Investing for retirement

Citation for published version:
Donnelly, C, Khemka, G & Lim, W 2022, 'Investing for retirement: Terminal wealth constraints or a desired wealth target?', *European Financial Management*, vol. 28, no. 5, pp. 1283-1307. https://doi.org/10.1111/eufm.12351

Digital Object Identifier (DOI):
10.1111/eufm.12351

Link:
Link to publication record in Heriot-Watt Research Portal

Document Version:
Publisher's PDF, also known as Version of record

Published In:
European Financial Management

DOI: 10.1111/eufm.12351

EUROPEAN FINANCIAL MANAGEMENT WILEY

ORIGINAL ARTICLE

Investing for retirement: Terminal wealth constraints or a desired wealth target?

Catherine Donnelly[1] | Gaurav Khemka[2,3] | William Lim[2]

[1]Department of Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh, UK

[2]Research School of Finance, Actuarial Studies and Statistics, College of Business and Economics, Australian National University, Canberra, Australian Capital Territory, Australia

[3]Data61, Commonwealth Scientific and Industrial Research Organisation, Docklands, Victoria, Australia

Correspondence
William Lim, Research School of Finance, Actuarial Studies and Statistics, College of Business and Economics, Australian National University, Canberra, ACT 2601, Australia.
Email: william.lim@anu.edu.au

Abstract

We investigate how well different investment strategies can give pre-retirees more certainty about their income in retirement, whilst allowing them to benefit from taking investment risk. Under an expected utility-maximizing framework, we find that a loss aversion utility function gives a high degree of certainty about its desired wealth target and is robust to different market models. Imposing terminal wealth constraints does not improve the certainty of achieving the desired target enough to counterbalance the increased chance of obtaining a lower income. The power utility function is not robust to different market models and becomes too risk-averse with wealth constraints.

The authors gratefully acknowledge the constructive comments of the Editor and the Associate Editor of European Financial Management journal along with the thorough comments of an anonymous referee. We also have benefited from comments made by Jenni Bettman, Alfred Chong, Marcos Escobar-Anel, Mogens Steffensen and Rudi Zagst. Furthermore, we also thank seminar participants at the 22nd International Congress on Insurance: Mathematics and Economics School of Risk & Actuarial Studies, University of New South Wales, and the PhD Symposium (2018), Research School of Finance, Actuarial Studies and Statistics, College of Business and Economics, Australia National University for their comments on an earlier version of the paper. Catherine Donnelly gratefully acknowledges funding from the Institute and Faculty of Actuaries' Actuarial Research Centre, as part of the research programme 'Minimising Longevity and Investment Risk while Optimising Future Pension Plans'. We further thank David Wilkie for providing some of the UK financial market data.

KEYWORDS
asset allocation, constrained optimization, loss aversion utility, numerical dynamic programming, retirement outcomes

JEL CLASSIFICATION
G11, G22, G51

1 | INTRODUCTION

How to invest one's pension savings before retirement is a decision faced by defined-contribution (DC) pension plan members. A typical DC pension plan member both saves and invests to build up a pension fund before retirement, from which they withdraw an income in retirement. The member's DC fund value at retirement can be volatile as it depends on several factors in the savings phase, such as their salary, the amount and frequency of contributions, and, how the funds are invested. To cite (Merton, 2014, p. 1402), the primary concern of the pre-retiree is 'Will I have sufficient income in retirement to live comfortably?'. This paper studies whether the investor is better off with investment strategies that constrain their retirement income, or equivalently fund value accumulated after the savings phase, so as to protect them against extreme (negative) scenarios of income in retirement.

We compare several ways of formulating the investor's problem, by allowing for different utility functions and terminal wealth boundary constraints and deriving optimal investment strategies through the maximization of expected utility. Using these optimal investment strategies, we calculate, analyse and compare the distributions of income at retirement. Our findings show that a loss aversion utility function gives a very attractive retirement income distribution, with the distribution peaked at the investor's chosen income goal with some level of robustness. In contrast, risk preferences expressed via a constant relative risk aversion (CRRA) utility function give a much more spread out income distribution, providing the investor less certainty on achieving a sufficient level of income in retirement. Imposing terminal wealth boundary constraints, in both the utility function settings, result in strategies that provide certainty of achieving the lower boundary but at the expense of significant reduction in the overall retirement outcome. We conclude that the investor can benefit from following a loss aversion-derived optimal investment strategy to target a sufficient level of income at retirement.

The investment problem with terminal wealth constraints at retirement has been extensively studied in the utility literature, in particular, a lower boundary constraint that guarantees the investor a minimum level of wealth at the cost of reducing the possible upside. For example, (Korn, 2005) and Kraft and Steffensen (2013) consider a static lower constraint while Teplá (2001) and Han and Hung (2012) consider minimum performance relative to a benchmark strategy. On the other hand, Donnelly et al. (2015) demonstrate the benefit of an upper constraint on terminal wealth in reducing the risk of poor retirement outcomes at the expense of giving up the possibility of higher returns. The merits of an upside constraint, in the form of capped investment earning rates, is also discussed by Mahayni and Schneider (2015). Schütte (2017) and Donnelly et al. (2018) consider outcomes with both upper and lower constraints on the terminal wealth based on an exponential utility function and a power utility function, respectively.

As evidenced by Kahneman and Tversky (1979) and Tversky and Kahneman (1992), investors tend to evaluate outcomes relative to a reference level. Several authors use loss aversion utility functions in investment problems pertinent to either DC pension plan members or plan managers (Berkelaar et al., 2004; Blake et al., 2013; Dong et al., 2020; Guan & Liang, 2016; He & Zhou, 2011; Jin & Zhou, 2008; Lin et al., 2017), and in conjuction with a lower terminal wealth constraint (Chen et al., 2017; Dong & Zheng, 2019).

However, the existing research has a number of limitations. There is a tendency to assume a wealth constraint that is either deterministic or a stochastic variable pegged to the labour income process. The literature with terminal constraints assume the labour income process to be either be a deterministic process or a tradable stochastic process where both can be replicated by a portfolio of risk-free and risky assets. Moreover, solving terminal wealth constraints often necessitate option-like payoffs as the wealth constraints are occasionally binding. To this end, the literature prescribes the investor to adopt option replicating strategies. While the above assumptions are required to construct an explicit mathematical solution, they may not be practical. Most investors have the ability to neither short-sell assets, invest more than their total wealth in the risky asset nor ability to trade continuously, which are all critical requirements in implementing these solutions.

This paper contributes to the literature in a number of ways. First, we extend the literature by incorporating a stochastic non-tradable labour income process which governs how both the preretirement investor's salary and the terminal wealth constraints evolve. The investor makes contributions during their working years which are invested in line with an optimal investment strategy. The optimal strategy is one which maximizes the expected value of the utility of retirement income. Two forms of the utility function are considered: power utility (i.e., constant relative risk aversion, or CRRA risk preferences) and loss aversion, and these are further considered in the presence of upper and lower terminal wealth constraints. To the best of our knowledge, this comparison has not been done before.

Second, we model inflation explicitly and study an investment market comprising of a nominal bond, an inflation-linked bond and a risky stock. Inflation is an important consideration for retirement planning and investment, since the time scales are generally over decades. Third, for greater realism, we allow the investors to only trade in short-term option contracts instead of long-term contracts or invest in a replicating strategy. It is not likely that long-term traded option contracts would be available in the market, and neither is it realistic to expect investors to construct an over-the-counter version or to do a replicating portfolio version. Finally, we compare the performance of the investment strategies both under the standard financial market model (geometric Brownian motion driven) from which the strategies are devised, and under historical simulations based on the financial market returns in the UK from December 1918 to December 2019.

We find that for both utility functions, adding in boundary constraints on the fund levels does not materially impact the asset allocation. The exception is when the fund level approaches the boundary constraints at either end. In that case the resultant strategies shift allocations from risky assets to the risk-free asset, to keep the fund level within the boundary constraints. Boundary constraints help improve the certainty of retirement income but do so at the expense of reducing the average retirement income. This is a consequence of reduced exposure to risky stock across all the possible paths of the market prices.

Second, we find that, under the standard financial market model, there is only a weak incentive for a CRRA investor to include boundary constraints. This is because the downside protection rarely comes into effect and the resultant strategies provide little increased certainty to the level of retirement income. Nonetheless, imposing boundary constraints remain a viable approach in preventing investor from achieving very poor retirement outcomes. It enables the CRRA investor to specify and maintain a minimum sufficient level of income at retirement that cannot be easily taken into account alone by the risk aversion parameter.

Compared to the CRRA risk preference, we find that there is an even weaker incentive for a loss aversion investor to include boundary constraints. With a reference level, the loss aversion risk preference enables the investor to target a sufficient level of retirement income without needing to impose boundary constraints. By adjusting the risk exposure throughout the investment phase, which is a feature of the loss aversion utility function, the investor may achieve their target retirement income. However, adding in the boundary constraints induces competing objectives, reduces allocation to the risky stock throughout the savings phase, and ultimately lowers the investor's chance of achieving their target.

Finally, we find that the optimal investment strategy derived from the loss aversion utility function appears more robust to a misspecification of the market model, compared to the that derived from CRRA preferences. The general shape of the retirement income distribution is largely preserved under the loss aversion derived strategy while the CRRA derived distribution changes, to the detriment of the investor.

Overall, our results suggest that an investor planning for their retirement, is better off following the loss aversion-derived optimal investment strategy without constraints on the utility function value at retirement. The remainder of the paper is structured as follows: in Section 2 the formulation and solution of the investor's problem are detailed; Section 3 discusses the potential outcomes resulting from the various investment strategies from Section 2; and, Section 4 provides further discussion and concludes the paper.

2 | THE INVESTOR'S PROBLEMS

The focus of our study is the distribution of an investor's real replacement ratio (RRR) at retirement. The RRR is the ratio of the annual real income bought at retirement to the annual real salary in the year before retirement. The advantage of using the replacement ratio is that it is more meaningful to the investor, as it tells them the percentage of the preretirement salary that they can expect postretirement and helps them in deciding their desired level of retirement income sufficiency. All real values are with respect to the purchasing power at time 0 when the investor is age 25. This avoids the investor having to do an inflation projection from age 25 to age 65 to determine the future purchasing power of their retirement income.

The constraints of no-short-selling of any asset, no borrowing against future income and discrete, annual re-balancing are included. Their inclusion necessitates the use of numerical methods to determine an investment strategy that maximizes the expected utility of retirement income. Additionally, we use interchangeably the terms ‘real fund value at retirement’ and ‘real retirement income.’ In our model, the accumulated fund value at retirement is used to purchase a life annuity. The life annuity has the same value in all future states of the world; see Equation (6).

2.1 | The financial market model

There are three asset classes available for investment. The uncertainty in the market is represented by a complete probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{0 \le t \le T}, \mathbb{P})$ that supports a two-dimensional Brownian motion $W(t) = (W_1(t), W_2(t))'$, where $W_1(t)$ is independent of $W_2(t)$ and the prime denotes transposition. Asset 0 is a nominal bond with price process S_0 satisfying

$$dS_0(t)/S_0(t) = r_{\mathrm{N}} dt, S_0(0) = 1, a.s. \tag{1}$$

The constant r_{N} is the annual nominal interest rate. The price of Asset 1, an inflation-linked bond, is driven by a price inflation index I with dynamics

$$dI(t)/I(t) = \mu_I dt + \sigma_I dW_1(t), I(0) = I_0 = 1, a.s, \tag{2}$$

with the constant μ_I representing the mean return on the index. The price process of Asset 1, S_1, satisfies

$$dS_1(t)/S_1(t) = (r_{\mathrm{R}} + \mu_I)dt + \sigma_I dW_1(t), S_1(0) = 1, a.s, \tag{3}$$

with the constant r_{R} denoting the real rate of return on the inflation-linked bond.

Asset 2 is a risky stock with price process S_2 with dynamics

$$dS_2(t)/S_2(t) = \mu_2 dt + \sigma_{21} dW_1(t) + \sigma_{22} dW_2(t), S_2(0) = 1, a.s, \tag{4}$$

with the constant μ_2 is the mean return on the stock, $\sigma_{21} > 0$ is its correlation with the price inflation index and $\sigma_{22} > 0$ is its correlation with other, non-modelled, random factors.

The expectation operator under $\mathbb{P}$ is denoted $\mathbb{E}$.

2.2 | Labour income

We model the labour income of the investor using the labour income structure seen in Cairns et al. (2006) and Blake et al. (2013) in a discretized setting. Specifically, we use a trinomial tree lattice to model the growth in labour income from integer time t to $t+1$ (Boyle, 1988; Hull, 2019). Figure 1 illustrates the discretized income labour process.

Consider an investor, aged 25 at time zero, who is saving for their retirement at age 65 (i.e., at time $T = 40$). As with Blake et al. (2013), the labour income of the investor has an age-dependent mean growth rate of $\mu_{L,t}$, whose components we detail shortly, and a constant volatility of σ_L. His labour income at the integer time t, when he is age $t + 25$, is given by Y_t. In our discretized setting, the labour income is determined every year so that there is a 1-year time step in the trinomial tree, that is, $\Delta t = 1$. From time t to $t+1$, the labour income moves from Y_t to one of the three values, $Y_t e^{u_t}$, $Y_t e^{m_t}$, and $Y_t e^{d_t}$. The movement from Y_t to $Y_t e^{u_t}$, $Y_t e^{m_t}$ and $Y_t e^{d_t}$ are ‘up’, ‘middle’ and ‘down’ movements, respectively. We choose the age-dependent movement factors $u_t = \mu_{L,t} + \sigma_L \sqrt{3\Delta t}$, $m_t = \mu_{L,t}$, and $d_t = \mu_{L,t} - \sigma_L \sqrt{3\Delta t}$. The mean growth rate is given by

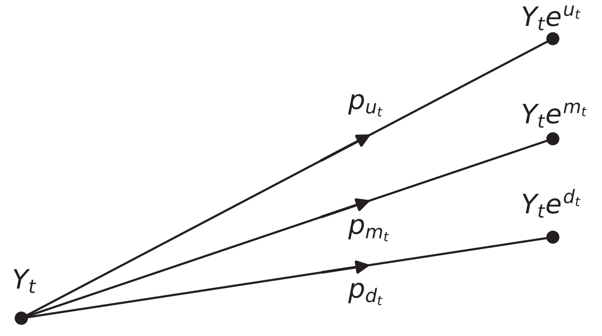

FIGURE 1 Illustrative example of a labour income process. This figure illustrates the labour income process modelled using a trinomial tree (not to scale). For each time step, the labour income moves from Y_t to one of the three values, $Y_te^{u_t}$, $Y_te^{m_t}$ and $Y_te^{d_t}$ with probabilities p_{u_t}, p_{m_t} and p_{d_t}, respectively

$$\mu_{L,t} = r_L + (P_{t+25+1} - P_{t+25})/P_{t+25},$$

in which r_L represents the long-term annual real growth rate in national earnings, P_{t+25} is the career salary profile at age $t + 25$ such that $(P_{t+25+1} - P_{t+25})/P_{t+25}$ reflects the promotional salary increase from time t to time $t + 1$. Let p_{u_t}, p_{m_t}, and p_{d_t} denote the probabilities of up, middle and down movements at each time, we choose the value of these probabilities so that the mean and standard deviation of the change in labour income in the discrete process matches $\mu_{L,t}$ and σ_L, respectively. Finally, we adopt the following quadratic function to model the salary profile P_x, at integer age $x = 25, ..., 65$:

$$P_x = 1 + h_1\left[\frac{(x-20)}{65-20} - 1\right] + h_2\left[-1 + \frac{4(x-20)}{65-20} - \left(\frac{\sqrt{3}\,(x-20)}{65-20}\right)^2\right],$$

where the parameters h_1 and h_2 are estimated from the historical data on promotional salary increases of UK male salary data from ages 20 to 65 (see Blake et al. 2007, 2013).

2.3 | Fund accumulation

The fund value of the investor at time t is denoted by the random variable F_t. Starting at age 25 ($t = 0$), the investor puts a constant proportion $\pi \in (0, 1]$ of their real labour income into the retirement fund at each integer time. For each integer year, the value F_{t-} represents the fund value at time t just *before* a lump-sum contribution is paid into the fund, and the value F_t represents the fund value at time t just *after* a lump-sum contribution is paid with

$$F_t = F_{t-} + \pi Y_t.$$

The investor decides upon their investment strategy once a year, immediately after their annual lump-sum contribution to their pension fund. The strategy is effectively a 1-year buy-and-hold strategy until the following year's new contribution and new investment strategy are implemented. Short-selling of any assets and borrowing against future labour income are prohibited.

The investment decision made at time t is represented by the random vector $\theta_t = (\theta_t^0, \theta_t^1, \theta_t^2)'$, in which the random variable θ_t^i represents the proportion of the fund value F_t invested in asset i, for $i \in \{0, 1, 2\}$, and the sum of the proportions satisfies $\sum_{i=0}^{2} \theta_t^i = 1$, for every $t \in \{0, 1, ..., T-1\}$. Furthermore, due to the no-short-selling constraint on all assets, $\theta_t^i \in [0, 1]$ for every $t \in \{0, 1, ..., T-1\}$ and $i \in \{0, 1, 2\}$.

At each integer time $t < T$, the investor makes a lump-sum contribution to the fund. As the investor implements a 1-year buy-and-hold strategy at the start of each year, the real fund value of the investor increases from the value F_t at time $t \in \{0, 1, ..., T-1\}$ to

$$F_{t+1} = (F_{t-} + \pi Y_t) \times \sum_{i=0}^{2} \theta_t^i \frac{S_i(t+1)}{S_i(t)} \frac{I(t)}{I(t+1)} \tag{5}$$

at integer time $t+1$.

At the terminal time T, the investor uses their accumulated retirement fund value, F_T, to buy a single life annuity whose income increases in line with the price inflation index. We assume that the underlying investment for the inflation-indexed life annuity is the inflation-linked bond. The single life annuity pays an annual income in advance, with the first income payment due at time T, until the investor dies. The expected cost of purchasing at time T (i.e., age 65) an inflation-indexed annuity income paid annually in advance for life, with the first payment being £1 per annum at age 65, is

$$\ddot{a}_{65} = \mathbb{E}\left(\sum_{k=0}^{x_{\max}} {}_k p_{65} \frac{S_1(T)}{S_1(T+k)} \frac{I(T+k)}{I(T)}\right) = \sum_{k=0}^{x_{\max}} {}_k p_{65} e^{-r_{\mathrm{R}} k}, \tag{6}$$

in which $x_{\max} < \infty$ is the maximum attainable age by the investor and ${}_k p_{65}$ is the real-world probability that the investor survives from age 65 to age $65 + k$. Thus, the investor can buy, at age 65, an annual real annuity income of $F_T / \ddot{a}_{65}$ to last for their retirement, with the annuity income increasing in line with the inflation index I.

2.4 | Risk preferences and the problems to solve

The risk preferences of the investor are modelled by a utility function υ. The utility function is either a power utility one, based on Expected Utility Theory, or a loss aversion utility function, based on Cumulative Prospect Theory. The general goal of the investor is to find an investment strategy that maximizes the expected utility of their real retirement income.

2.4.1 | Power utility investor with no retirement fund value constraints

First, assume the investor's risk preferences follows a power utility function

$$\upsilon_1(x) = x^\gamma / \gamma, \text{ for } x > 0. \tag{7}$$

The constant $\gamma \in (-\infty, 1)\setminus\{0\}$ represents the relative risk aversion of the investor.

The investor's problem is to determine the 1-year buy-and-hold investment strategies $\{\theta_t\}_{t=0,1,\ldots,T-1}$ which solve

$$\max_{\{\theta_t\}_{t=0}^{T-1}} \mathbb{E}\left[\beta^{T-t} \upsilon_1(F_T)\right], \tag{8}$$

in which $\beta \in (0, 1]$ is the investor's personal discount factor.

Numerical dynamic programming is used to derive an optimal investment strategy. With $V_t(\cdot)$ denoting the value function at time t, these optimal proportions are chosen to achieve the maximum in the Bellman equation

$$\begin{aligned} V_t(F_t, Y_t) = \max_{\theta_t \in [0,1]^3} (\beta\, \mathbb{E}[V_{t+1}(F_{t+1}, Y_{t+1})] | \mathcal{F}_t), \text{ for } t = 0, 1, \ldots, T-1, \\ \text{with terminal value } V_T(F_T, Y_T) = \upsilon_1(F_T), \end{aligned} \tag{9}$$

subject to budget constraint (5). From the terminal value, we work backwards to find an optimal investment strategy.

2.4.2 | Loss aversion utility investor with no retirement fund value constraints

The next utility function considered is a loss aversion utility function. Such utility functions arise from the prospect theory framework, which was proposed and developed in Kahneman and Tversky (1979) and Tversky and Kahneman (1992). The framework features an S-shape loss aversion utility function and a probability weighting function. The former models the preferences of investors who measure losses and gains relative to a reference point and are risk-seeking with respect to losses but risk-averse with respect to gains. The latter applies a nonlinear transformation to probability measure which inflates a small probability and deflates a large probability. The loss aversion utility is arguably more realistic than an Expected Utility Theory-based utility function as the loss aversion utility reflects the risk-seeking behaviour of investors when they have made a loss with respect to their reference point (Mitchell & Utkus, 2004). The problem of maximizing terminal wealth from the perspective of individual investors under a loss aversion utility function has been studied by various authors (e.g., Berkelaar et al., 2004; Blake et al., 2013; Chen et al., 2017; Dong & Zheng, 2019). Jin and Zhou (2008), He and Zhou (2011), and van Bilsen and Laeven (2020) study the implications of loss aversion with a probability weighting function. In this study, we adopt the modelling framework of Blake et al. (2013) as they tackled the investor's problem in a discrete-time rebalancing setting that is similar to ours.

In the framework of Blake et al. (2013), which we apply here, the investor has a target value that they would like their fund value to achieve at their retirement date, i.e., the level sufficient to support their desired retirement. The investor maximizes the weighted sum of the expected utility gained from gains and losses about a series of interim targets as well as the retirement date target. The idea of the interim targets is to keep the investor on track for the retirement date (final) target. The loss aversion utility function is given by

$$v_2(F,K) = \begin{cases} \frac{1}{\nu_1}(F-K)^{\nu_1}, & \text{if } F \geq K, \\ -\frac{\lambda}{\nu_2}(K-F)^{\nu_2}, & \text{if } F < K, \end{cases} \quad (10)$$

where F is the fund level, K is the target level, ν_1 is the degree of risk-aversion to gains, ν_2 is the degree of risk-seeking with respect to losses, and λ is the relative weighting of losses compared to gain.

Given an income of Y_t at time t, the final target, $K_k(T, Y_t)$, is the expected fund level at time T required to achieve a replacement ratio of k; and the interim target at time $s \geq t$, $K_k(s, Y_t)$, is the final target less the expected values of all future contributions to the fund from time s to time T, discounted to time s. Mathematically,

$$K_k(s, Y_t) = \begin{cases} \mathbb{E}[Y_T|\mathcal{F}_t] \times \ddot{a}_{65} \times k, & \text{for } s = T, \\ K_k(T, Y_t)e^{-r_R(T-s)} - \pi \sum_{j=s}^{T-1} \mathbb{E}[Y_j|\mathcal{F}_t]e^{-r_R(T-1-j)}, & \text{for } s = t, t+1, \ldots, T-1. \end{cases} \quad (11)$$

The investor's problem is to determine the 1-year buy-and-hold investment strategies $\{\theta_t\}_{t=0,1,\ldots,T-1}$ which solve

$$\max_{\{\theta_t\}_{t=0}^{T-1}} \mathbb{E}\left[\omega\left(\sum_{s=0}^{T-1} \beta^s v_2(F_t, K_k(s, Y_t))\right) + \beta^{T-t} v_2(F_T, K_k(t, Y_T))\right], \quad (12)$$

in which $\omega \in [0, 1)$ is the weighting factor applied to the utility derived from the interim targets, and $\beta \in (0, 1]$ is the investor's personal discount factor. The choice of $\omega < 1$ reflects the lower weight given to the interim targets compared to the retirement date target, which has an implicit weight of unity.

Denoting by $V_t(\cdot)$ the value function at time t, the investor's problem is solved using the Bellman equation

$$V_t(F_t, Y_t) = \max_{\theta_t \in [0,1]^3} (\omega\ v_2(F_t, K_k(t, Y_t)) + \beta\ \mathbb{E}[V_{t+1}(F_{t+1}, Y_{t+1})|\mathcal{F}_t]), \text{ for } t = 0, 1, \ldots, T-1,$$
$$\text{with terminal value } V_T(F_T, Y_T) = v_2(F_T, K_k(T, Y_T)), \quad (13)$$

subject to budget constraint (5). From the terminal value, we work backwards to find the optimal investment strategies.

2.5 | Adding in both lower and upper fund value boundary constraints

Here, we describe how fund value boundary constraints that apply only at retirement are incorporated into the problems detailed in Section 2.4 to answer one of our research questions: how does applying constraints on the fund value at retirement change the distribution of the real replacement ratio? Note, they are in addition to the no-short-selling and no borrowing constraints. We implement them using options.

First, the investor chooses two values, $l \geq 0$ and $u \geq l$, guided by their minimum and maximum levels of income sufficiency, between which they would like their replacement ratio at retirement to lie. Then the corresponding fund value boundary constraints at the time of retirement, $K_l(T, Y_t)$ and $K_u(T, Y_t)$, respectively (see Equations 15–14), are calculated. The fund value boundary constraints captures additional risk preferences of the investors that are not accounted for in the utility functions. The upper bound represents the level beyond which the investor forgoes to increase the lower quantiles of the retirement fund value (Donnelly et al., 2015). The lower bound is the minimum that the investor needs at retirement, to secure a minimal sufficient income in retirement. This is the same interpretation as in Donnelly et al. (2018).

The implementation of the investment strategy in the presence of the lower and upper bounds involves either trading in very long-term options or investing in the underlying replicating portfolio, both of which may be unrealistic for an individual investor. In our discrete-time framework where trading is done once annually, we try to achieve a realistic solution to the problem that may be attractive to investors. The solution has two parts.

One part involves the imposition of interim bounds to constrain the investor's preretirement fund value, in a similar fashion to the interim target values for the loss aversion utility. The idea is to ensure the investor's fund value is on track to be between the lower and upper bounds at retirement date.

The other part of the solution is that the investor trades in 1-year options at each integer time. In addition, this is a more realistic possibility than trading in 40-year options to meet the problem's solution as long-dated options are relatively rare and prohibitively expensive for individual investors. The 1-year options are chosen to keep the fund value between the interim bounds at the end of each year. The net fund value, after the cost of option trading has been deducted, is then fed into the appropriate system to solve, either Equation (9) or Equation (13). The procedure by which the 1-year options-based strategy is carried out is described next.

2.5.1 | 1-year options-based strategy

Here we impose a state-dependent constraint that the investor maintains his fund level at or above the level of $K_l(s, Y_t)$—the lower fund level bound for the entire remaining investment period, $s \in [t, T]$, given his current state at time t with an income of Y_t. The calculation of $K_l(s, Y_t)$ is detailed below. The intuition for this is that if the investor were to allow their fund level to fall below the value of $K_l(s, Y_t)$, no subsequent trading strategy could guarantee that they would meet the boundary constraints at retirement (Korn, 2005; Teplá, 2001).

Consider a situation where the investor can buy and sell European options, which are 1-year options written annually on the investor's 1-year buy-and-hold strategies. As described in the previous section, the investor chooses values for the lower replacement ratio l and upper replacement ratio u such that the upper fund level bound is higher than the lower fund level bound $K_u(T, Y_t) > K_l(T, Y_t)$. The investor would like the real value of their fund value at retirement to lie between these two boundary values.

The final and interim lower fund level bounds are the fund level required to achieve the replacement ratio of l at retirement. The fund level bounds evolve in accordance to the states of the investor:

$$K_l(s, Y_t) = \begin{cases} Y_M(s, Y_t) \times \ddot{a}_{65} \times l, & \text{for } s = T, \\ K_l(T, Y_t)e^{-r_R(T-s)} - \pi\sum_{j=s}^{T-1} Y_M(j, Y_t)e^{-r_R(T-1-j)}, & \text{for } s = t, t+1, \ldots, T-1, \end{cases} \tag{14}$$

in which

$$Y_M(s, Y_t) = \begin{cases} Y_t \exp\left(\sum_{j=t}^{s-1} u_j\right), & \text{for } s = t+1, \ldots, T, \\ Y_t, & \text{for } s = t \end{cases}$$

is the maximum income attainable at time s given the investor has an income of Y_t at time t.

The interim fund level bounds are constructed to ensure that the investor can achieve a minimum replacement ratio of l at retirement even when he is at the maximum labour income path throughout the investment period.

Similarly, the upper bounds, beyond which the investor forgoes the upside risk, also evolve in accordance to the state of the investor:

$$K_u(s, Y_t) = \begin{cases} Y_N(s, Y_t) \times \ddot{a}_{65} \times u, & \text{for } s = T, \\ K_u(T, Y_t)e^{-r_R(T-s)} - \pi\sum_{j=s}^{T-1} Y_N(j, Y_t)e^{-r_R(T-1-j)}, & \text{for } s = t, t+1, \ldots, T-1, \end{cases} \tag{15}$$

in which

$$Y_N(s, Y_t) = \begin{cases} Y_t \exp\left(\sum_{j=t}^{s-1} d_j\right), & \text{for } s = t+1, \ldots, T, \\ Y_t, & \text{for } s = t \end{cases}$$

is the minimum income attainable at time s given the investor has an income of Y_t at time t.

Given the current state with an income of Y_t at time t, the investor deals in 1-year options to keep their fund value at integer time $s > t$ between the interim bounds $(K_l[s, Y_t], K_u[s, Y_t])$. The investor's resultant annual investment strategy at each integer time t has three elements that can be broadly described as:

- Securing at least the interim lower bound $K_l(t+1, Y_t)$ at time $t+1$ by buying a 1-year put option with strike price $K_l(t+1, Y_t)$;
- Securing at most the interim upper bound $K_u(t+1, Y_t)$ at time $t+1$ by selling a 1-year call option with strike price $K_u(t+1, Y_t)$; and,
- Investing the remaining fund value to maximize the expected utility of the real fund value.

Next we describe the evolution of the investor's fund value process $\{F^{\text{opt}}{}_t\}_{t\in[0_-,T]}$ under this strategy. First, consider a portfolio which (i) has a value Z_t at time t, whose calculation we detail shortly; and (ii) invests at time t the amount $Z_t \times \theta_t^{\text{opt},i}$ in asset i, for $i = 0, 1, 2$.

The proportions $\left\{\theta_t^{\text{opt},i}\right\}_{i=0,1,2}$ are obtained by the maximization of either (8) or (12), according to which utility function is being applied.

At integer time $t \in \{0, 1, ..., T-1\}$, the investor buys a put option, with its strike price $K_l(t+1, Y_t)$ calculated according to Equation (14). Simultaneously, the investor sells a call option, written on the same portfolio as the put option, with its strike price $K_u(t+1, Y_t)$ calculated according to Equation (15). The fund value at integer time $t \in \{0, 1, ..., T-1\}$ is calculated recursively as follows. The starting fund value is defined to be $F^{\text{opt}}{}_0 = F_{0\_} + \pi Y_0$. For time $t = 0, 1, ..., T-1$, the value Z_t is the solution to the equation

$$F^{\text{opt}}{}_t = Z_t - \text{call}(Z_t, K_u(t+1, Y_t), r_{\text{R}}, \sigma_{Z_t}, 1) + \text{put}(Z_t, K_l(t+1, Y_t), r_{\text{R}}, \sigma_{Z_t}, 1), \quad (16)$$

where call(put)(S, K, r, σ, T) denotes the Black-Scholes value of a European call (put) option at inception with S refers to the value of the underlying asset at inception, K the strike price, r the risk-free rate of return,[1] σ_S the volatility of the underlying asset, and T the time to maturity. The underlying asset of the 1-year options is the portfolio with value Z_t and its volatility σ_{Z_t}, determined by the volatilities of the assets i and the proportions invested in each of the assets $\theta_t^{\text{opt},i}$ for $i = 0, 1, 2$:

$$\sigma_{Z_t} = \sqrt{(\theta_t^{\text{opt},0})^2\sigma_I^2 + (\theta_t^{\text{opt},2})^2\sigma_{22}^2 + 2\theta_t^{\text{opt},0}\theta_t^{\text{opt},2}(\sigma_{21} - \sigma_I)^2}. \quad (17)$$

After 1 year, the fund value just after the lump-sum contribution made at integer age $t+1$ is calculated as

$$F^{\text{opt}}{}_{t+1} = \pi Y_{t+1} + \min\left(\max\left(Z_t \times \sum_{i=0}^{2} \theta_t^{\text{opt},i} \frac{S_i(t+1)}{S_i(t)} \frac{I(t)}{I(t+1)}, K_l(t+1, Y_t)\right), K_u(t+1, Y_t)\right). \quad (18)$$

As no contribution is made at retirement $t = T$, we set

$$F^{\text{opt}}{}_T = \min\left(\max\left(Z_{T-1} \times \sum_{i=0}^{2} \theta_{T-1}^{\text{opt},i} \frac{S_i(T)}{S_i(T-1)} \frac{I(T-1)}{I(T)}, K_l(T, Y_t)\right), K_u(T, Y_t)\right). \quad (19)$$

We remark that this strategy cannot be implemented directly from individual exchange-traded options. The strike price applies to the value of a basket of assets. This means that the investor is unable to buy the equivalent of each 1-year put option through the purchase of a combination of 1-year European put options written separately on each of the three assets. There is no strike price to apply to each of the three options separately. Instead, the strike price operates across the three underlying assets in each option, rather than being applied individually.[2]

[1]We use the real rate of returns, r_{R}, because the fund levels and prices are specified in real terms (Fischer, 1978).
[2]We perform sensitivity analysis to allow for transaction costs by following the Leland (1985) approach to adjust volatility given in Equation 17, which reduces the (already weak) incentive of using the option-based strategies (see Supporting Information Appendix A).

The investor's problems with fund value constraints are to maximize Equation (8) or Equation (12) with the change that the fund value evolves as $F^{\text{opt}}{}_t$ instead of F_t. The adjusted problems are solved using the Bellman equations (Equations 9–13) subject to a new set of budget constraints (Equations 16–19).

2.6 | Numerical methods used

The numerical results for the discrete-time optimization problems detailed in Sections 2.4 and 2.5 are obtained by a numerical stochastic dynamic programming approach.

First, we discretize the state of the problems into three state variables, which include the investor's fund level, F_t, his labour income Y_t, and his age (or time until retirement t). The fund level is discretized into 201 equidistant grid points that spans the range of 0 and $50 \times Y_t$ for problems with no constraints on fund values and the range of $K_l(t, Y_t)$ and $K_u(t, Y_t)$ for problems with constraints on fund values. The labour income is discretized into up to 9 grid points that spans the range of income at each age and coincides with the nodes in the lattice (see Figure 1). Assuming a step in time corresponds to exactly 1 year, we solve the problem for each integer time.

For each grid point in the discretized state space, we solve the Bellman equations for value functions V_t recursively with backward induction, by finding the optimal asset allocations θ_t^i for $i \in \{0, 1, 2\}$ using a multi-dimensional optimization with constraints such that $\theta_t^i \in [0, 1]$ at time $T - 1, T - 2$ and so on. Once the backward induction has reached $t = 0$, the optimal investment strategy for an investor can be derived by following the optimal asset allocations that correspond with the fund level grid point for each integer time $t = 0, \ldots, \ T - 1$. For the problem with constraints on fund values, an additional root-finding procedure is embedded in the optimization process to obtain Z_t, which is necessitated by the option payoffs (see Equation 18).

Since the value function is only available at pre-defined grid points, the value function V_t between grid points is interpolated. The expectation with respect to the stochastic return in the Bellman equations is calculated with two-dimensional Gauss-Hermite quadrature with total number of nodes 10, 000. The large number of nodes ensures the shape of value functions in the presence of the option payoffs is adequately captured. For recent applications of quadrature integration and interpolation in lifecycle modelling, see Horneff et al. (2015), Andréasson et al. (2017) and Khemka and Butt (2017).

In our numerical procedure, we use the investor's normalized fund level (defined as the fund level divided by the labour income at time t) in place of the fund level. This numerical trick enables us to take advantage of the scale independence of the problems under the CRRA risk preference, allowing the problems to be solved with just two state variables (see Gomes & Michaelides, 2003).

3 | NUMERICAL ANALYSIS

We evaluate the potential retirement outcomes numerically from the perspective of an investor who starts investing at age 25 and retires at age 65. First, in Section 3.1, we discuss our baseline parameterizations under which we derive the optimal investment strategies implied by the problems considered in Section 2. Next, we provide some economic explanations to the asset allocations of these strategies in Section 3.2. Finally in Section 3.3, we compare the potential retirement outcomes following each of the investment strategies using simulated results.

3.1 | Parameterization of the model variables

We calibrate the financial market model detailed in Section 2.1 to UK market data.[3] The risky stock is represented by the FTSE 100 Total Return Index. We proxy for nominal bonds through the Thomson Reuters UK 10 Years Government Benchmark Index and for inflation we use the Retail Price Index (RPI). Using the data from January 1980 to December 2019 and adjusting for inflation, 12-month real rolling returns were calculated, resulting in 468 annual returns that were used to calibrate the parameters for the nominal bond and the risky stock. The real risk free rate is calibrated to the real rate of returns of the inflation-linked bonds (called index-linked gilts) issued by the UK government, which were first introduced in 1981. Specifically, we use the real returns of the FTSE British Government Index Linked (All Maturities) Index from January 1994 (earliest available) to December 2019. The parameter values are shown in Table 1.

The investor, who is exactly age 25 at $t = 0$, has a starting labour income of $Y_0 = 1$ unit per annum. His real labour income (i.e., the income expressed in terms of purchasing power at age 25) gets updated on each integer year following the tree process described in Section 2.2. The mean growth rate of the labour income has parameters $r_L = 0.019$, $h_1 = -0.1865$ and $h_2 = 0.7537$, which are obtained from Blake et al. (2013) who fitted the equations to UK male national income data. Finally, we assume a relatively low income volatility of $\sigma_L = 0.01$ such that the labour income progression fluctuates within a reasonable range as seen in well-functioning economies.

The investor contributes a constant fraction $\pi = 11.5\%$ of his income as a lump-sum annually on each integer year for 40 years (i.e., $T = 40$). This is between (i) the UK minimum mandatory contribution of 8% (Vanguard & Nest Insight, 2019) and (ii) the contribution rate of 16% required to achieve a two-thirds replacement ratio after 40 years of savings following a life-style investment strategy (Byrne et al., 2007). We assume that the investor starts with initial savings of $F_{0-} = 0$ units.

The goals and bounds are expressed as RRRs. To translate them into a fund value target, we use an annuity rate of $\ddot{a}_{65} := 14.3779$, based on a continuously-compounded, annual real rate of return $r_R = 0.026$ (consistent with Table 1) and the male annuitant mortality table S1PMA.[4] Note that the annuity rate is constant across all simulations and across all the studied problems.

The CRRA utility function is parameterized with $\gamma = -2$. For the loss aversion utility function, we adopt the same parameter values of $\lambda = 4.5$, $\nu_1 = 0.44$ and $\nu_1 = 0.88$ as those in Blake et al. (2013) who calibrate the parameters to UK data which is context of our investigation as well. The investor targets a RRR of $k = 2/3$. The personal discount factors of both utility functions (Equations 8 and 12) are assumed to be $\beta = 0.96$; and the weighting factor for interim utility derived under the loss aversion risk reference in Equation 12 is $\omega = 0.5$.

For the problems with constraints on the fund value at retirement time, we set the lower bound of the replacement ratio at $l = 0.3$, so as to satisfy the initial budget constraints (Equation 18) at time 0 for the baseline parameterization and our sensitivity analyses. The upper bound of the replacement ratio is assumed to be $u = 1.5$ (Table 2).

[3]The data is obtained from DataStream.
[4]S1PMA is a UK, male mortality model which can be downloaded from https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-mortality-and-morbidity-tables/s1-series-tables, S1 series table, CMI mortality and morbidity tables (accessed on 10 November, 2020).

TABLE 1 Parameters for the stochastic financial market model

This table reports the parameters for the stochastic financial market model as described in Section 2.1, in which r_N is the annual nominal interest rate; r_R is the real rate of return; μ_I are σ_I are the mean return and volatility on the price inflation index, respectively; μ_2 is the mean return on the risky stock, σ_{21} is its correlation with the price inflation index and σ_{22} is its correlation with non-modelled risks. The parameters are calibrated to the UK market using data from January 1981 to December 2019.

r_N	r_R	μ_I	σ_I	μ_2	σ_{21}	σ_{22}
0.073	0.026	0.038	0.078	0.091	0.096	0.142

TABLE 2 Parameters for labour income and other investment specification

This table reports the parameters for labour income and other investment specification. For labour income, we follow Blake et al. (2013)'s model specification, in which r_L and σ_L are the mean growth rate and volatility of the labour income, respectively; h_1 and h_2 are the curvature of parameters of the salary profile; Y_0 is the initial labour income. As with Blake et al. (2013), the parameters are fitted to the UK male national income data. The investor has an initial savings of F_{0-} and contributes a constant fraction π of his income annually. The annuity rate $\ddot{a}_{65}$ is calculated from the male annuitant mortality table S1PMA by the Institute and Faculty of Actuaries. l and u are the lower and upper bound of the replacement ratio, respectively.

Labour income		**Investment specification**	
r_L	0.019	F_{0-}	0
σ_L	0.1	π	11.5%
h_1	−0.1865	$\ddot{a}_{65}$	14.3779
h_2	0.7537	l	0.3
Y_0	1	u	1.5

3.2 | Investment strategies

Here we present the optimal asset allocations of the four different strategies:

- CRRA preferences with no lower and upper bounds on the RRR;
- CRRA preferences with the RRR at retirement constrained to lie in the range $[l, u]$;
- Loss aversion preferences with no lower and upper bounds on the RRR; and
- Loss aversion preferences with the RRR at retirement constrained to lie in the range $[l, u]$.

These optimal strategies are calculated from the problems described in Sections 2. The investment strategies depends on three state variables of the investor: his fund level, his labour income, and his age (or time until retirement). We simplify the discussion by focusing on the normalized fund level of the investor, which is defined as the fund level divided by the labour income at time t. Taking advantage of the scale-independence property of CRRA utility function, this enables us to fully characterize the investment strategies under the CRRA risk preference with just two state variables: the investor's normalized fund level and his age.

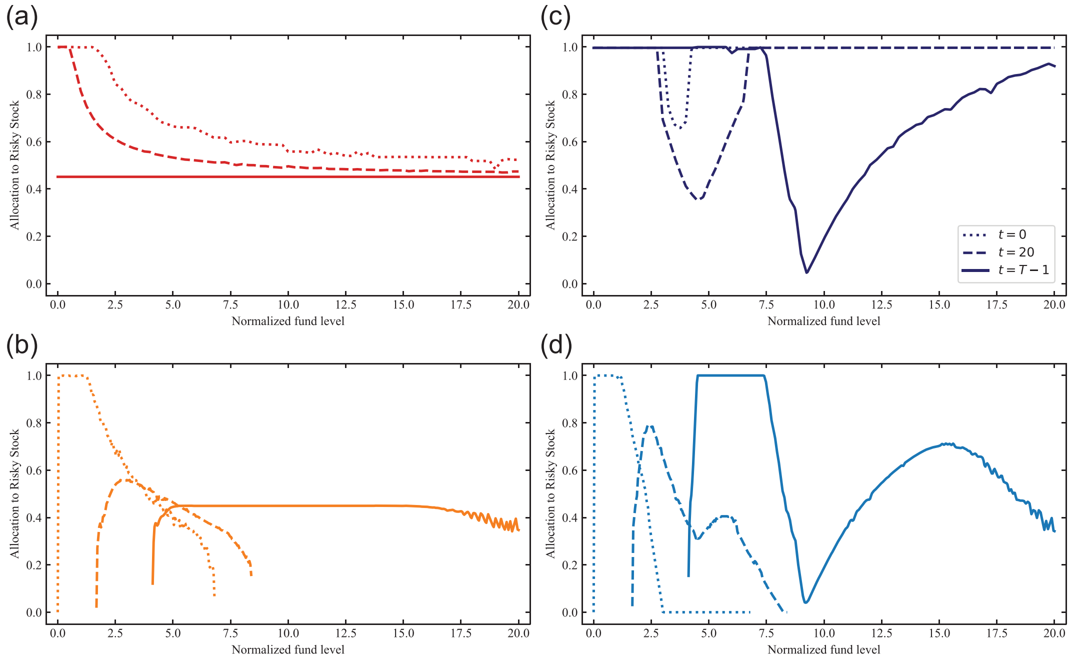

FIGURE 2 Optimal allocations to the risky stock. (a) Unconstrained (CRRA), (b) unconstrained (loss aversion), (c) 1-year option (CRRA), and (d) 1-year option (loss aversion). These figures show the optimal allocation to the risky stock for the four strategies in three different points of the investment period, assuming the investor's labour income is at its expected value. Due to the time-specific budget constraints, the constrained strategies are only applicable for a specific range of normalized wealth for each point in time. CRRA, constant relative risk aversion

Although the problems under the loss aversion risk preference are not scale-independent and its state space cannot be reduced, using the normalized fund level better reveals the asset allocation decisions implied by the problem.

Given the investor's normalized fund level and his age, we note that the asset allocation decisions remain largely unchanged across different levels of labour income. In particular, a slightly higher allocation to the risky stock is observed when the labour income is lower, and vice versa.

For brevity we assume that the investor's labour income matches the expected value across all ages. In addition, we focus our discussion on the optimal allocation to the risky stock because it best illustrates the risk appetite of the investor prescribed by the respective investment strategies. Figure 2 shows the allocations to the risky stock against the normalized fund level for each of the four strategies over three different points in the investment period with varying time until retirement.[5]

Under the CRRA utility function, at time $T-1$, the allocation to the risky stock is a constant value at about 45% without the boundary constraints. For time $t < T$, the allocation tends asymptotically to about 45% as the fund level increases. This is because an investor with

[5]For allocation to nominal bonds and additional visualizations of the risky stock allocation, please see Supporting Information Appendix B.

CRRA risk preferences maintains a constant optimal allocation proportion to the risky stock after accounting for his future contributions (Korn & Krekel, 2003).

Under the loss aversion utility function, the allocation to the risky stock are V-shaped curves. The investor tends to take risks when their fund level is lower (i.e., they are in the risk seeking domain) to try and achieve the target. Conversely, when their fund is higher than their target (i.e., when they are in the risk averse domain), the likelihood of falling below the target is low and the investor can afford to take additional risks to generate extra utility. Moreover, the V-shapes are narrower and shallower for younger ages. This indicates a higher overall risk appetite for a larger range of balances when the investors are further away from retirement. As the investor approaches retirement, the propensity to 'lock-in' the target imposes a significant penalty on falling below the target. In turn, this leads to the wider and deeper V-shaped allocation to the risky stock (Warren, 2019). Our findings are consistent with the results of Blake et al. (2013) and Butt et al. (2019).

For both utility functions, adding in boundary constraints on the fund levels does not materially impact the asset allocation except when the fund level approaches the boundary constraints. Under the 1-year option strategy, as the fund level approaches the boundary, the optimal allocation to the risky assets (stock and nominal bond) reduce towards zero (and thus volatility tends to zero) to reduce the cost of options and ensure that the fund is able to satisfy the boundary constraints. A closer inspection reveals that the reduction in the allocation to the risky stock is offset by increasing the allocation first to the nominal bond, which is a safer risky asset, and later, to the inflation-indexed bond, which is the real risk-free asset. As fund level reaches the boundary, the investor invests entirely in the inflation-indexed bond. Our findings are consistent with the results of Donnelly et al. (2018) albeit they study only the problem with power utility in a continuous time framework with only one risky asset.

Finally, we note that optimal allocation to the nominal bond is the complement of the optimal allocation to the risky stock whenever the strategy prescribes no allocation to the inflation-indexed bond. This is the case for most fund levels at most ages, with a few exceptions when a risk-free investment is desired. The exceptions happen in two instances. First, when the fund level of the 1-year option strategies approaches their bounds at both ends. In other words, as discussed earlier, when the investor has no other choice but to invest fully in the risk-free asset to meet the boundary constraints. Second, when a loss averse investor's fund level is near his target level in the last 10 years of his working life, due to the 'lock-in' nature of the loss aversion utility function.

3.3 | Distribution of retirement outcomes

Here we present the distributions of the RRR under the four different strategies. For each strategy, we report the distributions of the RRR at retirement time T, or equivalently age 65. The outcomes under these strategies are simulated first, under the stochastic financial market model of Section 2.1 with parameter values shown in Table 1, and second, under historical market return data in the UK from December 1918 to December 2019.

3.3.1 | Distribution under the stochastic financial market model

For each strategy, 10, 000 simulations were generated to obtain the distribution of the RRR. Histograms of the distribution of the RRR are shown in Figure 3 and the corresponding summary statistics are reported in Table 3.

Without the boundary constraints, the loss aversion utility function leads to a distribution of retirement outcomes that centres around the investor's target (Figure 3c). Following the optimal investment strategy under the unconstrained loss aversion preferences, the investor achieves the highest mean RRR and a sufficiently high probability of achieving the target. However, this exposes the investor to have a relatively high chance of failing to meet the lower boundary at 2% probability. The investor also faces relatively high variability (as measured by standard deviation) in their outcome distribution and is most prone to achieving low RRRs (as measured by Conditional Value at Risk (CVaR), which we define as the probability-weighted average of the lowest 1% of the RRRs). Although the loss aversion strategy leads to highest chance of achieving their retirement income level, it offers the least protection against the worst outcomes.

Adding in the boundary constraints to the loss aversion risk preference problem cuts off the distribution below l and above u (Figure 3d) with a probability mass at each end-point. The 1-year option strategy provides a downside protection against the worst outcomes by having a higher CVaR value at 0.304, eliminates the outcomes below l and above u, and reins in the standard deviation. However, there is a concomitant reduction in both the mean outcomes and

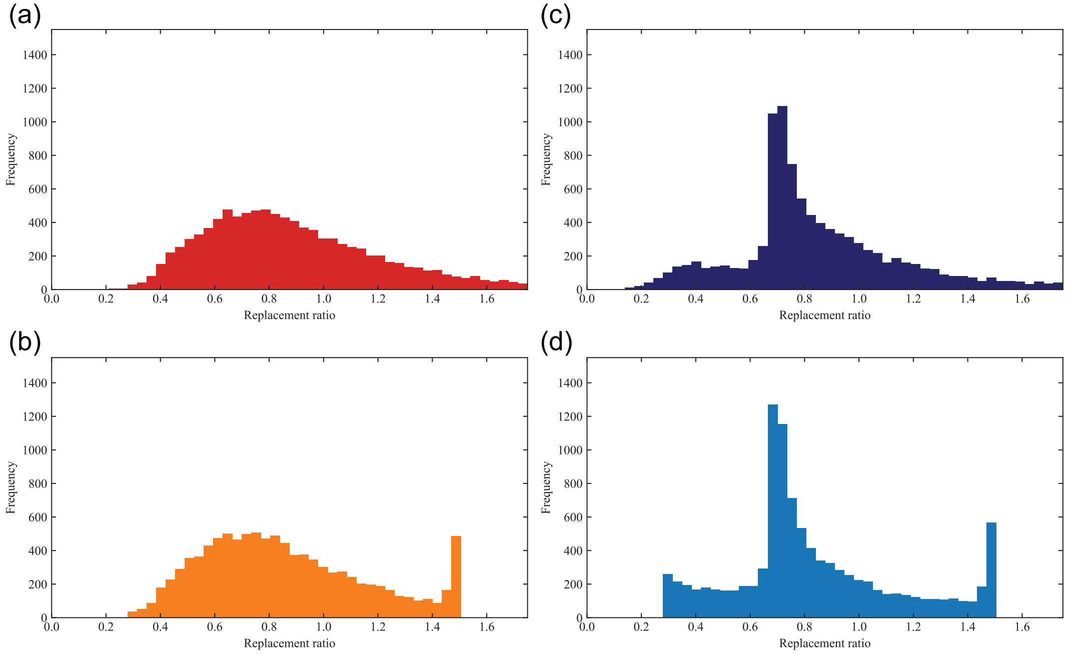

FIGURE 3 Histrograms of simulated real replacement ratios under the stochastic financial market model. (a) Unconstrained (CRRA), (b) unconstrained (loss aversion), (c) 1-year option (CRRA), and (d) 1-year option (loss aversion). These figures show the histograms of 10, 000 simulated real replacement ratios at time T for an investor following either one of the four strategies under the stochastic financial market model (the figures are right-truncated at the value of 1.75). CRRA, constant relative risk aversion

TABLE 3 Comparison of outcomes of investment strategies under the stochastic financial market model

This table reports the summary statistics of the outcomes distributions following each of the four investment strategies under the stochastic financial market model specified in Section 2.1. We report the mean, the standard deviation (SD), CVaR (Conditional Value at Risk, a risk measure which we define as the probability-weighted average of the lowest 1% of the RRRs), the probability of failing to achieve the lower boundary, and the probability of achieving the target RRR. Abbreviations: CRRA, constant relative risk aversion; RRR, real replacement ratio. [a]Since there is no explicit target for the CRRA investor, we report the probability of achieving the same target as the loss averse investor, namely $k = 2/3$.

	Mean	SD	CVaR	$\mathbb{P}[\text{RRR} < l]$	$\mathbb{P}[\text{RRR} \geq k]$
CRRA, unconstrained	0.916	0.374	0.324	0.2%	73.2%[a]
CRRA with 1-year options	0.856	0.295	0.327	–	69.9%[a]
Loss aversion, unconstrained	0.925	0.487	0.222	2.0%	82.0%
Loss aversion with 1-year options	0.832	0.302	0.304	–	78.1%

the chance of attaining a RRR near the target level. This is attributed to the interference of boundary constraints on the natural targeting mechanism of the loss aversion risk preference. The boundary constraints limit the risk taking tendency of a loss averse investor when the fund level is lower than the target (Figure 2d). Adding in the boundary constraints results in a more risk averse strategy that shifts an average of about 20% of allocation from the risky stock to the nominal bond each year across the investment period (Figure 4). This impairs the investor's ability to catch up, eventually costing them a reduction of 0.093 in mean RRR and a reduction of 3.9% in the probability of achieving the target replacement ratio k at retirement.

On the other hand, the simulated outcomes show that the use of the well-studied CRRA utility function to express the investor's risk preferences give a much more dispersed range of outcomes (Figure 3a). There is no strong concentration of outcomes about a particular value. Instead, the distribution peaks gently in the range of [0.7, 0.9] of RRRs. Without the boundary constraints, the CRRA utility function also results in a distribution with relative high values of mean and standard deviation that are not dissimilar to the unconstrained loss aversion strategy. Compared to the unconstrained loss aversion strategy, the unconstrained CRRA investor is less vulnerable to the risk of achieving poor retirement outcomes with $\mathbb{P}[\text{RRR} < l]$ at 0.2% and a CVaR value of 0.324. This is because when there are significant investment losses (relative to the target), the loss aversion risk preference, in search of higher potential gains, prescribes a high allocation to risky assets which also increases the chance of incurring greater investment losses. Contrarily, the CRRA risk preference seeks an optimal allocation to risky assets independent of the fund level.

Constraining the RRRs to lie in the range $[l, u]$ at age 65 implies that the investor is more risk averse than the risk preference implied by the CRRA utility function. By completely eliminating the chance of outcomes below l and above u, the 1-year option strategy also reduces the standard deviation and marginally improves the lower quantiles (e.g., CVaR) of the RRR at retirement for the CRRA investor (see Donnelly et al., 2015). However, this is accompanied by an increase in the allocation to the nominal bond over the investment period and costs the investor a reduction of 0.060 in mean RRR at retirement.

Based on the simulation results, the incentive of including the boundary constraints is not strong for investors under either risk preference structure. The downside protection rarely

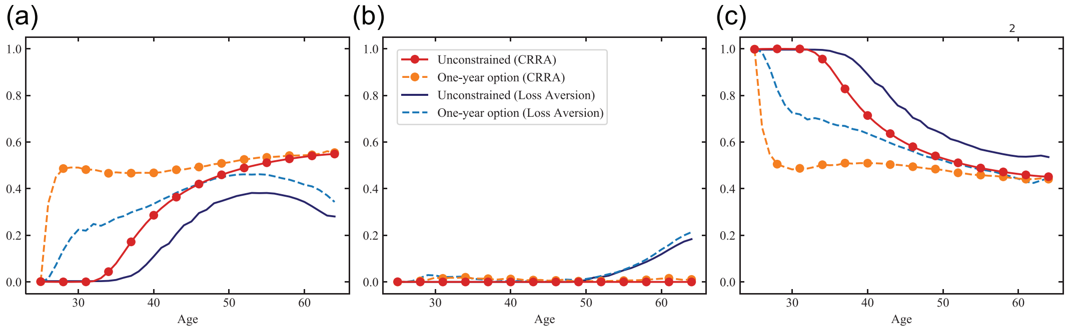

FIGURE 4 Mean allocations over the investment phase. (a) Mean allocation to nominal bond (S_0), (b) mean allocation to inflation-linked bond (S_1), and (c) mean allocation to risky stock (S_2). These figures show the mean allocations to each asset over the investment phase of the investor for the four investment strategies. The initial increase in allocation to the nominal bond and reduction in allocation to the risky stock observed in the 1-year option strategies (for both CRRA and loss aversion utility preferences) is a result of our simulation parameter of $F_{0-} = 0$ which places the investor closer to the lower boundary constraint at the left-side of Figure 2. The mean allocations of the constrained CRRA strategy moves closer towards its unconstrained counterpart compared to the loss aversion strategies because over time fewer of the CRRA simulations are bounded by the constraints (see Section 2.5.1). CRRA, constant relative risk aversion

comes into effect and the resultant strategies provide little improvement in the lowest percentiles of RRRs at retirement, especially under the CRRA risk preference. We also conduct sensitivity analyses to examine the impact of boundary constraints have on the retirement outcomes with different levels of the boundary constraints and different utility parameters (see Supporting Information Appendix A). The boundary constraints have impact on the shape as well as the key statistics of the simulated distributions. A smaller distance between the boundary constraints results in larger changes in the retirement outcomes, and vice versa. Furthermore, adding in constraints has a greater effect on less risk averse investors implied by the utility function parameters alone, since it limits their otherwise higher exposure to the risky asset.

3.3.2 | Distribution under historical UK market return data

This section analyzes the performance of the optimized strategies when the market returns do not conform to the model set out in Section 2.1. Specifically, we use the moving block bootstrap method to simulate a series of retirement outcomes using the historical UK return data. This method preserves the cross-sectional correlation and the serial dependence of the original data. To begin, we randomly draw a starting date with replacement. Next, for each investment strategy, we simulate the retirement outcomes (RRRs) after 40 years of investment following the starting date, with a starting fund level of $F_{0-} = 0$ units and starting labour income of $Y_0 = 1$ unit at age 25. This procedure is then repeated 10,000 times. Histograms of the distribution of the RRR are shown in Figure 5 and in Table 4 their summary statistics.

Since there were no index-linked gilts issued before 1981, we have assumed the existence of inflation-indexed bonds before 1981 which have a constant real rate of return $r_R = 0.026$ as discussed in Section 3.1. Furthermore, for consistency with this assumption and with the

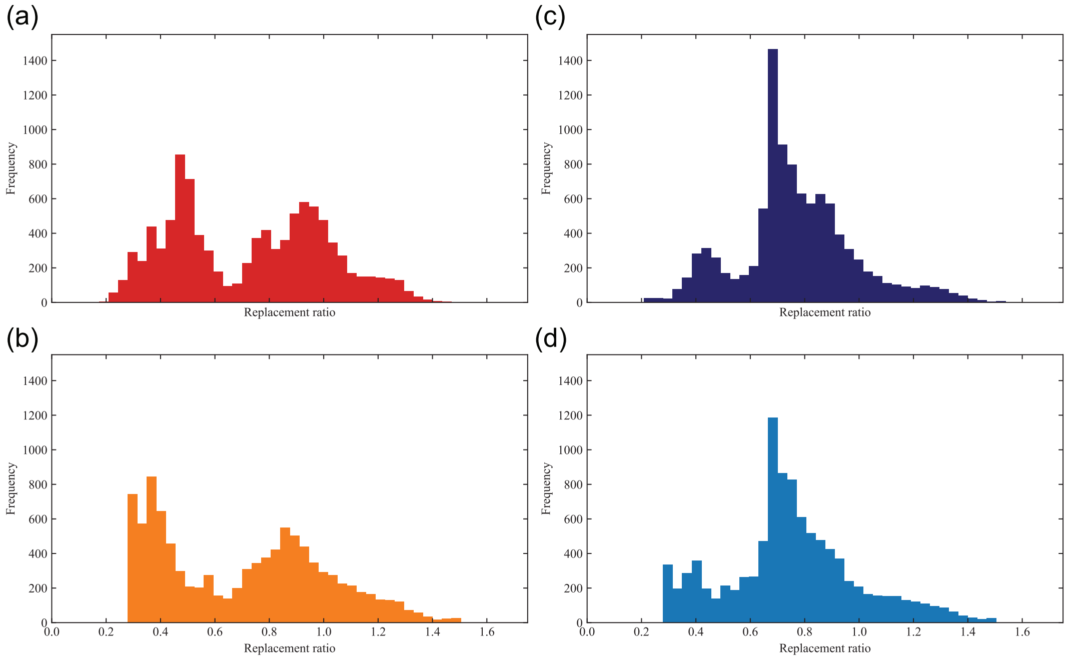

FIGURE 5 Histrograms of simulated real replacement ratios under historical simulations. (a) Unconstrained (CRRA), (b) unconstrained (loss aversion), (c) 1-year option (CRRA), and (d) 1-year option (loss aversion). These figures show the histograms of 10, 000 simulated real replacement ratios at time T for an investor following either one of the four strategies under historical simulations (the figures are right-truncated at the value of 1.75). CRRA, constant relative risk aversion

constant annuity rate assumption (which we detail in the next paragraph), we assume that the real rate of return for inflation-indexed bonds issued since 1981 is also r_R. This is a critical assumption that allows us to do the historical simulation, as otherwise we do not have a single set of 40-year data which rely on actual inflation-indexed bond returns. However, we acknowledge that this limits the strength of the conclusions drawn from the historical simulation exercise.

The second assumption is that the annuity rate is a constant value $\ddot{a}_{65}$ of 14.3779 at all times (see Section 3.1). We were not able to obtain data on the cost of inflation-indexed annuities over any reasonable time period. Keeping the same annuity rate for all simulations imply that the resultant real replacement ratio (RRR) is a constant multiple of the investor's real fund value. With these two assumptions, we use the historical UK returns data for the period from December 1918 to December 2019 that encompasses a wide range of market conditions, including periods with different levels of interest rates as well as the financial market crisis of 2008.[6]

Under the historical simulations, the CRRA risk preferences lead to bimodal distributions. The bimodality is largely attributed to the investment period of the simulations. Specifically, the simulations whose retirement date is before mid-1980s typically achieve lower RRRs and

[6]The historical simulation data is created by splicing the DataStream data from 1980 onwards with long-term monthly indices supplied by David Wilkie, of whom the authors are very appreciative, from 1918 to 1980. We do this for the inflation data, the nominal bond data and the risky stock data. The data is available from the authors upon request.

TABLE 4 Comparison of outcomes of investment strategies under historical simulations

This table reports the summary statistics of the outcomes distributions following each of the four investment strategies under historical simulations. We report the mean, the standard deviation (SD), CVaR (Conditional Value at Risk, a risk measure which we define as the weighted average of the lowest 1% of the RRRs), the probability of failing to achieve the lower boundary, and the probability of achieving the target RRR. Abbreviations: CRRA, constant relative risk aversion; RRR, real replacement ratio. [a]Since there is no explicit target for the CRRA investor, we report the probability of achieving the same target as the loss averse investor, namely $k = 2/3$.

	Mean	SD	CVaR	$\mathbb{P}[\text{RRR} < l]$	$\mathbb{P}[\text{RRR} \geq k]$
CRRA, unconstrained	0.729	0.280	0.239	3.6%	55.2%[a]
CRRA with 1-year options	0.705	0.301	0.300	–	54.5%[a]
Loss aversion, unconstrained	0.770	0.214	0.277	0.6%	76.0%
Loss aversion with 1-year options	0.754	0.243	0.304	–	70.3%

cluster around the first peak of the distribution because of the low real returns of the nominal bonds over the period. The investor with CRRA risk preferences is susceptible to the risk of achieving a low RRR at retirement, with a probability of 3.6% failing to achieve at least the lower bound l and a low value of CVaR at 0.239. Adding in the boundary constraints provides some extent of downside risk protection by removing the chance of outcomes below l and lifting the CVaR to the value of l. It shifts the distribution slightly towards the left (as seen in a reduction of 0.024 in the mean) and pushes it against the lower boundary of l.

On the other hand, the outcomes distributions of loss aversion risk preferences show mild sign of bimodality whilst retaining much resemblance of their counterparts as seen in Section 3.3.1, with bunching of outcomes around the investor's target. This demonstrates that the loss aversion investment strategy successfully keeps the majority of the outcomes around the target, including simulations that retire before mid-1980s. Having the highest mean, the lowest standard deviation and the lowest $\mathbb{P}[\text{RRR} < l]$, the unconstrained loss aversion strategy appears to be the best performing strategy among the four strategies. The relative low risk of achieving RRRs below l can be eliminated by imposing the boundary constraints, which also increases the CVaR to a value slightly above l.

In a market model that does not conform to the model set out in Section 2.1, the risk averse investor is more susceptible to achieving very poor RRRs.[7] This may incentivize a risk averse investors to include the boundary constraints—the downside protection is more likely to come into effect preventing the worst of retirement outcomes. In contrast, the loss averse risk preference prescribes asset allocation depending on the cumulative investment gain and loss (relative to the target). This investment strategy leads to largely robust retirement outcomes even when the realized market returns deviate from the assumptions under which the strategy is derived (comparing Figures 3c and 5c). In both cases, it gives the investor a higher chance of

[7]We perform a robustness check by drawing the starting dates from a smaller range of dates. Under CRRA risk preference, the risks of achieving very poor RRRs are related to the (cumulative) asset returns during the simulation period—the risks are particularly high for simulations that end before mid-1980s. This relationship is not apparent under the loss aversion risk preference.

achieving the target RRR. As a result, the need of downside protection is not nearly as material for the investor following a loss aversion utility function compared to the CRRA investor.

4 | CONCLUSION

In this paper, we have studied the investment strategies that allow the investor to constrain their retirement incomes, or equivalently financial wealth accumulated after the savings phase, within a framework of expected utility maximization. We contribute to the literature by extending Donnelly et al. (2018) to allow the investor to have a non-tradable stochastic labour income, to have wealth boundary constraints that evolve according to the labour income process, and to use short term option contracts in managing their investment risks. This differentiates our study from most literature in which a tradable labour income process is assumed and continuously-traded replicating strategies are devised to replicate the option-like payoffs for the investors. In addition to the traditional risk aversion utility preferences, we also investigate the investment behaviour implied by the loss aversion risk preference that has been attracting the interest of both practitioners and academic researchers.

We find that there is no strong incentive for CRRA investor to include terminal wealth constraints under the standard market model. Constraints are useful for the very risk averse investor, since they protect them from extremely adverse outcomes. Thus they give more certainty in the retirement outcome. Yet this comes at a potential cost of investing in low return bonds for long periods of time, to avoid breeching the constraints.

On the other hand, the loss aversion risk preferences naturally lead to a high chance of achieving the investment target, when the latter evolves according to the labour income of the investor. Our analysis shows that the investor's retirement outcomes are robustly centred around the investment target even when the realized market returns do not conform to the assumptions under which the strategies are derived. The terminal wealth constraints do little to increase the certainty of income in retirement. Indeed, similarly to the CRRA framework, they result in a strategy that is too risk averse. Instead, the investor can benefit from following a loss aversion-derived strategy in conjunction with an appropriate target income. This leads to an optimized strategy which naturally focuses on the desired income level and keeps the retirement income above the lower constraint in the vast majority of scenarios.

ACKNOWLEDGMENTS

Open access publishing facilitated by Australian National University, as part of the Wiley - Australian National University agreement via the Council of Australian University Librarians.

DATA AVAILABILITY STATEMENT

The data was shared by Prof. David Wilkie (Department of Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University) on request.

REFERENCES

Andréasson, J. G., Shevchenko, P. V., & Novikov, A. (2017). Optimal consumption, investment and housing with means-tested public pension in retirement. *Insurance: Mathematics and Economics*, *75*, 32–47.

Berkelaar, A. B., Kouwenberg, R., & Post, T. (2004). Optimal portfolio choice under loss aversion. *Review of Economics and Statistics*, *86*(4), 973–987. https://doi.org/10.1162/0034653043125167

Blake, D., Cairns, A., & Dowd, K. (2007). The impact of occupation and gender on pensions from defined contribution plans. *The Geneva Papers on Risk and Insurance—Issues and Practice*, *32*(4), 458–482. https://doi.org/10.1057/palgrave.gpp.2510141

Blake, D., Wright, D., & Zhang, Y. (2013). Target-driven investing: Optimal investment strategies in defined contribution pension plans under loss aversion. *Journal of Economic Dynamics and Control*, *37*(1), 195–209. https://doi.org/10.1016/j.jedc.2012.08.001

Boyle, P. P. (1988). A lattice framework for option pricing with two state variables. *The Journal of Financial and Quantitative Analysis*, *23*(1), 1. https://doi.org/10.2307/2331019

Butt, A., Khemka, G., & Warren, G. J. (2019). What dividend imputation means for retirement savers. *Economic Record*, *95*(309), 181–199. https://doi.org/10.1111/1475-4932.12468

Byrne, A., Blake, D., Cairns, A., & Dowd, K. (2007). Default funds in U.K. defined-contribution plans. *Financial Analysts Journal*, *63*(4), 40–51. https://doi.org/10.2469/faj.v63.n4.4748

Cairns, A. J., Blake, D., & Dowd, K. (2006). Stochastic lifestyling: Optimal dynamic asset allocation for defined contribution pension plans. *Journal of Economic Dynamics and Control*, *30*(5), 843–877. https://doi.org/10.1016/j.jedc.2005.03.009

Chen, Z., Li, Z., Zeng, Y., & Sun, J. (2017). Asset allocation under loss aversion and minimum performance constraint in a DC pension plan with inflation risk. *Insurance: Mathematics and Economics*, *75*, 137–150. https://doi.org/10.1016/j.insmatheco.2017.05.009

Dong, Y., Lv, W., Wei, S., & Gong, Y. (2020). Optimal investment of DC pension plan under incentive schemes and loss aversion. *Mathematical Problems in Engineering*, *2020*, 1–14. https://doi.org/10.1155/2020/5145848

Dong, Y., & Zheng, H. (2019). Optimal investment of DC pension plan under short-selling constraints and portfolio insurance. *Insurance: Mathematics and Economics*, *85*, 47–59. https://doi.org/10.1016/j.insmatheco.2018.12.005

Donnelly, C., Gerrard, R., Guillén, M., & Nielsen, J. P. (2015). Less is more: Increasing retirement gains by using an upside terminal wealth constraint. *Insurance: Mathematics and Economics*, *64*, 259–267. https://doi.org/10.1016/j.insmatheco.2015.06.003

Donnelly, C., Guillén, M., Nielsen, J. P., & Pérez-Marín, A. M. (2018). Implementing individual savings decisions for retirement with bounds on wealths. *ASTIN Bulletin*, *48*(1), 111–137. https://doi.org/10.1017/asb.2017.34

Fischer, S. (1978). Call option pricing when the exercise price is uncertain, and the valuation of index bonds. *Journal of Finance*, *33*(1), 169–176. https://doi.org/10.1111/j.1540-6261.1978.tb03396.x

Gomes, F., & Michaelides, A. (2003). Portfolio choice with internal habit formation: Alife-cycle model with uninsurable labor income risk. Review of Economic Dynamics, *6*(4), 729–766. https://doi.org/10.1016/S1094-2025(03)00059-0

Guan, G., & Liang, Z. (2016). Optimal management of DC pension plan under loss aversion and value-at-risk constraints. *Insurance: Mathematics and Economics*, *69*, 224–237. https://doi.org/10.1016/j.insmatheco.2016.05.014

Han, N. W., & Hung, M. W. (2012). Optimal asset allocation for DC pension plans under inflation. *Insurance: Mathematics and Economics*, *51*(1), 172–181. https://doi.org/10.1016/j.insmatheco.2012.03.003

He, X. D., & Zhou, X. Y. (2011). Portfolio choice under cumulative prospect theory: An analytical treatment. *Management Science*, *57*(2), 315–331. https://doi.org/10.1287/mnsc.1100.1269

Horneff, V., Maurer, R., Mitchell, O. S., & Rogalla, R. (2015). Optimal life cycle portfolio choice with variable annuities offering liquidity and investment downside protection. *Insurance: Mathematics and Economics*, *63*, 91–107. https://doi.org/10.1016/j.insmatheco.2015.03.031

Hull, J. C. (2019). Options, futures and other derivatives (10 ed.). Pearson.

Jin, H., & Zhou, X. Y. (2008). Behavioral portfolio selection in continuous time. *Mathematical Finance*, *18*(3), 385–426. https://doi.org/10.1111/j.1467-9965.2008.00339.x

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263. https://doi.org/10.2307/1914185

Khemka, G., & Butt, A. (2017). Non-parametric integral estimation using data clustering in stochastic dynamic programming: An introduction using lifetime financial modelling. *Risks*, *5*(4), 57. https://doi.org/10.3390/risks5040057

Korn, R. (2005). Optimal portfolios with a positive lower bound on final wealth. *Quantitative Finance*, *5*(3), 315–321. https://doi.org/10.1080/14697680500167927

Korn, R., & Krekel, M. (2003). Optimal portfolios with fixed consumption or income streams. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.938215

Kraft, H., & Steffensen, M. (2013). A dynamic programming approach to constrained portfolios. *European Journal of Operational Research*, *22*(2), 453–461. https://doi.org/10.1016/j.ejor.2013.02.039

Leland, H. E. (1985). Option pricing and replication with transactions costs. *Journal of Finance*, *40*(5), 1283–1301. https://doi.org/10.1111/j.1540-6261.1985.tb02383.x

Lin, H., Saunders, D., & Weng, C. (2017). Optimal investment strategies for participating contracts. *Insurance: Mathematics and Economics*, *73*, 137–155. https://doi.org/10.1016/j.insmatheco.2017.02.001

Mahayni, A., & Schneider, J. C. (2015). Minimum return guarantees, investment caps, and investment flexibility. *Review of Derivatives Research*, *19*(2), 85–111. https://doi.org/10.1007/s11147-015-9116-5

Merton, R. C. (2014). The crisis in retirement planning. Harvard Business Review.

Mitchell, O. S., & Utkus, S. P. (2004). *Pension design and structure: New lessons from behavioral finance*. Oxford University Press.

Schütte, L. (2017). *Retirement wealth under fixed limits: The optimal strategy for exponential utility* (Master's thesis). University of Basel. https://arxiv.org/abs/1712.00463

Teplá, L. (2001). Optimal investment with minimum performance constraints. *Journal of Economic Dynamics and Control*, *25*(10), 1629–1645. https://doi.org/10.1016/s0165-1889(99)00066-4

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. https://doi.org/10.1007/bf00122574

van Bilsen, S., & Laeven, R. J. (2020). Dynamic consumption and portfolio choice under prospect theory. *Insurance: Mathematics and Economics*, *91*, 224–237.

Vanguard, & Nest Insight. (2019). *How the UK saves 2019*. Technical report.

Warren, G. J. (2019). Choosing and using utility functions in forming portfolios. *Financial Analysts Journal*, *75*(3), 39–69. https://doi.org/10.1080/0015198x.2019.1618109

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Donnelly, C., Khemka, G., & Lim, W. (2022). Investing for retirement: Terminal wealth constraints or a desired wealth target? *European Financial Management*, 28, 1283–1307. https://doi.org/10.1111/eufm.12351