-cycle is encountered by every investor. Popular finance books (e.g. Malkiel, 1996) and financial counselors generally give the advice to shift the portfolio composition towards relatively safe assets, such as T-bills, and away from risky stocks as the investor grows older and reaches retirement. But what could be the economic justification for doing so?

A seminal reference addressing the problem of portfolio choice over the life-cycle is Samuelson (1969). The concept of 'businessman's risk' (i.e. holding risky stocks is only advisable for young businessmen, not for widows) is explored and rejected as invalid. However this conclusion is reached under the assumptions of independently and identically distributed returns and requires frictionless markets and the absence of labor income.

A crucial element one needs to consider when discussing portfolio choice over the life-cycle is labor income and the risk associated with it. To the extent that the level and risk of the labor income stream change over the life-cycle, and to the extent that portfolio choice depends on these factors, the presence of labor income can provide a rationale for age-varying investment strategies, without relying on predictability in asset returns. This is the route we explore in this paper.

Of course, if markets are complete so that labor income can be capitalized and its risk insured, the introduction of labor income is well understood analytically from the seminal work by Merton (1971). However, market incompleteness seems to be an important feature to consider when analyzing portfolio choice in a quantitatively focussed study. Because of moral hazard issues, many investors face borrowing constraints that prevent them from capitalizing future labor income. Moreover, explicit insurance markets for labor income risk are not well-developed so that many investors face uninsurable labor income risk.

We solve numerically for the optimal portfolio and savings decisions using a realistically and quantitatively calibrated model. We consider a finitely-lived investor facing mortality risk, borrowing and short-sale constraints, and receiving labor income. The labor income profile and its risk characteristics are estimated using the Panel Study of Income Dynamics (PSID), the largest longitudinal U.S. dataset containing careful information on labor income and individual control variables. The agent can invest her savings in two assets: a riskless and a risky asset, the return to which may be positively correlated with labor income shocks.[1] One can also think of our model as extending the recent consumption literature on buffer-stock saving (Deaton, 1991, Carroll, 1997 and Gourinchas and Parker, 2002) to include an asset allocation decision.

In order to understand the effects of labor income risk on portfolio allocation, it is important to have in mind that a labor income stream constitutes the implicit holding of an asset. We show that labor income acts as a substitute for riskfree asset holdings, if the correlation between labor income risk and stock market risk is set at the (insignificantly positive) value we estimated. These results confirm the earlier results obtained by Heaton and Lucas in an infinite-horizon setting (1997), and follow the intuition presented in Jagannathan and Kocherlakota (1996).

The shape of the labor income profile over life induces the investor in our model to reduce her proportional stock holdings when aging, and thus provides a rationale for the advice given in the popular financial literature. The implicit riskfree-asset holdings in the form of labor income lose importance as the investor ages, leading her to hold more riskfree assets explicitly, i.e. in her financial portfolio. All else equal, investors subject to more labor income risk hold a smaller share of their portfolio in equities so that labor income risk crowds out asset-holding risk.[2] Moreover, we find that allowing for an empirically calibrated (small) probability of a disastrous labor income draw has a large crowding out effect, particularly for young households, an important element for explaining the data.

We also examine the life-cycle portfolio implications of endogenous borrowing in our incomplete markets setting, with a realistic and empirically parameterized penalty for default.[3] We show that a crucial determinant of borrowing capacity and portfolio allocation is the lower bound for the income distribution. This highlights the importance of the extent to which social insurance bounds labor income away from zero. Investors with a bounded income process face a positive endogenous borrowing limit and, as a result, they hold negative wealth when young and do not invest in equities. To further enhance the realism of the model we also examine the portfolio implications of allowing for uncertainty in retirement income, a bequest motive, and recursive preferences.

In order to assess the importance of non-tradable labor income and its risk for portfolio decisions we compute the utility cost (measured in consumption equivalent units) associated with suboptimal portfolio decisions. In particular we compute the utility loss incurred by investors who ignore their labor income and invest a constant fraction of wealth in equities, as would be optimal in the complete-markets, no-labor-income situation. The losses are substantial, and up to 2% of annual consumption. However, the loss resulting from behavior that only ignores the risk features of the labor income stream, is an order of magnitude smaller than the penalty for ignoring labor income altogether, except when we allow for the possibility of a disastrous labor income realization.

There are several papers that study the effects of labor income risk on portfolio composition.[4] Heaton and Lucas (1997), Koo (1998) and Viceira (2001) consider infinite-horizon models of portfolio choice with uninsurable labor income risk. Viceira (2001) captures retirement effects through a constant probability of zero labor income forever. By their stationary nature, infinite-horizon models are less suited to address life-cycle issues. More precisely, one of our findings is that an important determinant of portfolio composition is the ratio of accumulated wealth to expected future labor income, which is clearly not stationary over the life-cycle.

Bertaut and Haliassos (1997), Davis and Willen (2000) and Gakidis (1999) consider finite horizon models and to this extent are closer to our paper. Dammon, Spatt and Zhang (2001) introduce taxes in this framework, while Constantinides, Donaldson and Mehra (2002) and Storesletten, Telmer and Yaron (2000) explore the asset pricing implications of these models. The papers by Cocco (2004), Hu (2002), and Yao and Zhang (2004) study the implications of introducing housing in a life-cycle model.

The main contribution of our paper is to solve a realistically calibrated life-cycle model of consumption and portfolio choice with uninsurable labor income risk, which allows us to obtain a measure of the importance of market-incompleteness and labor income risk for investment behavior. Moreover, we quantify the utility cost associated with alternative portfolio rules, for realistic heterogeneity in investors.

The rest of the paper is organized as follows. Section 2 discusses the model's assumptions and set-up. The calibration and parametrization is presented in section 3. Section 4 looks at the solution of our model in terms of the optimal portfolio and consumption rules. Section 5 gives the simulation results for the benchmark parametrization, explores the effects of heterogeneity in labor income and preferences, and considers a number of extensions of the benchmark model, whose empirical predictions are then compared to the data. The utility cost computations for several alternative investment strategies are reported in section 6. Finally, section 7 concludes.

1 The Model

1.1 Model Specification

1.1.1 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 uncertainty in T in the manner of Hubbard, Skinner and Zeldes (1995). Let p_t denote the probability that the investor is alive at date $t+1$, conditional on being alive at date t. Investor i's preferences are described by the time-separable power utility function:

$$E_1 \sum_{t=1}^{T} \delta^{t-1} \left(\prod_{j=0}^{t-2} p_j \right) \left\{ p_{t-1} \frac{C_{it}^{1-\gamma}}{1-\gamma} + b \left(1 - p_{t-1}\right) \frac{D_{it}^{1-\gamma}}{1-\gamma} \right\}, \tag{1}$$

where $\delta < 1$ is the discount factor, C_{it} is the level of date t consumption, $\gamma > 0$ is the coefficient of relative risk aversion, and D_{it} is the amount of wealth the investor bequeaths to his descendants at death. For simplicity, we assume that the utility function applied to the bequest is identical to the utility function applied to the investor's own consumption when alive. The parameter b controls the intensity of the bequest motive and is set to zero in the benchmark analysis. In section 5.2, we make the bequest motive operational to investigate the sensitivity of our benchmark results. Later on we relax the assumption of time-additive utility and consider Epstein-Zin preferences.

1.1.2 The labor income process

Before retirement, 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 of a vector of 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 before retirement, 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 transitory and one persistent. The process for v_t is taken to be a random walk, following Carroll (1997) and Gourinchas and Parker (2002). Hubbard, Skinner and Zeldes (1995) estimate a general first-order autoregressive process and find the autocorrelation coefficient to be very close to one. 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 (distributed as $N(0, \sigma_\xi^2)$) and an idiosyncratic component ω_{it} (distributed as $N(0, \sigma_\omega^2)$):

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

This decomposition implies that the random component of aggregate labor income follows a random walk, an assumption made in the finance literature by Fama and Schwert (1977) and Jagannathan and Wang (1996). At the same time, empirical time series for individual labor income exhibit less persistence and this is achieved by adding the idiosyncratic transitory shock ε_{it}. Finally, we will allow for correlation between innovations to excess stock returns and labor income shocks through the aggregate component ξ_t, as will be described in the next section.

The assumption that labor income is exogenous is made primarily for simplicity. In reality individuals must decide how many hours to work and how much effort to put on the job, decisions that will influence the amount of labor income received. In particular, by having exogenous labor income we rule out the possibility that an individual who has had a bad portfolio return (or labor income) realization works more hours to compensate for it.[5]

Retirement income is modeled as a constant fraction λ of permanent labor income in the last working-year:[6]

$$\log(Y_{it}) = \log(\lambda) + f(K, Z_{iK}) + v_{iK} \qquad \text{for } t > K \tag{5}$$

Although oversimplified, this specification considerably facilitates the solution of the model, as it does not require the introduction of an additional state variable.

1.1.3 Financial assets

We assume that there are two assets in which the agent can invest, a riskless and a risky assets. The riskless asset, which we call *Treasury bills*, has a constant gross real return of $\overline{R}_f$. We denote the dollar amount of T-bills the investor has at time t by B_{it}. The risky asset has a gross real return R_t, and its excess return is given by:

$$R_{t+1} - \overline{R}_f = \mu + \eta_{t+1}, \tag{6}$$

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 ρ. We call the risky asset *stocks* and denote the dollar amount the investor has in stocks at time t by S_{it}.

We assume that the investor faces the following borrowing and short-sales constraints:

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

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

The borrowing constraint (7) ensures that the investor's allocation to bills is non-negative at all dates. It prevents the investor from capitalizing or 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. If we let α_{it} denote the proportion of savings invested in stocks at time t, then constraints (7) and (8) imply that $\alpha_{it} \in [0, 1]$ and wealth is non-negative.

These constraints can be motivated using the standard moral hazard and adverse selection arguments. It is straightforward to allow for a negative limit in (7) or (8). What is important for our results is that the individual is to some extent liquidity constrained in the early years of his adult life. We believe that this is the case for most households. In section 5, we extend the analysis and consider endogenous borrowing constraints.

1.2 The investor's optimization problem

In each period t the timing of the events is as follows. The investor starts the period with wealth W_{it}. Then labor income Y_{it} is realized. Following Deaton (1991) we denote *cash-on-hand* in period t by $X_{it} = W_{it} + Y_{it}$. We will also refer to X_{it} as wealth: it is understood that this includes labor income earned in period t. Then the investor must decide how much to consume, C_{it}, and how to allocate the remaining cash-on-hand (savings) between stocks and T-bills. Next period wealth, before earning period $t+1$'s labor income, is then given by:

$$W_{i,t+1} = R^p_{i,t+1}(W_{it} + Y_{it} - C_{it}), \tag{9}$$

where R^p_{t+1} is the return on the portfolio held from period t to period $t+1$:

$$R^p_{i,t+1} \equiv \alpha_{it} R_{t+1} + (1 - \alpha_{it})\overline{R}_f. \tag{10}$$

The problem the investor faces is to maximize (1) subject to constraints (2) through (8), in addition to the non-negativity constraint on consumption. The control variables of the problem are $\{C_{it},\ \alpha_{it}\}_{t=1}^{T}$. The state variables are $\{t,\ X_{it}, v_{it}\}_{t=1}^{T}$. The problem is to solve for the policy rules as a function of the state variables, i.e., $C_{it}(X_{it}, v_{it})$ and $\alpha_{it}(X_{it}, v_{it})$. Given the set-up we assumed, the value function is homogeneous with respect to current permanent labor income.[7] Exploiting this scaleability allows us to normalize v_{it} to one and to reduce the dimensionality of the state space.

The Bellman equation for this problem is given by:

$$\begin{aligned} V_{it}\left(X_{it}\right) &= \underset{C_{it}\geq 0, 0\leq \alpha_{it}\leq 1}{Max} \left[U(C_{it}) + \delta p_t E_t V_{i,t+1}(X_{i,t+1})\right] \text{ for } t < T \\ \text{where } X_{i,t+1} &= Y_{i,t+1} + \left(X_{it} - C_{it}\right)\left(\alpha_{it} R_{t+1} + (1 - \alpha_{it})\overline{R}_f\right) \end{aligned}$$

The problem cannot be solved analytically. We derive the policy functions numerically by using backward induction. In the last period the policy functions are trivial (the agent consumes all available wealth) and the value function corresponds to the indirect utility function. We can now substitute this value function in the Bellman equation and compute the policy rules for the previous period. We optimize using grid search and we discretize the state-space for the continuous state variable (cash-on-hand). Given these policy functions we can obtain the corresponding value function and this procedure is then iterated backwards until $t = 1$. More details on the numerical solution technique are given in appendix A.

2 Calibration

2.1 Labor Income Process

We used the PSID to estimate equations (2) and (3) which give labor income as a function of age and other characteristics. In this section we will give a brief description of the sample selection and the estimation method. More details are given in appendix B.

We took a broad definition of labor income so as to implicitly allow for (potentially endogenous) ways of self-insuring against pure labor income risk. Therefore we defined 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.

The estimation controls for family-specific fixed effects. To control for education the sample was split in three groups: the observations without high school education, a second group with high school education but without a college degree, and finally college graduates. The reason for doing this is the well-established finding that age-profiles differ in shape across education groups (see e.g. Attanasio, 1995 and Hubbard, Skinner and Zeldes, 1995). For each education group we assume that the function $f(t, Z_{it})$ is 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 size.[8]

Table 1 and Figure 1 report the results for the three education groups. The coefficients of the age dummies are clearly significant and the results match intuition and stylized facts (see Attanasio (1995), Gourinchas and Parker (2002) and Hubbard, Skinner and Zeldes (1995)). We fitted a third-order polynomial to the age dummies to obtain the profiles for the numerical solution (see Table 2 and Figure 1). The replacement ratio λ used to determine the amount of retirement income, was calibrated as the ratio of the average of our labor income variable defined above for retirees in a given education group to the average of labor income in the last working-year prior to retirement. The result is also reported in Table 2.

We estimate the error structure of the labor income process by following closely the variance decomposition method described by Carroll and Samwick (1997). We use a similar procedure to estimate the correlation between labor income shocks and stock returns, ρ. The results are reported in Table 3 and (as mentioned above) the details are given in appendix B.

2.2 Other Parameters

Table 4 reports our benchmark parameter values. 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 10. This is the upper bound for risk aversion considered reasonable by Mehra and Prescott (1985). We also consider lower values. The mean equity premium μ is 4.00%.[9] The risk-free rate is 2.00% and the standard deviation of innovations to the risky asset is set to its historical value of 0.157.

3 Policy Functions

Before looking at simulated life-cycle paths for consumption and portfolio choice, we briefly discuss the policy functions underlying these results. This allows us to highlight the main forces at work and to gain intuition on the determinants of consumption and portfolio choice. All the results presented in this section are for the benchmark case of the second education group (see Table 4 for a summary of the parameters). The policy functions behave in a similar manner for the other parameterizations we consider, unless explicitly stated.

Let us start with the portfolio rules. In the complete-markets setting and ignoring labor income, the optimal portfolio rule for an investor with power utility facing a constant investment opportunity set is straightforward. As Samuelson (1969) and Merton (1969) showed, the optimal fraction of wealth invested in the risky asset is constant, independent of wealth and age, and depends only on risk aversion and the moments of the asset's excess return:

$$\alpha = \frac{\mu}{\gamma \sigma_\eta^2} \tag{11}$$

In a realistic life-cycle setting however, risky labor income cannot be capitalized, and the portfolio rule is a function of the relevant state variables: financial wealth (cash-on-hand)

and age. We plot the optimal fraction of the portfolio invested into the risky asset ($\alpha(X,t)$), as a function of cash-on-hand X and for a given age t. Then we consider how this rule shifts as the agent ages.

3.1 Portfolio Rules during Retirement

It is easiest to consider first the retirement stage. In this phase of the life-cycle, we model 'labor income' as being constant and certain. Figure $2A$ shows the optimal portfolio rule for year $T-1$. This policy function is presented first because it solves a simple two-period problem, for which the intuition is well explained in Jagannathan and Kocherlakota (1996). The optimal portfolio rule is decreasing in wealth.[10] This can be understood as follows. During the last period, the investor receives a nonrandom amount of labor income. Clearly, this future retirement income acts as a substitute for risk-free asset holdings and induces the investor to hold more stocks. The agent with little wealth will then tilt her financial portfolio more aggressively towards equities than the agent with a large amount of financial wealth, simply because the poorer investor already has a *relatively* larger risk-free asset position from her retirement income. For very high values of cash-on-hand, retirement income becomes trivial and the proportional demand for stocks asymptotes to the complete-markets solution given by equation (11), and shown by the straight line in Figure $2A$.

Similar reasoning explains the behavior of the optimal investment strategy as a function of age during retirement. It is useful to recall Merton's solution (1971) for the case when labor income is constant and riskless, markets are complete and time is continuous. Denoting the present value of a constant labor income stream by $PDV_t(FY_t)$, Merton's result can be rewritten as:

$$\frac{\alpha_t W_t}{W_t + PDV_t(FY_t)} = \frac{\mu}{\gamma\sigma_\eta^2} \tag{12}$$

Thus the fraction of *total wealth*, consisting of financial wealth W_t and of human capital $PDV_t(FY_t)$, invested in the risky asset, equals the familiar ratio $\frac{\mu}{\gamma\sigma_\eta^2}$. For a given value of W_t, as the investor ages, the amount of future retirement income (and of the risk-free asset holdings implicit in it) decreases, and for a given level of cash-on-hand the agent therefore holds a larger proportion of her financial portfolio in the riskless asset. This implies that the portfolio rule shifts inwards as one ages in retirement.

3.2 Portfolio Rules before Retirement

Before retirement, when the labor income stream is stochastic, three interesting lessons can be learned from the analysis of the portfolio rules. First, although it is not obvious that the risky labor income stream still mimics the payoff of a riskless asset more closely than the one of the risky asset, Figure $2B$ shows that this is the case since the policy function is still decreasing in cash-on-hand. This happens because the labor income stream is not highly correlated with the innovations to stock returns.[11]

Second, with respect to age effects, Figure $2B$ shows that the portfolio rules still become less aggressive as the middle-aged agent grows older. In addition to the obvious fall in the present value of future labor income due to the shortening of the income stream (as in the retirement phase of the model), the capitalized value of labor income also drops with age because of the negative slope of the labor income profile during this part of the life-cycle.

A third interesting finding is that the steepness of the labor income stream early in life leads the agent to *increase* her demand for the risky asset as she ages, for a given amount of financial wealth. Graphically, the optimal investment strategy shifts out with age in Figure $2B$, leading to a more aggressive relative equity position. This result is remarkable because it implies that present discounted value of future labor income, $PDV_t(FY_t)$, must rise, not fall with age. What drives this is the fact that the earliest years are characterized by very low earnings and high earnings growth.

3.3 Consumption Decisions

Another important ingredient to our understanding of the simulation results in the next section is the optimal consumption rule. Because the optimal portfolio weight in equities depends on the state variable wealth, the consumption-savings decision will determine where the portfolio rules identified above are evaluated. The consumption function, giving optimal consumption as a function of current cash-on-hand, is concave as derived analytically by Carroll and Kimball (1996). In the first phase of the life-cycle (roughly until age 35 to 40, see Figure $2C$), the consumption function shifts upward as the agent ages. The reason is that her permanent income increases during this part of the life-cycle, due to the steep slope of the labor income profile. As households approach retirement and as their labor income profile becomes negatively sloped this pattern is reversed in Figure $2C$. The policy functions are similar to the ones in Gourinchas and Parker (2002): their results are therefore robust to the introduction of a portfolio decision.

4 Simulation Results

Using the policy functions derived above, we simulated the consumption and asset allocation profiles of 10000 agents over the life-cycle. Below we present and discuss the cross-sectional means of these simulated profiles. We start with the benchmark case for which we discussed the policy functions. Then we analyze the importance of heterogeneity in human capital and preferences for these benchmark results.

4.1 Benchmark Case

In Figure 3A the simulated income, wealth and consumption profiles are plotted. We see that households are liquidity constrained during, roughly, the first 15 years of their working lives. Consumption tracks income very closely and a small level of savings (around 6 months of labor income during the first decade) is accumulated to use as insurance-cushion against negative labor income shock. As labor income increases and this profile becomes less steep the agent starts accumulating wealth for retirement. The consumption profile ceases to be increasing as the agent gets older, reflecting the fact that the liquidity constraint becomes less binding. Finally, during retirement effective impatience increases due to mortality risk and the consumption path slopes down, while wealth is decumulated at a fast rate. The standard hump-shaped consumption profile emerges.

In Figure 3C we present the mean simulated portfolio allocation. Early in life, most agents invest fully in stocks and hit the borrowing constraint. Only in the very first years of the life-cycle do some investors choose to hold the riskless asset. This is easily explained from the behavior of the policy functions presented above: the very steep labor income profile shifts out the portfolio rule because the implicit riskless asset holdings represented by labor income increase rapidly initially (see Figure 3B which plots the ratio of the present discounted value of future labor income to cash on hand). In midlife, saving for retirement becomes a crucial determinant of the agent's behavior. The downward-sloping portfolio rule is then evaluated at higher values of wealth so that the investor tilts her portfolio towards the riskfree asset. Finally, during retirement the portfolio rule shifts in. At the same time wealth is run down very quickly. The net effect is a slight increase in the optimal stock holdings due to the rapid pace at which the old agent decumulates wealth, motivated by mortality-enhanced effective impatience. We examine whether this particular result is robust to the introduction of additional uncertainty during retirement (e.g. stochastic medical expenses along the lines of Hubbard, Skinner and Zeldes (1995)), since this may slow down the pace at which wealth is being depleted.

Figure 3C also plots the 5th and 95th percentile of α. As these percentiles show, the model is able to generate substantial heterogeneity in α from mid-life onwards, but not early in life when virtually every investor is fully invested in equities. In section 5.2 we consider some extensions that generate more heterogeneity, also early in life, and thus make the model empirically more appealing.

4.2 Heterogeneity and Sensitivity Analysis

Not only is human capital a crucial asset for many investors, it is also characterized by substantial heterogeneity across investors. In particular, differences in the properties of labor income and retirement income might have important implications for the optimal investment strategy. Different agents work in different sectors of the economy, and are therefore exposed to different amounts and different sorts of labor income risk. Also, as the estimation in section 3 shows, the shape of the income profile depends significantly on educational attainment. In this section we solve for the optimal portfolio and consumption rules for some of these cases in order to explore the importance of labor income heterogeneity for optimal investment strategies. Finally, we also analyze the sensitivity of our results to some crucial assumptions regarding preferences (bequest motive, risk aversion, time-separability) and financial markets (equity premium and endogenous borrowing constraints).

4.2.1 Labor Income Risk

Depending on the sector or industry that the investor works in, the risk aspects of her labor income might differ substantially from the benchmark case analyzed before. Parameters of interest are the variance of the temporary and permanent shocks to labor income, σ_ε^2 and σ_u^2 respectively, and the correlation between the permanent shocks to labor income and the innovations to excess returns, ρ. To illustrate the effects, we will focus on some extreme cases as identified in Campbell, Cocco, Gomes and Maenhout (1999). In particular, we solve the dynamic program for investors in industries with large standard deviations for permanent and temporary income shocks, i.e. with a lot of career and layoff risk respectively. Construction and especially Agriculture are sectors that yield large estimates for σ_ε^2 and σ_u^2, but interestingly for both cases the ratio $\frac{\sigma_u^2}{\sigma_\varepsilon^2+\sigma_u^2}$ is very similar to the one used in the benchmark analysis. Public Administration on the other hand is characterized by a large value for $\frac{\sigma_u^2}{\sigma_\varepsilon^2+\sigma_u^2}$, but a total variance of shocks to labor income ($\sigma_\varepsilon^2+\sigma_u^2$) that is about half the benchmark estimate. The parameters used in this section are summarized in Table 5. Apart from these realistic parameterizations, we also solve the consumption-portfolio problem for a hypothetical investor subject to zero labor income risk. This exercise demonstrates the effect of 'normal' labor income risk, i.e. as faced by the typical highschool graduate.

Let us start with the problem for the investor facing no labor income risk. Relative to the benchmark, we expect two effects. First of all, the lack of any labor income risk eliminates the precautionary savings motive. Secondly, the intuition that labor income risk crowds out portfolio risk suggests that the fraction of the portfolio allocated to the risky asset should increase, for a given amount of financial wealth. It can be seen from Figure $4A$ that both effects are at work. Until age 33, the investor saves nothing so that α is not defined. After that, we find that the agent invests significantly more in the risky asset.

Investors in Agriculture on the other hand, never choose to invest 100% in stocks. The outward shift in the portfolio rule early in life, discussed in the previous section, is now much stronger and results in a roughly hump-shaped portfolio profile. One might argue that the case of agricultural workers with risk aversion of 10 is somewhat extreme. In the same figure we plot the results for $\gamma = 3$ for which the background risk effect is significantly reduced. The agents in Construction and especially Public Administration have portfolio profiles that are very similar to the one obtained for the benchmark calibration. The difference in risk characteristics of their labor income is simply too small to yield any substantial effects.[12]

The empirical evidence on the value of the correlation between labor income innovations and equity returns is mixed. Indeed Davis and Willen (2000) find considerably higher correlations between labor income innovations and broad measures of equity. Heaton and Lucas (1999a) also report positive correlation for entrepreneurs. We explore the implications of positive correlation in Figure $4B$. The portfolio effects are significant, particularly early in life: the benefits of investing in equities are now lower so that the average α is lower. However, the benefits of saving (under the form of equities) are also lower so that the investor saves less and accumulates less wealth. This explains the somewhat higher allocation to equities in mid-life.

4.2.2 Disastrous Labor Income Shocks

Carroll's (1997) version of the buffer-stock model of savings explicitly allows for the possibility of a disastrous labor income shock. In particular, labor income is modeled as being zero with some probability and following (2) otherwise. This makes labor income substantially more risky and may constitute a powerful source of background risk affecting portfolio choice. Given the relatively moderate effects of labor income risk obtained so far in the benchmark analysis (and elsewhere in the literature, with the exception of entrepreneurial risk as shown by Heaton and Lucas (2000a)), allowing for the possibility of a disastrous labor income draw is an important robustness check.

In the estimation of the labor income process above we deliberately dropped zero income observations, since our measure of labor income is broadly defined and includes unemployment compensation, welfare and transfers, and since remaining zero income observations may well be due to measurement error. In our dataset the frequency of zero-income observations is equal to 0.495 percent. Therefore, we consider a 0.5 percent annual probability of a zero labor income draw.

Figure 5 shows that a 0.5% probability of a zero labor income draw dramatically lowers the optimal equity share. The qualitative features of our benchmark results are preserved: very young and old investors choose less equity exposure than midlife investors. Quantitatively especially the first twenty years are affected and the average optimal α even drops below 60% during the first ten years. These results are similar to what we obtained for investors working in Agriculture as presented in Figure 4A. The same effects are at work: due to substantial background risk stemming from labor income, optimal portfolio rules shift in (as discussed in the previous section) and simultaneously the investor accumulates more wealth for precautionary reasons. Both effects result in lower portfolio shares in stocks. Older investors are less affected, since the wealth they accumulate for retirement purposes in the benchmark scenario also serves as a buffer to smooth out temporary drops in labor income.

The 5th and 95th percentile of α plotted in Figure 5 show that allowing for the possibility of a disastrous labor income draw, yields far more heterogeneity in asset allocation than in the benchmark case, especially for young investors.[13]

Labor income risk has the potential to constitute a potent source of background risk, even in the absence of any correlation between adverse income shocks and stock market performance. Whether this explanation for cautious investment behavior is empirically relevant remains a challenging question, especially if one subscribes to the perspective of a 'Peso-problem': investors could act according to a nontrivial probability of a disastrous idiosyncratic income shock even when, empirically, these seem to occur with extremely small probability.

4.2.3 Uncertain Retirement Income

In the benchmark model, retirement income is a constant fraction (depending on educational attainment) of the permanent component of labor income in the last year of working life. As an extension we now allow for more uncertainty about retirement income, in the following two ways. First, we consider an investor whose retirement income is stochastic and correlated with the contemporaneous performance of the stock market. Second, we go back to the benchmark case where retirement income depends on labor income immediately prior to retiring, but allow for a disastrous retirement income draw (at 25% of the normal level). This is a convenient way of modeling the existence of medical and health-related expenses which are indispensable and needed to keep up utility.

In the first case, we make retirement income uncertain also after retirement by adding a stochastic component to the retirement income used in the benchmark analysis. The retirement income shock is assumed to be transitory and follows the same process as during working life (ε_{it} in (2)), except that it is now correlated with stock return innovations (correlation of 0.2). Making retirement income stochastic and correlated with the contemporaneous performance of the stock market is aimed at capturing the situation of (wealthy) investors and entrepreneurs who, during retirement, receive income from proprietary businesses which they do not wish to sell. The portfolio allocation (not reported) is almost identical to the benchmark scenario. The reason for this is twofold. First, as argued before, older investors are less vulnerable to background risks since they control a buffer of wealth that is available for insurance purposes. Second, unlike in section 5.2.1, the correlated income shocks are transitory, not permanent (since we want to capture retirement income risk due to stock-market-like risk, not labor income risk, and since stock returns are assumed i.i.d.).

In the second extension, we go back to the benchmark case where retirement income is certain (once retired), but allow for a disastrous retirement income draw at 25% of the mean level (with 0.5% probability). As was done for labor income during working life, this introduces significant background risk during retirement and could reveal the sensitivity of our benchmark results to the assumption of non-random retirement income. Alternatively, the low retirement income state can be thought of as resulting from an extreme health shock which necessitates medical expenses that are equal to 75% of normal retirement income and needed to maintain a certain utility level.[14] The results are shown in Figure 6. Not surprisingly, the presence of significant health shocks during retirement do not affect the young investor, who is liquidity-constrained and impatient. As retirement approaches, the portfolio share drops relative to the benchmark without health shocks due to additional precautionary savings. Upon retirement, the health shocks become operational and require a slightly larger contingency fund than otherwise. This leads to a more conservative equity share.

4.2.4 Endogenous Borrowing Constraints

So far, we have exogenously imposed tight borrowing constraints. In practice however, households do borrow.[15] In this section we present an extension of our life-cycle model that allows for some endogenously determined amount of borrowing and takes into account a variety of real-world imperfections in credit markets. These imperfections concern both the price and quantity of credit available to investors. While there are alternative ways of introducing borrowing in the model, this section builds on insights from new work on credit-market imperfections.[16] A recent literature in equilibrium asset pricing has successfully examined the implications of endogenous borrowing constraints in a complete-market setting (e.g. Alvarez and Jermann (2000) and Lustig (2001)). This section can therefore also be viewed as a study of the portfolio implications of these constraints in an incomplete-market life-cycle model.[17]

A first imperfection we incorporate concerns the price: borrowing, when possible, typically occurs at a rate exceeding the lending rate. A prime example is credit card borrowing at extremely high interest rates. This imperfection has important implications for our model. In a recent paper, Davis, Kubler and Willen (2002) point out that households facing sufficiently high borrowing rates would never borrow to hold leveraged equity portfolios. Also, when facing high borrowing rates, investors never borrow and simultaneously hold liquid assets. Of course, expensive credit card debt still has consumption-smoothing benefits, as shown empirically in Jappelli, Pischke and Souleles (1998). Denoting the borrowing rate by $\overline{R}_b$, we set for simplicity $\overline{R}_b = \overline{R}_f + \mu$, the expected equity return.[18]

Even at a high borrowing rate, investors' borrowing capacity is typically not unlimited. We restrict the quantity of borrowing endogenously by considering an important feature of real-world credit markets: imperfect enforcement of financial contracts when a full menu of state-contingent assets is lacking. Households do not always honor the promises made in financial contracts and can (and do) file for bankruptcy instead. As analyzed in a complete-market pricing model by Alvarez and Jermann (2000) and Lustig (2001), limited enforcement will endogenously constrain the amount that creditors are willing to lend. Naturally the incentives to default (and therefore the resulting endogenous borrowing capacity) depend on what happens in the event of default. Alvarez and Jermann assume a harsh punishment: creditors seize all financial assets and the debtor is denied access to financial markets forever. On the other extreme, Lustig assumes that investors lose their liquid assets, but maintain access to financial instruments in the future. Realistically, U.S. households that file for bankruptcy lose most liquid financial assets and are typically denied credit for a short period of time. We introduce this into our calibrated model by having investors that default be excluded from markets for one year[19] and lose cash-on-hand above a certain exemption level (Y^s). Fay, Hurst and White (2002) report a \$5000 average in 1995 for all nonhomestead exemptions under Chapter 7 bankruptcy. We use this as the value for

Y^s in our calibration. Therefore, at time $t-1$, for $t < K$, an investor is allowed to borrow an amount $L_{i,t-1}$ at rate $\overline{R}_b$ when the following incentive constraint holds:[20]

$$V_{it}\left(Y_{it} - L_{i,t-1}\overline{R}_b\right) \geq U(\min\{Y^s, Y_{it}\}) + \delta p_t E_t V_{i,t+1}(Y_{i,t+1}) \tag{13}$$

The left-hand side is simply the continuation value of not defaulting and repaying the amount borrowed. This continuation value should exceed the value of defaulting, which consists of consuming at most the exemption level Y^s for one period and starting over again in the subsequent period with zero assets, but a new labor income realization (which can no longer be seized by the creditors). Since we analyze an incomplete-market setting and do not allow for state-contingent borrowing, this incentive constraint must hold for all possible realizations of Y_{it}, so that the investor would always be worse off defaulting. This implies that the lower bound of the support of Y_{it} will play a crucial role in the determination of $L_{i,t-1}$, the endogenous borrowing capacity. Strictly speaking, the income process in (2) has zero as the lowerbound of the support since it is lognormal. That would make $L_{i,t-1} = 0$. In other words, our exogenously imposed borrowing constraints are actually endogenous borrowing constraints consistent with limited enforcement. While it is well known from the literature on consumption and income risk that this also happens when ignoring bankruptcy and limited enforcement (e.g. Carroll (1997)), limited enforcement will allow us to obtain quantitatively similar results without relying on a 'literal' zero-income shock, as will be clear from the subsequent analysis.

Alvarez and Jermann (2000) explicitly assume a strictly positive lower bound for the support of the labor income distribution, in order to obtain 'borrowing constraints that are not too tight'. Our numerical model is particularly useful in highlighting the economic relevance of this point. As mentioned before, numerical solutions in this literature typically replace continuous state spaces with discrete ones. The particular discretization chosen will now matter since it determines the lower bound of the support of the income distribution and hence the endogenous borrowing limit. We would like to emphasize that this is an important economic question, not just a numerical issue. What matters is to what extent for instance social insurance bounds labor income away from zero and allows debtors to convince creditors that their incentives to default are limited.

In order to emphasize and illustrate the sensitivity of the results to the effective lower bound of the income distribution, we solve a version of the model with endogenous borrowing constraints according to (13) for different discretizations and corresponding lower bounds on Y_{it}. In the benchmark analysis, where the results are not sensitive to the discretization of the income distribution, the discretization used corresponds to a worst income shock (in a given period) of 3.5 standard deviations. When introducing endogenous borrowing constraints in Figure 7A, the average investor borrows up to 5000 dollars and has negative wealth for most of working life. Before retirement, he pays off the loans and start saving for retirement. Because of the expensive borrowing (at 6%) and the lack of wealth accumulation early in life relative to the benchmark case, the investor reaches retirement with significantly less wealth. This explains the higher equity allocation during retirement in Figure 7B. Early in life, the investor no longer saves, but borrows (the endogenous borrowing capacity according to (13) is roughly 15% of next period's income), and as a result he no longer invests in equities, as explained above.

When considering a richer support for the distribution of income shocks (a worst-case income shock of 4.5 standard deviations), the borrowing capacity according to (13) decreases significantly and is at most 9% of next period's income. Interestingly, the typical investor no longer borrows in the very beginning of life. Given the small borrowing limit young investors still save a positive amount for precautionary reasons. The lower borrowing capacity translates into more wealth accumulation before retirement and consequently a lower equity share during the first twenty years of retirement.

Finally, when expanding the income shock support further to allow for a worst income shock of 5.5 standard deviations, the endogenous borrowing capacity shrinks to almost zero. Most investors never borrow and the results are very similar to what we obtained when exogenously ruling out borrowing. This shows that the exogenous borrowing constraints imposed throughout the paper can be interpreted as endogenously determined, without literally relying on a disastrous or excessively extreme zero-income shock. Only when 5-standard deviation income shocks can be ruled out as impossible is the investor able to borrow substantially.

4.2.5 Bequest Motive

Another important simplification in the benchmark analysis is the absence of any bequest motive ($b = 0$). Investors with a desire to bequeath wealth to their heirs would be expected to save more. This could affect the simulated optimal portfolio due to its dependence on wealth. To investigate this, we now consider $b > 0$ in (1). Loosely speaking, b can be interpreted as the numbers of years of consumption of his descendants that the investor wants to save for, or the number of years by which the investor's horizon is effectively increased. Calibrating this parameter is challenging as there is little consensus in the literature on the strength of the bequest motive. Hurd (1989) estimates the strength of the preference for (intentional) bequests to be essentially zero (as measured by the marginal utility of bequests, which is assumed to be independent of wealth) and stresses that virtually all bequests are therefore accidental. Dynan, Skinner and Zeldes (2002) argue that it may not even be meaningful to attempt to disentangle a pure bequest motive from other savings motives in models with substantial uncertainty and precautionary savings, since precautionary buffers can serve a bequest motive whenever the buffer is not needed ex post. Cagetti (2002) validates this in wealth simulations based on Gourinchas and Parker (2002) and shows that even strong bequest motives have little effect on the savings decisions of relatively prudent consumers since these already accumulate a substantial buffer-stock of wealth.

We therefore consider different values for b ranging from 1 to 5. Figure 8 shows the simulated optimal equity shares, along with the benchmark case ($b = 0$) for comparison. For $b \leq 2$, the largest effects obtain very early in life and after retirement. Very young investors are relatively impatient and therefore save little in the absence of a bequest motive. The bequest motive alters their savings behavior somewhat, since mortality risk is already present (unlike in Cagetti where consumers live at least until age 65). In midlife, the bequest motive has a small effect on savings and therefore on optimal portfolios, since these investors build large savings anyway, very much in line with the arguments of Dynan, Skinner and Zeldes. Finally, during retirement, introducing a bequest motive has a relatively stronger effect, because of two factors. First, retired investors dissave rapidly in the benchmark model due to a weaker precautionary motive. Since the buffer-stock is run down, investors now explicitly save for their descendants. Second, mortality risk rises significantly after age 65, which increases the effective strength of the bequest motive. Overall however, the effects are not very large, except for $b = 5$, which is in line with results in the consumption literature on buffer-stock saving.[21] For $b = 5$, the effects are quite pronounced. The bequest motive is now so strong that precautionary savings are no longer sufficient, and additional wealth accumulation is needed at all ages. This lowers the optimal equity share substantially.

4.2.6 Educational Attainment

Figure 9 plots simulated labor income profiles, invested wealth and portfolio allocation over the life-cycle for the different education groups. It is important to keep in mind that in our stylized model an education group is characterized solely by the age at which working life begins, a given labor income profile and the stochastic properties of the shocks to it (i.e. variance and correlation with return innovations). In particular, we ignore any informational costs of investing in stocks and how these might differ across education groups.

As Figure 9 shows, the share of savings invested in stocks is similar for all education groups. However, some interesting differences emerge. First, the maximum of the portfolio profile occurs much earlier in the life-cycle for education groups 1 and 2 than for education group 3. Remember the explanation for the increasing part of the investment profile in terms of portfolio rule shifts: the reason is the steepness of the labor income profile. As can be seen from Figure 1 the profile is especially steep for education group 3. In midlife, the share of savings invested in stocks is, for a given age, increasing in the level of education. For a given age, the importance of future labor income is increasing in the level of education (Figure 1) and this means that the implicit riskless asset holdings (in the form of future labor income) are higher for more educated households. Finally, around age 55, the profiles for investors with and without highschool degree cross. The reason for this phenomenon is the larger replacement ratio that characterizes the retirement income of education group 1. Investors in education group 2 have a relatively smaller implicit riskfree asset position when reaching retirement and tilt their financial portfolio more heavily towards the riskfree asset.

4.2.7 Risk Aversion and Intertemporal Substitution

The effect of decreasing risk aversion is presented in Figure 10A. Lowering risk aversion affects not only the portfolio share directly as in the complete-market solution of (11), but also lowers wealth accumulation. Less risk-averse investor accumulate less precautionary savings, raising α even more. This would lead one to expect a larger effect than in the case without labor income. However, the simulation results suggest that the effect is mitigated (Figure 10A). The reason is that the investor we analyze faces short-sale and borrowing constraints. The fact that many investors are constrained explains the much smaller effect of changes in γ on the average portfolio share.

The time-separable isoelastic preferences in (1) have the property that a single parameter controls the willingness to substitute over time and across states. Epstein and Zin (1989) propose the following recursive formulation of intertemporal utility which disentangles risk aversion ($\gamma > 0$) and intertemporal substitution ($\psi > 0$):

$$U_{it} = \left\{ (1-\delta) C_{it}^{1-\frac{1}{\psi}} + \delta \left(\mathrm{E}_t \, U_{i,t+1}^{1-\gamma} \right)^{\frac{1-\frac{1}{\psi}}{1-\gamma}} \right\}^{\frac{1}{1-\frac{1}{\psi}}}, \tag{14}$$

Power utility is nested as the special case where $\psi = 1/\gamma$. In the absence of a bequest motive, the terminal condition for the recursion is:

$$U_{i,T+1} = 0.$$

Increasing the elasticity of intertemporal substitution from $\psi = 0.1$ (as in the benchmark when $\gamma = 10$) to 0.2 or 0.5 makes the investor more willing to substitute intertemporally. He is less concerned with low-frequency consumption smoothing and therefore saves less for retirement. As expected, this makes the equity portfolio share higher, as is clear from Figure 10B.

4.3 Empirical Predictions of the Model and Survey of the Evidence

The model has the following robust empirical predictions. First, the presence of labor income that is uncorrelated with equity returns increases the demand for equities. Second, labor income risk as calibrated from the data diminishes the tolerance for financial risk. Third, middle-aged investors hold portfolios that are more tilted towards equities than do older retired investors. Finally, there is some tendency for very young investors to choose somewhat less equity exposure than middle-aged investors. This last effect shows up most strongly in the presence of substantial income risk or when a (moderate) bequest motive is operational. Only when young investors can borrow sufficiently, and when the borrowing rate is high, will they choose not to invest in equities at all.

As previously discussed our benchmark scenario predicts an average level for α that is too high and too little heterogeneity in portfolio choice early in life. However, some of the extensions we have considered allow us to obtain predictions that are more in line with what we observe in the data. An empirically calibrated (small) probability of a disastrous labor income draw decreases the average level of α substantially and therefore seems to be an important feature for explaining this aspect of the data. Furthermore, the possibility of a disastrous labor income shock is also more successful at generating heterogeneity in asset choices early in life. Another important extension of the model that can help in explaining the data is endogenous borrowing, since it leads to non-participation by young investors.

There is compelling evidence to support the second prediction of the model. Using a unique Italian data-set, Guiso, Jappelli and Terlizzese (1996) find that investors facing more income risk invest indeed more cautiously. Heaton and Lucas (2000a) show that wealthy entrepreneurs who derive income from risky business ventures tend to invest less in equities than other wealthy investors.

Concerning the life-cycle profile of the fraction of wealth invested in equities, the studies by Ameriks and Zeldes (2000), Bertaut and Haliassos (1997), Heaton and Lucas (2000a), Poterba (2001), Poterba and Samwick (1997) and Guiso, Jappelli and Terlizzese (1996) document that it is hard to obtain unambiguous and clear-cut results. As emphasized for instance by Ameriks and Zeldes and by Poterba, separate age-, cohort- and time-effects cannot be identified without further restrictions in longitudinal data. In our model, the life-cycle pattern is a mainly an age-effect. Barring any differences in preferences, the model predicts no cohort-effect.[22] This is the identifying assumption made by Heaton and Lucas. A result they stress is that older households decrease the share of their wealth invested in risky asset as they substitute stocks, cash and bonds for riskier proprietary business ownership. This can be interpreted as the retirement effect in our model. Also excluding cohort-effects, Ameriks and Zeldes obtain a hump-shaped life-cycle pattern for equity ownership. However, they argue that this result is mainly driven by non-participation of younger investors, since age effects disappear entirely when examining stock exposure conditional on participation. Non-participation by young investors is a robust empirical finding. It can be obtained in our model when borrowing is possible for young investors, but this result relies on a sufficiently high lower bound for the income distribution. A full explanation of non-participation by young investors probably requires some fixed participation cost, as is explored by Campbell, Cocco, Gomes and Maenhout (1999), Gomes and Michaelides (2004) and Vissing-Jorgensen (2002).

It is also interesting to observe that the predictions of the model are somewhat consistent with the advice given by popular financial advisers, at least at a qualitative level. In the next section, we analyze the merit of this advice quantitatively.

5 Utility cost calculations

5.1 Alternative asset allocation rules

The optimal portfolio profiles obtained above are intuitive, but differ from the rules identified under the assumption of complete markets as in (11) or (12) and from the rules of thumb proposed by popular financial advisors. Most popular advisors suggest portfolio rules that allocate a lower fraction of the portfolio to stocks for older investors. For instance, an explicit heuristic given in Malkiel (1996, p. 418) suggests to invest a fraction in equities equal to 100 minus the investor's age:

$$\alpha_t = \frac{100 - t}{100} \tag{15}$$

Although our results could also roughly be characterized as involving a decreasing equity share over the life-cycle, the above heuristic (15) involves some simplifications that seem extreme from a theoretical viewpoint. It ignores crucial parameters such as risk aversion, equity premium and the variance of the innovations to returns. More subtlely, the heuristic is independent of either wealth or labor income, or any other individual-specific characteristics.

The economic importance of the suboptimality associated with (15) can easily be analyzed in our framework by computing the utility cost relative to the optimal rule. This is an economically meaningful metric to gauge the importance of the differences between two portfolio strategies. We can study under which conditions or for which agents the utility loss is likely to be most substantial. A similar analysis can be conducted for the optimal rules derived by Merton under the assumption of complete-markets and of no labor income, (11). The results can be suggestive about the joint importance of labor income and its risk characteristics for individual financial decision-making, of course within the assumptions of our incomplete-markets model. Furthermore, computing additionally the utility loss associated with (12) allows us to quantify the relevance of labor income *risk* and of market incompleteness for the portfolio decision. Comparing the welfare cost of following (12) with the loss associated with (11) contrasts the importance of labor income *risk* for the portfolio decision with the relevance of labor income itself. Finally, given the literature on limited stock market participation (Mankiw and Zeldes (1991), Basak and Cuoco (1998), Heaton & Lucas (1999), Polkovnichenko (1999) and Vissing-Jorgensen (1999)) and strong empirical evidence, it is also of interest to compute the welfare loss associated with a zero equity share.

One could argue that our paper proposes the following replacement for Malkiel's rule which we also evaluate:[23]

$$\alpha_t = \begin{cases} 100\% & t < 40 \\ (200 - 2.5t)\% & t \in [40, 60] \\ 50\% & t > 60 \end{cases} \tag{16}$$

We first consider the utility cost in our benchmark case (Figure 11 plots the life-cycle patterns for different rules we study) and then analyze some of the cases we studied above in section 5.2 on heterogeneity.

The welfare calculations are done in the form of standard consumption-equivalent variations: for each rule we compute the constant consumption stream that makes the investor as well-off in expected utility terms as the consumption stream that can be financed by the investment rule. Relative utility losses are then obtained by measuring the change in this equivalent consumption stream when deviating from the optimal rule towards the rule considered. More details are given in appendix C.

5.2 Results

Before analyzing the results, some explanation is in order with respect to how rule (12) is implemented in our model. Remember that equation (12) was derived under the assumption of complete markets, no mortality or labor income risk and a constant labor income stream. We therefore consider an investor who ignores labor income risk and discounts future wages at the riskfree rate (appropriately adjusted for mortality risk). This is in the spirit of the motivation underlying the analysis conducted. We compute the present discounted value of future labor income at each age and set the portfolio share invested in stocks equal to (12). However, W appears in this formula and is a function of an endogenous state variable. The optimization with respect to consumption takes this interdependence into account. Also the current level of the permanent shock is included in the information set used by the investor, conform the timing assumption of the rest of the paper. Finally, the constraint that the portfolio share belong to the unit interval is also imposed on the rule (12) used in the subsequent exercise.

Table 6 reports the results for these utility cost calculations.

For the benchmark parameters, rule (15) dominates (11), which in turn is preferred to nonparticipation in equity markets. The size of the welfare losses resulting from investing a constant share (i.e. ignoring labor income altogether) and especially from nonparticipation are substantial: even impatient, liquidity-constrained investors lose 1.5% and almost 2% of annual consumption respectively. The loss associated with the Malkiel heuristic is somewhat smaller but still quite substantial. Interestingly, the cost of following (12), i.e. of treating labor income as certain, is an order of magnitude smaller than the cost of ignoring labor income altogether and investing according to (11). We can interpret this as saying that labor income itself is crucial for optimal portfolio decisions, not so much the typical risk associated with it, at least when we do not allow for the possibility of a disastrous labor income realization. Finally, portfolio rule (16) does very well in terms of the utility costs as shown in the last column of Table 6, indicating that it is good description of our benchmark results.

We obtain similar results for investors without highschool degree and for college graduates. Relative to the highschool graduates, investors with a college degree lose less from each suboptimal rule. This is surprising since college-graduates invest almost always more than highschool-educated investors. Of course, what matters for expected utility, is not the level of the portfolio share per se, but rather how important investing is for the optimal intertemporal allocation of life-time resources. As was motivated before, highschool-graduates have a labor income profile that induces them to save a substantial amount of wealth (because it is less steep early in life, has a sharp drop at retirement and because they face more temporary income uncertainty). They therefore stand to gain much more from investing optimally.

Another remark is important. If one thinks that investing in equities involves substantial fixed costs, e.g. informational costs, then these results do not imply that highschool graduates are more likely to bear this cost in order to gain access to financial markets. If the cost is fixed across education groups in dollar terms, college graduates might still be more inclined to paying this cost than people without a college degree even though the welfare cost of nonparticipation is relatively higher for the latter. The reason is of course that college graduates are simply wealthier so that a given dollar cost has a smaller weight for them. To the extent that the fixed information cost is decreasing in education (a reasonable assumption), stock market participation would definitely be expected to be increasing in educational attainment.

Turning to investors subject to substantially more labor income risk,[24] we find that the portfolio rule given by (15) has a very small, indeed negligible welfare cost for young investors. This can be understood by simply comparing Figures $4A$ and 11: (15) is strikingly similar to the optimal rule during the first part of the life-cycle. Of course, the fraction of the utility cost of (11) that can be attributed to the risk characteristics of labor income rises. Ignoring labor income risk becomes more costly when the investor faces more of it.

Judging from the welfare losses to young investors from following (15) in all of these cases, one might be tempted to conclude that the heuristic does fairly well: the loss is well below 1% of annual consumption. However, the lack of any dependence on preference or asset return parameters in (15) easily disproves this conclusion. The welfare losses become quite substantial when risk aversion is lowered or the equity premium increased, as in both of these cases the investor wants far more exposure to equities than prescribed by (15). Indeed, ignoring labor income and simply following (11) does much better for the $\gamma = 2$ investor.[25]

It is interesting to compare the welfare loss of implementing (12) across different values of γ. The penalty for ignoring labor income risk drops sharply as the investor becomes less risk averse. This is intuitive: the less risk-averse investor is less concerned about the risk properties of her labor income stream and is hurt less when forced to ignore those.

Finally, we consider the sensitivity of our results with respect to the time-preference parameter. A more patient investor would be expected to save more, which drives down the share invested in stocks. It is unclear what the net effect is on the dollar amount invested in equities, which is what matters for the importance of the portfolio decision and therefore for the utility loss resulting from suboptimal behavior. However, another obvious effect would be expected in terms of welfare calculations: as the future is discounted less, the losses to young investors would be expected to be larger when considering $\delta = 0.98$. This effect clearly dominates in the last row of Table 6.

6 Conclusion

In this paper we develop a quantitative and realistically calibrated model to solve numerically for the optimal consumption and portfolio decisions of a finitely-lived individual who faces labor income uncertainty and can invest in either a risky or a riskless asset. Even though labor income is risky, the optimal portfolio rules indicate that labor income that is uncorrelated with equity returns is perceived as a closer substitute for riskfree-asset holdings than for equities. Therefore the presence of labor income increases the demand for stocks, especially early in life. Given the quantitative focus of our paper, we investigate what can reduce the average allocation to stocks and thus bring the empirical predictions of the model closer to what we observe in the data. Of all the extensions explored, we find that an empirically calibrated (small) probability of a disastrous labor income draw substantially decreases the average allocation to equities, and therefore seems to be an important element for explaining the data.

We also examine the life-cycle portfolio implications of endogenous borrowing in our incomplete markets setting, with a realistic and empirically parameterized penalty for default. We show that a crucial determinant of borrowing capacity and portfolio allocation is the lower bound for the income distribution. This is an important economic question. Just like in the disastrous labor income draw scenario, what matters is the extent to which social insurance bounds labor income away from zero. Investors with a bounded income process face a positive endogenous borrowing limit and, as a result, they hold negative wealth when young and do not invest in equities.

In terms of the life-cycle pattern of portfolio allocation, the share invested in equities is roughly decreasing with age. This is driven by the fact that the labor income profile itself is downward sloping. When aging, labor income becomes less important and hence the implicit riskfree-asset holdings represented by it. The investor reacts optimally to this by shifting her financial portfolio towards the riskfree asset.

Our results roughly support and rationalize the investment advice given by popular finance books and financial counselors, namely to shift the portfolio composition towards relatively safe assets as one ages. However the advice is quite imprecise and independent of risk-aversion and of the riskiness of labor income. This shows up in our welfare calculations: although the utility cost of following Malkiel's rule of thumb is quite small for some parameterizations, it rises as the investor becomes less risk-averse. We also report substantial penalties for investment strategies that ignore the presence of labor income and a fortiori for not investing in equities at all. Interestingly, we find that ignoring only labor income risk is associated with utility costs that are an order of magnitude smaller than when ignoring labor income altogether, except when we allow for the possibility of a disastrous labor income realization.

Figure Legends

Figure 1: Labor Income Processes estimated from the PSID for the three different education groups: households without high school education, households with high school education but without a college degree, and college graduates. For each group, the figure plots the estimated age dummies and a fitted third-order polynomial.

Figure 2A: Policy functions for the portfolio share invested in stocks in the next-to-last year of life, and in the complete markets case.

Figure 2B: Policy functions for the portfolio share invested in stocks at different stages of the life cycle.

Figure 2C: Policy functions for consumption at different stages of the life cycle.

Figure 3A: Simulated consumption, income and wealth profiles for the benchmark case.

Figure 3B: Present discounted value (PDV) of future labor income and ratio of the PDV of future labor income to (simulated) current financial wealth.

Figure 3C: Simulated portfolio share invested in stocks for the benchmark case.

Figure 4A: Simulated portfolio share invested in stocks for households in different sectors (i.e. with different labor income risk) and for households with zero labor income risk.

Figure 4B: Simulated portfolio share invested in stocks for different degrees of correlation between labor income shocks and equity returns.

Figure 5: Simulated portfolio share invested in stocks with a 0.5% probability of zero income realization, and for the benchmark case.

Figure 6: Simulated portfolio share invested in stocks with a probability of a drop in retirement income to 25% of its value.

Figure 7A: Simulated consumption and wealth profiles for the model with endogenous borrowing constraints. "#stdev" denotes the lower bound on labor income realizations, expressed as the number of standard deviations below the mean.

Figure 7B: Simulated portfolio share invested in stocks for the model with endogenous borrowing constraints. "#stdev" denotes the lower bound on labor income realizations, expressed as the number of standard deviations below the mean.

Figure 8: Simulated portfolio share invested in stocks for different degrees of bequest preference. "b" denotes the weight on the bequest term in the value function.

Figure 9: Simulated portfolio share invested in stocks for different education groups (i.e. different labor income profiles and different labor income risk).

Figure 10A: Simulated portfolio share invested in stocks for different levels of risk aversion.

Figure 10B: Simulated portfolio share invested in stocks with Epstein-Zin utility for different values of the elasticity of intertemporal substitution (EIS) and with relative risk aversion equal to 10.

Figure 11: Simulated (or predicted) portfolio share invested in stocks for different alternative investment strategies. "Optimal" denotes the optimal share predicted by the model; "100-age" refers to the common recommendation given by several financial advisers; "No Income" is the optimal allocation for an household without labor income; "No Income Risk" is the optimal allocation for an household with riskless labor income.

Appendix A: Numerical Solution

The model was solved using backward induction. In the last period the policy functions are trivial (the agent consumes all available wealth) and the value function corresponds to the indirect utility function. We can use this value function to compute the policy rules for the previous period and given these, obtain the corresponding value function. This procedure is then iterated backwards.

To avoid numerical convergence problems and in particular the danger of choosing local optima we optimized over the space of the decision variables using standard grid search. The sets of admissible values for the decision variables (consumption and portfolio allocation), were discretized using equally spaced grids. The state-space was also discretized. We used an equally spaced grid for cash-on-hand and, following Tauchen and Hussey (1991), approximate the density function for returns in the risky asset using Gaussian quadrature methods. The density function for both innovations to the labor income process were also approximated using Gaussian quadrature to perform the necessary numerical integration. The upper and lower bounds for cash-on-hand and consumption were chosen to be non-binding in all periods.

In order to evaluate the value function corresponding to values of cash-on-hand that do not lie in the chosen grid we used a cubic spline interpolation in the log of the state variable. This interpolation has the advantage of being continuously differentiable and having a non-zero third derivative, thus preserving the prudence feature of the utility function. The support for labor income realizations is bounded away from zero due to the quadrature approximation. Given this and the non-negativity constraint on savings, the lower bound on the grid for cash-on-hand is also strictly positive and hence the value function at each grid point is also bounded below. This fact makes the spline interpolation work quite well given a sufficiently fine discretization of the state-space.

Appendix B: Labor Income Calibration

We estimate equations (2) and (3) from the PSID. We used the family questionnaire since it provides a disaggregation of labor income and asset income. The families that were part of the Survey of Economic Opportunities subsample were dropped to obtain a random sample. Because the age-profile is potentially different for households with a female head of household and therefore requires a separate estimation, the sample was split according to the gender of the head of household. However there were too few observations for the subsample with female head of household, so that the estimation was only done for households with male head of household. From this subsample we eliminated retirees, non-respondents, students and housewives.

We took a broad definition of labor income so as to implicitly allow for (potentially endogenous) ways of self-insuring against pure labor income risk. If one were to include only labor income, the risk an agent faces would be overstated for several reasons: multiple welfare programs effectively set a lower bound on the support of non-asset income available for consumption and savings purposes, both the agent and his spouse can vary their labor supply endogenously, help from relatives and friends might be used to compensate for bad labor income shocks and so on. For this reason we defined 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 reported zero for this broad income category were dropped (326 observations, corresponding to 0.495 percent of the sample). Labor income defined this way was then deflated using the Consumer Price Index, with 1992 as baseyear.

The estimation controls for family-specific fixed effects. We chose this technique over the synthetic-cohort approach, because the latter one might overstate the variance of the income shocks as many sources of heterogeneity are not properly accounted for.[26] Starting the sample in 1970, a household appears at most 23 times in our sample. We do not remove households with less observations and estimate an unbalanced panel.

To control for education the sample was split in three groups: the observations without high school education, a second group with high school education but without a college degree, and finally college graduates. Doing so in a fixed-effects context is potentially problematic if education changes endogenously over the life-cycle. However we have only three different education groups and found few households switching from one education group to another. Consequently we considered the household as a new entity once its education changes.

For each education group we assume that the function $f(t, Z_{it})$ is 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) 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. Obviously, modelling unemployment as a switch in occupation is not appropriate for our purposes as we believe that the possibility of getting laid off is one of the main sources of labor income risk.

The logarithm of labor income was then regressed on dummies for age, family and marital status, and on household composition. We used households whose head was between

20 and 65 years old (except for the third education group where the lowest age included in the sample was 22). We fit a third-order polynomial to the age dummies to obtain the labor income profiles for the numerical solution. The results are similar for a fifth-order polynomial. The income profile generated (see Tables 1 and 2, and Figure 1) mimics the results of Attanasio (1995), Gourinchas and Parker (2002) and Hubbard, Skinner and Zeldes (1995).

Finally, the replacement ratio λ used to determine the amount of retirement income, was calibrated as the ratio of the average of our labor income variable defined above for retirees in a given education group to the average of labor income in the last working-year prior to retirement. The result is also reported in Table 2.

Next we estimate the error structure of the labor income process. Our procedure follows closely the variance decomposition described by Carroll and Samwick (1997). Defining r_{id} as

$$r_{id} \equiv \log(Y^*_{i,t+d}) - \log(Y^*_{it}),\ d \in \{1, 2, ..., 22\}\,, \tag{17}$$

where Y^*_t is given by

$$\log(Y^*_{it}) \equiv \log(Y_{it}) - \widehat{f}(t, Z_{it}), \tag{18}$$

then

$$Var(r_{id}) = d * \sigma^2_u + 2 * \sigma^2_\varepsilon. \tag{19}$$

We can then combine any two different series of r_{id}'s to get estimates of σ_u^2 and σ_ε^2, by running an OLS regression of $Var(r_{id})$ on d and a constant term (for all d). Note that we constrain the estimates of σ_u^2 and σ_ε^2 to be the same across all individuals. In our estimation we included all possible series of r_d's to maximize efficiency gains. The results are reported in Table 3.

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

$$r_{i1} = \xi_t + \omega_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}. \tag{20}$$

Averaging across individuals ($\overline{x}$ denotes the cross-sectional sample mean of x_i) gives

$$\overline{r_1} = \xi_t. \tag{21}$$

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

$$\overline{r}_1 = \beta \left(R_{t+1} - \overline{R}_f - \mu \right) + \psi_t. \tag{22}$$

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. The hypothesis of zero correlation cannot be rejected (see Table 3).

Appendix C: Welfare Metric

The welfare calculations are done in the form of standard consumption-equivalent variations: for each rule we compute the constant consumption stream that makes the investor as well-off in expected utility terms as the consumption stream that can be financed by the investment rule. Relative utility losses are then obtained by measuring the change in this equivalent consumption stream when deviating from the optimal rule towards the rule considered.

More precisely, we first solve the optimal consumption/savings problem for an agent who follows a rule of thumb $\{\alpha_t^R\}_{t=1}^T$.[27] We therefore allow the investors to control the optimal wealth dynamics given an exogenous porfolio weight. Denoting the optimal consumption stream solving this constraint problem by $\{C_t^R\}_{t=1}^T$, we compute expected life-time utility from implementing $\{\alpha_t^R\}_{t=1}^T$ as follows:

$$V^R = E_1 \sum_{t=1}^{T} \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{C_t^{R^{1-\gamma}}}{1-\gamma}, \tag{23}$$

where we drop the argument and time-subscript of $V_1^R(X_1)$ for notational simplicity and where superscript R indexes the portfolio rule followed.

Thus, V^R represents the maximal life-time utility for someone who will use rule $\{\alpha_t^R\}_{t=1}^T$ throughout her life and is now at the beginning of adult life. Then we can convert this discounted (remaining) lifetime utility into consumption units by computing the equivalent constant consumption stream $EC^R \equiv \{\overline{C}^R\}_{t=1}^T$ that leaves the investor indifferent between EC^R and between the consumption stream attained when implementing $\{\alpha_t^R\}_{t=1}^T$, i.e. $\{C_t^R\}_{t=1}^T$:

$$V^R = E_1 \sum_{t=1}^{T} \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{\overline{C}^{R^{1-\gamma}}}{1-\gamma}. \tag{24}$$

Therefore:

$$\overline{C}^R = \left[\frac{(1-\gamma)V^R}{\sum_{t=1}^{T} \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right)} \right]^{\frac{1}{1-\gamma}}. \tag{25}$$

Similarly, the constant consumption stream $EC^* \equiv \{\overline{C}^*\}_{t=1}^{T}$ that is equivalent in expected utility to our optimal solution, indexed by $*$, is defined by:

$$\overline{C}^* = \left[\frac{(1-\gamma)V^*}{\sum_{t=1}^{T} \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right)} \right]^{\frac{1}{1-\gamma}}. \tag{26}$$

The utility cost L_i^R to investor i associated with rule of thumb R is then simply computed as the percentage loss in equivalent consumption when adopting the rule of thumb rather than the optimal decision rule:

$$L^R = \frac{\overline{C}^* - \overline{C}^R}{\overline{C}^R} = \frac{V^{*\frac{1}{1-\gamma}} - V^{R\frac{1}{1-\gamma}}}{V^{R\frac{1}{1-\gamma}}}. \tag{27}$$

References

Alvarez, F. and U. Jermann, 2000, "Efficiency, Equilibrium and Asset Prices with Risk of Default," *Econometrica*, 68, 775-797.

Ameriks, J. and S. Zeldes, 2000, "How do Household Portfolio Shares Vary with Age?," working paper, Columbia University.

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

Basak, S. and D. Cuoco, 1998, " An Equilibrium Model with Restricted Stock Market Participation," *The Review of Financial Studies*, 11, 309-341.

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

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

Cagetti, M., 2002, "Interest Elasticity in a Life-Cycle Model with Precautionary Savings," *American Economic Review*, 91, 418-421.

Campbell, J. Y., J. Cocco, F. Gomes, P. Maenhout, 1999, "Investing Retirement Wealth: A Life-Cycle Model," *NBER Working Paper*, 7029.

Carroll, C. D., 1997, "Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1-55.

Carroll, C. D. and M. S. Kimball, 1996, "On the Concavity of the Consumption Function," *Econometrica*, 64, 981-92.

Carroll, C. D. and A. A. Samwick, 1997, "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40, 41-71.

Cocco, J. F., 2004, "Portfolio Choice in The Presence of Housing," working paper, London Business School; forthcoming in *The Review of Financial Studies.*

Constantinides, G. M., J. B. Donaldson, R. Mehra, 2002, "Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle," *Quarterly Journal of Economics,* 117, 269-296.

Dammon, R., C. Spatt and H. Zhang, 2001, "Optimal Consumption and Investment with Capital Tax Gains," *The Review of Financial Studies,* 14, 583-616.

Davis, S., F. Kubler and P. Willen, 2002, "Borrowing Costs and the Demand for Equity over the Life Cycle," working paper, University of Chicago.

Davis, S. and P. Willen, 2000, "Using Financial Assets to Hedge Labor Income Risks: Estimating the benefits," working paper, University of Chicago.

Davis, S. and P. Willen, 2001, "Occupation-level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice," working paper, University of Chicago.

Deaton, A., 1991, "Savings and Liquidity Constraints," *Econometrica*, 59, 1221-1248.

Dynan, K., J. Skinner and S. Zeldes, 2002,"The Importance of Bequests and Life-Cycle Saving in Capital Accumulation: A New Answer," *American Economic Review*, 92, 274-278.

Epstein, L. and S. Zin, 1989, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937-69.

Fama, E. F. and G. W. Schwert, 1977, "Human Capital and Capital Market Equilibrium," *Journal of Financial Economics*, 5, 115-146.

Fay, S., E. Hurst and M. White, 2002, "The Household Bankruptcy Decision," *American Economic Review*, 92, 706-718.

Gakidis, H., 1999, "Earnings Uncertainty and Life-Cycle Portfolio Choice," working paper, MIT.

Gollier, C. and J. W. Pratt, 1996, "Risk Vulnerability and the Tempering Effect of Background Risk," *Econometrica*, 64, 1109-1123.

Gomes, F. and A. Michaelides, 2004, "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," working paper, London Business School; forthcoming in *Journal of Finance.*

Gourinchas, P.-O. and J. Parker, 2002, "Consumption Over the Life Cycle," *Econometrica*, 70, 47-91.

Gross, D. and N. Souleles, 2002, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *Quarterly Journal of Economics*, 117, 149-185.

Guiso, L., T. Jappelli, and D. Terlizzese, 1996, "Income Risk, Borrowing Constraints and Portfolio Choice," *American Economic Review*, 86, 158-172.

Heaton, J. and D. J. Lucas, 1997, "Market Frictions, Saving Behavior and Portfolio Choice," *Macroeconomic Dynamics*, 1, 76-101.

Heaton, J. and D. J. Lucas, 1999, "Stock Prices and Fundamentals," in B. Bernanke and J. Rotemberg, eds., *Macroeconomics Annual 1999*, National Bureau of Economic Research, MIT Press, Cambridge.

Heaton, J. and D. J. Lucas, 2000a, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance*, 55, 1163-1198.

Heaton, J. and D. J. Lucas, 2000b, "Portfolio Choice in the Presence of Background Risk," *Economic Journal*, 110, 1-26.

Hu, X., 2002, "Portfolio Choice for Home Owners," working paper, Northwestern University.

Hubbard, G., J. S. Skinner and S. Zeldes, 1995, "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103, 360-399.

Hurd, M., 1989, "Mortality Risk and Bequests," *Econometrica* 57, 779-813.

Jagannathan, R. and N. R. Kocherlakota, 1996, "Why Should Older People Invest Less in Stocks Than Younger People?," Federal Reserve Bank of Minneapolis *Quarterly Review*, 20, 11-23.

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

Jappelli, T., S. Pischke and N. Souleles, 1998, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," *Review of Economics and Statistics*, 80, 251-262.

Kennickell, A., M. Starr-McCluer and B. Surette, 2000, "Recent Changes in U.S. Family Finances: Results from the 1998 Survey of Consumer Finances," *Federal Reserve Bulletin,* 86, 1-29.

Kimball, M. S., 1993, "Standard Risk Aversion," *Econometrica*, 61, 589-611.

Koo, H. K., 1998, "Consumption and Portfolio Selection with Labor Income: A Continuous Time Approach," *Mathematical Finance*, 8, 49-65.

Lustig, H., 2001, "The Market Price of Aggregate Risk and the Wealth Distribution," working paper, Stanford University.

Malkiel, B. G., 1996, *A Random Walk Down Wall Street: Including a Life-Cycle Guide to Personal Investing*, 6th ed. New York: Norton.

Mankiw, N. G. and S. Zeldes, 1991, "The Consumption of Stockholders and Non-Stockholders," *Journal of Financial Economics*, 29, 97-112.

Mehra, R. and E. C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145-161.

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

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

Polkovnichenko, V., 1999, "Heterogeneity and Proprietary Income Risk: Implications for Stock Market Participation and Asset Prices," working paper, Northwestern University.

Poterba, J. M., 2001, "Demographic Structure and Asset Returns," *Review of Economics and Statistics*, 83, 565-584.

Poterba, J. M. and A. A. Samwick, 1997, "Household Portfolio Allocation Over the Life Cycle," *NBER Working Paper*, 6185.

Pratt, J. W. and R. J. Zeckhauser, 1987, "Proper Risk Aversion," *Econometrica*, 55, 143-154.

Samuelson, P. A., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51, 239-46.

Staten, M., 1993, "The Impact of Post-Bankruptcy Credit on the Number of Personal Bankruptcies," working paper, Purdue University.

Storesletten, K., C. Telmer and A. Yaron, 2000, "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," working paper, Carnegie Mellon University.

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

Viceira, L., 2001, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable

Labor Income," *Journal of Finance*, 55, 1163-98.

Vissing-Jorgensen, A., 2002, "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," working paper, University of Chicago.

Yao, R. and H. Zhang, 2004, "Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints," working paper, University of North Carolina; forthcoming in *The Review of Financial Studies.*

Zhang, H., 1997, "Endogenous Borrowing Constraints with Incomplete Markets," *Journal of Finance*, 52, 2187-2209.

Notes

[1]Heaton and Lucas (2000a) emphasize the importance of proprietary or entrepreneurial income as a determinant of savings behavior and portfolio choice. One important characteristic of proprietary income is that it is positively correlated with the return on stocks.

[2]This extends the results from the static theoretical literature on background risk (Pratt and Zeckhauser (1987), Kimball (1993), Gollier and Pratt (1996)) to a life-cycle setting.

[3]Alvarez and Jermann (2000) and Lustig (2001) study the implications of endogenous borrowing constraints in a complete-market setting, while Zhang (1997) analyzes endogenous borrowing constraints in an incomplete-market model with a single asset (a pure discount bond).

[4]See Heaton and Lucas (2000b) for a survey of the literature.

[5]The issue of labor supply and portfolio choice has been studied in the context of a life-cycle model by Bodie, Merton and Samuelson (1992).

[6]In section 5.2.3 we relax this assumption and allow for uncertain retirement income.

[7]This is the consideration that motivates our choice of (5).

[8]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. Obviously, modelling unemployment as a switch in occupation is not appropriate for our purposes as we believe that the possibility of getting laid off is one of the main sources of labor income risk.

[9]As will be clear from the results, this conservative equity premium was chosen in combination with a high degree of risk aversion because the presence of (even risky) labor income substantially increases the demand for stocks.

[10]Note that the share invested in risky assets is not defined for low values of wealth, because the agent chooses not to save anything at these points in the state space.

[11]Indeed, it can be shown that the policy rules become increasing for low values of wealth (i.e. where a given labor income stream is large relative to wealth) if labor income shocks and stock return innovations are sufficiently positively correlated.

[12]Heaton and Lucas (1997) also report that realistic labor income uncertainty has only minor effects on portfolio choice in the context of their infinite-horizon model. This happens because agents obtain effective insulation from labor income shocks by simply accumulating enough wealth. This way investors self-insure by building up resources when labor income shocks are positive and by running down their assets in the face of adverse shocks.

[13]One may object that these results are driven merely by the fact that marginal utility approaches infinity as consumption goes to zero with power utility, and conjecture that they are therefore of limited economic interest. To investigate this, we have considered a disastrous labor income shock that is less extreme and bounded away from zero. Remarkably, very similar results obtain when labor income drops with the same 0.5% probability to 10% of its mean level.

[14]The structure of the health shocks is admittedly oversimplified and purely illustrative. For instance, health shocks may well exhibit persistence. We make the i.i.d. health shock

in our model particularly extreme to partially capture this, while maintaining tractability. A more elaborate analysis is beyond the scope of this paper. Hubbard, Skinner and Zeldes (1995) carefully calibrate medical expenses based on the National Health Care Expenditure Survey in their model of buffer-stock saving without portfolio choice.

[15]For example, according to the 1998 Survey of Consumer Finances (Kennickell, Starr-McCluer and Surette, 2000), low-income households do have non-collateralized debt. Given a mean pre-tax family income of USD 53100 in 1998 (the median is USD 33400), we can define very low-income households as the ones with income below USD 10000 (12.6% of households) and low-income households as the ones with income between USD 10000 and USD 25000 (24.8%). Of the very-low income households, 20.6% have outstanding credit card balances, while 37.9% of low-income households report credit card debt. Conditional on having credit card debt, the median value of debt is USD 1100 and USD 1000 respectively. Other forms of non-collateralized borrowing are negligible for low-income households. However, Gross and Souleles (2002) argue that credit card debt is underreported in the SCF, so that these numbers should be interpreted as lower bounds.

[16]We are grateful to an anonymous referee for suggesting this extension.

[17]Zhang (1997) analyzes endogenous borrowing constraints in an incomplete-market model with a single asset (a pure discount bond).

[18]Davis, Kubler and Willen (2002) investigate in detail the effect of alternative parameter values for the borrowing rates.

[19]Staten (1993) reports that 73.7% of bankruptcy filers in his sample could access at least

one line of credit within one year (see also footnote 14 in Fay, Hurst and White (2002)). The results we present below are robust to a change of the exclusion period from one to two years.

[20]Since retirement income is assumed to be riskless in this version of the model, we exogenously (but realistically) rule out borrowing after age 65 ($t \geq K$).

[21]Dammon, Spatt and Zhang (2001) consider the effect of taxation on life-cycle portfolio choice. They demonstrate how older investors shift their portfolios towards equities since capital gains on bequests are not taxed.

[22]Poterba (2001) makes the identifying assumption of zero time-effects in order to estimate age- and cohort-effects. When looking at desired wealth accumulation he finds very modest age-effects.

[23]We are grateful to an anonymous referee for suggesting this rule as a way to summarize the portfolio rule we obtain for the benchmark results.

[24]For brevity, we only report the results for investors in the most extreme sector, Agriculture. The portfolio shares and welfare costs obtained for agents in Construction and Public Administration were found to be very similar to the benchmark case.

[25]To the extent that we exclude leveraging from the analysis, these welfare losses are to be interpreted as conservative lower bounds. Investors with low risk aversion or facing generous excess returns frequently find these constraints binding. When lifting the borrowing constraints, the investor would suffer even more from investing according to $\alpha = (100 - age)/100$.

[26]We estimated the income profiles with the synthetic-cohort technique as well. Although the number of degrees of freedom is substantially larger with this technique, the shape of the profiles are very similar to the ones obtained with family fixed-effects. The estimated variance is of course larger.

[27]For notational simplicity, we suppress the subscript i indexing the investor from here onwards.

Table 1: Labor Income Process: Fixed-Effects Regression

Independent Variable	No Highschool		Highschool		College	
Log Real Income	Coefficient	t-stat	Coefficient	t-stat	Coefficient	t-stat
Family Size	-0.0176	-3.12	-0.0236	-7.42	-0.0228	-4.63
Marital Status	0.4008	18.48	0.4437	43.21	0.4831	30.50
Constant	2.6275	56.63	2.7004	118.27	2.3831	44.38
n	1104		2816		1110	
T-bar	8.58		9.57		11.06	
σ_{ϵ}^{2}	0.1583		0.1161		0.1189	
R^2 within	0.0648		0.1395		0.2609	
F-stat	12.27		83.10		87.23	

Note to table 1: Labor income is defined 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. The table reports the results of the fixed-effects estimations for different education groups with age dummies.

Table 2: Labor Income Process: Coefficients in the age polynomial.

	No Highschool		Highschool		College	
	3rd order	5th order	3rd order	5th order	3rd order	5th order
Constant	-2.1361	0.0549	-2.1700	-7.5185	-4.3148	-29.6153
Age	0.1684	-0.1277	0.1682	0.9046	0.3194	3.4476
$Age^2/10$	-0.0353	0.1181	-0.0323	-0.4213	-0.0577	-1.5443
$Age^3/100$	0.0023	-0.0359	0.0020	0.1007	0.0033	0.3439
$Age^4/1000$	-	0.0046	-	-0.0121	-	-0.0377
$Age^5/10000$	-	-0.0002	-	0.0006	-	0.0016
Replac. rate	0.88983	0.88983	0.68212	0.68212	0.938873	0.938873

Note to table 2: The endogenous variable are the age dummies estimated in the first stage fixed effects regressions (shown in table 1). The exogenous variables are an age polynomial for the working life period (until age 65) and a constant for the retirement period.

Table 3: Variance Decomposition and Correlation with Stock Returns

	No High School	(t-ratios)	High School	(t-ratios)	College	(t-ratios)
σ_{ε}^2	0.1056	13.260	0.0738	21.962	0.0584	13.089
σ_u^2	0.0105	9.909	0.0106	24.258	0.0169	29.196
ρ	-0.0143	-0.217	0.0058	0.143	-0.0175	-0.540

Note to table 3: The table reports estimates of the variance of both permanent and transitory labor income shocks, and the correlation between permanent labor income shocks and stock return. The estimation is based on the error structure of the labor income process estimated by the regressions described in Table 1. The estimation method follows closely the variance decomposition method in Carroll and Samwick (1997), and we use a similar procedure to estimate the correlation between labor income shocks and stock returns (details in Appendix B).

Table 4: Baseline Parameters.

Description	Parameter Value
Retirement Age (K)	65
Discount Factor (δ)	0.96
Risk Aversion (γ)	10
Bequest Motive (b)	0
Variance of Transitory Shocks (σ_{ε}^2)	0.0738
Variance of Permanent Shocks (σ_u^2)	0.0106
Correlation with Stock Returns (ρ)	0
Riskless rate ($\overline{R}_f - 1$)	0.02
Mean return on stocks ($\mu - 1$)	0.06
Std. stock return (σ_η)	0.157

Note to table 4: Baseline parameters values for the model.

Table 5: Variance Decomposition for Different Sectors.

Industry	Total	Perm/Total
Agriculture	0.3094	0.1350
Construction	0.1030	0.1394
Public adm.	0.0470	0.1657

Note to table 5: The table reports, for households in three different sectors, estimates of the variance of both permanent and transitory labor income shocks, and the correlation between permanent labor income shocks and stock return.

Table 6: Utility Cost Calculation (Percentage Points)

Parameters	100-Age	No Income*	No Income Risk**	Zero	Approx.***
Benchmark	0.637	1.531	0.152	2.108	0.084
No high school	0.711	1.763	0.110	2.340	0.122
College	0.277	0.669	0.048	0.894	0.033
Agriculture, $\gamma = 10$	0.815	2.256	0.382	3.318	0.322
Agriculture, $\gamma = 3$	1.202	1.231	0.326	3.730	0.304
Disastrous Inc. Shock	0.517	0.657	0.810	1.038	0.804
$\gamma = 2$	1.304	0.548	0.035	2.395	0.672
$\gamma = 5$	0.986	1.516	0.135	2.541	0.318
$\mu = 5.75\%, \gamma = 5$	1.560	1.787	0.157	3.964	0.504
$\mu = 5.75\%, \gamma = 10$	1.069	2.140	0.271	3.314	0.194
$\delta = 0.98$	0.801	1.815	0.272	2.551	0.054

Note to table 6: This table reports welfare calculations in the form of standard consumption-equivalent variations. We start by computing, for each rule, the constant consumption stream that makes the investor indifferent in expected utility terms to the consumption stream that can be financed by the investment rule. We then obtain relative utility losses by measuring the change in this equivalent consumption stream when deviating from the optimal rule towards the rule considered. The rules considered are:

* The share invested in the risky asset is given by $\alpha = \frac{\mu}{\gamma \sigma_\eta^2}$

** The share invested in the risky asset is given by $\alpha_t = \frac{\mu}{\gamma \sigma_\eta^2} \frac{W_t + PDV_t(FY_t)}{W_t}$

*** The share invested in the risky asset is given by:

$$\alpha_t = \begin{cases} 100\% & t < 40 \\ (200 - 2.5t)\% & t \in [40, 60] \\ 50\% & t > 60 \end{cases}$$

Figure 1

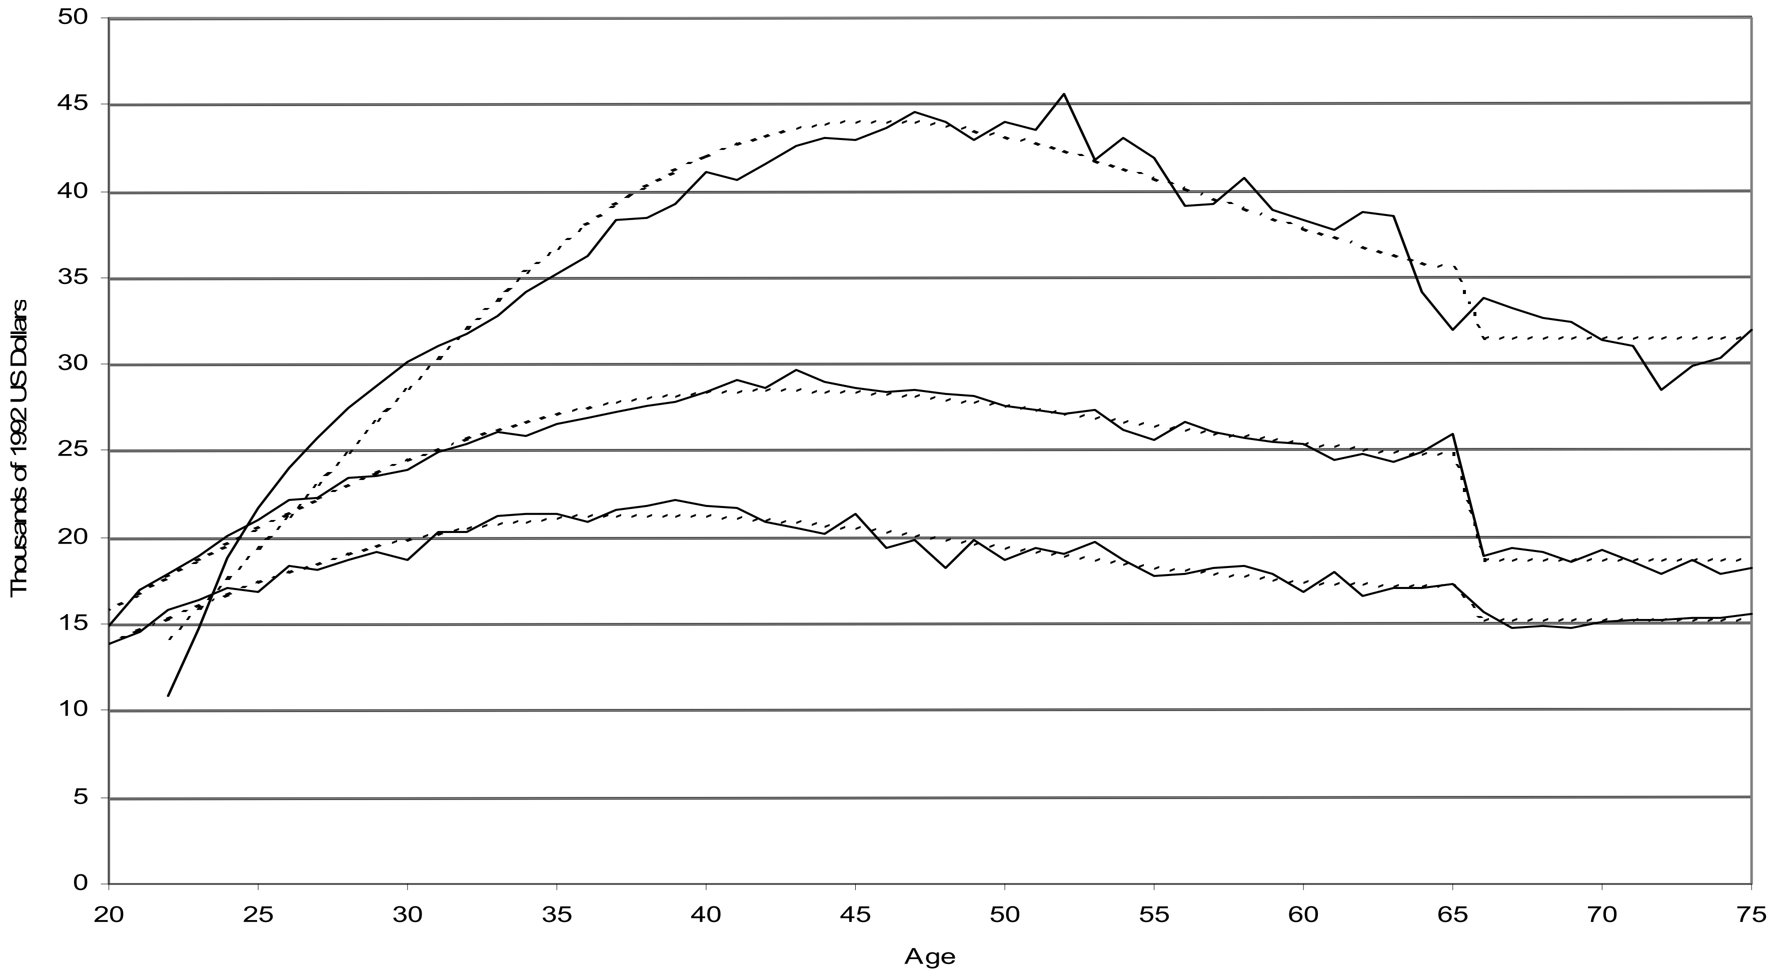
50
45
40
35
30
25
20
15
10
5
0
Thousands of 1992 US Dollars
20
25
30
35
40
45
50
55
60
65
70
75
Age

Figure 2A

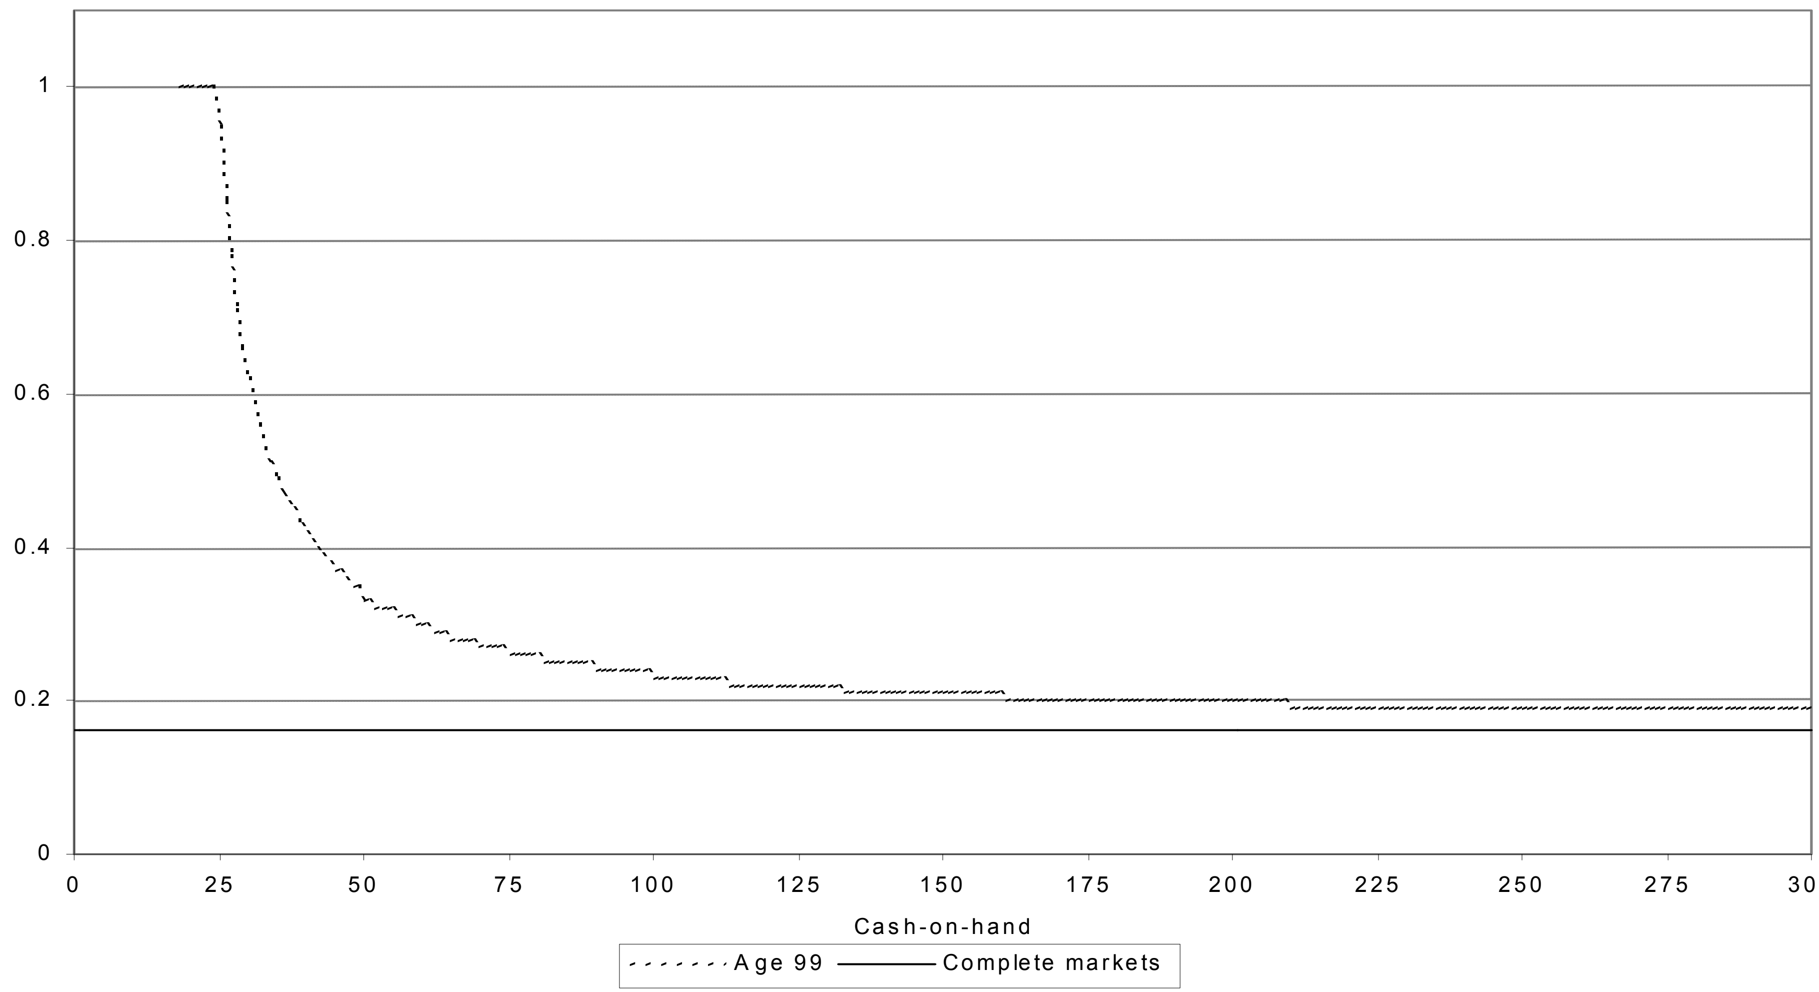

1
0.8
0.6
0.4
0.2
0
0
25
50
75
100
125
150
175
200
225
250
275
300
Cash-on-hand
Age 99
Complete markets

Figure 2B

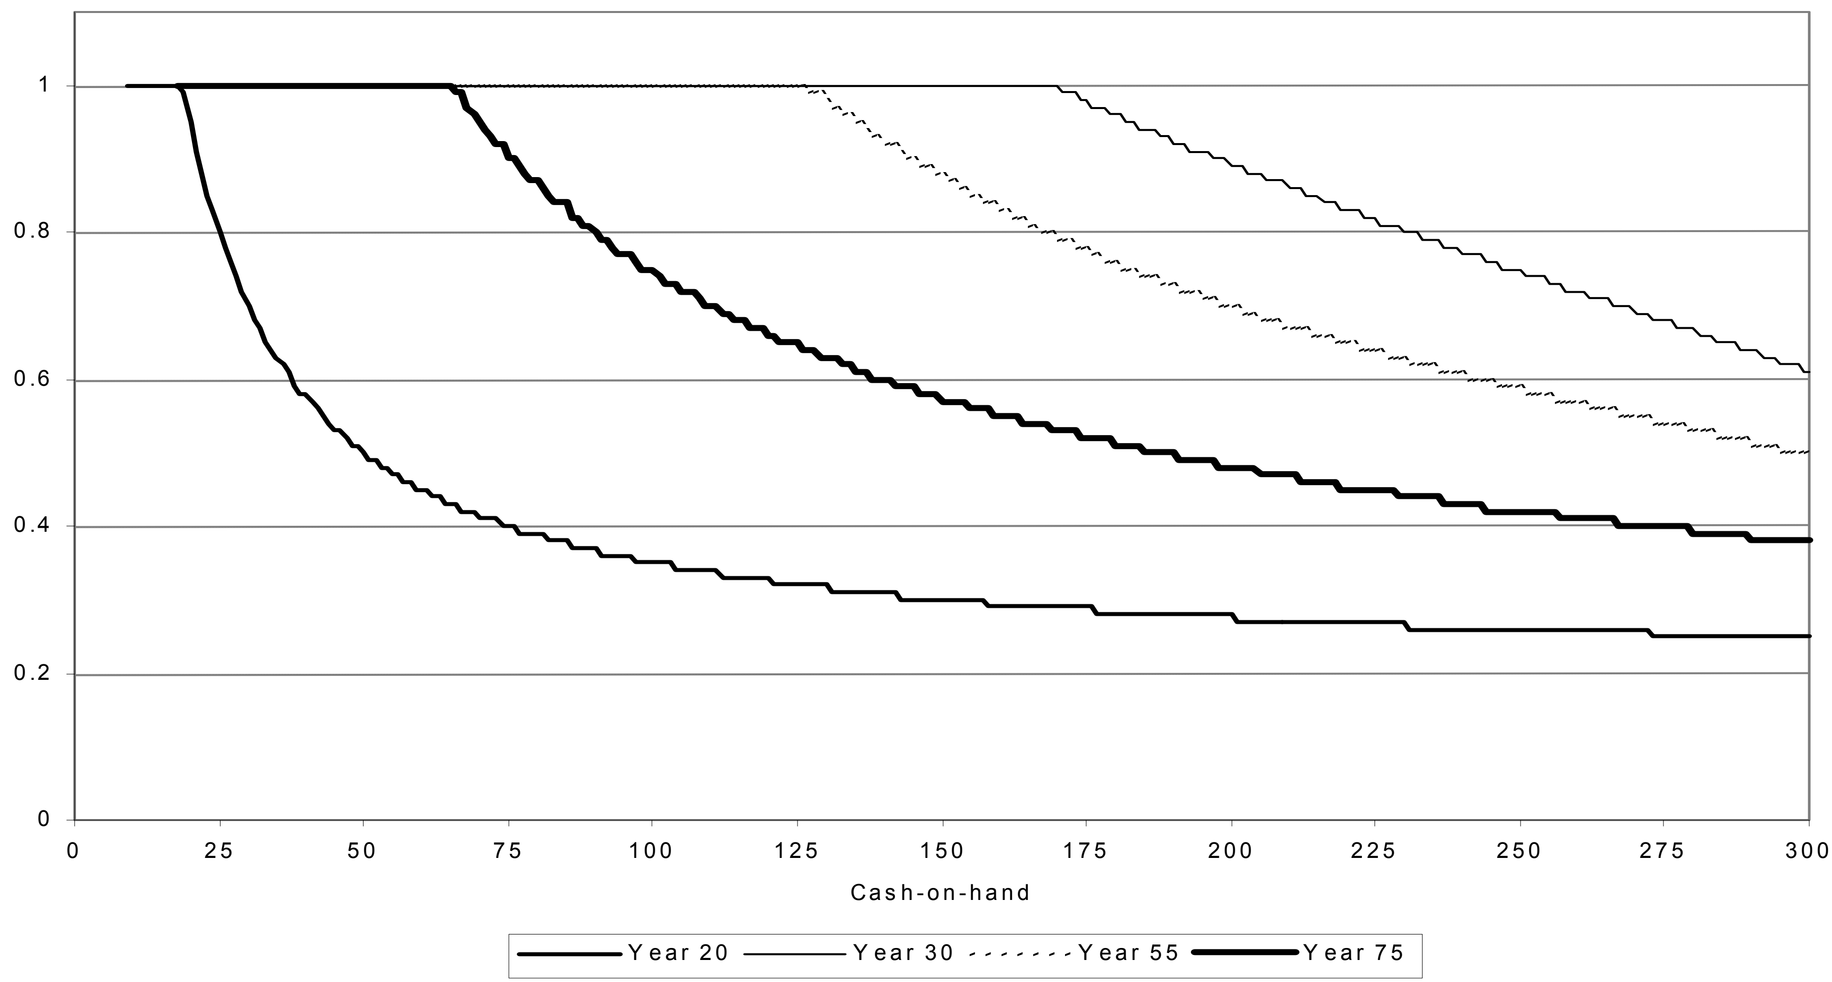

1
0.8
0.6
0.4
0.2
0
0
25
50
75
100
125
150
175
200
225
250
275
300
Cash-on-hand
Year 20
Year 30
Year 55
Year 75

Figure 2C

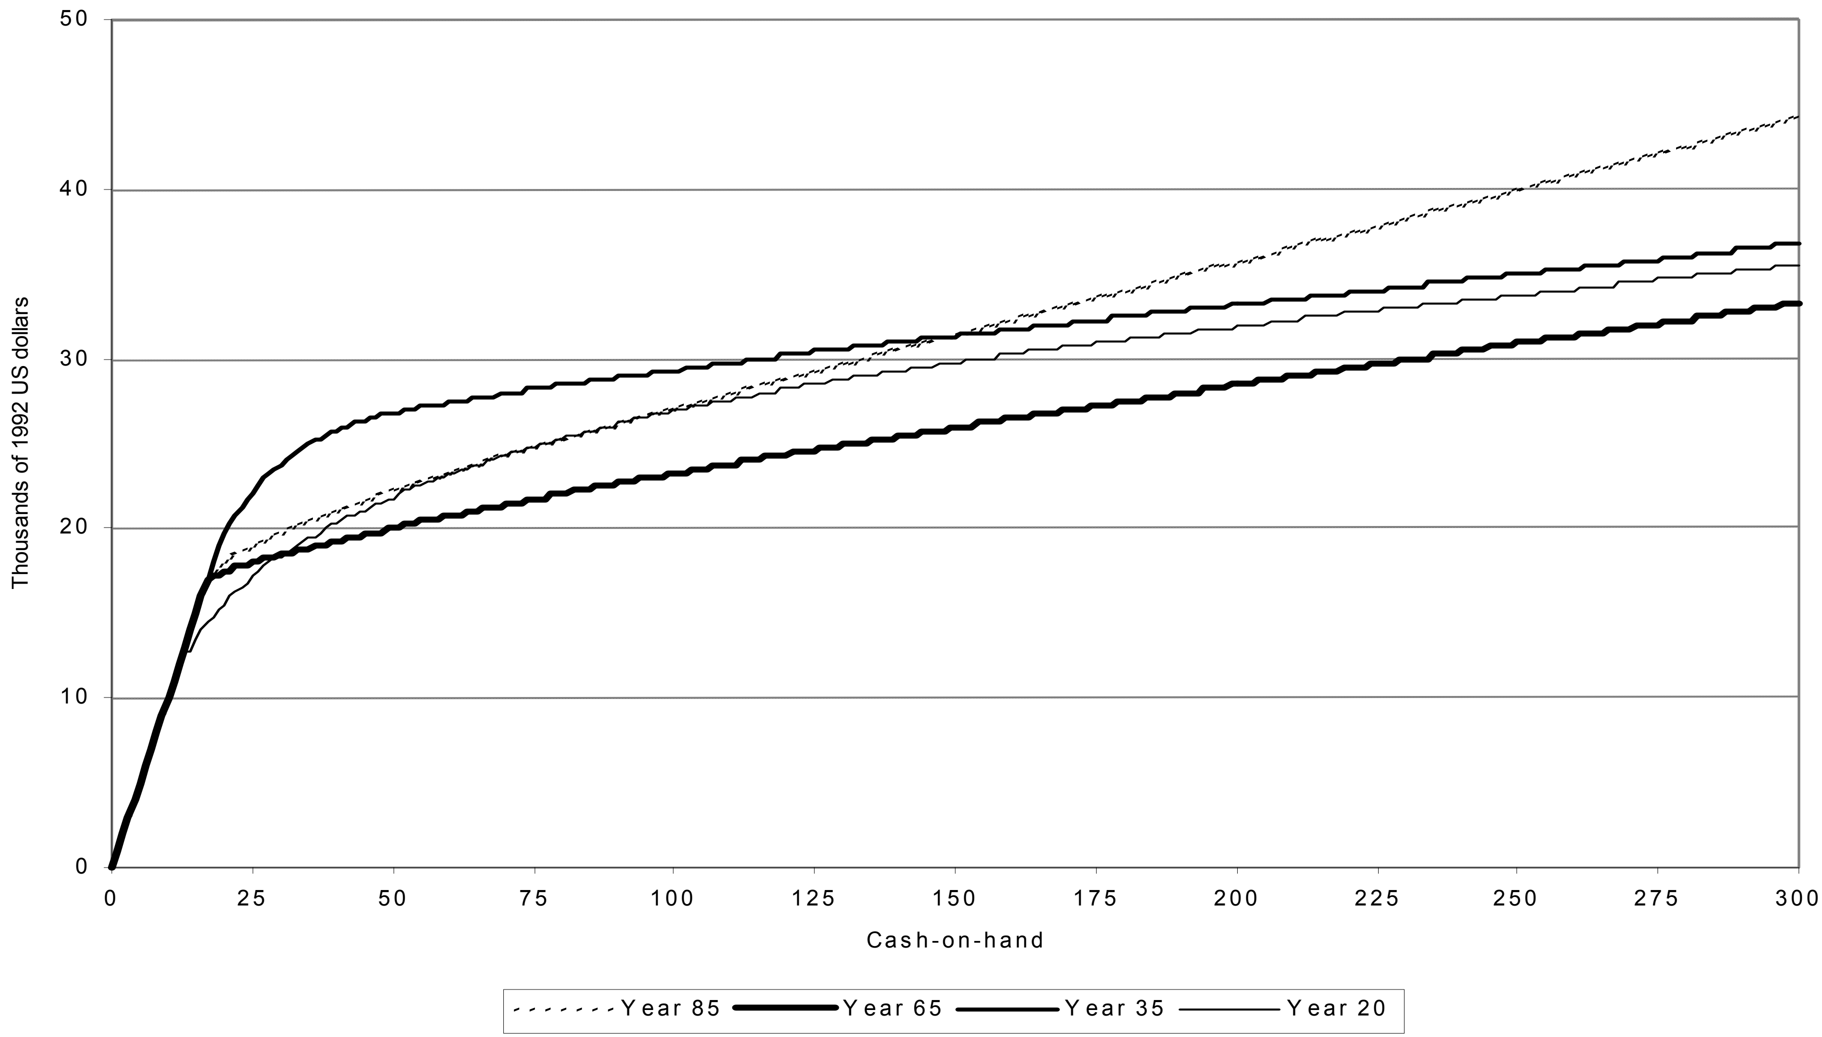

Thousands of 1992 US dollars
50
40
30
20
10
0
0
25
50
75
100
125
150
175
200
225
250
275
300
Cash-on-hand
Year 85
Year 65
Year 35
Year 20

Figure 3A

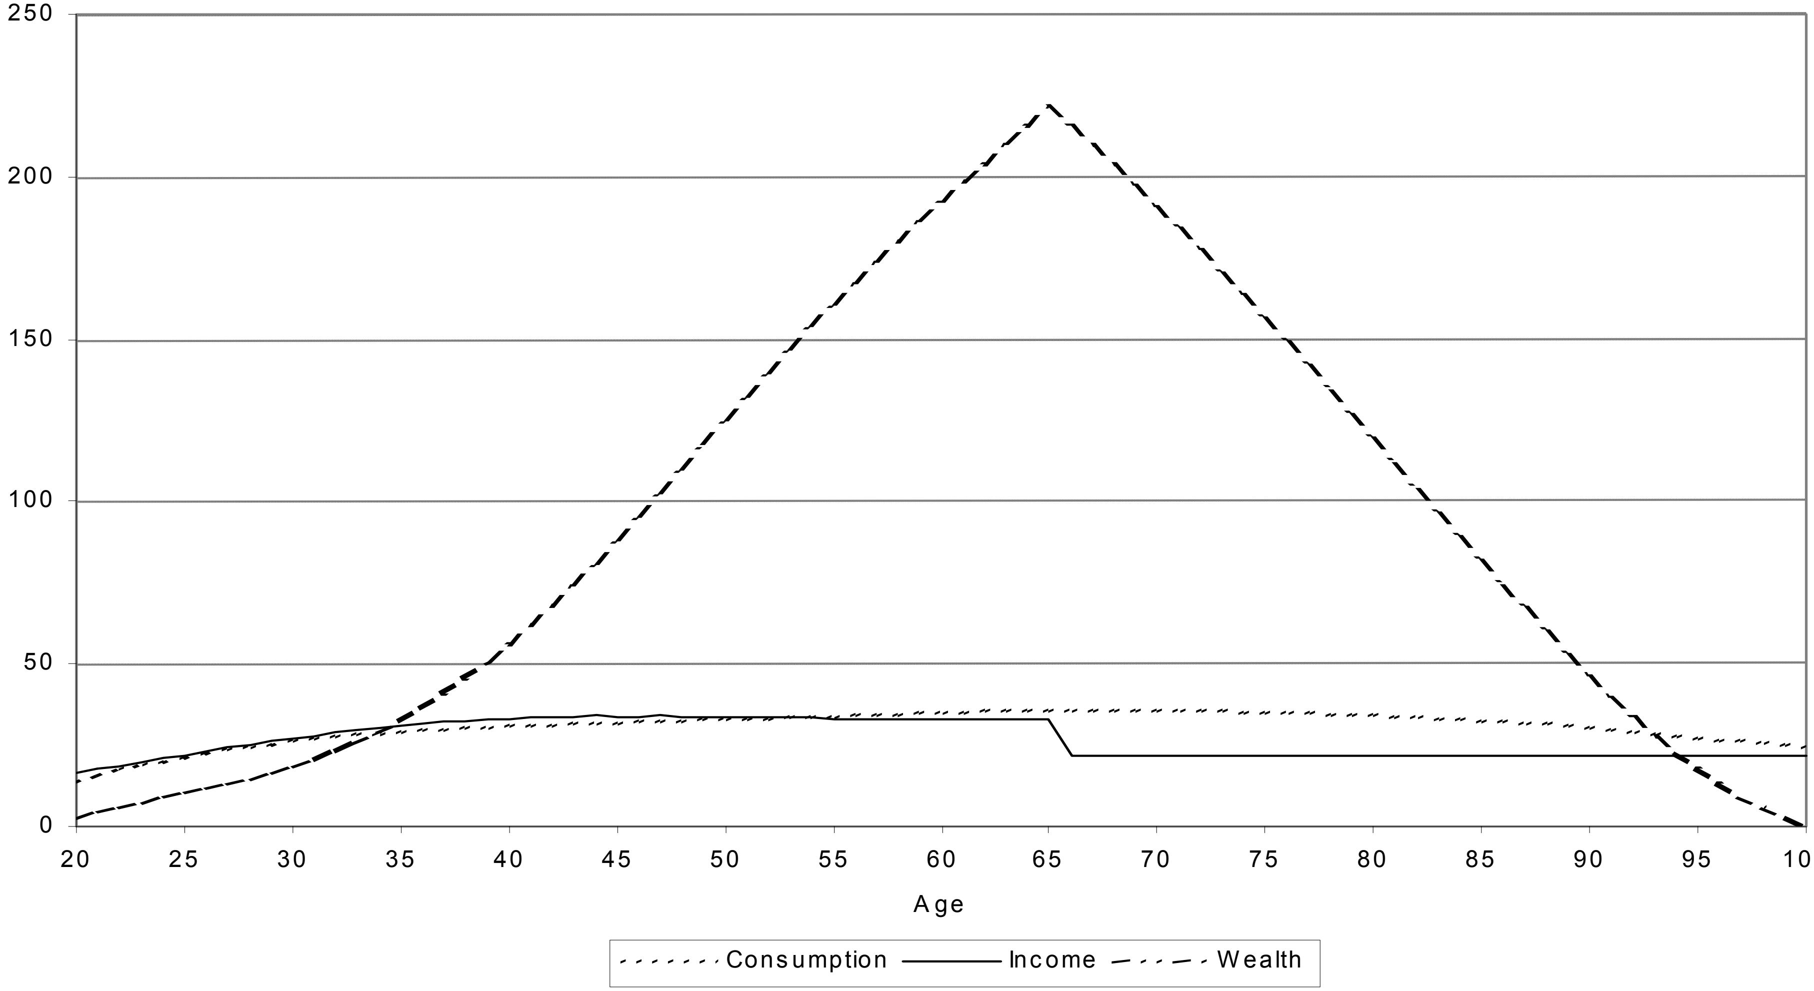
250
200
150
100
50
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
Age
Consumption
Income
Wealth

Figure 3B

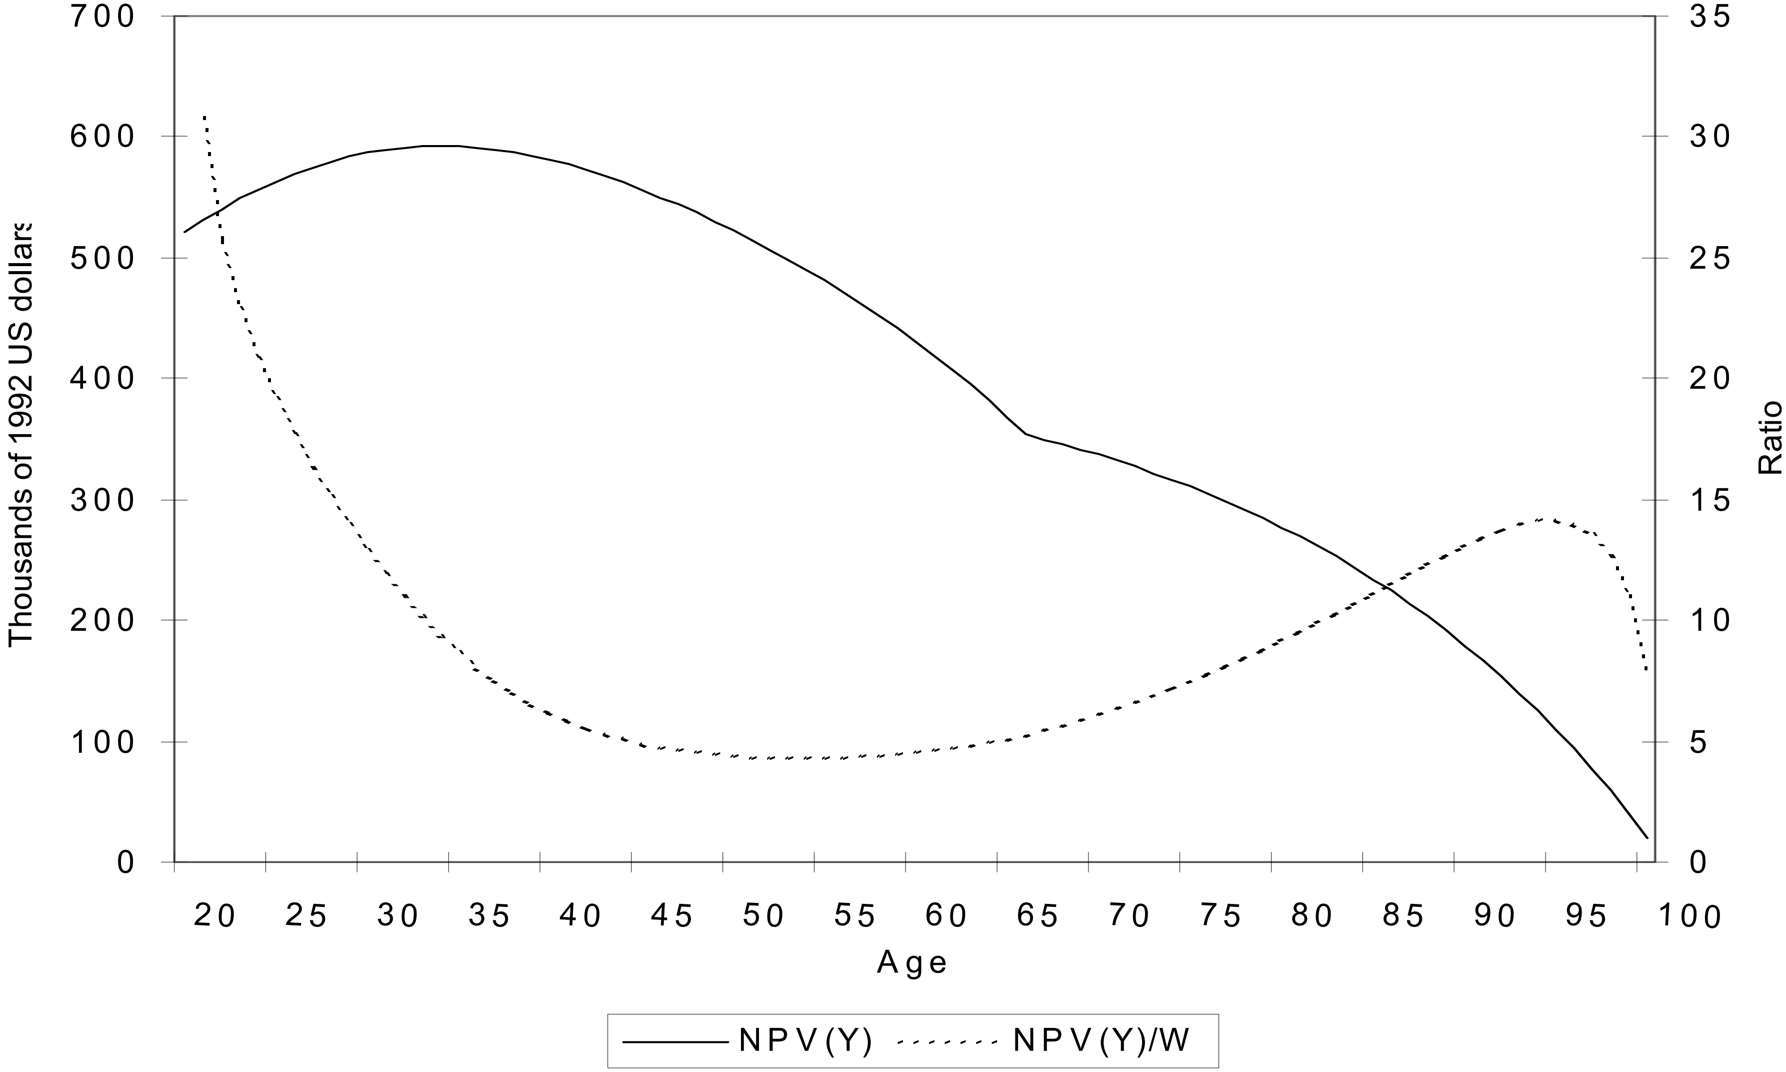

Thousands of 1992 US dollars
0
100
200
300
400
500
600
700
Ratio
0
5
10
15
20
25
30
35
20 25 30 35 40 45 50 55 60 65 70 75 80 85 90 95 100
Age
NPV(Y)
NPV(Y)/W

Figure 3C

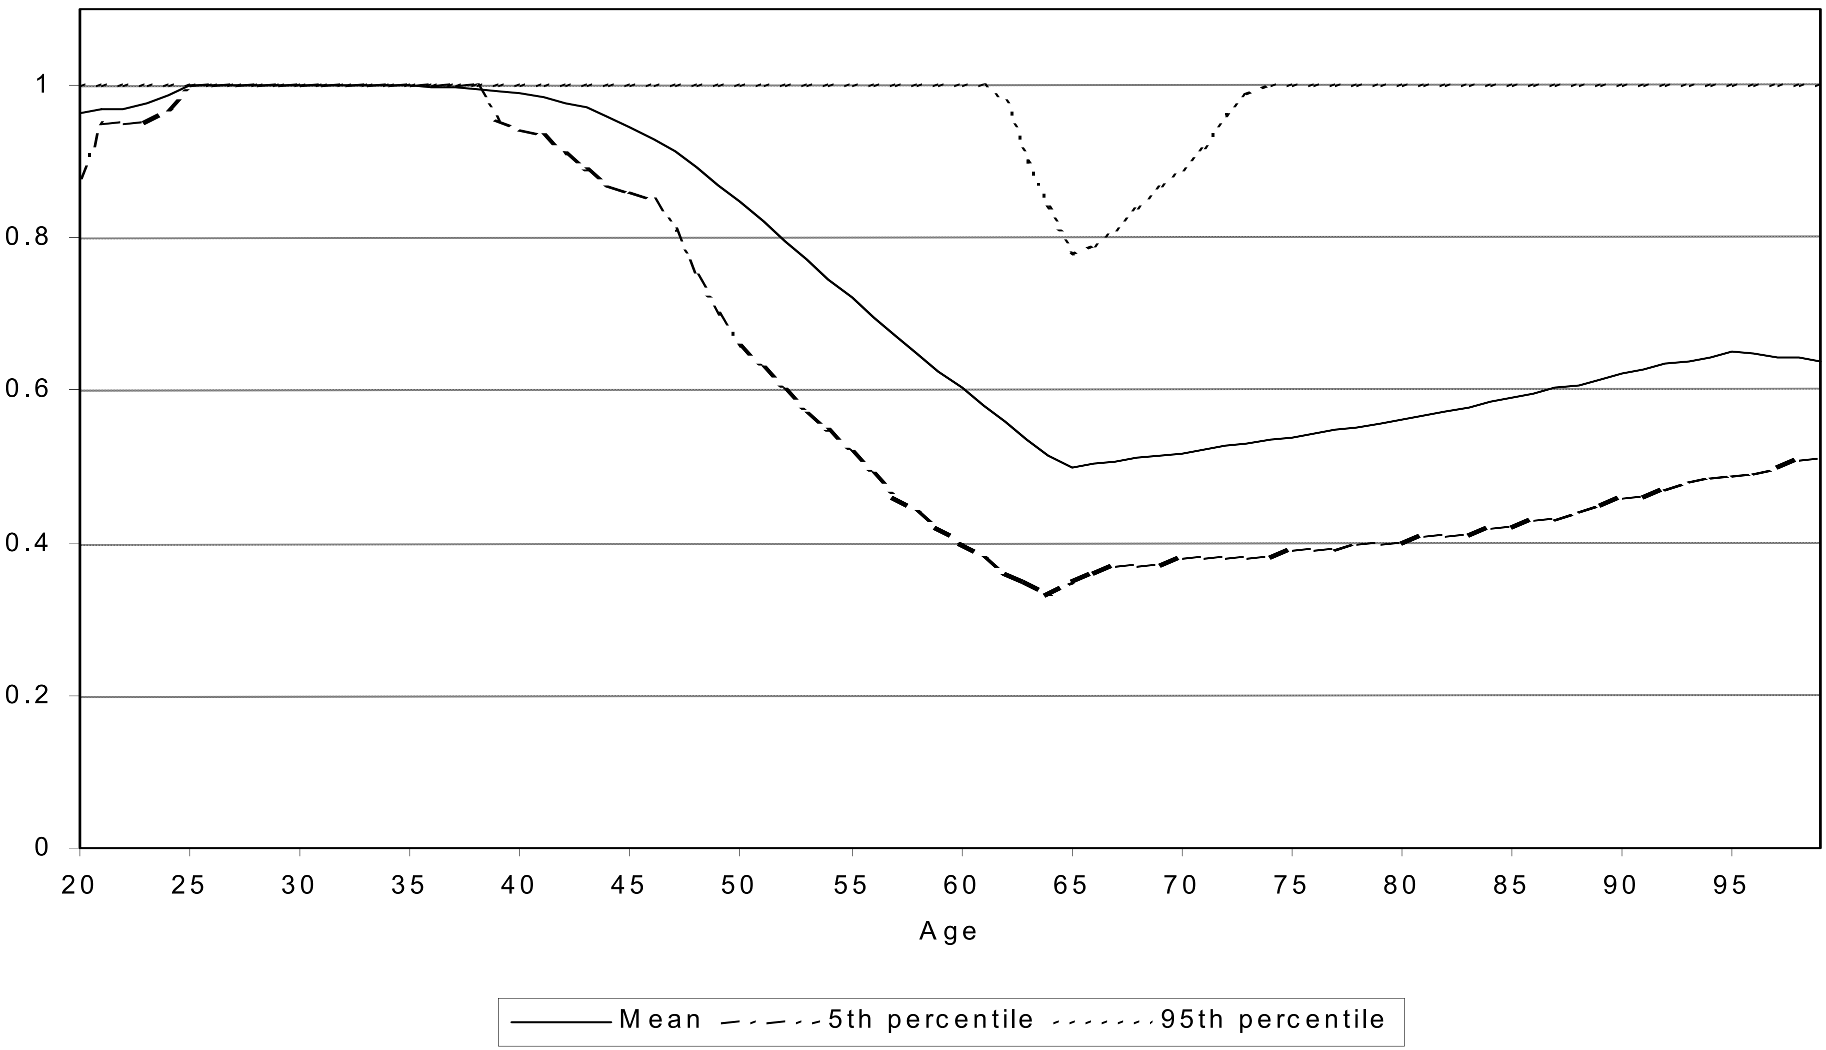
1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
Mean
5th percentile
95th percentile

Figure 4A

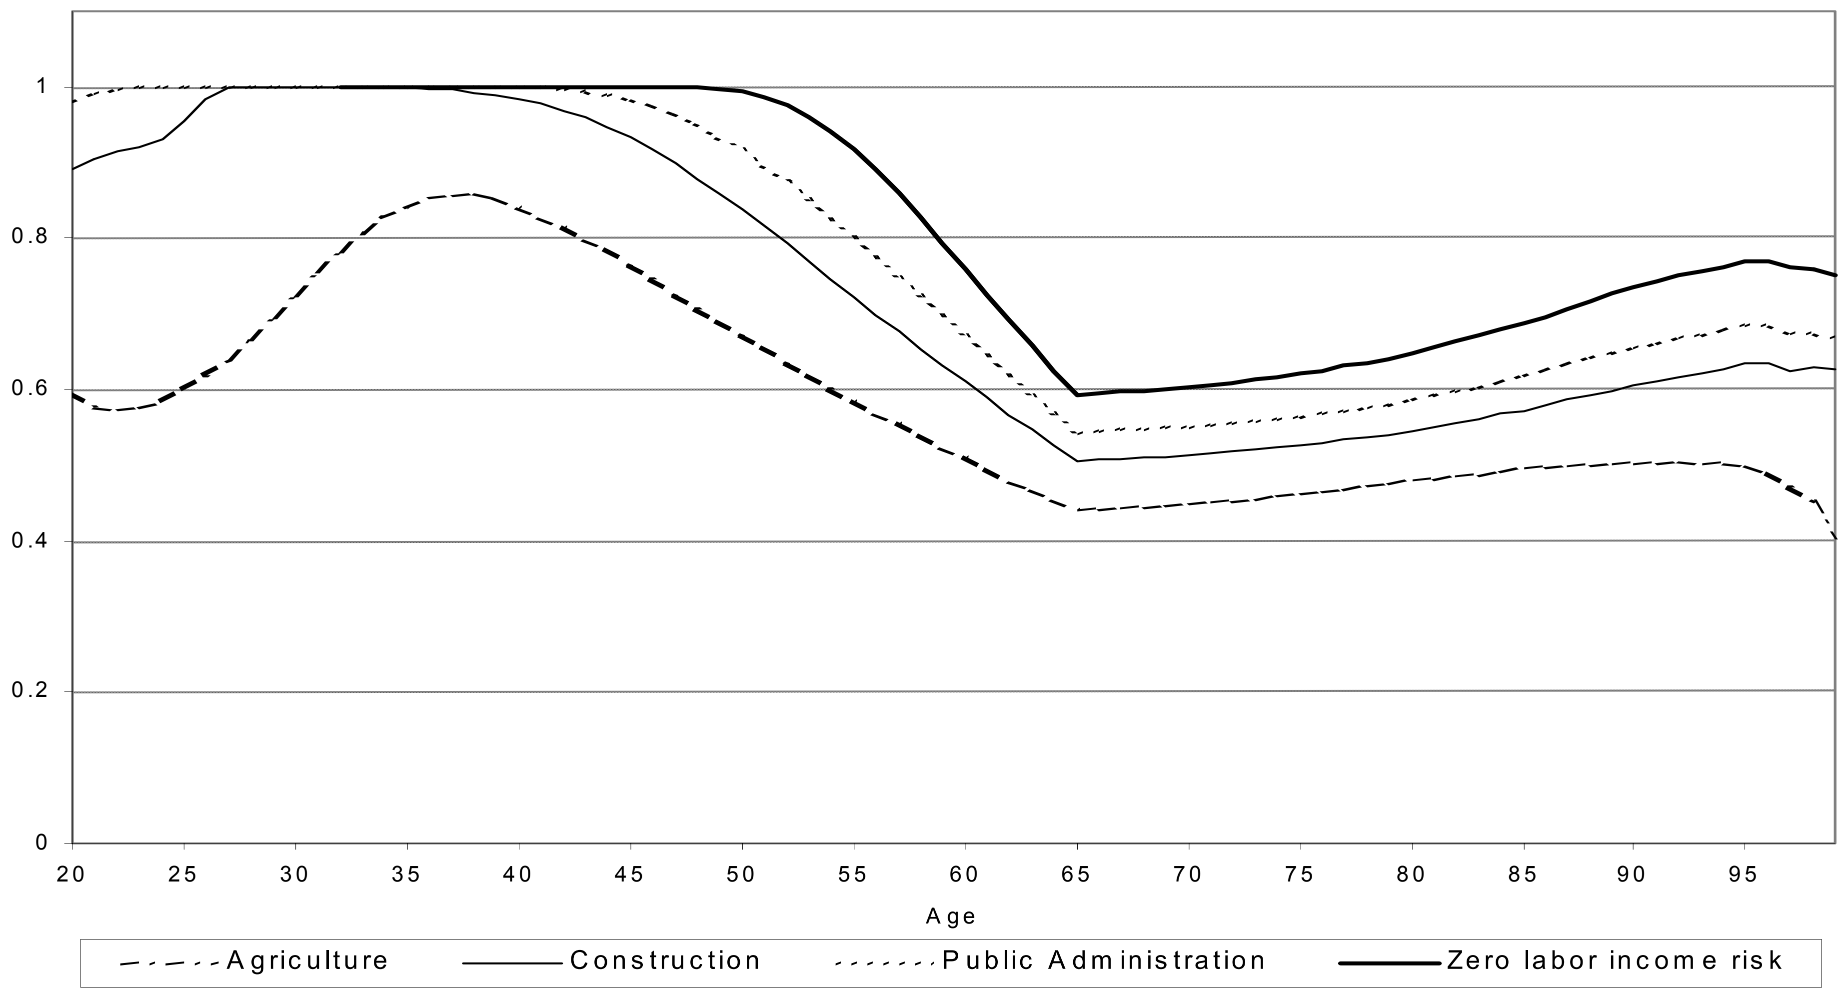

1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
Agriculture
Construction
Public Administration
Zero labor income risk

Figure 4B

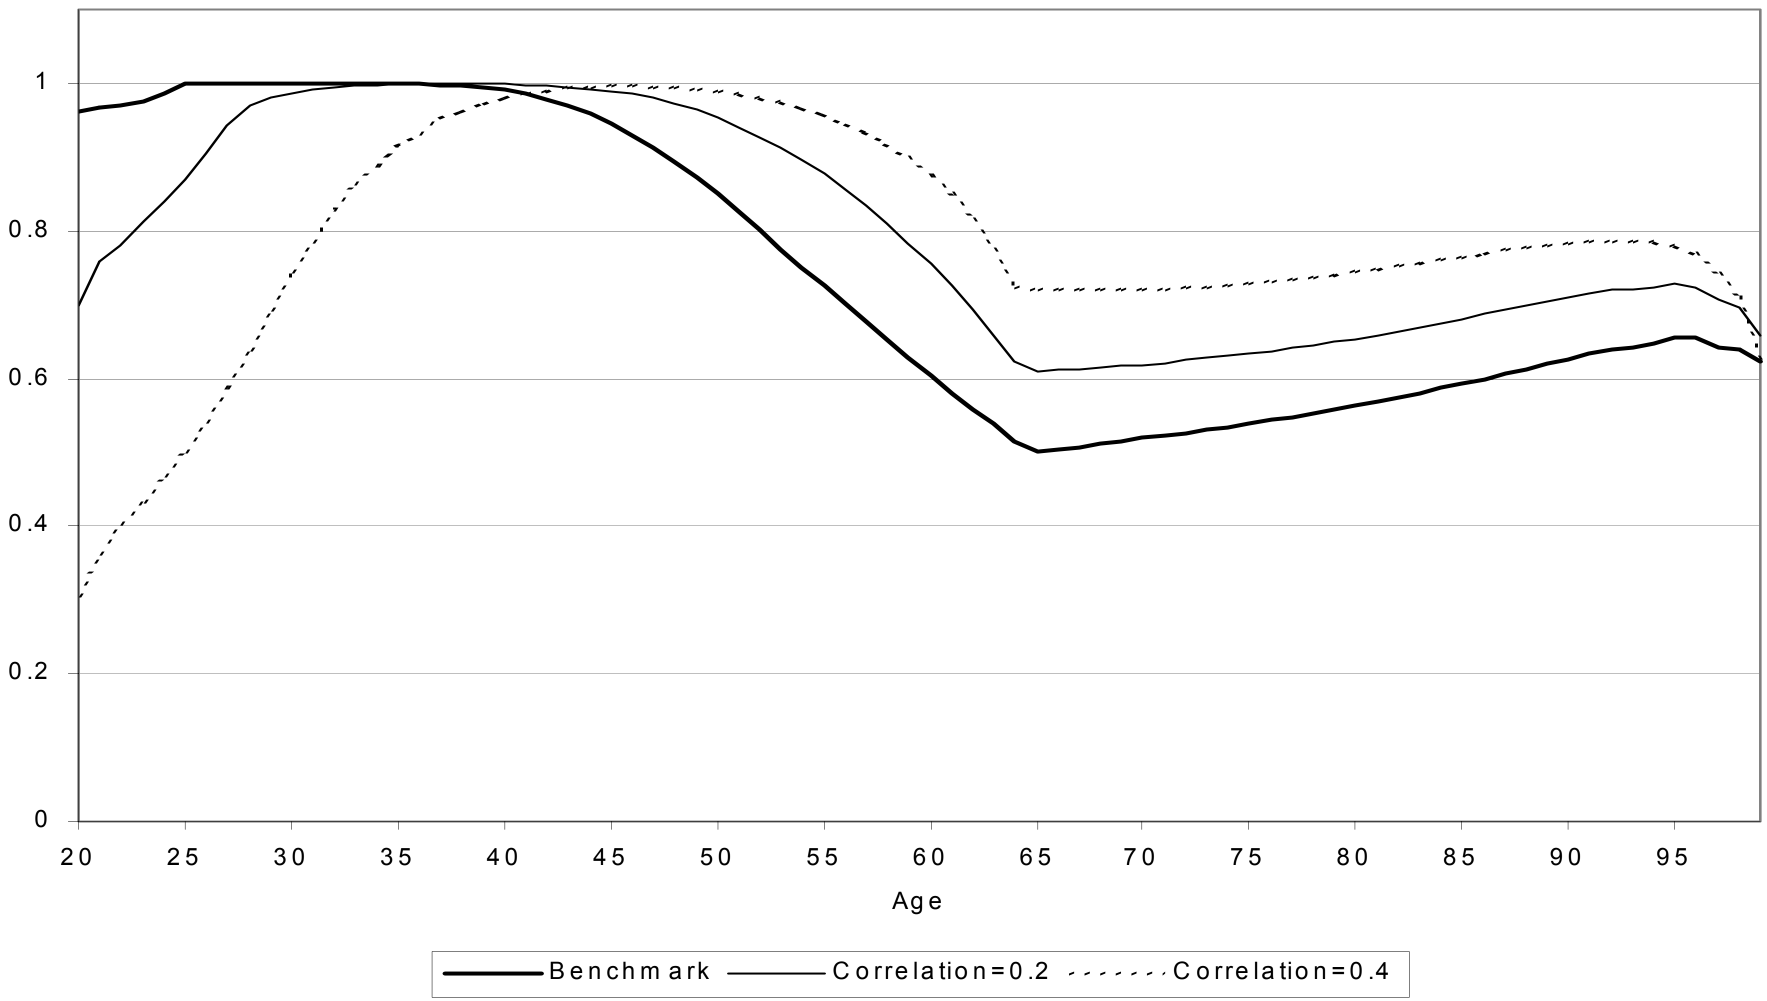

1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
Benchmark
Correlation=0.2
Correlation=0.4

Figure 5

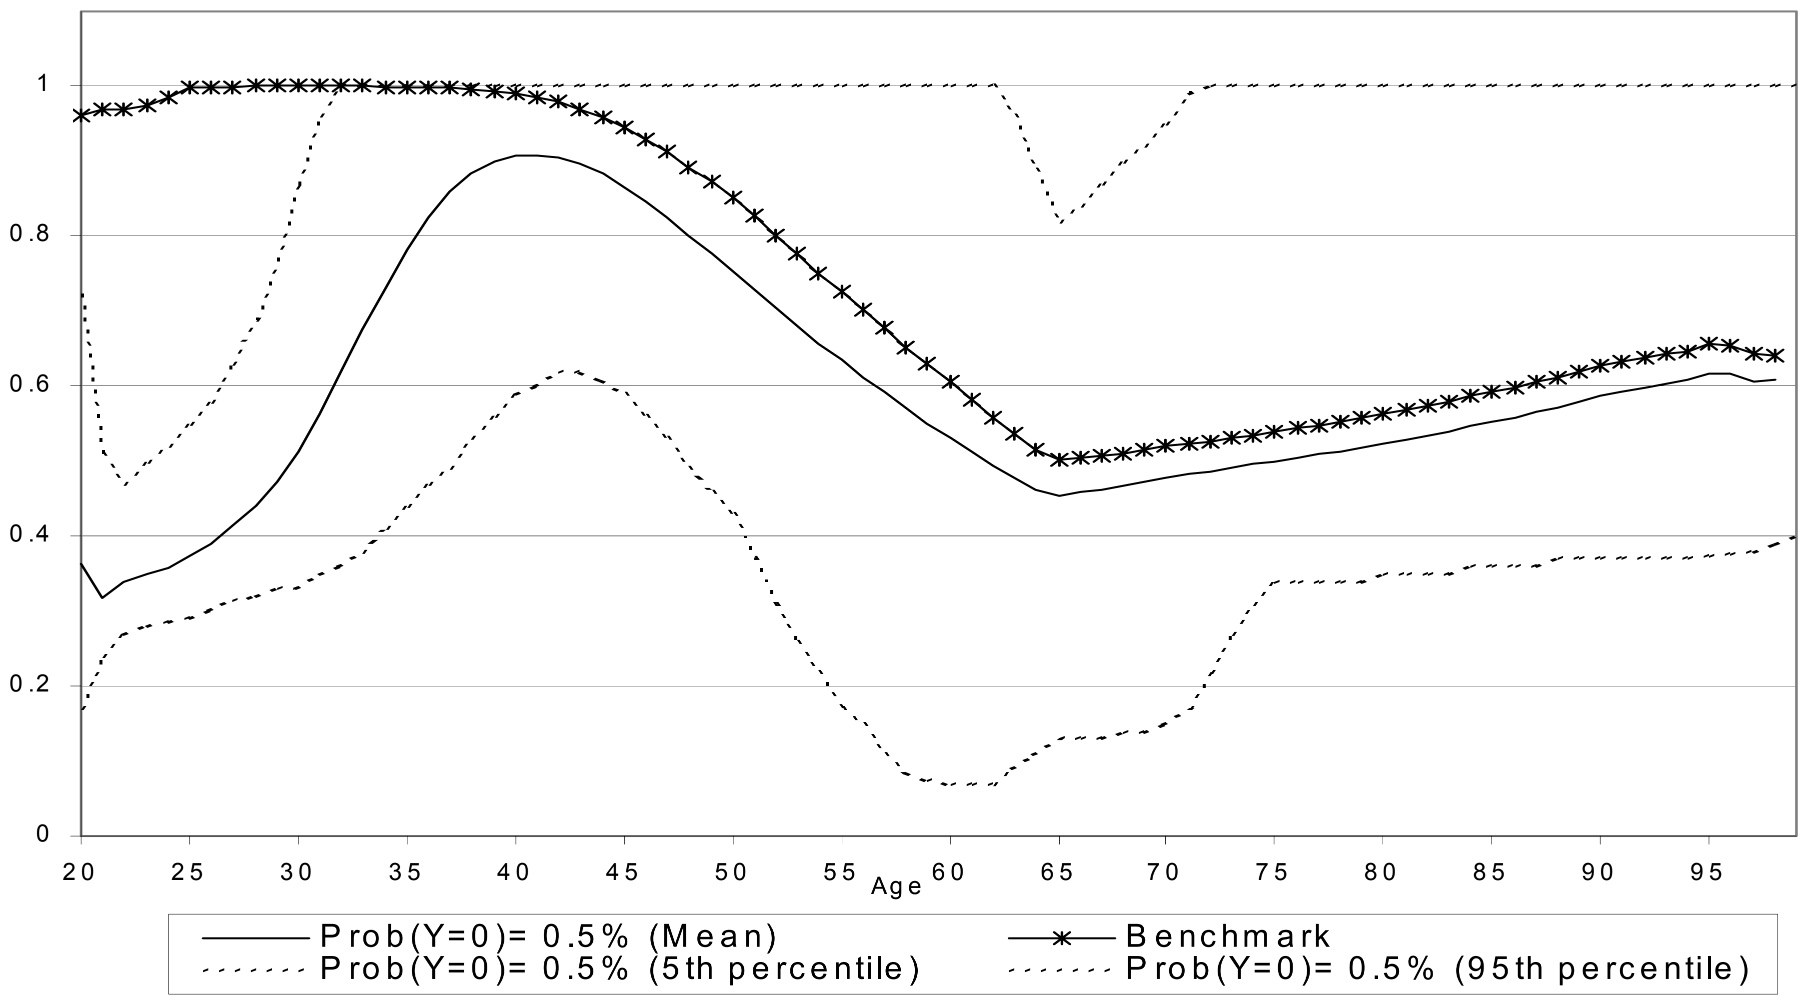

1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
Prob(Y=0)= 0.5% (Mean)
Benchmark
Prob(Y=0)= 0.5% (5th percentile)
Prob(Y=0)= 0.5% (95th percentile)

Figure 6

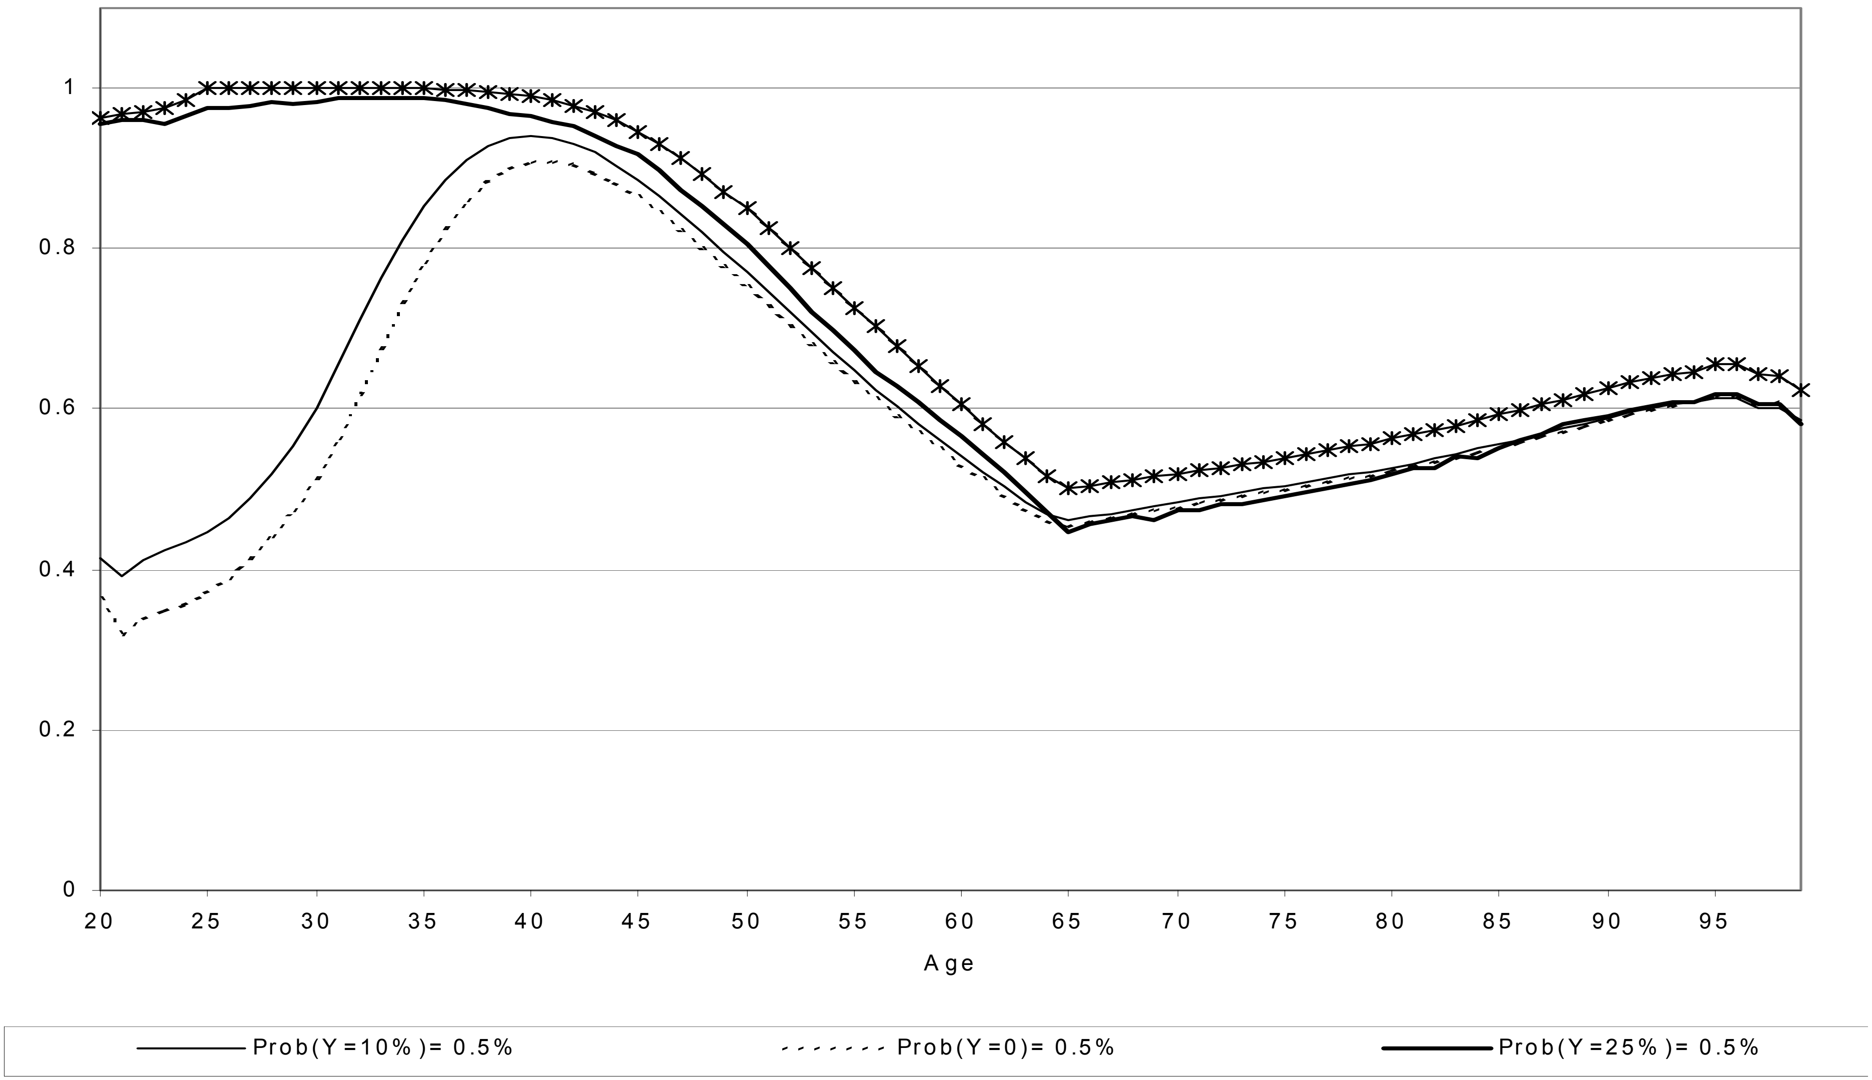

1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
Prob(Y=10%)= 0.5%
Prob(Y=0)= 0.5%
Prob(Y=25%)= 0.5%

Figure 7A

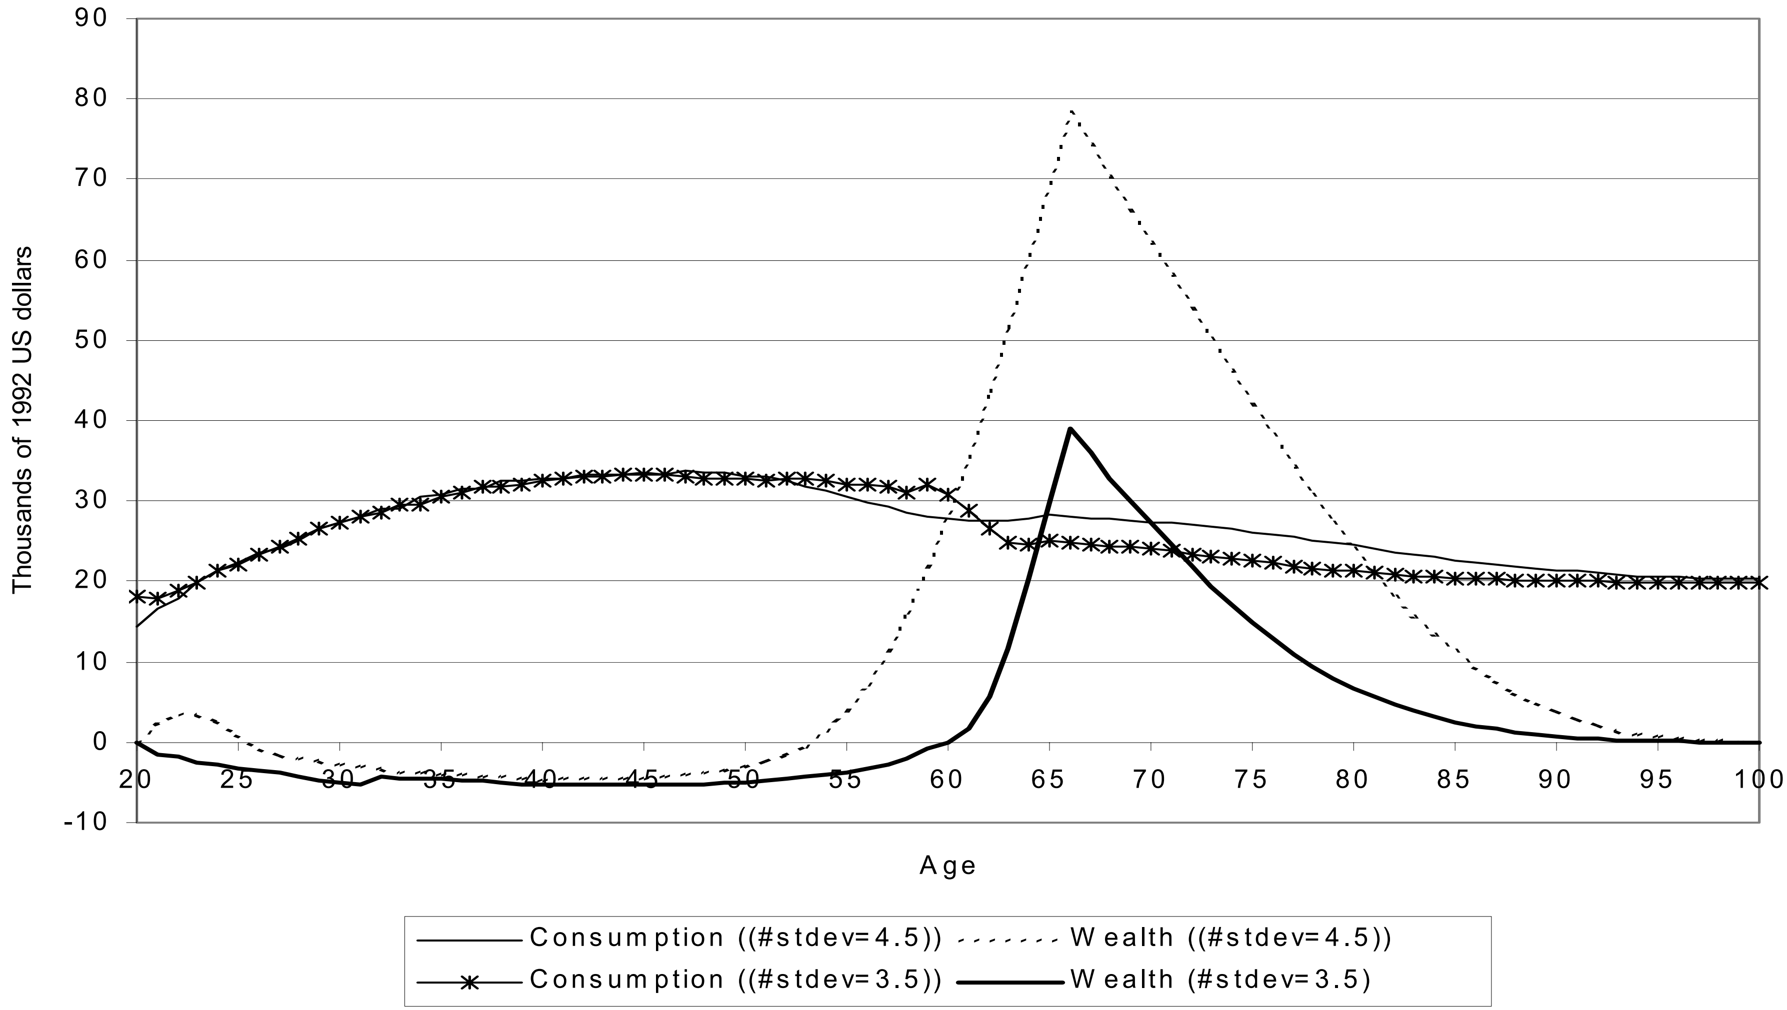

90
80
70
60
50
40
30
20
10
0
-10
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
Thousands of 1992 US dollars
Age
Consumption ((#stdev=4.5))
Wealth ((#stdev=4.5))
Consumption ((#stdev=3.5))
Wealth (#stdev=3.5)

Figure 7B

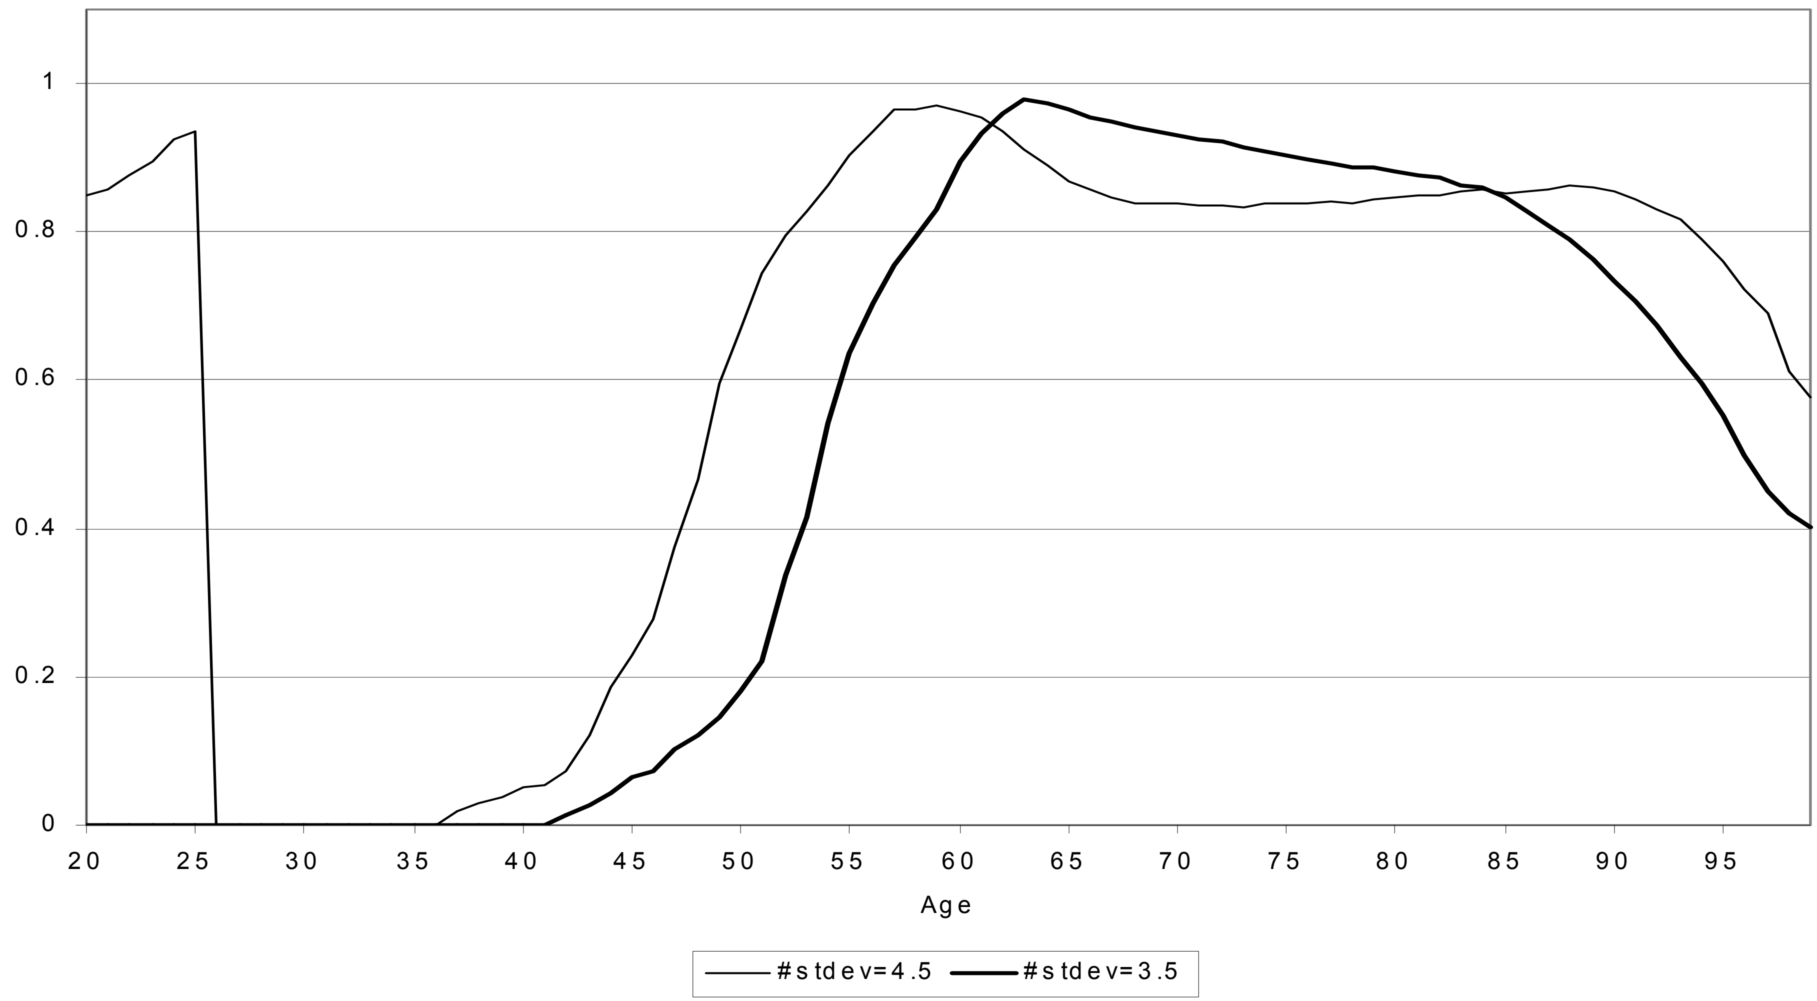

1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
#stdev=4.5
#stdev=3.5

Figure 8

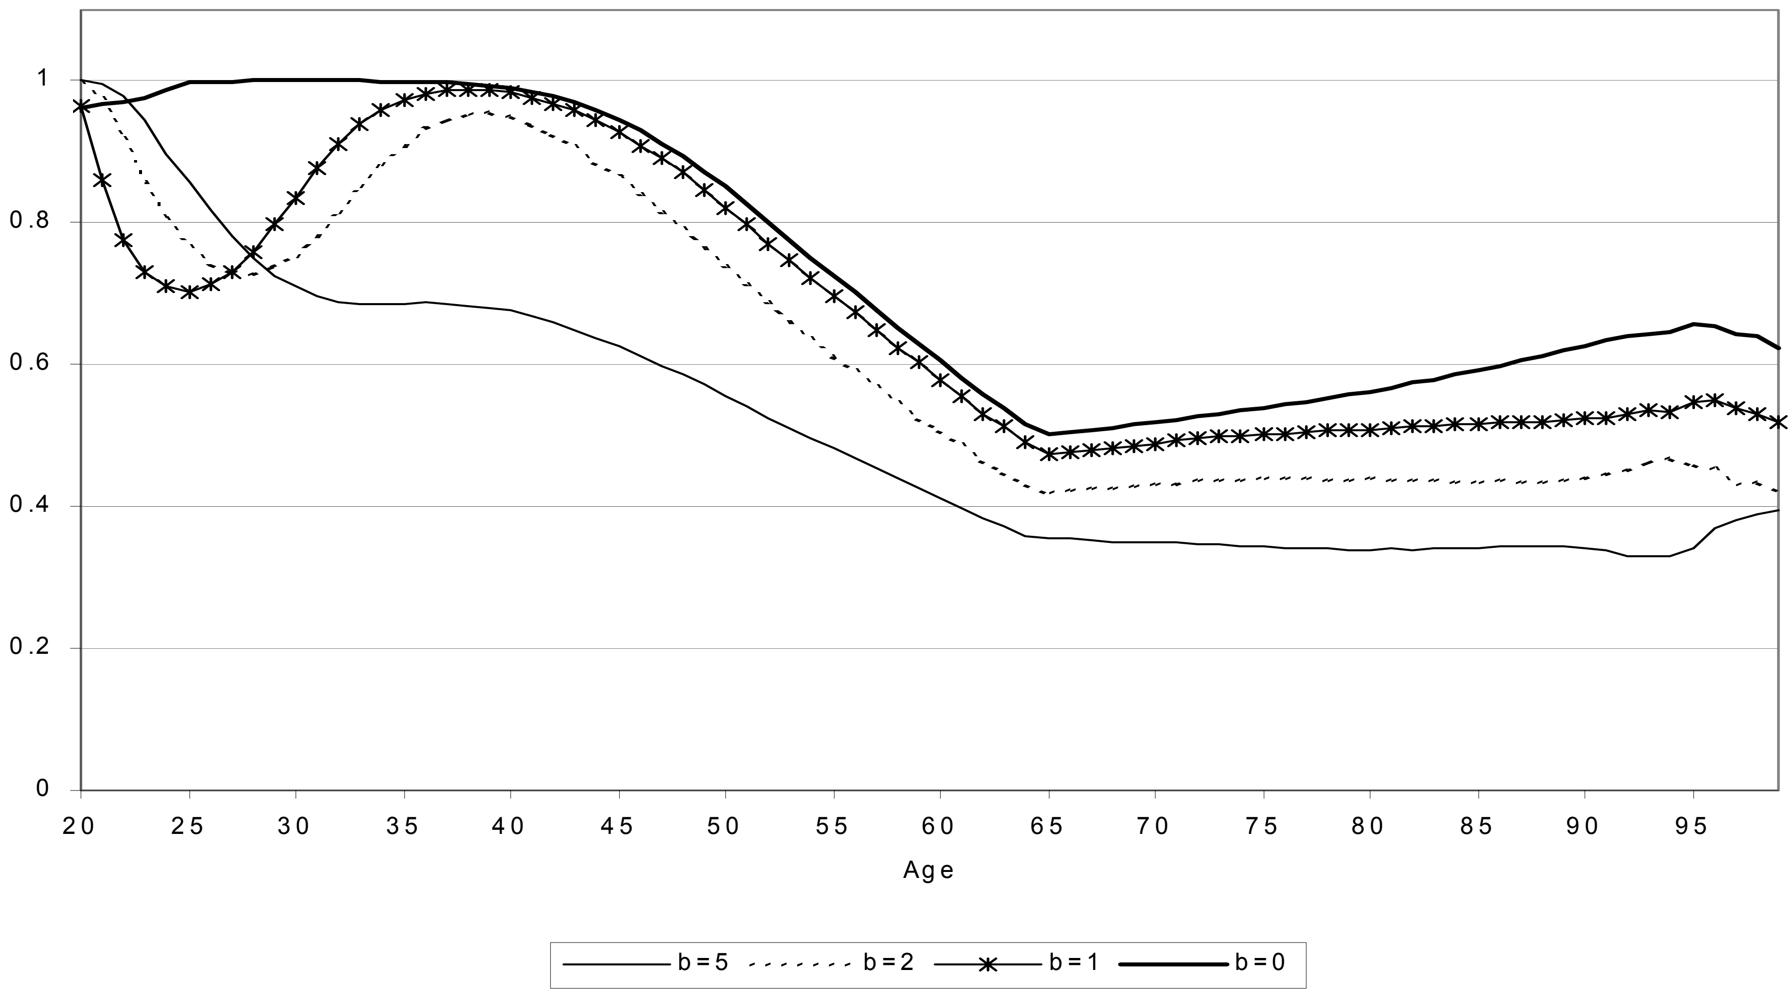
1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
b = 5
b = 2
b = 1
b = 0

Figure 9

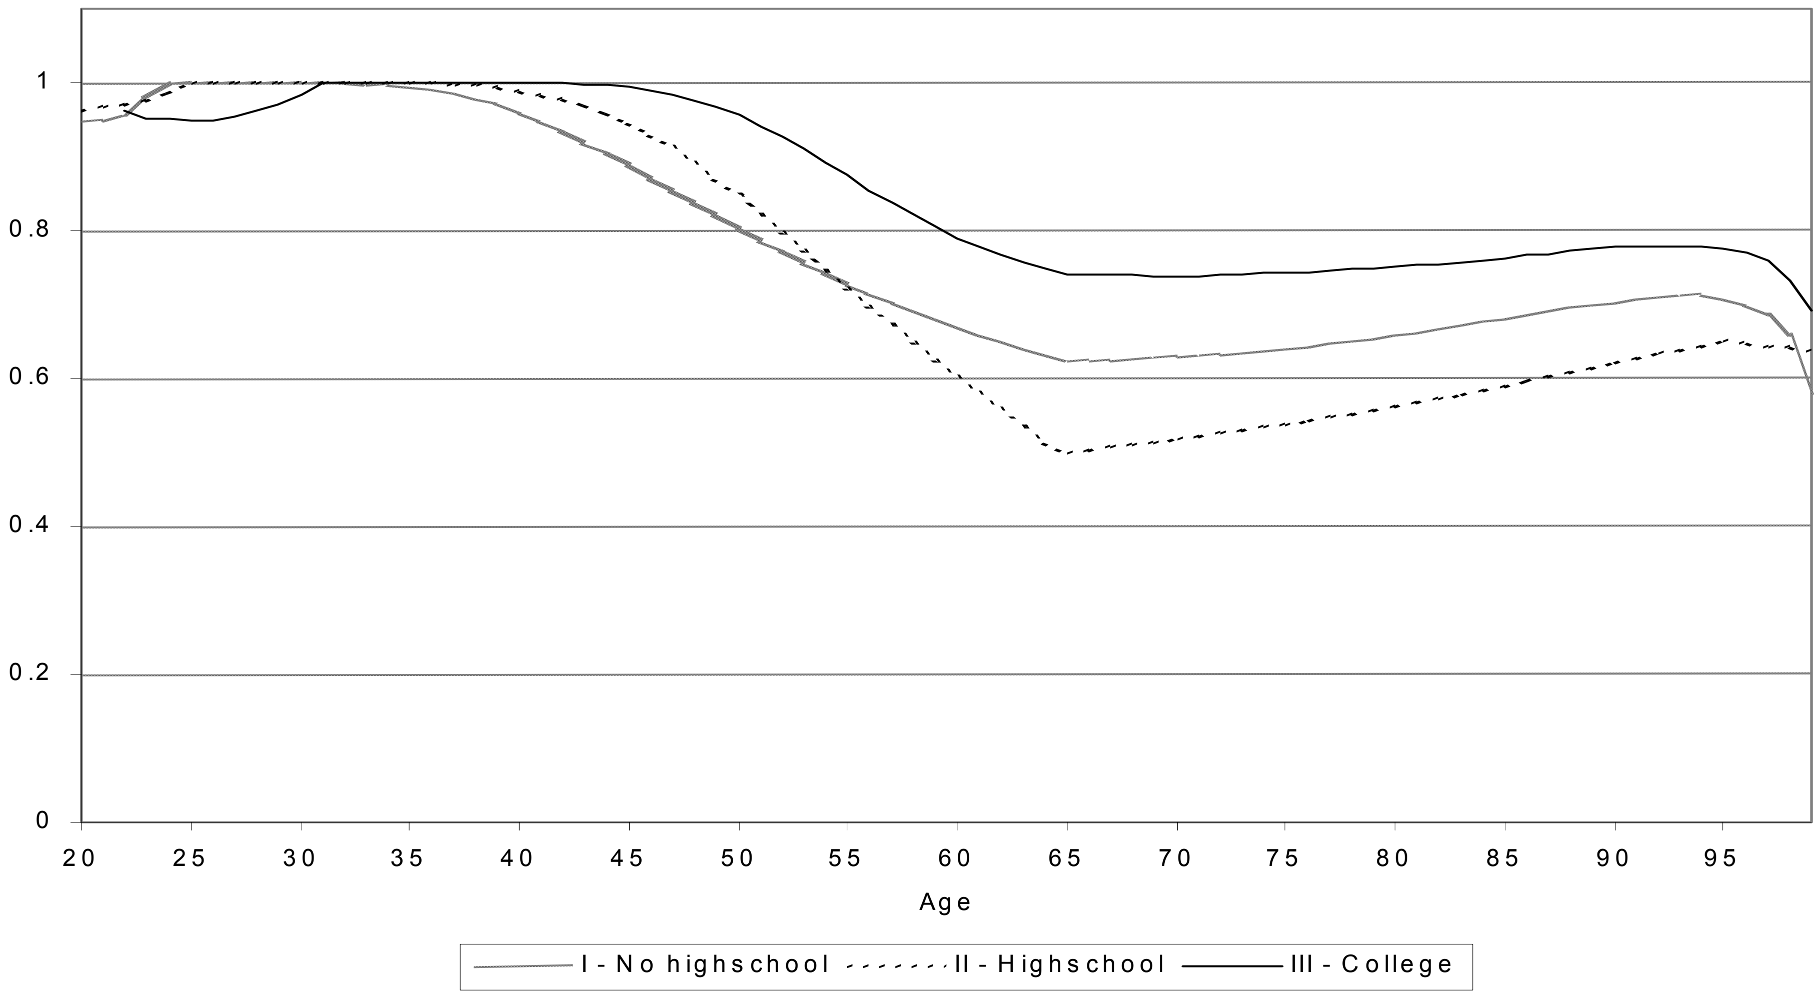
1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
I - No highschool
II - Highschool
III - College

Figure 10A

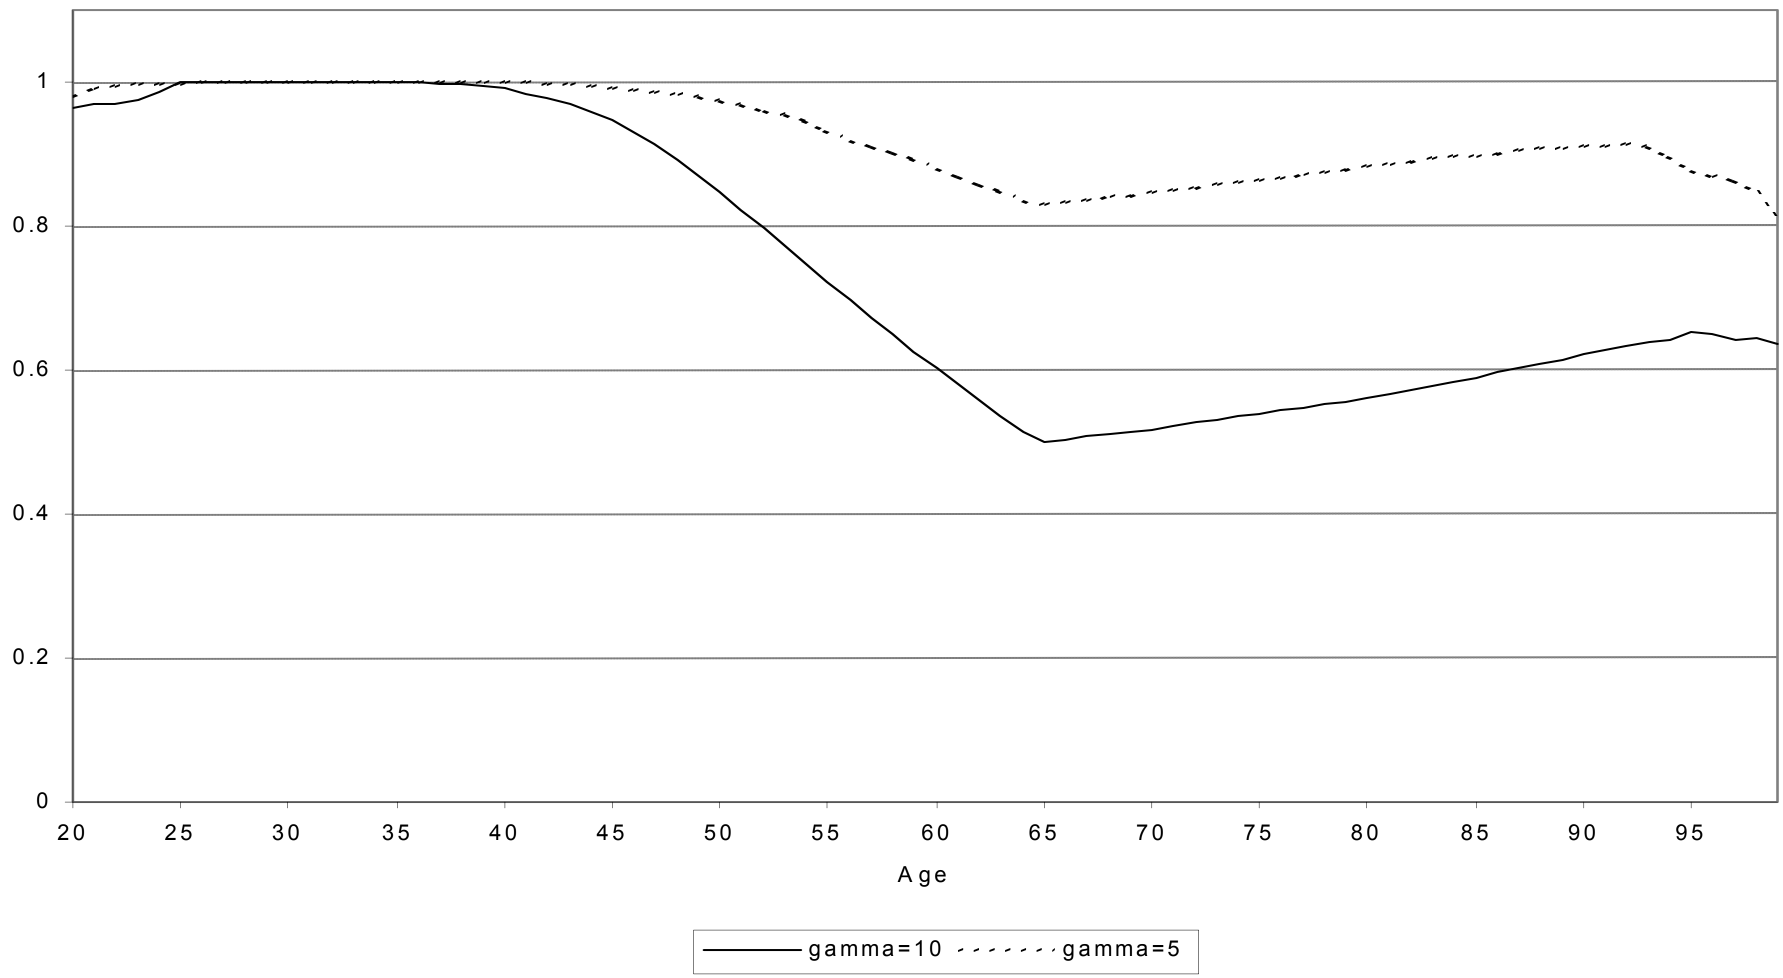

1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
gamma=10
gamma=5

Figure 10B

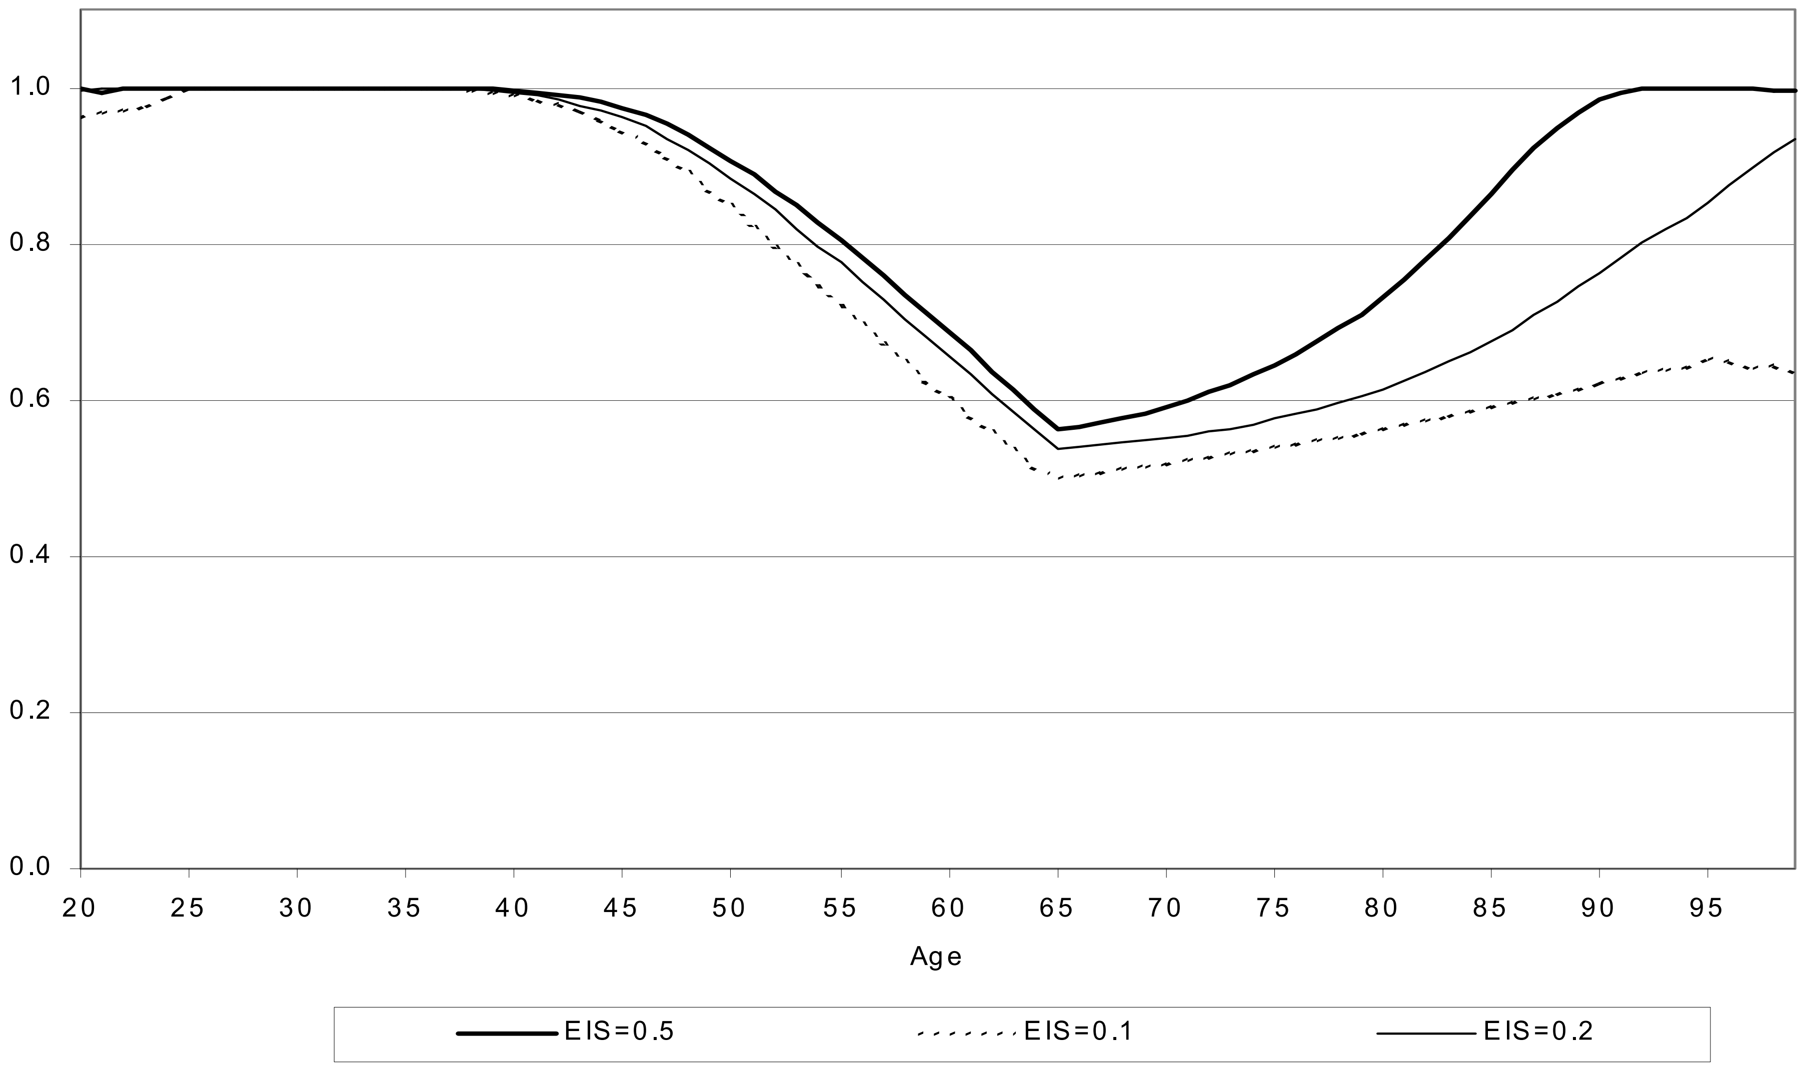
1.0
0.8
0.6
0.4
0.2
0.0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
EIS=0.5
EIS=0.1
EIS=0.2

Figure 11

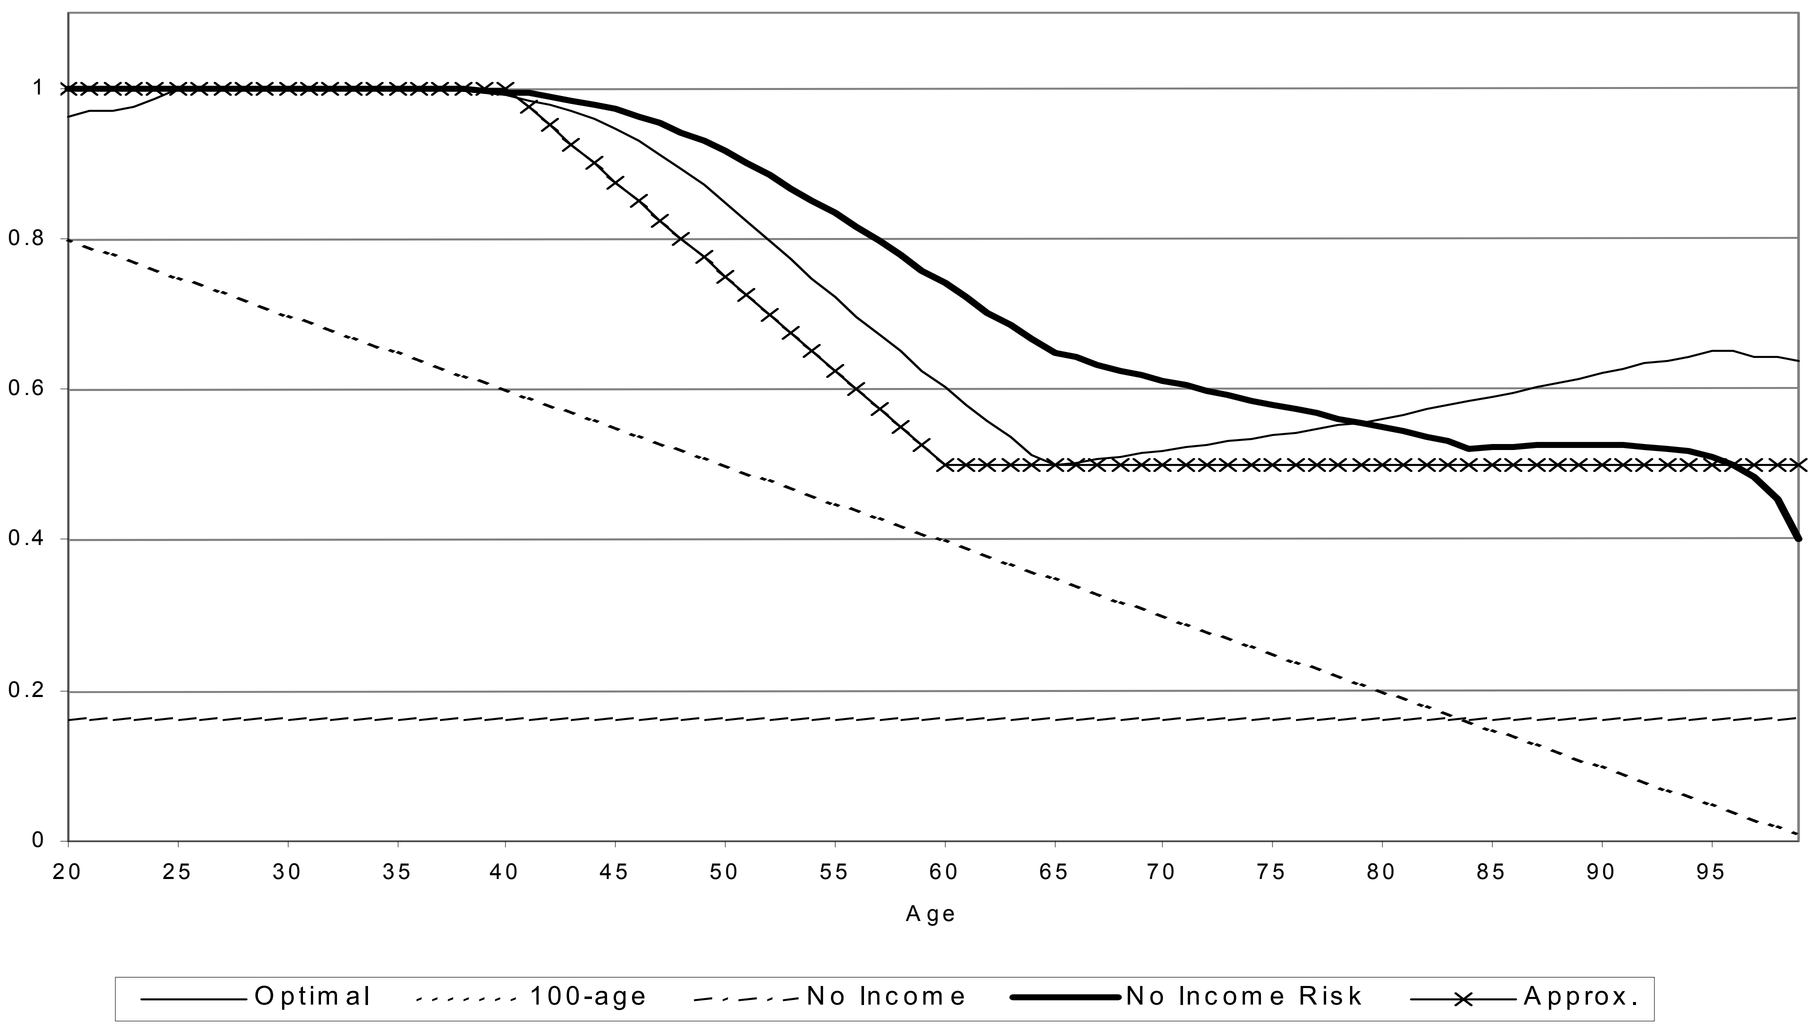

1
0.8
0.6
0.4
0.2
0
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
Age
Optimal
100-age
No Income
No Income Risk
Approx.