Stockholm School of Economics
Department of Finance
Master's Thesis

A Realistically Calibrated Life Cycle Model of Portfolio Allocation and Consumption with Stochastic Non-Financial Income

We analyze the effect of uninsurable stochastic non-financial income on the risky asset share in a life cycle model of portfolio allocation and consumption. The model is calibrated empirically using German data. We quantify the risk of non-financial income and show that this risk is largely idiosyncratic. Females face larger income shocks than males. Human capital has a significant effect on portfolio composition. Unless we allow non-financial income to drop to zero, the optimal risky share is close to 100% early in life as investors' wealth consists mostly of human capital. The optimal risky share decreases in subsequent years as agents build up financial wealth, but this effect partly reverses during retirement despite smaller human capital, because financial wealth is drawn down again. The utility costs of maintaining a constant risky asset share or of following rules of thumb for life cycle investing are economically large.

Key words: life cycle portfolio optimization, human capital, risky labor income, SOEP

Author: Johannes Haedicke (40016@student.hhs.se)
Tutor: Paolo Sodini
Presentation: 16 June 2011
Discussants: Svante Andersson, Tobias Helmersson

Note on authorship: This thesis has been written jointly by Andrejs Delmans (40008@student.hhs.se) and Johannes Haedicke. To allow work on the topic to continue as a thesis project, Andrejs Delmans does not appear as an author.

Acknowledgements: We would like to thank our tutor Paolo Sodini for his helpful advice and kind support.

Contents

1. Introduction

The seminal work of Merton (1971) gives a concise and intuitive answer to the optimal portfolio allocation problem in complete markets. The ideal weight of every risky asset in the portfolio is directly proportional to that asset's expected excess return over the risk-free rate and inversely proportional to its contribution to portfolio variance. In contrast to the simplicity of Merton's solution, real-world portfolio allocation advice for individuals and households is often more intricate.

This paper examines the optimal portfolio choice over the lifetime of agents who earn non-financial income and assesses the utility cost of deviations from that optimal allocation. For the most part, the analysis is based on the work by Cocco, Gomes, and Maenhout (2005), whose life cycle model is calibrated to U.S. micro data. We contribute by calibrating a model to German data. To the best of our knowledge, this is the first such calibration in the context of portfolio optimization outside the U.S.

In our model, non-financial income is exogenous and non-insurable so that agents cannot adjust their exposure to it. We calibrate the income process using data from the German Socio-Economic Panel Study (SOEP), the longest-running household panel in Germany. We obtain the common 'hump shape' in expected non-financial income profiles over agents' lifetime, albeit with some country-specific features. The risk characteristics of the income process are generally similar to those of U.S. income processes. Correlations between income innovations and stock returns are low. So, non-financial income is risky, but that risk is largely idiosyncratic. Interestingly, females face both larger income shocks and higher correlation between income innovations and stock returns than males.

We find that the value of human capital, an implicit non-tradable asset that yields a periodic 'dividend' in the amount of non-financial income, is large. On average, it accounts for more than 90% of total wealth, i.e. 90% of the sum of financial wealth and human capital, for agents younger than 45. The fact that human capital is large and non-tradable has a noticeable effect on the composition of the financial portfolio, and ignoring human capital is costly.

We identify three forces at work in shifting the risky asset share away from Merton's (1971) result. First, the motive of diversifying away the exposure to risky non-financial income creates demand for risky financial assets above the level predicted by Merton's optimal policy. Second, the amount of financial wealth, which we find to peak around the retirement age, affects the relative importance of human capital in the determination of the composition of the financial portfolio. Third, the riskiness of human capital decreases as agents approach retirement, which results in a more aggressive financial portfolio.

Given our calibration results and our choice of benchmark parameters, the model predicts that agents save at high rates early on, which is one of the main differences to Cocco, Gomes, and Maenhout's (2005) results, and that they draw down financial wealth quickly during retirement. The optimal risky asset share is 100% on average early in life unless we allow for a small probability of an extremely low income realization. The optimal risky share then decreases until midlife as the ratio of human capital to financial wealth falls, but it never falls substantially below 50%. It even bounces back before retirement as the remaining risk in human capital decreases rapidly in light of retirement income that is assumed riskless during retirement. Qualitatively, the asset allocation result is largely comparable to Cocco, Gomes, and Maenhout's (2005) result except for the pre-retirement rebound in the risky asset share.

We compare the utility achieved under different asset allocation and consumption policies to the utility achieved from following alternative policies on a constant consumption-equivalent basis. We find that ignoring the presence of labor income by following the classic Merton (1969) policy with a constant risky share or following a popular rule of thumb for life cycle investing (Malkiel, 1996) leads to utility losses equivalent to a reduction in constant certain consumption between 0.5% and 3%. We consider these losses economically significant.

Conclusions about the differences in optimal behavior between genders and education groups are impeded by the fact that the level of the optimal risky share is sensitive to changes in the estimation procedure for the variances of income innovations. Moreover, the results only hold under the assumption of a high coefficient of relative risk aversion.

The model's predictions do not match actual behavior. German households tend to save even during retirement (Börsch-Supan and Essig, 2005), and stock market participation rates as well as risky asset shares are very low (Barasinska et al., 2008, Börsch-Supan and Essig, 2002). If investors are rational and if the model employed here accurately incorporates investors' preferences and income processes, actual behavior leads to significant welfare losses. However, it is likely that the model does in fact not reflect all relevant features of investors' preferences and choices.

2. Theory and Previous Research

2.1. Classic portfolio optimization

Markowitz (1952) pioneers the formalization of the portfolio allocation problem. His framework is set up in a mean-variance space and is limited to a single investment period. Mossin (1968) and Samuelson (1969) show that under the assumptions of constant relative risk aversion (CRRA) utility, stationary asset return distributions, and complete and frictionless markets, optimal portfolio allocation is invariant to investment horizon, wealth, and optimal level of saving. Within the same setup, but in continuous time and with normally distributed returns, Merton (1969) finds the closed form solution to the problem for a two-asset investment universe as

$$w^* = \frac{\mu}{\sigma^2 \gamma}, \tag{1}$$

where w^* is the optimal share of the financial portfolio invested in the risky asset, γ is the coefficient of relative risk aversion, and μ and σ^2 are the first and second moments of the distribution of risky asset excess returns, respectively. The above solution is straightforwardly extendable to a multi-asset case. Assuming homogenous investor expectations, such an extension is, however, redundant: according to the mutual fund separation theorem, a case with many risky assets is equivalent to the two-asset case in Merton (1971), with the risk-free asset and the market portfolio.

The fact that investors earn labor and other non-financial income calls for the recognition of human capital, i.e. the present value of future non-financial income, as an embedded portfolio component. The efficient allocation to the risky asset, along the lines of Merton (1971), is then defined as

$$w^* = \frac{\mu}{\sigma^2\gamma}\left(1 + \frac{W^{HC}}{W^F}\right), \tag{2}$$

where W^F and W^{HC} stand for financial wealth and human capital, respectively. In the complete markets setup of Merton (1971), exposure to non-financial income could be replicated by some portfolio of marketable securities and turned into a synthetic W^{HC} position in the risk-free asset. Thus, everything else being equal, a larger exposure to non-financial income increases the optimal share of financial capital invested in the risky asset. This leads to the idea that as the value of human capital evolves over the life cycle, so does the optimal portfolio allocation. For example, Jagannathan and Kocherlakota (1996) argue that older people should invest less in stocks than younger ones since the present value of future labor income generally declines over the lifetime.

When exposure to non-financial income cannot be traded, either because markets are incomplete or because agents face liquidity constraints, Merton's (1971) portfolio allocation solution can no longer be used. For example, Gollier and Pratt (1996) show that exposure to independent uninsurable background risk induces CRRA investors to become more risk-averse with respect to risky assets. The optimal exposure to the risky asset portfolio is thus decreasing in the riskiness of human capital. This finding is confirmed by Elmendorf and Kimball (2000), Koo (1999), and Heaton and Lucas (2000), who estimate optimal saving and risky asset allocation for two-period, finitely, and infinitely lived investors, respectively.

2.2. Life cycle portfolio optimization

Contemporary life cycle portfolio optimization literature operates within a setup with finitely lived investors who are subject to uninsurable income risk. The allocation problem is defined in terms of maximizing the expected utility of consumption over the remaining lifetime given the specification of the income process, shape of the utility function, life span expectations, and investment opportunities. Such models do not normally have a closed form solution and are solved numerically by dynamic programming techniques. Although the mutual fund separation theorem does not hold when non-financial income is not tradable, two-asset models are common because of computational intensity.

To the best of our knowledge, Cocco, Gomes and Maenhout (2005) are the first to use an empirically calibrated non-financial income process in the context of life cycle portfolio optimization. Fitting a non-financial income process to U.S. data using the Panel Study of Income Dynamics (PSID), the authors discuss the evolution of realistic saving and portfolio allocation rules over the life cycle. The authors find that correlation of income shocks and stock returns is low and insignificant, implying that, in terms of its impact on the financial portfolio, human capital resembles an implicit holding of the risk-free asset rather than of the stock market portfolio. Cocco, Gomes and Maenhout (2005) estimate that it is optimal for investors to invest almost exclusively in the risky asset until approximately the age of 40. As the authors note, this is inconsistent with the actual allocation behavior of young households. Later in life, approximately until the assumed retirement age of 65, the optimal allocation to the risky asset gradually falls to around 50%. It rises again slightly during retirement, but overall, the optimal risky asset share is downward-sloping over the life cycle. The authors find that ignoring the effect of non-financial income on portfolio allocation leads to a loss equivalent to around 2% of annual consumption on a certainty-equivalent basis. The findings are robust to the introduction of a small probability of disastrous labor income shocks, uncertainty in retirement income, a bequest motive, Epstein-Zin utility, and endogenous labor supply.

The prediction of Cocco, Gomes and Maenhout's (2005) base case that investors should have large exposures to equities especially when young does not match the empirical evidence. First, most households were shown to hold at least part of their wealth in bonds (Poterba and Samwick, 2003). Second, nonparticipation in the stock market is common, particularly among young households (Heaton and Lucas, 2000). Only when the authors allow for disastrous labor income draws, the predicted optimal risky asset share for young investors decreases substantially.

Various other approaches to improving life cycle models with respect to this particular issue have been discussed. Gomes and Michaelides (2005) assume heterogeneity in investor preferences and a fixed cost that has to be borne by households when they first invest in the risky asset. In their setup, more risk-averse households start investing in the risky asset sooner, but hold more conservative portfolios. This matches the observed low stock market participation rates among young households. Davis, Kubler, and Willen (2006) use a wedge between borrowing and lending rates to reduce demand for equity. Polkovnichenko (2007) uses habit formation preferences to explain why young investors, who have not yet accumulated enough wealth to confidently sustain consumption above habit, invest more conservatively than middle-aged households. Benzoni, Collin-Dufresne and Goldstein (2007) explain late entry to the stock market by assuming cointegrated labor income and stock return processes, which results in human capital losing its 'stock-like' character as investors age.

Attempts to extend life cycle models so as to generally make them more realistic are diverse. One common approach is to allow for flexible labor supply. Bodie, Merton, and Samuelson (1992), Gomes, Kotlikoff and Viceira (2008), and Chan and Viceira (2000) add leisure time as another policy variable. The argument of the utility function is chosen as a linear combination or a Cobb-Douglas product of leisure and consumption. These papers find that labor supply flexibility acts as an insurance against financial losses and enables young investors to take significantly greater investment risks than the older investors.

Lynch and Tan (2009) develop a model in which the dividend yield predicts stock market returns and growth and volatility of labor income. Procyclical growth and countercyclical volatility of labor income explain small stock holdings of young investors with low ratios of financial wealth to income.

Munk and Sørensen (2010) add stochastic interest rates that follow Vasicek (1977) dynamics to the model of Cocco, Gomes and Maenhout (2005). They estimate the efficient allocation in cash, stocks, and long term bonds in the presence of risky labor income that varies over the business cycle.

2.3. Non-financial income process

It has become increasingly common in the literature on life cycle portfolio allocation and savings to calibrate a non-financial income process to longitudinal data. Elements of the processes employed have originally been developed in longitudinal studies of wages and earnings and have later been used for the calibration of life cycle models of savings and consumption. This section gives an overview of approaches and model specifications most often used.

Common specifications of non-financial income dynamics have a two-component structure: income consists of a deterministic and a stochastic component (MaCurdy, 1982). The deterministic component is often represented by a function of age and individual characteristics such as education and gender. With respect to the stochastic component, MaCurdy (1982) shows that common process specifications found in the literature (until the publication of his work) are all special cases of an ARMA process. The stochastic component is in turn often decomposed into a persistent, AR(1), and a transitory, MA(0), component. The persistent component is frequently modeled as having an autocorrelation coefficient of one, i.e. it is modeled as a random walk. Apart from distinguishing income innovations based on their persistence, authors sometimes distinguish stochastic elements of the income process based on whether they are aggregate or idiosyncratic (e.g. Benzoni and Chyruk, 2009).

Those process calibrations suitable for portfolio optimization that we are aware of are based on U.S. data. With differences in labor market structures and social security and welfare systems across countries, it appears to be a relevant question whether and how income processes in other countries differ from existing calibrations and whether the differences have an impact on optimal portfolio choice. A main contribution of this paper is to calibrate a non-financial income process to German data.

2.3.1. Deterministic component

There is a consensus that the deterministic component of non-financial income follows a hump-shaped trajectory over the life cycle. Non-financial income on average peaks during the middle age after a period of high growth in the early years of individuals' working lives. As retirement approaches, real income declines on average. During retirement, the deterministic part of non-financial income is relatively constant and generally lower than income before retirement. A smooth representation of the age-dependence of non-financial income is often obtained by fitting a polynomial to the loadings on age dummies from earnings regressions (for example, Hubbard, Skinner, and Zeldes, 1994, 1995, Storesletten et al., 2004, Cocco, Gomes, and Maenhout, 2005, Gomes and Michaelides, 2005).

Life cycle patterns of labor income are shown to differ depending on demographic parameters. Hence, many authors estimate processes separately for education groups (Hubbard, Skinner, and Zeldes, 1994, 1995, Gourinchas and Parker, 2002, Cocco, Gomes, and Maenhout, 2005, Gomes and Michaelides, 2005, Polkovnichenko, 2007) and/or occupation groups (Carroll, 1997, Davis and Willen, 2000, Gourinchas and Parker, 2002). Other characteristics such as household size or marital status are controlled for (Carroll and Samwick, 1997, Cocco, Gomes, and Maenhout, 2005). Cocco, Gomes, and Maenhout (2005) employ fixed effects to control for time-invariant heterogeneity among investors. Gender groups are formed in Davis and Willen (2000), but we are not aware of other income process calibrations for females. One contribution of our paper is to distinguish genders when calibrating the non-financial income process and to assess whether the results give rise to gender-specific optimal savings and portfolio allocation policies.

2.3.2. Transitory shocks

Transitory shocks are usually assumed to be i.i.d. normal or, following MaCurdy (1982), taken as an MA(2) as in Carroll and Samwick (1997). It is also possible to allow for a higher order MA process (Carroll and Samwick, 1997). Meanwhile, Hubbard, Skinner, and Zeldes (1994, 1995) caution that to the extent that transitory shocks may partly reflect measurement error, ignoring the variance of transitory shocks altogether prevents an overstating of income risk.

In some contributions, the distinction between transitory and persistent shocks is not made (e.g. Davis and Willen, 2000) and in other cases only permanent shocks are considered (Koo, 1998, Viceira, 2001).

Authors report varying estimates of the standard deviation of transitory shocks and find them to differ across education and occupation groups. The magnitude of estimated transitory shocks is economically high, and seems to generally decrease in the level of education. Cocco, Gomes and Maenhout (2005) estimate annualized standard deviations of 0.24 (i), 0.27 (ii) and 0.32 (iii) measures as a fraction of earnings for household heads with a college degree (i), with a high school degree but without a college degree (ii), and without a high school degree (iii), respectively. Gourinchas and Parker (2002) report standard deviations in the range of 0.18 to 0.26, and Heaton and Lucas (1997) estimate 0.24. Hubbard, Skinner, and Zeldes (1994) report estimates ranging from 0.12 to 0.20. Carroll (1992) estimates a standard deviation of 0.15 without distinguishing between education groups but uses only 0.10 in his model for reasons related to potential measurement error.

These estimates are not always directly comparable because they are found within different models and pertain to different sample periods; however, all of them are at least partly based on PSID data. In two cases, other U.S. datasets are used alongside the PSID.

One additional element of transitory income variability that authors often introduce is allowing for a possibility of zero or near-zero income observations. Hubbard, Skinner, and Zeldes (1994) obtain their somewhat below-average estimates despite including low income observations in their initial earnings regressions. All other analyses which we refer to above for standard deviations of transitory shocks exclude, as do several other authors, very low income observations when characterizing the distribution of transitory shocks and instead separately allow for the possibility of extremely low income draws based on the observed frequency thereof in the panel data.

2.3.3. Persistent stochastic component

The persistent stochastic component is commonly assumed to follow an AR(1) and to have normally distributed innovations which then are the permanent shocks to the income process. Many authors (e.g. Carroll (1992, 1997), Carroll and Samwick (1997), Gourinchas and Parker (2002), Cocco, Gomes, and Maenhout, 2005, Gomes and Michaelides, 2005, Polkovnichenko, 2007) restrict the coefficient of autoregression to 1, i.e. impose a random walk process. Authors usually justify the choice of the random walk over an AR(1) specification by referring to Hubbard, Skinner, and Zeldes' (1994) result that the first-order autocorrelation coefficient is close to one (0.95 – 0.96). Plus, relative to the assumption of an AR(1), the random walk assumption simplifies the solution of the consumption and portfolio choice model for which the income process is used. Campbell et al. (2001) argue in the context of a process specification that decomposes the permanent innovations into an aggregate and an idiosyncratic shock that the random walk assumption should have no material effect on the estimation of the optimal consumption and portfolio allocation policies.

Estimates of the standard deviation of permanent income shocks are generally smaller than those of transitory shocks. Cocco, Gomes, and Maenhout (2005) report standard deviations between 0.10 and 0.13 as a fraction of earnings, while Carroll's (1992) and Gourinchas and Parker's (2002) estimates over all education and occupation groups are around 0.15. Hubbard, Skinner, and Zeldes' (1994) estimates range from 0.13 to 0.18. In order to account for potential measurement error, Carroll (1992) again reduces the estimate from 0.15 to 0.10 when using it as parameter input to a buffer stock saving model. Unlike for transitory shocks, estimates of the variability of permanent shocks do not seem to be related to education.

2.3.4. Correlation with financial asset returns

In addition to the variability of shocks, the correlation between labor income shocks and innovations to risky asset returns is a factor in portfolio allocation decisions. The literature considers correlation between permanent labor income shocks and innovations to aggregate excess equity returns (Koo, 1998, Campbell et al., 2001, Cocco, Gomes, and Maenhout, 2005), correlation between occupation-level income innovations and aggregate equity returns, Fama-French factors and returns on industry portfolios (Davis and Willen, 2000) as well as correlation between human capital and equity market returns through cointegration of aggregate labor income and dividends (Benzoni, Collin-Dufresne, and Goldstein, 2007). In some cases, saving and portfolio choice models are based on hypothetic rather than empirically estimated correlations (e.g. Viceira, 2001). Results reported by authors who calibrate portfolio choice models realistically generally suggest that the contemporaneous correlation between income shocks and aggregate equity returns is close to zero (Cocco, Gomes, and Maenhout, 2005, Davis and Willen, 2000). Campbell et al. (2001) confirm this finding, but yet find a correlation of 0.15 between permanent labor income shocks and one-year lagged stock returns. Davis and Willen (2000) show that the Fama-French SMB factor is significantly correlated with occupation-level income shocks for a number of occupations in a setting where no distinction is made between transitory and permanent shocks.

3. Model

Most of the characteristics of the model we employ are chosen based on the model used in Cocco, Gomes, and Maenhout (2005).

3.1. Budget constraint

We specify the lifecycle consumption and saving model in discrete time, where agents' lives span for up to T periods. We set period length equal to one year, and throughout the rest of this paper refer to age t as the as the t-th year of life. Age as understood in this paper is one unit higher than the full number of years and individual has lived.

Agents start independently financing their consumption at age t_1 when they do not have any financial wealth. Up to age K, agents earn risky non-financial income, and, starting from age $K + 1$, they receive risk-free retirement income

($0 < t_1 \leq K \leq T$). At every point in time there is an infinite number of agents at all admissible ages.

In each period $t, t_1 \leq t \leq T$, an agent receives real non-financial income y_t. We suppress agent-specific subscripts throughout this section for clarity. We denote the sum of financial wealth and non-financial income realized in the beginning of period t by x_t and, following Deaton (1991), we call it cash on hand. In the beginning of period t, the agent chooses the amount c_t to spend on consumption in that period. After the amount c_t for consumption is set aside, the agent invests the remainder of cash on hand into financial assets. The investment yields a continuously compounded after-tax rate of return r_{t+1}^p in the beginning of period $t+1$. The intertemporal budget constraint then is

$$x_{t+1} = (x_t - c_t)e^{r_{t+1}^p} + y_{t+1}. \tag{3}$$

3.2. Investment constraints and financial asset universe

The agents allocate a proportion w_t of financial wealth to the risky asset and allocate the remainder to the risk-free asset. The risk-free asset yields a time-invariant continuously compounded after-tax rate of return r^f and the risky asset yields a continuously compounded after-tax rate of return r_{t+1} on investments from period t to $t+1$:

$$r_{t+1} = r^f + \mu + \eta_{t+1}, \tag{4}$$

where η_{t+1} is marginally distributed as $N(0, \sigma_\eta^2)$ so that the after-tax excess return on the risky asset follows geometric Brownian motion with a constant drift μ.

The portfolio return is

$$r_{t+1}^p = \ln\left(e^{r^f} + w_t\left(e^{r_{t+1}} - e^{r^f}\right)\right). \tag{5}$$

Agents face liquidity, borrowing, and short sale constraints. Thus financial wealth and positions in the investment assets are never negative:

$$x_t - c_t \geq 0 \text{ and } 0 \leq w_t \leq 1. \tag{6}$$

3.3. Non-financial income

Non-financial income is exogenous and non-tradable. The assumed income process is chosen to match Cocco, Gomes, and Maenhout's (2005) income process in most of the important aspects, but its form is standard and in that sense the process is also similar to the processes used in, for example, Polkovnichenko (2007), Hubbard, Skinner, and Zeldes (1994, 1995), or Gomes, Kotlikoff, and Viceira (2008). Log non-financial income in period t, before retirement, is

$$\ln(y_t) = f(t) + v_t + \varepsilon_t, \qquad t \leq K, \tag{7}$$

where $f(t)$ is a deterministic function of t that equals the expected value of an agent's log income at the corresponding age; $\varepsilon_{i,t}$ is an idiosyncratic temporary shock which is distributed as i.i.d. $N(0, \sigma_\varepsilon^2)$; v_t is a stochastic persistent component of income that follows

$$v_t = v_{t-1} + \zeta_t, \tag{8}$$

where ζ_t is marginally distributed as $N\left(0, \sigma_\zeta^2\right)$ and $v_{t_1-1} = 0$. The random walk assumption simplifies the solution of the life cycle model compared to the assumption of an AR(1). The innovation ζ_t to the stochastic persistent component is correlated with the innovation η_t to excess stock returns; the coefficient of correlation is $\rho_{\zeta,\eta}$. We do not decompose ζ_t any further as is done in Cocco, Gomes, and Maenhout (2005).

After retirement, non-financial income is constant and equal to a fraction λ of income before transitory shock in the last pre-retirement period:

$$\ln(y_t) = \ln(\lambda) + f(K) + v_K, \qquad t > K. \tag{9}$$

While ignoring a certain degree of uncertainty about real retirement income, the approach does account for the main properties of retirement insurance in Germany's social security system at least for the sample period used and it simplifies the solution of the life cycle model.

3.4. Optimization problem

In period $t, t_1 \leq t \leq T$, an agent's inter-temporal utility function is given by

$$U_t = \sum_{j=t}^{T} \delta^{j-t} \left(\prod_{k=t+1}^{j} p_k \right) E_t u(c_j), \tag{10}$$

where δ is a subjective risk-neutral discount factor, E_t is the expectation operator conditional on all information available in period t, and p_k is the probability of being alive in period k conditional on being alive in period $k-1$. Controlling for uncertainty in life expectancy in such a way, we implicitly assume that the force of mortality, p_k, is independent of financial or non-financial income realizations. We choose power utility for the indirect utility function, where γ is the coefficient of relative risk aversion:

$$u(c_t) = \begin{cases} \dfrac{c_t^{1-\gamma}}{1-\gamma}, & c_t > 0 \\ -\infty, & c_t \leq 0 \end{cases} \quad ,\gamma > 0. \tag{11}$$

The optimal saving and allocation problem faced by the agent in period t is formalized as maximizing U_t in (10) subject to constraints (3)-(9) and (11). The set of policy variables is $\{c_j, w_j\}_{j=t_1}^{T}$ and the state variables are t, $\{x_j, v_j\}_{j=t_1}^{T}$. The problem can be stated in the form of a stochastic dynamic programming equation:

$$V_t(x_t, v_t) = \max_{c_t, w_t} \{ u(c_t) + \delta p_{t+1} E_t \left(V_{t+1}(x_{t+1}, v_{t+1}) \right) \}, \tag{12}$$

where $V_t(x_t, v_t) = \max_{\{c_j, w_j\}_{j=t}^{T}} U_t$. Given $V_{t+1}(x_{t+1}, v_{t+1})$ and the contemporaneous state variables $\{x_t, v_t\}$, solving (12) amounts to optimization with respect to only the contemporaneous policy variables $\{c_t, w_t\}$ for each t, $t_1 \leq t$, at a time. Since consumption only up to period T matters, $V_\tau(x_\tau, v_\tau)|_{\tau > T} = 0$ and $V_T(x_T, v_T) = u(x_T)$, i.e. agents consume all their remaining wealth in the last period. Knowing this, we find $V_{T-1}(x_{T-1}, v_{T-1})$ by numerically solving (12) and continue the induction backwards until we find $V_t(x_t, v_t)$.

Here we give a brief description of the numeric solution method that we use to solve (12). First, using the fact that $V_t(x_t, v_t)$ is homogeneous with respect to x_t, we normalize x_t and c_t by e^{v_t}:

$$\dot{x}_t = x_t e^{-v_t},$$

$$\dot{c}_t = c_t e^{-v_t}.$$

Doing so, we withdraw the v_t dimension from the state space and reduce computational intensity. The Bellman equation can then be expressed in terms of the modified value function $\phi_t(\dot{x}_t)$ that is equal to $V_t(x_t, v_t)e^{-v_t(1-\gamma)}$ (please refer to Appendix 1 for the proof):

$$\phi_t(\dot{x}_t) = \max_{\dot{c}_t, w_t}\left\{u(\dot{c}_t) + \delta p_{t+1} E_t\left(\phi_{t+1}(\dot{x}_{t+1}) e^{\zeta_{t+1}(1-\gamma)}\right)\right\} \tag{13}$$

Starting with $t = T - 1$ we approximate the next period's value function, $\phi_{t+1}(\dot{x}_{t+1})$, by discretizing it over the grid of $\dot{x}_{t+1}$ with grid gaps following geometric series. Using an unequally spaced grid allows us to cover a larger grid range with fewer grid points without compromising the relative precision of the solution. We use third-order spline interpolation in ${\dot{x}_{t+1}}^{1-\gamma}$ in-between the grid points for a closer approximation. We also approximate the multivariate normal distribution of η_{t+1}, ζ_{t+1}, and ε_{t+1} by a discrete distribution with equally probable realizations, as proposed by Adda and Cooper (2003), to estimate the stochastic integrals necessary for finding the expected value in (13). Finally, we optimize (13) numerically using grid search over the policy variables $\dot{c}_t$ and w_t.

We continue by incrementing t by -1 and repeating the steps above until t reaches t_1. At each step of induction, we record the estimated optimal consumption and allocation policies as functions of cash on hand: $\dot{c}_t^*(\dot{x}_t)$ and $w_t^*(\dot{x}_t)$, respectively.

4. Calibration

4.1. Data source

For the calibration of the income process, we use data on individual and household income and other variables from the German Socio-Economic Panel Study (SOEP). The SOEP is a research-driven, representative household panel study focused "on the analysis of the life course and well-being [...] measured by the concepts of income and life satisfaction" (Wagner et al., 2007). It is run under academic direction by the German Institute of Economic Research (DIW Berlin). The study contains, among other things, detailed information on individual and household income over extended periods of time. To date, the SOEP encompasses 26 years of annual survey responses (1984 – 2009), making it the "longest-running longitudinal survey of private households and persons in the

Federal Republic of Germany" (Frick, 2010). Interested readers find a discussion of the purpose, structure and development of the panel study in Wagner et al. (2007).

To obtain our sample, we use a SOEP data release containing income information for calendar years 1983 to 2008 (Socio-Economic Panel (SOEP), Data for years 1984-2009, Version 26, SOEP, 2010). Specifically, we obtain most of the required data from the beta release of a long-format version of the SOEP contribution to the so-called Cross-National Equivalent File (CNEF). The CNEF aims at facilitating comparison among household panels in several countries, including the PSID in the U.S., by harmonizing data through consistent variable definitions (see Frick et al. (2007) for a detailed description of the CNEF project). We do not access the CNEF directly but rather use SOEP data in CNEF format, the differences being that we access the full SOEP sample as opposed to the 95% sample included in the CNEF and that we have extended income information available.

4.2. Sample and variable definitions

We estimate $f(t)$, σ_ζ, σ_ε, λ, and $\rho_{\zeta,\eta}$ by calibrating the non-financial income process defined in (7) - (9) to the panel data sample described above. We start from a sample of German households that contains information both at the household level and at the individual level. The raw SOEP sample is non-representative due to selection by design and selective panel attrition. We take this into account in the empiric calibration of the income process by weighting observations by the inverse sampling probability. When we use series of differenced observations, we apply longitudinal weights that take into account attrition probabilities.

4.2.1. Sample identification and restriction

We identify observations by a combination of a unique personal identification number i and by the maximum age t achieved during a particular calendar year. We only consider observations pertaining to individuals who are a head of a household and whose age is between 20 and 100 or between 25 and 100 depending on education. We further require that reported or imputed household-level and individual-level income items are available and that valid educational achievement information is reported. Below, we describe the reasons for applying these restrictions and the consequences for the composition of our final sample.

First, we find that our measures of income risk are sensitive to the inclusion of observations with imputed income information. Including imputed income information increases the estimated variance of transitory shocks and lowers the estimated correlation of persistent shocks with excess stock returns. This could either be a consequence of the imputation procedure, which would be an undesirable effect, or it could be due to the risk characteristics of the income processes of those survey participants who fail to (fully) disclose income information, in which case imputed income information should not be excluded from the analysis. We exclude observations for which more than 25% of household-level post-government income is imputed. Table 1 shows descriptive statistics for the dropped observations as well as for the final sample. The observations we lose pertain to households whose heads are younger and more likely to be female than those in our final sample. Plus, the dropped households tend to be slightly larger and better educated and receive higher non-financial income.

Table 1 Descriptive statistics (1)

Variable	Final sample		Observations with high more 25% imputation in household income		Eastern federal states during 1990s	
	Estimate	(SE)	Estimate	(SE)	Estimate	(SE)
Number of observations	151,232		19,335		15,544	
Average age	52.5	(0.046)	47.9	(0.117)	52.4	(0.143)
Proportion female	0.371	(0.001)	0.480	(0.004)	0.514	(0.004)
Proportion married / living with partner	0.501	(0.001)	0.547	(0.004)	0.516	(0.004)
Average number of persons in household ex. head/spouse	0.619	(0.003)	0.716	(0.007)	0.596	(0.007)
Average household non-financial income in 2006 €	26,163	(42.8)	31,163	(151.3)	21,937	(98.2)
Proportion education low	0.193	(0.001)	0.160	(0.003)	0.110	(0.003)
Proportion education middle	0.656	(0.001)	0.659	(0.003)	0.654	(0.004)
Proportion education high	0.151	(0.001)	0.181	(0.003)	0.237	(0.003)

The table shows descriptive statistics for the final sample and for two groups that are dropped from the original sample. Observations underlying the right panel pertain to heads of households located in eastern federal states (i.e. former East Germany) between 1991 and 1999. The middle panel concerns observations for which more than 25% of household post-government income is imputed.

Second, while income information pertaining to households located in those federal states formerly belonging to East Germany are available in the SOEP from calendar year 1991 onwards, we start using observations for eastern federal states only in 2000. During the 1990s, differences in real income between eastern and western federal states decreased as real income growth in the East was considerably stronger than in the West (Grabka, 2000). While real income continues to be lower in the eastern part of Germany today, we argue that the structural factors underlying both the large initial income differential and the subsequent income development do not reflect the factors driving income processes today and in the future. The argument is supported by the fact that the shapes of the age-earnings profiles discussed below are sensitive to the inclusion of eastern German observations. We therefore exclude income observations for eastern federal states from 1991 to 1999. The end of this exclusion period is chosen simply as the year when the German Federal Statistics Office merged the CPI indices for former East and West Germany. Table 1 documents the differences between sample characteristics for the excluded eastern observations and for the final sample.

Third, inclusion of observations pertaining to self-employed household heads increases the estimated variability of income shocks despite the relatively small size of this subgroup. As is sometimes done in the literature (Carroll, 1992; Davis and Willen, 2000), we exclude these observations. Table 2 shows that this leads to a loss of a group of households with below-average aged and predominantly male heads who report considerably higher income than households with heads who are not self-employed.

Table 2 Descriptive statistics (2)

Variable	Final sample		Self-employed		Observations with insufficient education information	
	Estimate	(SE)	Estimate	(SE)	Estimate	(SE)
Number of observations	151,232		7,687		2,217	
Average age	52.5	(0.046)	46.2	(0.129)	51.0	(0.434)
Proportion female	0.371	(0.001)	0.180	(0.004)	0.504	(0.011)
Proportion married / living with partner	0.501	(0.001)	0.601	(0.006)	0.395	(0.010)
Average number of persons in household ex. head/spouse	0.619	(0.003)	0.916	(0.013)	0.502	(0.019)
Average household non-financial income in 2006 €	26,163	(42.8)	40,385	(291.2)	22,111	(326.0)
Proportion education low	0.193	(0.001)	0.096	(0.003)	N/A	N/A
Proportion education middle	0.656	(0.001)	0.723	(0.005)	N/A	N/A
Proportion education high	0.151	(0.001)	0.181	(0.004)	N/A	N/A

The table show descriptive statistics for the final sample and for two groups that are dropped from the original sample. In the middle panel, statistics for households with self-employed household heads are shown. The right panel presents the available statistics for observations that are not matched with sufficient education information.

Forth, we do not consider observations without accompanying education information. The small subgroup that is dropped as a result of this restriction exhibits low income, a relatively high degree of female household heads and relatively small household size (Table 2).

Finally, when constructing our measure of non-financial disposable household income (see below), we drop a significant fraction of observations for retirement income recipients. The reason for this is that for calendar years 1983, 1984, 2001 and 2002, the level of aggregation in retirement income items prevents a reliable separation of financial and non-financial income and/or a reliable allocation of taxes to these items. In Table 3, we present descriptive statistics for the dropped observations and, for comparison, for all retained observations with positive retirement income. Given the similarity of the estimates, we conclude that the observations are missing at random from the subgroup of retirement income recipients.

Table 3 Descriptive Statistics (3)

Variable	Retirement income recipients in final sample		Dropped retirement income observations	
	Estimate	(SE)	Estimate	(SE)
Number of observations	55,097		7,148	
Average age	70.2	(0.046)	69.5	(0.126)
Proportion female	0.477	(0.002)	0.470	(0.006)
Proportion married / living with partner	0.422	(0.002)	0.447	(0.006)
Average number of persons in household ex. head/spouse	0.200	(0.003)	0.236	(0.008)
Proportion education low	0.272	(0.002)	0.284	(0.005)
Proportion education middle	0.627	(0.002)	0.626	(0.006)
Proportion education high	0.101	(0.001)	0.090	(0.003)

The table show descriptive statistics for dropped observations with positive retirement income and for all retained observations with positive retirement income.

Overall, for the population of households with wage- and salary-earning heads not earning their income in eastern federal states during the 1990s, it is likely that female heads and young heads are underrepresented as a result of the applied restrictions. As we calibrate the non-financial income process separately for male and female household heads, the loss of observations for females is in itself not a major shortcoming. It appears, though, that for the female subsamples, and to a lesser extent also for the male subsamples, we are disproportionally dropping observations pertaining to both higher-income and lower-income households.

4.2.2. Subsamples

We split the sample into six subsamples by gender and the level of education. Differences in the life cycle income profiles among different education groups are significant for the U.S. (Cocco, Gomes, and Maenhout, 2005), and works based on German data suggest the same at least in terms of income levels for males (see Dustmann and van Soest (1998) for an example based on SOEP data). Regarding differences in earnings profiles between genders, gender wage gaps are widely discussed (Rosenfeld, Trappe, and Gornick, 2004; Beblo and Robledo, 2008).

We call the three education groups "low", "middle", and "high", according to the highest level of education achieved at the point in time when an income realization is reported. Low-level education refers to inadequate completion of basic schooling and to degrees from tier-2 and tier-3 high schools[1]. Middle-level education encompasses all forms of vocational degrees as well as tier-1 high school degrees (if they are the highest education achieved), and high-level education means that higher education has been completed. These education groups largely correspond to the broad education groups in the CASMIN classification ("Comparative Analysis of Social Mobility in Industrial Nations"), with the exception that we assign basic vocational degrees added to a tier-3 high school degree to the middle group and move tier-2 high school degrees to the group with low education. In the CASMIN scale, they are assigned to the low and middle level, respectively.

For the low and middle education groups, households heads aged 20 to 100 are considered, i.e. $t_1 = 20$ and $T = 100$; for the high-education group, $t_1 = 25$ and $\mathrm{T} = 100$.

4.2.3. Income definition

Following Cocco, Gomes, and Maenhout (2005), we include a wide range of income items in the calibration. We define household disposable non-financial income as the after-tax sum of all forms of labor income, including infrequent

[1] The German high school system distinguishes three types of high schools according to the educational achievement they are geared towards. For the purpose of our education classification, we refer to these schools as tier-1, tier-2 and tier-3 high schools.

payments such as bonuses and severance packages, household public transfers[2], private transfers and all social security pension income for household members aged 16 or older. Additionally, we include private pension income related to former employment, because - unlike other private pension income – it is not the outcome of an explicit savings and portfolio allocation decision but rather compensation for prior service provision and hence a form of labor income. The level of detail at which income information is provided generally allows us to make the necessary distinctions.

The preferences in the life cycle model are defined for individuals rather than for households, meaning that expected non-financial income fed into the optimization should be the equivalent of individual income in terms of the level of income. Yet, when calibrating the non-financial income process, we start with household income so as to capture risk sharing effects implicit in household formation, particularly in head/spouse relationships. Besides, several income items are only reported at the household level and allocating these items to different individuals within the household cannot be done reliably. A minor disadvantage of this approach is that the income of children, relatives and non-relatives with at least 16 years of age living in a household together with the household head (and possibly spouse) is included in our household income measure. By controlling for household characteristics when estimating the age-earnings profile, we seek to obtain income predictions for the life cycle model representing the equivalent of expected individual income for a person with a particular gender, age and education.

However, for certain pension income items during a few survey years, the necessary distinctions cannot be made for reasons of survey structure; as mentioned above, we drop the affected observations. In Appendix 2, we detail the income components available in the SOEP and show how we classify them.

[2] We exclude housing-support for owner-occupiers from public transfers. Given the subsidy's design, it constitutes a discount on an asset purchase and hence concerns portfolio allocation rather than non-financial income.

4.2.4. Estimating disposable non-financial income

The household income components are provided on a pre-tax basis in the SOEP data while income and social security taxes are estimated using a procedure outlined in Schwarze (1995) and then provided at the household level for all forms of financial and non-financial income combined. As a consequence, the disposable after-tax household income measure in the dataset combines financial and non-financial income. We subtract pre-tax financial income items from this disposable income measure and add back the estimated tax on financial income to obtain an estimate of household disposable non-financial income.

Detailed information on the applied tax estimation procedure, i.e. information that goes beyond that given in Schwarze (1995), is not available to us. Therefore, we use the relevant features of the German income tax code for the calendar years 1983 to 2008 in order to reverse-estimate taxes pertaining to financial income. Specifically, we estimate taxable income for each household and find the applicable income tax rate using the SOEP-estimated income taxes. We then apply the estimated tax rate to the taxable financial income components. In Appendix 3, we give more detail on this procedure.

The approach to cleaning the SOEP-provided income taxes adopted in this paper potentially introduces a bias. However, our calibration results are strongly robust to crude approaches to cleaning income taxes and to artificially induced noise in estimated tax rates.

The simplifying assumptions applied by the SOEP for the tax estimation possibly result in biased tax estimates, too. The main concern is an overestimation of taxes for households that itemize expenses, particularly with respect to labor income.

4.3. Zero non-financial income observations

The non-financial income process used in this paper assumes a lognormal distribution for non-financial income. However, income process calibrations based on PSID data sometimes find a small fraction of zero-income observations (e.g. Carroll, 1992, Cocco, Gomes, and Maenhout, 2005). The SOEP sample used in this paper exhibits a similar pattern (Figure 1).

Figure 1 Household non-financial income in 2006 terms

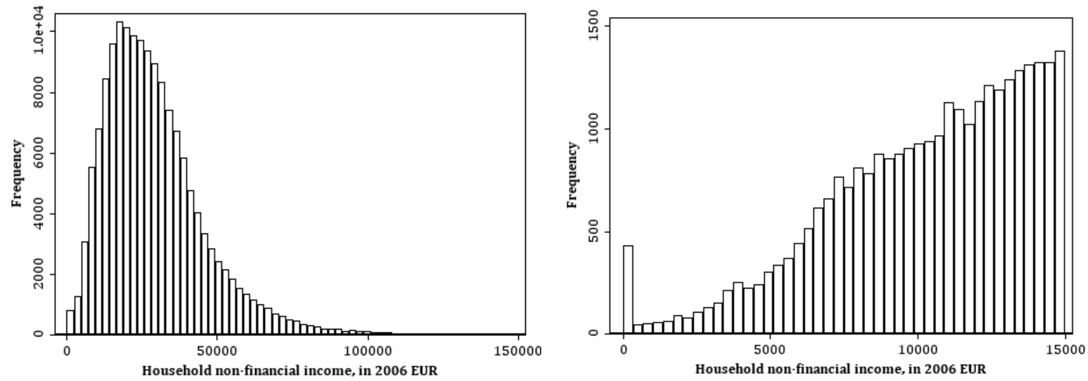

The figure shows household non-financial income in 2006 terms; left: complete distribution, capped at € 150,000 (unweighted); right: lower tail with zero non-financial income observations visible (unweighted).

Given that reported zero-income draws considered implausible by the SOEP would have been treated as item non-response and hence would have been imputed, we assume that the zero-income observations in our sample do not suffer from considerably higher measurement error than any other income observation. Yet, zero non-financial income that is not offset by high financial income is hard to reconcile with the German welfare system. We exclude zero non-financial income observations from the analysis in the base case and introduce them in an extension. There, we treat them as true zero non-financial income draws whenever the respective household's reported after-tax financial income does not exceed € 12,000 per year in 2006 terms. We apply this threshold because for households reporting financial income in excess of this limit, zero non-financial income is increasingly likely to reflect choice rather than non-financial income risk.

Table 4 summarizes the incidence p_z of zero non-financial income draws in our subsamples for different financial-income thresholds. The estimates suggest that females are more likely to experience disastrous income draws than males and that education reduces this risk at least for females. For males, the estimated zero-income probabilities are lower than the corresponding U.S. estimates in Carroll (1992) and Cocco, Gomes, and Maenhout (2005), which is consistent with intuition regarding differences in welfare policies between the two countries.

Table 4 Probability of zero non-financial income draws

	Females			Males		
	Edu: low	Edu: middle	Edu: high	Edu: low	Edu: middle	Edu: high
p_z	0.675%	0.609%	0.448%	0.184%	0.216%	0.334%
p_z if financial income < € 100,000 / year	0.675%	0.609%	0.448%	0.184%	0.216%	0.299%
p_z if financial income < € 50,000 / year	0.675%	0.609%	0.437%	0.184%	0.211%	0.275%
p_z if financial income < € 25,000 / year	0.675%	0.561%	0.430%	0.163%	0.200%	0.180%
p_z if financial income < € 12,000 / year	0.644%	0.413%	0.298%	0.106%	0.164%	0.129%

The table shows zero non-financial income draws as a weighted percentage of the number of observations at household level by subsample for different maximum levels of financial income.

The non-financial income process underlying Section 8.2 treats zero non-financial income draws in a discrete way:

$$\begin{aligned} \text{for } z_t = 0: \quad & y_t = 0, \\ \text{for } z_t = 1: \quad & \ln(y_t) = \begin{cases} f(t) + v_t + \varepsilon_t, & t \leq K \\ \ln(\lambda) + f(K) + v_K, & t > K, \end{cases} \end{aligned} \tag{14}$$

where $P(z_t = 0) = p_z$ and $P(z_t = 1) = 1 - p_z$.

4.4. Deterministic component: estimation procedure and results

The income process in equations (7) to (9) is defined for income y_t of an average household head (in a certain education group), while in our sample we observe income $y_{i,t}^h$ of household h in which a particular household head i lives at age t.

We specify the log non-financial income of a particular household head i at age t in our sample as

$$\ln(y_{i,t}) = \alpha_i + f(t) + v_{i,t} + \varepsilon_{i,t}, \qquad t \leq K,$$

$$\ln(y_{i,t}) = \alpha_i + f(K) + v_{i,K}, \qquad t > K,$$

which differs from (7) and (9) only by the household-head fixed effect α_i. We further assume that the relation between $y_{i,t}^h$ and $y_{i,t}$ is given by

$$\ln(y_{i,t}^h) = \ln(y_{i,t}) + \beta_m m_{i,t} + \beta_s s_{i,t},$$

where $m_{i,t}$ is marital status, which takes value 1 if individual i is married or lives with a partner and zero otherwise, and $s_{i,t}$ is family size, which is the number of people other than head and spouse in the household (Cocco, Gomes, and Maenhout, 2005). This assumption is needed for us to employ the same first-stage regression specification as Cocco, Gomes, and Maenhout (2005), which we seek to do to ensure comparability.

Then, to estimate $f(t)$ for $t \leq K$, i.e. to estimate the deterministic component of non-financial income as a function of age at every age prior to retirement, and to obtain an estimated error structure, we run a fixed-effects OLS regression of log household non-financial income $y_{i,t}^h$ on a set of age dummies and on the household characteristics $m_{i,t}$ and $s_{i,t}$ defined above:

$$\ln\left(y_{i,t}^h\right) = a + \varphi_t + \alpha_i + \beta_m m_{i,t} + \beta_s s_{i,t} + \epsilon_{i,t}. \tag{15}$$

In (15), the φ_t are the true age dummy loadings. The cross sectional mean of α_i is restricted to 0 so that $f(t) = a + \varphi_t = E\left(\ln(y_{i,t}^h) \mid m_{i,t} = 0, s_{i,t} = 0\right)$. For the error $\epsilon_{i,t}$, we have $\epsilon_{i,t} = v_{i,t} + \varepsilon_{i,t}$. Table 5 presents the results of the first-stage regressions.

Table 5 First-stage fixed-effects regression for household non-financial income

A: Female	Education: low		Education: middle		Education: high	
	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic
Constant	8.75	59.89	8.69	85.40	10.05	76.91
$s_{i,t}$ (family size)	0.27	11.31	0.11	6.37	0.07	2.99
$m_{i,t}$ (marital status)	0.49	10.26	0.39	13.61	0.42	7.22
R^2	0.769		0.785		0.818	
F-statistic for joint test of age dummy coefficients	1.60		3.56		4.99	
p-value	0.001		0.000		0.000	
Number of observations	13,223		29,897		7,676	
Number of persons in subsample	2,293		5,276		1,416	
Average number of obs. per person	5.77		5.67		5.42	

B: Male	Education: low		Education: middle		Education: high	
	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic
Constant	8.91	48.76	9.07	77.57	10.29	84.06
$s_{i,t}$ (family size)	0.17	14.18	0.13	19.72	0.05	4.50
$m_{i,t}$ (marital status)	0.15	3.42	0.20	13.80	0.23	8.66
R^2	0.799		0.812		0.844	
F-statistic for joint test of age dummy coefficients	2.92		127.30		6.72	
p-value	0.000		0.000		0.000	
Number of observations	16,183		64,399		19,854	
Number of persons in subsample	2,385		9,447		3,231	
Average number of obs. per person	5.77		5.67		5.42	

The table shows the results of the regression of log non-financial household income $\ln\left(y_{i,t}^h\right)$ on age dummies, controlling for marital status and family size and individual fixed effects; age dummy loadings are not shown; standard errors underlying the test statistics are heteroskedasticity-robust and cluster-robust, where clusters are household heads; F-statistics and p-values pertain to a test of joint significance of the age dummy coefficients; Panel A: female household heads; Panel B: male household heads.

With heteroskedasticity-robust and cluster-robust standard errors, where data is taken to be clustered by household head, the coefficient estimates in the first-stage regressions are significant both for the controls and for the age dummies (not shown) except for the high-education subsamples, where a large fraction of the age dummy coefficients is not significantly different from zero at the 95% level. However, for all subsamples, the age dummy coefficients are jointly significant (Table 5). Compared to Cocco, Gomes, and Maenhout (2005), we estimate considerably higher coefficients for the family size variable, which we attribute mainly to relatively more generous child benefit payments in Germany.

The age-earnings profile resulting from the estimates of $\varphi_t + a$ is hump-shaped for all subsamples (Figure 2). As expected, the profiles differ in their overall level both across education groups and across gender. Achieving a higher level of education and and/or being a male coincides with higher expected income. The shape of the profiles is similar for the low-education and middle-education subsamples. Persons with higher education degrees show considerably stronger income growth during the early years of their working lives; the hump in their earnings profile is more pronounced than for the other education groups. The profiles suggest that, for females and males in the two lower education groups, there is a period of low or zero income growth between age 25 and 35 when household characteristics are controlled for, leading to a saddle in the hump shape.

We fit 3rd- to 7th-order polynomials to the estimates of $\varphi_t + a$ for $t \leq K$ in order to obtain a smooth age-earnings profile $f(t)$ for the time before retirement. The width of the 95% confidence interval around the estimated age dummy coefficients suggests that the 3rd-order polynomial be used for the life cycle model, but the saddle-like feature described above is captured only by the 5th-order polynomial. Given that four of the six earnings profiles exhibit this feature, we take to the life cycle model the fitted 5th-order polynomial. This polynomial, in turn, shows an undesired feature for females in the middle education group right before retirement. We use a 6th-order polynomial for this subsample. Figure 2 shows the polynomials used in the life cycle model. The results of the polynomial regression are presented in Table 6.

Figure 2 Estimates of loadings on age dummies in the first-stage regression and fitted polynomials

A: Females

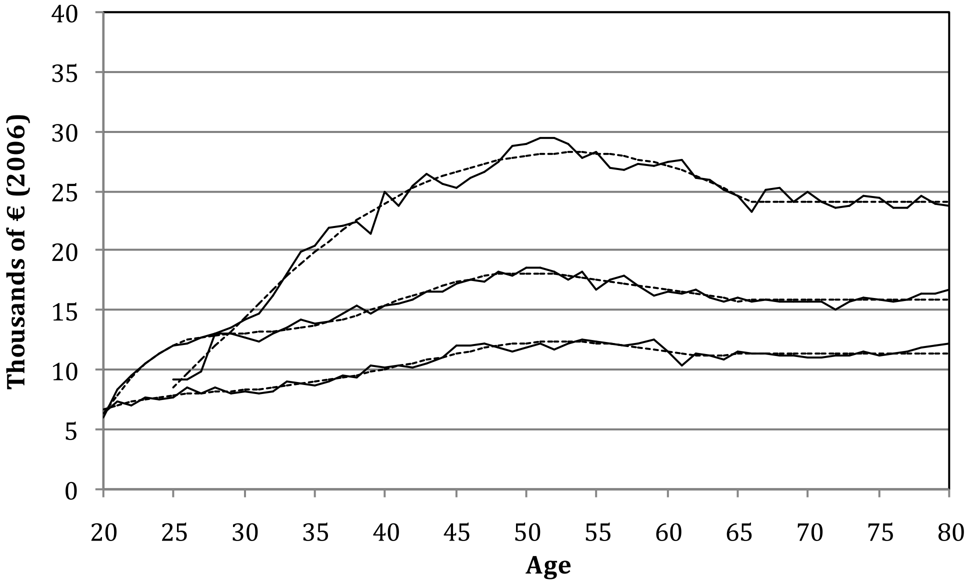

B: Males

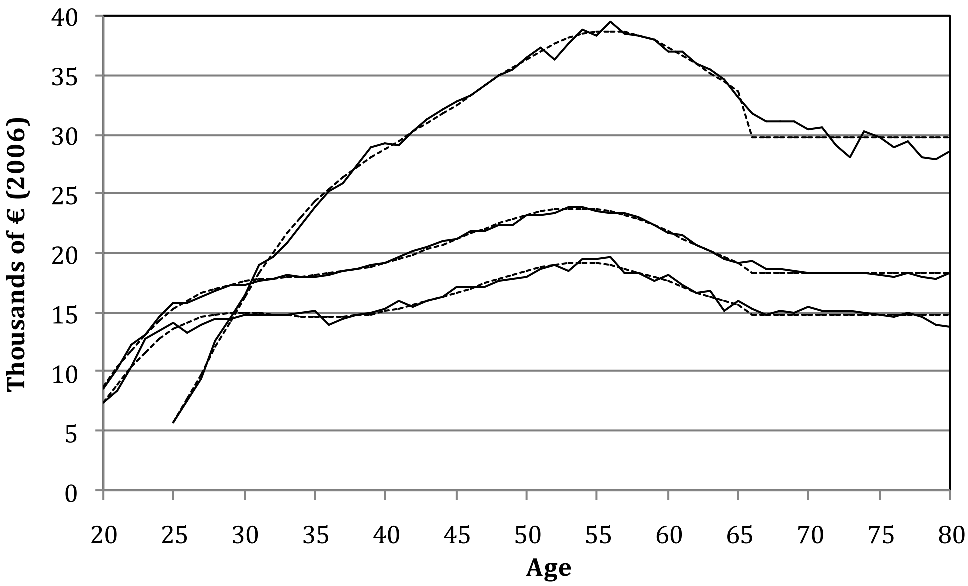

The graphs show the estimates of loadings on age dummies (in levels) in the first-stage regression and fitted deterministic components (5th-/6th-order polynomial) of income trajectories by education and gender; the overall level of the dummies for each education group corresponds to the education level.

In the polynomial regression, we include a constant for the retirement period until age 80 in order to obtain the mean of $a + \varphi_t$ for $t > K$, i.e. the expected log retirement income. We estimate the log of the replacement ratio λ as the difference between the mean of $a + \varphi_t$ for $K < t \leq 80$ and the fitted value for $t = K$ from the polynomial regression. For females, replacement ratios are remarkably high. For the lower two male education groups, the replacement ratios are higher than the corresponding estimates for the U.S. in Cocco, Gomes, and Maenhout (2005). This is mainly due to relatively high replacement ratios in statutory pension insurance.

Table 6 Coefficient estimates for the fitted age polynomials

A: Females	Education: low coefficient	Education: middle coefficient	Education: high coefficient
Constant	-1.168	-46.419	-0.937
Age	1.354	8.160	0.825
Age^2 / 10	-0.709	-4.839	-0.248
Age^3 / 100	0.180	1.488	0.038
Age^4 / 1000	-0.022	-0.250	-0.003
Age^5 / 10000	0.001	0.022	0.000
Age^6 / 100000	not included	-0.001	not included
R^2	0.961	0.983	0.984
λ (replacement ratio)	1.009	1.004	0.983

B: Males	Education: low coefficient	Education: middle coefficient	Education: high coefficient
Constant	-17.182	-11.929	-44.790
Age	3.287	2.612	5.680
Age^2 / 10	-1.566	-1.231	-2.345
Age^3 / 100	0.361	0.282	0.481
Age^4 / 1000	-0.040	-0.031	-0.049
Age^5 / 10000	0.002	0.001	0.002
R^2	0.968	0.995	0.996
λ (replacement ratio)	0.949	0.954	0.875

The table shows the results of the regression of age dummy coefficients (including the constant term from the first-stage regression) estimated in the first-stage regression on a 5th-/6th-order polynomial in age and a constant for the retirement period. The coefficient estimate for the retirement constant is reported after conversion into the replacement ratio λ.

To facilitate comparison, we plot in Figure 3 the fitted polynomial for the middle education group against the fitted polynomial for the corresponding education group in the U.S. reported by Cocco, Gomes, and Maenhout (2005). The absolute levels of the profiles are not comparable as the units of measurement differ. The series are scaled such that their peaks appear at the same level in the graph, allowing for a comparison of the respective patterns.

The German profile starts at a lower relative level and shows stronger growth in the first five years, possibly due to low relative earnings levels for individuals still completing apprenticeships. The U.S. profile keeps rising steadily until it reaches its peak between age 40 and 45, while the German profile first exhibits the saddle-like feature described above and then peaks at age 55. Generally, the location of the peaks with respect to age does not seem to vary considerably among education groups in the German earnings profiles (Figure 2), while U.S. profiles suggest that lower education coincides with earlier earnings peaks (not shown). In Figure 3, the transition into retirement in the U.S. profile is characterized by a small decline in income before retirement and a large drop upon retirement, while the German profile exhibits the opposite. Comparability is limited with respect to this particular pattern, though. While the difference is driven partly by higher replacement ratios in German statutory pension insurance, the treatment of households with a head retiring before age 65 differs in the two analyses.

Figure 3 Fitted polynomials for the middle education group in Germany and in the U.S.

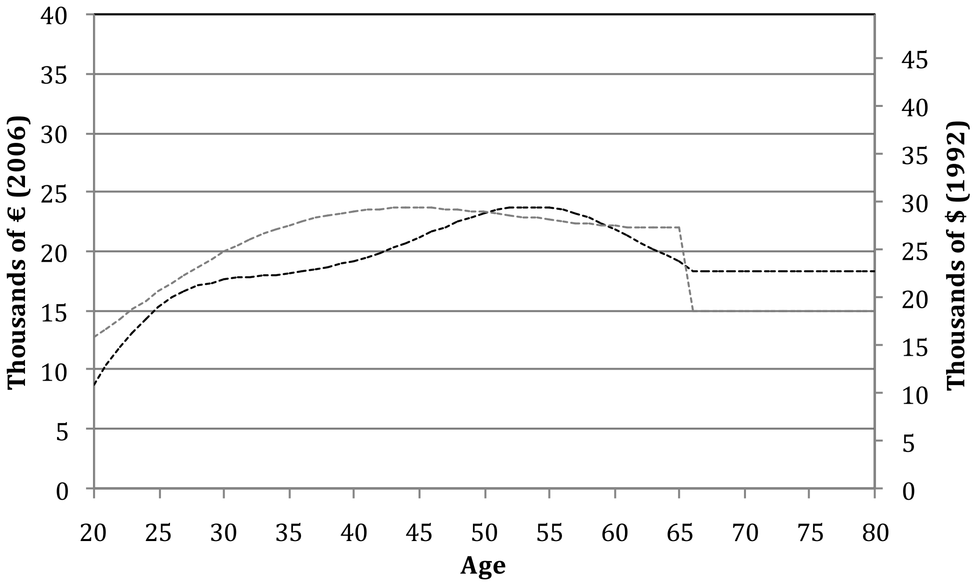

The graph shows polynomials fitted to estimated age dummy loadings for the middle education group in our analysis (black dotted line, left scale, in thousands of € and in 2006 prices) and in Cocco, Gomes, and Maenhout (2005) (gray dotted line, right scale, in thousands of $ and in 1992 prices). The units of measurement differ and hence the absolute levels of the profiles are not comparable. The axes are scaled such that the peaks of the profiles appear at the same level in the graph.

4.5. Error structure: estimation procedure and results

As is common in the literature, we closely follow the variance decomposition method proposed by Carroll and Samwick (1997) to estimate the variances of the transitory shocks and the innovations to the stochastic persistent component of the income process. Let $d_{i,t,l}$ denote the l-period differenced residual in (15), i.e. $d_{i,t,l} = \epsilon_{i,t} - \epsilon_{i,t-l}$. It is straightforward to see that

$$d_{i,t,l} = \sum\nolimits_{j=t-(l-1)}^{t} \zeta_{i,j} + \varepsilon_{i,t} - \varepsilon_{i,t-l} \,,$$

so that

$$Var\left(d_{i,t,l}\right) = 2\sigma_{\varepsilon}^{2} + l\sigma_{\zeta}^{2}. \tag{16}$$

When estimates of $Var\left(d_{i,t,l}\right)$ for more than two distinct values of l are available, (16) is over-specified. We use all available $d_{i,t,l}$ to obtain the sample variance estimates $\widehat{Var}\left(d_{i,t,l}\right)$ for $l = 1, 2, \dots, L$ and regress these on the vector $\{1, 2, \dots, L\}$ and a vector of constants $\{2, 2, \dots, 2\}$:

$$\widehat{Var}\left(d_{i,t,l}\right) = \sigma_{\varepsilon}^{2} 2 + \sigma_{\zeta}^{2} l + \tau_l, \tag{17}$$

where τ_l is the residual in the regression. Assuming that errors in the estimates of $Var\left(d_{i,t,l}\right)$ are i.i.d., the OLS regression coefficients are efficient estimates of σ_{ε}^{2} and σ_{ζ}^{2}.

We use $L = 10$, that is, series of differenced first-stage regression residuals with up to a ten-period lag. If the persistent stochastic component follows a random walk and if σ_{ε}^{2} and σ_{ζ}^{2} are the same both across individuals and across life cycle stages in each subsample, the choice of L only affects the efficiency of the estimates of the shock variances. However, we caution that one or several of these assumptions are not fulfilled in our sample. Increasing the maximum lag length L tends to shift variance from persistent to transitory shocks. Multiple reasons for this effect, including age-variation in shock variances, are conceivable. To allow for comparison with previous research, however, we keep the standard specification of the error structure (see 3.3) in the base case and, in 8.3, show the effect of a change in L on optimal behavior. Here we choose $L = 10$ so as not to use variance estimates from either end of the continuum of estimates that can be obtained from our sample.

We allow for correlation $\rho_{\zeta,\eta}$ between labor income shocks and excess stock returns through the innovation $\zeta_{i,t}$ to the stochastic persistent component. As we do not estimate $\zeta_{i,t}$ but rather $d_{i,t,l}$, we first estimate the correlation $\rho_{d,\eta}$ between $d_{i,t,l}$ and η_t for $l = 1$ using all available i and t. Suppressing the subscripts i and t,

$$\rho_{d,\eta} = COV(d_1, \eta)/(\sigma_{d1} \times \sigma_\eta)$$

and

$$\rho_{\zeta,\eta} = COV(\zeta, \eta)/(\sigma_\zeta \times \sigma_\eta).$$

With the idiosyncratic shock ε assumed to be uncorrelated with ζ and η,

$$COV(\zeta, \eta) = COV(d_1, \eta)$$

so that

$$\rho_{\zeta,\eta} = (\rho_{d,\eta} \times \sigma_{d1})/\sigma_\zeta, \tag{18}$$

where σ_{d1} and σ_ζ are estimated in the variance decomposition above. The stock returns are for the CDAX, a broad, value-weighted index for the German equity market (Thomson Reuters, 2011). Table 7 presents the results of the variance decomposition regression and estimates of $\rho_{\zeta,\eta}$ by education and gender.

Table 7 Variance decomposition

A: Females	Education: low		Education: middle		Education: high	
	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic
σ_ζ^2 (variance of innovation to persistent component)	0.021	11.46	0.015	7.36	0.022	9.93
σ_ε^2 (variance of transitory shocks)	0.076	13.30	0.076	11.77	0.057	8.42
R^2 of variance decomposition regression	0.997		0.995		0.995	
$\rho_{\zeta,\eta}$ (contemporaneous corr. with real excess stock returns)	-0.029	-2.17	0.037	4.98	-0.115	-8.17
$\rho_{\zeta,\eta}$ (1-period lagged corr. with real excess stock returns)	0.061	4.53	0.053	7.07	0.117	8.34

B: Males	Education: low		Education: middle		Education: high	
	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic
σ_ζ^2 (variance of innovation to persistent component)	0.011	5.90	0.011	12.82	0.009	10.02
σ_ε^2 (variance of transitory shocks)	0.056	9.64	0.037	13.86	0.038	13.50
R^2 of variance decomposition regression	0.993		0.997		0.997	
$\rho_{\zeta,\eta}$ (contemporaneous corr. with real excess stock returns)	-0.017	-1.84	-0.050	-10.21	-0.115	-13.19
$\rho_{\zeta,\eta}$ (1-period lagged corr. with real excess stock returns)	0.044	4.65	0.036	7.32	0.045	5.10

The table reports (1) the results of the variance decomposition regression of sample variance estimates of differenced residuals for different lags on the vectors $\{1, 2, \ldots, L\}$ and $\{2, 2, \ldots, 2\}$ for $L = 10$ and (2) both the contemporaneous and one-period lagged correlation between the innovation $\zeta_{i,t}$ to the persistent stochastic component and log real excess returns on a broad German equity index.

For females, the variance estimates in Table 7 correspond to standard deviations of 12 - 15% for the persistent shocks and around 24 - 28% for transitory shocks. For males, the ranges are 10 - 11% and 19 - 24%, respectively.

Consistent with the evidence referred to in the literature review, the variability of transitory shocks is larger than that of the innovations to the persistent component. The relative magnitude of transitory shocks seems to decrease in education. Persistent shocks do not seem to exhibit such a pattern. For all education groups, females seem to face both larger transitory shocks and larger persistent shocks than males. For both genders, the levels of the estimates are generally in line with the estimates for the U.S. discussed earlier.

The estimates of contemporaneous correlation between the innovation to the persistent component and log real excess stock returns range from negative 0.12 to positive 0.04 and are significant at the 95% level with the exception of the low education subsample for males. As done in Campbell et al. (2001), we also estimate the correlation with one-period lagged stock returns to allow for delayed reactions in non-financial income. Estimates are positive and significant for all education groups and range from 0.05 to 0.12 for females and from 0.04 to 0.05 for males. The fact that these estimates are positive is consistent with the notion that capital and labor income are associated at the aggregate level over longer periods of time. Meanwhile, the small magnitude of the coefficient estimates suggests that labor income risk is largely idiosyncratic in our sample.

Like Campbell et al. (2001), we use the estimates of one-period lagged correlation as contemporaneous correlation when solving the life cycle model.

4.6. Other parameters and benchmark parameter set

Agents enter the optimization at $t_1 = 20$ for low- and middle-level education ($t_1 = 25$ for high education) and have a probability p_k of being alive at time period k conditional on being alive at time $k - 1$. They retire at age $K = 65$ (the first period in retirement in the optimization is hence $K + 1$), which is the mean and median retirement age in our sample, and die at age 100 with probability one. The p_k are taken from the mortality tables published by the German Statistics Office (Statistisches Bundesamt, 2010). Following Cocco, Gomes, and Maenhout (2005), we set the base case risk aversion to $\gamma = 10$ and the intertemporal discount factor to $\delta = 0.96$.

In the benchmark case, we solve the model for male and female subsamples with middle-level education. These groups are by far the largest education groups in our sample and their relative education level is comparable to the relative education level of the group used in Cocco, Gomes, and Maenhout's (2005) benchmark optimization.

As mentioned before, all financial asset returns are continuously compounded and understood on an after-tax basis. Taxes on capital income are set to 25%, in line with current interest, dividend and capital gains withholding taxes in Germany (ignoring additional surcharges and exemptions for simplicity). Given these assumptions, we use a real after-tax log risk-free rate r_f of 1%. Using Dimson et al.'s (2003) historical estimate adjusted for taxes, we set μ to 3%. We follow their argument that the historical standard deviation of excess returns is unlikely to be an indicator of future variability given the nature of the historical events included in their sample period and we adopt their approach of limiting σ_η to 15% (20% pre-tax).

Table 8 summarizes the parameters used as a benchmark case parameter set.

Table 8 Benchmark case parameter set

Parameter	Parameter value	
	Females	Males
r^f (real risk-free rate)	1.0%	same as for females
μ (equity risk premium)	3.0%	same as for females
σ_η (STDEV of the innovation to excess returns)	0.15	same as for females
σ_ε^2 (variance of transitory shocks)	0.076	0.037
σ_ζ^2 (variance of persistent shocks)	0.015	0.011
$\rho_{\zeta,\eta}$ (correlation with excess stock returns)	0.053	0.036
t_1 (starting age)	20	same as for females
K (retirement age)	65	same as for females
T (terminal year)	100	same as for females
δ (intertemporal discount factor)	0.96	same as for females
γ (risk aversion)	10	same as for females

5. Policy Functions

The results presented throughout Sections 5 to 7 pertain to the middle education group and are based on the benchmark parameter set presented in 4.6. We refer to these results as the benchmark case.

5.1. Consumption policy

Figure 4 plots optimal consumption policy functions $\dot{c}_t^*(\dot{x}_t)$ for selected ages. The function is concave and has a general shape similar to that in Deaton's (1991) precautionary savings model. For low values of cash on hand, the liquidity constraint $x_t \geq 0$ is binding. In such cases, the impatience motive dominates the behavior, and it is optimal to spend all of the cash on hand on consumption. When cash on hand is relatively high, the precautionary savings motive is dominant, and it is optimal to save some wealth with an intention to smooth consumption in later periods.

As agents age and their life horizon shortens, saving becomes less attractive. This is why the optimal consumption function becomes less concave and approaches a straight line with a slope of 1 as t approaches T: consider, for example, the dotted ($t = 60$) and gray ($t = 90$) lines in Figure 4.

Figure 4 Optimal consumption policy

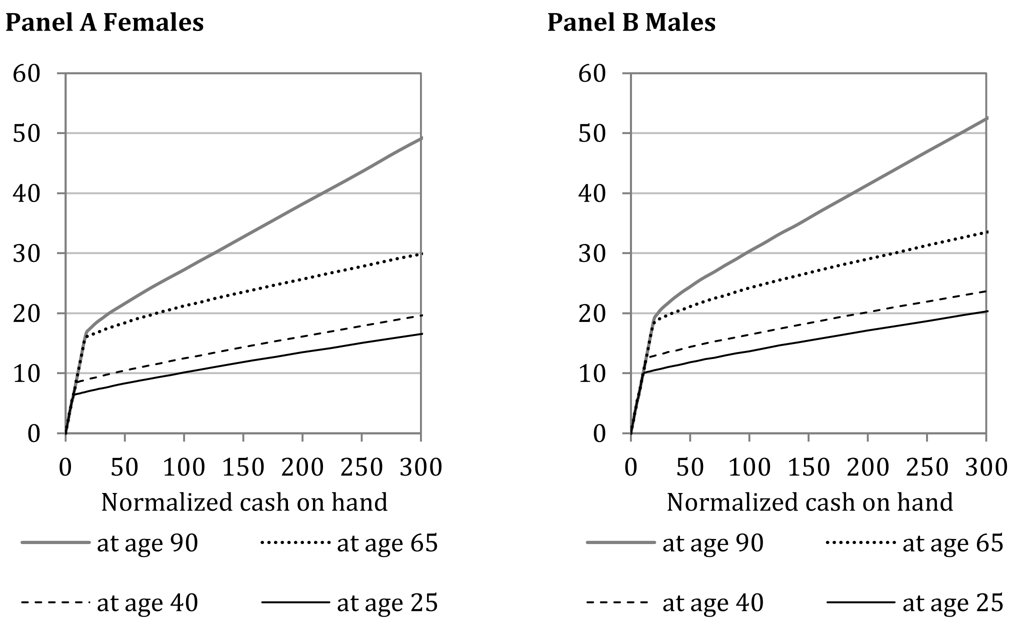

The figure plots $\dot{c}_t^*(\dot{x}_t)$, the optimal normalized consumption, as a function of normalized cash on hand (in thousands of € and 2006 prices) for females and males in the middle education group.

Before T, at any given level of cash on hand for which the liquidity constraint is not binding, optimal consumption for females is lower than for males because of lower expected non-financial income (lower human capital). This effect decreases as t approaches T because as the ratio of human capital to total wealth (cash on hand plus human capital) decreases, human capital becomes less important for consumption choices.

5.2. Allocation policy

Figure 5 plots the optimal allocation policy functions $w_t^*(\dot{x}_t)$ for selected ages. The optimal allocation to the risky asset is decreasing in cash on hand at every age. When the level of cash on hand is relatively low, cet. par., human capital accounts for a large fraction of total wealth. As we establish empirically, persistent innovations to non-financial income are only weakly correlated with risky asset returns - non-financial income risk is largely idiosyncratic – and hence human capital covaries with the risky asset only to a very limited extent. Therefore, the diversification motive induces agents to allocate a larger fraction of their financial wealth to the risky asset when the ratio of human capital to total wealth is large. Human capital 'crowds out' the risk-free asset from the financial portfolio. These findings are supported by the popular opinion in the literature (for example, Jagannathan and Kocherlakota 1996, Cocco, Gomes, and Maenhout, 2005, Calvet and Sodini, 2010) that human capital affects allocation as if it was comprised of the risk-free rather than the risky asset.

In the absence of borrowing constraints, agents with sufficiently low cash on hand and sufficiently large human capital would find it optimal to borrow and hold leveraged positions in the risky asset to diversify away exposure to human capital. In our model, we find that the borrowing constraint, $w_t \leq 1$, is binding for economically non-negligible levels of cash on hand at all ages. In other words, agents with low cash on hand – young and/or poor agents - invest 100% of their financial wealth in stocks.

As agents age, the L-shaped allocation policy graphs 'expand' to the right approximately until the age of retirement as it is illustrated by the solid ($t = 25$), dashed ($t = 40$), and dotted ($t = 65$) lines in Figure 5. After retirement, the optimal allocation function 'retreats' back to the left. Thus, the willingness to take financial risks increases approximately until the retirement age and decreases afterwards.

Figure 5 Optimal allocation policy

Panel A Females **Panel B Males**

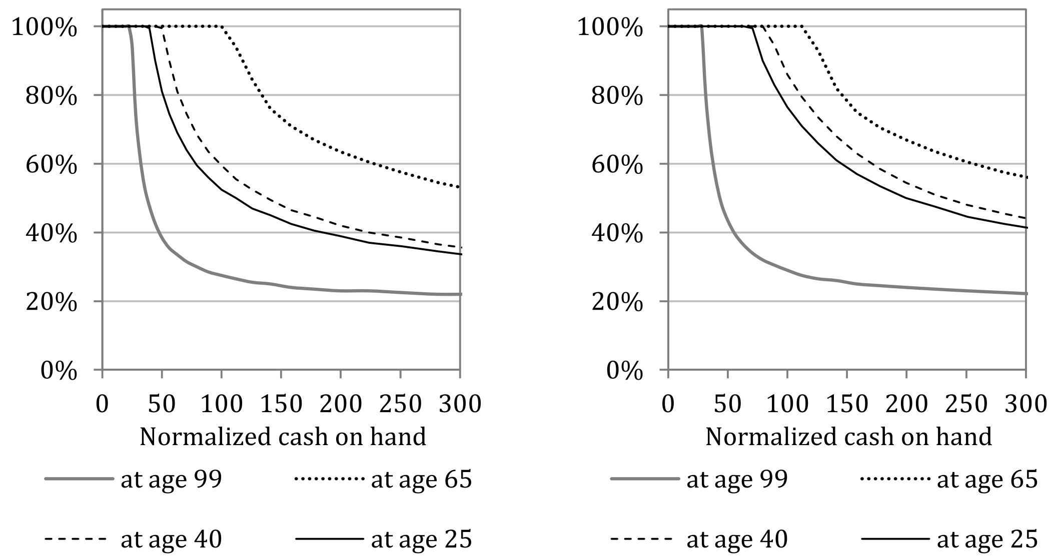

The figure plots $w_t^*(\dot{x}_t)$, the optimal allocation to the risky asset, as a function of normalized cash on hand (in thousands of € and 2006 prices) for females and males in the middle education group.

To understand the intuition behind this result, consider first an agent aged 99 (gray line in Figure 5). The last remaining retirement income stream at $t = 100$ is deterministic; exposure to it is equivalent to holding the risk-free asset. Hence, the optimal allocation function equals Merton's (1971) result as in (2) where $W^F = x_{T-1} - c_{T-1}$ and $W^{HC} = y_T p_T e^{-rf}$. Since non-financial income is risk-free throughout the whole retirement period, Merton's result holds for every age from K, the last pre-retirement year, to $T - 1$. The optimal allocation functions for these ages are therefore hyperbolas with a horizontal asymptote of $\mu/\sigma^2\gamma$: as cash on hand increases, human capital plays a less important role in the allocation of the financial portfolio. Younger retirees have larger human capital (i.e. more future expected pension income), causing the allocation functions in Figure 5 to compress as retirees age.

Prior to retirement, non-financial income is stochastic and has two sources of randomness: transitory and persistent components. The conditional variance of the persistent component v_t is larger for more distant income flows. The conditional variance of the transitory component, in turn, is constant. A significant part of human capital consists of the present value of retirement income, uncertainty in which is solely dependent on v_K, the conditional variance of which decreases as t approaches K. This is why the riskiness of human capital reduces as agents approach retirement. As Gollier and Pratt (1996) show, this reduction in 'background risk' reduces the risk-aversion towards the risky asset. Consequently, agents closer to retirement find it optimal to invest more heavily in the risky asset for each level of cash on hand.

For any given level of cash of hand before $T-1$, human capital is smaller for females than it is for males, and, before K, it is also riskier. Following the intuition laid out above, the optimal risky share is then lower for females than it is for males. The difference gradually reduces as human capital diminishes with t approaching T.

6. Simulation Results

We run Monte-Carlo simulations of our model for $N = 100{,}000$ agents that live from period t_1 to T. For each agent i and at each age t we draw a vector of 'core' random variables $\{\eta_t, u_{i,t}, \varepsilon_{i,t}\}$ from a multivariate normal distribution using Cholesky decomposition of the correlation matrix. This allows us to calculate the simulated realizations of non-financial income, $\left\{\{y_{i,t}\}_{t=t_1}^{T}\right\}_{i=1}^{N}$, and the returns on the risky asset $\{r_t\}_{t=t_1}^{T}$. The agents start independent lives with no financial wealth and in the first period their cash on hand equals the first realization of non-financial income. Agents then make saving and allocation decisions according to the estimated optimal policy functions. Given the allocation choice, we find the rate of return on the financial portfolio (from (5)) and the corresponding realizations of cash on hand (from (3)) in the subsequent period. We iterate the calculation over age and in the end obtain a set of realizations of cash on hand, consumption, and optimal allocation: $\left\{\{x_{i,t}, c_{i,t}, w_{i,t}\}_{t=t_1}^{T}\right\}_{i=1}^{N}$. The average life cycle trajectories of non-financial income and consumption for agents in the middle education group are plotted in Figure 6.

Figure 6 Expected trajectories of non-financial income and consumption

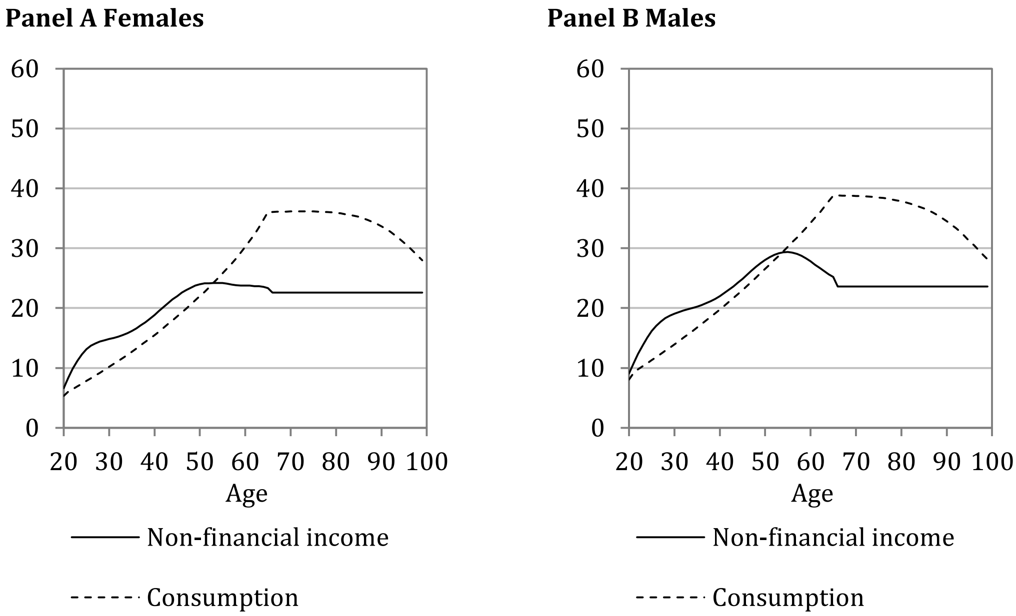

The graphs show average simulated realizations of non-financial income and consumption (both in thousands of € and 2006 prices) over the life cycle.

We find the pattern of average consumption to be generally similar to the one reported by Cocco, Gomes, and Maenhout (2005). Even though expected non-financial income in the second half of the life cycle is high relative to its level early in life, it is on average optimal to sustain a moderate level of consumption in the early years. Average consumption follows the growth in income and, as uncertainty about future income shrinks, gradually increases towards the retirement age. Consistent with the classic microeconomic life cycle consumption model, agents in our model accumulate financial wealth in the first half of their lives and use it to finance consumption in the second half.

Our findings regarding the savings behavior over the life cycle differ from those of Cocco, Gomes, and Maenhout (2005) on one account. Cocco, Gomes, and Maenhout (2005) find that early in the life cycle, when income is low and agents have not yet accumulated substantial savings, consumption closely mirrors income. On the contrary, we find that it is optimal to make sizeable savings during the early years. The average ratio of savings to non-financial income is about 20% for males and 30% for females from age 20 to 40.

6.1. Pattern of human capital and financial wealth

We find the approximate expected value of human capital by discounting the expected future non-financial income streams at the risk-free rate:

$$E_{t_1}(W_t^{HC}) \approx \sum_{j=t+1}^{T} E_{t_1}(y_j) \left(\prod_{k=t+1}^{j} p_k \right) e^{-r^f(j-t)}, \tag{19}$$

where $E_t(y_j)$ is estimated as a sample mean of simulated income realizations for age j across agents. Since retirement income is risk-free, (19) gives the precise estimate for $K \leq t \leq T-1$ and overestimates it for $t_1 \leq t < K$. Figure 7 plots the expected trajectories of human capital and financial wealth over the life cycle. We find that human capital follows a downward-sloping trajectory as opposed to a hump-shaped trajectory as in Cocco, Gomes, and Maenhout (2005). While the present value of an agent's expected high-income realizations occur far from the beginning of the life cycle in our sample, too, the low level of our discount rate prevents the hump shape from emerging. We also find that the ratio of human capital to total wealth, $W^{HC}/(W^{HC}+W^F)$, is decreasing all the way from its maximum early in life until approximately the age of 85 when increasing force of mortality encourages rapid consumption of financial wealth.

Figure 7 Expected trajectories of human capital and financial wealth

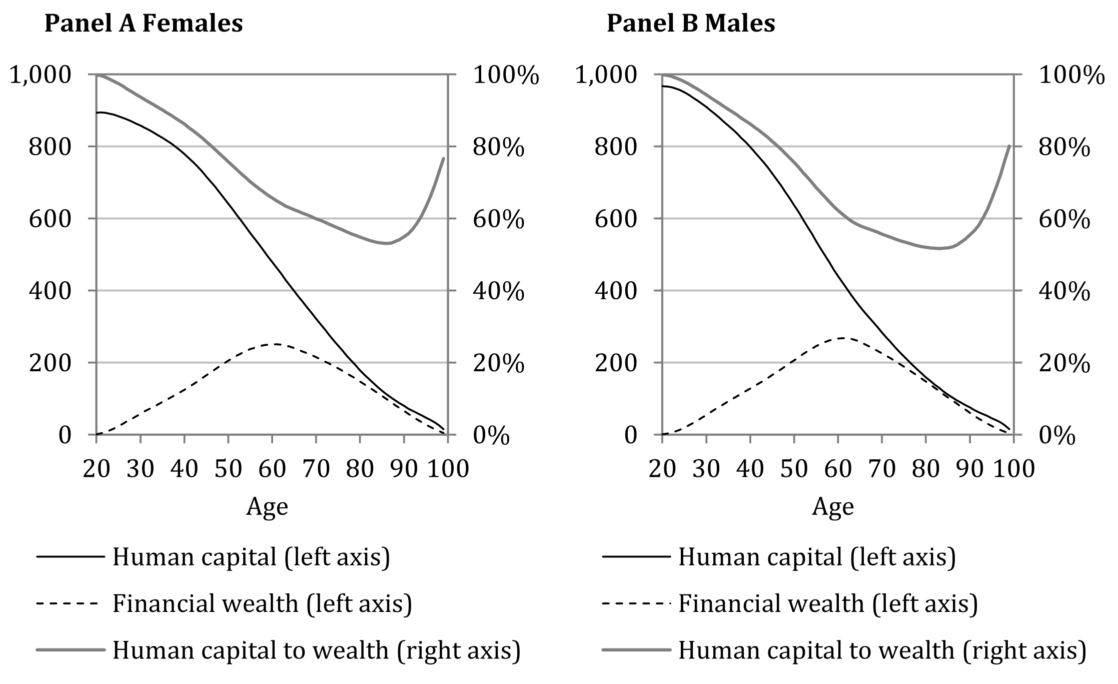

The graph shows average simulated realizations of financial wealth, human capital (in thousands of € and 2006 prices), and the ratio of average human capital to total wealth over the life cycle.

6.2. Allocation pattern

Figure 8 plots the mean and 90% confidence bounds of the simulated allocation to the risky asset over the life cycle. Young agents with relatively large values of human capital and low financial wealth aggressively invest in the risky asset, which follows directly from the optimal allocation policy in Figure 5. As agents age and accumulate financial assets, the allocation to the risky asset reduces and reaches its lowest point of approximately 45% around the age of 55 for females and of approximately 55% around the age of 60 for males. For both genders, the average allocation follows a U-pattern, but females are predicted to start reducing the risky share slightly earlier in life and to do so more quickly. During retirement, the predicted risky share increases gradually again.

Figure 8 Expected trajectory of the allocation to the risky asset

Panel A Females **Panel B Males**

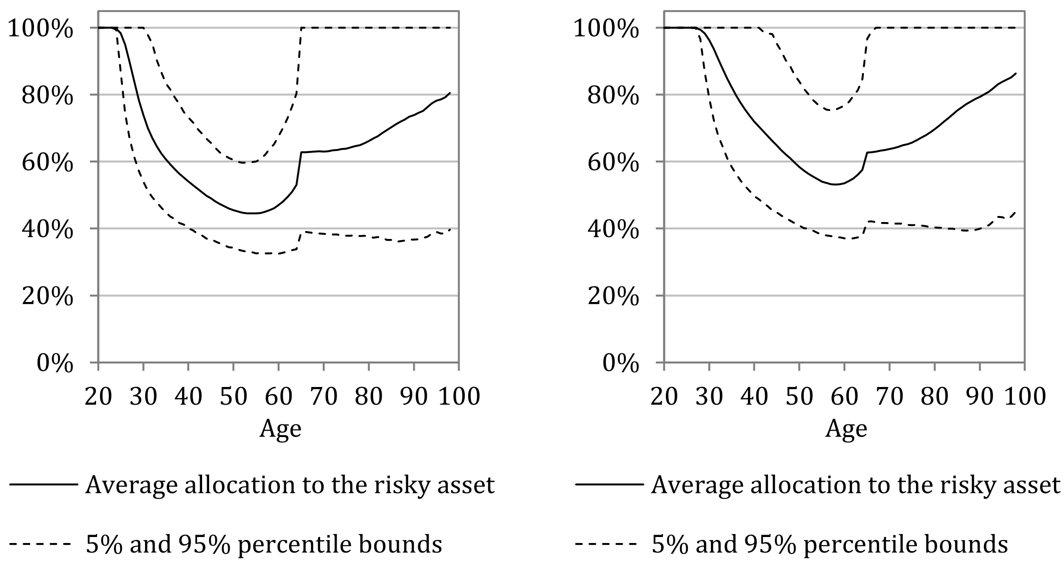

The graph shows the mean and the 90% confidence bounds of the simulated allocation to the risky asset over the life cycle.

Three forces affect the optimal allocation to the risky asset prior to retirement. First, as the level of human capital decreases, the diversification demand for the risky asset decreases. Second, as agents accumulate a stock of financial wealth towards the middle of the life cycle, the impact of human capital on the allocation of the financial portfolio becomes weaker, and allocation tilts towards Merton's (1969) policy in (1), that is towards a lower allocation to the risky asset. Third, as the riskiness of human capital decreases, agents become less risk averse towards the risky asset and allocate more financial wealth to it. The last effect is more pronounced for agents close to the retirement age, when each passing year reduces the conditional variance of v_K (the persistent stochastic component at age K) by an increasingly large fraction. Thus, before retirement, when the third effect dominates the two first ones, average optimal allocation bounces back.

During retirement only the first two forces are active: retirement income is risk-free and the 'riskiness' of human capital does not change any more. As both human capital and financial wealth decrease during retirement, the two forces affect allocation in opposite directions, the net effect being an increase in average allocation to the risky asset during retirement.

With the ratio of human capital to financial wealth being very similar for males and females over the entire life cycle, the differences in the predicted allocation to the risky asset between genders before retirement can be explained by the difference in the riskiness of human capital. The variance of persistent shocks is higher for females, which increases the riskiness of their human capital and hence leads to a lower optimal risky share.

Compared to the average optimal risky asset share in Cocco, Gomes, and Maenhout (2005), first, the predicted risky asset share in our model (for males) starts deviating from 100% about a decade earlier. With the higher savings rates early in life predicted in our model, agents' build up financial wealth more quickly. Then, with a higher level of cash on hand, the optimal allocation policy discussed in 5.2 leads to a lower risky asset share. Second, the rebound in the optimal risky asset share before retirement discussed above does not occur in Cocco, Gomes, and Maenhout (2005).

7. Welfare Metrics

We are interested in measuring the effect of divergence from the optimal allocation policy on welfare. The methodology we employ for this purpose is taken from Cocco, Gomes, and Maenhout (2005). We examine and compare three suboptimal policy rules apart from the estimated optimal consumption and allocation policies $\dot{c}_t^*(\dot{x}_t)$ and $w_t^*(\dot{x}_t)$. First, we consider a nonparticipation policy, i.e. $w_t = 0$ for every t. Second, we examine Merton's (1969) policy as in (1). This policy ignores any effect of human capital on portfolio choice. Finally we consider a rule of thumb, proposed in Malkiel (1996). The policy is defined as $w_t = (100 - t)/100$, so that the allocation to equities is 80% at age 20 and it linearly decreases to 0% by the age of 100.

Each of the suboptimal policies assigns w_t exogenously, which leaves consumption as the only policy variable in the utility maximization problem. For each of the suboptimal allocation rules we solve the dynamic programming equation in (12), treating w_t as given. We find the optimal consumption policies subject to the given allocation policy and estimate $V_{t_1}(x_{t_1})$.

For the optimal and each of the suboptimal policies we evaluate the maximum achievable utility at age t_1 for the expected level of cash on hand:

$$U^{max} = V_{t_1}\left(E(y_{t_1})\right) = V_{t_1}\left(f(t_1) + \frac{1}{2}\left(\sigma_\zeta^2 + \sigma_\varepsilon^2\right)\right).$$

We then find a constant certain level of consumption which results in the same level of utility as the maximum achievable utility for each of the allocation policies considered. From (10), the constant level of consumption that results in the required value of the utility function is

$$\bar{c} = \left(\frac{(1-\gamma)U^{max}}{\sum_{j=t_1}^{T} \delta^{j-t_1}\left(\prod_{k=t_1+1}^{j} p_k\right)}\right)^{\frac{1}{1-\gamma}}. \tag{20}$$

We compare the $\bar{c}$ that result from each of the suboptimal allocation policies to the $\bar{c}$ for the optimal policy in Table 9.

Table 9 Allocation policy comparison by constant-consumption equivalent

	nonparticipation	Merton (1969)	Malkiel (1996)
Females	93.4%	97.5%	99.0%
Males	93.5%	96.9%	99.4%

The table presents the ratios of utility-equivalent certain consumption levels $\bar{c}$ of three suboptimal allocation policies to $\bar{c}$ of the optimal policy.

We also consider uniform $\pm$ 5% and $\pm$10% adjustments to the optimal consumption policy for all but the very last age. Again, we treat the adjusted consumption policies as given and solve (12) by finding optimal conditional allocation policies. We report the corresponding utility-equivalent constant levels of consumption in Table 10.

Table 10 Allocation policy comparison by constant-consumption equivalent

	Adjustment to the optimal consumption function			
	-10%	-5%	+5%	+10%
Females	96.6%	99.0%	98.2%	90.1%
Males	96.4%	98.9%	98.4%	91.8%

The table presents the ratios of utility-equivalent certain consumption levels $\bar{c}$ of the adjusted consumption policies to $\bar{c}$ of the optimal policy.

We find that suboptimal allocation decisions are costly. Ignoring the effect of human capital and following Merton's (1969) portfolio rule results in around 2.5 - 3% loss on a constant-consumption equivalent basis, which we consider economically large. A life-long deviation between -10% and +5% from the optimal savings rule results in forgone welfare of similar magnitude. A deviation of +10% from the optimal consumption policy leads to a considerably higher utility loss.

Malkiel's (1996) portfolio rule reflects the idea that as agents become older, their human capital decreases and their capacity to endure large financial losses diminishes as the financial portfolio plays an increasingly important role in funding consumption. The rule, however does not take into account the change in riskiness of the human capital on portfolio allocation. We find that Malkiel's (1996) rule is around 1% inferior to the optimal policy on a constant-consumption equivalent basis. To put this result into perspective, we estimate that a life-long deviation of $\pm$ 5% from the optimal consumption policy results in a comparable or slightly larger utility loss.

Compared to the results for deviation from the optimal allocation in Cocco, Gomes, and Maenhout (2005), losses are comparable for the rule of thumb, but we obtain larger losses for the Merton (1969) policy and particularly for equity market nonparticipation.

Overall, the magnitude of utility losses from following suboptimal policies is economically large. In the context of both the model employed here and the portfolio allocation rules considered, maintaining the optimal risky asset share in the financial portfolio and following the optimal consumption policy are both important for welfare, and for deviations from the optimal consumption policy in the range of −10% to +5%, neither of the two forms of deviation considered here is more or less important than the other. Meanwhile, large positive deviation from the optimal consumption policy as well as nonparticipation in equity markets are particularly costly.

8. Sensitivities and Extensions

The income process parameters estimated in Sections 4.4 and 4.5 differ by education. Below, we examine the extent to which these differences justify education-specific portfolio advice. Plus, we incorporate extremely low income draws into the model and assess the robustness of the main results to changes in the coefficient of relative risk aversion and to changes in the estimation of the variances of income innovations.

8.1. Education

The benchmark case shows optimal behavior for the middle education group. It is a relevant question whether the differences in age-earnings profiles and income risk characteristics across education groups lead to variation in optimal risky shares across these groups. As certified education is relatively easily observable, incorporating it into asset allocation advice would be straightforward unless there exist education-dependent information costs or other forms of transaction costs with respect to achieving a certain asset allocation. Like Cocco, Gomes, and Maenhout (2005), we ignore any such frictions in the life cycle model.

Figure 2 and Table 7 summarize the parameters that distinguish education groups in the life cycle model. The high education groups exhibit considerably higher expected non-financial income than the two lower education groups. Education-related patterns in the variances of income innovations are not obvious.

Figure 9 Expected trajectory of the allocation to the risky asset for different education groups

Panel A Females

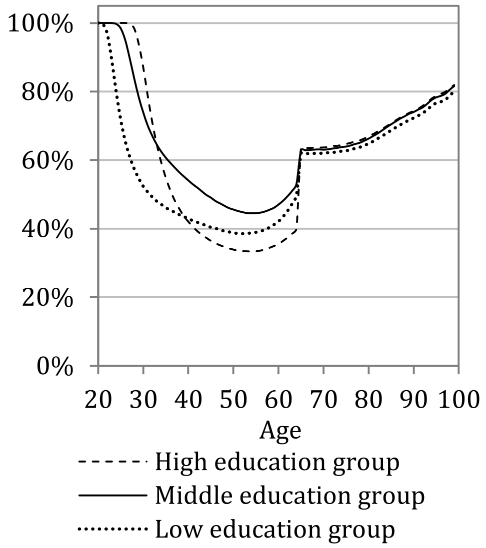

Panel B Males

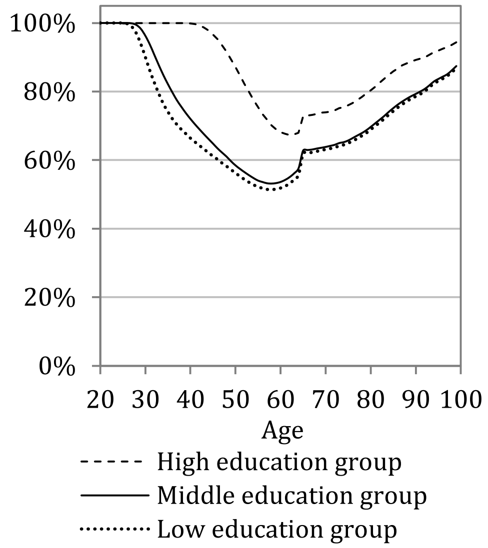

The graph shows the mean of the simulated allocation to the risky asset over the life cycle for different education groups.

Figure 9 shows that the overall life cycle pattern of the optimal risky share is similar to that in the benchmark case. The difference in optimal allocation between females with low and middle education is driven by the larger magnitude of persistent shocks faced by household heads with low education. For males, the differences in optimal allocation between the low and middle education groups are negligible. The larger transitory shocks faced by the low education group do not seem to induce a lower risky share. For males in the high education group, however, the model predicts a considerably higher risky share over the entire life cycle. Two effects drive this result. First, the high education group faces slightly smaller persistent shocks; its human capital is less risky. Second, at any given age, the high education group has higher human capital. This induces diversification demand for risky financial assets.

At the education group level, the welfare losses of following suboptimal policies (not reported) are similar to those in the benchmark case. Yet, given the similarity in and the confidence intervals around the predicted allocations, particularly for the two lower education groups, it is not obvious that education-specific portfolio advice would be justified. We consider the benchmark results to be robust to differences in income process parameters across education groups.

8.2. Zero non-financial income draws

In Section 4.3, we discuss the incidence and treatment of zero non-financial income observations in our sample. In the benchmark case, we ignore these observations when characterizing the distribution of income innovations. As we cannot preclude that these extreme income draws constitute real non-financial income risk rather than measurement error, we examine the sensitivity of the benchmark results to the introduction of a positive probability of zero non-financial income realizations.

We treat zero non-financial income draws in a discrete way according to (14). As in the benchmark case, we limit the analysis to the middle education group. The probability of a zero non-financial income event then is $\mathrm{p_z} = 0.41\%$ for females and $\mathrm{p_z} = 0.16\%$ for males during each year (see Table 4) for $t_1 \leq t \leq K$.

Figure 10 Optimal consumption policy with positive probability of zero non-financial income realizations

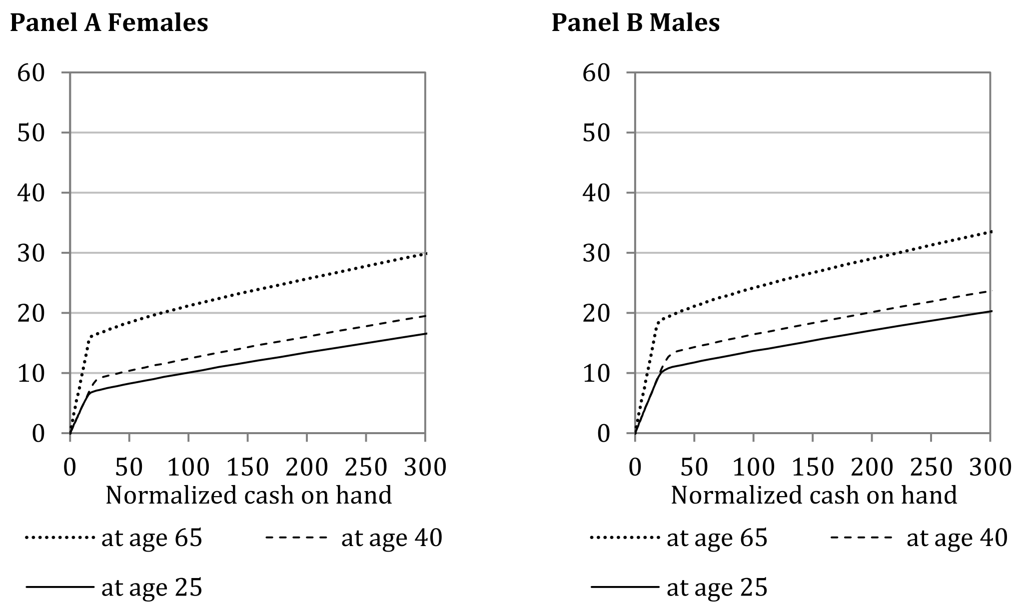

The graph shows optimal normalized consumption as a function of normalized cash on hand (in thousands of € and 2006 prices) for the middle education group when the probability $\mathrm{p_z}$ of a zero non-financial income draw is 0.41% for females and 0.16% for males each year during the pre-retirement period.

As retirement income continues to be risk-less, the optimal consumption policy changes only for the pre-retirement period. This can be seen by a comparison of the consumption policy for age 65 and the consumption policies for age 40 and 25 in Figure 10. Before retirement, the slope of the consumption policy is decreased for low levels of cash on hand. Agents with little cash on hand save considerably more to achieve a level of cash on hand that enables them to maintain consumption if they realize zero non-financial income. Figure 11 shows the corresponding allocation to the risky asset. Again, the allocation policy only changes for low levels of cash on hand. As cash on hand is insufficient or only barely sufficient for agents to maintain consumption in the face of zero-income realizations, the risky asset share goes down considerably relative to the benchmark case.

Figure 11 Optimal allocation policy with positive probability of zero non-financial income realizations

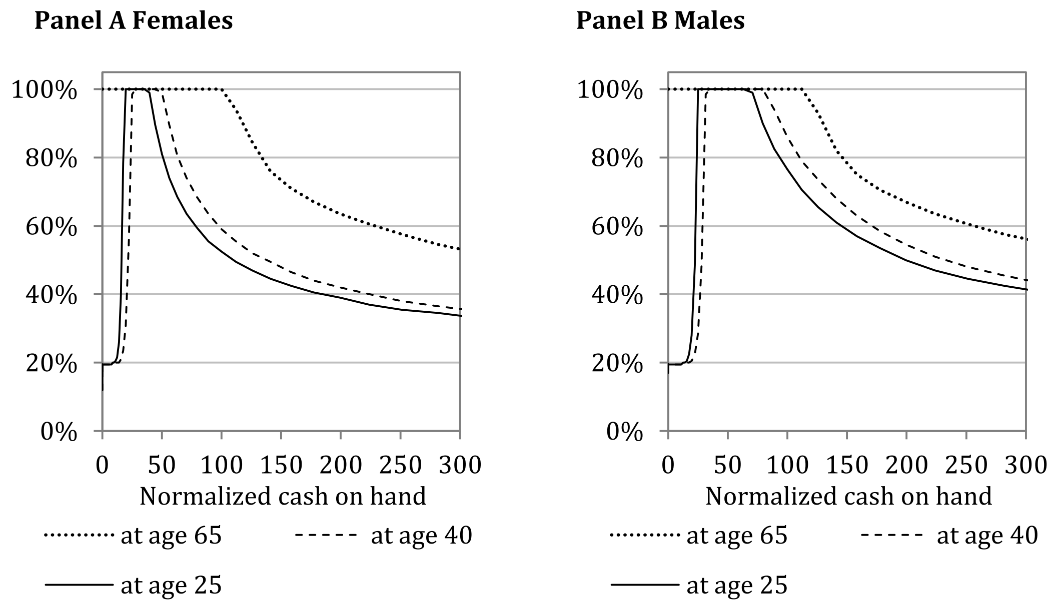

The graph shows the optimal allocation to the risky asset as a function of normalized cash on hand (in thousands of € and 2006 prices) for the middle education group when the probability p_z of a zero non-financial income draw is 0.41% for females and 0.16% for males for each year during the pre-retirement period.

Figure 12 compares the optimal risky share over the life cycle for the benchmark case and the case where disastrous income draws are possible. Given that in a life cycle context, low levels of cash on hand are typical for young agents, the result is consistent with the changes to the consumption and allocation policies. Agents in their 20s start out with a heavily reduced allocation to risky assets relative to the benchmark case, but the risky asset share quickly reverts to the benchmark policy as cash on hand grows sufficiently large for agents not to reduce consumption in the face of an extremely low income realization. From agents' late 20s onwards, the benchmark result is strongly robust to the introduction of a small, positive p_z.

Figure 12 Expected trajectory of the allocation to the risky asset with positive probability of zero non-financial income realizations

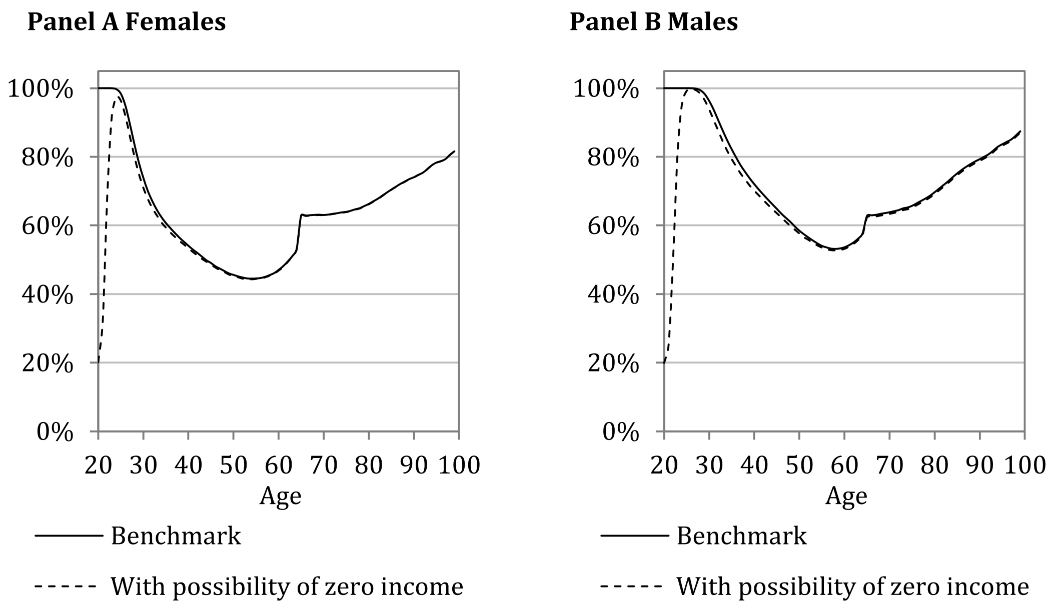

The graph shows the mean of the simulated allocation to the risky asset over the life cycle when the probability p_z of a zero non-financial income draw is 0.41% for females and 0.16% for males for each year during the pre-retirement period.

8.3. Maximum lag length in the variance decomposition

The choice of L determines the number of sample variances of differenced first-stage regression residuals used in the variance decomposition regression by setting the maximum lag length for the differenced residuals (see 4.5). Contrary to the theoretic result in (17), which is based on our specification of the non-financial income process, we find that the series of $\widehat{Var}(d_{i,t,l})$ are not linear in l. The relationship appears to be concave for both males and females in all education groups.

The slope of a straight line fitted to the $\widehat{Var}(d_{i,t,l})$ series estimates σ_ζ^2; half of the vertical intercept measures σ_ε^2. In the context of a concave relationship, a larger L reduces the slope of the fitted line and increases its intercept, thus 'shifting' estimated variance from the persistent shocks to the transitory ones. Figure 13 illustrates this point. The graphs show the sample variances of differenced first-stage regression residuals for $L = 20$ and fitted straight lines for $L = \{5, 10, 20\}$ for males and females in the middle education group. Table 11 reports the corresponding estimates of σ_ζ^2 and σ_ε^2.

Figure 13 The influence of the choice of *L* on estimates of shock variance

Panel A Females **Panel B Males**

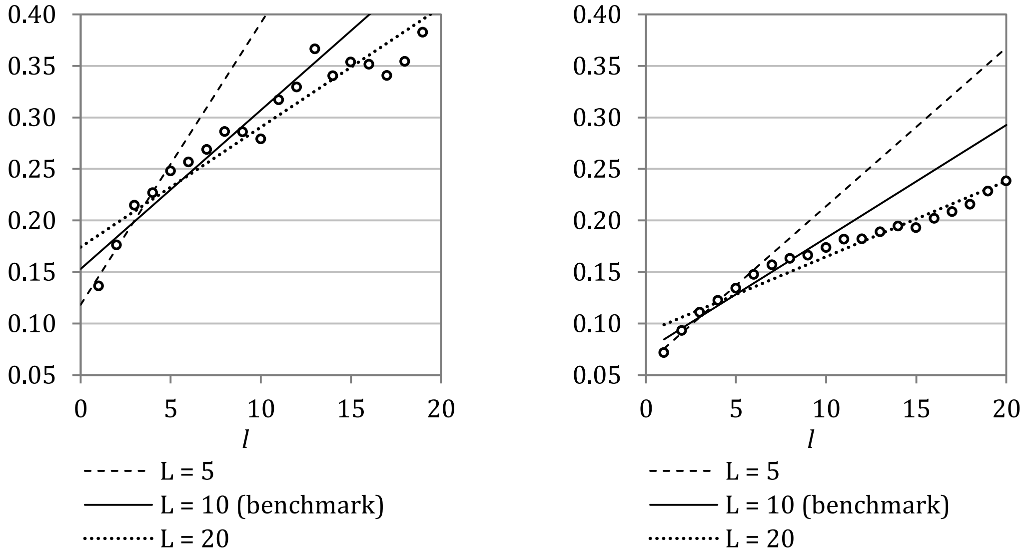

The graph shows OLS fits of equation (17) for males and females in the middle education group; the relationship between $\widehat{Var}(d_{i,t,l})$ and l is concave, and the choice of larger values of L results in larger intercepts and smaller slopes of the fitted lines, which corresponds to larger estimated variances of transitory shocks and lower estimated variances of persistent shocks.

Table 11 Shock variance estimates for different values of *L*

	Females			Males		
	L = 5	L = 10	L = 20	L = 5	L = 10	L = 20
σ_ζ^2 (variance of innovation to persistent component)	0.0274	0.0154	0.0116	0.0154	0.0110	0.0073
σ_ε^2 (variance of transitory shocks)	0.0591	0.0764	0.0871	0.0301	0.0368	0.0457
σ_ζ (STDEV of innovation to persistent component)	16.5%	12.4%	10.8%	12.4%	10.5%	8.6%
σ_ε (STDEV of transitory shocks)	24.3%	27.6%	29.5%	17.3%	19.2%	21.4%

The table shows estimates of the variance and standard deviation of the innovations to the non-financial income process for different values of L for females and males in the middle education group.

The larger variance estimates of persistent shocks obtained for smaller values of L result in higher riskiness of human capital. Increases in the riskiness of human capital, in turn, make agents less willing to invest in the risky asset. That is why, as it is shown in Figure 14, the average allocation to the risky asset is lower for lower values of L.

The overall pattern of the risky share over the life cycle is unchanged from the benchmark case. In that sense, the benchmark result is robust to the choice of L. However, the choice of L does affect the level of the risky share. The differences in predicted risky asset shares for $L = 5$ and $L = 20$ exceed 30 percentage points for certain ages.

Figure 14 Sensitivity of the expected allocation to the risky asset to the choice of *L*

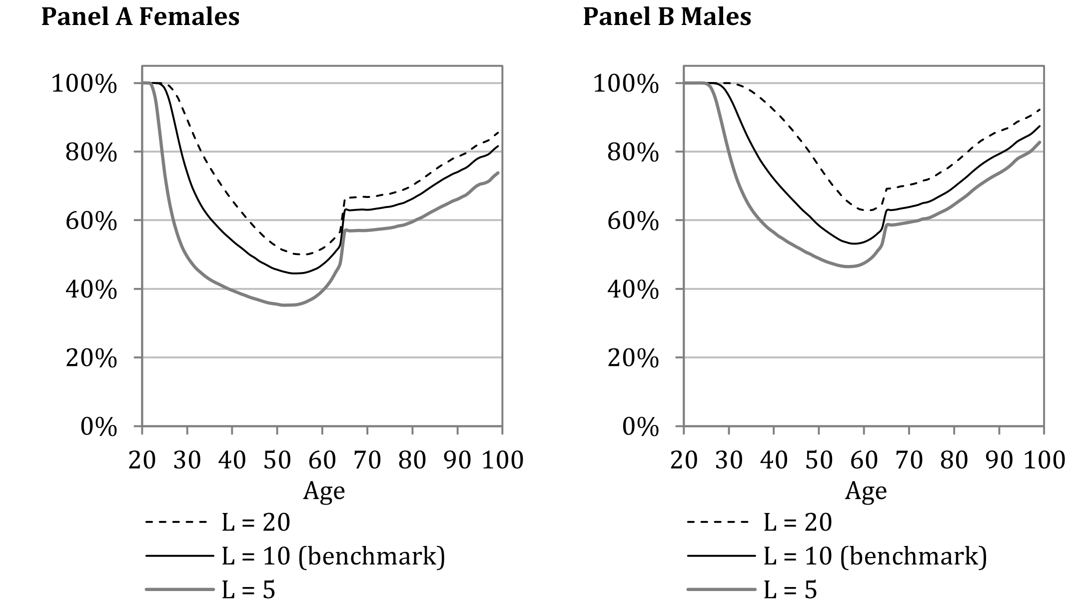

The graph shows the average simulated realizations of the allocation to the risky asset over the life cycle for different choices of L for males and females in the middle education group; higher values of L result in lower estimates of σ_ζ^2 and hence in a higher allocation to the risky asset.

8.4. Risk-aversion

One of the most widely-discussed phenomena in finance is the so-called "equity premium puzzle". With CRRA utility, estimated risk premia in equity markets can only be rationalized under the assumption of risk aversion coefficients that are considered empirically implausible. The γ employed in the benchmark case of this paper is chosen based on Cocco, Gomes, and Maenhout (2005). At $\gamma = 5$, the optimal risky asset share is very close to or equal to 100% over the entire life cycle. At $\gamma = 2$, it equals 100%. The benchmark results are not robust to a change in the assumed coefficient of relative risk aversion.

9. Discussion

Three effects underlie the results presented above. First, cet. par., larger human capital that is only weakly correlated with stock returns increases the optimal risky share in the financial portfolio as it induces a diversification demand for the risky asset. Second, the risky-share effect of human capital at any point during the life cycle depends on the ratio of human capital to financial wealth. Larger financial wealth reduces the impact of human capital on portfolio allocation. Third, cet. par., higher riskiness of human capital (higher background risk), which is to a large extent driven by the variance of persistent shocks, exerts downward pressure on the optimal risky share; a higher variance of transitory shocks has a qualitatively similar but quantitatively smaller effect.

The direction in which the optimal risky share moves at any point during the life cycle can be explained by the interaction of these effects. Young agents, whose total wealth consists almost entirely of human capital, are predicted to allocate 100% of their financial portfolio to the risky asset. During this phase, the risky asset share only deviates from 100% if zero-income realizations are allowed for, in which case the combination of low cash on hand and the prospect of disastrously low income induces both higher savings and a lower risky asset share in the financial portfolio. Regardless of whether zero-income realizations are allowed for, the risky asset share then follows a pronounced U-shape until retirement and keeps rising gradually during retirement. The downward movement in risky asset share 'into the U' is driven by the fast build-up of financial wealth, while the dominating factor behind the rebound before retirement is the rapid reduction in the riskiness of human capital. The gradual rise in the predicted risky asset share is driven by the reduction in financial wealth during this period.

Qualitatively, the life cycle pattern for the optimal risky share is robust to gender-related and education-related heterogeneity in non-financial income characteristics as well as to parameter choice in the variance decomposition, but not to changes in the coefficient of relative risk aversion.

Differences in risky shares across education groups follow directly from differences in non-financial income risk. This is also why females, whom we find to face larger income shocks than males across all education groups, are predicted to maintain lower risky shares than males. The differences in average risky asset shares across gender and education groups are for the most part not larger than the differences in optimal allocation that can be produced by changing the estimation procedure for obtaining the shock variances. Therefore, we cannot reliably quantify differences in investment policies potentially mandated by investor heterogeneity. A more detailed approach to modeling the variances of income innovations might increase confidence in the differences in predicted asset allocation.

The optimal policy leads to economically significant differences in utility relative to classic portfolio policies and rules of thumb, but (not) adhering to the optimal consumption policy is just as influential. So, the results suggest that rational investors whose preferences and choices are accurately reflected in the model stand to lose substantially from not following the predicted optimal behavior. As part of a discussion of the question whether this mandates changes in actual behavior, we examine the extent to which the life cycle model predicts actual investor behavior. We discuss this separately for savings and asset allocation behavior.

9.1. Savings behavior

The results of the life cycle optimization exercise are largely in line with the classic life cycle savings and consumption model in the sense that agents first save and then consume out of savings. Borrowing by very young agents is ruled out by the imposed constraints.

Yet, observed savings and consumption behavior in Germany does not seem to match that prediction. Börsch-Supan and Essig's (2005) discussion of household savings behavior in Germany suggests that the prediction of our model is broadly consistent with observed savings behavior for the first half of agents' lives, even though actual savings rates seem to be lower than predicted by the model. For the second half of agents' lives, however, the optimal consumption policy represents the opposite of the observed behavior. High-income and median-income households in Germany appear to keep saving even during retirement, albeit at a considerably lower rate. Börsch-Supan et al. (2001) call this "the German savings puzzle". This behavior might partly be due to the fact that in the past households were faced with high replacement ratios in statutory pension insurance. The availability of sufficient levels of non-financial income should have reduced the expected need to dissave during retirement, but this only raises the question why substantial savings are accumulated in the first place and why consumption is not higher during retirement. In a cross-sectional analysis discussed by Börsch-Supan and Essig (2005), saving for retirement is – together with saving for "unexpected events" - rated as "important" by a larger fraction of households than any other savings motive, which is in line with the precautionary savings behavior in our model for the period before retirement. A bequest motive, which our model does not consider but which could theoretically bring the model prediction more in line with observed behavior for retired agents, is rated as being of very low importance. This suggest that the omission of such a motive from the model does not constitute a major shortcoming of the analysis.

One potential explanation for the difference between predicted and actual behavior during the second half of the life cycle then is that the precautionary savings motive remains very strong because agents perceive income and/or expenditures during retirement as risky while our model assumes risk-less retirement income and ignores expenditures. Incorporating risks during retirement could allow the model prediction to match actual savings behavior more closely. At the same time, it is possible that actual behavior is driven by biased perceptions of risks during retirement and that altering savings policies would increase utility.

9.2. Asset allocation behavior

Information on risky asset shares in German household portfolios is limited, particularly with respect to current life cycle allocation behavior. Barasinska et al. (2008) report very low stock market participation rates of about 30% (for both direct and indirect equity holdings) of households using SOEP data for 2001 to 2005. Börsch-Supan and Essig (2002) report lower rates for the 1990s. They report that participation increases with education, but they also show that most of the effect is driven by higher income and higher wealth as opposed to education itself. In our model, agents always participate in the stock market. For those households that do participate, Börsch-Supan and Essig (2002) show data on risky shares for 1993 and 1998 that do in fact suggest a U-shaped pattern of the risky share as function of age. However, the authors themselves caution that due to its purely cross-sectional character, the data is likely to suffer from cohort effects and representativeness issues. In any case, the level of the risky asset share at any age is considerably lower than predicted by the model. Slicing the data by the amount of financial wealth, the authors find a pattern that is consistent with the general shape of the optimal allocation policy (as a function of cash on hand) in our model, but the level of the risky share is again considerably lower than predicted by the model. Allowing for zero non-financial income realizations leads to recommended risky shares that are more in line with actual behavior for young households, though.

Overall, while a complete assessment of the model's ability to describe actual allocation behavior is not feasible, we conclude based on the available evidence that large discrepancies exist between predictions and observed behavior. It might be possible to achieve a closer match by making popular assets such as building society saving contracts available and allowing for real estate investment decisions or modeling government incentive policies and subsidies for certain forms of saving. Yet, it is not obvious that this would improve the analysis. It is conceivable that actual behavior cannot easily be rationalized because it might be suboptimal. Equity market nonparticipation, which seems to be the most common policy actually followed by German households, underperforms remarkably on a consumption-equivalent basis in our model even when compared to policies like the classical Merton rule that ignores non-financial income altogether. Yet, as mentioned before, this result is obtained in the context of a two-asset universe.

10. Conclusion

We numerically solve a life cycle model of portfolio allocation and consumption in a two-asset universe with CRRA-utility maximizing agents receiving exogenous stochastic non-financial income. The non-financial income process is calibrated to German data. For the most part, we follow Cocco, Gomes, and Maenhout (2005) in terms of methodology.

Despite perceived differences in labor market structure and welfare policies between Germany and the U.S., we obtain comparable estimates for income risk. We confirm that this risk is largely idiosyncratic, even though we do estimate correlations between labor income innovations and excess stock returns that are significantly different from zero. Meanwhile, we caution that the estimated variance of income innovations depends on the estimation procedure.

The life cycle model shows that larger human capital leads to a higher risky asset share in the financial portfolio while more risky human capital leads to a lower share. Given our calibration results and our choice of benchmark parameters, this leads to very high risky asset shares on average early in life unless we allow for a small probability of an extremely low income realization. The average risky share then decreases until midlife as human capital decreases and as growing financial wealth reduces the impact of human capital on portfolio allocation, but it never falls substantially below 50% in the benchmark case. It even bounces back before retirement as the remaining risk in human capital decreases rapidly as investors approach retirement. Qualitatively, this life cycle pattern is robust to investor heterogeneity related to education and gender as well as to differences in shock variance estimation. As is common, the result only holds under the assumption of a relatively high coefficient of relative risk aversion. Quantitatively, the predicted life cycle asset allocation policy does change with education and gender, but we do not translate this into precise allocation advice as the level of risky asset shares depends as much on choices with respect to shock variance estimation.

We confirm Cocco, Gomes, and Maenhout's (2005) finding that classic portfolio rules and rules of thumb lead to economically significant utility losses relative to the optimal asset allocation policy as these policies do not or not fully take into account the properties of human capital.

However, the implications for real-world investing are not obvious at this stage. Predicted behavior is inconsistent with observed savings behavior during the second half of agents' lives and with observed risky asset shares over the entire life cycle. Actual equity market participation rates are exceptionally low and the risky asset share of those households which do participate is considerably lower than predicted by the model. There are two possible explanations, and we deem it likely that both are relevant to some extent. First, the model might need to be extended or altered to reflect investor preferences and choices more accurately. Second, if investors were fully rational and their preferences and choices were captured by the model, observed behavior would be suboptimal and investors could gain from following the life cycle consumption and allocation policies discussed above.

With respect to this discussion, the paper can serve as a basis for further analysis. A relevant question is, for example, what form and level of retirement income risk or retirement expenditure risk is needed for the model to reproduce actual savings behavior during retirement. This could help understand how likely it is that investors are actually following suboptimal policies.

Moreover, the analysis of income risk and its impact on optimal portfolio choice could be extended to incorporate correlations between income innovations and equity returns at the industry level. An industry portfolio, for which short-selling constraints would be relaxed, could then be added to the investment universe to allow for conclusions about whether industry-specific characteristics of non-financial income and hence human capital imply that the optimal risky portfolio differs from the market portfolio.

Appendices

Appendix 1

We define the adjusted value function in the dynamic programming equation, which we are in fact solving instead of (12), as

$$\phi_t(x_t e^{-v_t}) = \max_{c_t, w_t}\{u(c_t e^{-v_t}) + \delta p_{t+1} E_t(\phi_{t+1}(x_{t+1} e^{-v_{t+1}}) e^{(v_{t+1}-v_t)(1-\gamma)})\}$$

Let us show that $\phi_t(x_t e^{-v_t}) = V_t(x_t, v_t) e^{-v_t(1-\gamma)}$, $t_1 \leq t \leq T$.

The identity is true in the last period, i.e. for $t = T$:

$$\phi_T(x_T e^{-v_T}) = u(x_T e^{-v_T}) = u(x_T) e^{-v_T(1-\gamma)} = V_T(x_T) e^{-v_T(1-\gamma)}.$$

Then, by induction, it is also true in all preceding periods, i.e. for $t_1 \leq t < T$:

$$\phi_t(x_{t-1} e^{-v_{t-1}}) =$$

$$= \max_{c_t, w_t}\{u(c_t e^{-v_t}) + \delta p_{t+1} E_t(\phi_{i,t+1}(x_{i,t+1} e^{-v_{i,t+1}}) e^{(v_{i,t+1}-v_{i,t})(1-\gamma)})\} =$$

$$= \max_{c_t, w_t}\{u(c_t) e^{-v_t(1-\gamma)} + \delta p_{t+1} E_t(V_{i,t+1}(x_{t+1}, v_{t+1}) e^{-v_{i,t}(1-\gamma)})\} =$$

$$= V_t(x, v) e^{-v_t(1-\gamma)}.$$

In the first period, $v_{t_1} = 0$, and $\phi_{t_1}(x_{t_1} e^{-v_{t_1}}) = V_t(x_t, v_t) = U_t$.

Appendix 2

The first-stage regressions that we use to characterize $f(t)$ use a measure of non-financial disposable income as dependent variable. We construct this income measure based on an SOEP-provided measure of household post-government income. Table 12 presents the income and tax items that household post-government income is based on and documents whether we classify these items as financial or non-financial. Income items classified as financial are excluded from our measure of non-financial disposable income. Tax items classified as financial are added back.

Several items have been aggregated for the purpose of this appendix. For exact references to each individual SOEP variable used, please contact the authors.

Table 12 Classification of income items

SOEP Category	Income item	Classification	Comment
Labor earnings	Wages/salary from primary job	Non-financial	
	Income from secondary employment	Non-financial	
	Income from self-employment	Non-financial	If the household head reports being primarily self-employed, we exclude the head (and hence the household) from our sample
	Military/community service pay	Non-financial	
	13th/14th monthly salary	Non-financial	
	Christmas bonus, vacation bonus	Non-financial	
	Profit-sharing, other bonuses	Non-financial	
	Severance payments	Non-financial	
	Commuting expense allowances	Non-financial	
Private Transfers	Alimony payments	Non-financial	
	Other private transfers	Non-financial	
Public Transfers	All forms of public unemployment benefits including related payments	Non-financial	
	Subsistence allowance	Non-financial	
	Old-age transition benefit	Non-financial	
	All forms of social welfare ("social assistance") benefits and related payments	Non-financial	
	Maternity benefit	Non-financial	
	Student grants	Non-financial	
	Child allowance	Non-financial	
	Nursing allowance	Non-financial	
	Housing support for owner-occupiers	Financial	A discount on an asset purchase; assumption: home ownership as a portfolio allocation decision
Social Security Pensions	Statutory old-age and disability pension, including widows/orphans payments	Non-financial	
	Civil servants pensions,including widows/orphans payments	Non-financial	
	Miners and farmers pensions, including widows/orphans payments	Non-financial	
	War victims pensions, including widows/orphans payments	Non-financial	
	Statutory accident insurance, including widows/orphans payments	Non-financial	
Private Retirement Income	Company pensions, including widows/orphans payments	Non-financial	Related to prior service provision; a form of deferred labor income
	Supplementary benefits for civil servants, including widows/orphans payments	Non-financial	Related to prior service provision; a form of deferred labor income
	Private and other pensions, including widows/orphans payments	Financial	Outcome of a portfolio allocation decision
Asset Income	Income from rental and leasing net of operating expenses	Financial	
	Interest/dividend income, capital gains	Financial	
Taxes	Federal Taxes (= straight income taxes)	Partly financial, partly non-financial	Separated into a part pertaining to financial income and a part pertaining to non-financial income; see separate appendix
	Social Security Taxes	Non-financial	Social security taxes are related exclusively to those income items that we classify as non-financial

The table shows the classification of (categories of) income items into financial and non-financial income. The item "federal taxes" is separated into taxes pertaining to financial income and taxes pertaining to non-financial income. The required disaggregation of income taxes is described in a separate appendix.

Appendix 3

The income taxes in the SOEP data are estimated by the SOEP and reported for all forms of income and all household members combined, i.e. they are reported as one number per household per year. The same applies to social security taxes, but we do not discuss them since they are related exclusively to income items that we classify as non-financial and hence do not require adjustment.

When calculating household disposable non-financial income, we seek to remove financial income from total household disposable income. This requires that we allocate part of the SOEP-provided income taxes to the financial income items we remove.

We adopt the simplifying assumptions described in Schwarze (1995). Most importantly, we assume that filers never itemize expenses (except for expenses related to rental income) but rather take standard deductions wherever possible. We ignore losses from capital and rental income. However, unlike Schwarze (1995), we treat each household as a single filing unit.

Under German tax law, taxable income is determined in a multi-step procedure outlined in Schwarze (1995). The major steps are (1) the determination of the taxable part of each form (category) of income a filer receives and (2) the subtraction of deductions or itemized expenses that apply to overall income (as opposed to applying to any particular form of income dealt with under (1)). Depending on the form of income, (1) may involve applying some fraction determining the taxable portion of income reported under a particular category, applying a tax-exemption and/or subtracting either an expense deduction or a sum of itemized expenses from income in that category. For our purposes, only step (1) is relevant, because considering deductions applied after the sum of the taxable parts of each form of income is calculated only complicates the computations while having no impact on our allocation of SOEP-provided taxes to individual income categories.

To implement step (1), we apply standard deductions and/or exemptions to different forms of labor income, to interest and dividend income and capital gains and to civil servants pensions. For rental income, we deduct reported operating expenses from gross rental income. In case of civil servants pensions, we extrapolate one tax exemption for years for which the correct parameter value is not known. For all other forms of taxable pension income, we calculate taxable income using specific ratios that determine the taxable part of each type of pension. We choose the according ratios adopting all simplifying assumptions explicitly stated in Schwarze (1995) so that there is no variation in those ratios across individuals and households within a specific pension category. For later survey years, major changes in pension taxation require similar simplifying assumptions. We ignore one minor standard expense deduction related to non-civil servants social security pensions and private pensions.

This procedure yields taxable income by income category. The sum of these estimates is overall taxable income before any deductions that would usually be applied in step (2). We then allocate the SOEP-provided, estimated taxes proportionally to different forms of income, which allows us to quantify after-tax financial income and hence to move from household disposable income to household non-financial disposable income.

Our calibration results are robust to alternative approaches to allocating taxes to financial income and to higher variation in tax rates.

For calendar years 1983, 1984, 2002 and 2003, certain pension items subject to different taxation rules are reported in aggregated form, preventing the estimation of taxable income. We drop the affected observations. The effects of this are discussed in 4.2.1.

References

Barasinska, Nataliya, Dorothea Schäfer, and Andreas Stephan, 2008, Do Risk Attitude and Diversification Match? Evidence from German Household Portfolios, DIW Berlin.

Beblo, Miriam and Julio Robledo, 2008, The wage gap and the leisure gap for double-earner couples, Journal of Population Economics 21, 281-304.

Benzoni, Luca, and Olena Chyruk, 2009, Investing over the Life Cycle with Long-Run Labor Income Risk, Working paper, Federal Reserve Bank of Chicago.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2007, Portfolio Choice over the Life cycle when the Stock and Labor Markets Are Cointegrated, The Journal of Finance 62, 2123-2167.

Bodie, Zvi, Robert C. Merton, and William F. Samuelson, 1992, Labor supply flexibility and portfolio choice in a life cycle model, Journal of Economic Dynamics and Control 16, 427-449.

Börsch-Supan, Axel, and Lothar Essig, 2002, Stockholding in Germany (Mannheim Research Institute for the Economics of Aging (MEA), University of Mannheim, Germany).

Börsch-Supan, Axel, and Lothar Essig, 2005, Household Saving in Germany, Results of the First SAVE Study, in David A. Wise, ed.: Analyses in the Economics of Aging (University of Chicago Press).

Börsch-Supan, Axel, Anette Reil-Held, Ralf Rodepeter, Reinhold Schnabel, and Joachim Winter, 2001, The German Savings Puzzle, Research in Economics 55, 15 - 38.

Calvet, Laurent E., and Paolo Sodini, 2010, Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios, NBER Working paper 15859, HEC Paris and Stockholm School of Economics.

Campbell, John, João Cocco, Francisco Gomes, and Pascal Maenhout, 2001, Investing Retirement Wealth: A Life cycle Model, in John Campbell and Martin Feldstein, eds:. Risk Aspects of Social Security Reform (University of Chicago Press, Chicago).

Carroll, Christopher D., 1992, The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence, Brookings Papers on Economic Activity 2, 61-156.

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

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

Chan, Lewis Y., and Luis M. Vicera, 2000, Asset Allocation with Endogenous Labor Income: The Case of Incomplete Markets, Preliminary, Hong Kong

University of Science and Technology, Graduate School of Business Administration, Harvard University.

Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout, 2005, Consumption and Portfolio Choice over the Life Cycle, The Review of Financial Studies 18, 491-533.

Deaton, Angus, 1991, Saving and Liquidity Constraints, Econometrica 59, 1221-1248.

Davis, Steven J., and Paul Willen, 2000, Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice, The Center for Research in Security Prices Working Paper No. 523, University of Chicago Graduate School of Business.

Davis, Steven J., Felix Kubler, and Paul Willen, 2006, Borrowing Costs and the Demand for Equity over the Life Cycle, The Review of Economics and Statistics 88(2), 348-362.

Dustmann, Christian, and Arthur van Soest, 1998, Public and private sector wages of male workers in Germany, European Economic Review 42, 1417 – 1441.

Elmendorf, Douglas W., and Miles S. Kimball, 2000, Taxation of Labor Income and the Demand for Risky Assets, International Economic Review 41, 801-832.

Frick, Joachim R., 2010, Introduction to the German Socio-Economic Panel (SOEP), DIW Berlin, retrieved on January 30, 2010, from http://www.diw.de/documents/dokumentenarchiv/17/diw_01.c.353304.de/soep-intro_march2010.pdf

Frick, Joachim R., Stephen P. Jenkins, Dean R. Lillard, Oliver Lipps, and Mark Wooden, 2007, The Cross-National Equivalent File (CNEF) and its Member Country Household Panel Studies, Schmollers Jahrbuch 127, 627 – 654.

Gollier, Christian, and John W. Pratt, 1996, Risk Vulnerability and the Tempering Effect of Background Risk, Econometrica 64, 1109-1123.

Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira, 2008, Optimal Life cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life cycle Funds, American Economic Review 98, 297-303.

Gomes, Francisco, and Alexander Michaelides, 2005, Optimal Life cycle Asset Allocation: understanding the Empirical Evidence, The Journal of Finance 60, 869-904.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker, 2002, Consumption over the Life Cycle, Econometrica 70, 47-89.

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

Heaton, John, and Deborah Lucas, 2000, Portfolio Choice in the Presence of Background Risk, The Economic Journal 110, 1-26.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1994, The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving, in Carnegie-Rochester Conference Series on Public Policy, Vol. 40, ed. by A. H. Meltzer and C. I. Plosser. Amsterdam: North Holland, 59-125 .

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1995, Precautionary Saving and Social Insurance, Journal of Political Economy 103, 360-399.

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

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

Koo, Hyeng K., 1999, Consumption and portfolio selection with labor income: A discrete-time approach, Mathematical Methods of Operations Research 50, 219-243.

Lynch, Anthony W, Sinan Tan, 2009, Labor income dynamics at business-cycle frequencies: implications for portfolio choice, Working paper, New York University and NBER, Fordham University.

MaCurdy, Thomas E., 1982, The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis, Journal of Econometrics 18, 83-114.

Malkiel, Burton G., 1996, A Random Walk Down Wall Street: The Best and Latest Investment Advice Money Can Buy 6th edition, W. W. Norton & Company, New York.

Markowitz, Harry, 1952, Portfolio Selection, The Journal of Finance 7, 77-91.

Merton, Robert C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, Review of Economics and Statistics 51, 247-257.

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

Mossin, Jan, 1968, Optimal Multiperiod Portfolio Policies, The Journal of Business 41, 215-229.

Polkovnichenko, Valery, 2007, Life cycle Portfolio Choice with Additive Habit Formation preferences and Uninsurable Labor Income Risk, The Review of Financial Studies 20, 83-124.

Poterba, James M., and Andrew A. Samwick, 2003, Taxation and household portfolio composition: US evidence from the 1980s and 1990s, Journal of Public Economics 87(1), 5-38.

Rosenfeld, Rachel A., Heike Trappe, and Janet C. Gornick, 2004, Gender and Work in Germany: Before and after Reunification, Annual Review of Sociology 30, 103-124.

Samuelson, Paul A., 1969, Lifetime Portfolio Selection By Dynamic Stochastic Programming, The Review of Economics and Statistics 51, 239-246.

Schwarze, Johannes, 1995, Simulating German Income and Social Security Tax Payments Using the GSOEP, Cross-National Studies in Aging, Program Project Paper No. 19, Syracuse University.

Socio-Economic Panel (SOEP), Data for years 1984-2009, Version 26, SOEP, 2010.

Statistisches Bundesamt, 2010, Bevölkerung und Erwerbstätigkeit - Sterbetafel Deutschland 2007/09, Wiesbaden, retrieved on May 11, 2011 from http://www.destatis.de/jetspeed/portal/cms/Sites/destatis/Internet/DE/Content/Statistiken/Bevoelkerung/GeburtenSterbefaelle/Tabellen/Content100/SterbetafelDeutschland,property=file.xls.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 2004, Cyclical Dynamics in Idiosyncratic Labor Market Risk, Journal of Political Economy 112, 695-717.

Thomson Reuters Datastream, 2011, data on financial returns.

Wagner, Gerd G., Joachim R. Frick, and Jürgen Schupp, 2007, The German Socio-Economic Panel Study (SOEP) – Scope Evolution and Enhancements, Schmollers Jahrbuch 127, 139 – 169.

Vasicek, Oldrich, 1977, An equilibrium characterization of the term structure, Journal of Financial Economics 5, 177-188.

Viceira, Luis M., 2001, Optimal Portfolio Choice for Long-Horizon Investors with Non-Tradable Labor Income, Journal of Finance 56, 433-470.