



DISCUSSION PAPER PI-1112

Target-Driven Investing: Optimal Investment Strategies in Defined Contribution Pension Plans under Loss Aversion

David Blake, Douglas Wright and Yumeng Zhang

September 2011

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University London
106 Bunhill Row
London EC1Y 8TZ
UNITED KINGDOM

<http://www.pensions-institute.org/>

Target-Driven Investing: Optimal Investment Strategies in Defined Contribution Pension Plans under Loss Aversion

David Blake,⁺ Douglas Wright* and Yumeng Zhang[#]

September 2011

Abstract

Assuming loss aversion, stochastic investment and labour income processes, and a path-dependent target fund, we show that the optimal investment strategy for defined contribution pension plan members is a target-driven ‘threshold’ strategy. With this strategy, the equity allocation is increased if the accumulating fund is below target and is decreased if it is above. However, if the fund is sufficiently above target, the optimal investment strategy switches discretely to ‘portfolio insurance’. We show that under loss aversion, the risk of failing to attain the target replacement ratio is significantly reduced compared with target-driven strategies derived from maximising expected utility.

Keywords: defined contribution pension plan, investment strategy, loss aversion, target replacement ratio, threshold strategy, portfolio insurance, dynamic programming

JEL: C63, D91, G11, G23

[Typos corrected 15/9/11]

⁺ Professor David Blake, Director of the Pensions Institute, Cass Business School, City University, London (email: d.blake@city.ac.uk);

*Dr. Douglas Wright, Senior Lecturer, Faculty of Actuarial Science and Insurance, Cass Business School, City University, London (email: i.d.wright-1@city.ac.uk);

[#] Dr. Yumeng Zhang, Structurer, Solutions Group, Legal & General Investment Management.

1 Introduction

The purpose of this paper is to determine the optimal dynamic investment strategies for defined contribution (DC) pension plans when plan members experience loss aversion.

The concept of ‘loss aversion’ was first proposed by Kahneman and Tversky (1979) within the framework of prospect theory (PT), the foundation stone of behavioural finance. The recent literature on behavioural finance has provided powerful evidence that the standard optimisation paradigm, expected utility maximisation within a framework of risk-averse economic agents, does not correspond well with how economic agents actually behave in real world risk situations.¹ Real world investors are prone, among other things, to overconfidence in their investment abilities, regret and, especially, loss aversion. They also tend to monitor the performance of their portfolios (particularly their long-term portfolios) ‘too frequently’. As a result, they tend to become risk averse when winning and sell winning investments too quickly, and avoid cutting losses and even take extra risks when they have made losses.²

Loss aversion (LA) is defined in terms of gains and losses in wealth relative to a pre-defined reference or endowment point, rather than in terms of changes in the absolute level of total wealth, as with expected utility theory (EUT). LA has been applied in a number of recent studies of investment behaviour, largely because it can explain better many observed behavioural traits in investment decision making that are hard to rationalise in an expected utility setting.

For example, Benartzi and Thaler (1995) offer an explanation of the equity premium puzzle³ using LA, arguing that, when viewed myopically, volatile equity returns make equity investments look particularly unattractive. They called this behaviour ‘myopic loss aversion’

¹ Kahneman and Tversky (1979) developed this theory to remedy the descriptive failures of subjective expected utility theories of decision making.

² See Mitchell and Utkus (2004) for a review of investment decision making in pension plans in a behavioural context.

³ The equity premium is defined as the difference in expected returns between stocks and a risk-free asset such as treasury bills. Empirical evidence shows that stocks have outperformed bonds over the last century by a significant margin. In an EUT framework, this leads to an ‘equity premium puzzle’, since investors would have to be unreasonably risk averse to require such a high risk premium to hold stocks. Mehra and Prescott (1985) estimated that investors would have to have coefficients of relative risk aversion in excess of 30 to explain the historical equity premium, whereas previous estimates and theoretical arguments in EUT suggest that the actual figure is close to 1 (implying that many individuals are, essentially, indifferent to a fair gamble).

(MLA), where ‘myopic’ refers to the short-sightedness that induces a decision maker to evaluate each alternative within a sequence independently (whereas a rational decision maker would evaluate the sequence as a whole). Gneezy and Potters (1997) tested the MLA hypothesis in laboratory experiments and found that the more frequently returns on a portfolio are evaluated, the higher the level of LA exhibited (as measured by a lower average proportion of assets allocated to equity investments). Thaler et al. (1997) arrived at the same conclusions based on an analysis of how individuals, faced with different frequencies of evaluation, chose to split funds between two assets with different levels of risk (i.e., a bond fund and an equity fund).

Rabin and Thaler (2001) have argued that EUT is manifestly not a suitable explanation for most observed risk attitudes: ‘we have also often been surprised by economists’ reluctance to acknowledge the descriptive inadequacies of [the] theory’. They suggest that LA and the tendency to isolate each risky choice should replace EUT as the foremost descriptive theory of risk attitudes.

Given the behavioural traits exhibited by many investors, it is important to investigate the consequences of using a PT utility function to determine the optimal investment strategy in a DC plan and to compare the results with those implied by the traditional expected utility model.

In a DC plan, members contribute part of their income each year to building a pension fund for retirement.⁴ The accumulated fund is then used to buy a life annuity to provide a pension income after retirement. Members are assumed to have a target replacement ratio⁵ at retirement at age 65. This translates into a target pension fund at retirement which will depend, in part, on their longevity prospects during retirement. Members are assumed to be loss averse with respect to the target retirement pension fund and to a series of annual interim target fund levels prior to retirement. The interim targets reflect the discounted value of the final target retirement fund level and, for convenience, we will treat the interim targets as being age-related.⁶ Members are

⁴ Typically, the employer also makes a contribution. If so, we count the total contribution into the plan.

⁵ The replacement ratio is defined as the ratio of pension income immediately after retirement to labour income immediately before retirement.

⁶ However, they could equally be treated as time-related (in the sense of years from final target).

also assumed to have an investment strategy, i.e., make asset allocation decisions, which aims to maximise the expected total discounted value of PT utility over the period until retirement. To do this, we use a two-asset, dynamic-programming-based numerical solution method with stochastic labour income and borrowing and short selling constraints.

Within this proposed framework, it will be shown that the optimal dynamic asset allocation strategy is the target-driven strategy, known as ‘threshold’, ‘funded status’ or ‘return banking’, that was discussed in Blake et al. (2001). With this strategy, the weight in risky assets such as equities is increased if the accumulating fund is below the relevant interim target and is decreased if the fund is above target. This is because, under loss aversion, the member is risk seeking in the domain of losses and risk averse in the domain of gains. Close to each target (whether above or below), the plan member has the lowest equity weighting (for that target) in order to minimise the risk of a significant loss relative to the target. However, if the fund is sufficiently above the target, there is a discrete change in the investment strategy and the equity weighting is increased (subject to the member’s degree of risk aversion in the domain of gains), since the risk of the fund falling below the target is now considered to be acceptably low. This strategy of increasing the equity weight as the fund value continues to rise above the target is consistent with the investment strategy known as ‘portfolio insurance’ and its role in portfolio choice under loss aversion has been noted by other researchers (e.g., Berkelaar et al. (2004) and Gomes (2005)).

If the threshold strategy is successful in the sense that the series of interim targets has been met, the overall equity weight will tend to fall with age, since the fund is in line to meet the final target fund level at retirement. Although this is similar to what happens in conventional (deterministic) ‘lifestyle’ strategies,⁷ the target-driven strategy is very different. In particular, whilst conventional lifestyle strategies typically involve switching mechanically from 100% equities only in the last 5 to 10 years before retirement and often end up holding 100% of the fund in bond-type assets at retirement, the optimal strategy under loss aversion involves a much more gradual reduction in the equity holding if the fund remains close to the sequence of targets. If, however, the fund is either well below or well above a particular target, even one near to the

⁷ Also known as ‘lifecycle’ or ‘age-phasing’ strategies (Samuelson (1989)).

retirement date, the optimal equity holding will be high for reasons given in the previous paragraph. We also show that under loss aversion, the risk of failing to attain the desired replacement ratio at retirement is significantly reduced in comparison with a traditional risk aversion framework aimed at maximising a power utility function on retirement.

We assume that the PT utility function is defined as follows (see Tversky and Kahneman (1992)):

$$U(F) = \begin{cases} \frac{(F-f)^{v_1}}{v_1} & \text{if } F \geq f \\ -\lambda \frac{(f-F)^{v_2}}{v_2} & \text{if } F < f \end{cases} \quad (1)$$

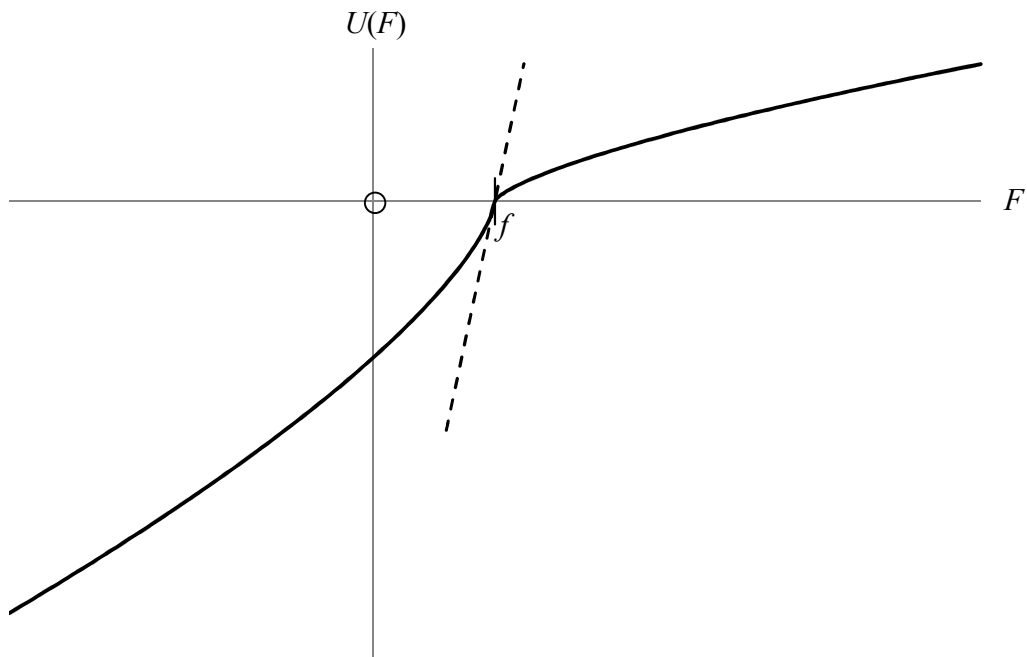
where

- F is the actual value of the pension fund when the plan member is a given age,
- f is the pre-defined target value of the pension fund at the same age,
- v_1 and v_2 are the curvature parameters for gains and losses, respectively, and
- λ is the loss aversion ratio.

As shown in Figure 1 below, the two key properties of the PT utility function are:

- the PT utility function is ‘S’-shaped (i.e., convex below the reference point and concave above it) when $0 < v_1 < 1$ and $0 < v_2 < 1$, implying that individuals are risk seeking in the domain of losses and risk averse in the domain of gains (this contrasts with the concave shape in standard utility functions, where individuals are assumed to be risk averse for all levels of wealth and have diminishing marginal utility of wealth), and
- the PT utility function is steeper below the reference point than above when $\lambda > 1$, implying that individuals are λ times more sensitive to a unit loss than to a corresponding unit gain.

Figure 1 – PT utility function (with $\lambda > 1$, $0 < v_1 < 1$ and $0 < v_2 < 1$)



Conventional lifestyle investment strategies are currently widely used by many such pension plans as the default investment option. However, as will be shown below, there can be substantial uncertainty over the size of the fund at retirement when a lifestyle investment strategy is used and this makes it difficult for the plan member to be confident about the level of retirement income.⁸ Thus, for DC plan members seeking greater certainty in their retirement planning, the plan's investment strategy needs to be far more focused on achieving the target pension.

We will assume that plan members evaluate the plan's investment performance on an annual basis. Members have a final target replacement ratio at retirement and a series of corresponding interim targets before retirement. They are assumed to be 'loss averse' with respect to these targets (which define the reference points in the PT framework outlined above) and to make asset

⁸ This is based on the assumption that members pay a constant contribution rate each year. In the real world, members would be able to assist in achieving the desired target replacement ratio by adjusting the annual contribution rate. However, previous behavioural research findings (e.g., Madrian and Shea (2001)) suggest that, once enrolled, plan members make few active changes to the contribution rates or investment decisions. This is known as the 'inertia effect'.

allocation decisions to maximise the total discounted value of the PT utility over the period until retirement. They are also, as mentioned, assumed to be ‘risk-averse’ above the target (since $v_1 < 1$) and ‘risk-seeking’ below the target (since $v_2 < 1$).

Having target fund levels when formulating the optimal investment strategy for a DC pension plan is not a new idea. Vigna and Haberman (2001) and Haberman and Vigna (2002) derive a dynamic-programming-based formula for the optimal investment allocation in DC plans. In their model, members are assumed to face a quadratic cost (or disutility) function each year based on actual and targeted fund levels and to make investment decisions that minimise the cost of deviations of the fund from these corresponding targets. Their analysis suggests that a lifestyle investment strategy remains optimal for a risk-averse member and that the age at which the member begins to switch from equities to bonds depends on both the member’s risk aversion and age when the plan started: the more risk averse the member or the longer the accumulation period prior to retirement, the earlier the switch to bonds. However, as the authors acknowledge, one obvious limitation of this approach is that the quadratic cost function penalises equally both under- and over-performance relative to the specified targets.

In summary, our proposed model differs from the existing literature in three significant respects:

- **Loss aversion:** Most existing studies (e.g., Haberman and Vigna (2002), Gerrard et al. (2004)) assume that the individual plan member has a quadratic cost function with respect to deviations in the actual fund from the appropriate targets. However, given the behavioural traits exhibited by many investors, we believe it is more appropriate to reflect this by using a PT function to determine the optimal asset allocation model in a DC pension plan.
- **Stochastic labour income:** Previous studies (e.g., Haberman and Vigna (2001)) used a simple deterministic model, whereas our study uses a more realistically-calibrated stochastic model.
- **Choice of investment targets:** The choice of an appropriate investment target is crucial in a target-driven model. Previous studies oversimplified the problem by assuming fixed targets (derived by assuming a fixed investment return over time). In our model, the final

fund target and, hence, the series of corresponding interim targets are path-dependent (and, thus, vary over time in accordance with the evolution of the member's income).

The rest of the paper is organised as follows. Section 2 formulates the target-driven asset allocation problem for a DC pension plan under loss aversion and outlines our model, including the optimisation method used. Section 3 calibrates the model's parameters. Section 4 presents the output from the optimisation exercise and conducts a sensitivity analysis. The main conclusions are presented in Section 5.

2 The asset allocation problem for a DC pension plan under loss aversion

In this section, we describe the two-asset discrete-time model with a constant investment opportunity set (i.e., a constant risk-free rate of interest, constant risk premium and constant volatility of return on the risky asset) used in the simulation process.

A number of assumptions have to be made:

- Members are assumed to join the pension plan at age 20 (without bringing a transfer value from a previous plan) and the retirement age is fixed at 65.
- A fixed contribution of 15% of current labour income is paid annually in advance.⁹
- Members desire a replacement ratio of two-thirds of income at retirement,¹⁰ which then enables us to estimate the target value for the pension fund at retirement (based on a projected final income and the expected purchase price for a life annuity at the date of retirement).
- Members are assumed to evaluate the investment performance of the portfolio annually before retirement. At each age prior to retirement, the final target fund is adjusted to reflect current income (and, hence, the new projected final income). Then, the

⁹ The justification for this particular contribution rate will be given later.

¹⁰ This is the standard replacement ratio in UK private sector defined benefit pension plans.

corresponding interim target is determined as the discounted value of the current final target fund, allowing for future contributions.

- In the baseline calculations, we give the same weight to each of the current interim targets and a higher weight to the final target at retirement.
- Members are assumed to be loss averse and to make investment decisions with the aim of maximising the expected total discounted value of the PT utility function over the period until retirement.

2.1 Financial assets

We assume that there are two underlying assets in which the pension fund can be invested:

- a risk-free asset (a bond fund), and
- a risky asset (an equity fund).

The risk-free asset is assumed to yield a constant real return of r per annum and the annual return on the risky asset during the year of age x to $(x+1)$, denoted by R_x , is given by:

$$R_x = r + \mu + \sigma Z_{1,x} \quad (2)$$

where

- μ is the annual risk premium on the risky asset,
- σ is the annual volatility (standard deviation) of return on the risky asset, and
- $\{Z_{1,x}\}$ is a series of independent and identically distributed standard normal random variables.

2.2 Labour income

We adopt the stochastic labour income process used in Cairns et al. (2006), where the growth rate in labour income over the year of age x to $(x+1)$ is given by:

$$I_x = r_I + \frac{S_x - S_{x-1}}{S_{x-1}} + \sigma_1 Z_{1,x} + \sigma_2 Z_{2,x} \quad (3)$$

where

- r_t is the long-term average annual real growth rate in national average earnings (NAE), reflecting productivity growth in the economy as a whole,
- S_x is the career salary profile (CSP) at age x , so that the term $(S_{x+1} - S_x)/S_x$ reflects the promotional salary increase during the year of age x to $(x+1)$,
- σ_1 represents the volatility of a shock from equity returns, which allows for possible correlation between labour income growth and equity returns,
- σ_2 represents the volatility of a shock to labour income growth, and
- $\{Z_{2,x}\}$ is a series of independent and identically distributed standard normal random variables which are uncorrelated with $\{Z_{1,x}\}$.

Then, assuming initial labour income at age 20 of $Y_{20} = 1$, the labour income at age x is given by $Y_x = Y_{x-1} \times \exp(I_x)$.

2.3 Pension fund accumulation and targets

2.3.1 Pension fund accumulation

Given an initial pension fund at age 20 of $F_{20} = 0$, the level of the accumulated fund at age x , denoted by F_x , is given by:

$$F_x = (F_{x-1} + \pi Y_{x-1}) \times [1 + r + \theta_{x-1} (\mu + \sigma_1 Z_{1,x})] \quad \text{for } x = 21, 22, \dots, 65 \quad (4)$$

where

- Y_{x-1} is the labour income received at age $(x-1)$,
- π is the fixed contribution rate, payable annually in advance, and
- θ_{x-1} is the proportion of the fund invested in equities at age $(x-1)$.

Note that no-borrowing and no-short-selling constraints are imposed, so that $0 \leq \theta_{x-1} \leq 1$.

2.3.2 Setting the final target

Given a risk-free real interest rate of $r = 2\%$ p.a. and assuming mortality in accordance with the projected PMA92 table,¹¹ the price of a life annuity (paying one unit annually in arrears so long as the annuitant is alive) on retirement at exact age 65 is $\ddot{a}_{65} = 15.87$. Assuming real labour income growth, r_l , of 2% p.a. and using the CSP described below in Section 3.2, the expected income at retirement is $E_{20}(Y_{65}) = 5.63$ units. Given a target replacement ratio of two-thirds of income at retirement, the initial value (i.e., at age $x = 20$) of the expected final target fund at retirement is given by:

$$f_{20}(65) = \frac{2}{3} \times 5.63 \times 15.87 = 59.6.$$

At each age x up to (and including) retirement (i.e., for $x = 20, 21, \dots, 65$), the final target fund will depend both on the income history prior to age x and the expected future income growth after age x . Thus, given an income of Y_x units at age x , the expected income at retirement is given by:

$$E_x(Y_{65}) = Y_x \times \frac{E_{20}(Y_{65})}{E_{20}(Y_x)} \quad \text{for } x = 20, 21, \dots, 65.$$

At age x , the value of the final target fund at retirement is given by:

$$f_x(65) = \frac{2}{3} \times E_x(Y_{65}) \times 15.87 = \frac{2}{3} \times \left(Y_x \times \frac{E_{20}(Y_{65})}{E_{20}(Y_x)} \right) \times 15.87 = 59.6 \times \left(\frac{Y_x}{E_{20}(Y_x)} \right)$$

implying that the final target fund is adjusted to reflect the difference between actual and expected labour income growth up to age x .

2.3.3 Setting the discounted interim target

By assuming a suitable discount rate, denoted by r^* , and a fixed annual contribution rate of π , the interim target at each age $x = 20, 21, \dots, 64$, denoted by $f_x(x)$, can be derived recursively from the final target, $f_x(65)$, using:

¹¹ PMA92 is a mortality table for male pension annuitants in the UK based on experience between 1991 and 1994. We use the projected rates for the calendar year 2010, i.e., the table PMA92(C2010), published by the Continuous Mortality Investigation (CMI) Bureau in February 2004. We assume there are no longevity improvements in this version of the model.

$$\left[f_x(s) + \pi E_x(Y_s) \right] \times (1 + r^*) = f_x(s+1) \quad \text{for } s = 64, 63, \dots, x+1, x \quad (5)$$

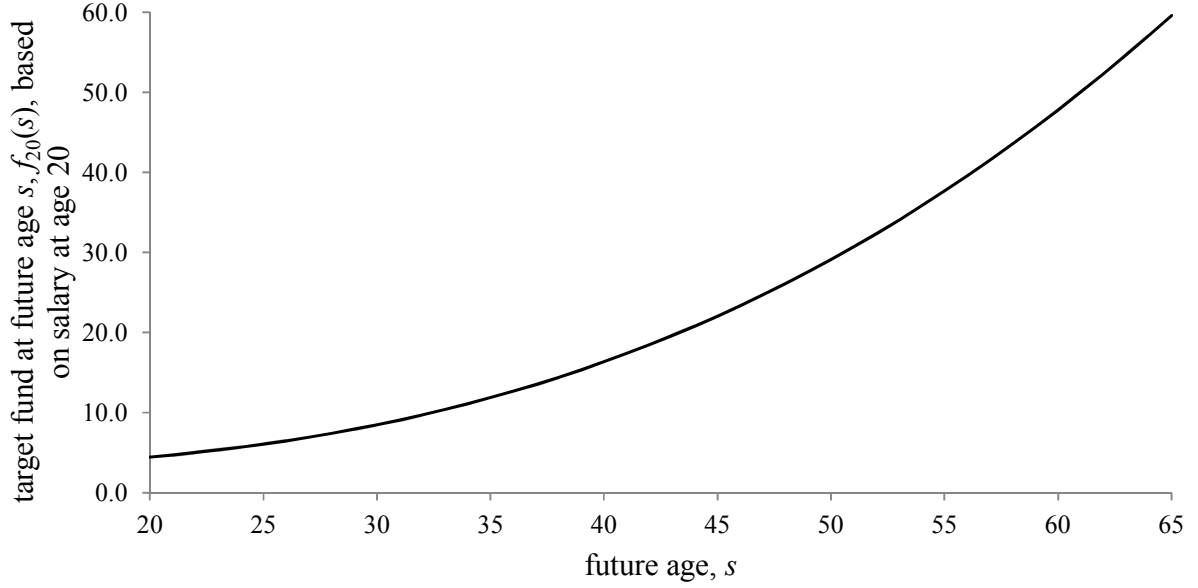
where $E_x(Y_s)$ is the expected income at age s , given an income of Y_x at age x , and is determined as for the expected income at retirement, $E_x(Y_{65})$, above.

We use the yield on AA-grade corporate bonds with a term in excess of 15 years as the discount rate, r^* , to calculate the current interim target. This is consistent with the method for valuing the liabilities of defined benefit pension plans applied by global pension accounting standards (e.g., FAS 158, FRS 17 or IAS 19). In January 2010, the credit spread on Iboxx Sterling denominated Corporate AA over 15 years was 1.1%, so we use $r^* = r + 0.011$ in our model.

Figure 2 shows the final target fund at initial age 20, $f_{20}(65) = 59.6$, and the corresponding discounted interim targets, $\{f_{20}(s); s = 20, 21, \dots, 64\}$, derived from Equation (5) above.

It is worth noting that, with a future annual contribution rate fixed at $\pi = 15\%$ and a discount rate of $r^* = r + 0.011 = 3.1\%$ p.a., the member has an initial interim target fund at age 20 of $f_{20}(20) = 4.45$. Since we assume an initial pension fund of $F_{20} = 0$, this implies that there is an apparent initial ‘deficit’. This is a direct and inevitable consequence of choosing a discount rate for liabilities that is independent of the return on assets in the fund, but it will be of no surprise to those familiar with pension accounting. The member could remove this ‘deficit’ by choosing a higher annual contribution rate of approximately 17% which would be sufficient to give an initial interim target fund of zero. However, it is not necessary to do this, since the plan member can instead (and is assumed to do so in the present study) use the asset allocation strategy to try and meet the desired final target fund level at retirement, $f_{20}(65) = 59.6$.

Figure 2 – Final and interim fund targets at age 20



2.4 Setting the target-driven objective function

For each age $x = 20, 21, \dots, 65$, given the actual fund level, F_x , and the target fund, $f_x(x)$, we assume that a plan member faces a PT utility function defined as:

$$U_x(F_x) = \begin{cases} \frac{(F_x - f_x(x))^{v_1}}{v_1} & \text{if } F_x \geq f_x(x) \\ -\lambda \frac{(f_x(x) - F_x)^{v_2}}{v_2} & \text{if } F_x < f_x(x) \end{cases} \quad \text{for } x = 20, 21, \dots, 65. \quad (6)$$

It is, however, reasonable to assume that, at any age prior to retirement, the final target is significantly more important to members than the interim targets. So, we apply a lower relative weighting coefficient, $\omega < 1$, to the interim targets than the final target which has an implicit weight of unity. Thus, at each age x prior to retirement, we define the total discounted PT utility up to retirement, V_x , as:

$$V_x = \left(\sum_{s=0}^{(65-x)-1} \beta^s \omega U_{x+s}(F_{x+s}) \right) + \beta^{65-x} U_{65}(F_{65}) = \omega U_x(F_x) + \beta V_{x+1} \quad (7)$$

where β is the member's personal discount factor.¹² We assume the member maximises the expected value of the total discounted PT utility up to retirement, appropriately weighted to account for interim and final targets.

Equation (7) generalises the optimal asset allocation problem of DC pension plans considered in much of the existing literature. Two different methodologies have traditionally been used to investigate the optimal dynamic asset allocation strategy of DC plans:

- Cairns et al. (2006) maximised the expected power utility of the terminal replacement ratio at retirement, which is a special case of this model with the weight applied to the interim targets, ω , set equal to zero and the target fund at each age x , $f_x(x)$, also set equal to zero (for $x = 20, 21, \dots, 65$): we consider this further in Section 4.2.5 below.
- Vigna and Haberman (2001) minimised the expected present value of the total cost at each age prior to retirement, which can be represented in the above framework with a quadratic cost function $U_x(F_x) = -(F_x - f_x(x))^2$.

The cost of the generalisation, however, is the loss of a closed-form solution that is available in each of these studies.

2.5 Optimisation

The optimal equity weight at each age x , θ_x for $x = 20, 21, \dots, 64$, can be determined recursively as follows:

$$\begin{aligned} \max_{\theta_x} E_x(V_x) &= \max_{\theta_x} E_x \left[\left(\sum_{s=0}^{(65-x)-1} \beta^s \omega U_{x+s}(F_{x+s}) \right) + \beta^{65-x} U_{65}(F_{65}) \right] \\ &= \max_{\theta_x} \left[\omega U_x(F_x) + \beta E_x(V_{x+1}) \right] \end{aligned} \quad (8)$$

subject to:

¹² Members with a high personal discount factor exhibit a high degree of patience and are more willing to forego consumption whilst in work in exchange for higher future consumption than members with a low personal discount factor.

- $Y_x = Y_{x-1} \times \exp(I_x)$, where $I_x = r_I + \frac{S_x - S_{x-1}}{S_{x-1}} + \sigma_1 Z_{1,x} + \sigma_2 Z_{2,x}$;
- $F_x = (F_{x-1} + \pi Y_{x-1}) \times [1 + r + \theta_{x-1} (\mu + \sigma_1 Z_{1,x})] \geq 0$; and
- $0 \leq \theta_x \leq 1$ for $x = 20, 21, \dots, 64$.

An analytical solution to this problem does not exist, because there is no explicit solution for the expectation term in Equation (8).¹³ Hence, we use stochastic dynamic programming methods to maximise the utility function, V_x , at age x and, thus, derive the optimal equity weighting at age x , θ_x . The idea is to use the terminal PT utility function on retirement at age 65, $U_{65}(F_{65})$, to compute the corresponding value function at age 64, denoted by $\max_{\theta_{64}} E_{64}(V_{64})$, and then to iterate this procedure backwards.

A crucial first step in the stochastic dynamic programming approach requires us to discretise the state space and, in particular, the shocks in the stochastic processes (i.e., equity returns and labour income growth). Wealth and labour income are discretised into 100 and 9 evenly-spaced intervals (with corresponding grid points), respectively. The normally distributed shocks in both equity returns and labour income growth are discretised into 9 nodes. The value of the PT utility function at age x , $U_x(F_x)$, is then computed using these nodes and the relevant weights attached to each.¹⁴ After determining the optimal value of the control variable, θ_x , at each grid point, we then substitute these values into Equation (8) and solve the optimisation problem for the previous time period. This process is then iterated backwards until age 20. Details of the dynamic programming and integration procedure are illustrated in the Appendix.

¹³ A path-dependent dynamic asset allocation strategy (such as a threshold strategy) combined with stochastic asset returns and labour income, plus short-selling and borrowing constraints means that there is no exact explicit solution to Equation (8). It might be possible to derive an approximate explicit solution as, for example, Campbell and Viceira (2002, Chapter 6) do in their modelling framework. However, their explicit solution comes at the cost of using a Taylor's approximation around the optimal implicit solution and it is not obvious that an approximate explicit solution dominates a numerical solution, especially if the approximation is a poor one.

¹⁴ This method is known as Gaussian quadrature numerical integration. For more details, see Judd (1998, pp. 257-266). Clearly, the choice of both the number of intervals spanning the state space and the number of nodes discretising the shocks is subjective, but, after some experimentation, we believe that our particular choice represents an appropriate trade-off between accuracy and speed of computation.

3 Calibrating the model parameters

3.1 Loss aversion parameters

Based on an experiment conducted on a group of 25 graduate students, Tversky and Kahneman (1992) concluded that, in the US, individuals are 2.25 times more sensitive to a unit loss than to a corresponding unit gain (i.e., $\lambda = 2.25$) and have gain and loss curvature parameters for their PT utility function of $\nu_1 = \nu_2 = 0.88$. We tested this parameterisation of the PT utility function but found that it resulted in much higher allocations to risky assets than was typical of DC plan members in the UK. We therefore decided to use the following parameterisation for our baseline case: $\lambda = 4.50$, $\nu_1 = 0.44$ and $\nu_2 = 0.88$. This involves a higher degree of loss aversion and a higher degree of risk aversion in the domain of gains than found amongst Tversky and Kahneman's students. This parameterisation is consistent with a relative risk aversion parameter of around 2 in a power utility framework.¹⁵

3.2 Career salary profile

For labour income in Equation (3), we adopt the following quadratic function to model the CSP (see Blake et al. (2007)):

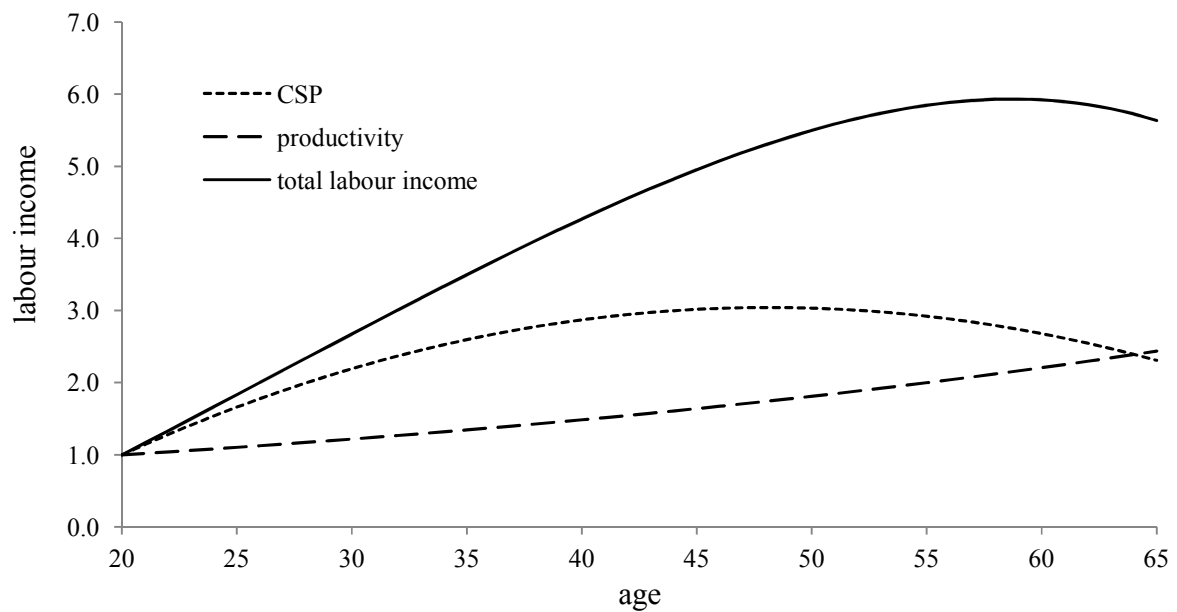
$$S_x = 1 + h_1 \times \left[-1 + \frac{(x-20)}{45} \right] + h_2 \times \left[-1 + \frac{4 \times (x-20)}{45} - \left(\frac{\sqrt{3} \times (x-20)}{45} \right)^2 \right] \quad (9)$$

The parameters h_1 and h_2 are estimated by least squares using average male salary data (across all occupations) reported in the 2005 *Annual Survey of Hours and Earnings*. The estimated parameter values are $h_1 = -0.1865$ and $h_2 = 0.7537$ for all male workers, which leads to the age-dependent CSP illustrated in Figure 3 (re-scaled to give a starting value of $Y_{20} = 1.0$). The figure shows the expected income from age 20 up to retirement at age 65, allowing for both

¹⁵ If we assume a relative risk aversion parameter of $\eta = 2$ and power utility, we have an optimal allocation to the risky asset at age 64 of $\alpha = \mu/\eta\sigma^2 = 0.04/(2 \times 0.2^2) = 0.5$ which is similar to the equity allocation with this parameterisation of the PT function at the same age (see Table 1 and Figure 16 below).

productivity growth (at a rate, r_l , of 2% p.a.) and promotional salary increases. The results suggest that individuals can expect to achieve maximum earnings (in real terms) in their early 50s, with income reducing thereafter (when many choose to work part-time in the lead up to retirement).

Figure 3 – Labour income



3.3 Baseline model parameters

For the baseline case, we use the set of parameter values presented in Table 1.¹⁶

¹⁶ These values have already been discussed previously in the text or are those commonly adopted in the UK. Using an equity risk premium of 4% p.a. (as opposed to the historical long-term average of closer to 6% p.a.) is a fairly common choice in the recent finance literature (e.g., Fama and French (2002), Gomes and Michaelides (2005)).

Table 1 – Baseline parameter values

Loss aversion parameters		Labour income	
Loss aversion ratio, λ	4.50	r_l	0.02
Curvature parameter for gains, v_1	0.44	σ_1	0.05
Curvature parameter for losses, v_2	0.88	σ_2	0.02
		h_1	-0.1865
		h_2	0.7537
Asset returns		Other parameters	
Real risk-free return, r	0.02	Contribution rate, π	15%
Equity risk premium, μ	0.04	Weight of interim target, ω	0.5
Volatility of equity returns, σ	0.20	Price of annuity on retirement, \ddot{a}_{65}	15.87
Discount factor, β	0.97		

4 Simulation results

4.1 Baseline case

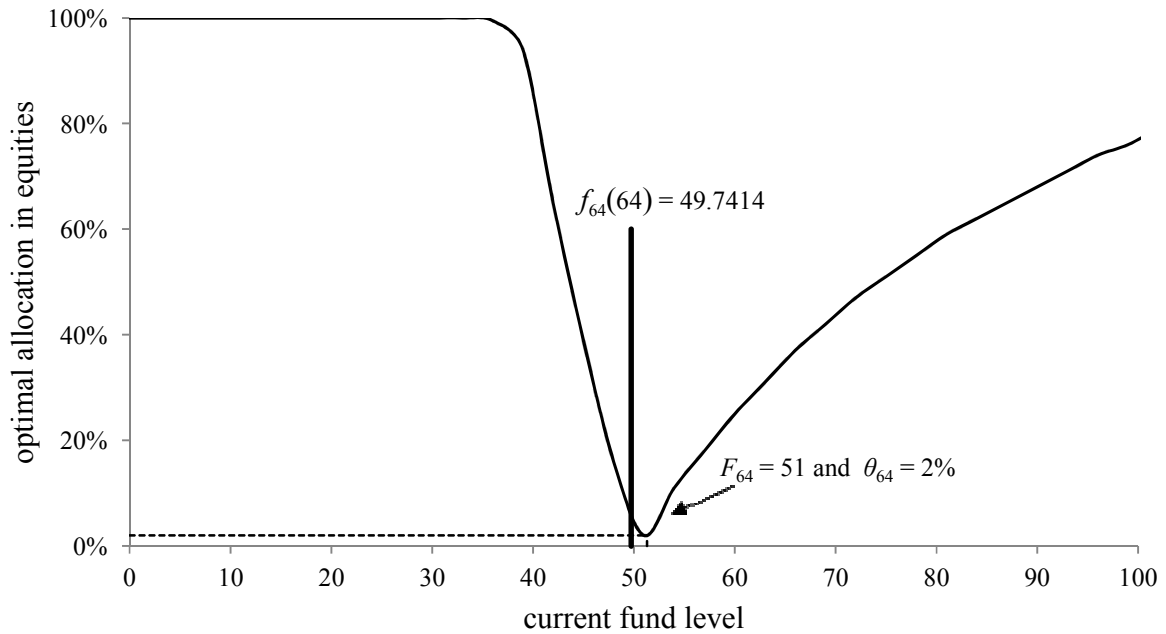
The optimisation problem is given in Equation (8) with $U_x(F_x)$ defined as in Equation (6). The outcome is a set of optimal equity allocation weightings at each age x , θ_x , for $x = 20, 21, \dots, 64$ at each grid point.

We can illustrate the outcome for the final year before retirement (i.e., age 64 to age 65). Suppose that the member's current income at age 64 happens to be 5 units. Then, from Section 2.3, the final target fund at age 65, denoted by $f_{64}(65)$, is 52.0566 units and, from Equation (5), the current interim target at age 64, denoted by $f_{64}(64)$, is 49.7414 units. Allowing for possible fund values at age 64 of $F_{64} = 0, 2, 4, \dots, 200$, the optimal asset allocation strategy adopted at each grid point is as shown in Figure 4.

The figure shows that if the current fund value is close to the interim target, the member adopts a conservative asset allocation strategy in order to protect or ‘bank’ the pension that will be received in a year’s time: the member becomes loss averse when the target has been achieved, since ‘a bird in the hand is worth two in the bush’. However, when the current fund value is below the interim target, the member rapidly increases the equity weighting and, if it is sufficiently below the target, the equity weighting is increased to 100%. This is because the member is risk seeking in the domain of losses: ‘I will do anything to get hold of a bird’. When the current fund value is above the interim target of 49.7414 units, but below 51 units, the member reduces the equity weighting to a negligible 2%: this is because, initially, the member is risk averse in the domain of gains. But the figure also shows that if the current fund value rises above 51 units (and, thus, the cushion between the current fund and the target fund increases to approximately 1.25 units), the member begins to raise the equity weighting again all the way to 100% if the fund value is above about 140 units. While always remaining risk averse in the domain of gains, the member becomes increasingly comfortable with accepting more equity risk (in exchange for the higher expected reward), since the risk of falling below the final target at age 65 becomes lower and lower: ‘I have eaten the bird in the hand and I now want the two in the bush as well!’.

Figure 5 shows the optimal asset allocation strategies at ages 44, 54 and again 64 under the assumption that the member’s income at each of these ages is 5 units. The current interim targets in the first two cases are $f_{44}(44) = 18.9561$ units and $f_{54}(54) = 29.5579$ units, respectively. The figure shows clearly that the bottom of the V-shaped asset allocation curve decreases with age, indicating that there is an age-related element to the optimal equity weighting, as in the case of deterministic lifestyling, *but only if* the actual fund at each age is close to the current interim target. The figure also shows that it is possible to have a very high equity weighting at any age if the fund value is either well above or well below the target for reasons given in Section 1.

Figure 4 – Optimal allocation to equities in year of age 64 to 65, assuming current labour income of 5 units



In order to determine the optimal asset allocation for each year of the plan, we generate 10,000 future scenarios over the period up to retirement, based on shocks in both risky asset returns, $Z_{1,x}$, and labour income, $Z_{2,x}$. For each year of age, since we can only derive the optimal asset allocation at each grid point in the 101-by-10 state variable matrix, bilinear interpolation is used to derive the optimal value for the control variables (i.e., the equity allocations, θ_x , at each age $x = 20, 21, \dots, 64$) for scenarios which lie within the state space but between the grid points.

Figure 5 – Optimal allocation to equities in years of age 44 to 45, 54 to 55 and 64 to 65, assuming current labour income of 5 units

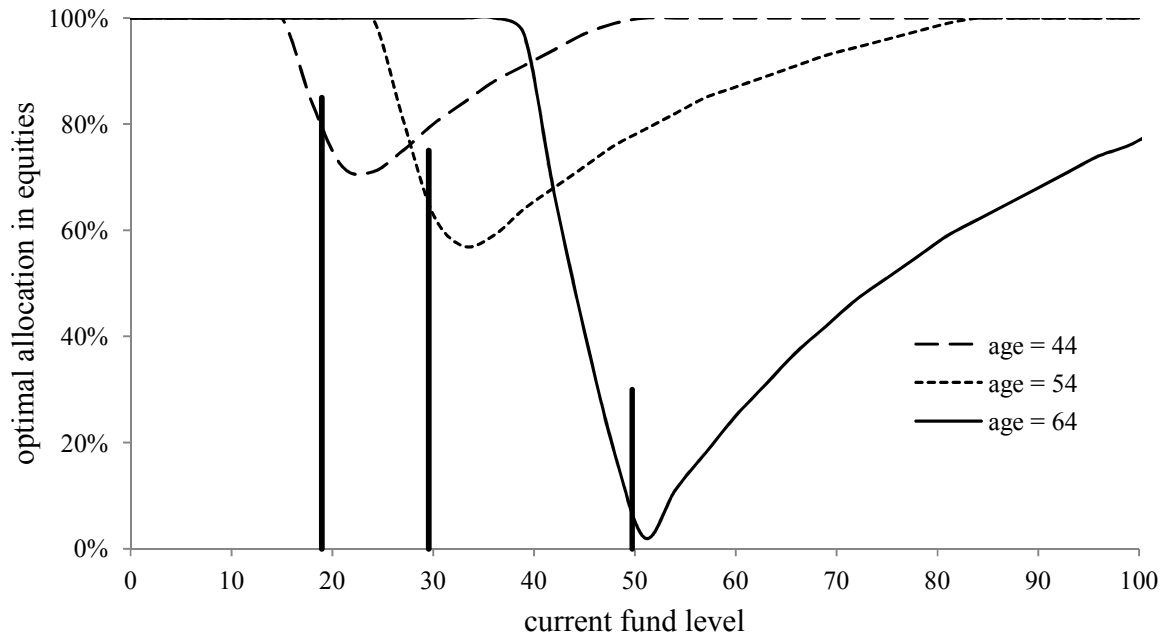
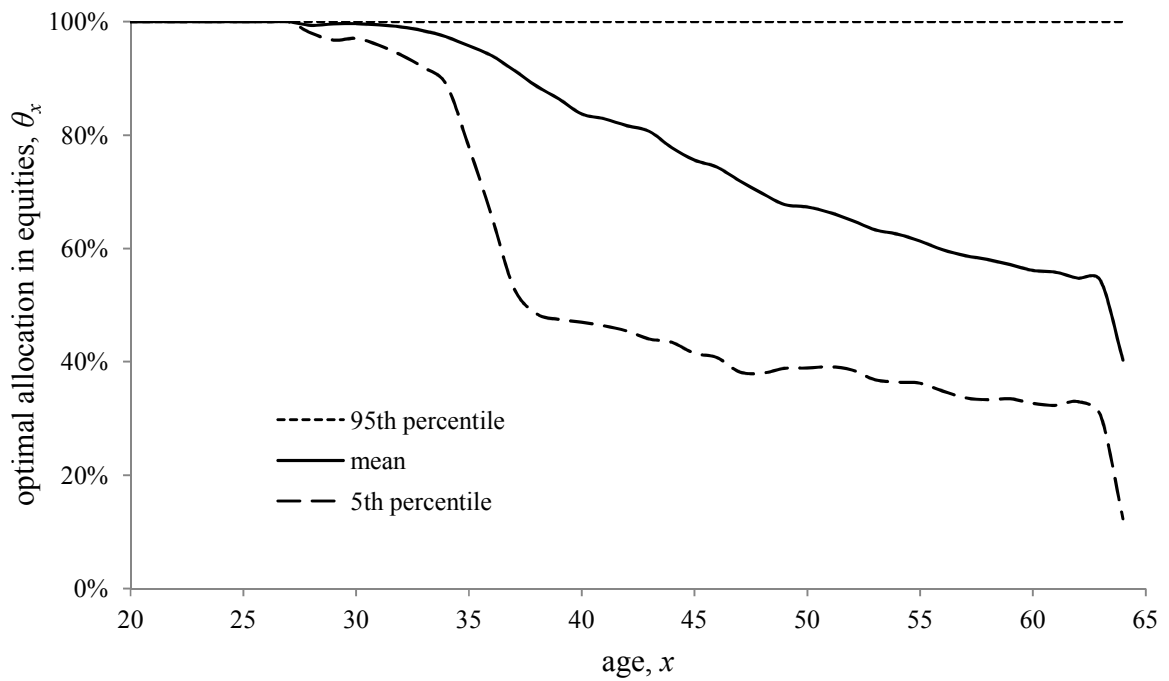


Figure 6 – Mean optimal allocation to equities under loss aversion



The mean optimal equity allocation at each age is shown in Figure 6 (along with the 5th and 95th percentiles in each case). Rather than depicting a conventional lifestyle strategy with its linear reduction in the equity weighting from 100% to 0% in the 5 to 10 years prior to retirement, the mean optimal equity allocation from the target-driven ‘threshold’ strategy involves a much more gradual move away from equities and, significantly, a far greater mean equity proportion at retirement (of around 40%, in this case).

We can compare the effect of adopting a threshold strategy with a conventional lifestyle strategy using Figure 7 which shows the distribution of the replacement ratio on retirement for each strategy. The range of potential replacement ratios at retirement from a conventional 10-year lifestyle strategy¹⁷ is very large, since there is no inherent mechanism for focusing on the desired target ratio. This makes it very difficult for a DC plan member to have any degree of certainty regarding the level of income after retirement. By contrast, the distribution of replacement ratios with a threshold strategy is more heavily concentrated around the target of two-thirds of final income. This is confirmed by Table 2 which shows that the dynamic target-driven strategy significantly improves the likelihood of achieving the target replacement ratio at retirement. The table also shows that under loss aversion there is a 75% chance of beating the target replacement ratio: we chose the contribution rate of 15% to achieve this result. If we had just wanted to beat the target replacement ratio on average, then we would have needed a contribution rate of only 13.2% (i.e., $15 \times 66.7 / 75.6$).

¹⁷ The member’s allocation to the equity fund will gradually be switched to a bond fund as follows: 100% in equities up to and including age 55, 90% at age 56, 80% at age 57, and so on to 10% at age 64.

Figure 7 – Distribution of the replacement ratio: loss aversion vs. 10-year lifestyling

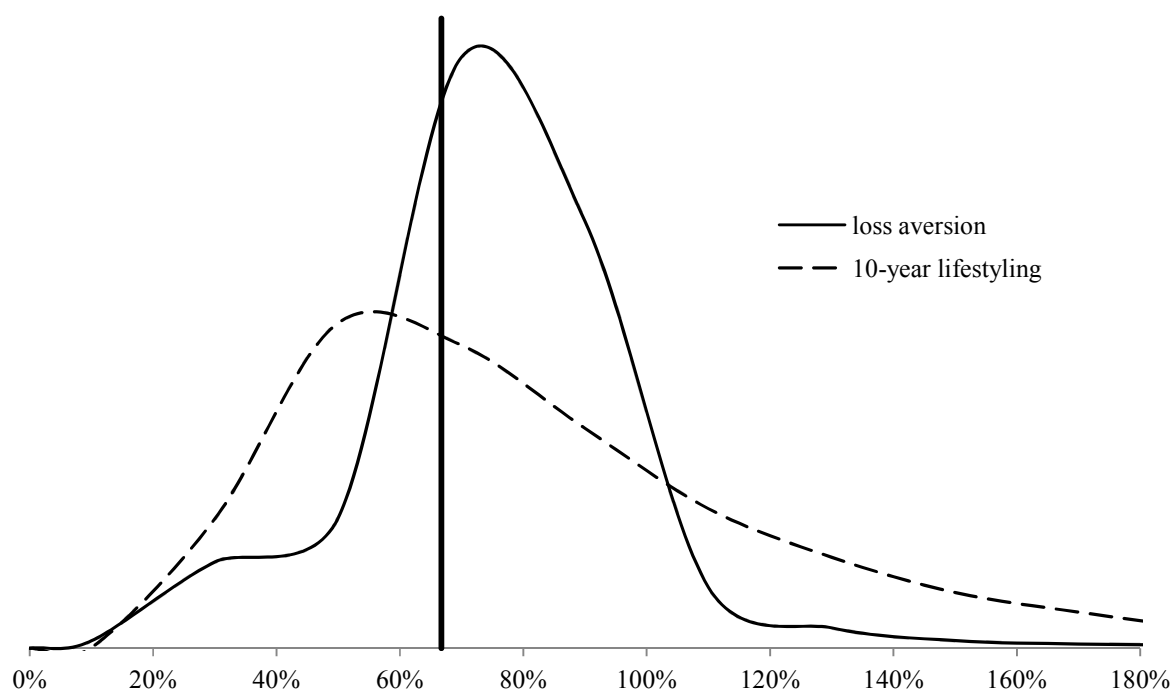


Table 2 – Replacement ratio under loss aversion and 10-year lifestyling

	Loss aversion	10-year lifestyling
Target replacement ratio	66.7%	66.7%
Mean replacement ratio	75.6%	83.0%
Standard deviation of replacement ratio	21.9%	40.1%
5 th percentile	35.2%	35.2%
25 th percentile	66.8%	53.6%
Median	76.9%	73.7%
75 th percentile	84.9%	103.0%
95 th percentile	106.5%	164.1%
Probability of achieving target	75.2%	57.8%
Expected shortfall	4.3%	7.3%

Note: Based on 10,000 simulations. Expected shortfall is defined as the mean of the difference between the target replacement ratio (of 66.7%) and the actual replacement ratio achieved, conditional on the actual replacement ratio being less than the target.

4.2 Sensitivity analysis

In this section, we conduct a sensitivity analysis of the baseline model. We change the loss aversion ratio, λ , the curvature parameters, ν_1 and ν_2 , the interim target weight, ω , and the discount rate used to determine the interim targets, r^* . We also compare what happens when only the final target matters and there is risk aversion rather than loss aversion with respect to this final target.

4.2.1 Loss aversion ratio

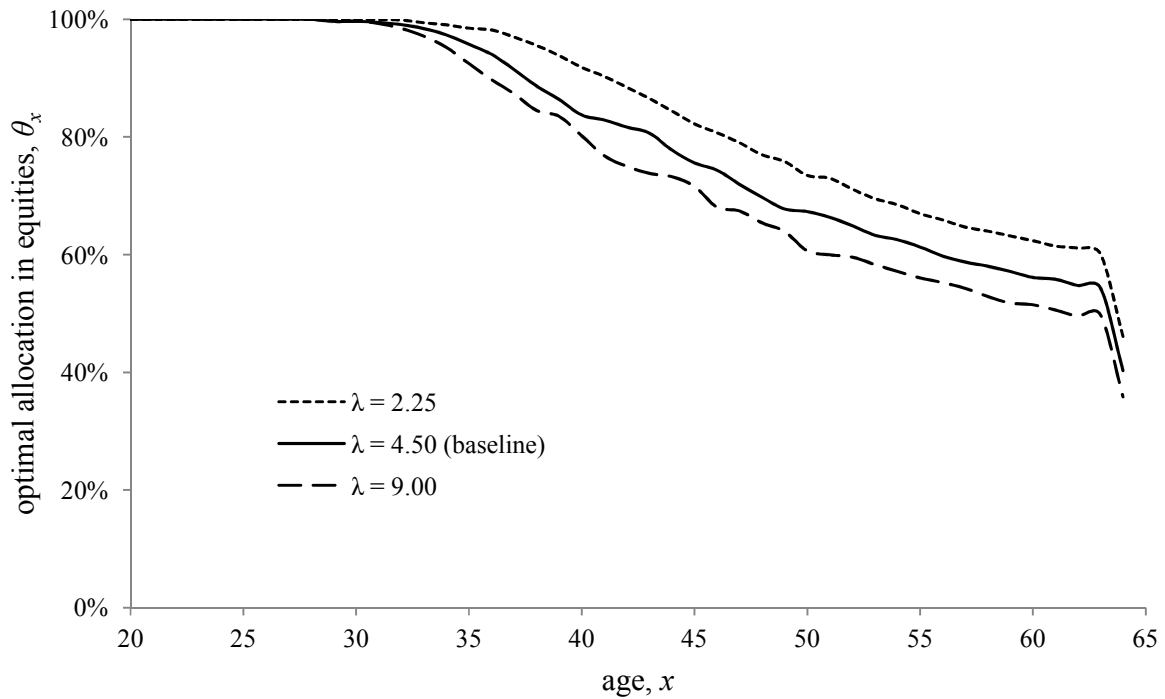
When the value of λ increases, the member becomes more loss averse. We consider the effect of increasing the loss aversion parameter from the baseline value of $\lambda = 4.50$ to $\lambda = 9.00$ and decreasing it to $\lambda = 2.25$ (the parameter value suggested by the original Tversky and Kahneman (1992) study).

Figure 8 shows the effect of these changes on the mean optimal allocation to equities. We can see that the higher the value of λ , the earlier the member begins to switch out of equities and the lower the proportion of the fund invested in equities on retirement. Thus, for a more loss averse member, there will be a lower allocation to risky assets at earlier ages, so the mean replacement ratio at retirement falls and the distribution of the replacement ratio at retirement is more concentrated around the target of two-thirds of final income as shown in Figure 9. This is confirmed by Table 3 which also shows that the expected shortfall (relative to the target) reduces as λ increases.

4.2.2 Curvature parameters

The parameters ν_1 and ν_2 control the curvature of the PT utility function in the domain of gains and losses, respectively. A lower value of ν_1 will increase the level of risk aversion when the fund is above target and, relative to the baseline case in Figure 6 above, the member will tend to switch back into equities more gradually in this domain (to protect the value of the accumulated fund). By contrast, a lower value of ν_2 will increase risk seeking when the fund is below target and, relative to Figure 6, the member will maintain a higher equity allocation in this domain (with the aim of eliminating the loss).

Figure 8 – Mean optimal allocation to equities: sensitivity analysis for changes in the loss aversion ratio, λ



In comparison with the baseline case, we will consider the following four cases:

- $v_1 = 0.22$ and $v_2 = 0.88$: the individual is significantly more risk averse with respect to gains; and
- $v_1 = 0.88$ and $v_2 = 0.88$: the individual is significantly more risk seeking with respect to gains;
- $v_1 = 0.44$ and $v_2 = 0.44$: the individual is significantly more risk seeking with respect to losses; and
- $v_1 = 0.44$ and $v_2 = 1.32$: the individual is significantly more risk averse with respect to losses.

Figure 9 – Effect on the replacement ratio of changing the loss aversion ratio, λ

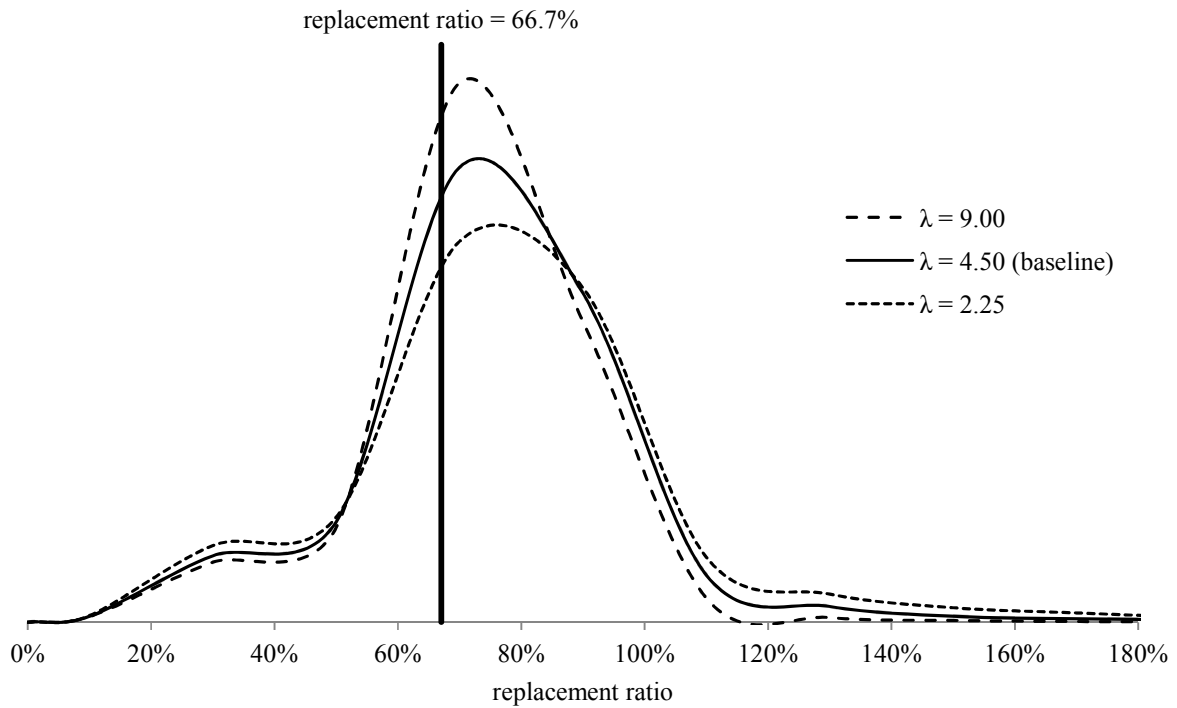


Table 3 – Effect on the replacement ratio of changing the loss aversion ratio, λ

Loss aversion ratio, λ	2.25	4.50 (baseline)	9.00
Target replacement ratio	66.7%	66.7%	66.7%
Mean replacement ratio	79.0%	75.6%	73.0%
Standard deviation of replacement ratio	28.3%	21.9%	17.3%
5 th percentile	33.6%	35.2%	36.5%
25 th percentile	66.9%	66.8%	66.4%
Median replacement ratio	78.4%	76.9%	75.4%
75 th percentile	88.0%	84.9%	82.2%
95 th percentile	126.5%	106.5%	95.3%
Probability of achieving target	75.3%	75.2%	74.3%
Expected shortfall	4.7%	4.3%	4.0%

Figure 10 – Mean optimal allocation to equities: sensitivity analysis for changes in the curvature parameters, v_1 and v_2

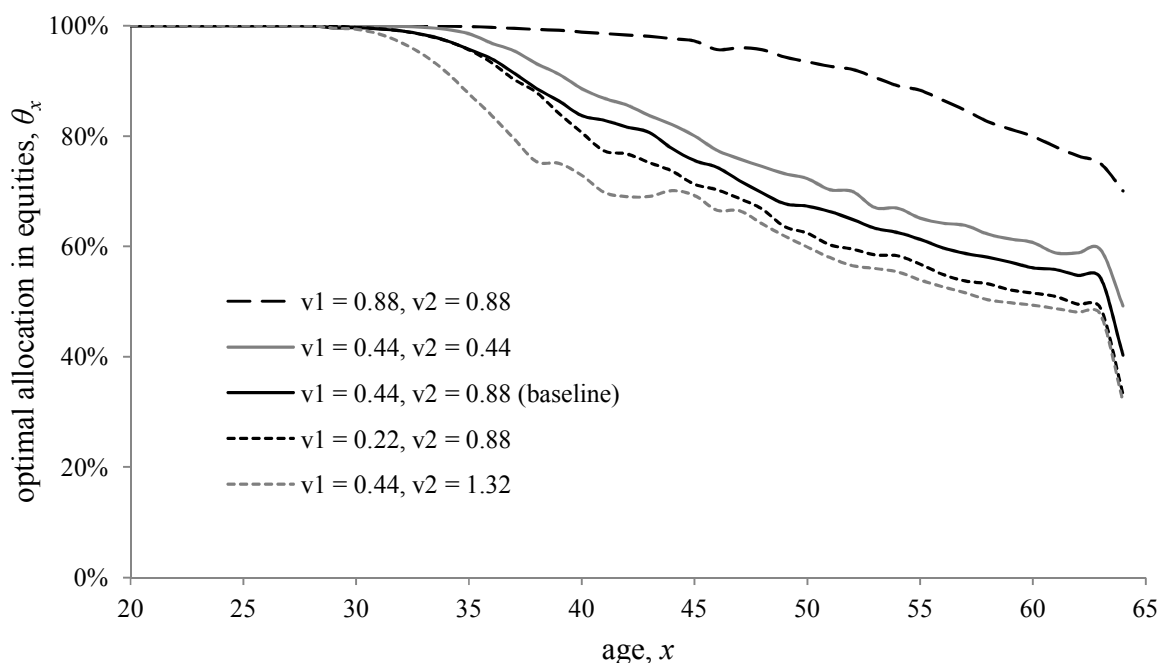


Figure 10 shows that, when risk aversion in the domain of gains is reduced (e.g., v_1 is increased from 0.44 to 0.88), the member can be expected to maintain a significantly higher equity allocation at each age. This is because the member is less inclined to ‘bank’ any previous investment gains by switching to bonds when above the current interim target. Similarly, when risk aversion in the domain of losses is reduced (e.g., v_2 is reduced from 0.88 to 0.44), the member can also be expected to maintain a higher equity allocation at each age. However, the effect is much less pronounced in this case, as the member will be more comfortable taking on extra risk when below the interim target (in the hope of making good the loss), but, should this increased investment risk taking prove successful, will then tend to ‘bank’ any gains more quickly when above the target. Not surprisingly, these effects are reversed when we increase the level of risk aversion in either the domain of gains (e.g., to $v_1 = 0.22$) or losses (e.g., to $v_2 = 1.32$).

Table 4 – Effect on replacement ratio of changing the curvature parameters, v_1 and v_2

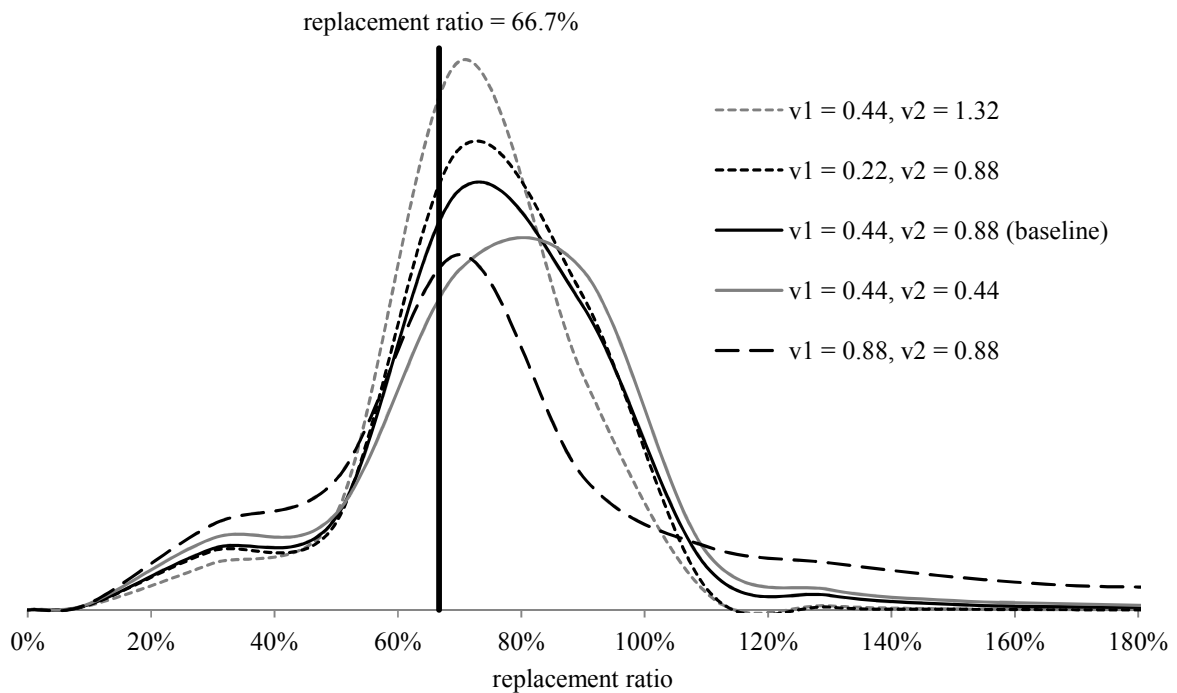
	0.44	0.22	0.88	0.44	0.44
Curvature parameter (gains), v_1					
Curvature parameter (losses), v_2	0.88	0.88	0.88	0.44	1.32
	(baseline)				
Target replacement ratio	66.7%	66.7%	66.7%	66.7%	66.7%
Mean replacement ratio	75.6%	73.2%	89.0%	77.8%	71.9%
Standard deviation of replacement ratio	21.9%	17.1%	50.9%	25.3%	16.0%
5 th percentile	35.2%	35.5%	31.6%	33.3%	39.1%
25 th percentile	66.8%	66.7%	61.4%	68.9%	65.0%
Median	76.9%	76.2%	74.2%	79.0%	73.6%
75 th percentile	84.9%	83.1%	100.6%	87.2%	80.6%
95 th percentile	106.5%	95.0%	203.7%	116.6%	92.9%
Probability of achieving target	75.2%	75.1%	67.5%	77.4%	71.1%
Expected shortfall	4.3%	4.1%	5.9%	4.8%	3.8%

Table 4 and Figure 11 show that, when risk aversion is increased in either the domain of gains (e.g., v_1 is reduced from 0.44 to 0.22) or losses (e.g., v_2 is increased from 0.88 to 1.32), the uncertainty regarding the final replacement ratio is reduced and, as a result, the expected shortfall is reduced. However, since the member is less inclined to take on extra risk (either when well above the current target or to make good a current shortfall), both the mean replacement ratio and the probability of achieving the target replacement ratio are reduced (and significantly so when increasing the degree of risk aversion in the domain of losses).

Conversely, reducing risk aversion in the domain of gains (e.g., when v_1 is increased from 0.44 to 0.88) leads to a very significant increase in the mean replacement ratio (as would be expected given the higher mean equity allocation noted above). However, as a result of this increased investment risk, the uncertainty regarding the final replacement ratio achieved is also significantly increased, leading to a much lower probability of achieving the desired target on retirement (67.5% compared with 75.2% for the baseline) and a higher expected shortfall.

If the risk aversion in the domain of losses is reduced (e.g., v_2 is reduced from 0.88 to 0.44), then the individual is more willing to take on extra investment risk to make good a current shortfall (relative to the current interim target). As a result, the mean replacement ratio and probability of achieving the desired target replacement ratio are both increased relative to the baseline. However, the downside is that, when this increased investment risk is unsuccessful in producing the desired extra return, the effect is more severe and the expected shortfall increases.

Figure 11 – Effect on the replacement ratio of changing curvature parameters, v_1 and v_2



4.2.3 Interim target weight

Having demonstrated the benefits of a target-driven approach, it is important to examine the effect of the weight applied to the interim targets (relative to that on the final target at retirement). This is because there is a real danger that, within such a framework, the member develops a myopic risk attitude, focusing too much on short-term performance (relative to the interim targets) to the detriment of long-term performance (and the ultimate goal of achieving the target

replacement ratio at retirement). In the baseline case above, the weight applied to the interim targets was $\omega = 0.5$ (i.e., the interim targets were given a weight of 50% of the final target).

Figure 12 – Mean optimal allocation to equities: sensitivity analysis for changes in the weight applied to the interim targets, ω

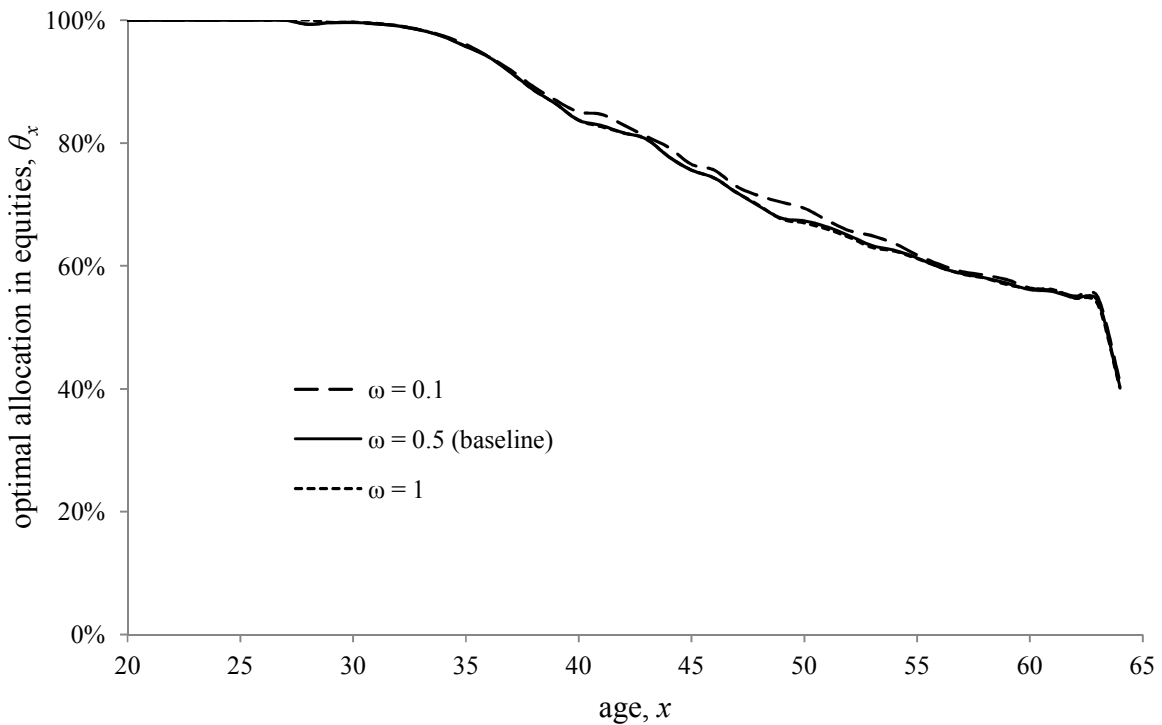


Figure 12 shows the mean optimal allocation to equities for different values of the weight applied to the interim targets, ω . It can be seen that the optimal asset allocation strategy is not significantly affected by changing the weight of the interim target. However, when the weight applied to the interim target is increased, the member tends to adopt a very slightly lower equity allocation, since the member becomes more concerned about short-term volatilities that may cause the fund to under-perform the interim targets.

Table 5 – Effect on the replacement ratio of changing the weight applied to the interim targets, ω

Interim target weight, ω	1.0	0.5 (baseline)	0.1
Target replacement ratio	66.7%	66.7%	66.7%
Mean replacement ratio	75.6%	75.6%	76.1%
Standard deviation of replacement ratio	21.9%	21.9%	22.4%
5 th percentile	35.2%	35.2%	35.0%
25 th percentile	66.7%	66.8%	66.7%
Median	76.8%	76.9%	77.4%
75 th percentile	84.7%	84.9%	85.7%
95 th percentile	106.5%	106.5%	108.1%
Probability of achieving target	75.1%	75.2%	75.1%
Expected shortfall	4.2%	4.3%	4.3%

Figure 13 – Effect on the replacement ratio of changing the weight applied to the interim targets, ω

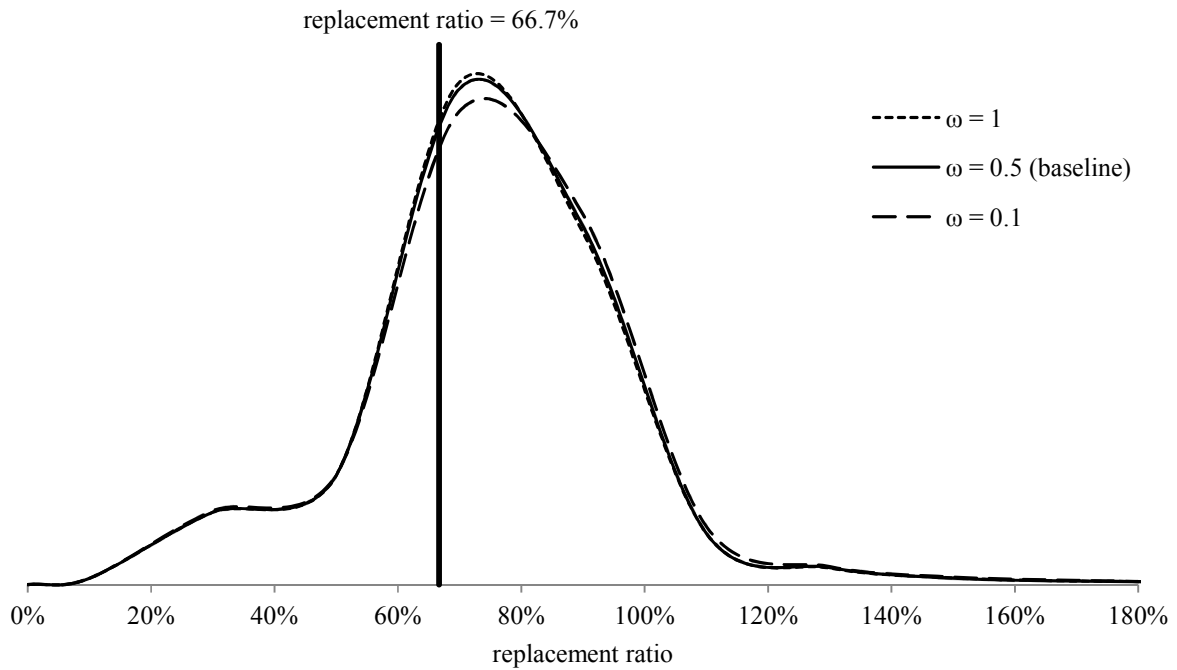


Table 5 and Figure 13 show that a myopic individual, who puts a greater relative weight on achieving short-term interim targets, has a slightly lower expected replacement ratio on retirement. The advantage of giving a higher weight to the interim targets, though, is a corresponding slight reduction in the volatility of the replacement ratio as a result of the more conservative asset allocation strategy adopted (e.g., with $\omega=1$, the inter-quartile range is $84.7\% - 66.7\% = 18.0\%$, whereas, with $\omega=0.1$, the inter-quartile range increases marginally to $85.7\% - 66.7\% = 19.0\%$). However, the distribution of the replacement ratio is largely insensitive to the weight applied to the interim target. Given that this is likely to be a difficult parameter for an individual to quantify, this robustness is welcomed.

4.2.4 Discount rate used to determine interim targets

The baseline case uses $r^* = r + 0.011 = 0.031$ and we now examine the effect of changing the credit spread as follows:

- $r^* = r = 0.02$: i.e., the credit spread on long-dated corporate bonds falls to 0%; and
- $r^* = r + 0.02 = 0.04$: i.e., the credit spread rises to 2%.

Figure 14 shows the mean optimal allocation to equities for the different values of r^* . Reassuringly, the optimal asset allocation is largely unaffected by changing the discount rate, particularly from about age 45 or so. Table 6 and Figure 15 show that the distribution of the replacement ratio is also fairly insensitive to the choice of the discount rate used to calculate the interim targets, although the fact that the interim target is higher (particularly at younger ages) when a lower discount rate is used, means that the switch from equities tends to begin later (as can be seen in Figure 14). As shown in Table 6, this leads to a slight increase in the level of risk accepted (as measured by size of the expected shortfall).

Figure 14 – Mean optimal allocation to equities: sensitivity analysis for changes in the discount rate for interim targets, r^*

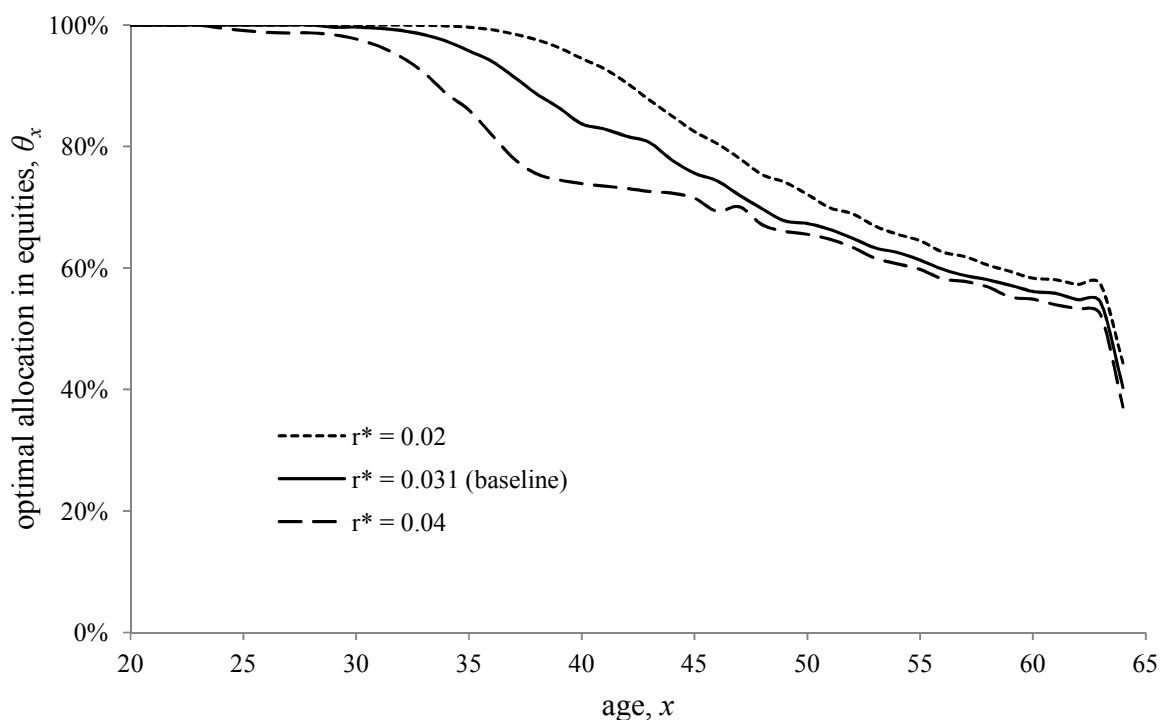
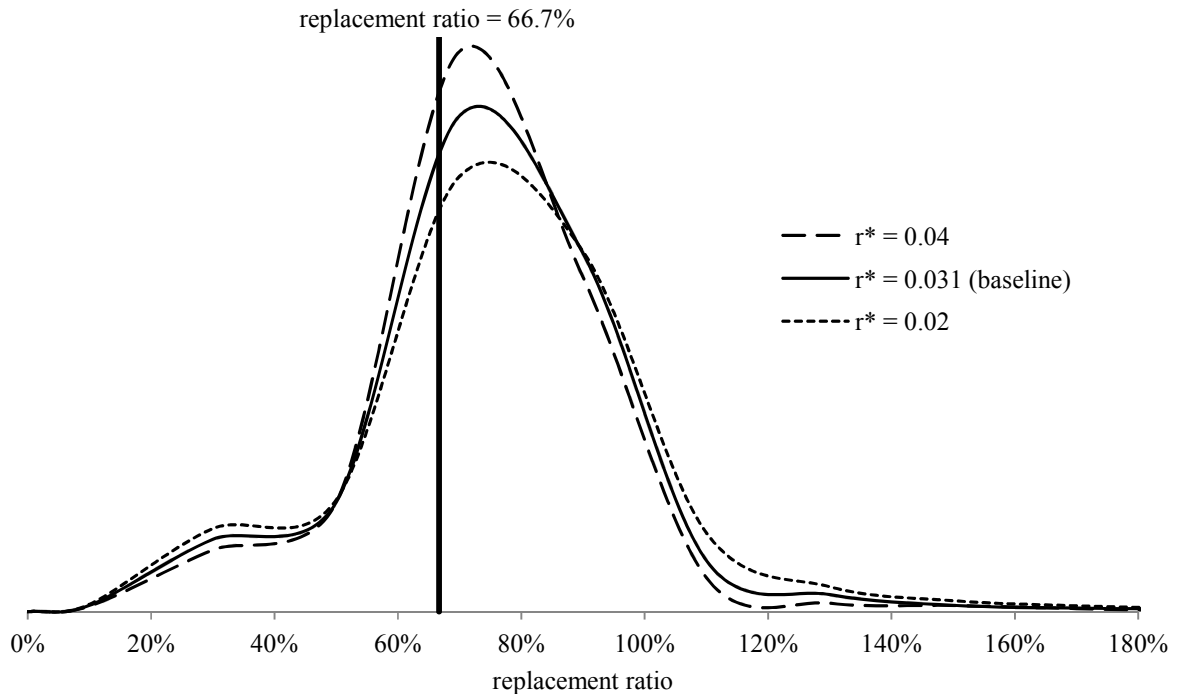


Table 6 – Effect on the replacement ratio of changing the discount rate used to determine the interim targets, r^*

Discount rate, r^*	2.0%	3.1% (baseline)	4.0%
Target replacement ratio	66.7%	66.7%	66.7%
Mean replacement ratio	77.7%	75.6%	73.7%
Standard deviation of replacement ratio	25.1%	21.9%	19.4%
5 th percentile	33.2%	35.2%	37.6%
25 th percentile	66.5%	66.8%	66.0%
Median	78.9%	76.9%	74.8%
75 th percentile	88.4%	84.9%	82.2%
95 th percentile	116.0%	106.5%	99.2%
Probability of achieving target	74.8%	75.2%	73.3%
Expected shortfall	4.7%	4.3%	4.0%

Figure 15 – Effect on the replacement ratio of changing the discount rate used to determine the interim targets, r^*



4.2.5 Comparison with the optimal investment strategy under risk aversion

Our final sensitivity test compares the LA framework with that where the member is assumed to focus only on the final target at retirement and is subject to risk aversion rather than loss aversion.

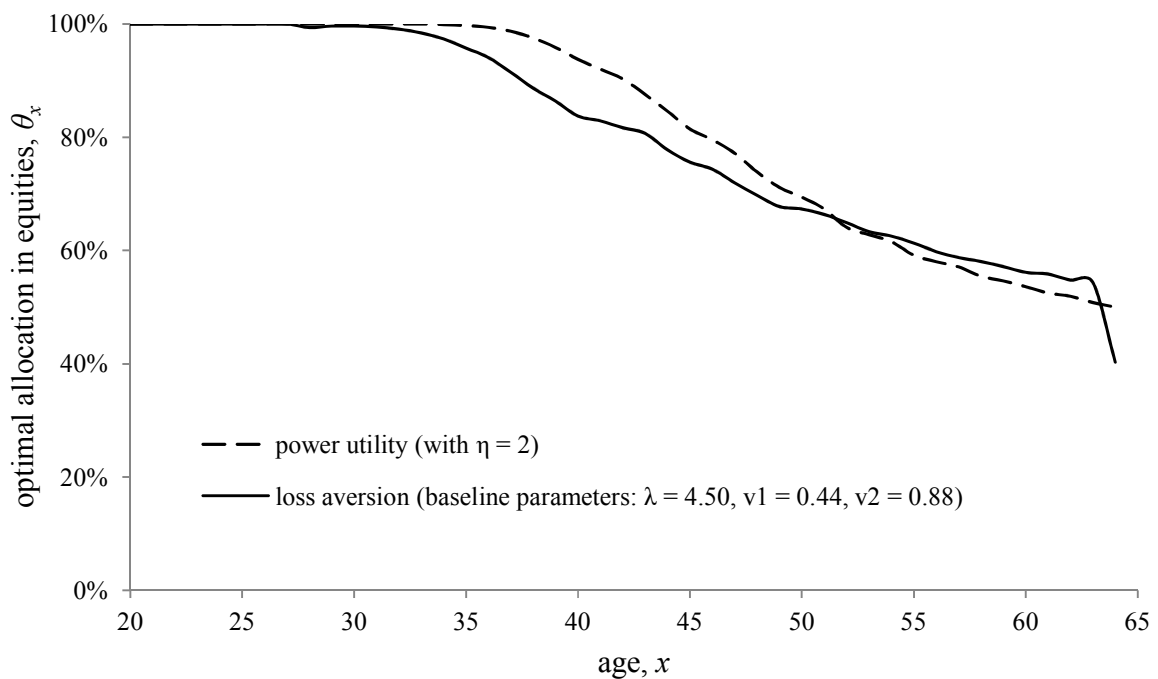
Boulier et al. (2001) and Cairns et al. (2006), for example, have proposed such a model where utility is represented by a power function in the final fund value (where η is the coefficient of relative risk aversion):

$$U_{65}(F_{65}) = \frac{(F_{65})^{1-\eta}}{1-\eta} \quad (10)$$

An inspection of Equation (6) shows that this is a special case of the PT utility function with $\omega = 0$ for $x = 20, \dots, 64$ and, for $x = 65$, we have $v_1 = 1 - \eta$, and $F_{65} > f_{65}(65) = 0$ (i.e., no explicit final target).

Figure 16 shows the mean optimal allocation to equities at each age for both the power utility and the LA frameworks. We can see that, under loss aversion, the switch out of equities begins at a younger age (as the member seeks to protect past investment gains relative to the target). By contrast, with a power utility framework, the member is focused only on maximising the expected utility of the terminal fund value and so maintains a higher equity allocation for longer, although the equity allocations on retirement are broadly similar under the two frameworks (by design).¹⁸

Figure 16 – Mean optimal allocation to equities: loss aversion vs. power utility



The effect of this difference on the distribution of the replacement ratio can then be seen in Table 7 and Figure 17. The higher equity allocation using power utility leads to a higher mean

¹⁸ See footnote 15.

replacement ratio, but also to a very significant increase in the uncertainty of the replacement ratio actually realised. Thus, such an approach increases the risk faced by the member, as both the probability of failing to reach the target replacement ratio on retirement and the expected shortfall (relative to this target) are significantly increased.^{19, 20}

Table 7 – Replacement ratio under loss aversion and power utility

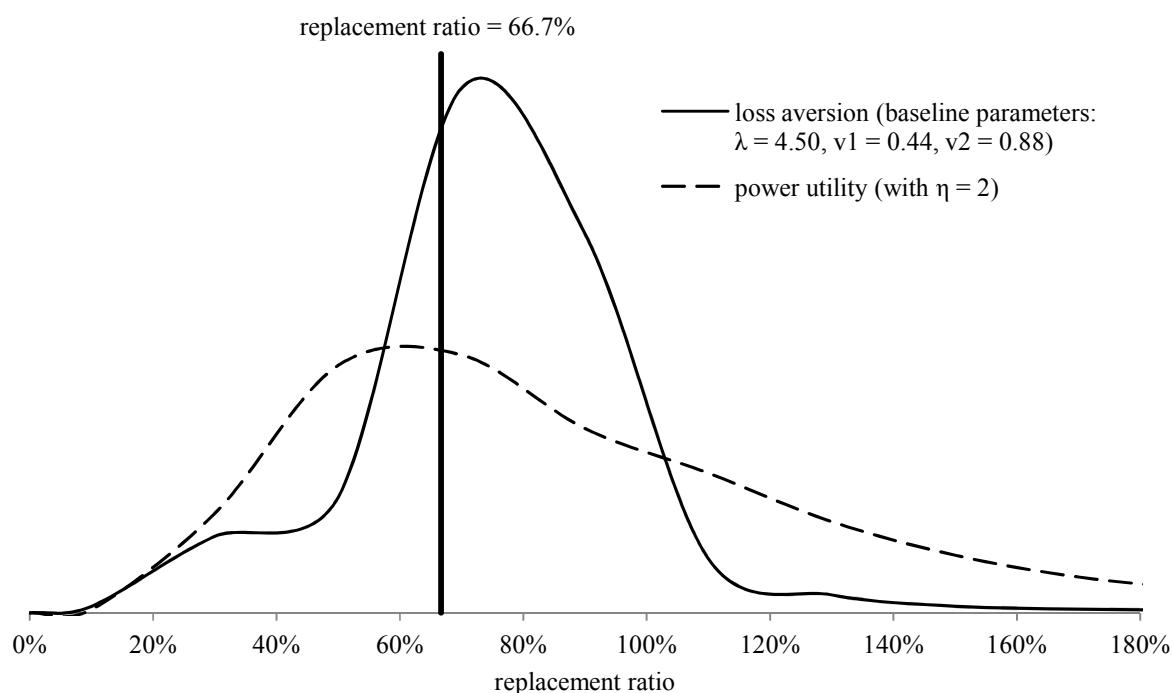
	Loss aversion (baseline)	Power utility
Target replacement ratio	66.7%	n/a
Mean replacement ratio	75.6%	91.2%
Standard deviation of replacement ratio	21.9%	50.4%
5 th percentile	35.2%	35.0%
25 th percentile	66.8%	55.8%
Median	76.9%	78.8%
75 th percentile	84.9%	113.2%
95 th percentile	106.5%	189.9%
Probability of achieving target	75.2%	62.7%*
Expected shortfall	4.3%	6.6%

Note: * While, there is no explicit target replacement ratio under risk aversion, we report the probability of achieving the same target replacement ratio as under loss aversion, namely 66.7%.

¹⁹ While there is no explicit target replacement ratio on retirement with power utility, this does not prevent us specifying an implicit target.

²⁰ The general similarity of shape and location of the two distributions of final fund values in Figure 17 is due to the common contribution rate of 15% which is sufficient to beat the target replacement ratio of 66.7% under both strategies. In the case of our LA parameterisation, the probability of beating the target is 75% (see Table 7). However, the target-driven strategy will have a higher probability of beating the target than one which simply maximises the expected utility of terminal pension wealth.

Figure 17 – Distribution of the replacement ratio: loss aversion vs. power utility



5 Conclusions

We have applied some of the most powerful lessons of behavioural finance to solve the optimal investment strategy in a defined contribution pension plan. The key lesson is that loss aversion with respect to a series of reference wealth levels in the form of increasing target fund values over the life of the plan might well provide a better representation of a typical plan member's attitude to risk taking than risk aversion in respect of the terminal fund value at retirement, the most commonly used approach from expected utility theory. The most popular investment strategy by practitioners, conventional lifestyling, does not even take the plan members' attitude to risk or loss into account. Nor does it take into account the stochastic nature of asset returns and the plan member's labour income.²¹

²¹ The stochastic lifestyling model of Cairns et al. (2006) does rectify the latter deficiencies, however.

Loss aversion leads to a new type of target-driven approach to deriving the dynamic optimal asset allocation. The key findings from using this approach are as follows:

- The optimal investment strategy under loss aversion is the ‘threshold’ strategy considered in Blake et al. (2001). With this strategy, the weight in equities is increased if the accumulating fund is below the interim target (since plan members are risk seeking in the domain of losses) and is decreased if the fund is above target (since plan members are risk averse in the domain of gains), unless the fund is very much above target, in which case plan members are comfortable with assuming more equity risk again and accordingly adopt a ‘portfolio insurance’ investment strategy. However, as the retirement date approaches and assuming the fund is on target, the overall equity weight begins to fall and the value of the fund is ‘banked’ by switching to lower risk investments, such as bonds. The strategy is highly focused on achieving a target replacement ratio at retirement.
- The switch to a more conservative asset allocation strategy is implemented at lower current fund values (relative to target) and at a lower age, the higher is the loss aversion ratio, λ . Although the mean replacement ratio falls as a consequence, the expected shortfall from the target decreases.
- The effect of higher risk aversion in the domain of gains (measured by a lower ν_1 in the PT utility function) leads, unsurprisingly, to an earlier switch out of equities and a lower mean replacement ratio, but also to a lower expected shortfall. The effect of greater risk seeking in the domain of losses (measured by a lower ν_2 in the PT utility function) leads to a later switch out of equities, a higher mean replacement ratio, a higher probability of achieving the target, but also a higher expected shortfall.
- The greater the weight attached to the interim targets (relative to the final target), the less aggressive the investment strategy adopted is, although the overall impact is fairly marginal. In practice, the key factors influencing the relative significance of the interim targets are likely to be the frequency and quality of the fund performance information given to the members.

- A discount rate is needed to find the value of the interim targets. Given the controversy surrounding the choice of discount rate in valuing pension liabilities,²² it is comforting to discover that the level of the discount rate has very little impact on the optimal asset allocation in a loss aversion framework.
- Compared with the standard framework of risk aversion, loss-averse plan members are committed to achieving interim and final target fund levels and, accordingly, adopt a more conservative asset allocation strategy. Although this leads to a lower mean replacement ratio at retirement, there is a greater likelihood of achieving the desired target replacement ratio and a lower expected shortfall.
- Compared with a conventional lifestyle investment approach, the optimal target-driven investment strategy significantly increases the likelihood of achieving the chosen target, thereby providing a much greater degree of certainty in retirement planning.

The risks inherent in the conventional lifestyle strategy appear to be much higher than generally understood. Thus, for DC plan members who seek greater certainty in retirement planning, the investment strategy adopted over time needs to be far more focused on achieving the specified target replacement ratio. Setting the investment strategy in a DC pension plan within the framework of loss aversion therefore has much to recommend it. However, the framework is not easy to implement since it requires the solution of a nonlinear dynamic programming problem whenever there is new information about the key state variables. Nevertheless, in practice, it should be possible to tabulate the optimal asset allocation in terms of member profile characteristics (such as age and occupation) and values of the key state variables (e.g., interim fund level and current labour income). Financial advisers would then be able to advise on the appropriate investment strategy for the coming year.

²² For example, see Blake et al. (2008).

References

- Benartzi, S., and Thaler, R. (1995). 'Myopic Loss Aversion and the Equity Premium Puzzle', *Quarterly Journal of Economics*, 110, 73-92.
- Berkelaar, A. B., Kouwenberg, R., and Post, T. (2004). 'Optimal Portfolio Choice under Loss Aversion', *Review of Economics and Statistics*, 86, 973-87.
- Blake, D., Cairns, A., and Dowd, K. (2001). 'PensionMetrics: Stochastic Pension Plan Design and Value at Risk During the Accumulation Phase', *Insurance: Mathematics and Economics*, 29, 187-215.
- Blake, D., Cairns, A., and Dowd, K. (2007). 'The Impact of Occupation and Gender on Pensions from Defined Contribution Plans', *Geneva Papers on Risk & Insurance*, 32, 458-82
- Blake, D., Khorasane, Z., Pickles, J., and Tyrrell, D. (2008) *An Unreal Number: How Company Pension Accounting Fosters an Illusion of Certainty*, Pensions Institute and Institute of Chartered Accountants in England and Wales.
- Boulier, J., Huang, S., and Taillard, G. (2001). 'Optimal Management under Stochastic Interest Rates: The Case of a Protected Defined Contribution Pension Fund', *Insurance: Mathematics and Economics*, 28, 173-189.
- Cairns, A., Blake, D., and Dowd, K. (2006). 'Stochastic Lifestyling: Optimal Dynamic Asset Allocation for Defined Contribution Pension Plans', *Journal of Economic Dynamics and Control*, 30, 843-877.
- Campbell, J., and Viceira, L. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*, Oxford University Press, Oxford.
- Fama, Eugene, and French, Kenneth (2002). 'The Equity Premium', *Journal of Finance*, 57, 637-659.
- Gerrard, R., Haberman, S., and Vigna, E. (2004). 'Optimal Investment Choices Post-retirement in a Defined Contribution Pension Scheme', *Insurance: Mathematics and Economics*, 35, 321-342.

- Gneezy, U., and Potters, J. (1997), 'An Experiment on Risk Taking and Evaluation Periods', *Quarterly Journal of Economics*, 112, 632-45.
- Gomes, F. (2005). 'Portfolio Choice and Trading Volume with Loss-Averse Investors', *Journal of Business*, 78, 675-706.
- Gomes, F., and Michaelides, A. (2005). 'Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence', *Journal of Finance*, 60, 869-904.
- Haberman, S., and Vigna, E. (2002). 'Optimal Investment Strategies and Risk Measures in Defined Contribution Pension Schemes', *Insurance: Mathematics and Economics*, 31, 35-69.
- Kahneman, D., and Tversky, A. (1979). 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica*, 47, 263-91.
- Madrian, B., and Shea, D. (2001). 'The Power of Suggestion: Inertia in 401(k) Participation and Savings Behaviour', *Quarterly Journal of Economics*, 116, 1149-1187.
- Mehra, R., and Prescott, E. (1985). 'The Equity Premium: A Puzzle', *Journal of Monetary Economics*, 15, 145-161.
- Mitchell, O., and Utkus, S. (2004) *Pension Design and Structure: New Lessons from Behavioural Finance*, Oxford University Press, Oxford.
- Rabin, M. and Thaler, R. (2001). 'Anomalies: Risk Aversion', *Journal of Economic Perspectives*, 219-232.
- Samuelson, P.A. (1989). 'A Case at Last for Age-phased Reduction in Equity', *Proceedings of the National Academy of Science*, 86, 9048-9051.
- Thaler, R., Tversky, A., Kahneman, D., and Schwartz, A. (1997). 'The Effect of Myopia and Loss Aversion On Risk Taking: An Experimental Test', *Quarterly Journal of Economics*, 112, 647-661.
- Tversky A., and Kahneman, D. (1992). 'Advances in Prospect Theory: Cumulative Representation of Uncertainty', *Journal of Risk and Uncertainty*, 5, 297-323.
- Vigna, E., and Haberman, S. (2001). 'Optimal Investment Strategy for Defined Contribution Pension Schemes', *Insurance: Mathematics and Economics*, 28, 233-262.

Appendix - Numerical solution for the dynamic programming and integration process

The optimisation problem at age x is:

$$\begin{aligned}\max_{\theta_x} E_x [V_x(F_x, Y_x)] &= \max_{\theta_x} E_x \left[\left(\sum_{s=0}^{(65-x)-1} \beta^s \omega U_{x+s}(F_{x+s}) \right) + \beta^{65-x} U_{65}(F_{65}) \right] \\ &= \max_{\theta_x} \left(\omega U_x(F_x) + \beta E_x [V_{x+1}(F_{x+1}, Y_{x+1})] \right)\end{aligned}$$

subject to:

- $Y_x = Y_{x-1} \times \exp(I_x)$, where $I_x = r_I + \frac{S_x - S_{x-1}}{S_{x-1}} + \sigma_1 Z_{1,x} + \sigma_2 Z_{2,x}$;
- $F_x = (F_{x-1} + \pi Y_{x-1}) \times [1 + r + \theta_{x-1} (\mu + \sigma_1 Z_{1,x})] \geq 0$; and
- $0 \leq \theta_x \leq 1$ for $x = 20, 21, 22, \dots, 64$.

Thus, at age 64, the value function is given by:

$$\max_{\theta_{64}} E_{64} [V_{64}(F_{64}, Y_{64})] = \max_{\theta_{64}} \left(\omega U_{64}(F_{64}) + \beta E_{64} [V_{65}(F_{65}, Y_{65})] \right)$$

To avoid choosing local optima, we discretise the control variable, θ_{64} (i.e., asset allocation to the risky asset in the year prior to retirement), into 20 equally spaced intervals (with corresponding grid points) and optimise across these grids using a standard grid search. As an important step in stochastic dynamic programming, we need to discretise the state space and shocks. Wealth and labour income level are discretised into 100 and 9 evenly-spaced intervals respectively (using grid points of 0, 2, 4, ..., 200 for the wealth level and 1, 2, ..., 10 for the income level), so that we can calculate the optimal control variable and value function for each grid point (e.g. as shown in Figure 4, for an income level of 5 units).

To solve the non-linear expectation component of the above equation at age $x=64$, namely $E_{64} [V_{65}(F_{65}, Y_{65})] = E_{64} [U_{65}(F_{65})]$, the Gauss-Hermite quadrature method is used, given the assumption of normally distributed equity returns and income growth rates. The idea is to approximate the value function by using several significant nodes of the distribution and their

relevant weights. The equity return shock in the final year, $Z_{1,64}$, and the labour income shock in the final year, $Z_{2,64}$, are both discretised into 9 nodes, by substituting $\{\sqrt{2} \times Z_{1,64}(m) : m = 1, 2, \dots, 9\}$ and $\{\sqrt{2} \times Z_{2,64}(n) : n = 1, 2, \dots, 9\}$, respectively.

Thus, we have (where π is the mathematical constant):

$$\begin{aligned} & E_{64} [U_{65}(F_{65})] \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U_{65}(F_{65}, Y_{65}) \phi(Z_{1,64}, Z_{2,64}) dZ_{1,64} dZ_{2,64} \\ &\approx \pi^{-1} \times \sum_{m=1}^9 \sum_{n=1}^9 w_{Z_{1,64}(m)} w_{Z_{2,64}(n)} \left[U_{65}(\sqrt{2} \times Z_{1,64}(m), \sqrt{2} \times Z_{2,64}(n); F_{65}, Y_{65}) \right] \end{aligned}$$

where:

- $\phi(Z_{1,64}, Z_{2,64})$ is the joint probability density function for the standard Normal random variables $Z_{1,64}$ and $Z_{2,64}$,
- $\{w_{Z_{1,64}(m)} : m = 1, 2, \dots, 9\}$ and $\{w_{Z_{2,64}(n)} : n = 1, 2, \dots, 9\}$ are the Gauss-Hermite quadrature weights, and
- $\{Z_{1,64}(m) : m = 1, 2, \dots, 9\}$ and $\{Z_{2,64}(n) : n = 1, 2, \dots, 9\}$ are the Gauss-Hermite quadrature nodes for $Z_{1,64}$ and $Z_{2,64}$, respectively.

The advantage of this definition of nodes is that the state variables can be computed more precisely and run time is significantly reduced. However, because we have quite fine grids for the control variable and much coarser grids for the shocks, we may have some state variable values not located exactly at the grid points in the previous time period. In this case, bilinear interpolation is employed to approximate the value function.

At each age x prior to retirement at age 65, we compute the value function and the corresponding optimal variables at each grid point. Substituting these values in the Bellman equation below, we can solve the maximisation problem of the previous period.

Thus, at age $x = 63, 62, \dots, 20$, the Bellman equation for this problem is given by:

$$\begin{aligned}
V_x(F_x, Y_x) &= \max_{\theta_x} E_x \left[\left(\sum_{s=0}^{(65-x)-1} \beta^s \omega U_{x+s}(F_{x+s}) \right) + \beta^{65-t} U_{65}(F_{65}) \right] \\
&= \max_{\theta_x} \left(\omega U_x(F_x) + \beta E_x [V_{x+1}(F_{x+1}, Y_{x+1})] \right)
\end{aligned}$$

We set up the same 100-by-9 intervals for state variables and bilinear interpolation is employed to approximate the value function. After we obtain the optimal control variable at each grid point, we then substitute this into the Bellman equation and use it to solve the maximisation problem at the previous time period. This process is then iterated backwards until the initial age 20. The computations were performed in MATLAB.²³

²³ <http://www.mathworks.com/products/matlab/>.