Costs of Suboptimal Investments ☆

Paulo Rodrigues
Finance Department, Maastricht University, Maastricht, Netherlands, 6211 LM, (p.rodrigues@maastrichtuniversity.nl)

Peter Schotman
Finance Department, Maastricht University, Maastricht, Netherlands, 6211 LM, (p.schotman@maastrichtuniversity.nl)

Abstract Risk aversion is the primary factor in determining an optimal life-cycle portfolio strategy for allocating investment risk across age cohorts within a pension fund. Since participants have varying levels of risk aversion, a pension fund must decide how to address the heterogeneity in preferences. Furthermore, the human capital of young participants enables them to take more risks, as poor returns can still be offset by future labor income and investment returns.

Welfare costs associated with sub-optimal strategies can be substantial. The optimal strategy for a low-risk-aversion individual imposes significant welfare costs on someone with high-risk-aversion. The loss represents 97% in terms of certainty-equivalent real benefits. Welfare losses are not symmetrical. Low-risk-aversion individuals experience significantly less welfare loss when the fund follows the life-cycle policy optimal for a high-risk-aversion participant. The lowest maximum loss occurs when the fund adheres to a moderate policy suitable for medium-risk-aversion individuals.

Key Words Pension Funds, Defined Contribution, Collective Investment Strategy, Lifecycle Model

JEL Classifications: G11, G29

☆ We greatly benefited from constructive suggestions during our presentation at a Netspar after-lunch webinar.

Costs of Suboptimal Investments

Abstract Risk aversion is the primary factor in determining an optimal life-cycle portfolio strategy for allocating investment risk across age cohorts within a pension fund. Since participants have varying levels of risk aversion, a pension fund must decide how to address the heterogeneity in preferences. Furthermore, the human capital of young participants enables them to take more risks, as poor returns can still be offset by future labor income and investment returns.

Welfare costs associated with sub-optimal strategies can be substantial. The optimal strategy for a low-risk-aversion individual imposes significant welfare costs on someone with high-risk-aversion. The loss represents 97% in terms of certainty-equivalent real benefits. Welfare losses are not symmetrical. Low-risk-aversion individuals experience significantly less welfare loss when the fund follows the life-cycle policy optimal for a high-risk-aversion participant. The lowest maximum loss occurs when the fund adheres to a moderate policy suitable for medium-risk-aversion individuals.

Key Words Pension Funds, Defined Contribution, Collective Investment Strategy, Lifecycle Model

JEL Classifications: G11, G29

1 Introduction

Pension funds invest on behalf of their participants. In a stylized model of their investment strategy, they adhere to the life-cycle recommendation, suggesting that young people initially have a high exposure to risky assets that gradually decreases with age. As people age, the importance of a hedge portfolio that protects against interest rate risk increases. A coefficient of risk aversion dictates how the proportions of risky investments and the hedge portfolio change with age.

Optimizing the investment decisions over the life-cycle using only age as a state variable creates an interesting methodological problem in a model where the investment opportunities vary over time. We show that an algorithm that relies on backward induction generates suboptimal solutions. This is because the decisions are not optimal conditional on the relevant states of the world. We, therefore, use an optimization algorithm that searches for the optimal portfolio strategy over the entire lifecycle and does not rely on backward induction.

As the investment strategy is collective and the risky investment share only varies with age, the pension participants are exposed to the same investment plan. The pension fund, therefore, needs to decide for which risk aversion level it optimizes the strategy. The goal of our analysis is to estimate the welfare costs for participants whose risk aversion differs from the one underlying the decisions of the pension fund. For example, what happens to the certainty equivalent outcomes of pension participants with a relative risk aversion coefficient of 10 if the fund uses the strategy that is optimal for participants with a coefficient of 2?

The importance of the decision on how much risk to take in the collective portfolio has increased considerably in the new Dutch pension contract. The Netherlands switched the second pillar to a fully defined contribution system. In this new setup, the pension fund determines the pension wealth of each participant by allocating a portion of the collective investment portfolio to the participant. At the time of retirement, the pension benefits are a function of the wealth accumulated by each participant. As this value is a function of the collective portfolio return, knowing the risk aversion of participants is of first-order importance to the pension fund.

Therefore, measuring participants' risk attitudes is a mandatory element in the new Dutch pension contract. Many studies have analyzed and compared methods for measuring risk aversion (Goossens et al., 2023). Using various risk aversion measurements the pension fund will obtain a distribution of risk aversion among all participants. When the pension fund differentiates risk allocation solely by age, assuming uniform risk aversion among all participants, many (if not all) may experience a suboptimal risk allocation. The first step in our analysis is to estimate these costs based on a given risk aversion. These costs compare the certainty equivalent of the optimal policy with that of the fund's actual policy.

The paper closest to us is the study by Balter et al. (2024b). They determine the optimal policy for a fund that considers the diverse preferences of its participants. In their stylized setting, the fund only chooses between equity and cash (the risk-free asset) with constant investment opportunities. For a specific level of relative risk aversion, the optimal policy entails the so-called 'Merton' weights, indicating that the fund allocates a constant fraction of total wealth to equity, determined solely by age. The fraction is inversely related to the level of constant relative risk aversion, a concept that pension funds must measure among their participants. They caution that funds must be careful not to underestimate the level of risk aversion. Establishing limits on the maximum exposure to equity risk can help to prevent severe losses for more risk averse participants. Balter and Schweizer (2024) further argue that a pension fund that is concerned about inequality between participants may tilt the collective portfolio toward the more risk averse participants.

We extend the cost calculations in Balter et al. (2024b) and certainty equivalent calculations in Balter and Schweizer (2024) in various ways. First, we incorporate long-term nominal bonds into the asset menu. Long-term bonds are essential in the investment portfolio of pension funds. In the new pension contract, benefits are tied to the actual yield curve. Funds can mitigate the volatility of benefits by investing in a portfolio of bonds with matching maturities. Second, we include more realistic dynamics for returns by taking into account inflation and interest rate risk. Retirees are concerned about the real purchasing power of their pension income, whereas the pension fund mainly invests in nominal assets. Covariances between inflation and the various asset classes will thus become important. But instead of constructing our own model for returns and inflation, we take the scenarios provided by the Dutch Central Bank (DNB) as the economic environment.

Third, we take into account the full life-cycle, including the retirement stage. Participants continue to face investment risk during the decumulation phase and remain susceptible to inflation and interest rate risk. With the extended menu and real ambition, we revisit inflation hedging in the new pension contract studied in van Gastel et al. (2024). Long-term nominal bonds can hedge interest rate risk in retirement benefits but they offer no protection against inflation risk. The short-term cash asset has a nominal risk-free interest rate as its return. The positive correlation between inflation and short-term interest rates offers some protection against inflation risk. We do not consider Inflation Linked Bonds. They would be the ideal instrument for hedging both interest rate and inflation risk if they were available in sufficient volumes, range of maturities and based on the relevant inflation measure.

Moving beyond the stylized model with log-normal returns and constant investment opportunities presents methodological challenges. As the investment opportunities — in the DNB scenarios interest rates, inflation, and volatility — change over time, so does the optimal portfolio. Instead of a static life-cycle profile, the optimal portfolio will be strategic in the sense of Campbell and Viceira (2002). The strategic portfolio weights, however, tend to vary wildly over time and across different economic scenarios. Moreover, these solutions are often non-robust, not only for the equity allocation (Diris et al., 2015), but even more when it comes to the nominal bond exposure under inflation risk (Balter et al., 2021a). Since pension funds appear to avoid such volatile portfolio strategies, our analysis does not consider strategic portfolio allocations, but instead focuses on a fixed life-cycle profile. We assume that the fund allocates exposures to various asset classes solely based on age, rather than on specific levels of interest rates and inflation. It does take into account the time-varying nature of returns in determining the optimal exposure, but assigns weights that maximize the average utility across possible states of the world. We are therefore presenting a representative life-cycle strategy that, while not optimal in each state of the world, is best on average for the available set of scenarios. Another methodological issue is model uncertainty. Important parameters such as the average return on equity and the persistence of inflation and interest rate shocks, are notoriously hard to estimate, which is one of the main reasons why optimized policies lack robustness. Our approach to mitigating this uncertainty involves mixing scenarios from different vintages of the DNB scenarios since 2022. This approach allows us to generate a large number of initial conditions, even though all scenarios still use the underlying model parameters. A second way to mitigating the effect of model uncertainty is to split the scenarios into one set for optimizing the portfolios and a separate validation set for monitoring their performance.

The rest of the paper begins with a model of the pension design. Next we will discuss some of the more salient properties of the DNB scenarios. For the main results we distinguish between preferences defined on either the second pillar income or on the combined benefits of the first and second pillars. The paper concludes with a discussion of limitations and methodological issues along with a summary of the main findings.

2 Life-cycle Portfolios in a Pension Fund

We consider a pension fund that invests on behalf of its participants. It implements a life-cycle portfolio policy, in which the allocation to equity, bonds, and cash is determined solely by age distribution of the participants. The fund employs a bottom-up strategy, meaning that the collective portfolio is the sum of the individual portfolios. Since the allocations are based solely on age, the selected portfolio allocation is unlikely to be fully optimal. Below, we state the portfolio problem for an individual participant in the pension fund. We focus exclusively on optimal portfolio strategies and do not consider individual decisions regarding savings, labor supply, and the timing of retirement.

2.1 Pension Scheme and Portfolio Problem

An individual pays nominal contributions A_t per year. The contributions are invested and earn a gross annual return of $R_{p,t+1}$. During the accumulation stage nominal pension wealth (F_t) grows as

$$F_{t+1} = (F_t + A_t)R_{p,t+1}. \tag{1}$$

Initial wealth at age 25 is $F_0 = 0$. The individual works $\tau = 42$ years, from age 25 to 67. The real income of an individual is deterministically fixed at h_t at age t. The percentage pension premium is also fixed such that $a_t = \kappa h_t$ is the real contribution. With a price index Π_t this gives the nominal premium $A_t = \Pi_t a_t$. Since the DNB scenarios do not have separate wage and price indices, we assume they are equal.

The gross return $R_{p,t+1}$ is defined as

$$R_{p,t+1} = x_{Ct}R_{C,t} + x_{St}R_{S,t+1} + x_{Bt}R_{B,t+1}, \tag{2}$$

where $R_{C,t}$ is the nominal return on cash, for which we take the 1-year nominal interest rate; $R_{S,t+1}$ is the return on equity; and $R_{B,t+1}$ is the return on a bond portfolio. The bond portfolio is represented by a long-term bond with a fixed duration of 10 years. The most convenient 10-year duration bond is the discount bond with 10-year maturity.[1] All returns are gross nominal returns.[2]

The portfolio weights $x_t = (x_{Ct}, x_{St}, x_{Bt})$ must sum to one. Moreover, x_{St} and x_{Bt} are non-negative and $x_{St} + x_{Bt} \leq 1.5$. In words: we do not allow short-selling of risky assets, but we do allow leverage up to 150%. The latter is relevant to allow young participants to have an equity exposure that can be 150% of their pension wealth. The leverage is motivated by the human capital of young participants. Because we allow leverage, we must also allow the fund to be short in the cash market or invest in a swap contract. The investment profiles are functions of age only.

The retirement date is fixed at time $\tau = 42$, *i.e.* 42 years of working life. Benefits during retirement follow the formulas for the new Dutch pension contract and are equal to

$$B_t = \alpha_t F_t \tag{3}$$

Here B_t are nominal benefits, F_t denotes the personal pension wealth at the beginning of year t, and α_t is an annuity factor that allows for the financing of a constant nominal pension through investments in matched-maturity nominal bonds. The fund does not need to invest in these matched-maturity nominal bonds.[3] Actual benefits can therefore fluctuate over time. The retirement income B_t is the second pillar pension income. Additionally, a retiree receives income from the first pillar. In the stylized model, an individual consumes all retirement income each year and does not possess private savings.

The annuity factor is based on assumed interest rates (AIR) equal to the current term structure, $\text{AIR}_t^{(n)} = Y_t^{(n)}$ for a payment at time $t+n$. In our model the AIRs are not adjusted for the investment portfolio's risk exposure. When investments have a higher expected return than the riskfree term structure, the projected nominal benefits will increase over time. Real benefits may increase, remain flat or decrease depending on expected inflation. In the special case where expected inflation equals the risk premium of the investment portfolio, the projected real benefits will remain constant over time. Given the AIR's the annuity factor α_t follows as

$$\alpha_t = \left(\sum_{s=\tau}^{T} q_{s-\tau} \left(\text{AIR}_t^{(s-t)} \right)^{t-s} \right)^{-1}, \tag{4}$$

with q_n the probability of surviving n periods after retirement (with $q_0 = 1$), and $T = 75$ implying a maximum age of 100. Longevity risk is covered by the fund. If fund returns were equal to the AIR, the individual could maintain a flat income independent of age. Actual retirement income will

1 The annual log return of the 10-year bond is $r_{B,t+1} = -9y_{t+1}^{(9)} + 10y_t^{(10)}$, where $y_t^{(n)} = \ln Y_t^{(n)}$ is the continuously compounded yield on a discount bond with n-year maturity. These yields follow directly from the data in the DNB scenarios.

2 By 'gross' we mean that $R = \mathrm{e}^r$, with r the continuously compounded logarithmic rate of return. In (2), $R_{C,t} = Y_t^{(1)}$ is thus a gross return, which can be read as one plus an arithmetic rate. Similarly, for the bond return we have $R_B = \mathrm{e}^{r_B}$.

3 Since we don't include the complete set of zero-coupon bonds in our asset menu, the fund cannot fully replicate the hedge portfolio that would guarantee fixed nominal benefits. In the model it can get close by duration hedging using the 1-year and 10-year bonds.

change year by year, based on the actual returns of the fund and changes in the annuity factor (due to changes in interest rates). During retirement, pension wealth, at the level of an age cohort in the fund, evolves as

$$\frac{F_{t+1}}{F_t} = (1 - \alpha_t q_{t-\tau}) R_{p,t+1}. \tag{5}$$

Here, the survival probabilities must be understood as the fraction of participants of an age cohort that are still alive.

2.2 Expected Utility

Similar to Grebenchtchikova et al. (2017), we evaluate all portfolio strategies using the expected utility during retirement,

$$U = \sum_{t=\tau}^{T} \beta^{t-\tau} q_{t-\tau}\, \mathrm{E}[u(b_t)], \tag{6}$$

where $b_t = B_t/\Pi_t$ are real benefits and β is a time preference parameter. We take the CRRA utility function, expressed as $u(b) = b^{1-\gamma}/(1-\gamma)$, featuring different values for the risk aversion parameter γ. We set $\beta = 1$ and will omit it from the utility specification going forward. Although the benefits are nominal, we assume that individuals assess the real purchasing power of their benefits. Total benefits are b_t plus AOW, with AOW denoting the real first pillar benefits.[4] We will discuss two variations. First, we focus solely on the second pillar. In that case the first pillar AOW serves as a real subsistence level S, and we therefore subtract it inside the utility function, such that we have $u\left((b_t + \text{AOW}) - S\right) = u(b_t)$. Second we consider the CRRA specification on total, first and second, benefits and evaluate $u\left(b_t + \text{AOW}\right)$ without a subsistence level. For ease of interpretation we convert the expected utility into a certainty equivalent (CE). The CE is the inverse of the expected utility reflecting a constant level of *real* benefits during retirement that is equivalent to the expected utility derived from the implemented life-cycle portfolio. In formulas, the certainty equivalent is the level b_{CE} that solves

$$U = u(b_{\text{CE}}) \times \left(\sum_{t=\tau}^{T} q_{t-\tau} \right). \tag{7}$$

Higher values of γ correspond to lower values of the certainty equivalent b_{CE} for a given stream of benefits b_t.

Since the expectation in (6) is intractable, we evaluate it by simulation. We assume a flexible functional form $\mathbf{f}(t|\theta)$ of age, given a set of parameters θ to describe the evolution of the portfolio weights $x_t = (x_{Ct}, x_{Bt}, x_{St})$ over time. As a functional form we choose a neural network (see Appendix D). We consider N paths with portfolio returns $R_{p,i,t+1}$ in scenario i over the life cycle, leading to scenarios F_{it} for wealth and b_{it} for real benefits during retirement. The empirical average utility over the N scenarios is

$$\hat{U} = \sum_{t=\tau}^{T} q_{t-\tau} \left(\frac{1}{N} \sum_{i=1}^{N} u(b_{it}) \right) \tag{8}$$

In the optimisation we search for the parameter vector θ to find the function $x_t = \mathbf{f}(t|\theta)$ that maximises the average utility. The optimizer only requires simulated scenarios. It is unaware of the underlying

4 AOW stands for "Algemene Ouderdomswet", the first pillar state pension in the Netherlands.

model used to generate the scenarios. To prevent local optima and overfitting, we optimize using a random subsample of the scenarios with random starting values for the parameters, and we evaluate the performance of the optimized portfolio weights on a validation sample of the remaining scenarios. We repeat this optimization ten times and save the result that provides the best performance across the full sample of all scenarios.

Leverage can result in negative wealth within a discrete-time simulation model. With 150% of financial wealth in equity, there is a positive probability that next year's return could fall below -100%, effectively wiping out all existing pension wealth (for example when the risk-free rate is zero and the equity return is -67%). Due to the substantial growth in labor income in the early years, we may encounter scenarios where individuals experience negative wealth at a young age, but eventually end up with positive wealth at retirement. Strong labor income may not always be enough when returns are really poor. The problem will not arise in a Black-Scholes-Merton framework, where continuous-time dynamic trading strategies can prevent negative wealth. Similarly, when equity returns follow a log-normal distribution, the likelihood of such an event is negligible. Many authors therefore ignore the issue. However, such scenarios do occur when N is sufficiently large. Since the CRRA utility function is undefined for negative wealth, the optimization algorithm will avoid such scenarios and refrain from using leverage. To enable leverage, we add a put option to the portfolio, ensuring that F_t remains positive at all times. The cost of this put option, which is very small, is deducted from the returns. Option risk and valuation details are in Appendix A.

2.3 Suboptimal Life-cycle Investment Strategies

Optimizing the portfolio weights for different values of γ will lead to alternative optimal life-cycles. The cost of a suboptimal strategy is defined as the percentage loss in certainty equivalent. If $b^*_{\mathrm{CE},\gamma}$ is the certainty equivalent of the optimal strategy for some γ, and b_γ the certainty equivalent evaluated with the same γ of some other life-cycle strategy, then the costs are defined as

$$c = 1 - \frac{b_\gamma}{b^*_{\mathrm{CE},\gamma}} \tag{9}$$

We consider various alternative strategies. First, if we optimize for some value γ_1, what is the cost of this strategy for participants with a different γ_2?

As a second benchmark, we choose a linear life-cycle portfolio commonly found in traded life-cycle products. We have taken the equity weights in the study provided by the Netherlands Bureau for Economic Policy Analysis (CPB) on different variations of the new pension contract (Metselaar et al., 2020). Equity decreases linearly from 150% at age 20 to 35% at age 71, and stays constant at that level as age continues to increase. The CPB study also includes long-term nominal bonds with a weightthat increases linearly from 25% to 50%. We deviate from the CPB bond portfolio; their implementation has a bond portfolio with age-dependent duration, whereas we keep the duration fixed at ten years.[5]

5 An age-dependent duration, depending on the duration of the projected retirement benefits, generates a very strong interest rate exposure for young participants. In our optimizer we don't see such strong interest rate exposure at a young age, even if we drop the restrictions on the bond holdings.

By combining the 10-year bond and the 1-year cash position, the fund can still obtain any desired duration. This linear life-cycle portfolio is not fully optimal for any γ, but it effectively captures the main idea of life-cycle investing.

Our third benchmark considers the so-called 'Merton' fractions.[6] In the Merton model investment opportunities do not change over time. A life-cycle pattern is then solely due to the gradual decline of human capital. The standard Merton rule for life-cycle investing states that the optimal allocation to risky assets is a constant fraction of total wealth, where total wealth is the sum of human and financial capital. The solution is optimal under the classic assumptions of real riskfree labour income, the existence of a real risk-free rate, constant investment opportunities, and constant relative risk aversion. Since the pension fund cannot invest in a true real risk-free financial asset, the Merton model does not strictly apply. Human capital is the only risk-free asset. Cash is only nominally risk-free, but subject to inflation risk. The Merton fractions therefore serves as a rough benchmark. Detailed formulas are provided in Appendix B. As they only depend on the unconditional expected returns and covariances of the ex-post real returns, they are easy to compute and thus remain a relevant benchmark. Almost all other cases require numerical solutions, which can become computationally intensive.

3 Data

As our data, we use the so-called CP2022 uniform scenario sets published by De Nederlandsche Bank on their website for Supervision of Pension Funds. These scenarios are generated in accordance with the methodological recommendations of the *Commissie Parameters 2022*. Each set contains scenarios that start from a fixed initial condition and span a 100-year horizon. For example, the scenario set `cp2022-p-scenarioset-100k-2024q2` is based on the yield curve and inflation expectations that were in effect in the second quarter of 2024. If we would train our model on the scenarios in this set, starting at the initial condition defined by inflation and the yield curve as of 2024Q2, the solution functions would also depend on these specific initial conditions. Our goal, however, is to establish a generalizable portfolio rule, that provides the optimal portfolio as a function of age without conditioning on the particular state of the world prevailing at the starting date. Our aim is to determine and assess life-cycle portfolios that are best across a wide range of initial conditions. For long-term pension policies, one would hope that advice on the optimal investment life-cycle does not rely on the initial conditions in one particular quarter.

Assuming an individual works for 42 years and enjoys a maximum of 34 years in retirement, we need only 75 years of data for each scenario. For this reason we drop the first 25 years of the scenario set, and only simulate portfolio returns and pension outcomes for the 75 years starting in year 25. For further randomization, we also combine the scenarios from the DNB sets spanning 2023Q4 to 2024Q3. Each of these four scenario sets contains 100,000 scenarios, resulting in a total of 400,000 scenarios

6 The quotation marks around 'Merton' indicate that it is a bit of a misnomer to label the constant investment opportunities result as the Merton model, since Merton (1973) analysed intertemporal models with time-varying investment opportunities. We now refer to the special case without any time-varying risk-return trade-off as the Merton model.

Table 1: Return summary statistics

	avg	stdev	skew	kurt	1%	med	99%
equity	7.08	18.03	-0.35	0.95	-42.0	8.60	47.9
equity (log)	5.26	18.38	-1.23	3.52	-54.4	8.25	39.1
equity(excess)	5.86	17.99	-0.45	0.94	-44.0	7.65	45.1
10-y bonds	2.66	9.37	0.06	0.78	-20.2	2.60	26.0
10y bonds (log)	2.21	9.22	-0.33	1.05	-22.6	2.56	23.1
10-y bonds (excess)	1.45	8.63	-0.18	0.80	-21.0	1.66	21.6
1-y rate	1.21	2.00	0.91	2.15	-2.54	0.99	7.25
inflation	2.01	1.59	1.02	2.65	-0.97	1.81	6.95
1-y ex-post real rate	-0.80	2.30	0.41	1.43	-5.99	-0.92	5.55

Annual stock returns and inflation aresourced directly from the DNB scenarios for 2023Q4 to 2024Q3. The 1-year nominal interest rate and 10-year nominal bond returns are derived from yield curves constructed using the parameters and state variables in the DNB scenarios. All returns are expressed as arithmetic percentage per year. Excess returns are measured relative to the 1-year interest rate. The ex-post real rate is the difference between the nominal rate and inflation. Skewness and kurtosis are defined as $\mathrm{E}[(x-m)^3]/s^3$ and $\mathrm{E}[(x-m)^4]/s^4 - 3$, respectively, where m is the mean and s is the standard deviation. The last three columns report quantiles, with 'med' indicating the median. The total number of observations is 3 million (4 sets, each 75 years and 100,000 scenarios).

with random starting values at year 25. By randomizing over different scenario sets, we ensure that our optimized policies are robust to incidental peculiarities in the scenarios. After 25 years state variables and yield curves no longer depend on the initial conditions of that specific quarter. We view the resulting optimized portfolios as *unconditional* life cycle strategies.

Table 1 provides summary statistics of the scenarios. These are unconditional sample statistics that pool all time periods (after year 25) over all 400k scenarios. The arithmetic average return on equity is 7%, which aligns well with the 5.3% geometric mean recommended by the *Commissie Parameters 2022*. The geometric expected return on equity is fixed across all years and scenarios, implying that the expected excess returns vary with the 1-year nominal interest rate. The bond premium is less than 1.5%. Comparing the Sharpe ratios of equity (0.33) and bonds (0.17), equity is a far more attractive investment. The correlation between the *excess* returns of the two asset classes is 0.49. Unlike previous DNB scenarios, both equity and nominal bonds exhibit non-normal log returns due to stochastic volatility in the model. Returns less than minus 42%, as in the financial crisis, occur approximately once every 100 years. The 'fat-tailedness' of equity will increase the risk associated with equity as severe negative returns are very costly, especially for investors with high risk aversion. The average real interest is negative over the 75-year planning period. Furthermore, the real rate is more variable than both the nominal rate and inflation. For the Merton fractions (see Section 2.3 above and Appendix B) we can distill the inflation hedging potential of equity and nominal bonds using the regression equation ($R^2 = 0.21$),

$$\pi_{i,t+1} - R_{C,it} = 0.10 + 0.004(R_{S,i,t+1} - R_{C,it}) - 0.127(R_{B,i,t+1} - R_{C,it}) + \epsilon_{i,t+1}$$

Since the regression is based on three million observations, the standard errors of the parameter estimates are nearly zero and meaningless. The regression shows that equity is almost unrelated to inflation, whereas nominal bonds carry a negative weight in an inflation hedge portfolio. The largest

Figure 1: Term structure of risk

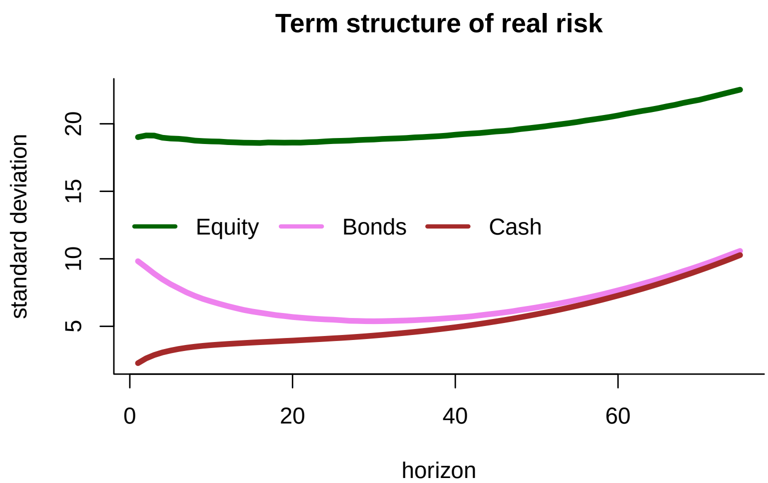

The figure shows the term structure of risk as defined in Campbell and Viceira (2005). The vertical axis is the percentage annualized standard deviation of the cumulative real returns over the horizon on the x-axis. All returns are logarithmic in excess of inflation.

weight is for cash, calculated as $1 - 0.004 + 0.127 = 1.123$, even though cash has a negative expected real return. Moreover, most of the variation in inflation (or the real interest rate) is unrelated to the excess returns of the other asset classes. Therefore, inflation hedging with these instruments and scenarios will be quite limited, regardless of the risk aversion parameter.

Using the regression coefficients and the summary statistics for equity and nominal bonds, we find the optimal mean-variance portfolio for investors with varying γ as shown below:

Mean-variance weights

Risk aversion (γ)	2	5	10
Equity	90%	36%	18%
Nominal Bonds	0	0	0

For $\gamma = 5$, the implied unconditional equity weight in retirement, when human capital is zero, would be 36% in these DNB scenarios. That number is very close to the 35% equity weight in the CPB benchmark portfolio. Nominal bonds obtain a zero weight.[7]

As a final element of the descriptive statistics we analyze the dynamic properties of the returns. A standard tool for this is the term structure of risk (Campbell and Viceira, 2005), which is defined as the annualized standard deviation of cumulative returns over different investment horizons. If this curve is downward sloping as a function of the horizon, then long-term risk is less than short-term risk, a phenomenon referred to as mean reversion. In Figure 1 the term structure is shown for the real returns on the three asset classes. Ultimately, all term structures are upward-sloping at long horizons.

7 The ‘Merton’ weight entries in the table are from a mean-variance optimisation with short-sell constraints on equity and bonds. Since the short-sell constraint for bonds is always binding, a short-cut for an intuitive understanding is the simple calculation $x_S = \frac{\mu}{\gamma\sigma^2}$. For $\gamma = 2$ we thus find $\frac{0.0586}{2\times 0.1799^2} = 0.90$ using the numbers in Table 1.

The term structure of risk for the real return on equity is upward-sloping due to inflation risk. Only nominal bonds exhibit some mean reversion over short to medium horizons. Long-term bonds and cash have comparable risks over long horizons.

4 Results

4.1 Second Pillar Benefits

We start by presenting the main results regarding individuals who place value on second pillar benefits, considering the first pillar as a subsistence level. Figure 2 shows the median and lowest 1% ('bad luck') quantile of the real retirement benefits using the optimal life-cycle portfolios. Our first observation is that the rankings of the median and the 'bad luck' quantile are as expected from the different risk aversion parameters. With low risk aversion ($\gamma = 2$), median benefits are the highest, but the 'bad luck' quantile is the worst. The opposite is true for the high risk aversion ($\gamma = 10$) quantile. Only individuals with low risk aversion are able to maintain median benefits at a relatively constant level as they age. With higher risk aversion, the median benefits are not only initially lower, but also decrease with age. The age effect on the benefit curves is partly attributable to the fund's payout policy. With a more conservative, *i.e.* lower, AIR in (4), benefits are paid more slowly, such that the median curves for medium ($\gamma = 5$) and high risk aversion become flatter, while the curve for low risk aversion will become more upward sloping. Changing the AIR's will, however, also have an effect on the optimal investment policy.

In absolute terms, participants with a complete pension (42 years) and low risk aversion have a median replacement ratio above 100% but with a 'bad luck' outcome below 30%.[8] With high risk aversion, median initial benefits are approximately 70% of income and the 'bad luck' outcome level is significantly above the level of the low-risk-aversion outcome. The benchmark *CPB* curves in Figure 2 are in between the results for low and medium risk aversion, both for the median as well as the 'bad luck' outcomes.

The distribution of benefits results in a certainty equivalent (CE) real benefit that varies depending on the risk aversion parameter. The cost of a suboptimal life-cycle strategy for a participant with risk aversion level of γ is the CE loss of that strategy relative to her optimal strategy. The CE loss is highly asymmetric. For example, a high-risk-aversion individual considers the optimal life-cycle policy for the low-risk-aversion individual as devastating: CE loss in the first line of Table 2 is 97% relative to what is optimal for the high-risk-aversion participant. The opposite effect, the loss of the low-risk-aversion individual forced to follow the optimal high-risk-aversion life-cycle, is less severe. In the third line of Table 2, it is 43% relative to the best in the column for $\gamma = 2$.

From Table 2 we also learn that the simple *Merton* life-cycle strategies are about 12% to 26% worse than the optimal life-cycles for the same level of risk aversion. The *CPB* benchmark shows CE losses between 14% and 57%. Once more, both for the *Merton* life-cycle and for the *CPB* benchmark, a

8 Note that these are second pillar pensions only. In the model all premiums are used to save for retirement income. We don't have a provision for partners or any other insurance attached to a pension scheme.

Table 2: Certainty Equivalent Costs

Strategy	γ	CE γ 2	5	10	% loss γ 2	5	10
optimal	2	43.0	5.4	0.7	–	84	97
optimal	5	31.7	33.7	17.7	**26**	–	**20**
optimal	10	24.6	30.6	22.1	43	9	–
Merton	2	37.9	2.3	0.3	12	93	99
Merton	5	31.6	27.5	11.3	26	18	49
Merton	10	21.6	26.1	16.4	50	23	26
CPB benchmark	–	37.0	27.4	9.5	14	19	57

Rows report the certainty equivalent real retirement benefit as a percentage of the last income for the optimized portfolio, incorporating three levels of risk aversion, the Merton fractions, and the benchmark portfolio. The left panel displays valuations by individuals with the risk aversion indicated in the column heading. In the right panel, the same data is represented as a loss relative to the best in the column. **Bold** entries indicate the minimax loss.

high-risk-aversion individual is most sensitive to any deviation from her optimal strategy.

Costs for the optimal life-cycle for the medium-risk-averse participant have been entered in bold in Table 2. From the strategies under consideration this is the life-cycle strategy that has the lowest maximum costs. Participants with either low or high risk aversion do not view it as optimal, but both groups similarly affected. Among the portfolio strategies in the table, this middle ground may be the better choice. Without assumptions about which group of participants the fund prioritizes, and which group is the most sizable, it is impossible to offer a more specific collective life-cycle portfolio strategy. Balter and Schweizer (2024) suggest that the pension may give more weight to the high-risk-aversion individuals in order to reduce inequality in CE outcomes.

With very different CEs, the underlying portfolio strategies must also vary considerably. Equity and bond exposures for different levels of γ are shown in Figure 3. Even though we allow for a very flexible functional form, the optimized life-cycles look very simple. They primarily differ in how quickly the large exposure to equity is decreased to which fixed level at a later age. The equity life-cycles closely

Figure 2: Real retirement benefit quantiles

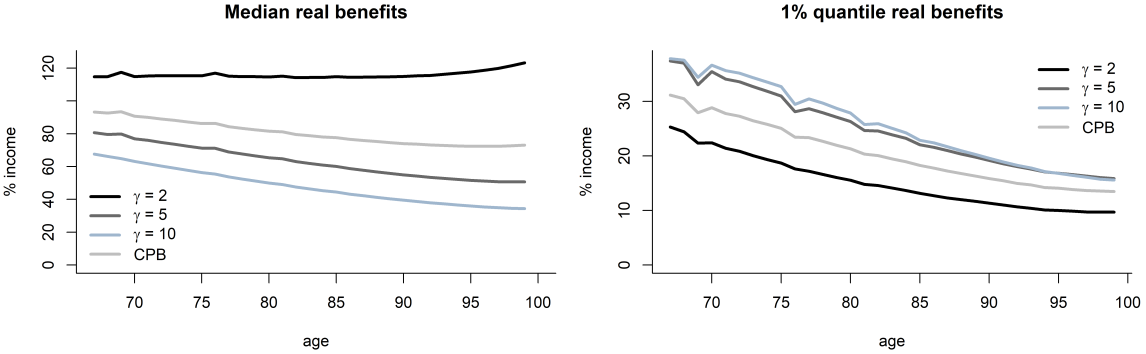

The figures show the median and 1% quantiles of real benefits for the optimised life-cycle portfolio strategies with different values of γ. The additional 'CPB' line represent the benchmark linear life-cycle strategy.

Figure 3: Portfolio weights

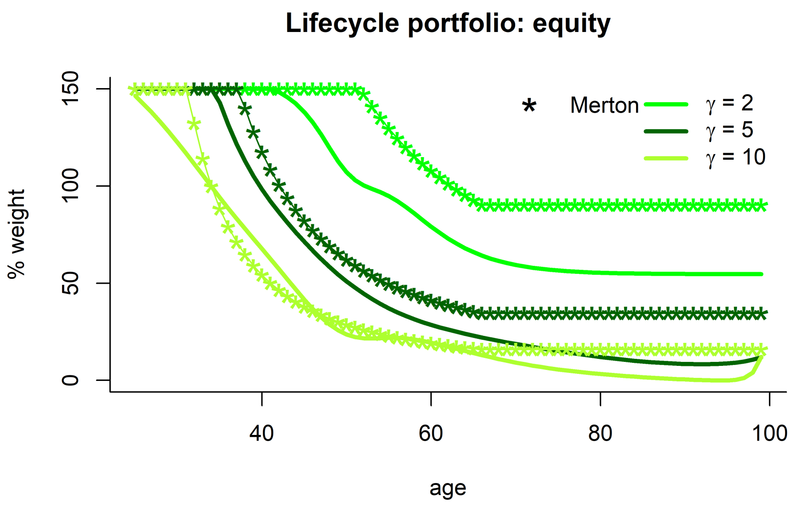

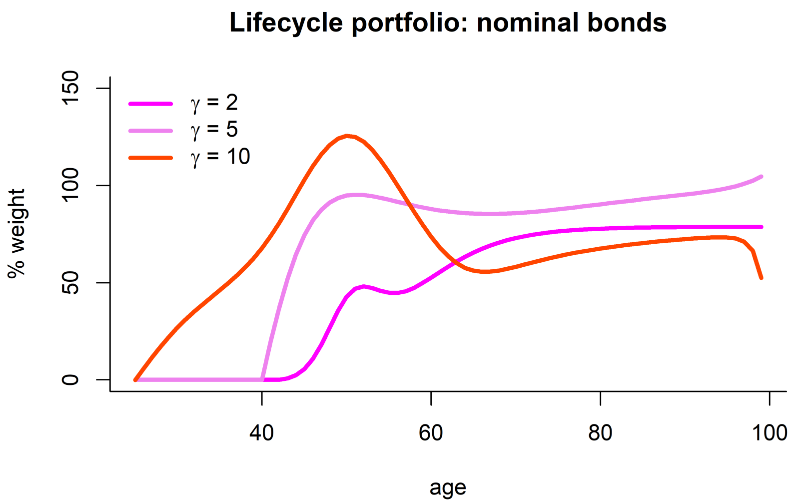

The two graphs show the optimized portfolio weights for equity and bonds as a function of age. The Merton weights for nominal bonds are not shown as these are all equal to zero.

resemble the standard 'glide paths'. At a young age the maximum equity weight of 150% is binding for everyone except individuals with $\gamma = 10$, for whom it is exactly their optimal exposure. The differences lie in the age at which the exposure to equity is reduced. The *Merton* equity weights are consistently larger than the optimal counterparts. This is most evident at retirement age when human capital is zero. Here, the optimal weights are much smaller than the *Merton* weights. The is mainly because equity is fat-tailed, which is an additional risk not accounted for by the *Merton* mean-variance weights. The put option in the portfolio is most effective for the $\gamma = 2$ optimal portfolio. Since the equity weights remain at 150% for an extended period, the probability of ever hitting a bad shock that triggers exercise of the put option during the lifetime is 2.5%. In contrast, for the $\gamma = 10$ optimal portfolio, the put option is exercised with a probability of less than 1%.

Bond exposures are more complex. Because equity is at its maximum at a young age, and total

Table 3: Certainty Equivalent Costs with First Pillar AOW

		2^{nd} pillar γ			$1^{st}+2^{nd}$ pillar γ		
optimized strategy	γ	2	5	10	2	5	10
2^{nd} pillar	2	–	84	97	**10**	**0**	**3**
2^{nd} pillar	5	26	–	19	33	10	2
2^{nd} pillar	10	43	9	–	41	18	8
$1^{st}+2^{nd}$ pillar	2	27	99	99	–	6	12
$1^{st}+2^{nd}$ pillar	5	1	64	90	13	–	3
$1^{st}+2^{nd}$ pillar	10	14	11	54	26	5	–
CPB benchmark	–	14	19	57	25	5	1

Rows report the certainty equivalent percentage losses for each strategy in the first column when evaluated using preferences in the headings of the other columns. Empty cells are the optimal against which there is zero loss. The entries in the upper left quadrant (3 rows, columns) are identical to the losses in the first 3 rows of table 2. For this table, the AOW is a real benefit equal to 64% of the last income. **Bold** entries indicate the minimax solution for the strategies and preferences considered.

leverage is restricted to 150%, nominal bonds have a zero weight in the portfolio.[9] Unlike the *Merton* weights, bond exposure increases with age. The optimizer recognises that retirement benefits are closely tied to the nominal yield curve, creating a powerful hedging incentive.

The exact optimal policy remains numerically imprecise. The expected utility function is extremely flat around the optimal policy. Even with 50,000 scenarios, we observe a lot of sampling variation in the optimized portfolios. The portfolio weights at the end of the life-cycle are ill-determined. At an advanced age, very few participants are still alive. That small fraction of survivors has minimal impact on the overall life-cycle optimization. The decisions may be important for those still alive, but are negligible for the assessment of the policy over the entire planning horizon.

4.2 First Pillar Benefits

Up to this point we have assumed that the CRRA utility pertains to the level of second pillar income, rather than the total retirement benefits including the first pillar AOW. Formally, this means we utilize a preference function with a subsistence level equal to the first pillar. While probably a reasonable behavioral assumption, and definitely a convenient one for the analysis, we have not seen much research justifying AOW as the subsistence level.[10] The experiments in eliciting a risk aversion parameter are focused on the pure CRRA specification without AOW as a subsistence level or any other lower bound.[11]

With the addition of AOW, we need to be more specific about income. The first pillar is more

9 If allowed, the weight of nominal bonds would be positive, leading to an even more negative position in cash. See Figure 7 in appendix Appendix E.

10 It is convenient for the analysis because with CRRA utility the optimal allocations do not depend on the level of income or wealth. We thus do not have to be specific about the level of income and construct separate portfolio strategies for each level of income. The first pillar income is an absolute number, however.

11 See, for example, the 'multiple lottery comparison' questionnaire in Alserda et al. (2016, 2019, Figure 1), which specifically emphasizes that the utility calculation is based on total income including the first pillar. Another example is the 'Beleggingbalans' in Joseph et al. (2021).

important for low-income participants than for high-income individuals. We will thus also find different optimal life-cycle portfolio strategies depending on the level of income. For low-income participants, the second pillar benefits represent a relatively small portion of their total retirement benefits. This is why with CRRA preferences, they should be inclined to take on more risk in the second pillar. To obtain a single life-cycle strategy for collective investments, a fund would then have to average expected utility over participants with diverse life-cycle income patterns.

In Table 3 we evaluate the costs of suboptimal investments, but now from the perspective of an individual who evaluates total retirement benefits. For the cost calculations, we assume that the pension participant's last income results in AOW benefits being 64% of their last income in real terms. That income level is at the low end of the income distribution.[12] The comparison also adds three more optimized life-cycle portfolio strategies, one for each level of risk aversion ($\gamma \in \{2, 5, 10\}$), that are optimal for CRRA on total benefits.[13] Contrary to the analysis of 'second pillar only' benefits, participants with the highest risk aversion do not incur much CE loss regardless of the portfolio strategy the fund pursues. Even the strategy for the low-risk-aversion individual has a maximum loss of 12% CE. The low risk participant, however, now faces substantial costs from suboptimal portfolio strategies.

As we add three additional potential life-cycle strategies, we also evaluate their costs using the 'second pillar only' utility function, which reflects the preferences outlined in the previous section, with AOW as the subsistence level. The entries in the last three rows of the left panel of Table 3 indicate that now it is the high-risk-aversion individual again who very strongly dislikes the strategies that are optimized when AOW is included. The CPB benchmark life-cycle portfolio achieves almost the same CE values as the high gamma strategy.

The cost calculations are for individuals with low income. With higher levels of income, the optimized strategies and the associated CE losses will be in between what is on the left and on the right in Table 3. Moreover, AOW may not be the true subsistence level, and results will differ when we change this reference level to some lower value (below AOW). One conclusion from our analysis, however, is that a risk aversion coefficient alone does not determine the optimal life-cycle strategy. Both income level and additional parameters in the utility function (such as a subsistence level) have large effects on CE costs and optimal life-cycle portfolios.

5 Discussion

The portfolio life-cycles and their associated certainty equivalents depend on several assumptions. When presenting the main results, we have already discussed a few variations in the cost calculations. Below we highlight a number of additional issues. Many more are covered in the review of life-cycle portfolio choice by Gomes (2020).

Behavioral assumptions. Much of the behavioral economics literature suggests that economic

12 Since we do not model the tax system, assume that all amounts refer to the net, after-tax, income.

13 In Appendix E, the optimized life-cycle portfolios are shown in Figure 8 and the benefit quantiles in Figure 9.

agents do not assess the absolute value of their retirement benefits, but rather evaluate them relative to some benchmark. Examples include habit formation, regret aversion, and loss aversion. The benchmark can also be fixed or time-varying, or linked to a desired level of retirement income. Each of these alternative preferences, for which there exists supporting empirical evidence, leads to a distinct optimal life-cycle portfolio strategy. The result will be a wide range of potentially optimal strategies, which may be far from optimal for agents with differing preferences. Our results already demonstate considerable variety with just a few changes to the assumed preferences.

Intertemporal substitution. The retirement benefits are spread out over an extended period. In the evaluation, the CRRA utility function integrates two elements of risk into the single parameter γ. The first element is the risk aversion. At a given moment in time, it measures a participant's willingness to accept the risk of receiving a below average benefit. However, with a lengthy retirement horizon participants also face the risk of variations in their benefits over time. With CRRA preferences, this intertemporal substitution is measured by the inverse $1/\gamma$. The literature suggests that this intertemporal substitution parameter is just slightly below or above one. For this reason, most of the long-term portfolio models in the literature work with Epstein-Zin preferences that allow for the separation of risk aversion and intertemporal substitution.[14] Goossens et al. (2023) report a large difference between γ measured from a risk perspective and γ from a time perspective. Preference elicitation using risk focused methods results in a median γ of 5.5, whereas γ drops to 0.4 when assessed based on time budgets. Ideally, one should separate the two parameters in deriving the optimal life cycle and conducting the cost calculation. This is still a numerically hard problem.

Risk capacity. Like Balter et al. (2024b), we mainly concentrate on differences in risk aversion. In addition to heterogeneity in preferences, there will be heterogeneity in what is known as 'risk capacity', which takes into account individuals' income, income risk, and their wealth invested outside the pension plan. The latter also encompasses the distinction between owning and renting a home. Heterogeneity in income profiles matters, even in the standard Merton model with riskfree real income. In the model with human capital, a young individual holds a large natural position in the riskfree asset. As a person ages, their human capital decreases, while the value of the financial capital grows. For this reason, the proportion of equity within financial wealth decreases with age. The decrease will occur more quickly with shorter duration of the labor income cashflows. When labor income is relatively high in the early years, more pension premiums are paid in early years, and financial capital will be larger than when labor income is low during the early years. These effects are mechanical, yet they result in heterogeneity in the optimal risk exposure. The optimal 'glide path' for equity, going down from 150% to the most suitable level at retirement, will vary substantially among individuals. When human capital becomes risky, this further impacts the optimal portfolio decision. In the simplest case, when human capital is risky but unrelated to any financial returns, it is background risk that results in a lower optimal investment in stocks. In other cases, for example in Catherine (2022), the covariance between stock returns and labour income

14 Among the numerous studies, an early example is chapter 4 in Campbell and Viceira (2002).

could even imply an upward sloping life-cycle for the equity allocation.

Additional assets. We restricted the asset menu to the traditional asset classes of equity, nominal bonds and cash. Naturally, results will be quite different when the fund also has access to inflation-linked bonds. As we know since Campbell and Viceira (2002, Chapter 3) and Brennan and Xia (2002), a long-term inflation-linked bond is the appropriate riskfree asset for a long-term investor who has preferences defined on real benefits. In that case the inflation-indexed bond serves as the hedging instrument, while equity is primarily chosen for its expected return.

The two nominal bonds, representing the 1-year and 10-year interest rates, enable fixed-income portfolios to achieve a wide range of desired durations. When the yield curve is generated by multiple factors, as in the DNB scenarios, a duration based portfolio may be suboptimal. By adding a third bond with a much longer duration, for example a 30-year bond, one could also optimize the convexity of the bond portfolio. However, this will at most have a second-order effect on the overall results, which are dominated by the age-dependent exposure to equity throughout working life.

Since the duration of the expected benefits decreases with age, neither a 10-year nor a 30-year bond will have the optimal duration throughout the entire life-cycle. We redid our life-cycle calculations for different choices regarding the nominal bond. One specifc choice was a portfolio of matched-maturity bonds, designed to hedge the nominal interest rate risk in expected future benefits. At the start of the working life the duration of such a bond portfolio will be approximately 52 years (depending on the yield curve at the initial conditions), gradually decreasing to 12 years at retirement and reaching zero towards the end of the lifetime. When we construct the optimal life-cycle with this matched-maturity bond, the resulting allocation to equity remains relatively stable. However, the weights for the bond obviously do change due to its very different and time-varying duration. More importantly, the certainty equivalent losses are very similar to the baseline case we have discussed thus far.

All of the above add to broadening the set of personalized optimal life-cycles and increasing the number of entries in a cost matrix that details the costs of each potential strategy for individuals based on their unique preferences and economic circumstances. For a collective fund strategy, the board of the pension fund must define some weighted average of all these individual optimal strategies. Some participants, however, will always find the fund's collective life-cycle to be very costly. But although a single collective strategy may be very costly for at least some participants, clustering the different profiles to obtain a small number of investment strategies, maybe not more than four or five, seems feasible (Balter et al., 2024a).

We conclude the discussion by addressing some methodological issues.

Time consistency. For the main results, we searched for the optimal portfolio strategy over the entire life-cycle. The planning could also be viewed as being divided into two separate problems. Given wealth at retirement, what is the optimal portfolio strategy during the decumulation phase?

Figure 4: Portfolio Weights after Retirement

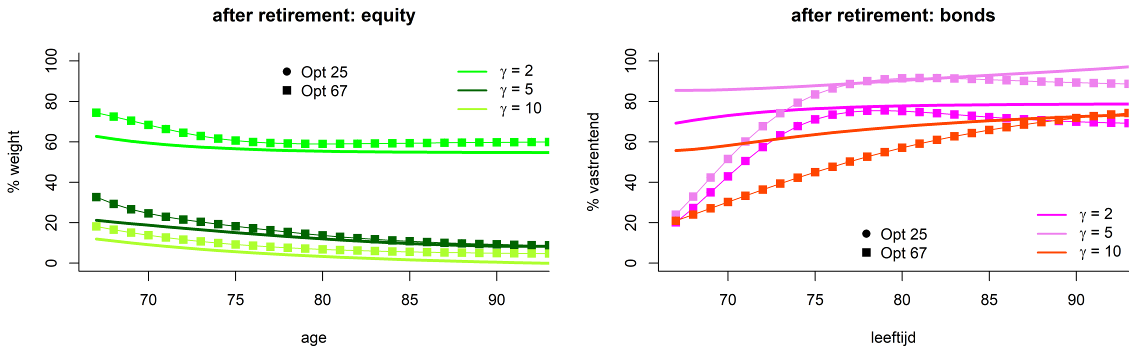

The figures show the optimized portfolio weights for equity and nominal bonds for an investor who starts planning at retirement. The lines with square markers are from optimisation starting at age 67. The solid lines, which are identical to Figure 3, are added for comparison.

This optimization problem may also be relevant for someone who joins the fund close to the retirement date. The initial planning problem is then planning during one's working life by using the implied utility function for retirement wealth. Using dynamic programming, the optimal portfolio strategy during retirement is computed first. Conditional on what is optimal during retirement, the fund can then ascertain what is optimal during working life. This is dynamically consistent, but suboptimal, and it will inevitably result in a CE loss compared to the full life-cycle plan we discussed earlier. This may seem counterintuitive, since optimal dynamic strategies are typically constructed by backward induction. It occurs in this special case because the life-cycle policy depends solely on age. Scenarios with bad outcomes at the retirement date may continue to deliver poor performance during retirement. When optimizing over the full life-cycle the fund will thus attempt to select a portfolio strategy that mitigates this accumulation of bad outcomes. When optimizing solely for the retirement stage, conditional on retirement wealth, bad outcomes will have a less severe impact compared to a situation when retirement wealth is already affected by earlier losses. Appendix C provides the technical explanation. Time-inconsistent portfolio rules arise in different settings as well, see for example Balter et al. (2021b).

The difference between the full life-cycle planning and focusing solely on the retirement stage can be large. Figure 4 shows equity and bond weights when optimized from the perspective of a 67-year-old. In the equity exposure we see the effect of reduced risk-taking due to the accumulation of bad outcomes.[15] For bonds, the differences are even more pronounced for all levels of risk aversion.

State dependence. By focusing on pure age-dependent life-cycle policies, we deviate from the vast academic literature on strategic asset allocation. We justified this choice by referencing the volatile fluctuations in portfolio weights when allowing state-dependence. For example, consider the effect of the short-term interest rate. Within the scenarios, the expected log return on equity is constant.

15 In much of the literature the effect would be the reverse. When equity is mean-reverting, bad outcomes tend to be followed by better than average outcomes. For the DNB scenarios, the term structure of risk in Figure 1 shows that instead of mean reversion the risk of equity increases with the investment horizon.

Thus, the equity premium varies in proportion to the 1-year nominal interest rate. This would lead to large swings in the equity allocation, even without dynamic hedge effects. Commissie Parameters (2022, p 35) rightly remarks that (large) changes in the real interest rate will have an impact on the expected return on equity. Therefore, a future *Commissie Parameters* will probably adjust the expected return on equity when the economic environment changes. This will make it awkward to fully optimise conditional on the currently available scenarios. In the current setting the age-dependent portfolio weights are derived from averaging over all economic scenarios, which encompass both high and low interest rates as well as a high and low equity premium.

Some degree of state-dependence in life-cycle strategies may still be welfare enhancing, but only if we restrict the optimizer to avoid large changes from one period to the next. One way to implement this is by adding adjustment costs (Bams et al., 2016). Optimizing with adjustment costs remains a numerical challenge. In ongoing work we are working to implement this by training a neural network to determine the policy function that sets portfolio weights as a function of the states of the economy and the planning horizon.

6 Conclusions

Below we summarize our main conclusions.

1. Optimal allocations are highly sensitive to how first pillar benefits are treated. At high income levels, or when preferences focus solely on second pillar benefits, participants in the fund with high risk aversion incur the greatest certainty equivalent costs from pursuing a suboptimal investment strategy. When the first pillar comprises a large fraction of total retirement benefits, the group of low-risk-aversion participants is most affected by suboptimal policies.

2. Simply knowing the level of relative risk aversion is not enough to define an optimal life-cycle portfolio plan. The expected life-cycle of real income is equally important.

3. In environments such as the DNB scenarios, the standard 'Merton' life-cycle strategy is far from optimal. The 'Merton' strategy typically assumes greater risk than optimal, except at the highest level of risk aversion, because it fails to account for the time-varying volatility and the persistence of the state variables in the DNB scenarios.

4. Optimal life-cycle portfolio strategies vary widely based on risk aversion, except at young age when the optimal equity allocation reaches its upper limit of 150% regardless of the risk aversion parameter.

5. In many cases the optimal life-cycle strategy for a medium risk averse individual balances the large certainty equivalent losses experienced by either high- or low-risk-aversion individuals.

6. A simple linear life-cycle strategy, like the CPB benchmark, is effective when preferences are defined on total benefits for low-income individuals.

References

Alserda, G., Dellaert, B., Swinkels, L., and van der Lecq, F. (2016). Pension Risk Preferences: A Personalized Elicitation Method and its Impact on Asset Allocation. Netspar Design Paper 62.

Alserda, G., Dellaert, B., Swinkels, L., and van der Lecq, F. (2019). Individual Pension Risk Preference Elicitation and Collective Asset Allocation with Heterogeneity. *Journal of Banking and Finance*, 101:206–225.

Balter, A., Coumans, L., and de Jong, F. (2021a). Robust Hedging of Terminal Wealth under Interest Rate Risk and Inflation Risk. Netspar DP 11/2021-023.

Balter, A., Garcia, J., and Schweizer, N. (2024a). Welfare Effects of Collective Investment for Heterogeneous Agents. Netspar Design Paper 251.

Balter, A., Mahayni, A., and Schweizer, N. (2021b). Time-inconsistency of Optimal Investment under Smooth Ambiguity. *European Journal of Operations Research*, 293:643–657.

Balter, A. and Schweizer, N. (2024). Robust Decisions for Heterogeneous Agents via Certainty Equivalents. *European Journal of Operations Research*, 317:171–184.

Balter, A., van den Goorbergh, R., and Schweizer, N. (2024b). The Impact of Uncertainty in Risk Preferences and Risk Capacities on Lifecycle Investment. Netspar Design Paper 237.

Bams, D., Schotman, P., and Tyagi, M. (2016). Asset Allocation Dynamics of Pension Funds. Netspar DP 03/2016-016.

Brennan, M. and Xia, Y. (2002). Dynamic Asset Allocation under Inflation. *Journal of Finance*, 57:1201–1238.

Campbell, J. and Viceira, L. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.

Campbell, J. and Viceira, L. (2005). The Term Structure of the Risk-Return Tradeoff. *Financial Analysts Journal*, 61:34–44.

Catherine, S. (2022). Countercyclical Labor Income Risk and Portfolio Choice over the Life Cycle. *Review of Financial Studies*, 35:4016–4054.

Commissie Parameters (2022). Advies Commissie Parameters. Rijksoverheid document.

Diris, B., Palm, F., and Schotman, P. (2015). Long-term Strategic Asset Allocation: An Out-of-sample Evaluation. *Management Science*, 61:2185–2202.

Gomes, F. (2020). Portfolio Choice Over the Life Cycle: A Survey. *Annual Review of Financial Economics*, 12:277–304.

Goossens, J., Knoef, M., Kuijpers, B., Potter van Loon, R., Ponds, E., Riedl, A., and Vos, S. (2023). A Comparison of Risk Preference Elicitation Methods and Presentation Formats. Netspar Design Paper 233.

Grebenchtchikova, A., Molenaar, R., Schotman, P., and Werker, B. (2017). Default Life-Cycles for Retirement Savings. Netspar Design Paper 70.

Joseph, A., Pelsser, A., and Werner, L. (2021). Beleggingsbeleid bij onzekerheid over risicobereidheid en budget. Netspar Design Paper 174.

Merton, R. C. (1973). An Intertemporal Asset Pricing Model. *Econometrica*, 41:867–887.

Metselaar, L., Nibbelink, A., and Zwaneveld, P. (2020). Het doorontwikkelde contract: beschrijving van varianten, opties en resultaten. CPB Achtergronddocument.

van Gastel, R., Kortleve, N., Nijman, T., and Schotman, P. (2024). Reële keuzes in het nieuwe pensioencontract. Netspar Design Paper 243.

Appendix A Protective Put

To prevent negative wealth during the accumulation stage when the equity portfolio weight exceeds 100% we add a protective put option. The put option ensures that the portfolio return stays above a lower bound $R_{\min}$, which we set at $R_{\min} = 0.1$. Participants can thus lose up to 90% of their pension wealth in a given year. The value of such a put is negligible because it is far out-of-the-money.
Without protection the gross portfolio return with maximum leverage of 150% equity is $R_p = -0.5R_C + 1.5R_S$. Pension wealth is negative whenever $R_S < \frac{1}{3}R_C$. The probability of this event in the DNB scenarios is 0.0011. This is the probability that pension wealth can become negative in any year when leverage is at its maximum. Although very small, it has a huge impact on utility calculations, as for large N there will always be scenarios in which pension wealth is negative at the retirement date, resulting in undefined utility.
We estimate an approximate premium based on the DNB scenarios. From the summary statistics we obtain a standard deviation of $\sigma = 27\%$ for a portfolio composed of 150% equity and -50% cash. The Black-Scholes value of a put option with strike (per unit of invested wealth) equal to $R_{\min}$ and maturity of one year is then negligible at 5×10^{-20}. The Black-Scholes value almost surely underestimates the value of the put option, since the returns are fat-tailed. Under a log-normal distribution, the probability that the put expires in-the-money is less than 6×10^{-18}, while the actual probability is the earlier mentioned 0.0011. Nevertheless, the approximate Monte-Carlo value of the put is only 0.0001, *i.e.* it costs one basispoint in returns to insure that losses will not exceed 90% in a given year. In our simulations we therefore work with the truncated return distribution $\max(R_p, 0.1)$ for which participants pay a price equal to one basispoint.
The one basispoint premium may actually overestimate the costs, as it assumes that the pension fund will rebalance only once a year. More frequent rebalancing will reduce the value of the option. The option value should also be lower when the equity position is below 150%. Even for the aggressive $\gamma = 2$ investor, who does not always allocate 150% to equity, the actual frequency of the option exercise in any given year is 0.0006. The probability of it occurring at least once in a lifetime is 0.024. An option with such a low strike price will not be available in the market, but it can be synthetically created within a pension fund. The amount of pension wealth at risk is limited, since the option primarily protects young participants with high leverage. At an older age the equity exposure is less than 100%. When all participants share the protection risk, their additional exposure to equity risk in bad times remains minimal. One way to implement this is by reserving the necessary capital in a solidarity buffer. The exposure is bounded since equity cannot fall below zero, *i.e.* the maximum payoff is equal to the leveraged amount.

Appendix B Merton Fractions

One of the benchmarks we consider is the unconditional mean-variance portfolio, which aims for the optimal trade-off between expected real return and variance. Let R be the vector of gross nominal returns and let $Y = R - R_C$ be the vector of excess returns in deviation of the nominally riskfree cash return R_C. The real return on the portfolio is,

$$R_{p,r} = x_r' Y + R_C - \pi, \tag{10}$$

where $x_r = (x_S, x_B)$ contains the weights of equity and nominal bonds, and where π denotes realized inflation. The mean-variance objective is to maximize

$$\mathrm{E}[R_{p,r}] = x_r'\mu - \tfrac{1}{2}\gamma(x_r'\Sigma x_r + \sigma_\pi^2 - 2x_r'\phi_\pi), \tag{11}$$

where μ are the expected excess returns, Σ is the covariance matrix of excess returns, σ_π^2 is the variance of inflation, and ϕ_π is the covariance of inflation with the excess returns.[16] The solution is:

$$\begin{aligned} x_r &= \frac{1}{\gamma}\Sigma^{-1}\left(\mu + \gamma\phi_\pi\right) \\ &= \frac{1}{\gamma}\Sigma^{-1}\mu + \beta_\pi \end{aligned} \tag{12}$$

The second line defines $\beta_\pi = \Sigma^{-1}\phi_\pi$ as the hedge portfolio of risky assets used to mitigate inflation risk. The hedge portfolio weights are the regression coefficients of inflation on excess returns.

In practice, because R_C varies over time and across scenarios, the β_π coefficients are obtained from a regression of $R_{C,i,t-1} - \pi_{it}$ on excess returns Y_{it}. Furthermore, the weights on equity and bonds are restricted to be non-negative. The latter restriction is often binding for nominal bonds, in which case we re-optimize subject to the short-sell constraint.

The mean-variance weights would be optimal for a myopic investor without human capital. Even after retirement, when human capital is no longer available, an investor must still take into account the interest and inflation risks for future pension benefits and may therefore optimally deviate from the mean-variance portfolio.

Before retirement the Merton model scales the mean-variance weights by the ratio of total wealth to financial capital. The standard transformation is (see, *e.g.*, Campbell and Viceira (2002, eq (6.1))):

$$x_{M,t} = x_r\left(1 + \frac{H_t}{F_t}\right), \tag{13}$$

where H/F represents the ratio of human capital to financial capital. The formula remains a crude approximation, because human capital offers riskfree real payoffs ('income'), while the cash investment is nominally riskfree and varies across scenarios. At the beginning of one's working life, once the first year's premium has been paid, the ratio is likely around of 40, implying a huge leveraged equity position, regardless of the level of risk aversion. Since we restrict the equity exposure to 150% of financial wealth, the leverage constraint will be binding at all levels of risk aversion. At a young age risk aversion has little impact on optimal equity investments. After five years, when the H/F ratio will typically have dropped to around 6, the leverage constraint stops being binding for the most risk averse participants. To make the benchmark easily computable we compute the ratio H_t/F_t under the simplifying assumption that the expected real rate is constant at zero percent. In that case the

16 Note that R_C has zero variance as it is nominally riskfree.

ratio reduces to

$$\frac{H_t}{F_t} = \frac{\sum_{s=t+1}^{\tau} a_s}{\sum_{s=0}^{t} a_s} = \frac{\text{future contributions}}{\text{past contributions}} \tag{14}$$

This is admittedly a very simplified benchmark. While it does not account for the dynamics of inflation and interest rates, it does utilize the information in the unconditional first and second moments of real returns.

Appendix C Strategic, Precommitment and Time-Inconsistent Solutions

In the paper we consider unconditional life-cycle strategies with portfolio weights solely depending on age. This appendix explains the difference with strategic portfolios and illustrates the time-inconsistency. We consider an example where an investor makes portfolio decisions at times 0 and 1, and cares about the expected utility at the end of period 2. The investor can choose between a riskfree asset and a risky asset. To keep the example simple, we assume that the return on the risky asset can take only two values each period. The process is shown in Figure 5.

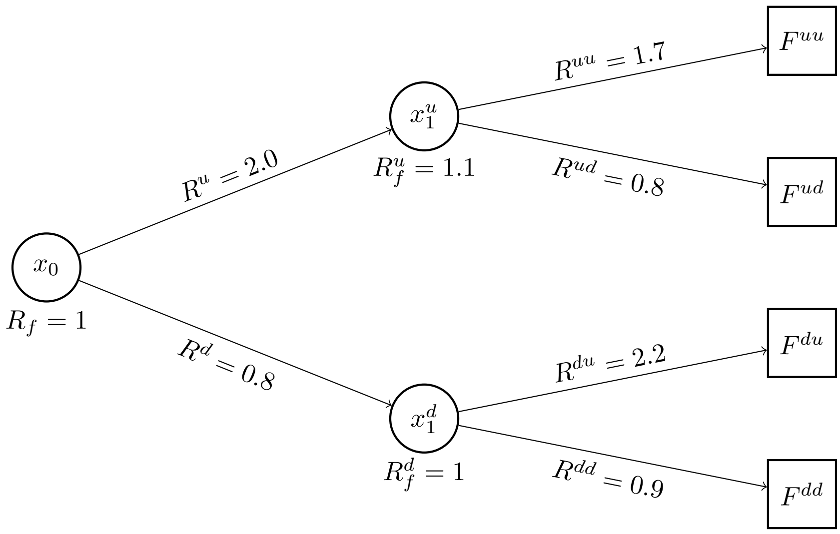

Figure 5: Binomial decision tree

The return in period 1 is $R_f + x_0(R_1 - R_f)$, where R_f is the riskfree (gross) return and R_1 the return on the risky asset. The return R_1 can take the values R^u or R^d; x_0 denotes the weight of the risky asset. Depending on the realized return in period 1, the return in period 2 can take four different values with risky returns R^{uu}, R^{ud}, R^{du} or R^{dd}. Terminal wealth can thus take four values, with for example $F^{uu} = (R_f + x_0(R^u - R_f)) \times (R_f^u + x_1^u(R^{uu} - R_f^u))$. We assume the investor has CRRA preferences with $\gamma = 5$. The goal is to maximize

$$\mathrm{E}[U] = \frac{1}{4}\left(U(F^{uu}) + U(F^{ud}) + U(F^{du}) + U(F^{dd})\right), \tag{15}$$

by choosing portfolio weights x_0, x_1^u and x_1^d. The first best is the strategic portfolio, when the investor solves the problem by backward induction. She will first find the optimal portfolio weights x_1^u and x_1^d at time 1, and given these optimal decisions select the optimal weight x_0. Using the numbers in Figure 5, the results are in Table 4 labeled as *strategic*. Investment opportunities in the lower part of

Table 4: Optimized Portfolios (Binomial Tree)

Strategy	x_0	x_1^u	x_1^d	r_{CEQ}
Strategic	28.6	17.3	47.2	15.82
Precommitment	29.8	39.3		15.30
Time consistent	28.8	31.5		15.09
Time inconsistent	29.8	17.3	47.2	15.81

The table shows optimal investments in the risky assets (in percent of wealth) using data in figure 5. The r_{CEQ} in the last column is the certainty equivalent percentage return of the optimal strategy.

the tree — after an initial bad return for the risky asset — are much better than in the upper part of the tree. The weight x_1^d is thus much larger than x_1^u. If investment opportunities were identical at every node in the tree, we would be back in the 'Merton' model with the same weight for the risky asset at every node at all times.

Typical for the strategic weights are the very active strategies with wildly changing weights depending on the state of the world. The life-cycle strategies discussed in the paper impose the restriction that weights remain constant across states of the world, *i.e.* $x_1^d = x_1^u$. Weights can still depend on the investment horizon, so $x_0 \neq x_1$. Without state dependence, the optimal strategy is found by simultaneously optimizing over x_0 and x_1. In Table 4, the result of this strategy is referred to as *precommitment*. The portfolio weight at time 1 falls between the two values of the strategic portfolio, because at time 0, when the investor is formulating the life-cycle strategy, she cannot know the state at time 1. By construction the expected utility is lower than that of the strategic portfolio. In terms of certainty equivalent returns, the difference is $15.82 - 15.30 = 0.52\%$ in favor of the strategic portfolio.

A problem with *precommitment* is that it leads to time-inconsistency. By simultaneously optimizing over x_0 and x_1 the strategy would be optimal when the investor commits to following the intended plan. However, when the investor arrives at time 1 and re-optimizes, she will make a different decision regarding x_1. We already know that the state-dependent strategic weights x_1^u and x_1^d would then be optimal. Indeed, this *time-inconsistent* strategy enhances the certainty equivalent and (in this example) closely resembles the strategic result.

Re-optimizing at time 1 is time-inconsistent, but implies that the portfolio weights do change with the investment opportunities. The final *time-consistent* strategy imposes that the weights are equal across states, *i.e.* $x_1 = x_1^u = x_1^d$, but solves the model backwards. We first find the optimal x_1 for a one period investor who does not know ex-ante whether she will be in the up- or down-state. Her objective function is to maximise the expected utility of one-period returns, assuming equal wealth in states u and d,

$$\mathrm{E}[\tilde{U}] = \tfrac{1}{4} \sum_{k=u,d} \sum_{\ell=u,d} U(R_2^{k\ell}), \tag{16}$$

where

$$R_2^{k\ell} = R_f^k + x_1(R^{k\ell} - R_f^k) \qquad\qquad (k, \ell = u, d)$$

Given the optimal x_1, the *time consistent* investor subsequently determines the optimal portfolio at time 0. These separate optimizations will, by design, achieve a lower expected utility than the *precommitment* strategy. For the numbers in the example, Table 4 shows that indeed the optimal portfolio differs. The difference in certainty equivalent return is $15.30 - 15.09 = 0.21\%$. The time-consistent investor ignores the information in the entire path from time 0 to the terminal wealth.

To grasp the difference between the *precommitment* and *time-consistent* portfolios at time 1, the return $R^{ud} = 0.8$ is key. This 20% drop in value is the worst of the period 2 returns, and the most infavorable outcome for an investor with a 1-period horizon. That same state is less frightening for the 2-period investor, who understands that this bad return will only happen after a very good first period return, and therefore will not be the worst overall outcome. The 2-period investor is worse off with two bad returns, R^d and R^{dd}. When re-optimizing at time 1 the short-term investor does not take into account that she will have seen good past returns before facing the risk of the bad return R^{ud}.

The strategy we consider in the paper corresponds to *precommitment*.

Appendix D Portfolio Weights

The function $\mathbf{f}(t|\theta)$ describes life-cycle portfolio weights as a function of age. Start by partitioning $\mathbf{f}(t|\theta)' = \begin{pmatrix} f(t|\theta)' & 1 - \iota' f(t|\theta) \end{pmatrix}$, where $f(t|\theta)$ is a vector of length 2 corresponding to the weights for equity and bonds, and ι is a 2-vector of ones, such that $1 - \iota' f(t|\theta)$ is the weight of cash. In our model age t is the only input. For the function $f(t|\theta)$ we then specify

$$\begin{aligned} f(t|\theta) &= \min\left(L, \max\left(0, h_3(t)\right)\right) \\ h_3(t) &= a_3 + C_3 h_2(t) \\ h_2(t) &= \sigma\left(a_2 + b_2 t + C_2 h_1(t)\right) \\ h_1(t) &= \sigma\left(a_1 + b_1 t\right), \end{aligned} \tag{17}$$

where the vectors a_1, a_2, a_3, b_1 and b_2, and matrices C_2 and C_3 together form the parameter vector θ. For the results reported in the text, $h_1(t)$ and $h_2(t)$ both have 3 elements. This leads to a total of 29 unknown parameters in θ. The activation function $\sigma(z) = (1 + \mathrm{e}^{-z})^{-1}$ is applied elementwise. The constant L is the maximum leverage, which we set to 1.5. Since the leverage also applies to the sum of the weights of the risky assets we add a penalty term that enforces $\iota' f(t|\theta) \leq L$. The leverage constraint applies both to equity and bonds individually as well as on the portfolio level.

Appendix E Additional Figures

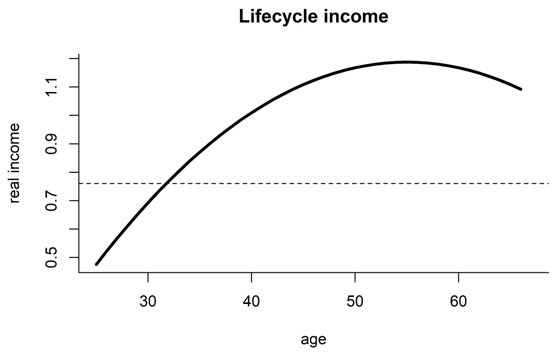

Throughout the paper we work with the above real age-earnings profile during working life. The dashed line is the first pillar income used in Section 4.2.

Figure 6: Cash life-cycle

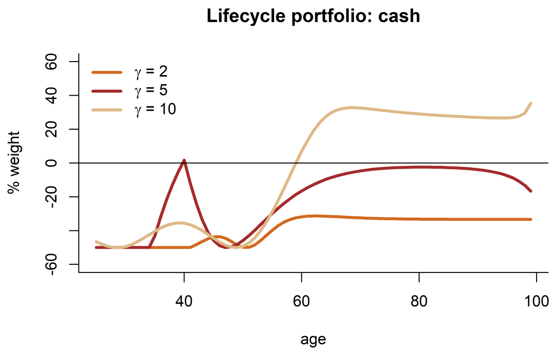

The figure shows the optimized weight of the cash investment (1-year nominal bonds) as a function of age. These are defined as $x_C = 1 - x_B - x_S$ with x_B and x_S shown in Figure 3.

Figure 7: Unrestricted Bond Weights

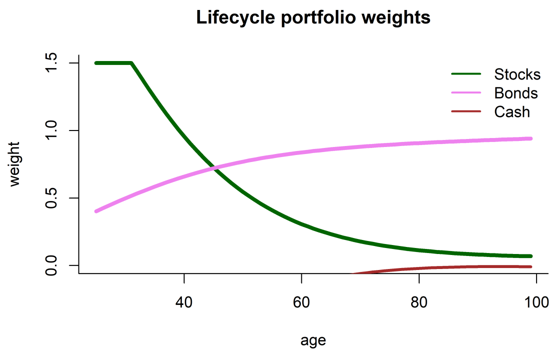

The graph shows the optimal life-cycle portfolio strategy for an investor with $\gamma = 5$ who does not face a restriction on bond investments. Only equity is restricted to a maximum of 150% (instead of the sum of equity and nominal bonds as in the baseline model).

Figure 8: Optimal Portfolio Weights with First Pillar Income

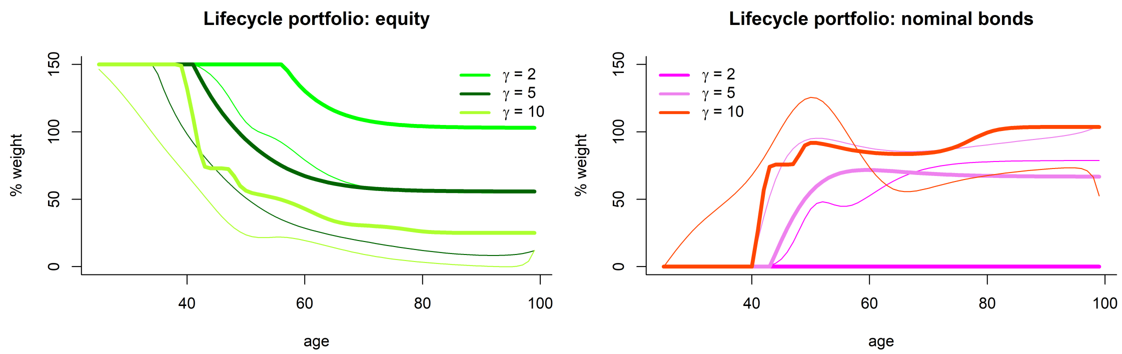

The graph shows the optimal life-cycle portfolio strategy for an investors that maximise expected utility of total lifetime retirement benefits (thick lines). For comparison the figure also shows the life cycle portfolio weights when utility is derived from second pillar income only (thin lines, identical to figure 3).

Figure 9: Real Retirement Benefit Quantiles including AOW

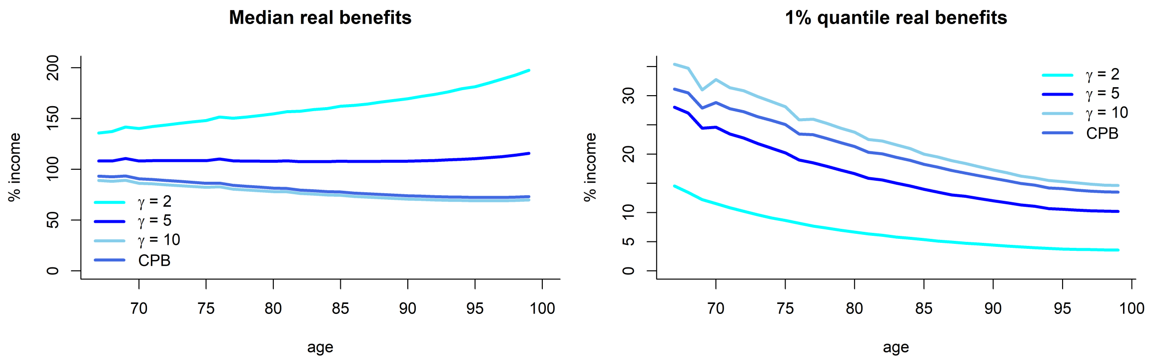

The figures show the median and 1% quantiles of real benefits for the optimized life-cycle portfolio strategies with different values of γ when preferences are defined for total benefits, including AOW. The 'CPB' line represents the benchmark linear life-cycle strategy.