Adaptive Strategies for Pension Fund Management

Raphael Chinchilla*[1], Thomas D. Rueter[1], Timothy R. McDade[1], Peter R. Fisher[1], Mykel J. Kochenderfer[1,2], Emmanuel Candès[1,2], Trevor Hastie[1,2], and Stephen Boyd[1,2]

[1]BlackRock, 50 Hudson Yards, New York, NY 10001, USA
[2]Stanford University, Stanford, CA 94305, USA

Abstract

This paper proposes a simulation-based framework for assessing and improving the performance of a pension fund management scheme. This framework is modular and allows the definition of customized performance metrics that are used to assess and iteratively improve asset and liability management policies. We illustrate our framework with a simple implementation that showcases the power of including adaptable features. We show that it is possible to dissipate longevity and volatility risks by permitting adaptability in asset allocation and payout levels. The numerical results show that by including a small amount of flexibility, there can be a substantial reduction in the cost to run the pension plan as well as a substantial decrease in the probability of defaulting.

Keywords: Retirement; Longevity; Pension; Decision-making under uncertainty; Simulation based design; Direct policy learning

1 Introduction

Retirement savers and investors confront a variety of risks, including market volatility, inflation, fund management cost, and longevity. The financial service industry uses annuities to manage longevity risk, but they are only available in select markets and are generally expensive. Individuals whose retirement is funded by Defined Contribution (DC) schemes rather than Defined Benefit (DB) must bear these four risks themselves. For most retirement savers, the risk of outliving accrued savings is not negligible.

The framework we present in this paper is applicable to a variety of retirement financing arrangements. Although it was motivated by considering a DB plan in runoff, the framework today presents a more compelling case for the DC space for two reasons. First, DC has become a far more common arrangement than DB. Second, in DC retirement plans, no beneficiary has a guaranteed floor on payouts for the duration of their lifetime; any floor on lifetime payouts, like the one our framework offers, is a great value indeed.

When DB pensions were more common, beneficiaries were insensitive to capital market returns, inflation, fund management fees, or longevity risk. The DB pension plan and its corporate sponsor absorbed the risk, guaranteeing the beneficiary a lifetime cash flow regardless of the behavior of stocks, bonds, currencies, interest rates, inflation, fund management fees, or longevity patterns.

Over the last several decades, pension beneficiaries have had to shoulder increasing risk. Pure DC schemes leaves risk to the individual investor. The financial services industry has developed ways to manage some of these risks. Different portfolio constructions can help insulate investors against capital market volatility. Inflation-linked bonds, or exposure to equities or real assets can help inoculate investors against inflation. Large asset managers offer funds for very low fees.

*Corresponding author: raphael.chinchilla@blackrock.com
This paper represents the views of the authors and not of BlackRock, Inc.

Longevity risk remains mostly unmanaged. Pension funds that confront mounting financial uncertainty, including that stemming from longevity risk, often consider undertaking a Pension Risk Transfer (PRT). In a PRT, the pension plan sells their assets and liabilities (payments to beneficiaries) to an insurance company, who then becomes responsible for paying their beneficiaries. The pension plan must pay a steep premium for the privilege of offloading this risk, at the price of a substantial proportion of their liabilities.

The insurance companies then pay the pensioners with an annuity. Having collected the assets plus a premium, insurance companies also secure themselves by trying to build a bulletproof box around their balance sheets. Their regulators require the insurers to hold an extremely conservative bond matching portfolio, hedge interest rates and exchange rates (FX) until the last beneficiary dies, and hold large amounts of capital (upon which their shareholders expect appropriate returns). The insurance company can choose to offload some longevity risk to a reinsurer. The rigidity of this structure is designed to protect the promises to the beneficiary but ends up being costly.

Longevity risk is salient to individual investors who are increasingly funding their own retirement, as well as to those institutional investors who still hold longevity risk on their balance sheet (*e.g.,*, pension funds or insurance companies that write annuities). Instead of relying on either a corporate sponsor or government balance sheet, or on transferring the uncertainties of future life expectancy and asset volatility to a fortress balance sheet of an insurer or reinsurer, schemes could be designed to absorb both longevity and volatility risks by being adaptable in both the assets held and in the level of payouts over time. In this way, adaptable schemes could provide higher expected returns while still largely maintaining predictability.

The key contribution of this paper is a theoretical framework to assess the performance of an adaptable scheme through simulation. We start by specifying a model that describes the yearly evolution of the pension's assets given the invested portfolio returns, the net liabilities, and external funding. These three inputs depend on choices of portfolio, liabilities' management, and calls for external funding. Because the future is uncertain, these choices should not be fixed, but rather they should respond to changes in the environment and projections. Instead of fixed actions, our simulator requires policies (or decision strategies) that determine which actions to take each year given the available information. The policies are then evaluated by simulating them over many randomly generated scenarios and collecting metrics about their performance. These metrics are quantities relevant to the stakeholders, such as default probability and investment return. If the performance metrics are not satisfactory, then the policies should be updated and re-simulated, iteratively until an acceptable policy is found. This framework is explained in section 3.

Section 4 outlines a simple implementation of the abstract framework. The portfolios are constructed using a long only Markowitz approach (Boyd et al., 2024) in which we include holding limits on certain asset classes. The liabilities are obtained from stochastic simulations of individual mortality, contributions paid into the fund, and pensions paid out. To manage these liabilities, the pensions paid out can be increased or decreased with respect to a baseline value. For the policy, we propose one that depends only on the value of the current assets under management divided by the total future (predicted) liabilities, commonly known as the asset to liability ratio. We discretize the range of this ratio into bins, and we assign which actions to take from each bin. Finally, we propose an automatic method optimize the policy with respect to the realized metrics obtained from simulation.

Section 5 summarizes numerical results. Because our goal is to showcase the power of including adaptable features, we simulate four policies, each with an increasing degree of flexibility. The numerical results show that by including a small amount of flexibility, there is a substantial reduction in the cost to run the pension plan as well as a substantial decrease in the probability of defaulting.

2 Related work

2.1 Life cycle investing

Economics literature has long studied an individual's ability to smooth consumption throughout their lifetime. Early work by Friedman (1957) on "permanent income" and Modigliani and Brumberg (1954) on the "life cycle" argued that individuals prefer stable consumption over their lifetime. This body of work has evolved into a robust analysis of multi-period optimization problems that individual economic actors can undertake, and the financial decisions that those processes imply (Cocco et al., 2005). Complex models in this field of study consider the way an individual balances current and future utility of consumption with mortality probabilities, a desire to bequeath assets to heirs, and risk aversion.

2.2 Annuities

Humans are not rational economic actors and do not perfectly smooth their consumption over the long run (Tirole, 2017). Many families do not have stable income or consumption characteristics over time (Morduch and Schneider, 2017). One way to help ensure consumption remains stable despite shocks is purchasing insurance, particularly an income-oriented insurance product such as an annuity.

Annuities pool mortality risks among beneficiaries, so those that live shorter than expected end up receiving an income stream equivalent to less than their initial contribution, leaving the remainder to subsidize those who outlive expectations. Annuities come in many varieties, such as term, life, deferred, real, nominal, variable, and bequeathable. All guarantee a stable cash flow for the duration of the covered time period. This guarantee precipitates considerable regulatory scrutiny to annuity providers and constrains the set of permissible assets an insurer can hold against their liabilities (promised current and future cash flows to beneficiaries). As a result of the stability that annuities promise, their cost is often viewed as high. Scholars have studied annuity pricing over recent decades (Bauer et al., 2010) and the social factors influencing demand for them (Boyer et al., 2020).

2.3 Mortality pooling

Because of the variety in household economic circumstances, there is an extensive literature considering the best way for households to balance portfolio investment with annuity purchases (Pang and Warshawsky, 2010). One important subset of this literature draws inspiration from history to analyze asset management solutions to lifetime income using mortality pooling, some of which are called tontines (Goldsticker, 2007; Forman and Sabin, 2014; Milevsky, 2015; Fullmer et al., 2025). The basic idea of these solutions is that individuals contribute to a fund which is invested over time, and pay the members an income during the retirement phase. When a member of the fund dies, some (or all) of the funds they contributed remain in the pool, increasing (or lengthening) the payout of the remaining members. The mortality pooling concept has been proposed for the modern-day financial services industry (Milevsky, 2022). One of the key conceptual complexities of this solution is the heterogeneity within a pool with regard to mortality likelihood and contribution size. In recent years, the industry has evolved different mechanisms to distribute mortality credits in an actuarially fair manner, even across a pool of heterogeneous individuals (Fullmer and Sabin, 2019).

2.4 Liability driven investment

Liability driven investment (LDI) seeks to address the asset-liability management problem by optimizing quantities involving both the current assets and some economic value of the liabilities of a plan. There has been much work done in this area (see, *e.g.*,, (Ezra, 1991; Ang et al., 2012)), but common approaches focus on surplus management (Sharpe and Tint, 1990). The goal of surplus management is to select an asset allocation by solving

$$\arg\max_w E[z] - \frac{\lambda}{2}\mathrm{var}(z)$$
$$z = \frac{A_1(w) - L_1}{A_0}$$

where z is the surplus, A_0 is the current value of the assets, $A_1(w)$ is the value of the assets in the next timestep, which depends on the asset allocation w selected in the current timestep, L_1 is the economic value of the liabilities in the next timestep, $E[\cdot]$ is the expected value, and λ is a risk-aversion parameter. The liabilities enter through their expected value and variance, as well as the covariance of the asset returns and the change in value of the liabilities $\mathrm{cov}(L_1, A_1)$. Since the asset allocator can only impact the distribution of the asset returns, this problem is equivalent to solving

$$\arg\max_{w} E[r_A] - \frac{\lambda}{2}\mathrm{var}[r_A] + \lambda \mathrm{cov}(r_A, r_L),$$

where $r_A = (A_1(w)/A_0) - 1$ is the return of the assets and $r_L = (L_1/L_0) - 1$ is the fractional change in the value of the liabilities.

This problem reduces to traditional mean-variance optimization when the asset returns and changes in the value of the liabilities are uncorrelated. If the covariance term is nonzero, then the asset manager is encouraged to select asset allocations that hedge changes in the value of liabilities, often in the form of fixed income or derivatives.

The LDI approach to asset management outlined above (and various extensions) are primarily sensitive to the covariance of the liabilities and asset returns, rather than to the full distribution of future liabilities. Other approaches attempt to immunize the plan to changes in interest rates by matching moments of the liability distribution and portfolio cash flows, often to match durations (Zheng et al., 2002).

2.5 Direct policy optimization

The core idea of using simulations to assess a policy performance is a standard procedure widely used in the field of control engineering (Ogata, 2020). Simulations can also be used to choose the best policy given the desired metrics. When a policy is defined by a set of parameters, which is usually the case, one can adjust these parameters while taking into consideration such metrics. There is an extremely wide literature on policy optimization (Kochenderfer, 2015; Kochenderfer and Wray, 2022). Due to the small complexity and low time constraints of the problem, we use a direct policy optimization (Maher et al., 2022).

3 Pension fund model

3.1 Dynamics of a pension fund

Let v_t be the total value of the assets of a pension plan at the beginning of year t. Our goal is to create a model of the value of the assets of the pension plan at the beginning of the next year. This model depends on four components.

The first component is the return of the investments. We assume a universe of n investable assets, denoting their returns during year t by $\alpha_t \in \mathbb{R}^n$. Given a choice of portfolio p_t, the portfolio return is given by $r(\alpha_t, p_t)$. The portfolio return may depend non-linearly on the assets' returns α_t.

The next component is the liabilities. Let ι_t be attributes (such as age and income) of an individual enrolled in the pension plan during year t. Given a plan's operational decisions q_t for year t, the cashflow from an individual is denoted by $\lambda(\iota_t, q_t)$. The cashflow is positive if the plan is collecting premia from the individual and negative if the plan is paying the individual (or their family) a pension. We denote by π_t the collection of the attributes of all individuals in the plan during year t. Then, the total net liabilities of the plan is $\ell(\pi_t, q_t) := \sum_{\iota_t \in \pi_t} \lambda(\iota_t, q_t)$.

The last two components are more straightforward. Some external funding may be injected into the plan in year t, which we denote e_t. Finally, third parties may charge a percentage annual fee m, to manage the assets.

The total value of the assets of the plan at the beginning of the year $t+1$ is

$$v_{t+1} = (1 + r(\alpha_t, p_t))(v_t + \ell_t(\pi_t, q_t) + e_t)(1 - m). \tag{1}$$

This model is agnostic on whether the values are expressed in real or nominal amounts (*i.e.*,, corrected by inflation). If the values are corrected by inflation, this would be reflected in the values of $r(\alpha_t, p_t)$, $\ell_t(\pi_t, q_t)$

and e_t. Moreover, this model assumes negligible transaction costs, but they could be included as an extra term.

3.2 Pension policy/decision strategy

At the beginning of each year, the pension plan's manager can take the following actions:

1. Change the portfolio p_t from a set of possible portfolios $\mathcal{P}$.
2. Ask the sponsor of the plan for an injection of external cash $e_t \in \mathbb{R}_{\geq 0}$.
3. Select an operational decision q_t that modifies the net liabilities. The set of permissible options $\mathcal{Q}$ is determined by the pension's rules. These options may include modification to premia and/or payout, the purchase of annuities or the offloading of liabilities via pension risk transfer.

For convenience, we will denote all the actions taken at time t by a_t, and the set of possible actions $\mathcal{A} \subset \mathcal{P} \times \mathbb{R}_{\geq 0} \times \mathcal{Q}$.

If there were no unknowns quantities, at $t = 0$ the pension plan manager could specify which actions to take at time $t = 1, 2, \ldots$ and guarantee that all of its liabilities would be met without ever risking default. In reality, pension plans manager react to the changing environment. Using a policy will allow us to model this recourse.

Let $\mathcal{Z}$ be the set of all possible information available at any time. We denote by $z_t \in \mathcal{Z}$ the specific information available at time t. In general, z_t will contain information both about known and unknown quantities. Examples of known quantities, are m, $v_t, v_{t-1}, \ldots$, $\alpha_{t-1}, \alpha_{t-2}, \ldots$ and $\pi_{t-1}, \pi_{t-2}, \ldots$. The information for unknown quantities (present or future), may either be an estimate or a probability distribution that models their behavior. Examples of unknown quantities are future returns $\alpha_t, \alpha_{t+1}, \ldots$ and future liabilities $\pi_t, \pi_{t+1}, \ldots$ A policy (also known as a decision strategy) $\varphi_\theta(\cdot)$ is a function parameterized by θ such that $\varphi_\theta(z_t) = a_t$. In other words, it determines which actions to take at each time given the available information given some choice of behavior captured in its parameters.

Before moving on, it is worth noting a few attributes of our framework. First, our approach involves changing actuarial liabilities and assets simultaneously. This is not the present approach of the industry, where current practice is (mostly) to model actuarial liabilities first, then convert them into contractual liabilities, against which assets are managed. Second, it is possible to apply our simple framework in a more complicated mortality pooling context, as in (Fullmer and Sabin, 2019). Third, as a contingent asset, a plan could also shed longevity or volatility risk for a price, for a period of time—as opposed to until the last beneficiary dies, at any price.

3.3 Simulating and evaluating a policy

To assess the effectiveness of a policy $\varphi_\theta(\cdot)$ parameterized by θ, we simulate the values of v_t according to (1) for $t = 1, 2, \ldots, T$ and use these simulations to compute various performance metrics. We start by selecting a probability distribution to model the unknown variables $\alpha_t, \alpha_{t+1}, \ldots$ and $\pi_t, \pi_{t+1}, \ldots$. Using this distribution, we generate S sample trajectories. These trajectories allow us to derive sample paths for v_t under the policy $\varphi_\theta(\cdot)$. With these sample paths, we can then calculate key metrics. For example, a pension plan manager might be interested in the expected value of the assets, the empirical probability of defaulting, and the average payout. We denote the metrics for a policy with parameters θ by $m_\theta \in \mathbb{R}^M$. This process is summarized by algorithm 1.

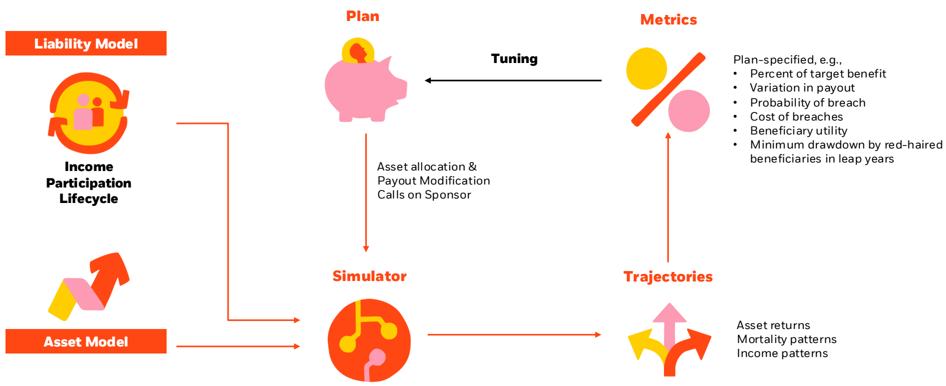

Figure 1: Schematic representation of the pension fund model.

Algorithm 1 Evaluating a policy using Monte Carlo

Start all world scenarios with an asset v_0
Sample S trajectories of asset returns $\alpha_1, \ldots \alpha_T$ and population $\pi_1, \ldots, \pi_T$
for $s = 1, \ldots, S$
  for $t = 1, \ldots, T$
    Determine all the information available $z_t^{(s)}$
    Evaluate $\varphi_\theta(z_t^{(s)})$ to obtain the actions $a_t^{(s)} = (p_t^{(s)}, e_t^{(s)}, q_t^{(s)})$.
    Sample from the distributions to obtain $r(\alpha_t^{(s)}, p_t^{(s)})$ and $\ell(\pi_t^{(s)}, q_t^{(s)})$.
    Compute $v_{t+1}^{(s)} = \left(1 + r(\alpha_t^{(s)}, p_t^{(s)})\right)\left(v_t^{(s)} + \ell(\pi_t^{(s)}, q_t^{(s)}) + e_t^{(s)}\right)(1 - m)$.
Compute metrics

3.4 Tuning the policy

Our objective in tuning the policy is to achieve acceptable outcomes according to the performance metrics. Since there are multiple performance metrics (or equivalently, we model the perfomance as a single vector), we introduce a cost function $h : \mathbb{R}^m \to \mathbb{R} \cup \{\pm\infty\}$ that generates a relative priority for improving different metrics. Tuning the policy means finding the set of parameters θ that finds a (possibly local) minimum of $h(m_\theta)$.

4 Implementation of the pension fund

The pension fund model of section 3 describes a general and abstract framework. In this section, we describe a specific implementation (or, more precisely, an instantiation) that we developed to illustrate these concepts. This is a notional implementation lacking some features that one would expect in an actual pension plan, such as using the duration and convexity as an input for the policy, but it is useful to show how even a simple implementation can lead to interesting results.

4.1 Model of portfolio returns

Our model of returns is given by

$$r(\alpha_t, p_t) = \alpha_t' w(p_t) \tag{2}$$

where α_t are returns of asset classes and $w_t(p_t)$ are the weights associated to portfolio p_t.

Let $\bar{\alpha}$ be the expected return of the asset classes and C the covariance matrix. We construct the weights of the portfolios $w(p_t)$ by using a Markowitz approach with risk aversion parameter $\mu_{p_t} > 0$ (Boyd et al., 2024). We define $w(p_t)$ to be the unique solution of the optimization problem

$$\begin{array}{ll} \text{maximize} & \bar{\alpha}'w - \frac{\mu_{p_t}}{2} w'Cw \\ \text{subject to} & \mathbf{1}'w = 1 \\ & w \geq 0 \\ & w \leq b_{in,up} \\ & \sum_{i \in \text{domestic}} w_i \in [b_{do,low}, b_{do,up}] \\ & \sum_{i \in \text{foreign}} w_i \in [b_{fo,low}, b_{fo,up}] \\ & \sum_{i \in \text{fixed income}} w_i \in [b_{fi,low}, b_{fi,up}] \\ & \sum_{i \in \text{equities}} w_i \in [b_{eq,low}, b_{eq,up}] \\ & \sum_{i \in \text{alternatives}} w_i \in [b_{alt,low}, b_{alt,up}] \end{array} \tag{3}$$

where $b_{in,up}$ is the maximum allocation for any asset class and $[b_{do,low}, b_{do,up}]$, $[b_{fo,low}, b_{fo,up}]$, $[b_{fi,low}, b_{fi,up}]$, $[b_{eq,low}, b_{eq,up}]$ and $[b_{alt,low}, b_{alt,up}]$ are lower and upper bounds on the total weights for asset classes that are domestic, foreign, fixed income, equities and alternatives, respectively.

4.2 Model of liabilities

As explained in section 3.1, the total liability is a sum over the cashflows of all the individuals in the pension plan. In our model, we simulate the life of each individual, or, more precisely, the change in the attributes ι_t through the life of each individual and when the individual dies. The individual's attributes we consider are biological factors (year of birth, biological sex) and socio-economic factors (income, place of residence, level of education). These attributes are used to construct an individualized mortality rate; we refer the reader to Appendix A for an explanation. Given each individual's mortality rate, we build a probability distribution on the trajectories of this individual's life over the next T years. The individual cashflows $\lambda(\iota_t, q_t)$ is a function of these trajectories.

Sampling life trajectories During the individual's life over the next T years, they can be in one of four states: alive and not working, alive and working, alive and retired, and dead. The probability of an individual dying is given by their mortality rate that year. If the individual is not dead, they can change state. If the individual is not working and is older than a minimum working age m_w, each year they have a probability of $1/y_w$ of starting to work (y_w is the average number of years it would take for them to start working). If the individual is working and is older than a minimum retiring age m_r, each year they have a probability of $1/y_r$ of retiring (y_r is the average number of years it would take for them to retire). Our model does not take into account unemployment, but this could easily be remediated by creating an extra state.

Individual cashflows During an individual life trajectory, if they are working, they pay $k_w\%$ of their income; if they are retired, they receive $k_r\%$ of the last salary they received.

Plan rules and operational strategies The set of permissible liability options $\mathcal{Q}$ (defined in section 3.2) is composed of m_w, m_r, y_w, y_r, k_w, and k_r. The values of m_w, m_r, y_w, y_r, and k_w are determined by the plan rules and cannot be changed from one year to the other. The value of k_r can be changed each year up to $\pm\tilde{k}_r\%$ from a baseline which we denote by $\bar{k}_r$ so $k_r = (100\% \pm \tilde{k}_r\%)\bar{k}_r$.

Example 1 (Example of rules for a pension plan)**.** The minimum age for starting to work is $m_w = 18$ years old, with an average of $y_w = 4$ years to start doing so. When working, someone contributes $k_w\% = 15\%$ of their income. The minimum age for retiring is $m_r = 63$ years old, with an average of $y_r = 2$ years for doing so. When retired, someone is promised a baseline $\bar{k}_r = 80\%$ of their salary before retiring. The actual value of the pension each year will vary by up to $\pm\tilde{k}_r\% = \pm 10\%$ meaning that a pensioner is guaranteed to receive $(100\% \pm 10\%)80\%$ of their salary before retiring.

We reference the net liabilities according to a difference in the payout level ($100\% \pm \tilde{k}_r\%$). For instance, $\ell(\pi_t, 100\%)$ denotes the net liability when the payments to retirees are at the baseline level.

4.3 Policy

The information available at time z_t is composed of:

1. the current portfolio value v_t,
2. the current baseline liabilities $\ell(\pi_t, 100\%)$,
3. the expected return of the asset classes defined by $\bar{\alpha}_t := \mathbb{E}[\alpha_t]$ provided by Janus,
4. the expected net baseline liability for the next T years given the population of year t, defined by $\bar{\ell}_{t+\tau} := \mathbb{E}[\ell(\pi_{t+\tau}, 100\%)],\ \tau = 1, \ldots, T$ provided by the model described in section 4.2,
5. the expected yield curve over the next T years as of year t, denoted by $\gamma_{t+\tau},\ \tau = 1, \ldots, T$, provided by Janus.

Let $L_t = \sum_{\tau=0}^{T} \gamma_{t+\tau} \bar{\ell}_{t+\tau}$ be the expected total net liabilities over the next T years. The first step of the policy is to compute an extended version of the asset to liability ratio, defined by

$$\rho_t = \begin{cases} v_t/L_t & \text{if } v_t \geq 0 \text{ and } L_t > 0 \\ +\infty & \text{if } v_t \geq 0 \text{ and } L_t \leq 0 \\ 0 & \text{if } v_t < 0 \end{cases} .$$

This formulation is motivated by the following observation. On one hand, if the total assets are negative, then the plan is in a dire state, independent of the future liabilities. On the other hand, when the assets are positive, if the future liabilities are not positive, then the plan is in a very good state.

We have already described in sections 4.1 and 4.2 the choices of portfolio p_t and liability options q_t. The last action we need to describe is how e_t is determined. The policy specifies a target asset to liability ratio at time $t+1$ denoted by $\widehat{\rho}_{t+1}$ and the external cash e_t is given by

$$e_t = \max \left\{ \frac{\widehat{\rho}_{t+1} L_{t+1}}{(1-m)(1 + w(p_t)'\bar{\alpha}_t)} - (v_t + \ell_t(q_t)), 0 \right\} \tag{4}$$

which is the value of e_t that would achieve an asset to liability ratio equal to $\widehat{\rho}_{t+1}$ at the beginning of year $t+1$ if the future liabilities and returns match the expected ones.

The policy is specified by discretizing the values of ρ_t and defining which actions to take in each bin. This leads to a table in which each row specifies the actions to be taken at time t. Table 1 is an example of what such table would look like. The parameters θ that indexes the policy are the values of the actions for each asset to liability ratio bin. The set of possible parameters Θ are the possible values of portfolio volatility, payout level and target asset to liability ratio.

4.4 Metrics and policy tuning

Policy evaluation The policy is evaluated according to the following metrics:

- yearly probability of performing a cash call c
- average payout level $\bar{q}$
- average change in payout level from one year to the next Δq

Asset to liability ratio ρ_t	**Portfolio** p_t indexed by volatility	**Liability option** q_t indexed by payout level	**External cash** e_t defined by $\widehat{\rho}_{t+1}$
Greater than 2.0	4.437%	110%	0
Between 1.8 and 2.0	7.891%	108%	0
Between 1.6 and 1.8	7.891%	107%	0
Between 1.4 and 1.6	7.891%	105%	0
Between 1.2 and 1.4	5.620%	102%	0
Between 1.0 and 1.2	5.620%	100%	0
Between 0.8 and 1.0	4.437%	97%	1.2
Less than 0.8	4.437%	95%	1.0

Table 1: Example of a table describing the policy

We scalarize the metrics using the following function. For each metric, the user specifies three parameters: lower acceptable bound, a high acceptable bound, and a priority. Then for each metric we associate a value:

$$\eta(\text{metric}) = \begin{cases} \text{priority(metric)} \cdot \text{outer\_slope(metric)} \cdot (\text{low(metric)} - \text{metric}) & \text{if metric} < \text{low(metric)} \\ \text{priority(metric)} \cdot \text{inner\_slope(metric)} \cdot (\text{low(metric)} - \text{metric}) & \text{if low(metric)} \leq \text{metric} < \text{midpoint(metric)} \\ \text{priority(metric)} \cdot \text{inner\_slope(metric)} \cdot (\text{metric} - \text{high(metric)}) & \text{if midpoint(metric)} \leq \text{metric} < \text{high(metric)} \\ \text{priority(metric)} \cdot \text{outer\_slope(metric)} \cdot (\text{metric} - \text{high(metric)}) & \text{if metric} \geq \text{high(metric)} \end{cases} \tag{5}$$

where

- $\text{midpoint(metric)} = 0.5\Big(\text{low(metric)} + \text{high(metric)}\Big)$
- $\text{inner\_slope(metric)} = 0.5\Big(\text{high(metric) - low(metric)}\Big)$
- $\text{outer\_slope(metric)} = 10 \text{ inner\_slope(metric)}$

The objective function that is tuned is $h(m_\theta) = \eta(c) + \eta(\bar{q}) + \eta(\Delta q)$.

Direct policy optimization We would like to find the optimal parameters θ that achieve the best cost $h(m_\theta)$. We do this using a derivative-free algorithm that achieves one-optimality, meaning that the policy cannot be improved by only changing the value of the action in one cell. We start by creating a discrete set of possible values of portfolio volatility, of payout levels, and of target asset to lability ratio $\widehat{\rho}$. All sets need to be organized in either increasing or decreasing order.

Algorithm 2 Algorithm to achieve one-optimality

Initialize table with values that seem reasonable.
Start all world scenarios with an asset a_0
counter = 0
while counter $\leq$ number of cells in the table
    for each row of the table
        for each column of the table
            for each allowed action corresponding to the current column
                Evaluate the policy using this action for this cell
            if any of the tested action produces a smaller value than the current best action
                Change θ to use this new action for this cell
                counter = 0
            else
                counter = counter + 1

5 Results

To demonstrate our framework, we choose implementations of the framework that take place in the United States. For the purpose of this research study, we use the demographic data from the Health and Retirement Study from the University of Michigan (RAND, 2024). The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant numbers NIA U01AG009740 and NIA R01AG073289) and is conducted by the University of Michigan (HRS, 2024). This is a nationally representative data set that includes extensive information about individuals socio-demographic and economic characteristics, including income and mortality. From this dataset, we use the information from income, level of education, residency, biological sex, data of birth and date of death which we combine with the mortality rate projections created by the United Nations Department of Economic and Social Affairs Population Division (United Nations, Department of Economic and Social Affairs, Population Division, 2024) in order to create a Cox proportional hazards model, as described in Appendix A.

The asset classes returns are provided by Janus (van Beek, 2020), a tool developed by the BlackRock Investment Institute. It generates random trajectories of annual returns for asset classes and macroeconomic projections for over 600 indicators for 50 years. Of these indicators, we selected a set of 25 asset classes denominated in U.S. Dollars. They include bond (*e.g.,*, U.S. Treasury bonds with maturities of 1, 3, 5, 7, 10, 20+ years), equities (*e.g.,*, U.S. large cap, U.S. small cap, China A Shares) and some alternatives (*e.g.,*, infrastructure debt, real estate).

The characteristics of the instantiation are:

- **Population:** The simulated population is composed of 1000 individuals with gender, ages and socio-economic background that are representative of the United States population. The plan is closed to new entrants that are born after the initial year, but is open to entrants that are already born.
- **Plan rules:** Each individual is allowed to start working when they are 18 years old, and, on average, start working when they are 22 years old. While they are working, they contribute 10% of their income. They are allowed to retire once they turn 60 years old, and, on average, retire when they are 65 years old. Once retired, they expect a baseline pension equivalent of 80% of their average income over the previous 20 years.
- **Simulation length:** The policy is simulated for 30 years, from 2025 to 2055. Each year, the total liabilities are computed by summing the projected liabilities over the following 30 years discounted by the projection of the yield curve at that date. So, for instance, in the year 2035, the total liabilities consist the projected liabilities from 2035 to 2065, discounted by the projected yield curve of 2035.

- **Initial capitalization:** All plans start with the same amount of assets. This initial amount of asset was chosen such that the plan start with an asset to liability ratio of 1.6.
- **Asset to liability bins:** We discretize the asset to liability ratio into the following bins: $[0.0, 0.2)$, $[0.2, 0.4) \dots [2.2, \infty)$.

By running our machine with different pension plan design features, we can test the effectiveness of various pension plan set-ups. Our ultimate aim is to be able to compare the efficiency of different pension plans across a spectrum from high cost (like a PRT-insurer provided annuity) to low cost (like a flexible "adaptable annuity" that dynamically changes asset allocation and payout level, subject to plan design constraints). We run four different instantiations of pension plan designs within our framework.

- **Plan A:** Designed to be an approximation of a Pension Risk Transfer annuity (fixed income only, static allocation, paying beneficiaries 100% of target always). Various costs associated with a realistic PRT, such as the costs of hedging, cost of capital, or pension insurance premia, are not included, but would only further inflate the costs to operate this plan.
- **Plan B:** Introduces flexibility on the payouts: fixed income only, static allocation, paying beneficiaries between 90–110% of target (with payout level varying by no more than 2% of target in any year).
- **Plan C:** Introduces flexibility on the asset allocation as well: fixed income only, variable allocation, paying beneficiaries between 90–110% of target (with payout level varying by no more than 2% of target in any year).
- **Plan D:** Introduces equities in the asset allocation, in addition to fixed income.

Using the framework described above, we find that compared to a plan with a static fixed income asset allocation and a fixed benefit payout (Plan A), a plan with adaptive strategies of a dynamic fixed income asset allocation and a variable payout within a limited range (Plans C and D) can achieve the same payouts while decreasing the probability of asking for external funding by half for the same initial capital contribution.

The plans are different in the following way. Plan A is used to model the general behavior of an annuity. The other plans increase the degree of freedom to make the plan more adaptable to uncertainties. Table 2 summarize the difference between the plans.

	Plan A	**Plan B**	**Plan C**	**Plan D**
Benefit % of target maximum year over year change	100% 0%	90% to 110% 2%	90% to 110% 2%	90% to 110% 2%
Assets available	Fixed income	Fixed income	Fixed income	Fixed income and equities
Asset allocation	Static, annual rebalance to portfolio	Static, annual rebalance to portfolio	Adaptive portfolio allocations	Adaptive portfolio allocations

Table 2: Difference between Plans A, B, C, and D

We run each plan over 10000 simulated world scenarios (*i.e.,* 10000 trajectories of liabilities and 10000 trajectories of asset returns) and compute metrics of these results. The results are available in table 3 and figs. 2a to 2d. Increasing the degrees of freedom of the plans, improves performance, either by distributing a larger benefit to the pensioners or by decreasing the risk of requiring a cash call. Notice that the level at which each plan is considered breached is automatically learned, and is the optimal given the plan's degrees of freedom. This happens even though all the plans share the same objectives. The automatically learned breach asset-to-liability thresholds are 1.2, 1.0, 0.6, and 0.6 for Plans A, B, C, and D, respectively.

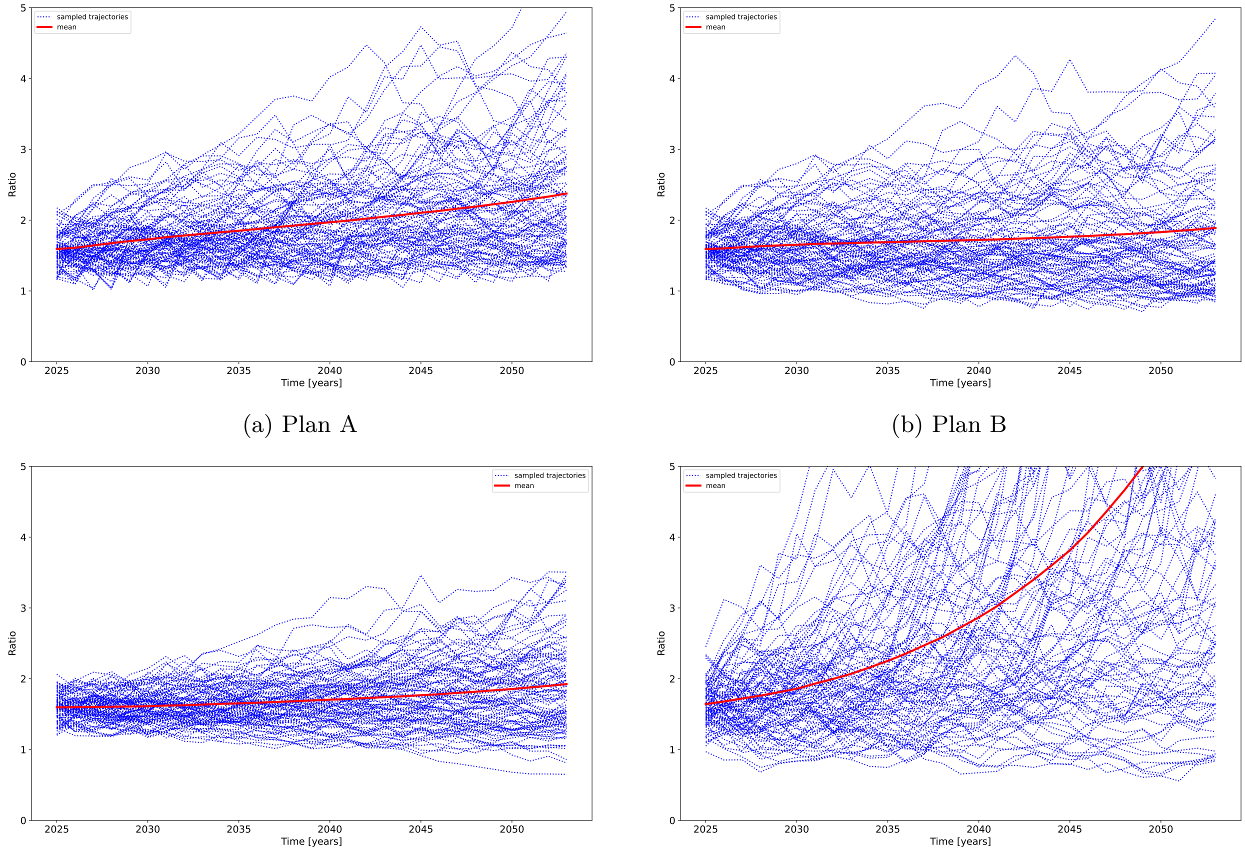

(a) Plan A (b) Plan B

(c) Plan C (d) Plan D

Figure 2: Simulations of asset to liability ratio under the different plans.

Metric	Plan A	Plan B	Plan C	Plan D
Mean percent of target benefit	100%	99.9%	102.4%	104.0%
Probability of breaching A:L floor in one year (30y)	1.32% (53%)	0.88% (25%)	0.21% (7.7%)	0.19% (6.0%)
Sum of ex post value of breaches as a percentage of initial funds	38%	13%	2.9%	3.0%

Table 3: Metrics

It is worth reiterating that the Plan A described here, even with its relative expense, is still cheaper than an annuity would be in practice. Our Plan A does not deduct from returns the cost of capital or hedging risks from interest rates or foreign exchange fluctuations. When these costs are included, we expect that Plans B, C, and D will be even cheaper by comparison.

6 Conclusion

The principal contribution of this paper is a general framework to address both longevity risk and return uncertainty in providing a low(er) cost source of lifetime income. The core idea is to create a simulator that generates trajectories of asset returns and liabilities and use them to tune a policy. While this idea is widespread in the field of control, applying an engineering solution of this type appears to be less common in finance.

Every country, every pension system, every individual will ultimately provide “adaptive” answers to the problem of sustaining consumption in the future from whatever pot of assets may have been accumulated. By directly approaching the idea of an adaptive pension seems an important place to start.

Our second contribution was a specific policy. While this policy is rather simple, it was used to illustrate a key result: including just a small amount of flexibility can substantially decrease risk and cost by allowing the pension plan to adapt to new circumstances. In our simple policy, this was restricted to only changes in portfolio allocation and pension payout, but could be used in an array of other decisions that are currently fixed by plan rules. If designed correctly, this approach would benefit both pensioners and sponsors because it trades small amounts of risk for substantial decreases in cost and reliability.

Throughout this paper, we have mentioned several possible extensions of our work. Some of these are necessary for an investment-grade solution:

- build out asset-liability duration matching,
- accurately model management and other fees,
- add derivatives to the list of eligible asset classes,
- implement a robust income model,
- increase the capability of the asset allocation and interest rate models, and
- give a portfolio manager the ability to implement a bond ladder for liabilities less than two years away and turn on the algorithm for longer-durated liabilities

Some extensions are more conceptual. For example, the modeling approach required to predict the socio-demographic characteristics of an open-ended beneficiary populations is non-obvious. Also, various retirement finance systems have established a reserve fund that must be paid back if the pension funds tap it; modelling this would be a helpful extension of the model.

7 Funding

The authors are employees of BlackRock and completed this as part of their employment there.

Appendix

A Mortality rate and income projections

A.1 Mortality rates through the Cox model

Each individual is characterized by three core characteristics: year of birth, biological sex, and country of residency. In addition to these core characteristics, individuals are also described by socio-economic covariate, such as income, level of education, etc. We assume we have at our disposal a dataset which contains individuals covariate over a period of some years and their dates of deaths or whether they are alive by the end of the period.

There are several institutions around the world, such as the United Nations, that provide a mortality rate for an individual given the three core characteristics. Our goal is to understand how this mortality rate is affected by covariates.

A possible way to model longevity for an individual with covariate $X \in \mathbb{R}^d$ is to use the Cox model (Cox, 1972). The hazard function is given by

$$P(T \in [t, t+dt] \mid T \geq t, X) = \lambda(t) \exp\{\theta^\top X\};$$

here,

1. $\lambda(t)$ is the baseline mortality and assumed to be given by the United Nations;
2. $\exp\{\theta^\top X\}$ is an exponential term that captures the effect of the covariates on the hazard function;
3. $\theta \in \mathbb{R}^d$ is the vector of regression coefficients which needs to be estimated from the data.

The survival function Let T be the random variable representing the time of death. Let $f_\theta(T, X)$ be the density function of T and $\bar{F}_\theta(T, X)$ be the survival function (the chance of making to at least T). Then by definition,

$$\frac{\mathrm{d} \log \bar{F}_\theta(t, X)}{\mathrm{d}t} = -\lambda(t) \exp\{\theta^\top X\}.$$

This implies that the survival function is given by

$$\bar{F}_\theta(t, X) = \exp\left\{ -\int_{t_0}^{t} \lambda(s) \exp\{\theta^\top X\} \,\mathrm{d}s \right\},$$

where t_0 represents the moment when the individual starts to be tracked. It also follows that the density function is given by

$$f_\theta(t, X) = \lambda(t) \exp\{\theta^\top X\} \, \bar{F}_\theta(t, X).$$

Maximum likelihood estimation We have i.i.d. data $(T_i, X_i), \ldots, (T_n, X_n)$, where T_i is either the time of death or the time of censoring for the ith individual and X_i is the covariate vector for the ith individual. Let δ_i be the indicator of whether the time of death of the ith individual is observed. The likelihood function for the Cox model is given by

$$\prod_{i=1}^{n} \{f_\theta(T_i, X_i)\}^{\delta_i} \left\{\bar{F}_\theta(T_i, X_i)\right\}^{1-\delta_i}.$$

The likelihood takes the form

$$L(\theta) = \prod_{i=1}^{n} \left\{\lambda(T_i) \exp\{\theta^\top X_i\}\right\}^{\delta_i} \exp\left\{ -\int_{t_0}^{T_i} \lambda(s) \exp\{\theta^\top X_i\} \,\mathrm{d}s \right\},$$

Therefore, the log-likelihood is given by

$$\ell(\theta) = \sum_{i=1}^{n} \left\{ \delta_i \left\{ \log \lambda(T_i) + \theta^\top X_i \right\} - \exp\{\theta^\top X_i\} \int_{t_0}^{T_i} \lambda(s)\, \mathrm{d}s \right\}.$$

Setting $\Lambda(t) = \int_{t_0}^{t} \lambda(s)\, \mathrm{d}s$, the maximum likelihood estimator $\hat{\theta}$ is the solution of

$$\text{minimize} \ \sum_{i=1}^{n} \Lambda(T_i)\ \exp\{\theta^\top X_i\} - \delta_i \theta^\top X_i.$$

This is a convex optimization problem and can be solved using standard optimization algorithms. (Recall that $\lambda(t)$ and, therefore, $\Lambda(t)$ are known.)

In our implementation, we are interested in understanding how the baseline mortality rate for a population in the United States changes according to income, years of education and location. They are treated as categorical:

- Income is separated in 4 quantiles (rich, upper middle class, lower middle class, poor) each representing 25% of the population.
- Education is separated into 4 levels (high school dropout, high school degree, college degree, advanced degree).
- Location is separated into the 4 census regions (Northeast, South, Midwest, West).

A.2 Model of income

The income of an individual is determined by their age and the income bin in which they categorize (see the end of the previous section). Let α be the individual age. Then their income is given by

$$\text{income}(\alpha) = \begin{cases} 25\beta & \text{if } \alpha < 25 \\ \alpha\beta & \text{if } \alpha \in [25, 45] \\ 45\beta & \text{if } \alpha > 45 \end{cases}$$

where β is a constant that depends on the individual's income bin (which is fixed throughout a person's life). These values of β are learned from our dataset. For our dataset, we obtained poor = 322, lower middle class = 748, upper middle class = 1394, and rich = 2880.

References

Ang, A., Chen, B., Sundaresan, S.M., 2012. Liability Driven Investment with Downside Risk .

Bauer, D., Börger, M., Ruß, J., 2010. On the Pricing of Longevity-Linked Securities. Insurance: Mathematics and Economics 46, 139–149.

van Beek, M., 2020. Consistent calibration of economic scenario generators: The case for conditional simulation. arXiv preprint arXiv:2004.09042 URL: https://arxiv.org/abs/2004.09042.

Boyd, S., Johansson, K., Kahn, R., Schiele, P., Schmelzer, T., 2024. Markowitz portfolio construction at seventy. arXiv preprint arXiv:2401.05080 .

Boyer, M.M., Box-Couillard, S., Michaud, P.C., 2020. Demand for Annuities: Price Sensitivity, Risk Perceptions, and Knowledge. Journal of Economic Behavior & Organization 180, 883–902.

Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2005. Consumption and Portfolio Choice over the Life Cycle. The Review of Financial Studies 18, 491–533.

Cox, D.R., 1972. Regression models and life-tables. Journal of the Royal Statistical Society: Series B (Methodological) 34, 187–202.

Ezra, D.D., 1991. Asset Allocation by Surplus Optimization. Financial Analysts Journal 47, 51–57.

Forman, J.B., Sabin, M.J., 2014. Tontine Pensions: A Solution to the State and Local Pension Underfunding Crisis. University of Pennsylvania Law Review 163.

Friedman, M., 1957. The Permanent Income Hypothesis, in: A Theory of the Consumption Function. Princeton University Press, pp. 20–37.

Fullmer, R.K., Garcia Huitron, M.E., Winter, P., 2025. Minimizing payout volatility in longevity risk-sharing pools: An asset. Matching Approach (April 20, 2025) .

Fullmer, R.K., Sabin, M.J., 2019. Individual Tontine Accounts .

Goldsticker, R., 2007. A Mutual Fund to Yield Annuity-Like Benefits. Financial Analysts Journal 63, 63–67.

HRS, 2024. Health and retirement study, (rand hrs longitudinal file 2020 (v2)) public use dataset. National Institute on Aging and the Social Security Administration. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant numbers NIA U01AG009740 and NIA R01AG073289).

Kochenderfer, M.J., 2015. Decision Making under Uncertainty: Theory and Application. MIT press.

Kochenderfer, Mykel J; Wheeler, T.A., Wray, K.H., 2022. Algorithms for decision making. MIT press.

Maher, G., Boyd, S., Kochenderfer, M., Matache, C., Reuter, D., Ulitsky, A., Yukhymuk, S., Kopman, L., 2022. A Light-Weight Multi-Objective Asynchronous Hyper-Parameter Optimizer. arXiv preprint arXiv:2202.07735 .

Milevsky, M.A., 2015. King William's Tontine: Why the Retirement Annuity of the Future Should Resemble its Past. Cambridge University Press.

Milevsky, M.A., 2022. How to Build a Modern Tontine: Algorithms, Scripts and Tips. Springer Nature.

Modigliani, F., Brumberg, R., 1954. Utility analysis and the consumption function: An interpretation of cross-section data, in: Kurihara, K.K. (Ed.), Post-Keynesian Economics. Rutgers University Press, New Brunswick, pp. 388–436.

Morduch, J., Schneider, R., 2017. The Financial Diaries: How American Families Cope in a World of Uncertainty. Princeton University Press.

Ogata, K., 2020. Modern Control Engineering.

Pang, G., Warshawsky, M., 2010. Optimizing the Equity-Bond-Annuity Portfolio in Retirement: The Impact of Uncertain Health Expenses. Insurance: Mathematics and Economics 46, 198–209.

RAND, 2024. Rand hrs longitudinal file 2020 (v2). National Institute on Aging and the Social Security Administration. URL: `https://hrsdata.isr.umich.edu/data-products/rand-hrs-longitudinal-file-2020`. produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration.

Sharpe, W.F., Tint, L.G., 1990. Liabilities-A New Approach. Journal of Portfolio management 16, 5–10.

Tirole, J., 2017. Economics for the Common Good. Princeton University Press.

United Nations, Department of Economic and Social Affairs, Population Division, 2024. World population prospects 2024: Mortality rate projections. URL: `https://population.un.org/wpp/downloads?folder=Probabilistic%20Projections&group=Mortality`. accessed: 2025-03-28.

Zheng, H., Thomas, L.C., Allen, D.E., 2002. The Duration Derby: A Comparison of Duration Based Strategies in Asset Liability Management, in: Proceedings of the 41st IEEE Conference on Decision and Control, 2002., IEEE. pp. 769–774.