

WHAT IS AT STAKE WHEN DETERMINING LIFETIME ASSET ALLOCATION?***

Gunter Löffler*

This version: December 11, 2000

Abstract

This paper examines the effects of deviating from the optimal life-cycle asset allocation. The magnitude of likely deviations is gauged through an analysis of differences between available life-cycle funds. The motivation is that due to search costs, the complexity of the decision problem and other factors, an investor may choose a fund with asset allocation profile A even though profile B would be optimal. The associated utility losses are quantified using a model of lifetime consumption and portfolio choice. The results suggest that utility losses are modest. In many cases, they are smaller than the utility loss which arises if the return on financial assets is reduced by 0.3% per annum. The results help to focus on the most important decisions when saving for retirement, and to balance the costs and benefits of customising saving products.

Key words: Asset allocation, portfolio selection, utility loss, financial planning, life-cycle funds, retirement savings

JEL classification: D91, G11

* Gunter Löffler, University of Frankfurt, Chair of Banking and Finance, P.O. Box 111932, 60054 Frankfurt (Main), Germany. Tel.: ++49-69-79823143; fax: ++49-69-79822143. E-mail address: gloeffler@wiwi.uni-frankfurt.de

*** I thank Niklas Siebenmorgen and participants of the 2000 annual INQUIRE Europe symposium in Venice for helpful comments and suggestions.

1 Introduction

In many countries, the role of defined contribution plans is increasing, and governments are encouraging personal saving to substitute diminishing social security benefits. In consequence, the investment of retirement assets is increasingly directed by households rather than by the state or employer sponsors.

Finding optimal investment rules for long horizons is a complex task. Expected long-run returns and risks are difficult to estimate¹, while solving intertemporal decision problems requires skills even economics curricula at universities seldom comprise. Investors thus have to determine the optimal amount of resources devoted to finding the 'right' asset allocation, while governments and employers may consider the costs and benefits of education programmes. The chosen asset allocation, however, is not the only determinant of investor welfare. James et al. (1999), for example, argue that standardising saving products, i.e. restricting choice, may actually increase investor utility as it may lower administration costs for these products.

When facing such trade-offs, estimates of utility gains and losses are useful. This paper therefore quantifies the effects of sub-optimal portfolio choices. The magnitude of likely departures from optimality is gauged through an analysis of existing life-cycle funds offered by mutual fund companies. These products are characterised through a predefined way of allocating an individual's fund portfolio through the life-cycle, with age-phased reductions of the equity investment being the typical element. Available products exhibit a wide range of possible allocation patterns, so that different investor needs can be met. However, due to imperfect information, search costs and other factors affecting investor decisions, it seems plausible that investors do not pick the product best suited to them. I thus examine the utility loss that stems from choosing a particular fund, which is available on the market, if another available fund would be optimal. Additionally, sub-optimality is defined as missing the optimal equity allocation consistently by a fixed percentage share.

¹ See Welch(2000) for evidence about differences in expectations about risk premia.

Having defined the extent of sub-optimality, its costs are evaluated using an intertemporal model of consumption and portfolio choice based on Bodie, Merton and Samuelson (1992). Utility losses are defined in terms of certainty equivalents of consumption. The model involves a large set of assumptions about investor preferences and circumstances, which are often hard to derive from empirical data. However, sensitivity analyses demonstrate the robustness of the results. For a representative parameterisation, they can be summarised as follows:

An individual who, from the age of 30 to the age of 65, saves roughly 16% out of current income to build up retirement assets, foregoes about 0.15% (1.5%) of consumption if the equity share in retirement assets is consistently 10 (30) percentage points above or below the optimal one. Choosing the wrong life-cycle product out of a group of products which differ only in the timing of risk reductions leads to a consumption loss in the range of 0% to 1.2%. This compares to a consumption loss close to 1.3% which arises if the return on financial wealth is lowered by 0.3% per annum during the accumulation phase. It is also illustrative to compare the results to those of Rodepeter and Winter (2000). They quantify the impact of non-optimisation in intertemporal consumption-saving decisions, and conclude that applying rules of thumb to solve the intertemporal consumption-saving problem can easily reduce consumption by 10%.²

Several conclusions can be drawn from the analysis. First, choosing a low cost mutual fund or making the right consumption-saving decision may be more important for the efficiency of retirement savings than picking the right asset allocation. Second, restricting the choice of investors to a few standardised products will often have only modest effects on investor utility. Nevertheless, losses from sub-optimality are not negligible, pointing to the importance of investment advice specifically tailored to an individual's preferences and circumstances.

The findings can help investors and investment advisors to focus on the most important decisions when saving for retirement. Mutual fund companies can use the results to find the right balance between offering customised products on the one

² Rodepeter and Winter (2000) do not address sub-optimality in portfolio choice.

hand, and saving costs through standardisation on the other hand. Similarly, the findings are relevant for the design of mandatory, funded social security systems. James et al. (1999) suggest that constraining individual fund choice may lower fund administration expenses, mainly through dampening marketing expenditures of investment companies. Based on an analysis of existing fee structures, they estimate annual expenses with constrained choice to be around 0.5% - 1% lower than in a system where investors are free to invest in the retail market. This points to another trade-off between standardisation and cost efficiency where estimates of utility gains and losses are useful.

Related papers are Goodfellow and Schieber (1997) and Bodie and Kahn (1997), who analyse the average composition of individual asset holdings and check whether age, wealth and other factors affect asset allocation in a way predicted by economic theory. Bodie and Kahn conclude that allocation decisions are 'reasonable', but their criteria are weak, making it difficult to quantify what 'reasonable' means in terms of consumer utility. The criterion most relevant for this paper, for example, is that 'the fraction of assets invested in equities should decline as the investor's age advances' (p.14). In addition, even if average asset holdings of people having the same age, wealth and other characteristics equal those derived from an optimisation model, individual asset holdings might still be far away from the optimum. Canner, Mankiw and Weil (1997) examine violations of the mutual-fund theorem, which states that the composition of risky assets should not vary with risk preferences. They find that violations of this rule evident in investment advisors' recommendations result only in small efficiency losses of about 20 basis points in annual return. Jensen and Sørensen (2000) analyse the utility loss from imposing a minimum interest rate guarantee as an exogenous constraint on asset allocation strategies. If the minimum interest is not too close to the risk-free rate, and risk aversion is similar to the one assumed in this paper, utility losses are also modest.

The remainder of the paper is organised as follows. Section 2 presents life-cycle products offered by German mutual fund companies and motivates the assumptions about likely departures from optimality. In Section 3, I describe the way in

which the effects of sub-optimal portfolio choice are quantified. Section 4 presents the results, which section 5 puts into perspective. Section 6 concludes.

2 Differences among life-cycle funds and likely departures from optimality

The main characteristic of life-cycle products, which are offered by many fund companies throughout the world, is an automatic, age-phased reduction of equity risk contained in an investor's mutual fund portfolio. I decided to analyse the German market for life-cycle products, which was boosted by a change in mutual fund legislation in 1998. The amendment created a new legal type of mutual fund meant to be especially suited for private retirement savings. The fund type was named 'Altersvorsorge-Sondervermögen', which can be loosely translated as 'pension fund'. One of the requirements to qualify for such a fund is a maximum equity share of 75%.³

Many Germany-based mutual fund companies launched such pension funds as soon as this was possible (November 1998). I restrict the analysis to the five biggest German mutual fund companies, which are in alphabetical order (parent bank in parentheses): ADIG (Commerzbank), DEKA (public sector banks), DIT (Dresdner Bank), DWS (Deutsche Bank), and Union (co-operative sector banks). At the end of 1999, the mutual fund volume managed by these five companies amounted to 80.6% of the volume of mutual funds domiciled in Germany.⁴ Written information material describing the funds was obtained directly from the mutual fund companies or from retail branches of the parent company.

Most of the companies launched several Altersvorsorge-Sondervermögen differing in their average equity share. In the sales brochures describing the pension funds, every company proposes how to reduce exposure to risk as the investor's age advances. The usual way is to recommend an exchange of the accumulated pension

³ So far, Altersvorsorge-Sondervermögen do not have a material advantage against other fund types when saving for retirement. The idea behind the creation of this fund type was to increase transparency and to add it to the list of tax-privileged saving vehicles.

⁴ Source: Frankfurter Allgemeine Zeitung, February 9, 2000, p. 31.

fund units into a fund with a smaller equity share. These recommendations determine the character of the life-cycle products offered by the fund companies. In addition, the companies propose to structure the investment in these funds as a saving plan which involves regular contributions until retirement.

In comparing the life-cycle profiles, I focus on those with an initial equity share close to 75%,⁵ and I make several simplifications. For the chosen first stage funds, the stated benchmark equity shares range from 70% to 75%. Given that this difference is not material, and noting that an equity share close to the maximum of 75% would require frequent rebalancing, I choose the equity share to be 70% for all of these funds. Differences in the strategic composition of equity holdings are ignored, as are differences in fixed income or real estate investments.

One company, DWS, states that reducing risk before retirement is not necessary with this particular fund. DIT recommends several switches into other, explicitly named funds (their benchmark equity share is known). The remaining three companies, while being specific on the timing of risk reductions, do not name a particular fund into which the shares of the pension fund should be exchanged. The scheme proposed by ADIG involves a stepwise move into a less risky fund after three quarters of the saving plan. DEKA and Union propose a gradual exchange into less risky funds in the last three and four years before retirement, respectively. To derive an allocation profile for these products, I choose the equity percentage of the 'less risky funds' to be 25%. As the analyses of this paper are conducted on an annual basis, I further assume that any changes in equity holdings come into effect at the beginning of the year.

The allocation profiles of the five life-cycle products are shown in figure 1, assuming the saving plan to last from the age of 30 to the age of 65. In addition, I include the profile stemming from a popular rule of thumb: The optimal investment in equities is 100 minus one's age (cf. Bodie and Crane, 1997).

There are various reasons for believing that the products bought by an individual do not implement the optimal lifetime asset allocation. Even if one of the available

⁵ These are the following funds: AS-AktivDynamik (ADIG), DEKA Privat-Vorsorge AS, DIT-Altersvorsorge 35, DWS Vorsorge AS Dynamik, and Geno AS (Union).

products were indeed the optimal one, the investor may wrongly choose another. This may be due to limited knowledge about other products and one's own optimal allocation profile. The associated information costs are one reason why product choice could be affected by marketing expenses (see Sirri and Tufano (1998) for evidence), or be biased towards companies one already owns funds of, an effect likely to be in force in Germany where much of the mutual fund distribution is still done through the local branches of the fund companies' parent banks.

One sensible way of quantifying likely departures from optimality thus seems to be the following: An individual for whom one of the six equity allocation profiles shown in figure 1 is optimal, wrongly chooses another one out of the remaining five profiles. Since the fund companies seldom give explicit advice on portfolio allocations after retirement, I further assume that the chosen equity share after retirement is optimal regardless of the chosen product.

Another reason for sub-optimal choices may be that mutual fund companies try to standardise products as well as investor advice in order to save costs. Typically, customers will be offered only a limited number of equity allocations to start off with, and only one path for risk reductions for a given initial equity share. Several of the five fund companies analysed here offer life-cycle products with an initial equity allocation of 50%, 70% or 100%. At least in the marketing brochures, the possibility of choosing yet another initial allocation, e.g. by combining two different products, is not mentioned. It thus seems likely that investors make sub-optimal choices because they pick one of the few products offered to them. They can then easily miss the optimal equity allocation by 10 percentage points. Judging from the divergent views on the equity risk premium (Welch, 2000) or on optimal long-run equity allocations,⁶ even greater mistakes seem possible. It is conceivable that an individual chooses 100% or 50% equity even though 75% would be optimal.

These observations motivate the second way of modelling sub-optimal choices. Before retiring, investors are assumed to miss the optimal equity allocation (again defined through the six allocation profiles from above) consistently by 10, 20, or 30

⁶ See, e.g., Thaler and Williamson (1994), and Bernstein (1996).

percentage points. After retirement, the portfolio structure is again taken to be optimal.

3 Methodology

3.1 Modelling optimal portfolio choice

Beginning with the papers of Samuelson (1969) and Merton (1969), many contributions to optimal portfolio selection in a lifetime setting have been made. In this literature, portfolio selection is analysed along with consumption and saving behaviour, labour supply, and demand for insurance.⁷

I follow the literature by rooting the portfolio problem in the modelling of optimal consumption, but do not attempt to incorporate all facets of intertemporal optimisation. Focusing on portfolio choice, I neglect some aspects to keep the analysis tractable, including mortality risk, labour flexibility, non-tradable income, and utility from bequests. The justification is not that these factors are irrelevant for portfolio choice. Labour flexibility, for example, increases an individuals' ability to bear financial risk (Bodie, Merton and Samuelson, 1992), while saving for bequests may increase the amount of money invested in financial assets. Since the focus of this paper is on optimal portfolio choice, however, I propose to take optimal decisions in related choice problems as given. The simplification will mainly affect the proportion of financial wealth to total wealth endogenously derived in the model. Obviously, this can change the costs of non-optimisation, but such effects can be gauged by varying the parameterisation of the simplified model. Furthermore, I reduce the asset allocation problem to the choice between a riskless asset and a risky one. The former can be thought of as an inflation-protected bond, the latter as a stock index.

In essence, the model used in this paper is a discrete-time variant of the lifetime consumption and portfolio choice model of Bodie, Merton and Samuelson (1992). It can be summarised as follows. An individual lives from year 0 to T. She receives

⁷ See, e.g Richard (1975), Bodie, Merton and Samuelson (1992), Jagannathan and Kocherlakota (1996) and Viceira (2001).

income from financial wealth and from labour. The present value of future labour income is the individual's human capital $H(t)$. Financial wealth can be invested in a riskless asset earning a real return of r or in an equity index offering an expected real return of a .

In each period, the individual's choice problem is to determine the optimal level of consumption $C(t)$ as well as the proportion of financial wealth invested in equity. Formally, the individual's objective is to maximise expected discounted utility from consumption:

$$E \left(\sum_{t=0}^T \frac{u(C(t))}{(1+d)^t} \right) \quad (1)$$

with u denoting the utility function and δ the rate of time preference. Consumption takes place at the beginning of the period. Maximisation is subject to the following wealth constraint:

$$C(t) = W(t) - \frac{W(t+1)}{1+r_w(t)} \quad (2)$$

where $W(t)$ denotes total wealth, which is the sum of financial wealth and human capital, and r_w is the return on total wealth. Human capital is thus assumed to be tradable. At the end of the individual's life, wealth is assumed to be zero ($W(T+1)=0$).

The return of the stock index is assumed to be log-normally distributed:

$$\ln(P(t)/P(t-1)) = m_p - \frac{s^2}{2} + s z, \quad z \sim N(0,1) \quad (3)$$

where z is a standard normal variate, s is volatility and m_p is set equal to $\ln(1+a)$ in order to render the expected simple return equal to a . As in Bodie, Merton and Samuelson, shocks to labour income $Y(t)$ are assumed to be permanent and perfectly correlated with the equity index. In such a setting, there is a closed form solution to the consumption and portfolio problem, which greatly simplifies the comparison of various strategies. While the income-equity correlation will typically not be perfect in the real world, empirical evidence suggests that it is significant. Using panel data on individual earnings, Campbell et al. (2001) estimate the correlation

between income changes and the lagged equity return. Their estimates range from 0.328 to 0.5155, depending on education. In addition, section 5 contains an analysis which suggests that the correlation assumption does not affect the results of this paper.

The perfect correlation can be achieved by making wage income depend on z , the shocks to equity prices. Formally, the logarithmic change in labour income $Y(t)$ is modelled as follows:

$$\ln(Y(t)/Y(t-1)) = \mu_Y - \frac{k(t)^2 \sigma^2}{2} + k(t) \sigma z \quad (4)$$

with the parameter $k(t)$ determining the variance of labour income in period t . The growth rate of Dollar income is equal to $\exp(\mu_Y)$, and wages are assumed to be paid at the end of the period. When calculating the present value of future income, its riskiness has to be taken into account. The correct discount rate is $r + k(t)(a - r)$, so that human capital $H(t)$ is defined as:

$$H(t) = \sum_{i=0}^T \frac{Y(t)}{[1 + r + k(t)(a - r)]^i} \quad (5)$$

As can be seen from this formula, adding idiosyncratic income risk would not affect the individual's behaviour, because it would leave discount rates and thus human capital unchanged. Rather, idiosyncratic income is treated as an implicit investment in the riskless asset.

3.2 Calibrating and solving the model

The purpose is to calibrate the model with reasonable parameter values such that various asset allocation paths can be obtained as the optimal solution to the individual's maximisation problem. Obviously, it is impossible to find a single set of parameter values generally accepted as representative. The characteristics of individuals and their environments differ considerably, and even for a well defined type of individual precise values for parameters such as income uncertainty are hard to pin down. When introducing the parameter choices, I will therefore refrain from extended justifications and discussions, but rather refer the reader to the sen-

sitivity analyses of section 5 which quantify the effects of individual parameters on the results.

As is standard in the literature, the individual's utility function is taken to be an isoelastic utility function of the form $u(C) = C^{1-\gamma} / (1-\gamma)$, with γ being the coefficient of risk aversion. This choice implies constant relative risk aversion, i.e., the individual's optimal exposure to risk is independent of her wealth. As demonstrated in Merton (1969) and Samuelson (1969) it is also independent of time. The intuition is that the optimisation problem can be solved recursively, reducing the dynamic problem to a sequence of independent one-period decisions.

In the framework chosen here, the individual cannot influence the riskiness of her labour income, which leaves the composition of financial wealth as a means of achieving the desired exposure to risk. I assume that the financial portfolio can be rebalanced at the beginning of each year. Since there is no analytic solution to the optimum risk exposure x_T , the problem is solved using numerical methods.⁸ Based on the solution for x_T , the optimum share of financial wealth invested into equity x_F is given through

$$x_F(t) = x_T + (x_T - k(t))H(t) / F(t) . \quad (6)$$

where $F(t)$ denotes financial wealth. With $k(t)$ increasing, human capital becomes more risky, and the optimum equity share goes down. As long as the individual wants to take some risk beyond that implied by human capital ($k(t)$ below x_T), equity investment increases with the ratio of human capital to financial wealth. In other words, great financial risks can be taken if financial wealth is small relative to total wealth.

The optimum overall exposure to risk also determines the expected return on the individual's wealth. Since the optimum risk exposure is constant, the expected return is constant as well. Together with the rate of time preference, the expected return on wealth determines how attractive consuming today is vis-a-vis saving for tomorrow. If both rates are equal, the increase in consumption made possible

through the positive return on savings is exactly offset by the accompanying discount in utility. In such a case, the optimal level of consumption is equal to permanent income, i.e. the annuity derived from investing current wealth at its expected return.⁹

For a given combination of risk preferences and total wealth, the solution of the intertemporal optimisation problem thus involves the following steps:

- 1) Determine x_T , the individual's optimal exposure to risk.
- 2) Determine the optimum level of consumption given today's wealth.
- 3) Determine x_F , the optimum equity allocation within the financial portfolio given the size and riskiness of today's human capital.
- 4) Reassess 2) and 3) at each decision point.

In the base case parameterisation of the model, which is used for most of the analyses, the parameters are set as follows. The optimisation is done from the age of 30 to the age of 85. The risk-free rate is set at 0.02, while the expected return on equity is 0.06 with a volatility of 0.175. Together with a risk aversion (γ) of 5 this makes the optimal risk exposure x_T equal to 25% of total wealth. From these assumptions, the expected return on total wealth is 0.03, which is also the chosen rate of time preference δ .

The individual earns labour income from the age of 30 to the age of 65. After that, financial wealth, which is arbitrarily set to 25,000 at the age of 30, is the only source of income. To model effects of mandatory public pension schemes I examine cases in which the individual has to contribute a fixed percentage of current income to a public pension fund. The fund is assumed to earn the risk-free rate of return. In addition, the individual's claims are tradable, meaning that they increase current total wealth and influence the asset allocation policy. Within the model, mandatory contributions and interest earned are recorded in a separate account

⁸ As in Jagannathan and Kocherlakota (1996), the problem is solved by approximating the standard normal distribution by a 10-point discrete distribution which has the same first 20 moments as the standard normal.

⁹ Although the planned level of consumption is constant, realised levels will differ due to unforeseen changes in wealth.

$S(t)$, which is dissolved at age 65. This necessitates a slight modification of the previous notation in (6). The rule for the optimal equity share in financial wealth becomes:

$$x_F(t) = \frac{x_T(F(t) + H(t) + S(t)) - k(t)H(t)}{F(t)} . \quad (7)$$

Mandatory savings affect the cost of non-optimisation by crowding out the accumulation of financial wealth. The costs of deviating by x percentage points from the optimum equity share in financial wealth decrease with the ratio of financial wealth to total wealth. In the base case, the mandatory contribution is set to zero.

Labour income is assumed to be 50,000 (net) at the age of 30, with a zero expected growth rate. This is consistent with the flat income profiles reported by Campbell et al. (2001) for education groups with no college degree.¹⁰ The riskiness of labour income is assumed to decrease linearly from 4% at age 30 to zero at age 65, meaning that the parameter $k(t)$ decreases from 0.229 to zero. As a result, the unexpected change in income from the age of 30 to the age of 65 has a standard deviation of 0.14; the 90% confidence interval for pre-retirement income ranges from 39,750 to 62,890. The initial income volatility matches empirical findings on systematic permanent income risk. For high school graduates, Campbell et al. (2001) estimate the annual volatility of permanent income shocks to be 0.103, and the correlation of aggregate income shocks (permanent plus transitory) with stock returns to be 0.3709. Modelling permanent shocks through a one-factor market model and assuming permanent and transitory shocks to have the same correlation with equity, these figures imply that the systematic component of labour income has a standard deviation of $0.3709 \cdot 0.103 = 0.038$, which is close to the value proposed above.

With these assumptions, human capital is 1,132,843 at the age of 30. At the expected rate of return, the planned level of consumption is 41,985 implying that the individual saves roughly 16% out of her current income in order to finance con-

¹⁰ One could also mention that liabilities (e.g. expenses for health care or children's education), which are not modelled here, typically rise with age, thus increasing the income necessary to support a given level of utility.

sumption during retirement. The evolution of human capital and financial wealth is shown in figure 2 for the case in which the realised return on equity is equal to the expected return. The diagram also contains the evolution of the optimum proportion of financial wealth invested in equity.¹¹ The decrease in the optimum equity percentage is due to the fact that the depletion of human capital and the accumulation of financial wealth boosts financial wealth relative to total wealth. Holding everything else constant, this increases the exposure to risk, which can be balanced by reducing the equity share in the financial portfolio.¹²

As in other studies (e.g. Bodie, Merton and Samuelson, 1992), derived optimum equity shares are above one early in the life-cycle. A common way of reconciling this pattern with observed behaviour is to note that many young households hold real estate financed with mortgage loans. Such a leveraged, risky investment substitutes equity investments.

This observation motivates the way in which optimum equity investments within the model are calibrated to the ones recommended by available life-cycle products. I assume that the individual has debt-financed investment in other risky assets M such that the optimum equity share in the financial portfolio is equal to a prescribed percentage $x_F^*(t)$. Assuming, for simplicity, that the risk characteristics of these other assets M are equal to equity¹³, the necessary investment in M is given by

$$M(t) = \left(x_F(t) - x_F^*(t) \right) F(t) . \quad (8)$$

Any losses and gains from this investment as well as interest paid on $M(t)$ are recorded in the separate account $S(t)$ introduced above.

¹¹ The optimum equity share is the one before expenses for current consumption.

¹² There is a countervailing effect in this model, which is the decreasing riskiness of human capital. With the parameters chosen here, however, this substitution effect is dominated by the wealth effects.

¹³ Modelling this other investment similar to the modelling of human capital, i.e. assuming a lower volatility than equity, would not change the results.

3.3 Comparing different asset allocation profiles

The effects of sub-optimal portfolio choices are analysed in the following way. Given a certain asset allocation profile A which is optimal for an investor, the aim is to quantify the utility loss stemming from choosing another profile B. For each profile, I choose identical parameters and adjust leveraged investments in other risky assets $M(t)$ such that the planned optimum life-cycle allocation within the model is equal to that recommended by A.

The different strategies are compared based on the expected utility of future consumption at the age of 30:

$$E_{30} \left(\sum_{t=31}^{85} \frac{u(C(t))}{(1+d)^{t-31}} \right) \quad (9)$$

I use Monte Carlo simulations to evaluate this expression. They have the following structure:

- 1) Beginning with the age of 30, consumption and portfolio composition are determined based on the expected return on wealth under the optimal strategy. For all allocation profiles under analysis, the leveraged investments in other assets $M(t)$ are set such that the optimum equity share in financial wealth is equal to that recommended by allocation A. If this would result in a risk exposure below 0 or above 1 for the sub-optimal strategies, $M(t)$ is capped to avoid the possibility of bankruptcy.
- 2) The return on the risky asset is drawn for the following year, affecting the value of human capital, financial assets and other assets. Gains or losses from other assets $M(t)$ are recorded in the separate account $S(t)$.
- 3) The individual revises her consumption plan; the investment in other assets $M(t)$ is adjusted such that, for each strategy, the optimum equity share in financial wealth is again equal to the one prescribed by A. Both optimum consumption and leveraged investment in $M(t)$ will differ across the various strategies when the recommended equity share differed in previous periods.
- 4) Steps 2) and 3) are repeated until the age of 85.

Steps 1) to 4) produce random consumption paths, which are translated into a utility value using the chosen utility function and the chosen rate of time preference. This exercise is repeated 20,000 times, and the resulting 20,000 utility values are averaged to obtain expected utility. To reduce simulation error, every second run of the simulations uses the negative values of the previous run's random numbers. Furthermore, I use the same set of random variables for all analyses conducted in this paper, thus reducing errors when comparing results across different parameterisations.

As a final step, I compute certainty equivalents for the obtained expected utility levels. The certainty equivalent is the risk-free consumption level which yields the same expected utility as the one obtainable under risk. If, for example, following the optimal profile A leads to a certainty equivalent of 40,000 (consumption p.a.), while the sub-optimal profile B leads to 36,000, making the sub-optimal choice is said to reduce consumption by 4,000 or 10%. In the base case, the certainty equivalent of the optimal strategy is 37,678.

4 Results

First, I use Monte Carlo simulation to find an answer to the following question: If one of the six allocation profiles presented in section 2 is optimal for an investor, what is the utility loss when she chooses to invest in one of the remaining five products instead? The analysis involves 6 times 20,000 simulations of consumption paths from age 30 to age 85. In each of the six settings, the model of section 3 is calibrated such that one particular allocation policy is optimal, while the remaining five are not.

The results of the analyses are summarised in table 1, which can be read as follows: the life-cycle product mentioned in the column header is the optimal choice, while the row headers denote the products actually chosen. Table entries are utility losses from sub-optimal choices in terms of percentage differences in certainty equivalents of consumption.

The consumption losses range from 0.01% to 1.26%. As should be expected, losses are small if the two life-cycle products under comparison are very similar.

This is, for example, the case with the DEKA and Union products, which reduce the initial equity share of 70% in the final four and three years before retirement, respectively. If one starts to reduce equity share with a delay of one year (i.e. chooses the DEKA product instead of the Union product), the utility loss is only 0.03%.

Utility losses are larger if the average equity shares of optimal and sub-optimal products differ significantly. For example, the DIT has an (unweighted) average equity share of 54%, while the DWS plan recommends a constant 70%. The loss from choosing DWS instead of DIT is 1.16%. Intriguingly, utility losses are negligible if a smooth allocation path is approximated through a stepwise change in equity allocations. This is evident from a comparison of the DIT plan and the 100-age rule. The utility loss from choosing DIT instead of following the 100-age rule is only 0.01%, even though the DIT captures the annual allocation changes proposed by the 100-age rule through only four reallocations.

How big are utility losses if the timing of risk reductions is right, but the equity allocation chosen before retirement is consistently below or above the optimal one? The answer is given in table 2. Again, the six different profiles from above are taken to be optimal. Sub-optimal choices are defined as consistently deviating by 10, 20, or 30 percentage points from the optimal allocation. Across the six profiles, utility losses are similar in magnitude. As is also evident from table 2, utility losses from underweighting or overweighting equity are largely symmetric. They amount to 1.32% - 1.58% if the chosen equity share is 30 percentage points away from the optimum. The losses drop to 0.15% - 0.19% if the departure from optimality is only 10 percentage points. Since many fund companies offer allocation products only in steps of 20 to 25 percent benchmark equity share, the maximum utility loss from such a standardisation would be less than 0.2%.

So far, I have assumed zero mandatory savings, making accumulated financial wealth the only source of retirement income. Table 3 shows the results from repeating the analysis of table 2 with mandatory saving levels changed to 5% and 10% of current income, respectively. The analysis is done for the ADIG product,

whose asset allocation profile is intermediate in terms of average equity allocation and the timing of risk reductions.

Increasing the mandatory savings rate from 0% to 5% roughly halves the losses from sub-optimal asset allocations. The higher the mandatory savings, the smaller is financial wealth, thus reducing the effects of misallocating financial wealth. At age 55, for example, the expected ratio of financial wealth to non-financial wealth drops from 95.4% to 54.2% if mandatory savings are raised from 0% to 5%. For a mandatory savings rate of 10%, utility losses amount to less than one third of the ones documented in table 2.

5 Putting the results into perspective

In the presentation of the results, two important questions were left open. How much do the obtained figures depend on the assumptions? Assuming that we can trust the figures, should the utility losses cause concern, or are they negligible?

As to the model of optimal portfolio choice, three aspects seem to merit particular attention: the choice of parameters, the introduction of an additional asset for calibrating the optimal equity allocation, and the assumed perfect correlation between labour income and equity returns.

To check the influence of individual parameters, I re-ran the analysis of the preceding section with model parameters changed. The variations from the base case are as follows (they are not accumulating and, apart from E, the rate of time preference is matched to the expected return on wealth):

A Increase both the risk-free and the equity return by 1 percentage point. This does not affect the optimum exposure to risk, but changes the level of consumption and the magnitude of financial wealth relative to human capital.

B Reduce equity volatility to 15% (from 17.5%). Holding other parameters constant, this increases the optimum exposure to risk from 25% in the base case to 34.1%.

- C Lower the risk aversion parameter γ from 5 to 3. This increases the optimum exposure to risk from 25% in the base case to 41.8%
- D Increase the risk aversion parameter γ from 5 to 7. This reduces the optimum exposure to risk from 25% in the base case to 17.8%
- E The expected growth of labour income is 1.5% rather than 0%, while the rate of time preference is set to 2% instead of 3% in the base case (at an expected return on wealth of 3%).¹⁴ The two changes have opposing effects on the incentives to save early in the life-cycle. The latter makes the individual more patient, while the former allows to increase current consumption by borrowing against future income. The overall effect is such that the base case savings rate (16%) is reached only at the age of 45.
- F Assume that income volatility stays constant at 1.5% until age 65, rather than decreasing from 4% to zero.

The sensitivity checks are applied to the mutual comparison of the six profiles (table 1). In each case, I take the ADIG product as the optimal one and determine the utility loss from following another profile under different parameter constellations. Results are presented in table 4.

In most cases, the parameter changes do not affect the derived utility loss in a substantial way. For example, estimates of the loss from choosing the DEKA product instead of the ADIG one are in five cases close to 0.26%. Changing the risk aversion coefficient from 5 to 7 (variation D) has the biggest impact, which is attributable to two causes. The higher risk aversion reduces expected return on wealth, thus inducing the individual to accumulate more capital in order to maintain consumption after retirement. Also, deviations from optimality hurt more with higher risk aversion. Accordingly, utility losses decrease with a smaller coefficient of risk aversion (variation C). In variation E, the individual saves much less in her thirties and forties. As a result, the differences between the ADIG product on the one hand, and DIT/Golden rule on the other, which are most pronounced during that

¹⁴ The solution for optimal consumption if the rate of time preference and the expected return on wealth diverge can be found in Samuelson (1969).

period, are diminished. Even such predictable changes, however, do not change the broad picture presented in section 4.¹⁵

To compare the allocation profiles recommended by mutual fund companies, the optimal equity shares had to be adjusted through the introduction of an additional asset. Moreover, this asset was assumed to be perfectly correlated with equity returns. To gauge whether my results depend on this device, I conduct simulations in which the optimum equity share derived within the base case model of section 3 is actually taken to be optimal. There is thus no need to introduce an additional asset. Through simulations, I analyse the impact of missing this optimal path consistently by 10, 20 or 30 percentage points.

In a similar fashion, I address the impact of making labour income perfectly correlated with stock income. I modify the base case model by assuming the other extreme to be true: income is riskless, so its correlation with equity is zero. This increases both human capital and the optimal equity share within financial wealth. In order to conduct a *ceteris paribus* comparison, I adjust annual labour income (down from 50,000 to 46,970) so that initial human capital is equal to that in the base case. This renders both the expected ratio of financial wealth to total wealth as well as the expected path for the equity allocation in financial wealth identical across the two parameterisations.

In both experiments, bankruptcy under sub-optimality is precluded by requiring the overall exposure to risk to lie between zero and one. The results are summarised in table 5. Apparently, both the introduction of additional assets and the assumptions about income risk do not have a significant impact on the results. There is no substantial difference between the two experiments presented in table 5, suggesting that the modelling of income risk does not effect the magnitude of utility losses. In addition, the results are close to those of table 3, in which the same degrees of sub-optimality were examined for the six asset allocations recommended by fund managers.

¹⁵ I also checked whether setting initial financial wealth to 1 (instead of 25,000) affects results. The differences to the base case are 0.01% or less.

Taking the documented utility losses to be reliable, they do not seem to be negligible even though one may argue that there are other choice problems which embody larger potential costs from erroneous decisions. Rodepeter and Winter (2000), for instance, show that applying rules of thumb to solve the intertemporal consumption-saving problem can easily reduce consumption by 10%.

To facilitate the interpretation of a given utility loss, it is useful to formulate the lifetime consumption loss in present value terms. A loss of 1.5% (0.5%) over 55 years, discounted at 2% (the real return in the base case), is equal to 51% (17%). Applying these figures to the results of table 1, an individual aged 30 should thus be willing to forego roughly 17% of current annual consumption to someone who guides him towards picking the optimal allocation out of the six profiles in question. Someone wondering whether 70% or 100% equity are optimal might sacrifice more than 50% of current consumption to get the correct answer. This supports the interpretation that utility losses are not negligible, and suggests that the advice of financial planners can be very valuable.

When choosing a fund for retirement savings, the fund's allocation should not be the only criterion to investors. Mutual funds often exhibit considerable differences in fees and other expenses. If higher fees negatively affect net fund performance (for evidence see Carhart, 1997), investors should watch out for lower fees. Among the five funds analysed in this paper, there are large differences in fees. For example, one of the funds charges an annual fee of 0.8% combined with an up-front fee of 4%, while another has a 0.5% annual fee and an up-front fee of 3%. This suggests that investors may forego more than 0.3% in annual return just through the choice of an expensive fund. In a competitive market, such differences should not persist. However, by increasing marketing expenditures, the very forces of competition could drive up average administration expenses by 0.3% or more (cf. James et al., 1999).

The utility loss associated with such administration costs can be evaluated by re-running the simulations with the rates of return changed. During the accumulation phase, I lower the annual return on financial wealth invested in the risk-free and risky asset by 0.3%. As an effect, the return on total wealth is no longer constant.

To find the optimum consumption path, I use a numerical search procedure which determines permanent consumption each year anew after the return on the risky asset has been realised. This is done for every run of the simulations. For each of the six allocation profiles, these are compared to the certainty equivalents under optimality in the base case. The difference is 1.3% on average, which is significant relative to the differences documented previously. For example, it is close to the utility loss from over- or underweighting equity by 30 percentage points.

6 Concluding remarks

Using a model of intertemporal consumption and portfolio choice, I have shown that likely departures from optimal equity allocation reduce attainable consumption by up to 1.6%. Small changes in the parameterisation do not change the results materially, even though it is obvious that the proportion of financial wealth relative to other assets affects the cost of non-optimisation. As shown in the paper, even a small mandatory savings rate can substantially reduce the costs of choosing sub-optimal allocations for financial wealth.

The documented utility losses are not negligible, stressing the importance of investment advice which takes individual preferences and circumstances into account. However, when fitting portfolio allocations to individual needs, costs and benefits should be balanced. Small deviations from the optimal allocation pattern do not cause material losses. On the other hand, they may serve to reduce administration and marketing costs incurred by the fund management companies and thereby increase attainable fund returns. An example in point would be that available life-cycle products offer to investors a stepwise reduction of equity risk even though phasing out equity smoothly will be more close to optimality in most cases.

Furthermore, getting the asset allocation right can be less important than reducing the cost of investing in funds. Assuming that observed differences in fund fees lead to equally sized differences in expected fund performance, choosing expensive funds may do as much harm as missing the optimal equity allocation by 20 to 30 percentage points. Together with the documented small costs of standardisation, the findings also support James et al. (1999) who argue that the benefits of con-

straining choice in mandatory funded security systems will outweigh the associated costs. James et al. estimate the saving potential to be up to 1% per annum; in the framework of this paper, saving 0.3% per annum creates utility gains which exceed the maximum loss from being forced to pick one out of six allocation profiles recommended by investment companies.

References

- Bernstein, P., 1996, Are stocks the best place to be in the long run? A contrary opinion. *Journal of Investing* 5, 9-12.
- Bodie, Z., and D. Crane, 1997, Personal investing: Advice, theory, and evidence. *Financial Analysts Journal* 53, 13-23.
- Bodie, Z., R. Merton and P. Samuelson, 1992, Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control* 16, 427-449.
- Campbell, J., J. Cocco, F. Gomes, and P. Maenhout (2001): Investing retirement wealth: a life-cycle model. Campbell, J. and M. Feldstein (Eds), *Risk Aspects of Social Security Reform*. Chicago.
- Canner, N., G. Mankiw, and D. Weil, 1997, An asset allocation puzzle. *American Economic Review* 87, 181-191.
- Carhart, M., 1997, On persistence in mutual fund performance. *Journal of Finance* 52, 57-82.
- Goodfellow, G. and S. Schieber, 1997, Investment of assets in self directed pension portfolios. Gordon, M., O. Mitchell and M. Twinney (Eds), *Positioning pensions for the twenty-first century*. Philadelphia, 67-90.
- Jagannathan, R. and N.R. Kocherlakota, 1996, Why should older people invest less in stocks than younger people? *Federal Reserve Bank of Minneapolis Quarterly Review* 20, 11-23.
- James, E., G. Ferrier, J. Smalhout, and D. Vitas, 1999, Mutual funds and institutional investments: what is the most efficient way to set up individual accounts in a social security system. NBER working paper.
- Jensen, B. and C. Sørensen, 2000, Paying for minimum interest rate guarantees: Who should compensate who? Working paper, Copenhagen Business School.
- Merton, R.C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics* 51, 247-257.

Richard, S.F., 1975, Optimal consumption, portfolio and life insurance rules for an uncertain lived individual in a continuous time model. *Journal of Financial Economics* 2, 187-203.

Rodepeter, R. and J. Winter, 2000, Rules of thumb in life-cycle savings models. Working paper, University of Mannheim.

Samuelson, P., 1969, Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics* 51, 239-246.

Sirri, E. and P. Tufano, 1998, Costly Search and Mutual Fund Flows. *Journal of Finance* 53, 1589-1622.

Thaler, R. and P. Williamson, 1994, College and university endowment funds: Why not 100% equities? *Journal of Portfolio Management* 21, 27-37.

Viceira, L., 2001, Optimal portfolio choice for long-horizon investors with nontradable income. *Journal of Finance*, forthcoming.

Welch, I., 2000, Views of financial economists on the equity premium and on professional controversies, Working Paper, UCLA.

Table 1

Utility losses from following sub-optimal lifetime asset allocations (in percent of certainty equivalents of consumption)

Chosen product	Optimal life-cycle product					
	ADIG	DEKA	DIT	DWS	Union	100-age
ADIG	0.00	-0.26	-0.33	-0.57	-0.15	-0.31
DEKA	-0.26	0.00	-0.73	-0.23	-0.03	-0.75
DIT	-0.45	-0.91	0.00	-1.24	-0.80	-0.01
DWS	-0.86	-0.36	-1.16	0.00	-0.47	-1.26
Union	-0.16	-0.02	-0.64	-0.30	0.00	-0.63
100 - age	-0.43	-0.92	-0.01	-1.24	-0.80	0.00

Key assumptions: The individual lives from age 30 to 85, receiving risky labour income from age 30 to 65. There are no mandatory savings. Based on an isoelastic utility function, the individual determines consumption and chooses between a safe asset (0.02 return) and a risky one (0.06 return with 0.175 volatility). The rate of time preference is 0.03. The products' asset allocation path is depicted in figure 1. Additional details are given in the text.

Table 2

Utility losses from following sub-optimal lifetime asset allocations (in percent of certainty equivalents of consumption)

Deviation from optimum equity share (percentage points)	Optimal life-cycle product					
	ADIG	DEKA	DIT	DWS	Union	100-Age
+ 30	-1.33	-1.32	-1.39	-1.32	-1.32	-1.39
+ 20	-0.60	-0.60	-0.61	-0.59	-0.60	-0.61
+ 10	-0.15	-0.15	-0.15	-0.15	-0.15	-0.15
- 10	-0.19	-0.19	-0.18	-0.19	-0.19	-0.18
- 20	-0.74	-0.75	-0.69	-0.76	-0.75	-0.69
- 30	-1.58	-1.58	-1.50	-1.58	-1.58	-1.50

Key assumptions: The individual lives from age 30 to 85, receiving risky labour income from age 30 to 65. There are no mandatory savings. Based on an isoelastic utility function, the individual determines consumption and chooses between a safe asset (0.02 return) and a risky one (0.06 return with 0.175 volatility). The rate of time preference is 0.03. The products' asset allocation path is depicted in figure 1. Additional details are given in the text.

Table 3

Utility losses from following sub-optimal lifetime asset allocations (in percent of certainty equivalents of consumption) for three levels of mandatory savings

Deviation from optimum equity share (percentage points)	Optimal life-cycle product is ADIG		
	No mandatory savings (base case)	5% mand. savings rate	10% mand. savings rate
+ 30	-1.33	-0.74	-0.32
+ 20	-0.60	-0.33	-0.14
+ 10	-0.15	-0.08	-0.03
- 10	-0.19	-0.10	-0.05
- 20	-0.74	-0.40	-0.17
- 30	-1.58	-0.89	-0.37

Key assumptions: The individual lives from age 30 to 85, receiving risky labour income from age 30 to 65.

Based on an isoelastic utility function, the individual determines consumption and chooses between a safe asset (0.02 return) and a risky one (0.06 return with 0.175 volatility). The rate of time preference is 0.03. The products' asset allocation path is depicted in figure 1. Additional details are given in the text.

Table 4

Utility losses from following sub-optimal lifetime asset allocations (in percent of certainty equivalents of consumption) under various parameterisations

Chosen product	Parameterisation (optimal product is ADIG)						
	Base case	A	B	C	D	E	F
DEKA	-0.26	-0.22	-0.21	-0.19	-0.45	-0.22	-0.28
DIT	-0.45	-0.35	-0.30	-0.23	-0.88	-0.25	-0.48
DWS	-0.86	-0.71	-0.69	-0.58	-1.50	-0.79	-0.93
Union	-0.16	-0.14	-0.12	-0.12	-0.28	-0.13	-0.17
100-Age	-0.43	-0.33	-0.28	-0.21	-0.86	-0.23	-0.45

The variations of the base case parameterisation are as follows:

A risk-free and equity return plus 1% vs. base case B equity volatility 15% instead of 17.5%

C risk aversion coefficient 3 instead of 5 D risk aversion coefficient 7 instead of 5

E wage growth 1.5%, rate of time preference = 2%

F riskiness of human capital constant at 1.5% instead of decreasing from 4% to zero

Table 5

Utility losses from following sub-optimal lifetime asset allocations (in percent of certainty equivalents of consumption)

Deviation from optimum equity share (percentage points)	Optimum derived in life-cycle model with stochastic income (Base case)	Optimum derived in life-cycle model with riskless income
+ 30	-1.46	-1.53
+ 20	-0.65	-0.70
+ 10	-0.16	-0.19
- 10	-0.19	-0.16
- 20	-0.75	-0.70
- 30	-1.61	-1.57

Key assumptions: The individual lives from age 30 to 85, receiving risky labour income from age 30 to 65. Based on an isoelastic utility function, the individual determines consumption and chooses between a safe asset (0.02 return) and a risky one (0.06 return with 0.175 volatility). The rate of time preference is 0.03. The optimum asset allocation path is depicted in figure 2. Additional details are given in the text. For the model with riskless income, annual income is reduced to match the initial human capital of the base case model.

Figure 1
 Lifetime equity allocation of five German mutual fund products and a popular rule of thumb.

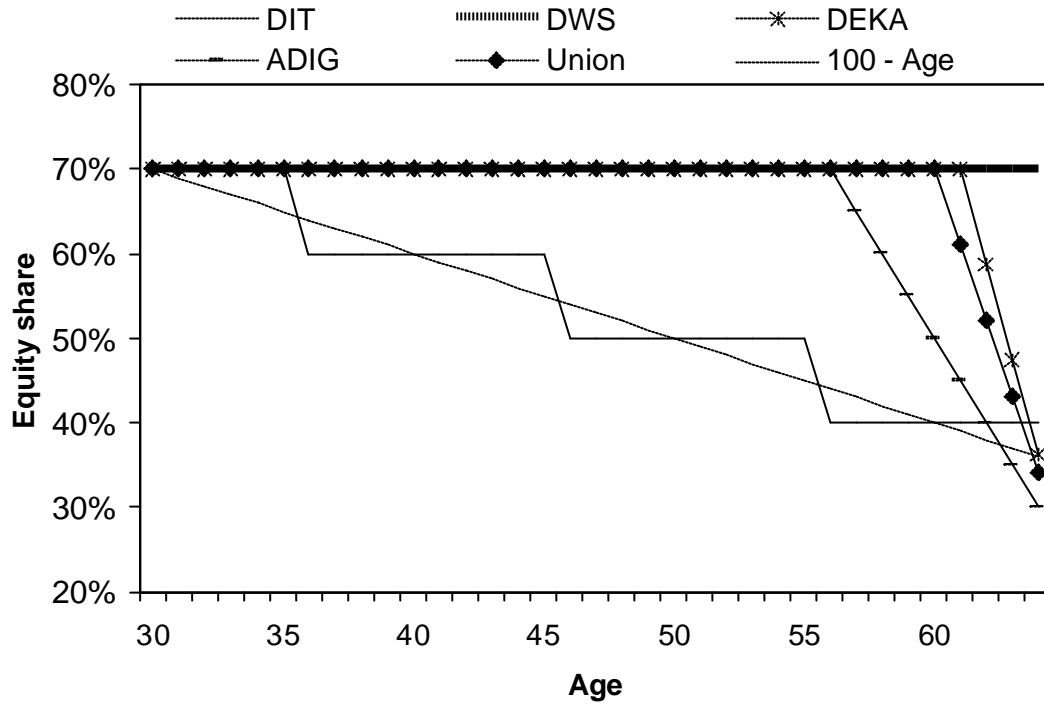


Figure 2
 Expected human capital, financial wealth and optimum equity investment in the base case

