

Money-Back Tontines for Retirement Decumulation: Neural-Network Optimization under Systematic Longevity Risk

German Nova Orozco *

Duy-Minh Dang[†]

Peter A. Forsyth[‡]

February 18, 2026

Abstract

Money-back guarantees (MBGs) are features of pooled retirement income products that address bequest concerns by ensuring the initial premium is returned through lifetime payments or, upon early death, as a death benefit to the estate. This paper studies optimal retirement decumulation in an individual tontine account with an MBG overlay under international diversification and systematic longevity risk. The retiree chooses withdrawals and asset allocation dynamically to trade off expected total withdrawals (EW) against the Conditional Value-at-Risk (CVaR) of terminal wealth, subject to realistic investment constraints. The optimization is solved under a plan-to-live convention, while stochastic mortality affects outcomes through its impact on mortality credits at the pool level. We develop a neural-network based computational approach for the resulting high-dimensional, constrained control problem. The MBG is priced ex post under the induced EW–CVaR optimal policy via a simulation-based actuarial rule that combines expected guarantee costs with a prudential tail buffer. Using long-horizon historical return data expressed in real domestic-currency terms, we find that international diversification and longevity pooling jointly deliver the largest improvements in the EW–CVaR trade-off, while stochastic mortality shifts the frontier modestly in the expected direction. The optimal controls use foreign equity primarily as a state-dependent catch-up instrument, and implied MBG loads are driven mainly by tail outcomes (and the chosen prudential buffer) rather than by mean payouts.

Keywords: defined contribution, tontine, money-back guarantee, stochastic mortality, portfolio optimization, Conditional Value-at-Risk, neural network

AMS Subject Classification: 93E20, 91G10, 91B30, 62P05, 68T07

1 Introduction

The worldwide shift from Defined Benefit (DB) pensions to Defined Contribution (DC) arrangements is well documented [40, 58]. While DC plans offer flexibility and portability, they also transfer much of the responsibility for managing key retirement risks to individuals. In particular, upon retirement, a DC member faces the *decumulation problem*: how to invest and draw down accumulated savings to support sustainable real (inflation-adjusted) spending under longevity risk (the possibility of outliving savings) and uncertain market returns [3, 5, 33]. Rapid global population ageing further intensifies these challenges [60].

*School of Mathematics and Physics, The University of Queensland, St Lucia, Brisbane 4072, Australia, email: g.novaorozco@student.uq.edu.au

[†]School of Mathematics and Physics, The University of Queensland, St Lucia, Brisbane 4072, Australia, email: duyminh.dang@uq.edu.au

[‡]David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, N2L 3G1, Canada, email: paforsyt@uwaterloo.ca

To manage longevity risk in practice, many retirees rely on simple spending and asset allocation heuristics, most notably the “4% rule” [4] (often paired with a fixed stock–bond mix) and performance-based adjustments to withdrawals and/or portfolio weights (e.g. [23]). These rules are transparent and easy to implement, and this reliance on rule-based decumulation is consistent with the evidence in [1], which reports a revealed preference for spending rules among retirees and wealth advisors in DC drawdown settings. However, because these rules are not derived from a risk–reward optimization, they can be far from efficient under realistic market and wealth conditions (see, e.g. [17, 19]).

As an alternative to rules-based drawdown strategies, retirees can transfer longevity risk to an insurer by purchasing a life annuity. In practice, voluntary annuitization remains limited (e.g. [42]), and low demand for life annuities can be rational once bequest motives, illiquidity costs, and product design are taken into account [33].

In parallel, there has been renewed interest in *pooled* retirement income products that share longevity risk among members while offering greater transparency and flexibility than traditional annuities. Modern tontines and related survivor pools achieve longevity pooling by redistributing the accounts of deceased members to survivors, generating *mortality credits* that can support higher sustainable real (inflation-adjusted) withdrawals for members who remain alive [15, 16, 22, 37]. The appeal comes with a clear trade-off: mortality credits raise payouts conditional on survival, but wealth is typically forfeited upon death.

In [21], a stochastic optimal control framework is developed for an individual tontine retirement account. It operationalizes the idea of modern tontines in an individual setting by adding a tontine overlay that redistributes the balances of deceased members to survivors. In this framework, the retiree has full control over both the withdrawal amount (subject to minimum and maximum constraints) and the asset allocation in the account. The optimization objective is defined using a risk–reward criterion: reward is measured by total expected accumulated real (inflation-adjusted) withdrawals (EW) over a fixed retirement horizon (30 years), while risk is measured by Conditional Value-at-Risk (CVaR), also known as expected shortfall, of end-of-horizon real wealth (assuming the retiree survives to the horizon, i.e. “plan-to-live, not to die”; see, e.g. [43]).

In that setting, the results in [21] show that longevity pooling can materially improve the withdrawal–risk trade-off relative to non-pooled drawdown strategies and simple constant withdrawal/allocation benchmarks. In particular, for a reasonable level of tail-risk tolerance (CVaR), the tontine overlay delivers substantially higher expected cumulative real withdrawals, even after allowing for “fees of the order of 50–100 basis points (bps) per year” [21][Section 1].

This efficiency gain, however, comes at the cost of forfeiting account wealth upon death, which can limit the appeal of “pure” tontine overlays for retirees with bequest or estate considerations [21]. Reflecting this practical constraint, recent industry innovation—particularly in the Australian superannuation market—has introduced *money-back guarantees* (MBGs), sometimes described in product disclosures as “money-back protection”, for pooled retirement income products (e.g. QSuper and MyNorth [38, 49]). An MBG overlay can be added to an otherwise pure tontine, as in [21], and ensures that the member receives at least the initial purchase price (a nominal dollar amount fixed at inception) through lifetime withdrawals and/or a death benefit paid to the estate. Because the MBG is settled only upon death and is paid to the estate, it does not affect the account dynamics while the retiree is alive. Importantly, in practice the funding mechanism for MBGs need not appear as a stand-alone member-level charge and can change over time (e.g. insured funding versus pool-funded self-insurance) [48, 49]. This practical reality motivates cost measures that are economically comparable across operational implementations, rather than tied to any particular fee design.

From the perspective of the entity that ultimately bears the guarantee, whether an insurer or the

retirement-income pool, the MBG is a low-frequency, high-severity death-benefit liability: it is triggered only upon death, but conditional on occurrence the payout can be large. Accordingly, the economic cost of an MBG depends not only on expected payouts, but also on tail outcomes and on any prudential buffer/reserving rule used to fund adverse scenarios. Crucially, the distribution of guarantee payouts is endogenous to retiree behavior, since withdrawal and asset allocation choices determine the account balance at death. Hence, MBG valuation cannot be carried out independently of the retiree’s decumulation policy, even though the MBG itself does not feed back into the retiree’s optimal controls under the plan-to-live convention.

From a modelling perspective, two aspects are central in tontine decumulation: the specification of the asset market and the treatment of mortality risk. While most individual account decumulation models assume a domestic stock-bond portfolio, in practice, for retirees, the investable universe is inherently richer and often includes international assets [11, 56]. International diversification is particularly important for investors in countries where the local equity market capitalization is small relative to the global market. Extending the individual account decumulation setting to allow international diversification is economically important, since it can affect the risk-return trade-offs and, in turn, optimal withdrawal and asset allocation decisions. Allowing multiple risky asset classes, however, makes the control problem high-dimensional and renders grid-based dynamic programming computationally prohibitive. Recent work in portfolio optimization and related control problems has shown that neural network (NN) parameterizations of policies can compute high-dimensional, state-dependent controls without suffering from the curse of dimensionality inherent in grid methods [8, 31, 39].

A second modelling aspect is mortality at the pool level: deaths determine the redistribution to surviving members and hence the mortality credits. Following the discussion in [21], an important caveat is systematic mortality risk (e.g. unexpected improvements in life expectancy), which is typically ignored when mortality credits are computed from a fixed life table. To capture this source of uncertainty, stochastic mortality specifications such as Lee-Carter (LC) [30] and Cairns-Blake-Dowd (CBD) [7] can be used to generate time-varying one-year death probabilities and hence stochastic mortality credits. If mortality improves unexpectedly, realized mortality credits can be lower than anticipated because lighter realized mortality implies fewer deaths and hence less redistribution to survivors. This separates idiosyncratic longevity risk (diversifiable within the pool) from systematic longevity risk (non-diversifiable).

Motivated by the above observations, this paper sets out to achieve three primary objectives. First, we formulate a multi-period EW-CVaR optimal decumulation framework for an individual retirement account with a tontine overlay in an internationally diversified setting with realistic long-horizon DC constraints. The individual optimization problem is solved conditional on survival over a fixed retirement horizon, while stochastic mortality enters through pathwise mortality-credit inputs. Conditioning on survival to the end of a 30-year horizon (e.g. from age 65 to 95) is a prudent stress test for sustaining real (inflation-adjusted) retirement spending and aligns with the common “plan-to-live, not to die” convention used in practice.¹

Second, we develop a computational approach capable of solving the resulting high-dimensional, constrained control problem using NN parameterizations of state-dependent policies, and integrate this with a simulation-based MBG pricing framework under endogenous retiree behavior and alternative prudential-buffer (tail-risk) funding rules for MBGs.

Third, using nearly a century of realized market data, we quantify the impact of international diversification (at the individual account level) and stochastic mortality (at the pool level) on optimal tontine decumulation outcomes, EW-CVaR frontiers, and MBG loads, and assess the relative importance of

¹Survival to age 95 from age 65 is far from certain under standard life tables; under recent Australian population life tables, it is about 0.11 for a 65-year-old male.

these modelling dimensions.

Concretely, the numerical illustration is conducted from the perspective of a representative domestic retiree with access to domestic and foreign equity and government bond indices. Foreign returns are converted into domestic-currency real (inflation-adjusted) returns so that outcomes are measured in units of real (inflation-adjusted) retirement spending. The empirical implementation reported in Section 9 uses an Australian-investor calibration with long-horizon data, but the modelling and computational framework is not country-specific and can be applied elsewhere given the corresponding asset and mortality inputs.

Our main conclusion is that international diversification is economically material even before tontines. The gains are most pronounced when diversification is combined with a tontine overlay, which is the primary driver of improvements in the EW–CVaR withdrawal–risk trade–off. Allowing for stochastic mortality at the pool level shifts outcomes in the expected direction—longevity improvement reduces mortality credits—but does not alter the qualitative picture. Finally, when MBGs are summarized in equivalent-load terms, the expected-cost (actuarially fair) component is modest in our analysis, while larger implied loads are driven primarily by how strongly the funding rule places weight on tail outcomes (prudential buffers) rather than by mean payouts.

The remainder of the paper is organized as follows. Section 2 introduces the tontine modelling framework and the mortality specifications (deterministic and stochastic). Section 3 describes the money-back guarantee (MBG) overlay mechanism. Section 4 presents the stochastic control framework for individual-account decumulation, and Section 5 formulates the associated EW–CVaR optimization problem. Section 6 develops the NN-based computational approach, and Section 7 integrates it with a Monte Carlo pricing method for the MBG. Validation results are reported in Section 8. Section 9 presents data construction and preprocessing details, the internationally diversified tontine results, MBG pricing implications, and sensitivity analyses. Section 10 concludes and outlines directions for future research.

2 Tontine modeling

We let $\mathcal{T}$ denote the set of pre-determined, equally spaced decision times in $[0, T]$, at which mortality credit distributions, withdrawals, and portfolio rebalancing occur:²

$$\mathcal{T} = \{t_m \mid t_m = m\Delta t, m = 0, \dots, M\}, \quad \Delta t = T/M, \quad (2.1)$$

where $t_0 = 0$ is the inception time and $T > 0$ is the finite investment horizon. Throughout we take annual decision times, so $\Delta t = 1$ year and $T = M$.

For later use, we define $t^- = t - \varepsilon$ and $t^+ = t + \varepsilon$ to represent the instants just before and after any time $t \in [0, T]$, with $\varepsilon \rightarrow 0^+$. For a generic time-dependent function $f(t)$ and any $t_m \in \mathcal{T}$, we write

$$f_{m-} = \lim_{\varepsilon \rightarrow 0^+} f(t_m - \varepsilon), \quad f_{m+} = \lim_{\varepsilon \rightarrow 0^+} f(t_m + \varepsilon),$$

as shorthand for $f(t_m^-)$ and $f(t_m^+)$, respectively.

2.1 Modeling of individual tontine accounts

We follow the framework of [21] for individual tontine accounts. The tontine pool consists of J members, indexed by $j = 1, \dots, J$. At each time point $t_m \in \mathcal{T}$, we identify the actions that occur at the three successive instants: $t_m^- \rightarrow t_m \rightarrow t_m^+$. These actions apply only in the event of solvency, that is, when the investor’s wealth is strictly positive at time t_m^- .

²The assumption of equal spacing is made for simplicity. In practice, rebalancing schedules are typically fixed (e.g., semi-annually or annually), rather than irregular.

- t_m^- (for $m = 1, \dots, M$): end-of-period portfolio balances are observed. Accounts of members who died in $[t_{m-1}^+, t_m^-]$ are forfeited, and the corresponding mortality credits are immediately distributed to the surviving members, prior to any investor actions (i.e. withdrawals and rebalancing).
- t_m and t_m^+ (for $m = 0, \dots, M-1$): each surviving investor withdraws an amount q_m at t_m , and then rebalances their portfolio at t_m^+ .

Between t_{m-1}^+ and t_m^- , $m = 1, \dots, M$, the underlying asset processes and mortality evolve continuously under their respective (stochastic) dynamics. At $t_M = T$, the portfolio is liquidated with no withdrawal or rebalancing.

To proceed with mortality credits, we introduce the notation for individual account balances and formalize the survival and forfeiture events over $[t_{m-1}^+, t_m^-]$. A detailed discussion of withdrawals, rebalancing, and the incorporation of mortality credits is provided in Section 5.

We denote by W_{m-}^j the real (inflation-adjusted) portfolio balance (or wealth) of member j at time t_m^- , $m = 0, \dots, M$. At the beginning of each period $[t_{m-1}^+, t_m^-]$, $m = 1, 2, \dots, M$, all pool members are assumed to be alive and each holds an account balance of $W_{(m-1)+}^j \geq 0$. If a member dies during the period, their entire account is forfeited and distributed to the surviving members of the pool as mortality credits. Mathematically, we define the survival indicator and its conditional expectation

$$\mathbf{1}_m^j = \begin{cases} 1, & \text{member } j \text{ survives to } t_m^-, \\ 0, & \text{member } j \text{ dies during } [t_{m-1}^+, t_m^-], \end{cases} \quad \mathbb{E}_{m-1}[\mathbf{1}_m^j] = 1 - \delta_{m-1}^j. \quad (2.2)$$

Here, δ_{m-1}^j is the probability that member j dies in $[t_{m-1}^+, t_m^-]$, namely

$$\delta_{m-1}^j = \text{Prob}(\text{member } j \text{ dies during } [t_{m-1}^+, t_m^-] \mid \text{alive at } t_{m-1}^+). \quad (2.3)$$

In addition, $\mathbb{E}_{m-1}[\cdot] = \mathbb{E}[\cdot \mid \mathcal{F}_{m-1}]$ denotes expectation conditional on the σ -algebra $\mathcal{F}_{m-1}$, which contains all information available at t_{m-1}^+ (the start of $[t_{m-1}^+, t_m^-]$). If $\mathbf{1}_m^j = 0$, the portfolio balance W_{m-}^j is forfeited. If $\mathbf{1}_m^j = 1$, member j receives a mortality credit c_m^j at time t_m .

We now describe how c_m^j is determined in two stages: (i) an individual-level fairness condition, and (ii) a pool-level budget constraint.

2.1.1 Individual fairness condition

Assuming no fees—that is, we ignore administration, investment-management, and transaction costs for simplicity—the tontine is structured so that, in every period, participation yields zero expected net gain in advance. In other words, a member's expected forfeited wealth in case of death during a period is exactly balanced by the expected mortality credit they receive at the end of that period if they survive. Specifically, conditioning on information available at time t_{m-1}^+ , the fairness condition for the period $[t_{m-1}^+, t_m^-]$ is

$$W_{m-}^j \mathbb{E}_{m-1}[1 - \mathbf{1}_m^j] = (1 - \delta_{m-1}^j) \mathbb{E}_{m-1}[c_m^j \mid \Omega_m^j], \quad \text{where } \Omega_m^j = \left\{ \mathbf{1}_m^j = 1, \{W_{m-}^k\}_{k=1}^J \right\}. \quad (2.4)$$

That is, $\mathbb{E}_{m-1}[\cdot \mid \Omega_m^j]$ denotes the expectation conditional on all information available at t_{m-1}^+ , together with (i) the event that member j survives the period ($\mathbf{1}_m^j = 1$), and (ii) the realized end-of-period wealth balances of all members $\{W_{m-}^k\}_{k=1}^J$, before mortality credits are calculated and distributed.

Solving (2.4) yields the actuarially fair mortality credit distribution formula

$$\mathbb{E}_{m-1}[c_m^j \mid \Omega_m^j] = \frac{\delta_{m-1}^j}{1 - \delta_{m-1}^j} W_{m-}^j. \quad (2.5)$$

According to this rule, each surviving member's expected mortality credit is proportional to their own account balance, scaled by the ratio of their death probability to their survival probability. While the

right-hand side of (2.5) depends only on $(W_{m-}^j, \delta_{m-1}^j)$, this independence is not exact for a finite and heterogeneous pool: it is possible to construct scenarios in which the available forfeitures are insufficient to support (2.5) simultaneously for all surviving members, leading to a bias that favours some members over others (see the discussion in [21]). Following [54], we therefore interpret (2.5) as a large-pool approximation whose accuracy hinges on a *small-bias condition*, stated explicitly in Subsection 2.1.3.

2.1.2 Pool-level budget constraint

While equation (2.5) guarantees fairness for each member, the tontine must also balance cash flows at the pool level: the total mortality credits paid to surviving members at t_m must equal the total wealth forfeited by deceased ones during the period $[t_{m-1}^+, t_m^-]$. This requirement gives the budget rule³

$$\sum_{j=1}^J \mathbf{1}_m^j c_m^j = \sum_{j=1}^J (1 - \mathbf{1}_m^j) W_{m-}^j, \quad (2.6)$$

which must hold ex post in every period. Because (2.6) must hold with the actual forfeitures and credits—determined by the random, realized number of deaths in the period, which could be different from the expected count—enforcing this budget rule exactly can be cumbersome in a finite pool. We therefore introduce a practical adjustment in the next subsection.

2.1.3 Large-pool approximation

To obtain a simple, tractable rule, [54] introduce a pool-wide adjustment factor referred to as the group gain Γ_m defined as

$$\Gamma_m = \frac{\sum_{k=1}^J (1 - \mathbf{1}_m^k) W_{m-}^k}{\sum_{k=1}^J \mathbf{1}_m^k \frac{\delta_{m-1}^k}{1 - \delta_{m-1}^k} W_{m-}^k}. \quad (2.7)$$

Here, the numerator is the realized forfeiture during $[t_{m-1}^+, t_m^-]$, while the denominator is the total expected mortality credit for the members who actually survive that period, obtained by summing the fair-credit expression (2.5) $\frac{\delta_{m-1}^k}{1 - \delta_{m-1}^k} W_{m-}^k$ over all k with $\mathbf{1}_m^k = 1$.

Multiplying the fair-credit expectation (2.5) by Γ_m gives the simplified mortality credit distribution formula

$$c_m^j = \frac{\delta_{m-1}^j}{1 - \delta_{m-1}^j} W_{m-}^j \Gamma_m, \quad j = 1, \dots, J, \quad (2.8)$$

which satisfies the budget rule (2.6) exactly, even for finite pools. However, for finite and heterogeneous pools, the sharing rule (2.8) need not satisfy the individual fairness condition (2.5) exactly; instead, it can introduce a (typically small) bias in expected gains across members. It is shown in [54] that this bias is negligible under the following *small-bias condition*: (a) the pool is sufficiently large, and (b) the expected aggregate forfeiture in the period is large compared to any member's nominal gain, i.e.

$$\frac{\delta_{m-1}^j}{1 - \delta_{m-1}^j} W_{m-}^j \ll \sum_{k=1}^J \delta_{m-1}^k W_{m-}^k, \quad j = 1, \dots, J. \quad (2.9)$$

Condition (2.9) is essentially a diversification requirement: no member has an abnormally large share of the pool capital. When (2.9) holds (and the pool is large enough that realized deaths are close to their expectation), the random group-gain factor satisfies $\mathbb{E}[\Gamma_m] \simeq 1$ with small variance [54]. Accordingly, we adopt the approximation $\Gamma_m \equiv 1$ for the remainder of the paper.

³More generally, the mortality credit can be written as $c_m^j = \frac{\delta_{m-1}^j}{1 - \delta_{m-1}^j} W_{m-}^j H_m^j$, where $H_m^j \geq 0$ is a sharing factor. Fairness is preserved by imposing $\mathbb{E}_{m-1}[H_m^j | \Omega_m^j] = 1$; see [21] for details.

With the large-pool approximation $\Gamma_m \equiv 1$, the member-level mortality credit distribution rule (2.8) becomes

$$c_m^j = \frac{\delta_{m-1}^j}{1 - \delta_{m-1}^j} W_{m-}^j, \quad j = 1, \dots, J. \quad (2.10)$$

The rule (2.10) continues to satisfy, with high accuracy under the large-pool/small-bias assumptions (Condition (2.9)): (i) Fairness: $\mathbb{E}_{m-1}[c_m^j | \Omega_m^j] = \frac{\delta_{m-1}^j}{1 - \delta_{m-1}^j} W_{m-}^j$, (ii) Budget constraint: Equation (2.6) (exactly when Γ_m is retained as in (2.8), and approximately under the large-pool approximation $\Gamma_m \equiv 1$), and (iii) Non-negativity: $c_m^j \geq 0$ for all j .

In our optimal control formulation, we suppress the superscript j and consider a representative surviving member. The mortality credit distribution rule (2.10) then reduces to a per-member expression that depends only on the member's own total wealth at time t_m^- and their mortality rate. Based on this, we define the mortality credit c_m in terms of the tontine gain rate, denoted by g_m , as follows:

$$c_m = g_m W_{m-}, \quad \text{where} \quad g_m = \frac{\delta_{m-1}}{1 - \delta_{m-1}}. \quad (2.11)$$

The tontine gain rate can be interpreted as $g_m = \frac{\text{fraction of the cohort expected to die}}{\text{fraction expected to survive}}$, so it represents the proportional uplift each surviving member receives from mortality pooling during $[t_{m-1}^+, t_m^-]$. Equation (2.11) is the key mechanism used to compute mortality credits in the optimal control formulation developed subsequently in the paper.

2.2 Mortality models

Recall from (2.2)–(2.3) that longevity risk enters the tontine through the one-year conditional death probabilities δ_{m-1}^j for each member j over the interval $[t_{m-1}^+, t_m^-]$. In a homogeneous pool, we write $\delta_{m-1}^j \equiv \delta_{m-1}$, so that the tontine gain rate g_m in (2.11) is fully determined by the annual sequence $\{\delta_{m-1}\}_{m=1}^M$. We now describe how this sequence is obtained under deterministic and stochastic mortality.

2.2.1 Deterministic mortality

In the deterministic case, we work with a standard period life table, such as the Canadian Pensioner Mortality Tables or Australian mortality from the Human Mortality Database (HMD) [25], which provide one-year death probabilities $q_{x,y}$ for an individual aged x in calendar year y . Let the retiree be aged x_0 at retirement, which occurs in calendar year y_0 . In our tontine framework the decision times are measured in years since retirement, so $t_m = m$ for $m = 0, \dots, M$ with $t_0 = 0$. At decision time t_m , the retiree is age $x_0 + m$ in calendar year $y_0 + m$.

We extract from the life table the corresponding one-year conditional death probabilities

$$\delta_{m-1} := q_{x_0+(m-1), y_0+(m-1)}, \quad m = 1, \dots, M. \quad (2.12)$$

In a homogeneous pool, these probabilities are common to all members, so that $\delta_{m-1}^j = \delta_{m-1}$ for every j , and both $\{\delta_{m-1}\}$ and the corresponding tontine gain rates $\{g_m\} = \{\delta_{m-1}/(1 - \delta_{m-1})\}$ are deterministic.

2.2.2 Stochastic mortality

To incorporate systematic longevity risk, we allow the mortality surface to be generated by a stochastic model in the generalized age-period-cohort (GAPC) family; see [62]. Let $D_{x,y}$ denote the number of deaths at age x in calendar year y and $E_{x,y}$ the corresponding central exposure. We assume

$$D_{x,y} \sim \text{Poisson}(m_{x,y} E_{x,y}), \quad (2.13)$$

with central death rate $m_{x,y}$. Its systematic component is captured by a linear predictor

$$\eta_{x,y} = \alpha_x + \sum_{i=1}^N \beta_x^{(i)} \kappa_y^{(i)} + \beta_x^{(0)} \gamma_{y-x}, \quad (2.14)$$

where α_x describes the age profile, $\kappa_y^{(i)}$ are period factors and γ_{y-x} is a cohort effect. A link function g relates $\eta_{x,y}$ to the mortality quantity of interest, e.g.

$$\log m_{x,y} = \eta_{x,y} \quad \text{or} \quad \text{logit } q_{x,y} = \eta_{x,y},$$

with $\text{logit}(u) = \log(u/(1-u))$ for $u \in (0,1)$. In this general GAPC formulation the cohort term is optional; the specific Lee–Carter [30] and Cairns–Blake–Dowd [7] models used below are special cases that omit the cohort component.

Lee–Carter (LC) model

In the Lee–Carter specification, the predictor has a single age–period term,

$$\eta_{x,y} = \alpha_x + \beta_x \kappa_y, \quad (2.15)$$

and we typically set $\log m_{x,y} = \eta_{x,y}$. The period index κ_y captures the overall mortality level and is commonly modelled as a random walk with drift,

$$\kappa_y = \delta + \kappa_{y-1} + \xi_y, \quad \xi_y \sim N(0, \sigma_\kappa^2) \text{ i.i.d.} \quad (2.16)$$

together with the usual identifiability constraints on $(\alpha_x, \beta_x, \kappa_y)$ as implemented in StMoMo [62].

Cairns–Blake–Dowd (CBD) model

The Cairns–Blake–Dowd model [7] describes mortality as approximately linear in age around a reference age $\bar{x}$:

$$\eta_{x,y} = \kappa_y^{(1)} + (x - \bar{x}) \kappa_y^{(2)}. \quad (2.17)$$

Here it is natural to work directly with one–year death probabilities and set

$$\text{logit } q_{x,y} = \eta_{x,y}.$$

The bivariate period factor $\boldsymbol{\kappa}_y = (\kappa_y^{(1)}, \kappa_y^{(2)})^\top$ is usually specified as a random walk with drift,

$$\boldsymbol{\kappa}_y = \boldsymbol{\delta} + \boldsymbol{\kappa}_{y-1} + \boldsymbol{\xi}_y, \quad \boldsymbol{\xi}_y \sim N(\mathbf{0}, \Sigma_\kappa) \text{ i.i.d.}, \quad (2.18)$$

In contrast to the LC model, no additional identifiability constraints are required for this basic two–factor CBD specification [62].

From stochastic mortality models to tontine gain rate. Once an LC or CBD model has been calibrated to historical deaths and exposures, its fitted and projected period factors determine a surface of one–year death probabilities $\{q_{x,y}\}$. For a representative retiree aged x_0 in calendar year y_0 , with annual decision times $t_m = m$ and corresponding calendar years $y_m = y_0 + m$, we define

$$\delta_{m-1} := q_{x_0+(m-1), y_0+(m-1)}, \quad m = 1, \dots, M, \quad (2.19)$$

and in a homogeneous pool we set $\delta_{m-1}^j \equiv \delta_{m-1}$ for all members j , exactly as in the deterministic life–table case. These probabilities feed directly into the tontine gain rate $g_m = \delta_{m-1}/(1 - \delta_{m-1})$ in (2.11). Thus the tontine mechanics developed above apply unchanged; only the sequence $\{\delta_{m-1}\}$ differs between deterministic life–table mortality and stochastic GAPC-based mortality.

Remark 2.1 (Mortality inputs in simulation). *The same one–year conditional death probabilities also underpin our simulation framework. In the deterministic case, the sequence $\{\delta_{m-1}\}_{m=1}^M$ is obtained from a period life table as in Subsection 2.2.1, and we take $\delta_{m-1}^{(k)} \equiv \delta_{m-1}$ on every simulation path, indexed by k . Under stochastic LC or CBD mortality, each simulated mortality surface $\{q_{x,y}^{(k)}\}$ yields a pathwise sequence $\{\delta_{m-1}^{(k)}\}_{m=1}^M$ via*

$$\delta_{m-1}^{(k)} := q_{x_0+(m-1), y_0+(m-1)}^{(k)}, \quad m = 1, \dots, M. \quad (2.20)$$

On path k , these probabilities are used to construct tontine gain rates $g_m^{(k)} = \delta_{m-1}^{(k)}/(1 - \delta_{m-1}^{(k)})$, $m = 1, \dots, M$, which enter the wealth recursion in the NN training (Section 6) and are also used to generate death times and payouts in the money-back guarantee valuation (Section 7).

3 Money-back guarantee overlay

“Your purchase price is always paid back as either income to you or a death benefit paid to your beneficiaries. If you die, the death benefit is equal to the amount you paid for your Lifetime Pension, less the payments that have gone to you...”

QSuper’s Lifetime Pension Product Disclosure Statement [49]

This promise reflects the essence of the MBG, a recent industry innovation introduced by leading Australian superannuation providers, including QSuper and MyNorth [38, 49]. The MBG can be added to an otherwise “pure” tontine and ensures that every dollar a member invests at inception (i.e. their initial contribution or purchase price) is returned—either through lifetime withdrawals or, if the member dies early, as a death benefit paid to their estate.

3.1 Description

Specifically, at inception, suppose a member invests an amount L_0 (ie. the purchase price) to the retirement product with MBG. At each decision time $t_m \in \mathcal{T}$, as in the regular tontine described in Section 2, if the member survives to t_m^- , they first receive the mortality credit c_m defined in (2.11), then withdraw an amount q_m , and rebalance their portfolio. Once cumulative withdrawals reach or exceed the purchase price L_0 , the MBG becomes inactive.

If the member dies at time $\tau \in (t_{m-1}^+, t_m^-]$, no further withdrawal occurs, including the scheduled withdrawal at t_m . If their cumulative withdrawals up to t_{m-1} fall short of the purchase price L_0 (a nominal dollar amount fixed at inception), the MBG activates. The member’s estate receives the shortfall:

$$\text{MBG-payout} = \max \left(L_0 - \sum_{\ell=0}^{m_\tau-1} q_\ell \frac{\text{CPI}_\ell}{\text{CPI}_0}, 0 \right), \quad (3.1)$$

where m_τ is the time index such that death occurs in the interval $[t_{m_\tau-1}^+, t_{m_\tau}^-]$, and the right-hand side is expressed in nominal dollars (unadjusted for inflation). Here, CPI_ℓ is the consumer-price index at decision time t_ℓ (and CPI_0 is the index level at inception). Because the guarantee compares nominal cash flows, each real dollar withdrawal q_ℓ is first expressed in nominal terms by multiplying by $\text{CPI}_\ell/\text{CPI}_0$ before being summed in (3.1). For valuation and reporting, this nominal shortfall is converted to a real (inflation-adjusted) amount at inception by multiplying by $\text{CPI}_0/\text{CPI}_{m_\tau}$; see Section 7 (Eqn. (7.1)).

While it is straightforward to incorporate time-value discounting, we follow [19] in setting the real discount rate to zero. This conservative assumption, common in retirement-income studies, typically implies a constant discount factor of 1 throughout.

For instance, suppose a retiree makes an initial contribution of \$200,000 at t_0 , and withdrawals are scheduled annually ($\Delta t = 1$ year). By the end of year 4 (t_4), the member has withdrawn a total of \$65,000 in nominal terms. If the member then dies during the subsequent interval $[t_4^+, t_5^-]$ —that is, sometime between the end of year 4 and the scheduled decision time at year 5—no withdrawal is executed at t_5 . The MBG is triggered and the member’s estate receives: $\text{MBG-payout} = \max(200,000 - 65,000, 0) = \$135,000$.

Remark 3.1 (Timing convention). *In our discrete-time model, we follow the convention that the estate of a member who dies during the interval $[t_{m-1}^+, t_m^-]$ receives the MBG payout at the end of that interval. As described in Section 2, the member’s account balance at t_m^- contributes to the mortality credits distributed to members surviving the period. However, the account is then removed prior to the distribution, and no further withdrawals are made, including the scheduled withdrawal at t_m . In the above example, the MBG amount of \$135,000 will be paid at year 5.*

This end-of-interval convention is widely adopted in actuarial asset–liability models (see, for example, [14, Section 4] and [36, Section 2]), as it simplifies implementation and introduces only a first-order bias

$\mathcal{O}(\Delta t)$ arising solely from using the end-of-year CPI level when converting between real and nominal amounts. With $\Delta t = 1$ year, this bias is negligible compared with market and longevity risk.

Importantly, the MBG operates as an overlay implemented at the product/pool level, for example, through pool-funded self-insurance or an external insurance arrangement. Therefore, it does not alter the tontine’s internal mechanics: the pooling of longevity risk, mortality credit distributions, withdrawals, and rebalancing all remain as specified in Section 2.

In practice, the MBG may be subject to regulatory constraints, such as Australia’s Capital Access Schedule (CAS), which limits the refundable portion of a pension’s purchase price. While it is straightforward to incorporate the CAS into the model, we do not do so here in order to maintain focus on the core mechanics of the tontine and the MBG overlay pricing. As a result, the recoverable amount under the MBG may be overstated relative to a regulated product.

3.2 MBG pricing load

Product disclosures for pooled retirement income products typically emphasize the presence of MBGs (sometimes described as money-back protection), but the mechanisms used to fund these guarantees need not appear as stand-alone member-level charges and can evolve over time.

For QSuper’s Lifetime Pension, used as our industry example, the Product Disclosure Statement (PDS) introduces “money-back protection” and does not present it as a stand-alone recurring member-level insurance fee [49]. A recent product update further notes that money-back protection was previously offered through an external insurance policy (ART Life Insurance Limited), and that from 1 July 2025 the Trustee ceased this insurance arrangement; instead, “the money-back protection benefit can be funded from the Lifetime Pension pool directly” [48][p. 2].

Consistent with this pool-level funding approach, the current PDS explains that “The Lifetime Pension pool’s annual financial results will affect the annual income adjustment in the following year. The results include: . . . the mortality experience of the pool . . . all fees and costs,” and that “The money-back protection benefit (if applicable) is paid from the pool” [49, p. 34, p. 113].

Accordingly, in this paper we treat the MBG funding mechanism as a modelling choice and summarize its economic cost using an equivalent up-front load factor $f_g \in (0, 1)$. The quantity $f_g L_0$ can be interpreted as a transparent one-time cost measure—equivalently, a proportional reduction in a notional starting income rate or benefit base—that would finance the MBG under the pricing rule adopted in this paper. This translation device allows the MBG cost to be reported in a comparable way across alternative practical implementations (e.g. pool-funded self-insurance versus insured funding versus implicit benefit adjustments). Full details of the pricing methodology and its numerical implementation are provided in Section 7.

4 Stochastic control framework

We now turn to the stochastic control framework, where we model the evolution of an individual member’s portfolio and formulate the associated dynamic optimization problem. We consider a portfolio held by a domestic investor with access to four real (inflation-adjusted) asset classes: (i) a domestic stock index fund, (ii) a domestic bond index fund, (iii) a foreign stock index fund (converted to domestic currency), and (iv) a foreign bond index fund (converted to domestic currency). This setup allows us to examine both an internationally diversified portfolio, which includes all four asset classes, and a non-diversified alternative that is restricted to domestic assets only. The construction of these indices, the inflation adjustment, and the data sources used for calibration are detailed in Sections 8 and 9.

4.1 Index dynamics

For simplicity and clarity, we establish the following notational conventions. A subscript $\iota \in \{d, f\}$ is used to distinguish quantities related to the domestic stock or bond ($\iota = d$) from those corresponding to the foreign counterparts ($\iota = f$). Additionally, a superscript “ s ” denotes quantities associated with stock indices, while a superscript “ b ” identifies those related to their bond index counterparts.

Let $S^d(t)$, $B^d(t)$, $S^f(t)$, and $B^f(t)$ denote the real (inflation-adjusted) *amounts* invested in the stock and bond indices of the domestic and foreign markets, respectively, at time $t \in [0, T]$. To avoid notational clutter, we occasionally use the shorthand notation: $S_t^d \equiv S^d(t)$, $B_t^d \equiv B^d(t)$, $S_t^f \equiv S^f(t)$, and $B_t^f \equiv B^f(t)$. We denote by $\{X_t\}_{0 \leq t \leq T}$ (resp. by x) the controlled underlying index process (resp. a generic state of the system), where

$$X_t \quad (\text{resp. } x) = \begin{cases} (S_t^d, B_t^d) & (\text{resp. } (s^d, b^d)), & \text{domestic-only,} \\ (S_t^d, B_t^d, S_t^f, B_t^f) & (\text{resp. } (s^d, b^d, s^f, b^f)), & \text{internationally diversified.} \end{cases} \quad (4.1)$$

Between decision times $t_m \in \mathcal{T}$, the process evolves passively, driven solely by index-return dynamics. For any $t \in [t_m^+, t_{m+1}^-]$, we write

$$X_t = \mathcal{M}_{t_m, t}(X_{t_m^+}, \varepsilon_{m+1}), \quad t \in [t_m^+, t_{m+1}^-], \quad (4.2)$$

where $\mathcal{M}_{t_m, t}$ is a transition operator and ε_{m+1} collects all exogenous drivers acting over the interval $[t_m^+, t_{m+1}^-]$. These drivers may include Brownian or jump shocks (in parametric models) or resampled blocks of historical returns (in bootstrapped settings). The operator $\mathcal{M}_{t_m, t}$ maps the post-decision state $X_{t_m^+}$ and the exogenous drivers on $[t_m^+, t]$ to the state X_t .

In this paper, we consider two specifications for the driver set ε_{m+1} as follows.

- Parametric (jump–diffusion) model: For benchmarking against PDE-based methods, we implement a two-asset domestic-only case in which the domestic stock and bond indices $\{S_t^d\}$ and $\{B_t^d\}$ follow Kou-type jump–diffusion dynamics with asymmetric double-exponential jumps, capturing empirically observed heavy tails [28]. The full SDE specification appears in Subsection 8.1
- Historical block bootstrap: In our main experiments, both domestic-only and internationally diversified, we simulate asset paths nonparametrically via the stationary block bootstrap [13, 41, 45, 46]. Implementation details of the bootstrapping techniques are provided in Subsection 9.1.

In both cases, mortality-credit realizations (when applicable) are simulated independently and applied only at t_m^- . The resulting dynamics $\{X_t\}_{0 \leq t \leq T}$ from (4.2) capture the passive evolution of the index values between decision times. Active decisions—mortality updates, withdrawals, and rebalancing—are applied only at $\{t_m\}_{m=0}^{M-1}$, as detailed in the next subsection.

4.2 Mortality updates and control framework

We define the investor’s total portfolio wealth, hereafter referred to as “total wealth,” at time t as

$$W_t = \begin{cases} S_t^d + B_t^d, & \text{domestic-only,} \\ S_t^d + B_t^d + S_t^f + B_t^f, & \text{internationally diversified.} \end{cases} \quad (4.3)$$

The term “total wealth” refers to the sum of the values of the investor’s (tontine) account plus any accumulated debt arising from insolvency due to the minimum required withdrawals. In the event of solvency, we impose the investment constraints that (i) shorting stock and (ii) using leverage (i.e. borrowing) are not permitted. In case of insolvency, the portfolio is liquidated and trading stops. Debt accumulates at the borrowing rate, and no further mortality credits are applied. Importantly, minimum withdrawals continue in the event of insolvency, so once the account is exhausted, they are funded by borrowing and contribute to the accumulated debt.

Recall the set of decision times $\mathcal{T}$ defined in (2.1). At each $t_m \in \mathcal{T}$, the sequence of actions at $t_m^- \rightarrow t_m \rightarrow t_m^+$ described in Subsection 2.1 applies. To simplify bookkeeping, we adopt a uniform event structure across all $t_m \in \mathcal{T}$, including time t_0 for mortality credit and t_M for withdrawal as follows: in the event of solvency at t_{m-} , namely the total wealth $W_{m-} > 0$ as defined in (4.3) the following actions occur:

- At each time t_m^- , $m = 0, 1, \dots, M$, the mortality credit c_m defined in (2.11) is applied, with the convention that $c_0 = 0$, since no forfeitures have occurred at t_0 .
- At each time t_m (for $m = 0, \dots, M-1$), the investor withdraws an amount q_m , followed by portfolio rebalancing at t_m^+ . At t_M , the portfolio is liquidated (i.e. no rebalancing or withdrawal occurs), and terminal wealth W_T is realized. For notational completeness, this is enforced by setting $q_M = 0$.

To enforce no mortality credit at $t = 0$ and for the case of insolvency, we modify the definition of tontine gain rate g_m in (2.11) as follows

$$g_m = \begin{cases} \left(\frac{\delta_{m-1}}{1-\delta_{m-1}} \right) & m = 1, \dots, M, \\ 0 & m = 0 \text{ or } W_{m-} \leq 0. \end{cases} \quad (4.4)$$

Let

$$\widetilde{W}_{m-} = \begin{cases} (1 + g_m) (S_{m-}^d + B_{m-}^d), & \text{domestic only,} \\ (1 + g_m) (S_{m-}^d + B_{m-}^d + S_{m-}^f + B_{m-}^f), & \text{internationally diversified.} \end{cases} \quad (4.5)$$

That is, $\widetilde{W}_{m-}$ is the value of the portfolio immediately after the mortality credit distribution at t_m^- but before any fee, withdrawal, or rebalancing.

Next, we introduce a baseline tontine (management) fee that is deducted once per year, at each decision time t_m , $m = 1, \dots, M$. In line with Australian superannuation practice, and in particular the administration fee structure in QSuper's Lifetime Pension PDS [49], we model this as a proportional charge on the account balance. Let $\varrho \in (0, 1)$ denote the yearly fee rate (e.g. $\varrho = 0.11\%$ for an 11 bp charge).

To ensure the tontine fee is applied only when the investor is solvent and is omitted at t_0 , we define

$$\varrho_m = \varrho \mathbf{1}_{\{m \geq 1, \widetilde{W}_{m-} > 0\}}, \quad m = 0, \dots, M, \quad \widetilde{W}_{m-} \text{ given by (4.5)}. \quad (4.6)$$

With the conventions established in (4.4)–(4.6), the total wealth at $t_m^- \in \mathcal{T}$ after applying the mortality credit and deducting the tontine fee, and before any investor actions, is given by

$$W_{m-} = (1 - \varrho_m) \widetilde{W}_{m-}, \quad \widetilde{W}_{m-} \text{ given by (4.5)} \quad (4.7)$$

Unless otherwise stated, we use W_{m-} to denote the investor's wealth immediately before withdrawal (and rebalancing), but after mortality credits have been distributed and tontine fees deducted, as defined in (4.7).

Following this, the total wealth after processing the withdrawal amount q_m is

$$W_{m+} = W_{m-} - q_m, \quad t_m \in \mathcal{T}, \quad W_{m-} \text{ given by (4.7)}. \quad (4.8)$$

We model the withdrawal amount q_m as a withdrawal control, for $m = 0, \dots, M$, representing a strategy that depends on the total wealth W_{m-} and time t_m . Recalling the convention that $q_M = 0$, we define the withdrawal control function as $q_m(\cdot) : (W_{m-}, t_m) \mapsto q_m = q(W_{m-}, t_m)$, where W_{m-} is given by (4.7).

At each rebalancing time t_m , for $m = 0, \dots, M-1$, we denote the rebalancing control by $\mathbf{p}_m(\cdot)$, the vector of proportions of total wealth allocated to the asset indices. This control depends on the

current time t_m and on the total wealth W_{m+} after the cash withdrawal q_m in (4.8). Formally, $\mathbf{p}_m(\cdot) : (W_{m+}, t_m) \mapsto \mathbf{p}_m = \mathbf{p}(W_{m+}, t_m)$, where

$$\mathbf{p}_m = \begin{cases} (p_{s,m}^d), & \text{domestic-only case,} \\ (p_{s,m}^d, p_{b,m}^d, p_{s,m}^f), & \text{internationally diversified case.} \end{cases} \quad (4.9)$$

We denote by X_{m+} the state of the system immediately after applying the rebalancing control $\mathbf{p}_m$, where

$$X_{m+} = \begin{cases} (S_{m+}^d, B_{m+}^d), & \text{domestic-only,} \\ \text{where } S_{m+}^d = p_{s,m}^d W_{m+}, \quad B_{m+}^d = p_{b,m}^d W_{m+}, \\ (S_{m+}^d, B_{m+}^d, S_{m+}^f, B_{m+}^f), & \text{internationally diversified.} \\ \text{where } \begin{cases} S_{m+}^d = p_{s,m}^d W_{m+}, & B_{m+}^d = p_{b,m}^d W_{m+} \\ S_{m+}^f = p_{s,m}^f W_{m+}, & B_{m+}^f = p_{b,m}^f W_{m+}, \end{cases} \end{cases} \quad (4.10)$$

Here, in the domestic-only case, we define $p_{b,m}^d = 1 - p_{s,m}^d$, while in the internationally diversified case, $p_{b,m}^f := 1 - p_{s,m}^d - p_{b,m}^d - p_{s,m}^f$.

We denote by $\mathcal{Z}_q$ and $\mathcal{Z}_p$ the sets of all admissible withdrawal controls and rebalancing controls, respectively. For every $t_m \in \mathcal{T}$ we require $q_m \in \mathcal{Z}_q$ and $\mathbf{p}_m \in \mathcal{Z}_p$. A control at time t_m is therefore the pair $(q_m(\cdot), \mathbf{p}_m(\cdot))$, and we write $\mathcal{Z}$ for the set of all such admissible pairs. The sets are determined by withdrawal constraints (for $\mathcal{Z}_q$) and investment constraints (for $\mathcal{Z}_p$).

Contrary to the common perception that retiree spending is largely inflexible from year to year, survey evidence suggests that retirees adjust their lifestyle in response to cash-flow changes, including in categories often perceived as “fixed” expenses [2]. This supports modelling withdrawals as a flexible control within bounds, rather than as a rigid rule.

To reflect this evidence, and consistent with industry practice, we impose lower and upper bounds, $q_{\min}$ and $q_{\max}$, on withdrawals. Formally, for $t_m \in \mathcal{T}$, with W_{m-} given by (4.7), we define the admissible withdrawal control set as

$$\mathcal{Z}_q(W_{m-}, t_m) = \begin{cases} [q_{\min}, q_{\max}], & t_m \neq t_M, W_{m-} \geq q_{\max}, \\ [q_{\min}, \max\{q_{\min}, W_{m-}\}], & t_m \neq t_M, W_{m-} < q_{\max}, \\ \{0\}, & t_m = t_M. \end{cases} \quad (4.11)$$

While somewhat complicated, the piecewise structure in (4.11) captures the intuition that the retiree aims to avoid insolvency as much as possible by tightening the upper withdrawal limit when wealth declines. If wealth falls below $q_{\min}$, the retiree can still withdraw the minimum amount, accepting that this may lead to insolvency (i.e. $W_{m+} < 0$). At the terminal time t_M , no withdrawal occurs, as the portfolio is liquidated.

This constraint also admits a natural economic interpretation: $q_{\max}$ represents the retiree’s desired annual level of real (inflation-adjusted) spending, while $q_{\min}$ is a contingency floor the retiree is willing to adopt, potentially temporarily, to reduce the risk of depletion. The optimal control, therefore, uses the flexibility in (4.11) to preserve higher withdrawals when the account is well funded, but cuts withdrawals toward $q_{\min}$ in adverse market/wealth states to mitigate the long-run consequences of drawing at the maximum rate during a downturn.

This state-contingent spending cut mechanism is closely related to practitioner rules such as the *Canasta* strategy, which advocates reducing withdrawals following poor market performance.⁴ Time-varying bounds (e.g. imposing $q_{\min} = q_{\max}$ over an initial period to force high early spending) can also be

⁴<https://www.theglobeandmail.com/report-on-business/math-prof-tests-investing-formulas-strategies/article22397218/>; <https://cs.uwaterloo.ca/~paforsyt/Canasta.html>.

accommodated, but would mechanically increase the likelihood of early depletion when adverse returns occur early.

Remark 4.1 (Minimum withdrawals under insolvency). *At first sight it might seem natural to cease withdrawals once the account is depleted. However, the lower bound $q_{\min}$ is intended to represent a required minimum level of (real) spending. Accordingly, when $W_{m-} < q_{\min}$ we still enforce the minimum withdrawal by allowing $W_{m+} = W_{m-} - q_m < 0$, which is equivalent to borrowing the shortfall; the resulting debt then accrues at the borrowing rate. This convention ensures that tail-risk measures of terminal wealth penalize insolvency and, in particular, penalize early depletion more than late depletion (since debt has more time to accumulate).*

This convention also admits a practical interpretation in which minimum spending can be funded from assets outside the managed account (e.g. housing equity monetized via a reverse mortgage) once financial wealth is exhausted [44]. This “hedge of last resort” interpretation is also consistent with mental accounting views in which housing wealth is treated as a separate bucket that is tapped only in extreme circumstances, and otherwise can form part of the retiree’s bequest [55].

We enforce no leverage and no shortselling when the portfolio is solvent, and once the account becomes insolvent, we halt trading entirely and direct all wealth to the domestic bond. To define the set of admissible rebalancing controls, we first introduce

$$\Delta^{(k)} := \left\{ (p_1, \dots, p_k) \in [0, 1]^k \mid \sum_{i=1}^k p_i \leq 1 \right\}, \quad \mathbf{o}^{(1)} := (0), \quad \mathbf{b}^{(3)} := (0, 1, 0),$$

where $\Delta^{(k)}$ denotes the k -dimensional simplex of explicitly parameterized asset proportions. We set $k = 1$ for the domestic-only portfolio and $k = 3$ for the internationally diversified case. Under these conventions, the admissible rebalancing control set at time $t_m \in \mathcal{T}$ is given by

$$\mathcal{Z}_p(W_{m+}, t_m) = \begin{cases} \left\{ \Delta^{(k)}, & t_m \neq t_M, W_{m+} > 0, \right. \\ \left. \left\{ \{\mathbf{o}^{(1)}\}, & t_m \neq t_M, W_{m+} \leq 0, \right. \right. \\ \left. \left\{ \{\mathbf{o}^{(1)}\}, & t_m = t_M, \right. \right\} & \text{domestic-only,} \\ \text{or} & \\ \left\{ \{\mathbf{b}^{(3)}\}, & t_m \neq t_M, W_{m+} \leq 0, \right. \\ \left. \left\{ \{\mathbf{b}^{(3)}\}, & t_m = t_M, \right. \right\} & \text{internationally diversified.} \end{cases} \quad (4.12)$$

Here, W_{m+} is the post-withdrawal wealth defined in (4.8). In the domestic-only case, $\mathbf{o}^{(1)}$ corresponds to zero allocation to the explicitly parameterized (domestic stock) asset, implying full allocation to the domestic bond. In the internationally diversified case, $\mathbf{b}^{(3)} = (0, 1, 0)$ sets the explicitly parameterized proportions on domestic stock (0), domestic bond (1), and foreign stock (0), thereby assigning zero allocation to the foreign bond. If withdrawals render the account insolvent ($W_{m+} \leq 0$), trading stops and the negative balance accrues at the borrowing rate (a spread above the bond rate). The portfolio is liquidated at the terminal time t_M .

The admissible control set for the pair $(q_m(\cdot), \mathbf{p}_m(\cdot))$, $t_m \in \mathcal{T}$, can then be written as follows

$$(q_m, \mathbf{p}_m) \in \mathcal{Z}(W_{m-}, W_{m+}, t_m) = \mathcal{Z}_q(W_{m-}, t_m) \times \mathcal{Z}_p(W_{m+}, t_m), \quad (4.13)$$

where $\mathcal{Z}_q(W_{m-}, t_m)$ and $\mathcal{Z}_p(W_{m+}, t_m)$ are defined in (4.11) and (4.12), respectively. Let $\mathcal{A}$ be the set of admissible controls, defined as follows

$$\mathcal{A} = \left\{ \mathcal{U} = (q_m, \mathbf{p}_m)_{0 \leq m \leq M} \mid (q_m, \mathbf{p}_m) \in \mathcal{Z}(W_{m-}, W_{m+}, t_m) \right\} \quad (4.14)$$

For any t_m , we define the subset of controls applicable from t_m onwards as

$$\mathcal{U}_m = \left\{ (q_\ell, \mathbf{p}_\ell)_{m \leq \ell \leq M} \mid (q_\ell, \mathbf{p}_\ell) \in \mathcal{Z}(W_{\ell-}, W_{\ell+}, t_\ell) \right\}. \quad (4.15)$$

5 EW–CVaR portfolio formulation

We use CVaR to quantify downside risk in terminal wealth. Let $\psi(w)$ denote the probability density function of terminal wealth W_T . For a given confidence level $\alpha \in (0, 1)$ (typically 0.01 or 0.05), the CVaR of W_T at level α is defined by

$$\text{CVaR}_\alpha(W_T) = \frac{1}{\alpha} \int_{-\infty}^{\text{VaR}_\alpha(W_T)} w \psi(w) dw. \quad (5.1)$$

Here, $\text{VaR}_\alpha(W_T)$ is the Value-at-Risk (VaR) of W_T at confidence level α , given by⁵

$$\text{VaR}_\alpha(W_T) = \{w \mid \mathbb{P}[W_T \leq w] = \alpha\}. \quad (5.2)$$

Intuitively, $\text{VaR}_\alpha(W_T)$ is the wealth threshold that is not exceeded with probability α , while $\text{CVaR}_\alpha(W_T)$ represents the expected value of wealth in the worst α -quantile of the distribution [52].

For subsequent use, we denote by $\mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+}[W_T]$ the expectation of terminal wealth W_T under the real-world measure, conditional on the system being in state X_{0+} at time t_0^+ , and using the control $\mathcal{U}_0$ over $[t_0, T]$. Let $x_0 = X_{0-}$ be the initial state.

Following [53], the CVaR expression in (5.1) can be reformulated as a more computationally tractable optimization problem involving a candidate threshold W that partitions the distribution of W_T into its lower tail and the remainder. This leads to

$$\text{CVaR}_{\alpha, \mathcal{U}_0}^{x_0, t_0} = \sup_W \mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} \left[W + \frac{1}{\alpha} \min(W_T - W, 0) \mid X_{0-} = x_0 \right]. \quad (5.3)$$

The feasible range of W in (5.3) coincides with the set of all possible values for W_T .

As a measure of reward, we consider the expected total withdrawals. For any admissible control $\mathcal{U}_0 \in \mathcal{A}$, we define expected withdrawals (EW)

$$\text{EW}_{\mathcal{U}_0}^{x_0, t_0} = \mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} \left[\sum_{m=0}^M q_m \right] \quad (5.4)$$

where $q_m = q(W_{m-}, t_m)$ and W_{m-} is defined in (4.7) (i.e. after mortality credit distribution). Here, we assume the investor survives the full decumulation period, consistent with [4].

Remark 5.1 (Discounting and mortality weighting). *We remark that (5.4) applies no time discounting. With all quantities expressed in real (inflation-adjusted) dollars, this is equivalent to assuming a zero real discount rate. Over long horizons, short-term real rates have often been modest on average in advanced economies, so adopting a zero real discount rate provides a reasonable and conservative baseline. Although it is straightforward to introduce discounting, we adopt this assumption; see [19] for further discussion.*

In addition, we do not mortality-weight withdrawals by survival probabilities. Such weighting is natural when valuing cash flows averaged across a population (e.g. in annuity pricing), but our objective represents the strategy of an individual retiree: conditional on being alive, the retiree requires the full spending amount, not its survival-probability-weighted value. For example, if the probability of surviving to a later age were 50%, mortality-weighting would mechanically halve the corresponding cash flows (ignoring discounting), even though the retiree would still require the full amount if alive. Hence, a conservative approach is to condition on survival to a fixed horizon (age 95 in our numerical illustration), which is consistent with the Bengen spending-rule scenario that has proved popular with retirees [4]. Accordingly, we condition on survival to the horizon, consistent with the practitioner “plan-to-live, not to die” convention; see [43].

⁵That is, $\int_{-\infty}^{\text{VaR}_\alpha(W_T)} \psi(w) dw = \alpha$.

The goal is to simultaneously maximize expected withdrawals (EW) and the CVaR of terminal wealth—two objectives that are inherently in conflict. Solving this trade-off requires solving a bi-objective optimization problem. To obtain Pareto-optimal solutions, we apply a scalarization approach, which converts the bi-objective problem into a single-objective one by combining CVaR and EW using a scalarization parameter $\gamma > 0$. Formally, for a given $\gamma > 0$, we seek a control $\mathcal{U}_0$ that solves

$$\sup_{\mathcal{U}_0 \in \mathcal{A}} \left\{ \text{EW}_{\mathcal{U}_0}^{x_0, t_0} + \gamma \text{CVaR}_{\alpha, \mathcal{U}_0}^{x_0, t_0} \right\}. \quad (5.5)$$

Remark 5.2 (Pre-commitment and induced time-consistent interpretation). *The scalarized objective in (5.5) is evaluated at t_0 and thus defines a pre-commitment control: the retiree selects the policy at inception and then follows it over the horizon. Because CVaR_α is generally not time-consistent in multi-period settings, the EW–CVaR optimization is formally time-inconsistent and is often viewed as non-implementable, since at time $t > 0$ the retiree would typically have an incentive to deviate from the t_0 pre-commitment strategy. Nonetheless, we use the pre-commitment formulation to determine the joint optimizer W^* in (5.3). Fixing W^* thereafter yields the equivalent expected-value control problem*

$$\sup_{\mathcal{U}_0 \in \mathcal{A}} \mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} \left[\sum_{m=0}^M q_m + \frac{\gamma}{1-\alpha} \min(W_T - W^*, 0) \right],$$

which satisfies the Bellman principle and is therefore time-consistent. Accordingly, we interpret the computed control $\mathcal{U}_0^*(\gamma)$ as the associated induced time-consistent strategy determined by (α, γ, W^*) (see, e.g. [10, 18, 57]).

Using (5.3)–(5.4), we recast (5.5) as a control problem involving both system dynamics, mortality credit distributions, withdrawals, and rebalancing actions. For a fixed scalarization parameter $\gamma > 0$, the pre-commitment EW–CVaR problem, denoted $PCEC_{t_0}(\alpha, \gamma)$, is defined via the value function $V(x_0, t_0^-)$ as

$$(PCEC_{t_0}(\alpha, \gamma)) : V(x_0, t_0^-) = \sup_{\mathcal{U}_0 \in \mathcal{A}} \sup_W \left\{ \mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} \left[\sum_{m=0}^M q_m + \gamma \left(W + \frac{1}{\alpha} \min(W_T - W, 0) \right) + \epsilon W_T \middle| X_{0-} = x_0 \right] \right\} \quad (5.6)$$

$$\text{subject to: } \left\{ \begin{array}{ll} \text{(dynamics)} & \{X_t\} \text{ evolves via (4.2)} \\ \text{(tontine)} & W_{\ell-} \text{ given by (4.7),} \\ \text{(withdrawal)} & W_{\ell+} = W_{\ell-} - q_\ell \text{ as in (4.8),} \\ \text{(rebalancing)} & X_{\ell+} \text{ as defined in (4.10), using } \mathbf{p}_\ell(\cdot), \\ & (q_\ell(\cdot), \mathbf{p}_\ell(\cdot)) \in \mathcal{Z}(W_{\ell-}, W_{\ell+}, t_\ell), \quad \ell = 0, \dots, M, \end{array} \right. \quad \begin{array}{l} t \notin \mathcal{T}, \\ t_m \in \mathcal{T}. \end{array} \quad (5.7)$$

We denote by $\mathcal{U}_0^* = \mathcal{U}_0^*(\gamma)$ the control that attains the supremum in (5.6)–(5.7), i.e. the optimal policy for a given scalarization parameter γ .

Note that we have added the extra term $\mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} [\epsilon W_T]$ to (5.6). If we have a maximum withdrawal constraint, and if $W_t \gg W$ as $t \rightarrow T$, the controls become ill-posed. In this fortunate state for the investor, we can break investment policy ties either by setting $\epsilon < 0$, which will force investments in bonds, or by setting $\epsilon > 0$, which will force investments into stocks. Choosing $|\epsilon| \ll 1$ ensures that this term only has an effect if $W_t \gg W$ and $t \rightarrow T$. See [19] for more discussion of this.

We interchange the supremum operators in (5.6) to express the value function in a more computationally tractable form:

$$V(x_0, t_0^-) = \sup_W \sup_{\mathcal{U}_0 \in \mathcal{A}} \left\{ \mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} \left[\sum_{m=0}^M q_m + \gamma \left(W + \frac{1}{\alpha} \min(W_T - W, 0) \right) + \epsilon W_T \middle| X_{0-} = x_0 \right] \right\}. \quad (5.8)$$

The inner optimization yields a continuous function of W , and the optimal threshold $W^*(x_0)$ is defined as

$$W^*(x_0) = \arg \max_W \sup_{\mathcal{U}_0 \in \mathcal{A}} \left\{ \mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} \left[\sum_{m=0}^M q_m + \gamma \left(W + \frac{1}{\alpha} \min(W_T - W, 0) \right) + \epsilon W_T \middle| X_{0-} = x_0 \right] \right\}. \quad (5.9)$$

Finally, the scalarization parameter γ reflects the investor's level of risk aversion. For a given $\alpha \in [0, 1]$, the EW-CVaR efficient frontier is defined as the following set of points in $\mathbb{R}^2$:

$$\mathcal{S}(\alpha) : \left\{ \left(\text{CVaR}_{\alpha, \mathcal{U}_0^*}^{x_0, t_0} [W_T], \text{EW}_{\mathcal{U}_0^*}^{x_0, t_0} \right) : \gamma > 0 \right\}, \quad (5.10)$$

traced out by solving (5.6) for each $\gamma > 0$. In other words, for any fixed level of risk aversion γ , the corresponding point in (5.10) cannot be improved in the EW-CVaR sense by any other admissible strategy in $\mathcal{A}$.

6 Neural network formulation

Our numerical approach to solving the EW-CVaR stochastic control problem builds on the growing literature that uses NNs to approximate the optimal control directly, avoiding dynamic programming methods [6, 31, 50, 51, 61]. These methods have been termed “global-in-time” machine-learning approaches to stochastic control [26]. This contrasts with the stacked NN approach, in which a separate NN is used to approximate the control at each rebalancing step [24, 59]. More generally, these techniques are special cases of “policy function approximation” for optimal stochastic control [47].

We begin by formulating the NN optimization problem based on the stochastic control structure of the EW-CVaR formulation $PCEC_{t_0}(\alpha, \gamma)$ given in (5.6)–(5.7). To this end, for an arbitrary admissible control $\mathcal{U}_0 \in \mathcal{A}$ and an arbitrary threshold W , we define the objective

$$V(x_0, t_0^-; \mathcal{U}_0, W) = \mathbb{E}_{\mathcal{U}_0}^{X_{0+}, t_0^+} \left[\sum_{m=0}^M q_m + \gamma \left(W + \frac{1}{\alpha} \min(W_T - W, 0) \right) + \epsilon W_T \middle| X_{0-} = x_0 \right], \quad (6.1)$$

subject to the dynamics and constraints specified in (5.7).

Then, the (exact) value function $V(x_0, t_0^-)$ of the $PCEC_{t_0}(\alpha, \gamma)$ problem is given by

$$V(x_0, t_0^-) = \sup_W \sup_{\mathcal{U}_0 \in \mathcal{A}} V(x_0, t_0^-; \mathcal{U}_0, W). \quad (6.2)$$

6.1 Approximation of admissible control

The essence of our NN approach is to directly approximate an admissible control $\mathcal{U}_0$ in (6.2)—that is, a sequence of withdrawal and rebalancing decisions $(q_m(\cdot), \mathbf{p}_m(\cdot))$ for $m = 0, \dots, M$ —using feedforward, fully connected neural networks. More specifically, given NN parameters (weights and biases) $\boldsymbol{\theta}_q$ and $\boldsymbol{\theta}_p$, we denote by $\widehat{q}(W_{m-}, t_m, \boldsymbol{\theta}_q)$ and $\widehat{\mathbf{p}}(W_{m+}, t_m, \boldsymbol{\theta}_p)$ the NN-based approximations of the withdrawal control $q_m(\cdot)$ and the rebalancing control $\mathbf{p}_m(\cdot)$, respectively. Formally, we write

$$\begin{aligned} \widehat{q}_m(\cdot) &:= \widehat{q}(W_{m-}, t_m, \boldsymbol{\theta}_q) \simeq q_m(W_{m-}, t_m), \quad m = 0, 1, \dots, M, \\ \widehat{\mathbf{p}}_m(\cdot) &:= \widehat{\mathbf{p}}(W_{m+}, t_m, \boldsymbol{\theta}_p) \simeq \mathbf{p}_m(W_{m+}, t_m), \quad m = 0, 1, \dots, M, \\ \widehat{\mathcal{U}}_0 &:= \{(\widehat{q}_m(\cdot), \widehat{\mathbf{p}}_m(\cdot)) \mid m = 0, \dots, M\} \simeq \mathcal{U}_0. \end{aligned} \quad (6.3)$$

Here, $\widehat{\mathcal{U}}_0$ denotes the NN-based approximation of the admissible control $\mathcal{U}_0 \in \mathcal{A}$.

We now approximate $\mathcal{U}_0$ in the objective $V(x_0, t_0^-; \mathcal{U}_0, W)$, as defined in (6.1), by a NN-based counterpart $\widehat{\mathcal{U}}_0$. This yields the NN-approximated objective $V_{NN}(x_0, t_0^-; \widehat{\mathcal{U}}_0, W)$, defined as

$$V_{NN}(x_0, t_0^-; \widehat{\mathcal{U}}_0, W) = \mathbb{E}_{\widehat{\mathcal{U}}_0}^{X_{0+}, t_0^+} \left[\sum_{m=0}^M \widehat{q}_m + \gamma \left(W + \frac{1}{\alpha} \min(W_T - W, 0) \right) + \epsilon W_T \middle| X_{0-} = x_0 \right], \quad (6.4)$$

subject to NN-induced system evolution:

$$\text{subject to } \left\{ \begin{array}{ll} \text{(dynamics)} & \{X_t\} \text{ evolves via (4.2),} \\ \text{(tontine)} & W_{m-} \text{ computed via (4.7),} \\ \text{(withdrawal)} & W_{m+} = W_{m-} - \widehat{q}_m(\cdot) \text{ as in (4.8),} \\ \text{(rebalancing)} & X_{m+} \text{ computed via (4.10) using } \widehat{\mathbf{p}}_m(\cdot), \\ & (\widehat{q}_m(\cdot), \widehat{\mathbf{p}}_m(\cdot)) \in \mathcal{Z}(W_{m-}, W_{m+}, t_m), \\ & m = 0, \dots, M, \end{array} \right\} \begin{array}{l} t \notin \mathcal{T}, \\ \\ \\ t_m \in \mathcal{T}. \end{array} \quad (6.5)$$

Then, using (6.2), $V(x_0, t_0^-)$ is approximated by a value function induced by the NN-based control, denoted $V_{NN}(x_0, t_0^-)$:

$$V(x_0, t_0^-) \simeq V_{NN}(x_0, t_0^-) := \sup_W \sup_{\widehat{\mathcal{U}}_0 \in \mathcal{A}} V_{NN}(x_0, t_0^-; \widehat{\mathcal{U}}_0, W). \quad (6.6)$$

6.2 Network architecture for controls

To approximate the withdrawal and rebalancing controls $(q_m, \mathbf{p}_m)$ at each decision time t_m , we use two fully connected feedforward neural networks: one for withdrawals $\widehat{q}(\cdot)$, and one for rebalancing $\widehat{\mathbf{p}}(\cdot)$, parameterized by $\boldsymbol{\theta}_q \in \mathbb{R}^{\nu_q}$ and $\boldsymbol{\theta}_p \in \mathbb{R}^{\nu_p}$, respectively. These networks take inputs of the form $(W_{m\pm}, t_m)$, but at different stages of investor actions: (i) $\widehat{q}(\cdot)$ uses pre-withdrawal wealth after mortality credits have been applied, i.e. (W_{m-}, t_m) ; and (ii) $\widehat{\mathbf{p}}(\cdot)$ uses post-withdrawal wealth, i.e. (W_{m+}, t_m) .

To enforce the control constraints in (4.11) and (4.12) directly within the networks, we apply suitable output activation functions:

- Withdrawal control: Let $z \in \mathbb{R}$ denote the pre-activation output of the withdrawal network's final layer. We apply a sigmoid transformation scaled by a wealth-dependent range:

$$\widehat{q}(W_{m-}, t_m, \boldsymbol{\theta}_q) = q_{\min} + \max(\min(q_{\max}, W_{m-}) - q_{\min}, 0) \cdot \frac{1}{1 + e^{-z}}. \quad (6.7)$$

Since the sigmoid function $\frac{1}{1+e^{-z}}$ maps into $(0, 1)$, this ensures that $\widehat{q} \in \mathcal{Z}_q(W_{m-}, t_m)$ without introducing additional optimization constraints.

- Rebalancing control: Let $\widehat{\mathbf{p}}_m = (\widehat{p}_m^{(1)}, \dots, \widehat{p}_m^{(k+1)})$ denote the vector of rebalancing proportions at time t_m , where $k = 1$ for the domestic-only portfolio and $k = 3$ for the internationally diversified portfolio. These proportions are obtained by applying a $(k+1)$ -way softmax to the logits $(z_1, \dots, z_{k+1}) \in \mathbb{R}^{k+1}$ output by the network's final layer:

$$\widehat{p}_m^{(i)} = \frac{e^{z_i}}{\sum_{\ell=1}^{k+1} e^{z_\ell}}, \quad i = 1, \dots, k+1. \quad (6.8)$$

This guarantees $\widehat{p}_m^{(i)} \in [0, 1]$ and $\sum_{i=1}^{k+1} \widehat{p}_m^{(i)} = 1$. The first k components $\widehat{p}_m^{(1)}, \dots, \widehat{p}_m^{(k)}$ correspond to explicitly parameterized proportions, while $\widehat{p}_m^{(k+1)}$ (e.g. $\widehat{p}_m^{(2)}$ when $k = 1$ and $\widehat{p}_m^{(4)}$ when $k = 3$) represents the remaining proportion. This transformation enforces the admissibility condition $\mathcal{Z}_p(W_{m+}, t_m)$ automatically, without the need for post-processing or constraints.

Using these activation functions, we reparameterize the NN-based control $\widehat{\mathcal{U}}_0$ in terms of its underlying network weights and biases $(\boldsymbol{\theta}_q, \boldsymbol{\theta}_p)$. That is, we write $V_{NN}(x_0, t_0^-; \widehat{\mathcal{U}}_0, W)$ as $V_{NN}(x_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W)$ to emphasize that the control policies are fully determined by these parameters.

Then, the control optimization problem (6.6) becomes an unconstrained optimization over the network parameters $\boldsymbol{\theta}_q, \boldsymbol{\theta}_p$, and the CVaR threshold W :

$$\begin{aligned} V_{NN}(x_0, t_0^-) &= \sup_W \sup_{\boldsymbol{\theta}_q \in \mathbb{R}^{\nu_q}} \sup_{\boldsymbol{\theta}_p \in \mathbb{R}^{\nu_p}} V_{NN}(x_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W) \\ &= \sup_{(W, \boldsymbol{\theta}_q, \boldsymbol{\theta}_p) \in \mathbb{R}^{\nu_q + \nu_p + 1}} V_{NN}(x_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W). \end{aligned} \quad (6.9)$$

We denote by θ_q^* , θ_p^* , and W^* the optimal network parameters and threshold.

We emphasize that, while the original control problem is constrained via the admissible set $\mathcal{A}$ (see (4.14)), the reformulated NN objective (6.9) is unconstrained in terms of the training variables. This allows the use of standard gradient-based methods for optimization. In our numerical implementation (Sections 9 and 8), we use Adam stochastic gradient descent to train the networks and determine the optimal parameters.

6.3 Approximation of the NN-induced objective

To evaluate the NN-induced objective in (6.9), we approximate its expectation using a finite set of N independent sample paths, each representing a full set of exogenous drivers for the asset indices (generated by (4.2)). Mortality enters only through the pathwise tontine gain rates (see Remark 2.1). These sample paths are indexed by $n = 1, \dots, N$, and all path-dependent quantities carry the superscript “(n)”. For example, $X_t^{(n)}|_{t \in [0, T]}$ is the simulated state trajectory for the n -th path, and $W_T^{(n)}$ is the corresponding terminal wealth.

Specifically, given any set of NN parameters (θ_q, θ_p) and threshold W , the NN-induced objective $V_{NN}(x_0, t_0^-; \theta_q, \theta_p, W)$ is approximated by

$$V_{NN}(x_0, t_0^-; \theta_q, \theta_p, W) \approx \widehat{V}_{NN}(x_0, t_0^-, \theta_q, \theta_p, W) \quad (6.10)$$

where $\widehat{V}_{NN}(x_0, t_0^-, \theta_q, \theta_p, W) := \dots$

$$\dots := \frac{1}{N} \sum_{n=1}^N \left[\sum_{m=0}^M \widehat{q}_m(W_{m-}^{(n)}, t_m; \theta_q) + \gamma(W + \frac{1}{\alpha} \min(W_T^{(n)} - W, 0)) + \epsilon W_T^{(n)} \right], \quad (6.11)$$

$$\text{s.t.} \quad \left\{ \begin{array}{ll} \text{(dynamics)} & \{X_t^{(n)}\} \text{ drawn from the } n\text{-th sample path of index returns,} \\ & \text{(generated by (4.2))} \\ & t \notin \mathcal{T}, \\ \text{(tontine)} & W_{m-}^{(n)} \text{ computed via (4.7),} \\ & \text{using the pathwise } g_m^{(n)} \text{ given by (4.4) (see Remark 2.1),} \\ \text{(withdrawal)} & W_{m+}^{(n)} = W_{m-}^{(n)} - \widehat{q}(W_{m-}^{(n)}, t_m, \theta_q), \\ \text{(rebalancing)} & X_{m+}^{(n)} \text{ computed from (4.10), using } \widehat{p}(W_{m+}^{(n)}, t_m, \theta_p), \\ & (\widehat{q}_m(\cdot), \widehat{p}_m(\cdot)) \in \mathcal{Z}(W_{m-}^{(n)}, W_{m+}^{(n)}, t_m), \quad m = 0, \dots, M, \end{array} \right\} \quad t_m \in \mathcal{T}. \quad (6.12)$$

For subsequent benchmark comparisons, we generate sample paths using the generic transition dynamics defined in (4.2), which encapsulate both parametric SDE models and nonparametric bootstrapped trajectories. In our two-asset benchmarks, the dynamics follow a Kou-type jump–diffusion model (see Subsection 8.4). For higher-dimensional or empirically calibrated settings, we adopt the block bootstrap methodology described in Section 9.

As in Remark 2.1, each simulation path n carries a sequence of one-year conditional death probabilities $\{\delta_{m-1}^{(n)}\}_{m=1}^M$. In the deterministic life–table case these are path-independent, $\delta_{m-1}^{(n)} \equiv \delta_{m-1}$, whereas under stochastic LC/CBD mortality they are read off from the simulated surface $\{q_{x,t}^{(n)}\}$. In both cases we convert $\delta_{m-1}^{(n)}$ to tontine gain rates $g_m^{(n)} = \delta_{m-1}^{(n)} / (1 - \delta_{m-1}^{(n)})$ which enter the pathwise versions of (4.4)–(4.7) when computing $W_{m-}^{(n)}$ in (6.12). Throughout the EW–CVaR optimization we adopt the “plan to live” convention: during NN training the member is assumed to survive the full horizon, so these mortality inputs affect only the size of the mortality credits, not the termination time of the path.

Given the sample-based objective $\widehat{V}_{NN}(x_0, t_0^-; \theta_q, \theta_p, W)$ in (6.11)–(6.12), we train the networks by

solving the empirical maximization problem⁶

$$(\boldsymbol{\theta}_q^*, \boldsymbol{\theta}_p^*, W^*) := \arg \max_{\boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W} \widehat{V}_{NN}(x_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W) \quad (6.13)$$

with a gradient-based optimizer, such as Adam stochastic gradient descent. The resulting parameters define the learned control policies

$$\widehat{q}^*(\cdot) := \widehat{q}(\cdot; \boldsymbol{\theta}_q^*), \quad \widehat{p}^*(\cdot) := \widehat{p}(\cdot; \boldsymbol{\theta}_p^*),$$

which yields the NN-based optimal control

$$\widehat{\mathcal{U}}_0^* := \{ (\widehat{q}_m^*(\cdot), \widehat{p}_m^*(\cdot)) \mid m = 0, \dots, M \}.$$

We evaluate the performance of these trained control policies on out-of-sample or out-of-distribution test paths. Detailed numerical results are presented in Sections 9 and 8.

7 Pricing of the MBG

The optimal withdrawal and rebalancing controls $q^*(\cdot)$ and $p^*(\cdot)$ of the pre-commitment EW-CVaR problem $PCEC_{t_0}(\alpha, \gamma)$ maximize expected withdrawals conditional on survival—that is, the “plan to live, not to die” objective defined in (5.4), which assumes the retiree remains alive throughout the decumulation horizon.

By contrast, the MBG overlay is triggered by early death and guarantees that the member’s cumulative withdrawals, expressed in nominal dollars, are never less than the initial investment L_0 , as reflected in the payout formula (3.1). Because the MBG pays any shortfall to the estate only upon death—and never credits the retiree’s account—the payout does not affect the account balance or the decision process. Hence, the control pair $(q^*(\cdot), p^*(\cdot))$ remains optimal for the EW-CVaR objective even when the overlay is present.

Since the MBG-payout (3.1) is defined in nominal dollars at the payment time t_{m_τ} , the nominal amount must be converted to real units before valuation. Specifically, at time t_{m_τ} , the nominal shortfall is multiplied by the CPI ratio $\text{CPI}_0/\text{CPI}_{m_\tau}$ to obtain its real value at t_0 . This adjustment reflects the end-of-interval convention in Remark 3.1, under which the MBG-payout is implemented at time t_{m_τ} .

With this in mind, we define the random real dollar MBG payout

$$Z_g := \frac{\text{CPI}_0}{\text{CPI}_{m_\tau}} \max\left(L_0 - \sum_{\ell=0}^{m_\tau-1} q_\ell^*(\cdot) \frac{\text{CPI}_\ell}{\text{CPI}_0}, 0\right). \quad (7.1)$$

We note that, equivalently, in (7.1), the combination of the two CPI ratios— $\text{CPI}_\ell/\text{CPI}_0$ inside the sum and $\text{CPI}_0/\text{CPI}_{m_\tau}$ outside—acts to express all terms in real dollars at time t_0 .

7.1 Actuarial pricing formula

As discussed in Section 3.2, money-back protection can be funded through different mechanisms (e.g. internal pool funding versus external insurance). To report the economic cost of the MBG in a transparent and implementation-agnostic way, we summarize it using an equivalent up-front load factor $f_g \in (0, 1)$ applied to the initial contribution L_0 .

All quantities in the pricing rule below are expressed in real (inflation-adjusted) dollars at t_0 . We choose f_g so that the real value of the load equals the real-world expected MBG payout plus an explicit prudential buffer based on a tail risk measure (a standard risk-measure/premium-principle approach in actuarial pricing; see, e.g. [12, 63]).

Formally,

$$f_g L_0 = \mathbb{E}_{\mathcal{U}_0^*}^{x_0, t_0} [Z_g] + \lambda \text{CVaR}_{\alpha_g}^{x_0, t_0} [Z_g], \quad \lambda \geq 0, \quad \alpha_g \in (0, 1). \quad (7.2)$$

⁶Equivalently, we minimize the empirical loss function $\mathcal{L}(\boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W) := -\widehat{V}_{NN}(x_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W)$.

- $\mathbb{E}_{\mathcal{U}_0^*}^{x_0, t_0}[Z_g]$ is the real-world expectation of the MBG payout Z_g (in real t_0 dollars), evaluated under the wealth process induced by the optimal control $\mathcal{U}_0^* = (q^*(\cdot), \mathbf{p}^*(\cdot))$ of the problem $PCEC_{t_0}(\alpha, \gamma)$.
- $\text{CVaR}_{\alpha_g}^{x_0, t_0}[Z_g]$ is the CVaR of the same real dollar payout at tail level α_g , computed under the same induced wealth process. The tail level α_g is a prudential-buffer choice and need not coincide with the confidence level α used in the retiree's EW-CVaR objective.⁷
- λ is the prudential-buffer coefficient: $\lambda = 0$ corresponds to an actuarially fair (expected-cost) equivalent load, while $\lambda > 0$ adds an explicit buffer for adverse outcomes (e.g. to reflect internal risk limits or capital-adequacy considerations).

Given (λ, α_g) and the induced payout distribution of Z_g under the optimal policy $\mathcal{U}_0^*$, (7.2) defines the equivalent load f_g . In practice, the same economic cost could be implemented through different mechanisms, such as an implicit reduction in the starting income rate, pool-level self-insurance, or external insurance premia; our use of f_g provides a common metric for comparing such designs.

We now give explicit expressions for the two valuation functionals in (7.2), both expressed in real dollars at time t_0 :

$$\begin{aligned}\mathbb{E}_{\mathcal{U}_0^*}^{x_0, t_0}[Z_g] &= \mathbb{E}_{\mathcal{U}_0^*}^{X_{0+}, t_0^+}[Z_g \mid X_{0-} = x_0], \\ \text{CVaR}_{\alpha_g}^{x_0, t_0}[Z_g] &= \frac{1}{\alpha_g} \mathbb{E}_{\mathcal{U}_0^*}^{X_{0+}, t_0^+}[Z_g \mathbf{1}_{\{Z_g \geq \text{VaR}_{\alpha_g}(Z_g)\}} \mid X_{0-} = x_0],\end{aligned}\tag{7.3}$$

subject to system evolution induced by $\mathcal{U}_0^* = (q^*(\cdot), \mathbf{p}^*(\cdot))$:

$$\text{subject to } \left\{ \begin{array}{ll} \text{(dynamics)} & \{X_t\} \text{ evolves via (4.2),} \\ \text{(tontine)} & W_{\ell-} \text{ computed via (4.7),} \\ \text{(withdrawal)} & W_{\ell+} = W_{\ell-} - q_{\ell}^*(\cdot) \text{ as in (4.8),} \\ \text{(rebalancing)} & X_{\ell+} \text{ computed via (4.10) using } \mathbf{p}_{\ell}^*(\cdot), \end{array} \right\} \begin{array}{l} t \notin \mathcal{T}, 0 \leq t < t_{m_{\tau}} \\ \\ \\ \ell = 0, \dots, m_{\tau} - 1. \end{array} \tag{7.4}$$

Remark 7.1 (Interpreting f_g in starting-rate terms). *Equation (7.2) produces an equivalent up-front load f_g on the initial contribution L_0 (i.e. a one-time cost measure expressed as a proportion of L_0). For ease of interpretation in the numerical results, we also translate f_g into an implied reduction in a notional starting payment rate. Fix a notional starting payment rate β_0 (per dollar of contribution) for the same tontine design without the MBG. Under the reporting convention that the MBG cost is expressed as an equivalent up-front adjustment to the initial contribution (or benefit base), the implied post-load starting rate is $\beta_g = (1 - f_g)\beta_0$, i.e. a reduction of $f_g\beta_0$ (or $10^4 f_g\beta_0$ basis points) from the reference rate. We emphasize that β_0 is used only as a translation device for reporting and need not correspond to any specific operational funding mechanism.*

7.2 Simulation-based numerical methods

The MBG price is estimated from K independent sample paths, each representing a full set of exogenous drivers for index returns as given by (4.2), and, when relevant, mortality realizations. These sample paths are indexed by $k = 1, \dots, K$, and all path-dependent quantities carry the superscript (k) .

The pricing procedure uses the trained control networks $\hat{q}^*(\cdot)$ and $\hat{\mathbf{p}}^*(\cdot)$ obtained from Section 6, which approximate the optimal withdrawal and rebalancing strategies for the EW-CVaR problem. These controls are held fixed during the MBG valuation and applied across all sample paths.

For each path $k = 1, \dots, K$, the wealth evolution follows the dynamics induced by the control pair $\hat{\mathcal{U}}_0^* = (\hat{q}^*(\cdot), \hat{\mathbf{p}}^*(\cdot))$ obtained in Section 6. We write $\{X_t^{(k)}\}_{t \in [0, T]}$ for the resulting state trajectory on path

⁷The MBG payout is a low-frequency, high-severity liability with a heavy-tailed profile, so tail measures such as CVaR provide a natural basis for a prudential buffer.

k ; all withdrawals and allocations at the rebalancing times $t_m \in \mathcal{T}$ are the network outputs $\hat{q}^*(W_{m-}^{(k)}, t_m)$ and $\hat{p}^*(W_{m+}^{(k)}, t_m)$ evaluated with the path-specific inputs.

As summarized in Remark 2.1, each simulation path k is equipped with a sequence of one-year conditional death probabilities $\{\delta_{m-1}^{(k)}\}_{m=1}^M$. These probabilities are converted to tontine gain rates $g_m^{(k)} = \delta_{m-1}^{(k)} / (1 - \delta_{m-1}^{(k)})$ when computing mortality credits via (4.4)–(4.7). In addition, they also drive the simulation of death times and MBG payouts along each path as outlined in Algorithm 7.1 below. Together with the domestic CPI series $\{\text{CPI}_m\}_{m=0}^M$ (one value per decision time), these mortality inputs determine the timing and real dollar size of the pathwise MBG payouts.

The simulation-based pricing details are described in Algorithm 7.1. This algorithm computes Monte Carlo estimates $\widehat{\mathbb{E}}[Z_g]$ and $\widehat{\text{CVaR}}_{\alpha_g}(Z_g)$ of the theoretical valuation functionals $\mathbb{E}_{\mathcal{U}_0^*}^{x_0, t_0}[Z_g]$ and $\text{CVaR}_{\alpha_g}^{x_0, t_0}[Z_g]$ in (7.3). For each path $k = 1, \dots, K$, we denote by $Z_g^{(k)}$ the simulated real dollar MBG payout, obtained by applying (7.1) with the path-specific withdrawal history and death time.

Algorithm 7.1 Simulation-based MBG pricing under trained NN controls

- 1: initialize the index list $\mathcal{K} \leftarrow \{1, \dots, K\}$;
 - 2: initialise the list $\mathcal{Z} \leftarrow [0, \dots, 0]$ of length K ;
 - 3: initialize MBG-payouts $Z_g^{(k)} \leftarrow 0$, $k = 1, \dots, K$;
 - 4: initialize cumulative nominal withdrawals $\mathcal{W}^{(k)} \leftarrow 0$, $k = 1, \dots, K$;
 - 5: initialize $X_{0-}^{(k)} = x_0$ and compute $W_{0-}^{(k)}$ for all $k \in \mathcal{K}$;
 - 6: **for** $m = 0$ **to** $M - 1$ **do**
 - 7: **for each** $k \in \mathcal{K}$ **do**
 - 8: compute pathwise $g_m^{(k)}$ via (4.4) (see Remark 2.1) and $W_{m-}^{(k)}$ via (4.7); {mortality credit}
 - 9: $q_m^{(k)} \leftarrow \hat{q}^*(W_{m-}^{(k)}, t_m)$; {real dollars}
 - 10: $\mathcal{W}^{(k)} \leftarrow \mathcal{W}^{(k)} + q_m^{(k)} \frac{\text{CPI}_m^{(k)}}{\text{CPI}_0}$; {nominal withdrawal accumulation}
 - 11: $W_{m+}^{(k)} \leftarrow W_{m-}^{(k)} - q_m^{(k)}$;
 - 12: obtain $X_{m+}^{(k)}$ via (4.10) using $\hat{p}^*(W_{m+}^{(k)}, t_m)$; {rebalancing}
 - 13: compute $X_{(m+1)-}^{(k)}$, $k = 1, \dots, K$, via (4.2); {simulation over $[t_m^+, t_{m+1}^-]$ }
 - 14: **if** the member dies in $(t_m^+, t_{m+1}^-]$ **then**
 - 15: $Z_g^{(k)} \leftarrow \frac{\text{CPI}_0}{\text{CPI}_{m+1}^{(k)}} \max(L_0 - \mathcal{W}^{(k)}, 0)$;
 - 16: set $\mathcal{Z}^{(k)} \leftarrow Z_g^{(k)}$;
 - 17: remove path k from list $\mathcal{K}$;
 - 18: **end if**
 - 19: **end for**
 - 20: **if** $\mathcal{K} = \emptyset$ **then**
 - 21: break; {all deaths processed}
 - 22: **end if**
 - 23: **end for**
 - 24: compute sample mean: $\widehat{E} = \frac{1}{K} \sum_{k=1}^K Z_g^{(k)}$;
 - 25: sort list $\mathcal{Z}$ in descending order to obtain $Z_{(1)} \geq Z_{(2)} \geq \dots \geq Z_{(K)}$;
 - 26: compute empirical $\widehat{\text{CVaR}}_{\alpha_g}$: $\widehat{\text{CVaR}}_{\alpha_g} = \frac{1}{\lceil \alpha_g K \rceil} \sum_{j=1}^{\lceil \alpha_g K \rceil} Z_{(j)}$;
 - 27: return $(\widehat{E}, \widehat{\text{CVaR}}_{\alpha_g})$;
-

With $\widehat{\mathbb{E}}[Z_g]$ and $\widehat{\text{CVaR}}_{\alpha_g}(Z_g)$ computed by Algorithm 7.1, the actuarial load factor f_g defined by (7.2) is approximated by

$$\widehat{f}_g := \frac{\widehat{\mathbb{E}}[Z_g] + \lambda \widehat{\text{CVaR}}_{\alpha_g}(Z_g)}{L_0}. \quad (7.5)$$

In the numerical results, we report $\widehat{f}_g$ (and the associated up-front deduction $\widehat{f}_g L_0$).

Remark 7.2. In Algorithm 7.1 the test “member dies during $(t_m^+, t_{m+1}^-]$ ” (Line 14) is implemented as a Bernoulli experiment based on the pathwise probabilities $\delta_m^{(k)}$ introduced in Remark 2.1. Concretely, for every path k and every year m in which the member is still alive we draw $U \sim \text{Uniform}(0, 1)$ and declare death in $(t_m^+, t_{m+1}^-]$ when $U < \delta_m^{(k)}$. The corresponding index is then $m_\tau = m + 1$, so that, under the timing convention of Remark 3.1, the MBG payment is effected at t_{m_τ} . If no death is recorded up to $m = M - 1$, the member is deemed to have survived the entire horizon and $Z_g^{(k)} = 0$ on that path.

Remark 7.3 (Inflation treatment in MBG pricing). In the numerical experiments, wealth is simulated in real (inflation-adjusted) units using bootstrapped real asset returns, so CPI does not enter the wealth recursion. CPI is used only in Algorithm 7.1 to implement the money-back test in nominal dollars at the death time and to express the resulting MBG payout back in real t_0 dollars.

To obtain a stochastic (pathwise) inflation adjustment consistent with the bootstrap, the CPI index is treated as an additional series in the resampling procedure: monthly CPI changes are bootstrapped jointly with the asset-return blocks (preserving the dependence structure), a CPI index path is reconstructed along each simulated path, and the corresponding pathwise CPI ratios are used in the nominal–real conversion step of Algorithm 7.1. See $\text{CPI}_m^{(k)}$ on Line 10 and $\text{CPI}_{m+1}^{(k)}$ on Line 15 of Algorithm 7.1.

8 Benchmark validation

In this section we validate our numerical implementation by replicating, as closely as possible, the synthetic-market results of [21]. Their study solves an EW–CVaR optimal stochastic control problem for a decumulation portfolio with and without a tontine overlay, using dynamic programming and a PIDE solver, and reports the corresponding efficient frontiers and optimal controls.

We follow the modelling assumptions, data calibration, and base-case investment scenario in [21], which are described below, and assume mortality given by the 2014 Canadian Pensioner Mortality Table (CPM2014), treated as a deterministic period life table. In our notation this corresponds to a sequence of one-year conditional death probabilities $\{\delta_{m-1}\}_{m=1}^M$ taken directly from CPM2014, with $\delta_{m-1}^{(n)} \equiv \delta_{m-1}$ on every simulation path n , exactly as in the deterministic case discussed in Subsection 2.2.1 and Remark 2.1. Within this common framework, we compare the EW–CVaR efficient frontiers produced by our NN method with the reference frontiers reported in Figure 1 of [21].

In the remainder of the paper, we also consider a more general setting with an internationally diversified four-asset opportunity set and stochastic mortality, as developed in Subsection 2.2. However, the benchmark validation in this section is based solely on the original two domestic assets and deterministic mortality, in order to allow a direct comparison with [21].

8.1 Asset dynamics

For the benchmark we restrict attention to the two domestic assets considered in [21]: a broad domestic equity index and a constant-maturity domestic short-term government bond index. We follow the usual practitioner approach and directly model the returns of the constant-maturity bond index as a stochastic process (see, e.g. [32, 34]). Consistent with the stock index, we assume that the constant-maturity bond index also follows a jump–diffusion specification. Empirical justification for this modelling choice can be found in [19][Appendix A].

We denote by $\{S_t\}_{0 \leq t \leq T}$ the real value invested in the domestic stock index at time t , and $\{B_t\}_{0 \leq t \leq T}$

the real value invested in the domestic bond index. In the absence of investor actions (withdrawals or rebalancing), both assets follow correlated double-exponential jump-diffusion processes as in [21][Section 4].

Let ξ_s be a random variable representing the jump multiplier, such that a jump occurring at time t results in $S_t = \xi_s S_{t-}$. We assume $\ln(\xi_s)$ follows an asymmetric double-exponential distribution with probability density function (see [28, 29])

$$\varphi_s(y) = \zeta_s \eta_{s,1} e^{-\eta_{s,1}y} \mathbf{1}_{y \geq 0} + (1 - \zeta_s) \eta_{s,2} e^{\eta_{s,2}y} \mathbf{1}_{y < 0}, \quad \zeta_s \in [0, 1], \quad \eta_{s,1} > 1, \quad \eta_{s,2} > 0. \quad (8.1)$$

Between rebalancing times, and in the absence of active control, the stock index process $\{S_t\}$ evolves according to the jump-diffusion dynamics

$$\frac{dS_t}{S_{t-}} = (\mu_s - \lambda_s \kappa_s) dt + \sigma_s dZ_s(t) + d \left(\sum_{i=1}^{\pi_s(t)} (\xi_{s,i} - 1) \right), \quad t \in [t_{m-1}^+, t_m^-], \quad t_{m-1} \in \mathcal{T}. \quad (8.2)$$

Here, μ_s and σ_s are the (inflation-adjusted) drift and instantaneous volatility, respectively, and $\{Z_s(t)\}_{t \in [0, T]}$ is a standard Brownian motion. The process $\{\pi_s(t)\}_{0 \leq t \leq T}$ is a Poisson process with a constant finite intensity rate $\lambda_s > 0$. In (8.2), the jump amplitudes $\{\xi_{s,i}\}_{i=1}^\infty$ are independent and identically distributed (i.i.d.) random variables having the same distribution as ξ_s ; κ_s is the compensated drift adjustment given by

$$\kappa_s = \mathbb{E}[\xi_s - 1] = \frac{\zeta_s \eta_{s,1}}{\eta_{s,1} - 1} + \frac{(1 - \zeta_s) \eta_{s,2}}{\eta_{s,2} + 1} - 1. \quad (8.3)$$

The Brownian motion $\{Z_s(t)\}_{t \in [0, T]}$, the Poisson process $\{\pi_s(t)\}_{t \in [0, T]}$, and the jump multipliers $\{\xi_{s,i}\}$ are assumed to all be mutually independent.

Similarly, between rebalancing times, the bond index process $\{B_t\}$ evolves according to

$$\frac{dB_t}{B_{t-}} = (\mu_b - \lambda_b \kappa_b + \mu_b^c \mathbf{1}_{\{B_{t-} < 0\}}) dt + \sigma_b dZ_b(t) + d \left(\sum_{i=1}^{\pi_b(t)} (\xi_{b,i} - 1) \right), \quad t \in [t_{m-1}^+, t_m^-], \quad t_{m-1} \in \mathcal{T}, \quad (8.4)$$

where the parameters in (8.4) are defined similarly to those in (8.2). In particular, $\{\pi_b(t)\}_{0 \leq t \leq T}$ is a Poisson process with a positive, finite, constant jump intensity λ_b . The jump amplitudes $\{\xi_{b,i}\}_{i=1}^\infty$ are i.i.d. random variables, each distributed as ξ_b , where $\ln(\xi_b)$ follows an asymmetric double-exponential distribution with probability density function given by

$$\varphi_b(y) = \zeta_b \eta_{b,1} e^{-\eta_{b,1}y} \mathbf{1}_{y \geq 0} + (1 - \zeta_b) \eta_{b,2} e^{\eta_{b,2}y} \mathbf{1}_{y < 0}, \quad \zeta_b \in [0, 1], \quad \eta_{b,1} > 1, \quad \eta_{b,2} > 0. \quad (8.5)$$

In (8.4), the term $\mu_b^c \mathbf{1}_{\{B_{t-} < 0\}}$, the parameter $\mu_b^c \geq 0$ represents the borrowing spread applied when the bond amount becomes negative. The processes $\{Z_b(t)\}_{t \in [0, T]}$, $\{\pi_b(t)\}_{0 \leq t \leq T}$, and the jump multipliers $\{\xi_{b,i}\}$ are assumed to all be mutually independent.

The diffusion components are correlated, $dZ_s(t) dZ_b(t) = \rho_{sb} dt$, while all jump processes and jump sizes are mutually independent; namely, the Poisson processes $\{\pi_s(t)\}$, $\{\pi_b(t)\}$ and the corresponding jump multipliers $\{\xi_{s,i}\}$, $\{\xi_{b,i}\}$ are mutually independent, and independent of $\{Z_s\}$ and $\{Z_b\}$.

8.2 Synthetic-market calibration

Following [21], the synthetic market is calibrated to monthly real total-return series for the CRSP value-weighted equity index and the CRSP US 30-day T-bill index over the period 1926:1–2020:12, with both series deflated by the US CPI. In [21], the double-exponential jump-diffusion parameters $(\mu_s, \sigma_s, \lambda_s, \zeta_s, \eta_{s,1}, \eta_{s,2})$ for the stock index and $(\mu_b, \sigma_b, \lambda_b, \zeta_b, \eta_{b,1}, \eta_{b,2})$ for the bond index, together with the diffusion correlation ρ_{sb} , are estimated using the threshold-based jump filter described therein. In our validation we do not re-estimate these parameters. Instead, we adopt the annualized real parameter values reported in Table 1 of [21], which we reproduce in Table 8.1.

TABLE 8.1: *Annualized real parameter values for the double-exponential jump-diffusion model, taken from [21]. The stock asset is the CRSP value-weighted total return index; the bond asset is the 30-day US T-bill index, both deflated by CPI.*

Stock (CRSP)	μ_s	σ_s	λ_s	ζ_s	$\eta_{s,1}$	$\eta_{s,2}$	ρ_{sb}
	0.08912	0.1460	0.3263	0.2258	4.3625	5.5335	0.08420
30-day T-bill	μ_b	σ_b	λ_b	ζ_b	$\eta_{b,1}$	$\eta_{b,2}$	ρ_{sb}
	0.00460	0.0130	0.5053	0.3958	65.801	57.793	0.08420

8.3 Retirement scenario

The base-case retirement scenario, which underlies the synthetic-market efficient frontiers used for validation, mirrors the specification in Section 11 and Table 3 of [21]. All monetary quantities are expressed in thousands of real dollars. A concise summary is given in Table 8.2.

TABLE 8.2: *Base-case retirement scenario used for validation, matching the specification in [21]. Monetary units are thousands of real dollars.*

Item	Value
Retiree	65-year-old Canadian male
Mortality table	CPM2014 (deterministic life table)
Investment horizon T	30 years
Equity index	Real CRSP capitalization-weighted total return index
Bond index	Real US 30-day T-bill index
Initial wealth W_0	1000
Rebalancing / withdrawal times	$t = 0, 1, \dots, 29$ (annual)
Minimum annual withdrawal $q_{\min}$	40
Maximum annual withdrawal $q_{\max}$	80
Equity fraction range	$p_m \in [0, 1]$
Borrowing spread μ_b^c	0.02
Tontine gain rate g_m	as in Equation (2.11)
Group gain Γ_m	1.0 (as discussed in Subsection 2.1.3)
Tontine fee ϱ	$1 - e^{-0.005} \approx 0.00499$ (49.9 bps/year) ⁸
Risk tail level α	0.05
Stabilization parameter ε	-10^{-4}
Market parameters	as in Table 8.1

In this setting, the investor controls annual withdrawals $q_t \in [q_{\min}, q_{\max}]$ and the equity fraction $p_t \in [0, 1]$. Wealth can become negative if required withdrawals exceed available funds, in which case the portfolio is liquidated and subsequent withdrawals are financed as debt growing at the bond rate plus the borrowing spread μ_b^c . Mortality enters only through the deterministic life-table probabilities $\{\delta_{m-1}\}$ from CPM2014, so that all synthetic-market paths share the same mortality profile (Remark 2.1).

The risk-reward trade-off is quantified using expected total withdrawals (EW) and the expected shortfall (ES) at the 5% level of terminal wealth W_T , as in [21].

⁸Forsyth et al. [21] model the tontine fee as a continuously charged rate $T_f = 0.5\%$ per annum, which enters the wealth recursion through a factor $\exp(-T_f \Delta t)$ each year; see, for example, Equations (5.3)–(5.6) in [21]. In this paper, to align with industry practice, we instead deduct a single proportional yearly fee ϱ at each decision time. We choose ϱ such that $(1 - \varrho) = e^{-T_f}$ (with $\Delta t = 1$), i.e. $\varrho = 1 - e^{-0.005}$, so that the effective annual charge matches T_f .

8.4 Validation results

We validate our NN implementation by reproducing the synthetic-market EW-CVaR efficient frontiers reported in [21], both with and without a tontine overlay. Figure 8.1 compares the reference frontiers (computed using the PDE-based control method in [21]) with the NN frontiers obtained from our training procedure, together with the constant-withdrawal/constant-allocation benchmark. Overall, the NN approach yields EW-CVaR efficient frontiers in close agreement with the PDE-based reference and preserves its qualitative structure.

As shown in Figure 8.1, the NN frontiers closely track the PDE-based reference curves across the range of scalarization parameters γ used to trace the trade-off. Consistent with the figure, unreported summary errors confirm close agreement in both the annualized expected withdrawal $E[\sum_m q_m]/T$ and the corresponding $\text{CVaR}_{0.05}$ values along the frontier. The quality of the approximation is similar in the no-tontine case.

Beyond pointwise agreement, the NN frontiers preserve the expected qualitative behavior in γ : as risk aversion increases, the annualized expected withdrawal decreases and $\text{CVaR}_{0.05}$ moves in the corresponding direction. In particular, the region of the frontier most relevant in practice is nearly indistinguishable from the PDE-based reference in Figure 8.1.




FIGURE 8.1: *Validation of the EW-CVaR efficient frontiers in the synthetic market. Reference frontiers with and without a tontine overlay, labelled respectively as “No Tontine (PDE Control)” and “Tontine (PDE Control)”, are compared with the corresponding NN frontiers, together with the constant withdrawal/constant allocation benchmark. Units: thousands of dollars.*

9 Internationally diversified portfolios

9.1 Asset return data and preprocessing

For the internationally diversified experiments, we take the perspective of a domestic (Australian) investor and work with four indices: Australian equities and government bonds (domestic), and U.S. equities and government bonds (foreign) converted into AUD. All returns are expressed in real (inflation-adjusted) AUD at a monthly frequency.

Our empirical sample runs from 1935:1 to 2022:12. Where necessary, individual series are backcast and truncated so that all four indices and the Australian CPI are jointly available over a common horizon. This long historical panel is used both to describe the “historical market” and as input to the bootstrap simulations described below and in Section 6.

Domestic equity. As a proxy for the Australian equity market, we construct a capitalization-weighted total return series for the All Ordinaries index. Monthly price levels are taken from the Bloomberg All Ordinaries Price Index (AS30) and combined with trailing dividend yield data from [35] and the Reserve Bank of Australia (RBA) to impute reinvested dividends prior to the start of the Bloomberg All Ordinaries Accumulation Index (XAOA) in 1999. From 1999 onward, the constructed series is aligned to XAOA, so that the resulting index is a consistent long-horizon Australian equity total return series.

Domestic bonds. The domestic bond asset is represented by a 10-year Commonwealth Government bond total return index constructed from two sources. From December 1935 to December 2011 we use the annual Australian government bond total return series of [27], which is based on micro-level Commonwealth bond data targeting a 10-year maturity [9]. We convert this annual total return index to monthly frequency by interpolating the level of the total return index and then computing month-on-month returns. From May 2011 onward we splice in the S&P/ASX Australian Government Bond 7–10 Year Total Return Index (Bloomberg), available at a monthly frequency, to obtain a single long-horizon nominal total return bond index for the Australian market.

U.S. equity and bonds. For the U.S. market, we start from the same CRSP nominal total return series as in Subsection 8.2, namely the value-weighted equity index and the 30-day Treasury bill (T-bill) index, both in USD. These nominal USD returns are first converted into AUD using the end-of-month AUD/USD exchange rate, constructed from Federal Reserve Banking and Monetary Statistics archives (January 1935–December 1968) and the RBA historical exchange rate series thereafter, spliced to form a continuous monthly AUD/USD series, and are then placed on the same real AUD basis as the domestic assets via CPI deflation (see below).

Inflation and real returns. Inflation is measured by the Australian Consumer Price Index published by the RBA. The CPI series is obtained entirely from the RBA: from September 2017 onward it is available monthly, while before that date it is published quarterly. We treat the quarterly CPI as the benchmark series and linearly interpolate it to monthly frequency to match the financial data. For each asset we form a nominal total return index, convert U.S. series to AUD where appropriate, and then obtain real returns by deflating with the interpolated Australian CPI. In this way all four assets are expressed in real AUD terms, consistent with the objective of funding real retirement spending for an Australian investor.

Historical market and bootstrap sampling. Future index paths in the internationally diversified experiments are generated nonparametrically using the stationary block bootstrap [13, 41, 45, 46]. We apply geometrically distributed block lengths and sample blocks of four-asset returns jointly to preserve both cross-sectional correlations and serial dependence. The expected block length is chosen following the data-driven procedure of [41] applied to each series, and then harmonised across assets to obtain a common effective block size of the order of a few years. All bootstrap sampling is carried out on the preprocessed monthly real AUD return series described above.

Return characteristics and dependence structure. Tables 9.1 and 9.2 summarize the preprocessed monthly real AUD returns across all four assets. As indicated in Tables 9.1–9.2, U.S. equity has

TABLE 9.1: *Summary statistics of monthly real returns (1935:1–2022:12): annualized mean, annualized geometric mean, and annualized volatility. $\text{VaR}_{0.05}$ (m) and $\text{CVaR}_{0.05}$ (m) report the 5% Value-at-Risk and CVaR of monthly returns.*

Asset	Mean (ann.)	Geo. mean (ann.)	Vol (ann.)	$\text{VaR}_{0.05}$ (m)	$\text{CVaR}_{0.05}$ (m)
U.S. 30-day T-bill	0.003	-0.001	0.095	-0.041	-0.056
U.S. equity index	0.082	0.071	0.165	-0.064	-0.096
AU 10-year bond	0.012	0.012	0.036	-0.015	-0.025
AU equity index	0.069	0.059	0.151	-0.065	-0.100

the highest historical average real return in this sample and is only weakly correlated with Australian

TABLE 9.2: *Correlation matrix of monthly real returns (1935:1–2022:12).*

	U.S. 30-day T-bill	U.S. equity	AU 10-year bond	AU equity
U.S. 30-day T-bill	1.00	0.34	0.17	-0.22
U.S. equity	0.34	1.00	0.10	0.33
AU 10-year bond	0.17	0.10	1.00	0.11
AU equity	-0.22	0.33	0.11	1.00

equity (correlation ≈ 0.33 , well below 1)⁹. This comparatively low cross-country equity correlation suggests meaningful scope for international diversification, even among developed equity markets, a point we revisit later. By contrast, once expressed in real AUD, the U.S. 30-day T-bill has a near-zero average real return and substantial volatility inherited from exchange-rate fluctuations, making short-maturity foreign fixed income a less attractive defensive asset than domestic government bonds. These empirical patterns are revisited when interpreting the optimal controls below.

9.2 Mortality assumptions

In the internationally diversified experiments, we model the lifetime of a 65-year-old Australian male and consider both deterministic and stochastic mortality, consistent with the general framework in Subsection 2.2 and Remark 2.1.

For the deterministic specification, we use the most recent available single-year Australian male period life table (`mltper_1x1.txt`) from the Human Mortality Database (HMD, [25]), specifically the 1x1 Australian male life table covering calendar years 1921–2021.¹⁰ This table provides annual one-year death probabilities $q_{x,t}$ over age and calendar year. Fixing the retirement age x_0 and retirement date t_0 , and taking annual decision times $t_m = t_0 + m$, the corresponding conditional death probabilities $\{\delta_{m-1}\}_{m=1}^M$ are defined exactly as in equation (2.12). In a homogeneous pool we set $\delta_{m-1}^j \equiv \delta_{m-1}$ for all members j , and the associated tontine gain rates g_m entering the wealth recursion follow from equation (2.11).

For stochastic mortality, we work with the same HMD source but use the raw deaths and exposures required for model calibration. Specifically, we extract annual Australian male deaths $D_{x,t}$ and central exposures-to-risk $E_{x,t}^c$ for ages 55–95 and calendar years 1987–2021 from the HMD tables `Deaths_1x1.txt` and `Exposures_1x1.txt`.¹¹ When modelling one-year death probabilities $q_{x,t}$ (rather than central death rates), we follow [62] and approximate initial exposures by $E_{x,t}^0 \approx E_{x,t}^c + 0.5 D_{x,t}$, and then construct the corresponding empirical one-year death probabilities as $q_{x,t} := D_{x,t}/E_{x,t}^0$. These historical mortality series are used to calibrate a single-population model from the generalized age-period-cohort (GAPC) family (such as LC or CBD models), following Subsection 2.2. Calibration and forecasting are implemented in R using the `StMoMo` package [62], which casts these specifications in the GAPC framework. Model-specific identifiability constraints are those standard in `StMoMo`; for brevity we do not reproduce them here, and instead refer the reader to [62] for details.

On each simulated path k , the associated sequence of one-year conditional death probabilities $\{\delta_{m-1}^{(k)}\}_{m=1}^M$ is defined as in (2.20), and the corresponding pathwise tontine gain rates $g_m^{(k)}$ are obtained via (2.11). As summarized in Remark 2.1, these sequences $\{\delta_{m-1}^{(k)}\}$ (and hence $\{g_m^{(k)}\}$) are used through-

⁹Over 1990:1–2022:12, the correlation is 0.48 (vs 0.33 in the full sample), indicating higher dependence in the recent period.

¹⁰Available at https://www.mortality.org/File/GetDocument/hmd.v6/AUS/STATS/mltper_1x1.txt.

¹¹Available at https://www.mortality.org/File/GetDocument/hmd.v6/AUS/STATS/Deaths_1x1.txt and https://www.mortality.org/File/GetDocument/hmd.v6/AUS/STATS/Exposures_1x1.txt.

out our numerical experiments: they drive mortality credits in the NN training and efficient-frontier computation, and determine death times and MBG payouts in the pricing overlay (also see Remark 7.2).

9.3 Experimental setup and scenarios

For the internationally diversified experiments, we retain the same decumulation framework used in the benchmark validation (Table 8.2), but now from the perspective of a 65-year-old Australian male retiree investing in the four indices described in Subsection 9.1. All monetary quantities are interpreted as thousands of real Australian dollars, and the mortality assumptions follow Subsection 9.2.

More specifically, the retiree starts at time $t_0 = 0$ with initial wealth $W_0 = 1000$, corresponding to an initial purchase price $L_0 = 1000$ for the tontine product with the embedded MBG overlay described in Section 3. The planning horizon is $T = 30$ years with annual decision times $t_m = m$, $m = 0, \dots, M$, where $M = 30$. The minimum and maximum annual withdrawals are fixed at $q_{\min} = 40$ and $q_{\max} = 80$. In addition, we also enforce the no-shorting and no-leverage constraints as in Section 4. In these experiments, following QSuper’s Lifetime Pension PDS [49], we use an annual tontine fee of $\varrho = 0.11\%$ (11 bps).

Risk and reward are measured by the same EW-CVaR criterion as in Section 5: expected total real withdrawals over $[0, T]$ and CVaR $_{\alpha}$ of terminal wealth W_T at level $\alpha = 0.05$. For a given scalarization parameter $\gamma > 0$, we first train neural networks to approximate the optimal control $\mathcal{U}_0^*(\gamma)$ and compute the corresponding point on the EW-CVaR efficient frontier under either deterministic table mortality or stochastic GAPC-based mortality. The MBG overlay is then priced ex post, holding $\mathcal{U}_0^*(\gamma)$ fixed, using the actuarial pricing rule (7.2) and the simulation procedure in Section 7.

9.4 EW-CVaR efficient frontiers (deterministic mortality)

We now discuss the EW-CVaR efficient frontiers for both cases: two-asset (domestic-only) and four-asset (internationally diversified).

As in [21], we include a constant-withdrawal, constant-allocation benchmark as a reference point on the EW-CVaR frontiers. In line with the 4% rule of [4], the benchmark always withdraws 40 (thousand real dollars) per year and maintains fixed portfolio weights over time. For each case, we select the benchmark by grid search over asset weights in 10% increments, choosing the allocation that minimizes the ES (equivalently, maximizes CVaR $_{0.05}$) of terminal wealth.

In the two-asset case, the search is over equity fractions $p_s^d \in \{0, 0.1, \dots, 1\}$ with $p_b^d = 1 - p_s^d$; the best constant-weight benchmark has a 50% domestic equity allocation and $\text{ES} \approx -617.1$. In the four-asset case, we search over all weight vectors on the 10% grid whose components sum to one; the best benchmark has allocation $(p_s^d, p_b^d, p_s^f, p_b^f) = (0.1, 0.3, 0.2, 0.4)$ with $\text{ES} \approx -446.1$. These two constant-strategy benchmarks appear as single points in the efficient-frontier figures for the two-asset and four-asset experiments, respectively.

Figure 9.1 shows the EW-CVaR efficient frontiers obtained by NNs with and without a tontine overlay in the two-asset (domestic-only) and four-asset (internationally diversified) cases, together with the corresponding constant-weight benchmarks.

We make the following key observations about Figure 9.1:

- The internationally diversified (four-asset) constant-weight benchmark point lies essentially on the two-asset no-tontine frontier. Thus, at the 4%-rule withdrawal level, a simple buy-and-hold internationally diversified allocation can achieve an EW-CVaR trade-off comparable to that delivered by the domestic-only EW-CVaR optimal strategy without a tontine overlay. This highlights the strength of international diversification even before adding longevity pooling.
- In both the domestic-only and internationally diversified settings, the efficient frontiers obtained from our optimization lie well above and to the right of their corresponding constant-weight benchmark points. This confirms that “4%-rule-style” strategies with fixed withdrawals and static allocations are




FIGURE 9.1: *EW-CVaR efficient frontiers for domestic-only (two-asset) and internationally diversified (four-asset) portfolios, with and without a tontine overlay. Constant-weight benchmarks are shown as single points. Units: thousands of real AUD.*

substantially less efficient than the optimal dynamic policies obtained from our EW-CVaR optimization. For example, even at $\text{CVaR}_{0.05}(W_T) \approx 500$ (thousand real AUD) at age 95, a very conservative level of terminal-wealth downside protection, the four-asset tontine frontier supports expected annual withdrawals of about 70 (thousand real AUD), i.e. roughly 7% of initial wealth, compared with the fixed 40 (thousand, 4%) withdrawal implied by the 4% rule benchmark.

- Comparing the two “no tontine” curves, the four-asset frontier is generally shifted upwards and to the right relative to the two-asset frontier over the practically relevant range of $\text{CVaR}_{0.05}$. Allowing the retiree to invest in both Australian and U.S. assets therefore supports higher annualized withdrawals for a similar, or even improved, level of downside risk compared with a purely domestic portfolio.

This is consistent with Tables 9.1–9.2: in this sample, U.S. equity has a higher average real return than Australian equity, and the cross-country equity correlation is well below one, so adding foreign equity provides potential to improve the EW-CVaR trade-off even when used selectively rather than through uniformly higher equity exposure.

- For both asset settings, the tontine frontier dominates the corresponding no-tontine frontier, with especially pronounced gains in the four-asset case: over a wide range of scalarization parameters γ , the four-asset tontine strategies deliver both higher expected annual withdrawals and a much less adverse 5% tail of W_T than the corresponding four-asset no-tontine strategies, illustrating the substantial combined benefit of international diversification and longevity pooling under optimal dynamic controls.

In the next subsection, we examine how these frontiers shift when mortality is modelled stochastically.

9.5 Effect of stochastic mortality on EW-CVaR frontiers

We now examine how introducing stochastic mortality affects the EW-CVaR efficient frontier in the four-asset tontine setting. Throughout this subsection we keep the assets, fee structure, control constraints, and scalarization setup identical to Subsection 9.4; only the mortality specification changes. Figure 9.2 compares the EW-CVaR efficient frontiers for the four-asset tontine portfolio under three mortality specifications: (i) deterministic table mortality as in Subsection 9.2, (ii) stochastic LC mortality, and (iii) stochastic CBD mortality. In all cases we retain the “plan to live” convention: the retiree is assumed to survive to the end of the horizon, so mortality only affects the size of mortality credits through the one-year death probabilities (cf. Remark 2.1), not the termination time of a path.

A key point for interpretation is that the horizon corresponds to very advanced age (age 95 in the baseline calibration with a 30-year horizon). Consequently, extremely large tail targets for terminal wealth—for example, $\text{CVaR}_{0.05}(W_T) \approx 1000$ (thousand real AUD)—are arguably overly conservative, since they imply substantial residual wealth at age 95. In many practical settings, values around $\text{CVaR}_{0.05}(W_T) \approx 500$ already represent a sizeable buffer that could fund minimum spending well beyond




FIGURE 9.2: *EW-CVaR efficient frontiers for the four-asset tontine portfolio under deterministic table mortality and stochastic mortality (LC and CBD models). Units: thousands of real AUD.*

the model horizon. Therefore, the most decision-relevant part of Figure 9.2 is the region *to the left of* (or up to) roughly $\text{CVaR}_{0.05}(W_T) = 500$, where the controls correspond to higher withdrawals rather than preserving large terminal balances.

Over this economically relevant region, the impact of stochastic mortality is visible but modest: the stochastic frontiers track the deterministic table frontier closely, with only a relatively small leftward shift. This shift is consistent with systematic longevity improvement in the stochastic models: when death probabilities fall relative to the baseline table, survivors receive smaller mortality credits, which lowers terminal wealth in the adverse outcomes that drive $\text{CVaR}_{0.05}(W_T)$. While systematic longevity risk does erode the benefits of the tontine overlay to some extent, the four-asset tontine portfolio under stochastic mortality still offers an attractive combination of high expected withdrawals and improved $\text{CVaR}_{0.05}(W_T)$ relative to the domestic-only and constant-weight strategies discussed in Subsection 9.4 (Figure 9.1).

A further observation is that the LC and CBD frontiers are close to each other, with the CBD curve typically lying slightly further to the left. Conditional on modelling systematic longevity risk, the specific choice between LC and CBD therefore has only a relatively minor effect on the EW-CVaR trade-off.

9.6 Optimal control heatmaps


To illustrate the structure of the learned policies, Figure 9.3 plots the optimal rebalancing controls for the four-asset tontine with stochastic mortality (LC model) and a representative scalarization parameter $\gamma = 1.5$. Each panel shows, as a function of time and real wealth, the fraction of wealth invested in one of the four indices described in Subsection 9.1, together with the 5th, 50th, and 95th percentiles of the wealth distribution under the optimal policy. Control heatmaps under deterministic table mortality and the CBD mortality model are qualitatively very similar and are therefore omitted for brevity.

9.6.1 General comments about Figure 9.3


We make several qualitative observations from Figure 9.3:

- Along the median and upper wealth trajectories, the strategy exhibits a clear tilt towards domestic assets: a substantial fraction of wealth is invested in Australian securities, with the Australian 10-year government bond (Figure 9.3 (c)) acting as the main stabilising asset. Specifically, for moderate to high wealth levels, the allocation to the domestic bond is large (typically well above 50%), while the U.S. 30-day T-bill (Figure 9.3 (d)) is used only marginally. This aligns with Table 9.1: the domestic bond index has substantially lower volatility and markedly milder 5% monthly tail losses (VaR/CVaR) than the U.S. T-bill once returns are expressed in real AUD, supporting its role as the primary defensive allocation in the optimal policy.


Along the 95th-percentile path the equity allocations (Figures 9.3 (a) and (b)) are gradually reduced,




(a) Australian equity index (domestic)



(b) U.S. equity index



(c) Australian 10-year government bond index (domestic)



(d) U.S. 30-day T-bill index

FIGURE 9.3: Optimal rebalancing controls under the four-asset tontine with stochastic mortality (LC model) for a representative scalarization parameter $\gamma = 1.5$. Colours show the fraction of wealth invested in each asset as a function of time and real wealth. Units: thousands of real AUD.

so that the portfolio becomes dominated by the Australian bond index, effectively locking in favourable outcomes; any remaining equity risk is taken primarily via the domestic market.

- The U.S. equity allocation (Figure 9.3 (b)) is highly state dependent. When wealth falls into the lower region of the state space, particularly along and below the 5th-percentile path, the policy tilts aggressively toward U.S. equities, with weights close to 100% in that asset and very low exposure to domestic bonds. This behaviour is concentrated in the extreme low-wealth tail and can be interpreted as a low-probability, last-resort catch-up position, rather than a typical allocation across the state space. When wealth is low, the optimal EW-CVaR policy favours taking additional growth risk in order to improve the expected withdrawal profile, using foreign equity as the primary “catch-up” instrument.

Tables 9.1–9.2 provide a simple empirical rationale: U.S. equity combines the highest historical average real return in the sample with only moderate correlation to Australian equity, so it offers both growth upside and diversification potential when deployed as a state-dependent catch-up instrument in low-wealth regions. This is consistent with related optimal-decumulation evidence that moderate caps on the equity share have little effect on the efficient frontier, with the main control differences concentrated

in extreme low-wealth tail states [20].

- Comparing Figures 9.3 (a) and (b) shows that Australian and U.S. equities play different roles. At moderate wealth, the strategy holds a sizeable but balanced allocation to Australian equities, while U.S. equity exposure is more concentrated in low-wealth regions of the state space. This suggests that the optimal policy uses domestic equity as the core growth asset and foreign equity as a more opportunistic, high-return lever when the retiree is underfunded.

9.6.2 Role of equities in international diversification

To highlight how international diversification is used in the optimal policy, we compare the four-asset heatmaps in Figure 9.3 with the two-asset (domestic-only) counterpart in Figure 9.4.

In the domestic-only case (Figure 9.4), the equity fraction is primarily wealth-dependent rather than strongly time-dependent. Along the median trajectory the policy remains in a moderate equity region for most of the horizon (roughly in the 25%–45% range, corresponding to the teal/green colours), rather than shifting toward near-zero equity late in the horizon. By contrast, when wealth moves into the low-wealth region near and below the 5th-percentile path, the heatmap turns red and the policy allocates close to 100% to equities, using the only available growth asset to improve the expected withdrawal profile when wealth is small.




FIGURE 9.4: *Optimal fraction in domestic equities in the two-asset (domestic-only) tontine with stochastic mortality (LC model) for a representative point on the efficient frontier ($\gamma = 1.5$). Colours show the fraction of wealth invested in domestic equities as a function of time and real wealth. Units: thousands of real AUD.*

In the four-asset case (Figures 9.3 (a) and (b)), this role is split between Australian and U.S. equities. Domestic equity behaves much like the two-asset equity in relatively well-funded states, forming part of a growth core together with the Australian bond index. However, in low-wealth regions the allocation to U.S. equity becomes predominant, while domestic equity and bond exposures are reduced.

This comparison makes clear that the four-asset improvement in Figure 9.1 is driven by *how* equity risk is used, rather than by a uniformly higher equity allocation across all wealth levels. With only domestic assets, Australian equities must serve both as the core growth exposure and as the only available “catch-up” lever when wealth is low. Once U.S. equity is available, the optimal policy can separate these roles: in better-funded regions it maintains a more balanced, largely domestic mix, whereas in underfunded regions it shifts predominantly into U.S. equity to seek additional growth and support the expected withdrawal objective. In this sense, international diversification adds a second equity “catch-up” instrument that is activated mainly when the retiree is underfunded, instead of increasing equity exposure uniformly.

This interpretation is consistent with the return evidence in Tables 9.1–9.2, which shows that U.S. equity offers slightly higher average real returns than domestic equity and only moderate correlation

with Australian equity over the sample period.

9.7 MBG pricing results

We now report simulation-based MBG pricing results using the actuarial pricing formula (7.2) and the Monte Carlo procedure in Algorithm 7.1. In all cases, the MBG is priced *ex post* under the fixed EW-CVaR optimal control as discussed in the previous subsections. We report Monte Carlo estimates of $\widehat{\mathbb{E}}[Z_g]$ and $\widehat{\text{CVaR}}_{\alpha_g}(Z_g)$, and the implied equivalent up-front load factor $\widehat{f}_g$ computed as in (7.5).

For interpretability, $\widehat{f}_g$ can be read as an equivalent one-time deduction $\widehat{f}_g L_0$ from the initial purchase price (equivalently, as a reduction in a notional starting payment rate); see Remark 7.1.

9.7.1 Base-case MBG load

Table 9.3 reports MBG pricing results under a base-case parameter set: scalarization parameter $\gamma = 1.5$ (as used in the heatmap illustrations), $L_0 = 1000$ (thousand real AUD), pricing confidence level $\alpha_g = 5\%$, and prudential-buffer coefficient $\lambda = 0.5$. We apply this same parameter set to four settings—two-asset (domestic-only) versus four-asset (internationally diversified), each under deterministic table mortality and stochastic LC mortality. We also report the implied post-load notional starting payment rate $(1 - \widehat{f}_g)\beta_0$ for the reference value $\beta_0 = 0.05$.


Setting	$\widehat{\mathbb{E}}[Z_g]$	$\widehat{\text{CVaR}}_{0.05}(Z_g)$	$\widehat{f}_g$	$(1 - \widehat{f}_g)\beta_0$
2 assets, table mortality	65.63	736.61	0.43	2.85%
2 assets, LC mortality	66.67	755.47	0.44	2.80%
4 assets, table mortality	69.32	736.82	0.44	2.80%
4 assets, LC mortality	70.69	758.28	0.45	2.75%

TABLE 9.3: *Base-case parameter set for MBG pricing under EW-CVaR optimal controls: $\gamma = 1.5$, $L_0 = 1000$ (thousand real AUD), $\alpha_g = 5\%$, $\lambda = 0.5$. The load $\widehat{f}_g$ is computed from the Monte Carlo estimates via (7.5). The last column reports the illustrative post-load notional starting payment rate $(1 - \widehat{f}_g)\beta_0$ for the reference value $\beta_0 = 0.05$ (see Remark 7.1).*


Two features stand out. First, the actuarially fair expected-cost component $\widehat{\mathbb{E}}[Z_g]$ is modest (about 6%–7% of L_0). Thus, under actuarially fair pricing ($\lambda = 0$), the implied equivalent load would be only $\widehat{f}_g \approx \widehat{\mathbb{E}}[Z_g]/L_0 \approx 5\%$ –7%, which translates, for $\beta_0 = 5\%$, into an illustrative starting-rate reduction of roughly 25–35 bps. Second, the tail measure $\widehat{\text{CVaR}}_{0.05}(Z_g)$ is large (about 0.74–0.76 of L_0 , roughly an order of magnitude larger than $\widehat{\mathbb{E}}[Z_g]$), so adding a prudential buffer via $\lambda = 0.5$ yields a substantially larger equivalent load, $\widehat{f}_g \approx 0.43$ –0.45, corresponding to a post-load notional starting payment rate of about 2.75%–2.85%.

To visualize the distributional features behind this result, Figure 9.5 plots the empirical distribution of the MBG payout Z_g for the four-asset (internationally diversified) tontine with stochastic mortality (LC) and $\gamma = 1.5$; the pricing tail level is $\alpha_g = 5\%$. Figure 9.5 (a) shows the unconditional density of Z_g , which exhibits a pronounced spike near zero (many scenarios in which the guarantee is out of the money) and a long right tail. Figure 9.5 (b) shows the density conditional on $Z_g > 0$, highlighting that when the MBG is triggered the payout can still be large and widely spread. In Figure 9.5 (a), the pricing threshold is $\text{VaR}_{0.05}(Z_g) \approx 575$ (thousand real AUD) and the corresponding tail mean is $\text{CVaR}_{0.05}(Z_g) \approx 758$, consistent with Table 9.3.

For this representative policy and buffer choice, the differences between (i) domestic-only versus internationally diversified assets and (ii) deterministic table versus stochastic LC mortality are small in magnitude relative to the overall load: the resulting $\widehat{f}_g$ values are all close to 0.43–0.45. In other words, the dominant driver of the equivalent load is the tail term in (7.5), with only minor variation across the



(a) Unconditional density of Z_g



(b) Density conditional on $Z_g > 0$

FIGURE 9.5: Empirical densities of the MBG payout Z_g (in real dollars at t_0) under a representative base-case EW-CVaR optimal control: the four-asset (internationally diversified) tontine with stochastic mortality (LC), $\gamma = 1.5$, and $\alpha_g = 5\%$. Units: thousands of real AUD.

asset and mortality specifications considered here.

9.7.2 Sensitivity to risk loading and tail confidence level

We next study the sensitivity of the MBG load to (i) the prudential-buffer parameters (λ, α_g) and (ii) the retiree's risk-reward preference in the EW-CVaR optimization, indexed by the scalarization parameter γ . Table 9.4 presents the results for the four-asset (internationally diversified) tontine with stochastic mortality (LC model) across $\gamma \in \{1.5, 0.5, 0.2\}$, with $L_0 = 1000$ fixed. For each γ , we report the Monte Carlo estimates $\widehat{\mathbb{E}}[Z_g]$ and $\widehat{\text{CVaR}}_{\alpha_g}(Z_g)$ together with the implied load estimates $\widehat{f}_g$ computed from (7.5) for $\lambda \in \{0, 0.5, 1\}$. For fixed (γ, α_g) , $\widehat{f}_g$ depends linearly on λ , with slope $\widehat{\text{CVaR}}_{\alpha_g}(Z_g)/L_0$.

TABLE 9.4: Sensitivity of the Monte Carlo MBG load estimate $\widehat{f}_g$ in the four-asset tontine under stochastic mortality (LC), for different retiree risk-reward preferences indexed by γ . Parameters: $L_0 = 1000$ (thousand real AUD), $\beta_0 = 0.05$ (translation device; see Remark 7.1). Part (a): $\alpha_g = 5\%$. Part (b): $\alpha_g = 1\%$. $\widehat{f}_g$ is computed via (7.5).

(a) $\alpha_g = 5\%$					
γ	$\widehat{\mathbb{E}}[Z_g]$	$\widehat{\text{CVaR}}_{0.05}(Z_g)$	$\widehat{f}_g(\lambda = 0)$	$\widehat{f}_g(\lambda = 0.5)$	$\widehat{f}_g(\lambda = 1)$
1.5	70.69	758.28	0.07	0.45	0.83
0.5	58.12	729.37	0.06	0.42	0.79
0.2	47.50	664.38	0.05	0.38	0.71

(b) $\alpha_g = 1\%$					
γ	$\widehat{\mathbb{E}}[Z_g]$	$\widehat{\text{CVaR}}_{0.01}(Z_g)$	$\widehat{f}_g(\lambda = 0)$	$\widehat{f}_g(\lambda = 0.5)$	$\widehat{f}_g(\lambda = 1)$
1.5	70.69	917.55	0.07	0.53	0.99
0.5	58.12	915.19	0.06	0.52	0.97
0.2	47.50	871.95	0.05	0.48	0.92

Table 9.4 shows that the dominant drivers of the equivalent load are the prudential-buffer parameters (λ, α_g) . In particular, $\lambda = 0$ corresponds to actuarially fair (expected-cost) pricing and yields a modest load $\widehat{f}_g \approx 5\text{--}7\%$ across the reported γ values (about 25–35 bps off a notional $\beta_0 = 5\%$ starting rate). Introducing a buffer via $\lambda > 0$ increases the load sharply because $\widehat{\text{CVaR}}_{\alpha_g}(Z_g)$ is large relative to $\widehat{\mathbb{E}}[Z_g]$: for example, at $\alpha_g = 5\%$ and $\lambda = 0.5$ the load rises to $\widehat{f}_g \approx 0.38\text{--}0.45$. Tightening the tail confidence level from $\alpha_g = 5\%$ to $\alpha_g = 1\%$ further increases $\widehat{\text{CVaR}}_{\alpha_g}(Z_g)$ and therefore raises $\widehat{f}_g$ (e.g., to about 0.48–0.53 when $\lambda = 0.5$).

Varying γ affects the load in a secondary, policy-induced way: changing γ moves the EW-CVaR optimal policy along the efficient frontier and thereby shifts the distribution of Z_g (through withdrawal timing and the resulting account balance at death). Over the ranges considered here, this effect is modest compared with the direct impact of (λ, α_g) in (7.5).

Remark 9.1. *The parameters (λ, α_g) in (7.2)–(7.5) summarize an additional prudential buffer applied to the MBG tail (e.g. reflecting capital requirements, governance, or other conservatism beyond expected cost). These parameters are typically not directly observable from public disclosures, especially when the guarantee is embedded in the quoted payment schedule rather than shown as an explicit fee component. Consequently, the sensitivity results in Table 9.4 should be interpreted as a pricing stress test: they quantify how the implied load $\hat{f}_g$ changes across plausible buffer choices, holding the underlying EW-CVaR policy and the simulated payout distribution fixed.*

10 Conclusion and future work

This paper studies optimal retirement decumulation in an individual tontine account with a MBG overlay in a setting that allows both international diversification and systematic longevity risk at the pool level through stochastic mortality-credit inputs. The retiree’s withdrawal and asset-allocation decisions are determined by an EW-CVaR optimal control problem solved under a plan-to-live convention, and the MBG is priced ex post under the induced optimal policy using a simulation-based actuarial load that incorporates an explicit tail-risk prudential buffer. To compute state-dependent controls in this multi-asset, constrained setting, we develop a NN approach that scales beyond classical grid-based dynamic programming.

Numerically, international diversification improves the achievable EW-CVaR trade-off even before longevity pooling, and the largest gains arise when diversification is combined with a tontine overlay. Incorporating stochastic mortality shifts the efficient frontiers in the expected direction (systematic longevity improvement reduces mortality credits) but does not alter the qualitative structure of the optimal controls. The MBG results highlight a distinct economic mechanism: when expressed in equivalent-load terms, implied MBG loads are driven primarily by prudential buffers that place weight on adverse tail outcomes, rather than by actuarially fair expected payouts.

Several extensions are natural. First, the plan-to-live convention can be relaxed by modelling death as a stopping time and embedding bequest or estate preferences directly in the retiree’s objective. Second, MBG design and valuation could be made more fully endogenous by jointly choosing the guarantee terms and the associated funding/buffer rule, and by incorporating regulatory constraints. Third, the economic environment can be enriched by incorporating additional asset classes and hedging instruments, including alternative equity benchmarks (e.g. equal-weighted indices) and currency-hedged exposures. Finally, stress-testing performance under return and mortality misspecification would help assess the practical robustness of the learned controls.

References

- [1] A. Anarkulova, S. Cederburg, M. S. O’Doherty, and R. W. Sias. The safe withdrawal rate: Evidence from a broad sample of developed markets. *Journal of Pension Economics and Finance*, 24:464–500, 2025.
- [2] Sudipto Banerjee. Decoding retiree spending. <https://www.troweprice.com/institutional/us/en/insights/articles/2021/q1/decoding-retiree-spending-na.html>, 2021. T. Rowe Price Insights on Retirement.
- [3] Maximilian Bär and Nadine Gatzert. Products and strategies for the decumulation of wealth during retirement: insights from the literature. *North American Actuarial Journal*, 27(2):322–340, 2023.
- [4] W. Bengen. Determining withdrawal rates using historical data. *Journal of Financial Planning*, 7:171–180, 1994.

- [5] Thomas Bernhardt and Catherine Donnelly. Modern tontine with bequest: innovation in pooled annuity products. *Insurance: Mathematics and Economics*, 86:168–188, 2019.
- [6] Hans Buehler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- [7] Andrew JG Cairns, David Blake, and Kevin Dowd. A two-factor model for stochastic mortality with parameter uncertainty: theory and calibration. *Journal of Risk and Insurance*, 73(4):687–718, 2006.
- [8] M. Chen, M. Shirazi, P. A. Forsyth, and Y. Li. Machine learning and hamilton–jacobi–bellman equation for optimal decumulation: A comparison study. *Journal of Computational Finance*, 29(1):77–118, 2025.
- [9] Xiaoting Chen, Sherifa Elsherbiny, Ricardo Duque Gabriel, Òscar Jordà, Chi Hyun Kim, Moritz Schularick, and Alan M Taylor. Documentation on the jst database update 2016–2020, 6th release, july 2022, 2022.
- [10] Duy-Minh Dang and Chang Chen. Multi-period mean-buffered probability of exceedance in Defined Contribution portfolio optimization. *SIAM Journal on Financial Mathematics*, 2026. to appear.
- [11] Giorgio De Santis and Bruno Gerard. International asset pricing and portfolio diversification with time-varying risk. *The Journal of Finance*, 52(5):1881–1912, 1997.
- [12] Michel Denuit, Jan Dhaene, Marc Goovaerts, and Rob Kaas. *Actuarial Theory for Dependent Risks: Measures, Orders and Models*. John Wiley & Sons, 2005.
- [13] Hubert Dichtl, Wolfgang Drobetz, and Martin Wambach. Testing rebalancing strategies for stock-bond portfolios across different asset allocations. *Applied Economics*, 48(9):772–788, 2016.
- [14] David C. M. Dickson, Mary R. Hardy, and Howard R. Waters. *Actuarial Mathematics for Life Contingent Risks*. Cambridge University Press, 3 edition, 2020.
- [15] Catherine Donnelly. Actuarial fairness and solidarity in pooled annuity funds. *ASTIN Bulletin: The Journal of the IAA*, 45(1):49–74, 2015.
- [16] Catherine Donnelly, Montserrat Guillén, and Jens Perch Nielsen. Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, 56:14–27, 2014.
- [17] P. A. Forsyth, K. R. Vetzal, and G. Westmacott. Optimal asset allocation for a DC pension decumulation with a variable spending rule. *ASTIN Bulletin*, 50(2):419–447, 2020.
- [18] P.A. Forsyth. Multiperiod mean conditional value at risk asset allocation: Is it advantageous to be time consistent? *SIAM Journal on Financial Mathematics*, 11(2):358–384, 2020.
- [19] Peter A Forsyth. A stochastic control approach to defined contribution plan decumulation: “The Nastiest, Hardest Problem in Finance”. *North American Actuarial Journal*, 26(2):227–251, 2022.
- [20] Peter A. Forsyth and George Labahn. Numerical methods for optimal decumulation of a defined contribution pension plan. *Working paper*, 2026.
- [21] Peter A Forsyth, Kenneth R Vetzal, and Graham Westmacott. Optimal performance of a tontine overlay subject to withdrawal constraints. *ASTIN Bulletin: The Journal of the IAA*, 54(1):94–128, 2024.
- [22] Robert K. Fullmer. Tontines: A practitioner’s guide to mortality-pooled investments, 2019. Accessed: January 2026.
- [23] Jonathan T. Guyton and William J. Klinger. Decision rules and maximum initial withdrawal rates. *Journal of Financial Planning*, 19(3):48–58, 2006.
- [24] Jiequn Han et al. Deep learning approximation for stochastic control problems. *arXiv preprint arXiv:1611.07422*, 2016.
- [25] HMD. Human mortality database. <https://www.mortality.org/>. Max Planck Institute for Demographic Research (Germany), University of California, Berkeley (USA), and French Institute for Demographic Studies (France). Data downloaded on [date], DOI: 10.4054/HMD.Countries.[version code].
- [26] Ruimeng Hu and Mathieu Laurière. Recent developments in machine learning methods for stochastic control and games. *Numerical Algebra, Control and Optimization*, 14(3):435–525, 2024.
- [27] Òscar Jordà, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. The rate of return on everything, 1870–2015. *The quarterly journal of economics*, 134(3):1225–1298, 2019.
- [28] S.G. Kou. A jump diffusion model for option pricing. *Management Science*, 48:1086–1101, August 2002.
- [29] S.G. Kou and H. Wang. Option pricing under a double exponential jump diffusion model. *Management Science*, 50(9):1178–1192, September 2004.
- [30] Ronald D Lee and Lawrence R Carter. Modeling and forecasting us mortality. *Journal of the American statistical association*, 87(419):659–671, 1992.

- [31] Y. Li and P.A. Forsyth. A data-driven neural network approach to optimal asset allocation for target-based defined contribution pension plans. *Insurance: Mathematics and Economics*, 86:189–204, 2019.
- [32] Yijia Lin, Richard D MacMinn, and Ruilin Tian. De-risking defined benefit plans. *Insurance: Mathematics and Economics*, 63:52–65, 2015.
- [33] B.J. MacDonald, B. Jones, R.J. Morrison, R.L. Brown, and M. Hardy. Research and reality: A literature review on drawing down retirement savings. *North American Actuarial Journal*, 17(3):181–215, 2013.
- [34] Richard MacMinn, Patrick Brockett, Jennifer Wang, Yijia Lin, and Ruilin Tian. The securitization of longevity risk and its implications for retirement security. *Recreating sustainable retirement*, pages 134–160, 2014.
- [35] Thomas Mathews. A history of australian equities. Research Discussion Paper RDP 2019-04, Reserve Bank of Australia, 2019.
- [36] Moshe A. Milevsky and Thomas S. Salisbury. Financial valuation of guaranteed minimum withdrawal benefits. *Insurance: Mathematics and Economics*, 38(1):21–38, 2006.
- [37] Moshe A Milevsky and Thomas S Salisbury. Optimal retirement income tontines. *Insurance: Mathematics and economics*, 64:91–105, 2015.
- [38] MyNorth. MyNorth[®] Super and Pension, Product Disclosure Statement – Part A. <https://www.northonline.com.au/adviser/products/mynorth-lifetime/mynorth-lifetime-super>, October 2024. Issue date: 9 October 2024.
- [39] C. Ni, Y. Li, P.A. Forsyth, and R. Carroll. Optimal asset allocation for outperforming a stochastic benchmark target. *Quantitative Finance*, 22(9):1595–1626, 2022.
- [40] OECD. Pension markets in focus, 2019. Accessed: January 2026.
- [41] Andrew Patton, Dimitris N Politis, and Halbert White. Correction to “automatic block-length selection for the dependent bootstrap” by d. politis and h. white. *Econometric Reviews*, 28(4):372–375, 2009.
- [42] K. Peijnenburg, T. Nijman, and B.J.M. Werker. The annuity puzzle remains a puzzle. *Journal of Economic Dynamics and Control*, 70:18–35, 2016.
- [43] Wade D Pfau. An overview of retirement income planning. *Journal of Financial Counseling and Planning*, 29(1):114–120, 2018.
- [44] Shaun Pfeiffer, John Salter, and Harold Evensky. Increasing the sustainable withdrawal rate using the standby reverse mortgage. *Journal of Financial Planning*, 26(12):55–62, 2013.
- [45] Dimitris N Politis and Joseph P Romano. The stationary bootstrap. *Journal of the American Statistical association*, 89(428):1303–1313, 1994.
- [46] Dimitris N Politis and Halbert White. Automatic block-length selection for the dependent bootstrap. *Econometric reviews*, 23(1):53–70, 2004.
- [47] W Powell. A universal framework for sequential decision problems. *OR/MS Today February*. <https://tinyurl.com/PowellORMSfeature>, 2023.
- [48] QSuper. Product update: Lifetime pension. <https://qsuper.qld.gov.au/-/media/pdfs/qsuper-public/publications/ltp-sen-product-update-june-25.pdf>, July 2025. Accessed: January 2026.
- [49] QSuper. QSuper Product Disclosure Statement for Income Account and Lifetime Pension. <https://qsuper.qld.gov.au/pds>, July 2025. Issue date: 1 July 2025.
- [50] A Max Reppen, H Mete Soner, and Valentin Tissot-Daguette. Deep stochastic optimization in finance. *Digital Finance*, 5(1):91–111, 2023.
- [51] Anders Max Reppen and Halil Mete Soner. Deep empirical risk minimization in finance: Looking into the future. *Mathematical Finance*, 33(1):116–145, 2023.
- [52] R.T. Rockafellar and J.O. Royset. Random variables, monotone relations, and convex analysis. *Mathematical Programming*, 148:297–331, 2014.
- [53] R.T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.
- [54] Michael J Sabin and Jonathan Barry Forman. The analytics of a single-period tontine. *Available at SSRN 2874160*, 2016.
- [55] Hersh M Shefrin and Richard H Thaler. The behavioral life-cycle hypothesis. *Economic inquiry*, 26(4):609–643, 1988.
- [56] Bruno H Solnik. Why not diversify internationally rather than domestically? *Financial analysts journal*, 30(4):48–54, 1974.
- [57] M.S. Strub, D. Li, X. Cui, and J. Gao. Discrete-time mean-cvar portfolio selection and time-consistency induced term structure of the cvar. *Journal of Economic Dynamics and Control*, 108:103751, 2019.

- [58] Thinking Ahead Institute. Global pension assets study 2023. <https://www.thinkingaheadinstitute.org/research-papers/global-pension-assets-study-2023/>, 2023. Accessed: January 2026.
- [59] Ka Ho Tsang and Hoi Ying Wong. Deep-learning solution to portfolio selection with serially dependent returns. *SIAM Journal on Financial Mathematics*, 11(2):593–619, 2020.
- [60] United Nations, Department of Economic and Social Affairs, Population Division. World population prospects 2024, data sources. UN DESA/POP/2024/DC/NO. 11, 2024.
- [61] Pieter M van Staden, Peter A Forsyth, and Yuying Li. A global-in-time neural network approach to dynamic portfolio optimization. *Applied Mathematical Finance*, 31(3):131–163, 2024.
- [62] Andrés M Villegas, Vladimir K Kaishev, and Pietro Millossovich. Stmomo: An r package for stochastic mortality modeling. *Journal of Statistical Software*, 84:1–38, 2018.
- [63] Shaun S. Wang. A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67(1):15–36, 2000.