



Delft University of Technology

Rule-based strategies for dynamic life cycle investment

den Haan, T. R.B.; Chau, K. W.; van der Schans, M.; Oosterlee, C. W.

DOI

[10.1007/s13385-021-00283-0](https://doi.org/10.1007/s13385-021-00283-0)

Publication date

2022

Document Version

Final published version

Published in

European Actuarial Journal

Citation (APA)

den Haan, T. R. B., Chau, K. W., van der Schans, M., & Oosterlee, C. W. (2022). Rule-based strategies for dynamic life cycle investment. *European Actuarial Journal*, 12(1), 189-213. <https://doi.org/10.1007/s13385-021-00283-0>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Rule-based strategies for dynamic life cycle investment

T. R. B. den Haan^{1,2} · K. W. Chau³ · M. van der Schans¹ · C. W. Oosterlee^{4,5}

Received: 10 November 2020 / Revised: 14 April 2021 / Accepted: 5 May 2021 /
Published online: 31 May 2021
© The Author(s) 2021

Abstract

In this work, we consider rule-based investment strategies for managing a defined contribution pension savings scheme, under the Dutch pension fund testing model. We find that dynamic, rule-based investment strategies can outperform traditional static strategies, by which we mean that the investor may achieve the target retirement income with a higher probability or limit the shortfall when the target is not met. In comparison with dynamic programming-based strategies, the rule-based strategies have more stable asset allocations throughout time and avoid excessive transactions that may be hard to explain to an investor. We also study a combined strategy of a rule-based target with dynamic programming. A key feature of our setting is that there is no risk-free asset, instead, a matching portfolio is introduced for the investor to avoid unnecessary risk.

Keywords Life cycle investing · Pensions · Defined contribution · Rule-based strategies · Dynamic programming

1 Introduction

Nowadays, many people invest their retirement savings in a defined contribution pension scheme. In such a scheme, the contributions are agreed upon and are, e.g., a percentage of one's salary. The pension, however, is uncertain as it depends on the returns on investment. At retirement, the accumulated wealth is

✉ K. W. Chau
k.w.chau@rug.nl

¹ Ortec Finance, Rotterdam, The Netherlands

² Delft University of Technology, DIAM-Delft Institute of Applied Mathematics, Delft, The Netherlands

³ Department of Economics, Econometrics and Finance, University of Groningen, Groningen, The Netherlands

⁴ Mathematical Institute, Utrecht University, Utrecht, The Netherlands

⁵ CWI-Center for Mathematics and Computer Science, Amsterdam, The Netherlands

converted to a pension income that intends to replace a proportion of the investor's income, typically about 70%, which is referred to as the replacement ratio. In this paper, we propose a dynamic strategy that optimally steers the investor towards a replacement ratio target.

Our dynamic strategy will reduce risk after several years of good returns on investment. It presumes that upward potential concurs with downside risk. Our pension investor is only interested in reaching her replacement ratio target, i.e., not making the target is considered downside risk and she feels indifferent about any two values above the target. We will show that, in this sense, the designed dynamic strategy outperforms static life cycle strategies. By decreasing risk after several good years, our dynamic strategy prevents unnecessary risk taking.

A well-known static life cycle strategy is known as Bogle's rule [5], which prescribes to invest a percentage of 100 minus one's age in risky assets. Decreasing risk in the course of the life cycle in such a way is called a glide path. When the glide path is known in advance up to retirement, the strategy is static and does not adjust as events unfold. Therefore, static strategies may take unnecessary risk when returns on investment are better than anticipated, see [1, 12] for a discussion of drawbacks of static life cycle strategies. The strategy we propose is also rule-based, but it is dynamic as the prescribed rule depends on events that still have to unfold.

In the literature, dynamic strategies are often studied in the context of dynamic programming [3]. Dynamic programming optimizes the investment strategy backwards in time by optimizing decisions for the coming period given that consecutive decisions are already taken optimally. Merton [17] was the first to apply dynamic programming to an asset allocation problem with two assets, a risky and a risk-free asset, also allowing for consumption during the investment period. Optimal decisions were based on the constant relative risk aversion utility function. He showed in [17] that the optimal strategy continuously rebalances, i.e., the optimal allocation is constant.

The literature on optimal asset allocation is rich, and we cite here some contributions that influenced our work. Li and Ng [14] introduced mean-variance strategies with respect to a wealth target. The wealth target then allows the investor to identify a surplus: wealth up to the target may be invested in stocks, any remainder is invested in the risk-free rate. Zhang et al. [20] solved a similar problem, although utility-based, and combined dynamic programming with the least squares Monte Carlo method. Upper and lower bounds for the wealth were prescribed in that paper, showing that upward potential comes with downside risk. Terminal wealth is steered towards a desired range by investing the difference between a risk-free-discounted upper bound value and the current wealth in the risk-free asset. Forsyth and Vetzal [11] also applied dynamic programming and used a PDE solver to solve a so-called time-consistent mean-variance problem, meaning that similar mean-variance problems were solved at future times. In addition to mean-variance that balances the mean and variance of returns, they studied a problem with a fixed wealth target. To reduce risk, both [11] and [20] proposed to invest excess wealth in a risk-free asset.

Similarly, the rule-based strategies introduced in this paper will invest excess wealth into a so-called matching portfolio. Compared with static strategies, distributions of outcomes are more centered around the target value and the area below the target value will become smaller.

Besides many positive aspects, dynamic programming and its resulting strategies have several drawbacks including the following. Firstly, dynamic programming is computationally rather intensive. Secondly, the corresponding investment decisions can be sensitive to small changes in parameters and underlying assumptions. Because of this, the allocation may fluctuate over time resulting in large turnovers, of which, from a practical perspective, it is hard to explain why they are required. Intuitively defined rules, typically, do not suffer from these drawbacks. Finally, it is not straightforward to apply dynamic programming to the pension settings as an investor's replacement ratio target often depends on inflation influencing the future, which in turn also influences the future contributions.

Rule-based dynamic strategies fall in between the static and dynamic programming paradigms, when well constructed they aim for the best of both worlds. As shown by Basu et. al. [2], even simple rule-based strategies that reduce risk half way in the life cycle can outperform static life cycle strategies. Compared with the work in [2], our rule-based strategies can reduce risk annually, and consider the market price of future pension payments instead of a wealth target. In addition to the rule-based strategies, we will study an integrated approach in which we combine a rule-based strategy with dynamic programming.

2 The optimal asset allocation problem

2.1 The Dutch pension system

To demonstrate the rule-based strategy's practical value, we consider a Dutch pension investor, and use typical retirement data from the Netherlands. In the Netherlands, the pension system can roughly be described as follows. The system consists of three pillars. The first pillar is formed by government allowances for old age equal to about 70% of minimum wage. It is funded out of current payroll taxes. People earning more than a so-called franchise, roughly equal to minimum wage and amounting to 13,123 Euro ultimo 2016, can take part in the second or third pillar. The second pillar consists of participation in collective pension schemes. Contributions are equal to a percentage of one's pensionable salary, i.e., the salary above the franchise, and employers typically pay part of the contributions which makes opting out highly unfavourable. The third pillar consists of private pension products. At retirement, pension benefits and savings are converted to a period income, e.g. facilitated through the collective pension scheme, a life insurance product of an insurance company, or through periodic withdrawals from an investment account. The overall aim is to replace about 70% of one's income by a pension. In the second and third pillars, Dutch pension law facilitates tax free contributions to work towards this aim.

2.2 Model setting

In this paper, the Dutch pension investor doesn't take part in a collective pension scheme, but is assumed to have full control over her pension savings, either by taking part in an individually defined contribution pension scheme in the second pillar, or by means of an appropriate private pension product in the third pillar. At $t = 0$, the 26 year old investor will start saving annually the maximum allowed tax free amount (according to Dutch pension law) up to retirement at time $t = T$, coinciding here with a retirement age of 67 years. She intends to replace 70% of her salary by her pension (including government allowances for old age). As the government allowance for old age roughly equals 70% of the franchise, her pension savings intend to replace 70% of her pensionable salary, i.e., her salary in excess of the franchise. Although, in practice, an investor might be interested in insuring longevity risk or be interested in employing advanced withdrawal strategies, Blanchett et al. [4] illustrates that simple withdrawal strategies can perform well, e.g., based on an annuity with a maturity roughly equal to an investor's life expectancy. Therefore, as we focus on accumulating wealth before retirement, we simply assume the investor buys an annuity that indexes with the expected inflation, i.e., a bond which, apart from indexation for expected inflation, equals annual payoffs, for a period of N , say 20, years after retirement. Whichever withdrawal strategy an investor might follow, the assumption here is that this annuity gives a good estimate of, at least, the investor's income in her first year after retirement, and, thereby, to what extent she can replace her pensionable salary for 70% with a pension.

The investor can invest her wealth W_t in a risky, equity-like, asset, which is called the return portfolio, or in a safe, bond-like, asset with annual payoffs during retirement, the matching portfolio. In our setting, the strategy will use the matching portfolio to protect the current gains, and it grows with inflation. Therefore, the matching portfolio also carries risk. Put differently, we assume the investor doesn't hedge inflation risk with inflation protected securities as the market for inflation products is illiquid and strategies that hedge against inflation are not straightforward to follow in practice [16]. Finally, we assume there is no risk-free rate to invest money in.

The investor annually manages her portfolio, i.e., decisions, contributions and pension payments are made in discrete time, which runs up to retirement, from $t = 0$ to $t = T$. The pension payments start at $t = T$ and run up to $t = T + N - 1$. At time $t \leq T$ before retirement, she invests a fraction α_t of her wealth W_t in the return portfolio. The investor is not allowed to short-sell assets or borrow money, so that

$$0 \leq \alpha_t \leq 1. \quad (1)$$

In the dynamic programming literature, α_t is referred to as the control (as decisions intend to give the investor control over the outcome). A strategy maps information Z_t available at time t , e.g., past returns and current wealth W_t , to the desired allocation:

$$\alpha_t : \mathbb{R}^K \ni Z_t \mapsto \alpha_t(Z_t) \in [0, 1]. \quad (2)$$

Here Z_t is adapted to a filtration \mathcal{F}_t , governing the underlying stochastic processes. Before time t , the information Z_t is not yet available, and α_t is thus a stochastic

quantity. In a static strategy, such as Bogle’s rule, α_t only depends on time and is known, i.e., not stochastic, even when the information Z_t is not yet available. In practice, risk is reduced towards retirement, meaning that α_t typically decreases over time.

Just before rebalancing, the investor makes a contribution c_t to the portfolio equal to the maximum allowed tax free amount under Dutch pension law. These contributions resemble an age-dependent percentage p_t , see Table 3, of the investor’s pensionable salary s_t which she earned in the period $t - 1$ up to t . We assume that the investor’s salary follows a deterministic career path, i.e., it increases with age. The investor’s salary also increases stochastically with the wage inflation w_t . Altogether, Table 4 shows the expectation and standard deviation of the annual contribution to the portfolio for our investor, see also Appendix A. Note that the contributions merely depend on career path assumptions and inflation, and do not depend on investment decisions neither on return on investment.

The investor’s objective is to achieve a 70% replacement ratio target at retirement without incurring too much downside risk. The replacement ratio for the pensionable salary at retirement, RR_T , is given by

$$RR_T = \frac{W_T}{M_T} \cdot \frac{T + 1}{\sum_{t=0}^T s_t \prod_{\tau=t+1}^T (1 + \pi_\tau)}. \tag{3}$$

Here, the second term divides by the investor’s average wage in nominal amounts indexed with inflation π_t to retirement at $t = T$, and M_t is the market value factor. The market value factor M_t converts a nominal amount at time t , that is supposed to have similar purchasing power to the pension payments, to N pension payments at times $\tau \geq T$ after retirement by indexing with currently expected inflation, and then discounting those pension payments back to time $t \leq T$:

$$M_t = \sum_{\tau=T}^{T+N-1} (1 + r_t^{\tau-t})^{t-\tau} \mathbb{E}_t \left[\prod_{\tau'=t+1}^{\tau} (1 + \pi_{\tau'}) \right], \tag{4}$$

where \mathbb{E}_t is the expectation, conditional on \mathcal{F}_t (i.e., the information available at time t), and $r_t^{\tau-t}$ represents the market rates that discount payments from $\tau - t$ years into the future back to the present time. Put differently, dividing current wealth by the market value factor gives an amount that currently has similar purchasing power as the future pension. Using the market value factor M_T at retirement, the first term in (3) converts the accumulated wealth W_T into N annual income payments indexed for expected future inflation.

To measure whether a strategy achieves the investor’s objective, we use a utility function, U , which, whenever decisions are to be taken, intends to maximize the following expression in expectation:

$$\max_{\alpha_t, \dots, \alpha_{T-1}} \mathbb{E} [U(Z_T) | \mathcal{F}_t], \tag{5}$$

where \mathcal{F}_t again represents current market information, α_t is as in (2) and Z_T is a vector with outcomes including the terminal replacement ratio. Although other choices

are possible, we choose $U(\cdot)$ to be the shortfall below the investor's target replacement ratio of 70%:

$$U(Z_T) = \min(RR_T - 70\%, 0), \quad (6)$$

where there is no shortfall in replacement ratio if it ends above 70%. Note that this measure is not conditional on the shortfall. So, additionally, we will also evaluate a strategy's performance using the 10% conditional value at risk $CVaR_{0.1}(RR_T)$ of the replacement ratio, i.e., the expectation of the 10% worst case outcomes as defined by

$$CVaR_\alpha(RR_T) = \mathbb{E} \left[RR_T \mid RR_T \leq F_{RR_T}^{-1}(\alpha) \right], \quad (7)$$

where $F_{RR_T}^{-1}(\alpha)$ is the inverse cumulative distribution function of the terminal replacement ratio RR_T and represents the α -th quantile below which are the worst case outcomes.

2.3 Governing stochastic model

For general applicability, we require the designed strategies are not defined in terms of the governing stochastic model parameters. That is, the strategies can be applied when different governing stochastic models would be used. We merely assume that the governing stochastic model can be simulated by means of a Monte Carlo simulation. To make this explicit, we choose to use a standard model developed to make risk analyses comparable between Dutch pension funds, see [13]. The model and its calibration are well documented [10]. Calibration to recent market data and a Monte Carlo simulation of the model are publicly available at the website of the Dutch Central Bank [9]. In this paper, we use the data set of 2017 (quarter 1), which is calibrated to data up to ultimo 2016 and start simulating from there.

In discrete time, the model is a VAR(1) model with normally distributed increments, see [18] for a short summary of the model specification. In the calibration, some structure is imposed to achieve realistic market dynamics. Based on the model, sample paths are generated for the following variables:

- Equity returns x_t , which are used for the return portfolio;
- Inflation π_t ;
- Wage inflation w_t , which equals inflation π_t plus 0.5%;
- A yield curve with interest rates r_t^m containing rates for each maturity m .

The matching portfolio is tailored to the investor's retirement age. Its returns m_t equal the rate of change in the market value factor:

$$m_t = \frac{M_t}{M_{t-1}} - 1, \quad (8)$$

where M_t is defined in (4). Note that the matching portfolio protects the investor against *expected* future inflation. To determine the expected future inflation, we use the least squares Monte Carlo technique, as presented in Sect. B.3.

Table 1 gives the annual return statistics of the variables. Due to the fluctuating market price of future pension payments, the standard deviation of the matching returns is very similar to the one of the equity returns. Although the matching portfolio follows these fluctuations, it is considered less risky, in terms of the investor’s goals. By investing in the matching portfolio, the pensioner will receive the corresponding amount from the annuity, no matter the future market prices.

3 Rule-based strategies

In this section, we define three rule-based strategies: a cumulative target strategy that decreases risk once it reaches a cumulative target for the contributions paid so far, an individual target strategy that tracks the investments of the contributions separately and decreases risk once it reaches the target for that contribution, and a combination strategy that combines the two with dynamic programming. The strategies all intend to steer towards a replacement ratio of 70%, and decrease risk when return on investment develops well. The strategies differ in their views on when return on investment has been developing well enough to decrease risk.

3.1 Cumulative target strategy

The cumulative target strategy that we consider here has similarities with the strategies studied in [20] and [11]: risk is reduced once wealth exceeds a pre-defined wealth target. Contrary to [20] and [11], however, our investor saves for retirement and we relate the wealth target to the price of a bond with payoff equal to the desired pension.

Given a density forecast for the matching and return portfolios, see Sect. 2.3, the strategy depends on two parameters: a required real rate of return r (before retirement) and a discount rate δ (after retirement) to discount pension payments after retirement to the retirement date. At time t before retirement, i.e., $t = 0, \dots, T$, the investor contributes c_t to her pension savings, see Tables 3 and 4. The contributions c_τ up to time t , i.e., $\tau = 0, \dots, t$, are supposed to grow with inflation π_τ , plus the real rate of return r , to a target wealth $c_\tau \mathbb{E}_t[F_\tau]$ at retirement, where F_τ is given by

Table 1 Annual statistics of the underlying stochastic model calculated on a sample that combines all sample paths

	x_t	m_t	r_t^{10}	π_t	w_t
Mean	6.1%	3.4%	2.5%	1.6%	2.1%
Standard deviation	18.3%	18.5%	2.4%	1.5%	1.5%
Correlations					
Equity return (x_t)	1.00				
Matching return (m_t)	- 0.06	1.00			
10 year interest (r_t^{10})	0.17	- 0.17	1.00		
Inflation (π_t)	0.11	- 0.04	0.82	1.00	
Wage inflation (w_t)	0.11	- 0.04	0.82	1.00	1.00

$$F_\tau = \prod_{\tau'=\tau+1}^T (1 + r + \pi_{\tau'}), \tag{9}$$

and the conditional expectation, \mathbb{E}_t , enforces that the realized inflation is used before time t and the expected inflation is used beyond time t . The wealth targets at retirement for all contributions c_τ up to time t are combined and converted into a target pension using a discount factor \tilde{M}_T with discount rate δ :

$$\tilde{M}_T = \sum_{\tau=T}^{T+N-1} \frac{1}{(1 + \delta)^{\tau-T}}. \tag{10}$$

Note that, contrary to the market value factor M_t as defined in (4), \tilde{M}_T is merely an annuity with rate δ used to approximate the first pension given wealth at retirement. Using the market value factor M_t , this gives us the following current target wealth \tilde{W}_t :

$$\tilde{W}_t = \frac{M_t}{\tilde{M}_T} \sum_{\tau=0}^t c_\tau \mathbb{E}_t[F_\tau], \tag{11}$$

where the summation represents the combined wealth targets at retirement for all contributions c_τ up to time t . In short, (11) can be interpreted as follows: the summation converts contributions c_τ up to time t into a target wealth at retirement, \tilde{M}_T uses an annuity with rate δ to approximate the first pension, and, finally, the market value factor M_t indexes the first pension with expected inflation and discounts all future pensions back to time t . Altogether, this gives the current target wealth \tilde{W}_t .

The cumulative target strategy starts by investing new contributions c_t in the risky asset. If the current wealth W_t , including the current contribution c_t , exceeds the target wealth \tilde{W}_t , risk is reduced and W_t is transferred to the matching portfolio. For the matching portfolio, the investor follows a buy and hold strategy. New contributions invested in the risky asset, will also be transferred to the matching portfolio if the current wealth W_t , which consists of the current contribution c_t , the value of the matching portfolio and the value of the return portfolio, exceeds the target wealth \tilde{W}_t . In other words, at $t = 0$, the control α_0 is given by

$$\alpha_0 = \begin{cases} 0 & \text{if } W_0 \geq \tilde{W}_0, & (12a) \\ 1 & \text{otherwise,} & (12b) \end{cases}$$

and, for $t = 1 \dots T$, the control α_t is given by

$$\alpha_t = \begin{cases} 0 & \text{if } W_t \geq \tilde{W}_t, & (13a) \\ \frac{\alpha_{t-1}(1+x_t)}{\alpha_{t-1}(1+x_t)+(1-\alpha_{t-1})(1+m_t)} & \text{otherwise.} & (13b) \end{cases}$$

3.2 Individual target strategy

Contrary to the cumulative target strategy, the individual target strategy, which is the second strategy we will analyze here, defines a wealth target per contribution and invests each contribution separately, i.e., the wealth W_t is seen a sum of the individual wealth components resulting from investing the contributions separately:

$$W_t = \sum_{\tau=0}^t W_{t,\tau}, \tag{14}$$

where $W_{t,\tau}$ is the wealth component from investing the contribution c_τ . As in (11), a wealth target $\tilde{W}_{t,\tau}$, at time t for a contribution invested at time $\tau \leq t$, is given by

$$\tilde{W}_{t,\tau} = \frac{M_t}{\tilde{M}_T} c_\tau \mathbb{E}_t[F_\tau]. \tag{15}$$

Apart from this, the strategy works similarly: the individual contributions are invested in the risky asset until the invested amount exceeds the wealth target for that contribution, in which case they are transferred to the matching portfolio until retirement. Thus, the control $\alpha_{t,\tau}$ for investing contribution $c_{t,\tau}$ is given by

$$\alpha_{t,\tau} = \begin{cases} 0 & \text{if for any } \tau' = \tau \dots t \text{ we have } W_{\tau',t} \geq \tilde{W}_{\tau',\tau}, \\ 1 & \text{otherwise.} \end{cases} \tag{16a}$$

$$\tag{16b}$$

At the aggregated level, the control α_t is now given by

$$\alpha_t = \frac{1}{W_t} \sum_{\tau=0}^t W_{t,\tau} \alpha_{t,\tau}. \tag{17}$$

Conceptually, the difference between the cumulative target strategy and the individual target strategy is what triggers the risk reduction. Contrary to the individual target strategy, in the cumulative target strategy new investments have to make up for insufficient past returns before a transfer to the matching portfolio can take place. On the other hand, in the cumulative target strategy good past returns may cause new contributions to be transferred immediately to the matching portfolio. With the individual target strategy, each contribution has to generate sufficient return on investment before such a transfer takes place.

3.3 Combination strategy

Both the cumulative and the individual target strategy either reduce risk by switching completely to the matching portfolio or don't reduce risk at all. Instead of completely switching or not switching at all, the combination strategy, which is the third strategy considered, combines the individual target strategy with dynamic programming to dynamically steer the wealth $W_{t,\tau}$ resulting from the contribution c_τ above its wealth target $\tilde{W}_{t,\tau}$. For this, we define the following wealth to target ratio,

$$Z_{t,\tau} := \frac{W_{t,\tau}}{\tilde{W}_{t,\tau}}, \tag{18}$$

and solve

$$V(z, t, \tau) = \sup_{\mathcal{A}_{t,\tau}} \mathbb{E} [\check{U}(Z_{T,\tau}) | Z_{t,\tau} = z], \tag{19}$$

where \check{U} is a utility function, $V(z, t, \tau)$ is the value function in the dynamic programming problem and the control $\mathcal{A}_{t,\tau}$ consists of the future investment decisions:

$$\mathcal{A}_{t,\tau} = \{\alpha_{t,\tau}, \dots, \alpha_{T,\tau}\}. \tag{20}$$

Using the dynamic programming principle, it follows that the optimal control, $\mathcal{A}_{t,\tau}^*$, satisfies

$$\mathcal{A}_{t,\tau}^* = \{\alpha_{t,\tau}^*, \mathcal{A}_{t+1,\tau}^*\}, \tag{21}$$

which allows us to solve for the optimal control problem for $\mathcal{A}_{t,\tau}^*$, backwards in time.

In this context, we choose a utility function that steers the ratio $Z_{T,\tau}$ in between the bounds z_{\min}^* and z_{\max}^* . This is in line with the investor’s goal of minimizing downside risk, and with our assumption that upward potential comes with downside risk. The utility function should be positive concave and takes here the following functional form:

$$\check{U}(z) = \frac{-(z - \beta)^2 - (z - z_{\min}^*)^2}{z}, \tag{22}$$

where

$$\beta = \sqrt{2(z_{\max}^*)^2 - (z_{\min}^*)^2},$$

see Fig. 1; note that this utility function differs from $U(\cdot)$ as given by (6). Utility function $\check{U}(\cdot)$ is clearly concave and continuous on the domain $\mathbb{R}_{>0}$. We set $z_{\min}^* = 1$ and $z_{\max}^* = 3$, as this choice fits well with the investor’s replacement ratio target and, as we will show in Sect. 4.1, is sufficient to demonstrate the strategy’s added value.

Now, we will show that the ratio $Z_{t,\tau}$, between the current wealth $W_{t,\tau}$ and its target $\tilde{W}_{t,\tau}$, evolves in time by making returns on investment in the nominator and updating the inflation expectation in the denominator. Since this time evolution is independent of τ , we can show that the optimal control $\alpha_{t,\tau}^*$ is independent of τ , i.e., once the optimal control is found, it can be applied to all contributions.

Lemma 1 *If the optimal control $\alpha_{t,\tau}^*$ of dynamic programming problem (19) exists and is unique, then it is independent of the contribution c_τ and the time τ at which the contribution is made.*




Fig. 1 Plot of (22) with $z_{\min}^* = 1$ and $z_{\max}^* = 3$

Proof The portfolio wealth $W_{t,\tau}$, accumulated by investing contribution c_τ , increases with the return on investment and, therefore, satisfies

$$W_{t,\tau} = [(1 + x_t)\alpha_{t-1,\tau} + (1 + m_t)(1 - \alpha_{t-1,\tau})]W_{t-1,\tau}. \tag{23}$$

From (15), (9) and (8), it follows that the wealth target, $\tilde{W}_{t,\tau}$, satisfies

$$\tilde{W}_{t,\tau} = \frac{\mathbb{E}_t[F_{t-1}]}{\mathbb{E}_{t-1}[F_{t-1}]}(1 + m_t)\tilde{W}_{t-1,\tau}. \tag{24}$$

Substitution of (23) and (24) in (19) yields that the optimal controls $\alpha_{t,\tau}^*$ solve

$$\sup_{\mathcal{A}_{t,\tau}} \mathbb{E} \left[\check{U} \left(z \prod_{\tau'=t+1}^T \frac{(1 + x_{\tau'+1})\alpha_{\tau',\tau} + (1 + m_{\tau'+1})(1 - \alpha_{\tau',\tau})}{\frac{\mathbb{E}_{\tau'+1}[F_{\tau'}]}{\mathbb{E}_{\tau'}[F_{\tau'}]}(1 + m_{\tau'+1})} \right) \middle| Z_{t,\tau} = z \right]. \tag{25}$$

As the argument of \check{U} merely depends on the control $\alpha_{\tau',\tau}$, future returns on equity $x_{\tau'+1}$, and future returns of the matching portfolio $m_{\tau'+1}$, all optimal controls $\alpha_{t,\tau}^*$ solve, as a function of the state z , the same sequence of optimization problems, backwards in time. Consequently, both the value function $V(z, t, \tau)$ and the optimal control $\alpha_{t,\tau}^*$ are independent of τ . A similar argument shows the independence of the size of the contribution c_τ . \square

Lemma 1 implies that, theoretically, the dynamic programming problem has to be solved only once, i.e., the optimal control (as a function of the state) is the same for all contributions. Therefore, the optimal control of the first contribution c_0 can be used for all other contributions.

For the practical implementation for the dynamic programming algorithm, readers may refer to Appendix B.

3.4 Expected and target replacement ratio

The variable r , used in the construction of the wealth targets by means of Eq. (9), can be interpreted in several ways. First, it serves as a discount rate, which is used to compute the present value of contributions that are made in the future. It can also be interpreted as an annual return requirement: each contribution is required to have an average annual return of r . And third, it can be interpreted as a future expected annual return.

The third interpretation allows us to measure approximately a strategy’s progress towards the expected replacement ratio $\mathbb{E}_t[RR_T]$ by the quantity R_t :

$$R_t := \frac{(T + 1)\mathbb{E}_t[W_T]}{\mathbb{E}_t[M_T] \mathbb{E}_t\left[\sum_{t=0}^T s_t \prod_{\tau=t+1}^T (1 + \pi_\tau)\right]}, \tag{26}$$

where $\mathbb{E}_t[W_T]$ is approximated using:

$$\mathbb{E}_t[W_T] \approx [1 + r + I_{T,t}]^{T-t} W_t + \sum_{k=t+1}^{T-1} [1 + r + I_{T,t}]^{T-k} \mathbb{E}_t[c_k]. \tag{27}$$

Note that in the definition of R_t , the expectation of the denominator and nominator are, contrary to how $\mathbb{E}_t[RR_T]$ is defined, taken separately for computational reasons. Altogether, the computation of R_t requires four different estimators. The discount rate r is used as an estimator for the future expected annual return. The estimator for the annualized future inflation $I_{T,t}$, has to be estimated by regression between the future and the past cumulative inflation, as shown in Eq. (29). Future salaries are based on the information from Table 3. Lastly, the estimator for the market value factor at the end of the investment horizon, $\mathbb{E}_t[M_T]$, is based on regression between M_t and M_T , with $\Phi = \{1, x\}$. See Appendix B.3 for details of the regression method used.

Alternatively, one could use these approximations to try to steer towards a replacement ratio target R_t^* directly. Using (26) and (27), the replacement ratio target can be converted into a current wealth target:

$$\frac{R_t^* \mathbb{E}_t[M_T] \mathbb{E}_t\left[\sum_{t=0}^T s_t \prod_{\tau=t+1}^T (1 + \pi_\tau)\right] - (T + 1) \sum_{k=t+1}^{T-1} [1 + r + I_{T,t}]^{T-k} \mathbb{E}_t[c_k]}{(T + 1)[1 + r + I_{T,t}]^{T-t}}.$$

We found, however, that steering towards this approximate target in a dynamic programming framework didn’t yield satisfactory results. Therefore, we postpone exploring such a direction deeper to future research.

4 Numerical evaluation

In this section, we apply the rule-based strategies described in Sect. 3 to the pension investor introduced in Sect. 2.2, using the governing stochastic model described in Sect. 2.3.

4.1 Rule-based strategies

To illustrate the dynamics of the rule-based strategies, Fig. 2 shows one of the 2000 sample paths for the investor’s portfolio dynamics. In particular, the top left figure shows the investor’s wealth W_t when following the cumulative target strategy (orange) and when the investor’s wealth exceeds the target \tilde{W}_t (yellow). Note that when this occurs, the investments are transferred to the matching portfolio (orange line, bottom right figure). The individual target strategy (in green) works similarly, but, as discussed, uses a target per contribution, so that, typically, only part of the wealth is transferred to the matching portfolio (green line, bottom right figure). The bottom left figure illustrates that, in this sample path, the rule-based strategies outperform the optimal static strategy in terms of expected replacement ratio—although only in the first 10 years of the investment the rule-based strategies take substantially more risk, i.e., have a substantially higher allocation to the return portfolio. Therefore, in this particular sample path, one could argue that the better performance comes from the rule-based strategies and not from increased exposure to risk.




Fig. 2 Sample paths with cumulative target strategy (orange), individual target strategy (green) and optimal static strategy (blue) of: wealth W_t (top left) with the wealth target \tilde{W}_t of the cumulative target strategy (yellow), and cumulative contribution (dark blue); return (top right) of the matching (dark blue) and return portfolio (yellow), expected replacement ratio RR_T (bottom left) together with the 70% target (yellow); and (bottom right) allocation $1 - \alpha_t$ to the matching portfolio

The combination strategy is best illustrated by means of the resulting investment decisions, i.e., the optimal control $\alpha_{t,\tau}$ as defined by Eq. (25). Figure 3 illustrates the optimal allocation to the matching portfolio, $1 - \alpha_{t,0}$, for the first contribution c_0 as a function of time t and the wealth to target ratio, as defined by Eq. (18). In this example, allocations are restricted to multiples of 20%. Note that, contrary to the rule-based strategies, in the combination strategy, risk can be increased and investments can be transferred from the matching to the return portfolio. All together, this makes the combination strategy more refined than the rule-based strategies, which follow a “risk on” or “risk off” approach, in terms of their allocations.

Figure 4 compares the distribution of the terminal replacement ratio RR_T for the following best performing strategies in terms of the expected shortfall below the investor’s 70% replacement ratio target: two rule-based strategies, a combination strategy and a static strategy. The figure illustrates that, as intended, the dynamic strategies reduce downward risk at the expense of upward potential, i.e., the dynamic strategies are centered more around the target replacement ratio of 70%.

In the Dutch pension system, maximum allowed tax free contributions are set by law, and few people deviate from this. Nevertheless, as a sensitivity analysis, we investigated that with 40% higher contributions all strategies achieve the 70% replacement ratio target with about 90% certainty, see Fig. 5. The figure illustrates that the designed strategies perform especially well when one tries to steer towards a




Fig. 3 Optimal allocation $1 - \alpha_{t,0}$ to the matching portfolio, as a function of wealth to target ratio $w_{t,0}/\bar{w}_{t,0}$ (y-axis), as defined by Eq. (18), and time t (x-axis) for the first contribution of an investor following the combination strategy, discussed in Sect. 3.3, and using the stochastic model of Sect. 2.3. Allocations are restricted to multiples of 20%




Fig. 4 Distribution of the replacement ratio RR_T for, respectively, the cumulative target strategy with $r = 3.06\%$ (orange), individual target strategy $r = 2.99\%$ (green), combination strategy with $r = 1\%$ (yellow) and a static strategy with 46.02% constant allocation to the return portfolio

target that is challenging to achieve. Also, intuitively, for such a target, reducing risk after several years with better than anticipated return on investment makes sense, which is inherent in the design of the strategies. That being said, still the dynamic strategies reduce downward risk at the expense of upward potential.

A comparison of all strategies is best made by comparing the strategy successes, i.e., whether a strategy achieves the intended 70% replacement ratio target, versus its downside risk, and parametrize the strategies by the parameters that control




Fig. 5 Distribution of the replacement ratio RR_T with 40% higher contributions for, respectively, the cumulative target strategy with $r = 3.06\%$ (orange), individual target strategy $r = 2.99\%$ (green), combination strategy with $r = 1\%$ (yellow) and a static strategy with 46.02% constant allocation to the return portfolio.

the strategy’s risk appetite, see Fig. 6. From this figure, we conclude that all the dynamic strategies clearly outperform the traditional static strategies. Together with the intuitive rationale to reduce risk after several good years, we believe this sufficiently demonstrates the added value of these dynamic strategies. We do, however, find these simulations insufficient to rank the dynamic strategies based on their effectiveness. It is well-known that the relative performance of dynamic strategies can be sensitive to the characteristics of the underlying stochastic model. As such, the characteristics are not completely objective, and we believe that use of the strategies in practice is an appropriate way to test the strategies further (which lies beyond the scope of this research).

4.2 Discussion

One of the intended advantages of a static life cycle strategy is the reduced risk close to the retirement, meaning that one can provide the investor with an accurate estimate of her retirement income in the years before retirement. Table 2 provides a comparison of the dynamic strategies and traditional life cycle strategies. In particular, the table lists the standard deviations of the difference between the expected replacement ratio 5 years before retirement and the replacement ratio at retirement. We conclude that when following the rule-based strategies the investor can be provided with a similarly accurate estimate of the replacement ratio before retirement

Although the rule-based strategies outperform other strategies in our examples, we wish to point out that there are also disadvantages to the all-or-nothing approach, e.g, the portfolio remains 100% invested in the more risky return portfolio when




Fig. 6 Expected shortfall below the investor’s 70% replacement ratio target, see Eq. (6), versus 10% CVaR of the terminal replacement ratio RR_T , as defined by (7), for: the cumulative target strategy (orange), the individual target strategy (green), the combination strategy (yellow), all parametrized by the real rate of return r (yellow), and, also, several annually rebalanced allocations (blue), and several default life cycles reducing risk with the investor’s age

Table 2 Statistics for different investment strategies. Values in this table were calculated using $r = 3.06\%$ for the cumulative strategy, $r = 2.99\%$ for the individual strategy, and $r = 1\%$ for the combination strategy

	Static mix			Static life cycle			Dynamic strategies		
	0%	100%	46.02%	Def.	Neut.	Off.	Cum.	Indiv.	Comb.
Averages (R)									
Mean	0.56	1.13	0.77	0.66	0.71	0.79	0.70	0.70	0.75
10% CVaR	0.44	0.28	0.41	0.42	0.42	0.40	0.36	0.41	0.40
5% CVaR	0.42	0.24	0.37	0.39	0.39	0.36	0.28	0.34	0.35
Percentiles (R)									
Median	0.56	0.83	0.72	0.64	0.67	0.72	0.72	0.73	0.75
10% VaR	0.46	0.36	0.47	0.47	0.47	0.46	0.51	0.52	0.47
5% VaR	0.44	0.29	0.41	0.43	0.44	0.41	0.37	0.44	0.41
Goal (70% R)									
Shortage	0.139	0.093	0.070	0.088	0.078	0.073	0.044	0.045	0.063
Goal reached	6%	60%	53%	35%	44%	53%	65%	65%	57%
Estim. error (R)									
Mean	0.088	0.398	0.139	0.095	0.092	0.115	0.113	0.091	0.144
Std dev.	0.06	0.57	0.12	0.06	0.06	0.10	0.09	0.07	0.12

targets are not reached. Such truly worst case scenarios appear to have a minor influence, but are, e.g., illustrated in the far left lower tail in Fig. 4. The individual target strategy presented in Sect. 3.2 suffers less from the all-or-nothing disadvantages, as it defines a target per contribution. As a result, inferior past performance does not influence the required performance of current and future contributions.

Compared with the rule-based strategies, the combination strategy does not exploit the fact that the matching portfolio can grow an investment securely to its intended target (indexed by expected inflation) until retirement. As the rule-based strategies explicitly made use of this, the combination strategy could be further improved.

One advantage of the combination strategy in practical use is that the corresponding asset allocation is much more smooth than for the rule-based strategies. The necessity of large turnovers is difficult to explain and investors might be uncomfortable to follow such a drastic strategy to the end.

5 Conclusion

In this paper, we discussed several dynamic strategies, suitable for pension investors that aim to replace a proportion of their salary with a retirement income. The strategies reduce risk after several good years and steer the investor to her target. By having the allocation depend on return on investment, the approaches exploit a freedom which is typically not used by traditional static approaches. We

have shown that the dynamic approaches may outperform some traditional static approaches and prevent unnecessary risk taking.

Two simple and intuitive rule-based strategies were introduced that secure investments in a cash flow matching portfolio once they yielded sufficient return. Although both rule-based strategies can straightforwardly be implemented in practice, we recommend to also investigate alternatives where the investor, e.g., switches between an aggressive traditional life cycle and a matching portfolio to rule out very aggressive portfolios close to retirement.

The rule-based strategies were further refined into a combination strategy based on dynamic programming. In the current setup, the combination strategy may not be superior and we even found that the rule-based strategies outperform the combination strategy in a numerical example. We certainly believe that dynamic strategies based on dynamic programming can be further improved, as also this research clearly demonstrates their added value to pension investors.

A most suitable dynamic strategy is hard to determine objectively as its performance depends on the governing stochastic model. Also, such a dynamic strategy should fit well with practical requirements, such as whether an investor will follow through on the strategy or will feel the need to combine such a strategy with her own judgement, and whether such strategies comply with regulations. This research, however, demonstrates the added value of dynamic strategies to pension investors. In summary, such strategies exploit freedom that is not used by traditional approaches, can steer a pension investor to her target and prevent unnecessary risk taking.

Appendix A: Career path and contribution rate

In our model, salary increases by age with a rate to reflect one's career path. Maximum allowed (tax free) contributions under Dutch pension law also increase by age as closer to retirement there is less time to invest and, consequently, buying 1 Euro of pension becomes more expensive. Accordingly, Table 3 shows at $t = 0$, corresponding to ultimo 2016, the salaries and contributions by age. All is based on statistics reported for the Netherlands in 2016. The salary for a 25 year old is reported by Eurostat as the average Dutch salary under 30 years old in 2014 and is indexed with 2 years of wage inflation as reported by the Dutch Central Bureau of Statistics. The salary increases corresponding to a career path used in actuarial calculations of the Dutch government (available under document number blg-230018). Contribution percentages are as prescribed under Dutch law in 2016 and are applied to the pensionable salary, i.e., the difference between the salary and a franchise of 13123 Euro. For $t > 0$, nominal amounts are indexed with wage inflation and rates are kept fixed. Altogether, for a person 25 years of age in 2016 following the career path in Table 3 and indexing with inflation annually, Table 4

Table 3 Salary and contribution in 2016 by age

Age	Salary		Contribution	
	Rate (%)	Euro	Rate (%)	Euro
25	3	29,403	7.8	1270
26	3	30,285	7.8	1339
27	3	31,193	7.8	1409
28	3	32,129	7.8	1482
29	3	33,093	7.8	1558
30	3	34,086	9.0	1887
31	3	35,108	9.0	1979
32	3	36,162	9.0	2073
33	3	37,247	9.0	2171
34	3	38,364	9.0	2272
35	2	39,131	10.5	2731
36	2	39,914	10.5	2813
37	2	40,712	10.5	2897
38	2	41,526	10.5	2982
39	2	42,357	10.5	3070
40	2	43,204	12.2	3670
41	2	44,068	12.2	3775
42	2	44,950	12.2	3883
43	2	45,849	12.2	3993
44	2	46,765	12.2	4104
45	1	47,233	14.2	4844
46	1	47,705	14.2	4911
47	1	48,183	14.2	4978
48	1	48,664	14.2	5047
49	1	49,151	14.2	5116
50	1	49,642	16.5	6026
51	1	50,139	16.5	6108
52	1	50,640	16.5	6190
53	1	51,147	16.5	6274
54	1	51,658	16.5	6358
55	0	51,658	19.4	7476
56	0	51,658	19.4	7476
57	0	51,658	19.4	7476
58	0	51,658	19.4	7476
59	0	51,658	19.4	7476
60	0	51,658	23.0	8863
61	0	51,658	23.0	8863
62	0	51,658	23.0	8863
63	0	51,658	23.0	8863
64	0	51,658	23.0	8863
65	0	51,658	26.0	10,019
66	0	51,658	26.0	10,019

Table 4 Salary and contribution per year indexed by inflation for a person 25 year old person in 2016 (in Euro)

Year	Expected inflation (%)	Expected salary	Contribution	
			Expected	Stddev.
2016		29,403	1270	
2017	0.7	30,816	1362	9
2018	0.8	32,168	1454	18
2019	0.9	33,610	1551	31
2020	1.0	35,147	1909	53
2021	1.1	36,784	2036	73
2022	1.2	38,527	2171	97
2023	1.2	40,380	2315	124
2024	1.3	42,351	2469	154
2025	1.3	44,446	3070	219
2026	1.4	46,218	3225	259
2027	1.4	48,087	3389	303
2028	1.5	50,058	3562	351
2029	1.5	52,133	3744	403
2030	1.6	54,319	4574	534
2031	1.6	56,621	4810	606
2032	1.6	59,043	5058	683
2033	1.7	61,585	5320	763
2034	1.7	64,261	5596	851
2035	1.7	67,078	6852	1104
2036	1.7	69,347	7111	1206
2037	1.8	71,716	7382	1316
2038	1.8	74,185	7665	1433
2039	1.8	76,754	7960	1555
2040	1.8	79,428	9607	1956
2041	1.8	82,213	9979	2113
2042	1.8	85,112	10,368	2280
2043	1.9	88,142	10,775	2464
2044	1.9	91,297	11,199	2658
2045	1.9	94,580	13,687	3366
2046	1.9	97,022	14,041	3572
2047	1.9	99,540	14,405	3784
2048	1.9	102,134	14,780	4003
2049	1.9	104,806	15,167	4230
2050	1.9	107,554	18,453	5288
2051	1.9	110,402	18,942	5588
2052	1.9	113,314	19,441	5881
2053	1.9	116,312	19,956	6186
2054	2.0	119,400	20,486	6503
2055	2.0	122,588	23,776	7732
2056	2.0	125,859	24,411	8116
2057	2.0	129,229	25,064	8515

shows the annual salary and contribution to the pension plan. Note that future contributions are uncertain as inflation is stochastic in our model.

Appendix B: Dynamic programming implementation

In this section, we give a brief description of various computational insights from our implementation. The algorithm for the combination strategy includes the dynamic programming design, the selection of the state variable and the use local regression. The local regression technique in this section is also used by the rule-based strategies.

B.1 Dynamic programming algorithm

Asset allocations for a dynamic programming strategy follow from Algorithm 1. The algorithm runs backward in time, similar to the algorithms in [8] and [19]. The problem's solution space is discretized to efficiently solve the dynamic programming problem:

$$\alpha_t \in \{a_j \in [0, 1] | j = 1, \dots, k\}.$$

This means the investor can choose between at most k asset allocations at each time t .

Algorithm 1 solves the optimal control problem backward in time by calculating the expected utility, $\mathbb{E}_t[\check{U}(Z_T)]$, with the state variable Z_t (see Sect. B.2) by the local regression technique (see Sect. B.3). The use of local regression is similar to the use of bundling in [8]: neighborhood points are used in the local regressions for each step of the algorithm.

Solving the sub-problem at time t changes the future states Z_{t+1}, \dots, Z_T . These states have previously been used in the local regressions to find the optimal solution for the sub-problems at times $t + 1, \dots, T$. Thus, the optimal solutions for sub-problems at $t + 1, \dots, T$ may be different after solving the sub-problem at time t . This is why Algorithm 1 follows a snake-like pattern through time: after the sub-problem at time t is solved, future sub-problems are first updated in a forward fashion in time. Sub-problems are subsequently updated backwards in time at the next step until the sub-problem at time $t - 1$ is solved for the first time.

Once all sub-problems have been solved, the solution can be further improved by repeating the procedure. Algorithm 1 restarts at the beginning of the snake-like pattern through time. Each iteration of Algorithm 1 follows the snake-like pattern from T to t_0 once.

Algorithm 1: Dynamic programming investment strategy

Data: Scenario Set, asset allocations a_1, \dots, a_k
Result: Optimal admissible asset control \mathcal{A}^*
 Initialization: create initial solution by choosing $\alpha_t = a_1$ for all $t = 0, \dots, T$ for all scenarios.

```

1 for  $m = 1, \dots$ , number of iterations do
2   for  $j = 1, \dots, k$  do
3     Set  $\alpha_{T-1} = a_j$  for all scenarios
4     Determine the expected utility function  $f_{j,T-1}$ ,
        $f_{j,T-1}(z) = \mathbb{E} [\check{U}(Z_T) | Z_{T-1} = z, \alpha_{T-1} = a_j]$ ,
       by using local regression for  $Z_{T-1}$  and  $\check{U}(Z_T)$  (Least Squares Monte Carlo method)
5     Determine optimal asset allocation decision at time  $T - 1$  for all scenarios:
        $m = \arg \max_{j \in \{1, \dots, m\}} f_{j,T-1}(Z_{T-1})$ 
        $\alpha_{T-1}^* = a_m$ 
6   for  $t = T - 2, T - 3, \dots, 0$  do
7     Update optimal allocation decisions (backward update):
8     for  $\tau = T, \dots, t + 1$  do
9       Update expected utility function  $f_{j,\tau}$ , for  $j = 1, \dots, k$ 
10      Update optimal asset allocation decision  $\alpha_\tau^*$ , for all scenarios
11     Determine optimal asset allocation decision at time  $t$  for the first time
12     for  $j = 1, \dots, k$  do
13       Set  $\alpha_t = a_j$  for all scenarios
14       Determine the expected utility function  $f_{j,t}$ ,
          $f_{j,t}(z) = \mathbb{E} [\check{U}(Z_T) | Z_t = z, \alpha_t = a_j, \mathcal{A}_{t+1} = \mathcal{A}_{t+1}^*]$ ,
         by using local regression for  $Z_t$  and  $\check{U}(Z_T)$  (Least Squares Monte Carlo method)
15       Determine optimal asset allocation decision at time  $t$ , for all scenarios:
          $m = \arg \max_{j \in \{1, \dots, k\}} f_{j,t}(Z_t)$ 
          $\alpha_t^* = a_m$ 
16     Update optimal allocation decisions (forward update):
17     for  $\tau = t + 1, \dots, T - 1$  do
18       Update expected utility function  $f_{j,\tau}$ , for  $j = 1, \dots, k$ 
19       Update optimal asset allocation decision  $\alpha_\tau^*$ , for all scenarios
    
```

B.2 State variable

Individual wealth targets are not constant over the time horizon. They are defined using the expected inflation and the market value factor at time t (see also Sect. 3.4). The individual wealth targets are also not constant between the different scenarios, as they are dependent on the contributions. The dynamic programming approach requires, however, a fixed wealth target in order to evaluate the expected utility. This issue is resolved by using the ratio between the portfolio wealth and the wealth target Z_t , defined in Eq. (18), as the state variable of the dynamic programming algorithm. The target state is now constant in time: $Z_t^* = 1$.

An advantage of choosing this state variable is that, theoretically, the dynamic programming solution only has to be computed once. The investment decisions for the first contribution of the investor can be used for all other contributions. Dynamic programming results for the first contribution span the time horizon $0, \dots, T$. At each rebalancing time, the optimal investment decision, depending on state Z_t , has already been computed. Because a ratio is used, these investment decisions can also be used for the contributions in later years, see Lemma 1.

This is, however, not the case in practice. Algorithm 1 does not fully converge due to the limited sample size and due to sets of observations changing multiple times per iteration. Running Algorithm 1 separately for each contribution will give better decision rules. Separate runs of the algorithm will provide a better approximation of the optimal solution on average.

B.3 Least squares Monte Carlo method

At each optimization step in Algorithm 1, a least squares Monte Carlo method is used to avoid nested simulations (and the, related, exploding computation times). The least squares Monte Carlo method was introduced by Longstaff and Schwartz [15] as a simple method for pricing American options by simulation. The conditional expectation of the pay-off under the assumption of not exercising the option is estimated by using cross-sectional information already available in the simulation. Realized pay-offs from continuation (or, in the pension investment setting, from final utility $\check{U}(Z_T)$) are regressed on functions of the state variables. The fitted value provides an estimate of the conditional expectation.

The regression employed is based on a so-called *regress-now* strategy, and specifically on a local regression version. Regress-now estimates the expectation $\mathbb{E}[Y_{t+1}|X_t]$, $X_t \in \mathcal{F}_t$ by using a set of basis functions, Φ , with index set \mathcal{J} :

$$Y_{t+1} \approx \sum_{j \in \mathcal{J}} c_j \varphi_j(X_t),$$

with c_j coefficients found by using least squares regression and $\varphi_j \in \Phi$. Substitution gives

$$\begin{aligned} \mathbb{E}[Y_{t+1}|X_t] &\approx \mathbb{E}\left[\sum_{j \in \mathcal{J}} c_j \varphi_j(X_t) \middle| X_t\right] \\ &= \sum_{j \in \mathcal{J}} c_j \varphi_j(X_t). \end{aligned}$$

Local regression, introduced by Cleveland [6], estimates a linear or quadratic polynomial fit at x by using a weighted least squares regression. Weights for an observation (x_i, y_i) are dependent on the distance between x_i and x [7]. The smoothness of the fit is dependent on the percentage of observations that are taken into account when evaluating at x .

Let n be the number of observations and let $0 < d \leq 1$ be a neighborhood parameter, i.e., the share of observations used for the weighted least squares regression at the evaluation point. Furthermore, let $k = d \cdot n$, rounded up to an integer value, $\Delta_i(x)$ be the Euclidean distance of x to x_i , and $\Delta_{(i)}(x)$ be the values of these distances, ordered from smallest to largest.

The weight ζ_i for an observation (x_i, y_i) is then equal to

$$\zeta_i(x) = T\left(\frac{\Delta_i(x)}{\Delta_{(k)}(x)}\right),$$

with

$$T(u) = \begin{cases} (1 - u^3)^3, & \text{for } 0 \leq u < 1, \\ 0, & \text{for } u \geq 1, \end{cases}$$

also known as the tri-cube weight function.

Not only is the regress-now strategy used to approximate the expected utility, this strategy is also used to estimate the expected annual inflation $\mathbb{E}_t \pi_\tau$, with $\Phi = \{1, x\}$. The future cumulative inflation, Y_{t+1} , is regressed on the past cumulative inflation, X_t , to estimate the future annual inflation at time t :

$$x_j = \prod_{\tau=0}^t (1 + \pi_{\tau,j}),$$

$$y_j = \prod_{\tau=t+1}^T (1 + \pi_{\tau,j}),$$

for each j in the scenario set. Using a linear regression function of the form $\xi_1 x + \xi_2$ leads to:

$$\mathbb{E}_t \prod_{\tau=t+1}^T (1 + \pi_\tau) \approx \xi_1 x_k + \xi_2, \quad (28)$$

which can approximately be annualized to

$$I_{T,t} \approx \sqrt[T]{\xi_1 x_k + \xi_2} - 1, \quad (29)$$

so that we obtain an approximate future annual inflation rate.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Arnott R, Sherrerd K, Wu L (2013) The glidepath illusion... and potential solutions. *J Retire* 1(2):13–28. <https://doi.org/10.3905/jor.2013.1.2.013>
2. Basu A, Byrne A, Drew M (2011) Dynamic lifecycle strategies for target dateretirement funds. *J Portf Manag* 37(2):83–96. <https://doi.org/10.3905/jpm.2011.37.2.083>
3. Bellman R (1957) *Dynamic programming*. Dover Publications. <https://www.bibsonomy.org/bibtex/29cdd821222218ded252c8ba5cd712666/pcbouman>
4. Blanchett D, Kowara M, Chen P (2012) Optimal withdrawal strategy for retirement income portfolios. Morningstar Research Paper. https://www.morningstar.com/content/dam/marketing/shared/research/methodology/677951-Optimal_Withdrawal_Strategy_for_Retirement_Income_Portfolios.pdf
5. Bogle J (1994) *Bogle on mutual funds: new perspectives for the intelligent investor*. A Dell trade paperback. Random House Publishing Group. <https://books.google.nl/books?id=6A9sYcGT92sC>
6. Cleveland WS (1979) Robust locally weighted regression and smoothing scatterplots. *J Am Stat Assoc* 74(368):829–836. <https://doi.org/10.1080/01621459.1979.10481038>
7. Cleveland WS, Grosse E (1991) Computational methods for local regression. *Stat Comput* 1(1):47–62. <https://doi.org/10.1007/BF01890836>
8. Cong F, Oosterlee C (2016) Multi-period mean–variance portfolio optimization based on Monte-Carlo simulation. *J Econ Dyn Control* 64:23–38. <https://doi.org/10.1016/j.jedc.2016.01.001>. <http://www.sciencedirect.com/science/article/pii/S0165188916000026>
9. DNB: Scenarioset haalbaarheidstoets 2016q4. <https://www.toezicht.dnb.nl/2/50-233246.jsp>
10. Draper N (2014) A financial market model for the Netherlands. CPB Background Document. <https://www.cpb.nl/en/publication/a-financial-market-model-for-the-netherlands>
11. Forsyth PA, Vetzal KR (2019) Optimal asset allocation for retirement saving: deterministic vs. time consistent adaptive strategies. *Appl Math Finance* 26(1), 1–37. <https://doi.org/10.1080/1350486X.2019.1584534>
12. Graf S (2017) Life-cycle funds: much ado about nothing? *Eur J Finance* 23(11):974–998. <https://doi.org/10.1080/1351847X.2016.1151805>
13. Kojien RSJ, Nijman TE, Werker BJM (2009) When can life cycle investors benefit from time-varying bond risk premia? *Rev Financ Stud* 23(2):741–780. <https://doi.org/10.1093/rfs/hhp058>
14. Li D, Ng WL (2000) Optimal dynamic portfolio selection: multiperiod mean-variance formulation. *Math Finance* 10(3):387–406. <https://doi.org/10.1111/1467-9965.00100>
15. Longstaff FA, Schwartz ES (2001) Valuing American options by simulation: a simple least-squares approach. *Rev Financ Stud* 14:113–147
16. Martellini L, Milhau V, Tarelli A (2014) Hedging inflation-linked liabilities without inflation-linked instruments through long/short investments in nominal bonds. *J Fixed Income* 24(3):5–29. <https://doi.org/10.3905/jfi.2014.24.3.005>
17. Merton R (1969) Lifetime portfolio selection under uncertainty: the continuous-time case. *Rev Econ Stat* 51(3):247–57
18. Muns S (2015) A financial market model for the Netherlands: a methodological refinement. Netspar working paper. https://www.netspar.nl/assets/uploads/028_Muns-1.pdf
19. van Binsbergen JH, Brandt MW (2007) Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function? *Comput Econ* 29(3):355–367. <https://doi.org/10.1007/s10614-006-9073-z>
20. Zhang R, Langren N, Tian Y, Zhu Z, Klebaner F, Hamza K (2017) Sharp target range strategies for multiperiod portfolio choice by decensored least squares Monte Carlo

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.