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

German Nova Orozco *　　Duy-Minh Dang†　　Peter A. Forsyth‡

August 24, 2026

Abstract

Money-back guarantees (MBGs) address bequest concerns in pooled retirement income products by returning the initial purchase price through withdrawals or, after early death, through a benefit to the member's beneficiaries or estate. We study the distinct actuarial problem created by adding an MBG to an individual-account tontine with dynamic withdrawals, investment in domestic and foreign assets, and systematic longevity risk. The retiree trades expected withdrawals (EW) against the lower-tail Conditional Value-at-Risk (CVaR) of terminal wealth under a fixed-horizon plan-to-live convention. Neural networks parameterize admissible withdrawal and rebalancing controls; the MBG is then valued ex post under the learned policy through an equivalent up-front load combining the expected payout with an upper-tail CVaR prudential buffer. We also approximate the effect of contract pooling on per-contract tail risk and pricing.

Using long-horizon market and mortality data calibrated for an Australian retiree, we find that expected MBG payouts are modest, while the representative-contract payout has a severe upper tail. Contract pooling substantially reduces the upper-tail exposure on a per-contract basis and lowers the price per contract. International diversification materially improves the EW–CVaR retirement-income trade-off, while affecting the MBG payout distribution and equivalent load only modestly. Stochastic mortality likewise has a modest effect on the efficient frontier and MBG pricing.

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 [37, 55]. 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 and uncertain market returns [3, 5, 32]. Rapid global population ageing further intensifies these challenges [57].

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. [24]). These rules are transparent and

*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`

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. [18, 20]).

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. [39]), and low demand for life annuities can be rational once bequest motives, illiquidity costs, and product design are taken into account [32].

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 surviving members [16, 17, 23, 34]. The appeal comes with a clear trade–off: mortality credits raise payouts conditional on survival, but wealth is typically forfeited at death.

More broadly, mean-field game and mean-field control formulations provide a population-level approach in which individual decisions interact with the evolving population distribution and may accommodate heterogeneity across members [8, 42]. By contrast, the homogeneous large-pool framework considered in [22] reduces the pooling mechanism to a representative-member problem. In that setting, a stochastic optimal control framework is developed for a pure individual-account tontine with deterministic mortality and a two-asset stock–bond investment universe, but without a death-benefit guarantee.

Within the representative-member framework of [22], 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 expected cumulative real (inflation-adjusted) withdrawals (EW) over a fixed 30-year retirement horizon, while risk is measured by the lower-tail Conditional Value-at-Risk (CVaR) of terminal real (inflation-adjusted) wealth. The formulation assumes that the retiree survives to the horizon, following the "plan to live, not to die" convention; see, e.g. [40]. The resulting control problem is solved numerically using a grid-based method for the associated partial integro-differential equation (PIDE).

The results in [22] 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, at comparable and economically relevant levels of lower-tail terminal-wealth 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" [22, 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 [22]. 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 [35, 46]).

Under such protection, if the member dies before cumulative retirement-income payments reach the protected purchase price, the resulting shortfall is paid to the member's beneficiaries or estate. This return-of-purchase-price structure differs from the tontine-with-bequest design in [5], which uses a separate bequest account. Although [22] notes that an MBG could be added as an overlay and priced separately, it does not undertake that valuation.

From the perspective of the entity that ultimately bears the guarantee, whether an insurer or the retirement-income pool, the MBG is a contingent, right-skewed death-benefit liability: it is often zero, but can be large when triggered. Accordingly, the economic cost of an MBG depends not only on expected payouts, but also on tail risk and any prudential loading or reserving rule applied to adverse outcomes. Crucially, the distribution of guarantee payouts is endogenous to the decumulation policy: the withdrawal and asset-allocation controls jointly determine the account and withdrawal paths, while cumulative nominal withdrawals determine the guarantee shortfall at death. For an issuer holding a portfolio of broadly homogeneous MBG contracts, contract-specific death-timing risk may diversify across the book, so the price per contract can differ from the stand-alone representative-contract valuation.

This creates an actuarial pricing problem absent from the pure-tontine analysis. The MBG must be valued under the decumulation policy it accompanies, even though the death benefit is paid outside the member's account and does not feed back into the withdrawal and rebalancing decisions 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 [13, 53]. International diversification is particularly relevant 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.

Expanding the asset universe increases the dimensionality of both the state and control spaces, making grid-based dynamic programming computationally prohibitive. Recent work in portfolio optimization and related control problems has shown that neural network (NN) parameterizations can approximate high-dimensional, state-dependent controls without requiring a state-space grid, thereby mitigating the dimensionality limitations of grid-based methods [10, 31, 36].

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 [22], 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 distinguishes idiosyncratic longevity risk, which is diversifiable within the pool, from systematic longevity risk, which is non-diversifiable.

Motivated by the above observations, this paper sets out to achieve three primary objectives. First, building on the pure individual-account tontine setting of [22], we add a return-of-purchase-price MBG overlay and treat the resulting death benefit as a contingent liability with a policy-dependent payout distribution. We value this liability ex post under that policy. Because the MBG may be funded through different operational mechanisms, we report its economic cost as an implementation-agnostic equivalent up-front load rather than specifying a particular fee schedule. The load combines the expected payout under the physical measure with an explicit upper-tail CVaR prudential buffer. Using the representative-contract payout distribution, we then develop a practical approximation to the effect of contract pooling on per-contract tail risk and pricing.

Second, we extend the representative-member EW–CVaR control framework of [22] from a two-asset setting with deterministic mortality to one with domestic and foreign assets and mortality credits determined under either deterministic or stochastic mortality. We compute the wealth- and time-dependent withdrawal and rebalancing controls using NN parameterizations rather than the grid-based PIDE method used in [22]. The optimization follows a fixed-horizon plan-to-live convention; the 30-year horizon extends from age 65 to age 95, providing a prudent stress test of the sustainability of real (inflation-adjusted) retirement spending.

Third, using nearly a century of market data together with Australian mortality data, we quantify how international diversification at the individual-account level and systematic mortality at the pool level affect EW–CVaR frontiers and MBG pricing. In particular, we measure the magnitude of the frontier improvement relative to the domestic-only portfolio, examine the state-dependent use of domestic and foreign assets, and determine how the resulting controls and mortality specifications affect MBG payout distributions and equivalent loads. We also assess the relative importance of expected guarantee costs and upper-tail prudential buffers. In addition, we quantify how the approximate per-contract tail risk and price vary with MBG contract-book size.

Using an Australian-investor calibration with Australian and U.S. equity and government-bond indices, the numerical results show that expected MBG payouts are modest, while the representative-contract payout has a severe upper tail. Under the pooling approximation, upper-tail exposure is substantially reduced on a per-contract basis, lowering the approximate price per contract. International diversification materially improves the EW–CVaR retirement-income trade-off, with the largest frontier gains when combined with the tontine overlay. Stochastic mortality has only a modest effect on the efficient frontier, while both international diversification and stochastic mortality have limited effects on representative-contract MBG pricing.

The remainder of the paper is organized as follows. Section 2 introduces the MBG overlay and its payoff and timing conventions, while Section 3 summarizes the representative-member tontine mechanics and deterministic and stochastic mortality inputs. 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. Section 7 develops representative-contract MBG pricing, and Section 8 introduces the contract-pooling approximation. Section 9 presents the market and mortality data and numerical results on efficient frontiers, MBG pricing, and contract pooling. Section 10 concludes and outlines directions for future research.

2 Money-back guarantee overlay

> "*Your purchase price is always paid back as either income to you or a death benefit paid to your beneficiaries.*"
>
> QSuper's Lifetime Pension Product Disclosure Statement [46]

In a pure individual-account tontine, account wealth is forfeited at death and redistributed within the pool. A money-back guarantee (MBG) changes the actuarial accounting: early death may generate a death-contingent benefit payable to the member's beneficiaries or estate. Thus, the problem is no longer only how mortality credits support withdrawals while the member is alive, but also how to value and express the cost of the death-contingent shortfall determined by the member's cumulative withdrawals before death and the timing of death. This guarantee overlay is the main product-design departure from the pure tontine framework of [22].

The MBG is motivated by recent Australian pooled retirement income products, where similar features are described as money-back protection; examples include QSuper's Lifetime Pension and MyNorth Lifetime [35, 46]. We model this feature as an overlay on a tontine retirement account, under which any shortfall between the purchase price and cumulative withdrawals at death is paid as a death benefit to the member's beneficiaries or estate.

2.1 Guarantee payoff and timing

We define the MBG payout on the set of equally spaced decision times in $[0, T]$, denoted by $\mathcal{T}$:

$$\mathcal{T} = \{t_m \mid t_m = m\Delta t,\ m = 0, \ldots, M\}, \qquad \Delta t = T/M, \tag{2.1}$$

where $t_0 = 0$ is the inception time and $T > 0$ is the finite investment horizon. Throughout, time is measured in years and we take annual decision times, so $\Delta t = 1$ and $T = M$. For notational convenience, define $t_m^- = t_m - \varepsilon$ and $t_m^+ = t_m + \varepsilon$, with $\varepsilon \to 0^+$, as the instants immediately before and after t_m, respectively.

At inception, suppose a member invests an amount L_0 as the purchase price of the retirement product with MBG. Throughout this paper, L_0 denotes a nominal (unadjusted for inflation) dollar amount fixed at inception. At each decision time t_m, $m = 0, \ldots, M-1$, while the member is alive, the regular tontine account mechanics apply: mortality credits are distributed, a real (inflation-adjusted) withdrawal q_m is made, and the portfolio is rebalanced. Once cumulative withdrawals, converted to nominal (unadjusted for inflation) terms, reach or exceed L_0, the MBG becomes inactive.

Let m_τ denote the interval index associated with death before the horizon, so that $\tau \in (t_{m_\tau-1}^+, t_{m_\tau}^-]$, for some $m_\tau \in \{1, \ldots, M\}$. If the member dies in this interval, no further withdrawal occurs; in particular, the scheduled withdrawal at t_{m_τ} is not made if such a withdrawal would otherwise be scheduled. If cumulative withdrawals made up to and including $t_{m_\tau-1}$, converted to nominal (unadjusted for inflation) terms, fall short of L_0, the MBG activates. If the member survives to the horizon $T = t_M$, we set $m_\tau = M + 1$ and the MBG payout is zero. The nominal (unadjusted for inflation) MBG payout is

$$\text{MBG-payout} = \begin{cases} \max\left(L_0 - \displaystyle\sum_{\ell=0}^{m_\tau-1} q_\ell \frac{\text{CPI}_\ell}{\text{CPI}_0}, 0\right), & m_\tau \leq M, \\ 0, & m_\tau = M+1. \end{cases} \tag{2.2}$$

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 (inflation-adjusted) withdrawal q_ℓ is first expressed in nominal terms by multiplying by $\text{CPI}_\ell/\text{CPI}_0$ before being summed in (2.2).[1] Under this accounting convention, all withdrawals made before death are included in the cumulative nominal withdrawals in (2.2).

The member's account balance at death is treated separately from the MBG calculation: any outstanding negative balance is not deducted from the MBG payout, while any positive balance is forfeited to the tontine pool and supports mortality credits for surviving members. This differs from a variable-annuity death-benefit guarantee, for which the deceased member's remaining investment account is available to fund the benefit. The MBG death benefit must therefore be financed separately from the deceased member's tontine account, whether through the pool, external insurance, or another funding arrangement.

Remark 2.1 (Timing convention)**.** *In the discrete-time model, if a member dies at time $\tau \in (t_{m_\tau-1}^+, t_{m_\tau}^-]$, the MBG payout is made at the end of that interval, i.e. at the next decision time t_{m_τ}, to the member's beneficiaries or estate. The scheduled withdrawal at t_{m_τ}, if any, is not made. This convention fixes only the death interval and the timing of the MBG payout; the treatment of the deceased member's account within the tontine pool is specified separately in Section 3.*

The MBG does not change the account dynamics or the controls chosen for the individual tontine account while the member is alive. It is a contingent overlay whose cost is computed under the withdrawal and rebalancing policy used for that account. This separation is important: the guarantee does not feed back into the member's withdrawal or rebalancing decisions in our formulation, but its payout distribution is endogenous to those decisions.

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. We do not impose the CAS in the baseline model in order to isolate the core MBG overlay and pricing mechanism. As a

[1] When $m_\tau \leq M$, this nominal shortfall is converted to a real (inflation-adjusted) amount at inception by multiplying by $\text{CPI}_0/\text{CPI}_{m_\tau}$ for valuation and reporting; see Section 7 (Eqn. (7.1)).

result, the recoverable amount under (2.2) may overstate the amount payable under a regulated product when CAS limits bind.

2.2 Equivalent MBG load

Product disclosures for pooled retirement income products typically emphasize the presence of money-back protection, but the funding mechanism need not appear as a stand-alone member-level charge. The cost may instead be reflected through pool-level self-insurance, an external insurance arrangement, or adjustments to income rates.

Accordingly, we summarize the economic cost of the MBG using an equivalent up-front load factor $f_g \geq 0$. The quantity $f_g L_0$ is interpreted as a one-time cost measure, or equivalently, a proportional reduction in a notional starting income rate or benefit base, that represents the MBG cost under the pricing rule adopted in this paper. This translation device reports the guarantee cost in a form that is comparable across operational implementations. Full details of the pricing methodology and numerical implementation are provided in Section 7.

3 Tontine mechanics and mortality inputs

This section recalls the individual-account tontine framework of [22], which provides the account dynamics underlying the MBG overlay. In particular, we state the homogeneous, large-pool representative-member mortality-credit rule from that framework. We then specify the deterministic and stochastic mortality inputs needed for both the mortality-credit amounts in the stochastic-control formulation and the death times used in MBG valuation.

3.1 Representative-member tontine gain rule

The tontine-account mechanics use the decision-time grid $\mathcal{T} = \{t_m\}_{m=0}^{M}$ defined in (2.1). For any time-dependent quantity f, we use the shorthand $f_{m^-} := f(t_m^-)$ and $f_{m^+} := f(t_m^+)$.

The tontine pool consists of L members, indexed by $\ell = 1, \ldots, L$. For a solvent account, the account update at each decision time $t_m \in \mathcal{T}$ is ordered as $t_m^- \longrightarrow t_m \longrightarrow t_m^+$.

- t_m^-, $m = 1, \ldots, M$: end-of-period portfolio balances are observed. Within the tontine account, balances of members who died in $[t_{m-1}^+, t_m^-]$ are forfeited, and the corresponding mortality credits are distributed to surviving members before withdrawals and rebalancing.
- t_m and t_m^+, $m = 0, \ldots, M-1$: each surviving member withdraws an amount q_m at t_m and then rebalances their portfolio at t_m^+.

Between t_{m-1}^+ and t_m^-, $m = 1, \ldots, M$, asset returns and mortality evolve over the period. At $t_M = T$, the portfolio is liquidated with no withdrawal or rebalancing.

Let $\bar{W}_{m^-}^{\ell}$ denote the real (inflation-adjusted) account balance of member ℓ immediately before any mortality-credit distribution at time t_m^-, $m = 0, \ldots, M$. Over each period $[t_{m-1}^+, t_m^-]$, $m = 1, \ldots, M$, mortality outcomes are realized: if member ℓ dies during the period, their account balance is forfeited; if member ℓ survives to t_m^-, they receive a mortality credit c_m^{ℓ} before the scheduled withdrawal at t_m.

For $m = 1, \ldots, M$ and $\ell = 1, \ldots, L$, let δ_{m-1}^{ℓ} denote the conditional one-period death probability of member ℓ over $[t_{m-1}^+, t_m^-]$, given survival to t_{m-1}^+:

$$\delta_{m-1}^{\ell} = \mathbb{P}\big(\text{member } \ell \text{ dies during } [t_{m-1}^+, t_m^-] \,\big|\, \text{member } \ell \text{ is alive at } t_{m-1}^+\big) . \tag{3.1}$$

Thus, $1 - \delta_{m-1}^{\ell}$ is the corresponding conditional probability that member ℓ survives to t_m^-.

In [22], the mortality-credit rule is obtained by combining an individual fairness condition with a pool-level budget constraint. For finite heterogeneous pools, exact ex-post budget balance can be enforced through a group-gain adjustment; under the large-pool/small-bias approximation used below, this adjustment is set equal to one. Equivalently, no single member's expected mortality credit is material relative to aggregate expected forfeitures. The large-pool mortality-credit rule used here is

$$c_m^\ell = \frac{\delta_{m-1}^\ell}{1-\delta_{m-1}^\ell}\,\bar{W}_{m^-}^\ell, \qquad m = 1,\dots,M, \quad \ell = 1,\dots,L. \tag{3.2}$$

The derivation and the finite-pool group-gain construction are given in [22]; see also [51]. In the present paper, (3.2) serves as the large-pool mortality-credit rule for the subsequent stochastic-control and MBG pricing formulation.

In the optimal control problem, we suppress the member superscript ℓ and work with a representative surviving member in a homogeneous pool. The mortality credit can then be written as

$$c_m = g_m\,\bar{W}_{m^-}, \qquad \text{where} \quad g_m = \frac{\delta_{m-1}}{1-\delta_{m-1}}. \tag{3.3}$$

The tontine gain rate g_m is the proportional mortality-credit uplift received by a surviving member over $[t_{m-1}^+, t_m^-]$. Thus the mortality-credit component of the subsequent wealth recursion is determined by the annual death probabilities $\{\delta_{m-1}\}_{m=1}^M$ and the associated gain rates $\{g_m\}_{m=1}^M$.

3.2 Mortality inputs

The annual death-probability sequence $\{\delta_{m-1}\}_{m=1}^M$ is the mortality input for both the representative-member mortality-credit rule and the death-time simulation used in MBG valuation. We now describe how this sequence is specified under deterministic mortality and generated under stochastic mortality.

3.2.1 Deterministic mortality

In the deterministic case, we work with a fixed period life table, such as the Canadian Pensioner Mortality Tables or Australian mortality from the Human Mortality Database (HMD) [26]. Let q_{a,y_0} denote the one–year death probability for an individual aged a in the selected period table, where y_0 denotes its fixed reference year. Let the retiree be aged a_0 at retirement. 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 aged $a_0 + m$.

Holding the period-table year fixed at y_0, we define the corresponding one–year conditional death probabilities by

$$\delta_{m-1} \;:=\; q_{a_0+(m-1),\,y_0}, \qquad m = 1,\dots,M. \tag{3.4}$$

The deterministic sequence $\{\delta_{m-1}\}_{m=1}^M$ then determines the corresponding tontine gain rates $\{g_m\}_{m=1}^M$, where $g_m = \delta_{m-1}/(1-\delta_{m-1})$.

3.2.2 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 [59]. Let $D_{a,y}$ denote the number of deaths at age a in calendar year y and $E_{a,y}^0$ the corresponding initial exposure. We assume $D_{a,y} \sim \mathrm{Binomial}\big(E_{a,y}^0, q_{a,y}\big)$, where $q_{a,y}$ is the one–year death probability. Its systematic component is captured by a linear predictor

$$\eta_{a,y} = \alpha_a + \sum_{i=1}^{G} \beta_a^{(i)}\,\kappa_y^{(i)} + \beta_a^{(0)}\,\gamma_{y-a}, \tag{3.5}$$

where α_a describes the age profile, $\kappa_y^{(i)}$ are period factors and γ_{y-a} is a cohort effect; G denotes the number of GAPC period factors. For the numerical implementation used in this paper, the linear predictor is related directly to the one–year death probability through the logit link $\operatorname{logit} q_{a,y} = \eta_{a,y}$, with $\operatorname{logit}(u) = \log\big(u/(1-u)\big)$ 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,

$$\eta_{a,y} = \alpha_a + \beta_a\,\kappa_y, \tag{3.6}$$

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

$$\kappa_y = d_\kappa + \kappa_{y-1} + \xi_y, \qquad \xi_y \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\kappa^2). \tag{3.7}$$

Cairns–Blake–Dowd (CBD) model

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

$$\eta_{a,y} = \kappa_y^{(1)} + (a - \bar{a})\,\kappa_y^{(2)}. \tag{3.8}$$

As in the LC specification, we set $\operatorname{logit} q_{a,y} = \eta_{a,y}$. The bivariate period factor $\boldsymbol{\kappa}_y = (\kappa_y^{(1)}, \kappa_y^{(2)})^\top$ is usually specified as a random walk with drift,

$$\boldsymbol{\kappa}_y = \boldsymbol{d}_\kappa + \boldsymbol{\kappa}_{y-1} + \boldsymbol{\xi}_y, \qquad \boldsymbol{\xi}_y \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma_\kappa). \tag{3.9}$$

For both the LC and CBD specifications, model estimation, including the required normalizations and any model-specific identifiability constraints, follows the standard StMoMo implementation [59].

Death-probability inputs. Once an LC or CBD model has been calibrated to historical deaths and initial exposures, its fitted and projected period factors determine a surface of one–year death probabilities $\{q_{a,y}\}$. Let y_{proj} denote the first projected calendar year. For a representative retiree aged a_0 at the start of year y_{proj}, with annual decision times $t_m = m$, we define

$$\delta_{m-1} \;:=\; q_{a_0+(m-1),\, y_{\text{proj}}+(m-1)}, \qquad m = 1, \ldots, M. \tag{3.10}$$

In the homogeneous pool considered here, we set $\delta_{m-1}^{\ell} \equiv \delta_{m-1}$ for all members ℓ. Unlike the fixed-period deterministic specification, the stochastic sequence advances in both age and projected calendar year. These probabilities feed directly into the tontine gain rate $g_m = \delta_{m-1}/(1-\delta_{m-1})$ in (3.3). 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 3.1 (Mortality inputs in simulation)**.** *The same one–year conditional death probabilities also underpin our simulation framework. Let n index a simulated path. In the deterministic case, the sequence $\{\delta_{m-1}\}_{m=1}^{M}$ is fixed, so $\delta_{m-1}^{(n)} \equiv \delta_{m-1}$ for all n. Under stochastic LC or CBD mortality, each simulated mortality surface $\{q_{a,y}^{(n)}\}$ yields*

$$\delta_{m-1}^{(n)} \;:=\; q^{(n)}_{a_0+(m-1),\, y_{\text{proj}}+(m-1)}, \qquad m = 1, \ldots, M. \tag{3.11}$$

These probabilities determine the pathwise tontine gain rates through (4.4)*; when the account is solvent, the corresponding gain rate is $\delta_{m-1}^{(n)}/(1-\delta_{m-1}^{(n)})$. They enter the wealth recursion in the NN training (Section 6) and are also used to generate death times and payouts in the MBG valuation (Section 7).*

4 Stochastic control framework

We now turn to the stochastic-control framework for the representative member in the homogeneous large-pool tontine. This member is modelled as a domestic investor with access to four 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. Throughout the stochastic-control formulation, unless otherwise stated, account and portfolio quantities are measured in real (inflation-adjusted) terms. 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 B 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 $\{X_t\}_{0\le t\le T}$ (resp. by x) the controlled underlying index process (resp. a generic state of the system), where

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

Recall the set of decision times $\mathcal{T}$ defined in (2.1). 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

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

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 PIDE-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 [29]. The full SDE specification appears in Subsection B.2
- Historical block bootstrap: In our main experiments, both domestic-only and internationally diversified, we simulate asset paths nonparametrically via the stationary block bootstrap [15, 38, 43, 44]. 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 $\{X_t\}_{0\le t\le T}$ from (4.2) capture the passive evolution of the index values between decision times. Active decisions, namely mortality updates, withdrawals, and rebalancing, are applied only at $\{t_m\}_{m=0}^{M-1}$, as detailed in the next subsection.

4.2 Mortality-credit and fee updates

We define the investor's portfolio wealth at time t by

$$\bar{W}_t = \begin{cases} S_t^d + B_t^d, & \text{domestic-only,} \\ S_t^d + B_t^d + S_t^f + B_t^f, & \text{internationally diversified.} \end{cases} \tag{4.3}$$

At a decision time t_m^-, $\bar{W}_{m^-}$ denotes the investor's pre-update wealth, immediately before the mortality-credit and tontine-management-fee updates. The account is solvent at t_m^- when $\bar{W}_{m^-} > 0$.

At each $t_m \in \mathcal{T}$, the sequence of actions at $t_m^- \to t_m \to t_m^+$ described in Subsection 3.1 applies. To simplify bookkeeping, we use a uniform event structure across all $t_m \in \mathcal{T}$, with the following initial- and terminal-time conventions:

- At each time t_m^-, $m = 0, 1, \ldots, M$, the mortality credit c_m defined in (3.3) is applied when $m \geq 1$ and $\bar{W}_{m^-} > 0$. No mortality credit is applied at t_0, since no forfeitures have yet occurred, or after the account becomes insolvent.
- At each time t_m, $m = 0, \ldots, M-1$, 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 $q_M = 0$.

Accordingly, the tontine gain rate in (3.3) is defined by

$$g_m = \begin{cases} \dfrac{\delta_{m-1}}{1-\delta_{m-1}}, & m = 1, \ldots, M, \quad \bar{W}_{m^-} > 0, \\ 0, & \text{otherwise}. \end{cases} \tag{4.4}$$

Let

$$\widetilde{W}_{m^-} = (1 + g_m)\bar{W}_{m^-}. \tag{4.5}$$

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

We introduce a baseline tontine management fee that is deducted once per year, at each decision time t_m, $m = 1, \ldots, M$. In line with Australian superannuation practice, and in particular the administration fee structure in QSuper's Lifetime Pension PDS [46], 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 that the tontine fee is applied only when the account is solvent and is omitted at t_0, define

$$\varrho_m = \varrho\, \mathbf{1}_{\{m \geq 1,\, \bar{W}_{m^-} > 0\}}, \qquad m = 0, \ldots, M. \tag{4.6}$$

With the conventions established in (4.4)–(4.6), the investor's pre-withdrawal wealth, i.e. wealth immediately before withdrawal, after applying the mortality credit and deducting the tontine fee, is

$$W_{m^-} = \big(1 - \varrho_m\big)\widetilde{W}_{m^-}, \qquad \widetilde{W}_{m^-} \text{ given by (4.5)}. \tag{4.7}$$

Unless otherwise stated, we use W_{m^-} to denote this pre-withdrawal wealth.

4.3 Withdrawal and rebalancing controls

At each decision time $t_m \in \mathcal{T}$, the withdrawal q_m reduces W_{m^-} to the post-withdrawal wealth

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

We model q_m, for $m = 0, \ldots, M$, as a withdrawal control that depends on pre-withdrawal 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 $m = 0, \ldots, M-1$, we denote the rebalancing control by $\boldsymbol{p}_m(\cdot)$, the vector of proportions of post-withdrawal wealth allocated to the asset indices. This control depends on t_m and on the wealth W_{m^+} after the cash withdrawal q_m in (4.8). Formally,

$$\boldsymbol{p}_m(\cdot) : (W_{m^+}, t_m) \mapsto \boldsymbol{p}_m = \boldsymbol{p}(W_{m^+}, t_m),$$

where

$$\boldsymbol{p}_m = \begin{cases} \big(p^d_{s,m}\big), & \text{domestic-only case}, \\ \big(p^d_{s,m},\, p^d_{b,m},\, p^f_{s,m}\big), & \text{internationally diversified case}. \end{cases} \tag{4.9}$$

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

$$X_{m^+} = \begin{cases} \big(S^d_{m^+},\, B^d_{m^+}\big), & \text{domestic-only,} \\ \quad \text{with } S^d_{m^+} = p^d_{s,m} W_{m^+}, \qquad B^d_{m^+} = p^d_{b,m} W_{m^+}, & \\ \big(S^d_{m^+},\, B^d_{m^+},\, S^f_{m^+},\, B^f_{m^+}\big), & \text{internationally diversified,} \\ \quad \text{with } \begin{cases} S^d_{m^+} = p^d_{s,m} W_{m^+}, \qquad B^d_{m^+} = p^d_{b,m} W_{m^+}, \\ S^f_{m^+} = p^f_{s,m} W_{m^+}, \qquad B^f_{m^+} = p^f_{b,m} W_{m^+}. \end{cases} & \end{cases} \tag{4.10}$$

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

4.4 Insolvency and debt convention

The account becomes insolvent at t_m^+ when a withdrawal results in $W_{m^+} \leq 0$. In the event of insolvency, the portfolio is liquidated and trading stops. The negative balance is carried in the domestic bond position and accumulates at the borrowing rate (a spread above the bond rate), and no further mortality credits or tontine management fees are applied. Minimum withdrawals nevertheless continue, so once the account is exhausted, they are funded by borrowing and contribute to the accumulated debt.

4.5 Admissible control sets

We denote by $\mathcal{Z}_q(W_{m^-}, t_m)$ and $\mathcal{Z}_p(W_{m^+}, t_m)$ the wealth- and time-dependent admissible sets for the withdrawal and rebalancing controls, respectively, and by $\mathcal{Z}(W_{m^-}, W_{m^+}, t_m)$ the admissible set for the control pair $\big(q_m(\cdot), \boldsymbol{p}_m(\cdot)\big)$. We discuss these admissible sets below.

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

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

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.[2] Time-varying bounds (e.g. imposing $q_{\min} = q_{\max}$ over an initial period to force high early spending) can also be 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 (inflation-adjusted) 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. If a withdrawal renders the account insolvent, the insolvency convention in Subsection 4.4 applies. This treatment 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 [41]. 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 [52].

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 introduce

$$\Delta^{(k)} := \left\{ (p_1, \ldots, p_k) \in [0,1]^k \;\middle|\; \sum_{i=1}^{k} p_i \le 1 \right\}, \qquad \mathbf{0}^{(1)} := (0), \qquad \mathbf{b}^{(3)} := (0,1,0),$$

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

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

where W_{m^+} is the post-withdrawal wealth defined in (4.8). In the domestic-only case, $\mathbf{0}^{(1)}$ assigns the entire balance to the domestic bond. In the internationally diversified case, $\mathbf{b}^{(3)} = (0,1,0)$ assigns the entire balance to the domestic bond and zero to the other assets. When the balance is negative, this fixed allocation is the bookkeeping convention under which debt accrues at the borrowing rate. At t_M, these singleton controls are notational conventions only, since the portfolio is liquidated and no rebalancing occurs.

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

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

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:

$$\mathcal{A} = \left\{ \mathcal{U} = (q_m, \boldsymbol{p}_m)_{0 \le m \le M} \;\middle|\; (q_m, \boldsymbol{p}_m) \in \mathcal{Z}(W_{m^-}, W_{m^+}, t_m) \right\}. \tag{4.14}$$

For any t_m, we define the subset of controls applicable from t_m onward by

$$\mathcal{U}_m = \left\{ (q_\ell, \boldsymbol{p}_\ell)_{m \le \ell \le M} \;\middle|\; (q_\ell, \boldsymbol{p}_\ell) \in \mathcal{Z}(W_{\ell^-}, W_{\ell^+}, t_\ell) \right\}. \tag{4.15}$$

[2] `https://www.theglobeandmail.com/report-on-business/math-prof-tests-investing-formulas-strategies/article22397218/`; `https://cs.uwaterloo.ca/~paforsyt/Canasta.html`.

5 EW–CVaR portfolio formulation

We use CVaR to quantify downside risk in terminal wealth. 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.

For a tail probability $\alpha \in (0, 1)$, typically $\alpha = 0.01$ or $\alpha = 0.05$, we define the lower-tail CVaR of terminal wealth directly through the optimization representation of [50]:

$$\mathrm{CVaR}_{\alpha,\mathcal{U}_0}^{x_0,t_0} = \sup_{W\in\mathbb{R}} \mathbb{E}_{\mathcal{U}_0}^{X_{0^+},t_0^+}\Big[W + \tfrac{1}{\alpha}\min\big(W_T - W, 0\big) \,\big|\, X_{0^-} = x_0\Big]. \tag{5.1}$$

Here, W is a candidate threshold that partitions the distribution of W_T into its lower tail and the remainder. Intuitively, $\mathrm{CVaR}_\alpha(W_T)$ represents the average terminal wealth in the worst α fraction of outcomes [49].

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)

$$\mathrm{EW}_{\mathcal{U}_0}^{x_0,t_0} = \mathbb{E}_{\mathcal{U}_0}^{X_{0^+},t_0^+}\left[\sum_{m=0}^{M} q_m\right] \tag{5.2}$$

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 the convention underlying [4].

Remark 5.1 (Discounting and mortality weighting)**.** *We remark that* (5.2) *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.*

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 [40].

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

$$\sup_{\mathcal{U}_0\in\mathcal{A}} \left\{ \mathrm{EW}_{\mathcal{U}_0}^{x_0,t_0} + \gamma\, \mathrm{CVaR}_{\alpha,\mathcal{U}_0}^{x_0,t_0} \right\}. \tag{5.3}$$

Remark 5.2 (Pre-commitment and induced time-consistent interpretation)**.** *The scalarized objective in* (5.3) *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* (5.1). *Fixing W^* thereafter yields the equivalent expected-value control problem*

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

This equivalent control problem 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. [12, 19, 54]).*

Using (5.1)–(5.2), we recast (5.3) 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

$$(PCEC_{t_0}(\alpha,\gamma)) : V\left(x_0, t_0^-\right) = \sup_{\mathcal{U}_0\in\mathcal{A}} \sup_{W} \Bigg\{ \mathbb{E}_{\mathcal{U}_0}^{X_{0^+},t_0^+}\Bigg[\sum_{m=0}^{M} q_m + \gamma\left(W + \frac{1}{\alpha}\min\left(W_T - W, 0\right)\right) + \epsilon\, W_T \Bigg| X_{0^-} = x_0\Bigg]\Bigg\} \tag{5.4}$$

$$\text{subject to:} \begin{cases} \text{(dynamics)} & \{X_t\} \text{ evolves via (4.2)} \qquad t \notin \mathcal{T}, \\ \text{(tontine)} & W_{m^-} \text{ given by (4.7)}, \\ \text{(withdrawal)} & W_{m^+} = W_{m^-} - q_m \text{ as in (4.8)}, \\ \text{(rebalancing)} & X_{m^+} \text{ as defined in (4.10), using } \boldsymbol{p}_m(\cdot), \\ & (q_m(\cdot),\, \boldsymbol{p}_m(\cdot)) \in \mathcal{Z}(W_{m^-}, W_{m^+}, t_m), \quad m = 0,\ldots,M, \end{cases} \quad t_m \in \mathcal{T}. \tag{5.5}$$

We denote by $\mathcal{U}_0^* = \mathcal{U}_0^*(\gamma)$ the control that attains the supremum in (5.4)-(5.5), 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 (5.4). If we have a maximum withdrawal constraint, and if $W_t \gg W$ as $t \to T$, the investment controls may become non-unique. In such well-funded states, we can break investment policy ties either by setting $\epsilon < 0$, which selects investments in bonds, or by setting $\epsilon > 0$, which selects investments in stocks. Choosing $|\epsilon| \ll 1$ ensures that this term only has an effect if $W_t \gg W$ and $t \to T$. See [20] for more discussion of this.

We interchange the supremum operators in (5.4) to express the value function in a more computationally tractable form:

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

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

$$W^*(x_0) = \arg\max_{W} \sup_{\mathcal{U}_0\in\mathcal{A}} \Bigg\{ \mathbb{E}_{\mathcal{U}_0}^{X_{0^+},t_0^+}\Bigg[\sum_{m=0}^{M} q_m + \gamma\big(W + \frac{1}{\alpha}\min\left(W_T - W, 0\right)\big) + \epsilon\, W_T \Bigg| X_{0^-} = x_0\Bigg]\Bigg\}. \tag{5.7}$$

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 $\mathbb{R}^2$:

$$\mathcal{S}(\alpha): \quad \left\{ \left(\mathrm{CVaR}_{\alpha,\mathcal{U}_0^*}^{x_0,t_0}\left[W_T\right],\ \mathrm{EW}_{\mathcal{U}_0^*}^{x_0,t_0} \right) : \gamma > 0 \right\}, \tag{5.8}$$

traced out by solving (5.4) for each $\gamma > 0$. In other words, for any fixed level of risk aversion γ, the corresponding point in (5.8) cannot be improved in the EW–CVaR sense by any other admissible strategy in $\mathcal{A}$.

6 Neural network formulation

We numerically solve the EW–CVaR stochastic-control problem for the representative member by parameterizing the withdrawal and rebalancing controls with NNs. Our approach builds on the growing literature that uses NNs to approximate the optimal control directly, avoiding dynamic programming methods [6, 31, 47, 48, 58]. These methods have been termed "global-in-time" machine-learning approaches to stochastic control [27]. This contrasts with the stacked NN approach, in which a separate NN is used to approximate the control at each rebalancing step [25, 56]. More generally, these techniques are special cases of "policy function approximation" for optimal stochastic control [45]. Related NN-based approximation methods have also been developed for certain classes of mean-field games and mean-field control problems [8, 42].

For any admissible control $\mathcal{U}_0 \in \mathcal{A}$ and any CVaR threshold W, we define the objective

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

subject to the dynamics and constraints specified in (5.5).

The value function $V(x_0, t_0^-)$ of the $PCEC_{t_0}(\alpha, \gamma)$ problem can be written as

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

6.1 NN parameterization of admissible controls

The essence of our NN approach is to directly parameterize the withdrawal and rebalancing components of an admissible control policy $\mathcal{U}_0$ in (6.2) 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 $\widehat{\boldsymbol{p}}(W_{m^+}, t_m; \boldsymbol{\theta}_p)$ the NN parameterizations of the withdrawal control $q_m(\cdot)$ and the rebalancing control $\boldsymbol{p}_m(\cdot)$, respectively. Formally, we write

$$\begin{aligned} \widehat{q}_m(\cdot) &:= \widehat{q}\left(W_{m^-}, t_m; \boldsymbol{\theta}_q\right) \simeq q_m(W_{m^-}, t_m), \quad m = 0, 1, \ldots, M, \\ \widehat{\boldsymbol{p}}_m(\cdot) &:= \widehat{\boldsymbol{p}}\left(W_{m^+}, t_m; \boldsymbol{\theta}_p\right) \simeq \boldsymbol{p}_m(W_{m^+}, t_m), \quad m = 0, 1, \ldots, M, \\ \widehat{\mathcal{U}}_0 &:= \left\{ \left(\widehat{q}_m(\cdot), \widehat{\boldsymbol{p}}_m(\cdot)\right) \,\middle|\, m = 0, \ldots, M \right\} \simeq \mathcal{U}_0. \end{aligned} \tag{6.3}$$

Thus, $\widehat{\mathcal{U}}_0$ denotes the NN-parameterized control policy. The collection of all such policies satisfying the admissibility conditions defining $\mathcal{A}$ is denoted by

$$\widehat{\mathcal{A}} := \left\{ \widehat{\mathcal{U}}_0 \in \mathcal{A} \,\middle|\, \widehat{\mathcal{U}}_0 \text{ has the form (6.3)} \right\}.$$

We approximate the original policy search by restricting attention to $\widehat{\mathcal{A}}$. Specifically, for fixed $\widehat{\mathcal{U}}_0 \in \widehat{\mathcal{A}}$ and fixed W, define the NN-induced objective $V_{NN}(x_0, t_0^-; \widehat{\mathcal{U}}_0, W)$ by

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

subject to the system evolution under the NN-parameterized policy given by:

$$\begin{cases} \text{(dynamics)} & \{X_t\} \text{ evolves via (4.2)}, \qquad\qquad t \notin \mathcal{T}, \\ \left.\begin{array}{ll} \text{(tontine)} & W_{m^-} \text{ computed via (4.7)}, \\ \text{(withdrawal)} & W_{m^+} = W_{m^-} - \widehat{q}_m(\cdot) \text{ as in (4.8)}, \\ \text{(rebalancing)} & X_{m^+} \text{ computed via (4.10) using } \widehat{\boldsymbol{p}}_m(\cdot), \\ & \left(\widehat{q}_m(\cdot), \widehat{\boldsymbol{p}}_m(\cdot)\right) \in \mathcal{Z}\left(W_{m^-}, W_{m^+}, t_m\right), \end{array}\right\} & t_m \in \mathcal{T}. \end{cases} \tag{6.5}$$

The value function $V(x_0, t_0^-)$, defined in (6.2), is approximated by the NN-induced value function $V_{NN}(x_0, t_0^-)$, obtained by restricting the policy search to $\widehat{\mathcal{A}}$:

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

6.2 Network architecture for controls

We use two fully connected feedforward neural networks: one for withdrawals $\widehat{q}(\cdot)$, and one for rebalancing $\widehat{\boldsymbol{p}}(\cdot)$, 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) $\widehat{\boldsymbol{p}}(\cdot)$ uses post-withdrawal wealth, i.e. (W_{m^+}, t_m).

To enforce the control constraints in (4.11) and (4.12) within the NN-parameterized policy, we use output activation functions together with deterministic terminal and insolvency branches.

- Withdrawal control: Let $z_q \in \mathbb{R}$ denote the pre-activation output of the withdrawal network's final layer. For non-terminal decision times, we apply a sigmoid transformation scaled by a wealth-dependent range; at the terminal date, we impose the deterministic convention $q_M = 0$:

$$\widehat{q}_m(W_{m^-}, t_m; \boldsymbol{\theta}_q) = \begin{cases} q_{\min} + \max\left(\min(q_{\max}, W_{m^-}) - q_{\min},\, 0\right) \dfrac{1}{1+e^{-z_q}}, & m = 0, \ldots, M-1, \\ 0, & m = M. \end{cases} \tag{6.7}$$

  Since the sigmoid function $1/(1+e^{-z_q})$ maps into $(0,1)$, the first branch enforces the non-terminal withdrawal bounds in $\mathcal{Z}_q(W_{m^-}, t_m)$. The second branch enforces the terminal convention $q_M = 0$. Hence, $\widehat{q}_m(\cdot) \in \mathcal{Z}_q(W_{m^-}, t_m)$ for all $m = 0, \ldots, M$, without introducing optimization constraints on the trainable parameters.

- Rebalancing control: Let $(z_1, \ldots, z_{k+1}) \in \mathbb{R}^{k+1}$ denote the logits output by the rebalancing network's final layer, where $k = 1$ for the domestic-only portfolio and $k = 3$ for the internationally diversified portfolio. Define the auxiliary $(k+1)$-dimensional softmax vector $\widetilde{\boldsymbol{p}}_m = (\widetilde{p}_m^{(1)}, \ldots, \widetilde{p}_m^{(k+1)})$ by

$$\widetilde{p}_m^{(i)} = \frac{e^{z_i}}{\sum_{\ell=1}^{k+1} e^{z_\ell}}, \qquad i = 1, \ldots, k+1.$$

  For solvent non-terminal decision times, the first k softmax components are used as the rebalancing control; at the terminal date or after insolvency, the relevant deterministic convention is imposed:

$$\widehat{\boldsymbol{p}}_m(W_{m^+}, t_m; \boldsymbol{\theta}_p) = \begin{cases} \left(\widetilde{p}_m^{(1)}, \ldots, \widetilde{p}_m^{(k)}\right), & m = 0, \ldots, M-1, \quad W_{m^+} > 0, \\ \mathbf{0}^{(1)}, & \text{domestic-only, if } m = M \text{ or } W_{m^+} \leq 0, \\ \mathbf{b}^{(3)}, & \text{internationally diversified, if } m = M \text{ or } W_{m^+} \leq 0. \end{cases} \tag{6.8}$$

  Therefore, $\widehat{\boldsymbol{p}}_m(\cdot) \in \mathcal{Z}_p(W_{m^+}, t_m)$ for all $m = 0, \ldots, M$. These deterministic branches are part of the NN-parameterized policy map, not post-processing steps in the optimization.

With these output transformations in place, an NN-parameterized control policy $\widehat{\mathcal{U}}_0$ is fully determined by its trainable network parameters $(\boldsymbol{\theta}_q, \boldsymbol{\theta}_p)$. Accordingly, 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 make this dependency explicit.

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:

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

We denote by $(\boldsymbol{\theta}_q^*, \boldsymbol{\theta}_p^*)$, 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-induced objective (6.9) is unconstrained in terms of the training parameters. This allows the use of standard gradient-based methods for optimization. In our numerical implementation (Sections B and 9), 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 3.1). These sample paths are indexed by $n = 1, \ldots, N$, 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 $(\boldsymbol{\theta}_q, \boldsymbol{\theta}_p)$ and threshold W, the NN-induced objective $V_{NN}\left(x_0,\, t_0^-; \boldsymbol{\theta}_q,\, \boldsymbol{\theta}_p,\, W\right)$ is approximated by

$$V_{NN}\left(x_0,\, t_0^-; \boldsymbol{\theta}_q,\, \boldsymbol{\theta}_p,\, W\right) \approx \widehat{V}_{NN}\left(x_0,\, t_0^-,\, \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W\right) \tag{6.10}$$

where $\widehat{V}_{NN}\left(x_0,\, t_0^-,\, \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W\right) := \ldots$

$$\ldots := \frac{1}{N}\sum_{n=1}^{N}\Bigg[\sum_{m=0}^{M}\widehat{q}_m\big(W_{m^-}^{(n)}, t_m;\, \boldsymbol{\theta}_q\big) + \gamma\big(W + \tfrac{1}{\alpha}\min(W_T^{(n)} - W,\, 0)\big) + \epsilon W_T^{(n)}\Bigg], \tag{6.11}$$

$$\text{s.t.}\quad \left\{\begin{array}{lll} \text{(dynamics)} & \{X_t^{(n)}\} \text{ drawn from the } n\text{-th sample path of index returns,} & \\ & \text{(generated by (4.2))} & t \notin \mathcal{T}, \\ \text{(tontine)} & W_{m^-}^{(n)} \text{ computed via (4.7),} & \\ & \text{using the pathwise } g_m^{(n)} \text{ given by (4.4)(see Remark 3.1),} & \\ \text{(withdrawal)} & W_{m^+}^{(n)} = W_{m^-}^{(n)} - \widehat{q}\big(W_{m^-}^{(n)}, t_m; \boldsymbol{\theta}_q\big), & t_m \in \mathcal{T}. \\ \text{(rebalancing)} & X_{m^+}^{(n)} \text{ computed from (4.10), using } \widehat{\boldsymbol{p}}\big(W_{m^+}^{(n)}, t_m; \boldsymbol{\theta}_p\big), & \\ & \big(\widehat{q}_m(\cdot),\, \widehat{\boldsymbol{p}}_m(\cdot)\big) \in \mathcal{Z}\left(W_{m^-}^{(n)}, W_{m^+}^{(n)}, t_m\right), & \end{array}\right. \tag{6.12}$$

Given the sample-based objective $\widehat{V}_{NN}\big(x_0, t_0^-;\, \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, W\big)$ in (6.11)–(6.12), we train the networks by solving the empirical maximization problem[3]

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

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

$$\widehat{q}_m^*(\cdot) := \widehat{q}(\cdot;\, \boldsymbol{\theta}_q^*), \qquad \widehat{\boldsymbol{p}}_m^*(\cdot) := \widehat{\boldsymbol{p}}(\cdot;\, \boldsymbol{\theta}_p^*), \quad m = 0, \ldots, M.$$

The learned NN-parameterized control policy is given by

$$\widehat{\mathcal{U}}_0^* := \left\{\, \big(\widehat{q}_m^*(\cdot),\, \widehat{\boldsymbol{p}}_m^*(\cdot)\big) \;\middle|\; m = 0, \ldots, M \,\right\}.$$

The NN architecture, training parameters, sample sizes, and sequential transfer-learning procedure over the ordered γ-grid are reported in Appendix A.

We highlight that the learned policy may be affected by approximation, sampling, and optimization errors, and that gradient-based training is not guaranteed to attain the global optimum of the empirical NN problem. Under the assumption that this global optimum is attained, [9] establishes convergence in probability of the sample-based NN value $\widehat{V}_{NN}$ to the population value V as both the number of training samples and the NN capacity increase. The benchmark comparison with the PIDE-based frontiers of [22] in Appendix B provides numerical validation of the learned policies, but does not establish their optimality for the original EW–CVaR stochastic-control problem.

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

7 Representative-contract MBG pricing

This section first develops the stand-alone valuation of the MBG for a representative contract. Section 8 then uses the resulting payout distribution to construct an approximate pooled per-contract valuation. The MBG is valued ex post under the learned NN-parameterized control policy $\widehat{\mathcal{U}}_0^*$ obtained in Section 6. This policy is obtained by optimizing the NN-induced approximation of the pre-commitment EW–CVaR problem $PCEC_{t_0}(\alpha, \gamma)$ in (5.4), under the plan-to-live objective used for the withdrawal and rebalancing problem. Death-contingent MBG benefits are therefore not fed back into the member's account balance, withdrawal decisions, or rebalancing decisions.

The MBG overlay is then valued under this fixed policy. If the member dies before cumulative nominal withdrawals reach the initial purchase price L_0, the overlay pays the shortfall in (2.2) to the member's beneficiaries or estate. Because the payout occurs only at death and is not credited to the member's account, valuation is based on the payout distribution induced by $\widehat{\mathcal{U}}_0^*$.

Under the end-of-interval convention in Remark 2.1, when $m_\tau \leq M$ the MBG payout is made at time t_{m_τ}. If the member survives to the horizon, then $m_\tau = M + 1$ and the MBG payout is zero. We therefore define the random real (inflation-adjusted) MBG payout at inception by

$$Z_g := \begin{cases} \dfrac{\mathrm{CPI}_0}{\mathrm{CPI}_{m_\tau}} \max\Big(L_0 - \displaystyle\sum_{\ell=0}^{m_\tau - 1} \widehat{q}_\ell^*(\cdot)\, \frac{\mathrm{CPI}_\ell}{\mathrm{CPI}_0}, 0\Big), & m_\tau \leq M, \\ 0, & m_\tau = M + 1. \end{cases} \tag{7.1}$$

In the first branch of (7.1), the ratio $\mathrm{CPI}_\ell/\mathrm{CPI}_0$ converts each real (inflation-adjusted) withdrawal $\widehat{q}_\ell^*(\cdot)$ into the nominal amount counted by the guarantee, while $\mathrm{CPI}_0/\mathrm{CPI}_{m_\tau}$ converts the resulting nominal payout made at t_{m_τ} into real (inflation-adjusted) dollars at inception. Consistent with Remark 5.1, MBG payouts are valued using a zero real discount rate. Hence, the real discount factor is one and no additional discount term appears in (7.1).

7.1 Actuarial pricing formula

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 $f_g \geq 0$ 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 physical-measure expected MBG payout plus an explicit prudential buffer based on an upper-tail risk measure (a standard risk-measure/premium-principle approach in actuarial pricing; see, e.g. [14, 60]).

Formally,

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

- $\mathbb{E}^{x_0,t_0}_{\widehat{\mathcal{U}}_0^*}\big[Z_g\big]$ is the physical-measure expectation of the MBG payout Z_g (in real (inflation-adjusted) t_0 dollars), evaluated under the models for index dynamics, mortality, and CPI, with the wealth process generated by the learned NN-parameterized policy $\widehat{\mathcal{U}}_0^*$.
- $\mathrm{CVaR}^{+,\,x_0,t_0}_{\alpha_g}\big[Z_g\big]$ is the upper-tail CVaR of the same real (inflation-adjusted) dollar payout at tail level α_g, computed under the same models and learned policy. The tail level α_g is a prudential-buffer choice and need not coincide with the lower-tail probability α used in the retiree's EW–CVaR objective.[4]
- 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).

[4] The MBG payout is a low-frequency, high-severity liability with a right-skewed distribution, so upper-tail measures such as CVaR provide a natural basis for a prudential buffer.

For fixed (λ, α_g), (7.2) defines the equivalent MBG load f_g from the payout distribution induced by $\widehat{\mathcal{U}}_0^*$. Because f_g is computed ex post and does not alter the initial state x_0 or enter the EW–CVaR control problem, no load–policy fixed point is required.

The valuation functionals in (7.2) are written explicitly as follows, with both quantities expressed in real (inflation-adjusted) dollars at time t_0:

$$\begin{aligned} \mathbb{E}_{\widehat{\mathcal{U}}_0^*}^{x_0,t_0}\big[Z_g\big] &= \mathbb{E}_{\widehat{\mathcal{U}}_0^*}^{X_{0^+},t_0^+}\big[Z_g \mid X_{0^-} = x_0\big], \\ \mathrm{CVaR}_{\alpha_g}^{+,x_0,t_0}\big[Z_g\big] &:= \inf_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{\alpha_g} \mathbb{E}_{\widehat{\mathcal{U}}_0^*}^{X_{0^+},t_0^+}\big[(Z_g - \eta)^+ \mid X_{0^-} = x_0\big] \right\}. \end{aligned} \tag{7.3}$$

For compactness, and under the convention that $m_\tau = M + 1$ when the member survives to T, set $\bar{m}_\tau := \min(m_\tau, M)$. Thus $\bar{m}_\tau = m_\tau$ when death occurs before the horizon, while $\bar{m}_\tau = M$ when the member survives to T. The valuation functionals in (7.3) are evaluated under the system evolution generated by the learned NN-parameterized policy $\widehat{\mathcal{U}}_0^*$ up to this capped interval index:

$$\text{subject to} \begin{cases} \text{(dynamics)} & \{X_t\} \text{ evolves via (4.2)}, \qquad\qquad t \notin \mathcal{T},\ 0 \le t < t_{\bar{m}_\tau} \\ \left.\begin{array}{ll} \text{(tontine)} & W_{m^-} \text{ computed via (4.7)}, \\ \text{(withdrawal)} & W_{m^+} = W_{m^-} - \widehat{q}_m^*(\cdot) \text{ as in (4.8)}, \\ \text{(rebalancing)} & X_{m^+} \text{ computed via (4.10) using } \widehat{\boldsymbol{p}}_m^*(\cdot), \end{array}\right\} & m = 0, \ldots, \bar{m}_\tau - 1. \end{cases} \tag{7.4}$$

Remark 7.1 (Interpreting f_g in starting–rate terms)**.** *Equation* (7.2) *defines an equivalent up-front load $f_g \ge 0$ 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 income 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), for $0 \le f_g < 1$, 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.*

If $f_g = 1$, the equivalent load exhausts the initial contribution and the translated starting rate is zero. If $f_g > 1$, the required load exceeds the initial contribution and therefore cannot be funded solely through an up-front deduction from L_0. Accordingly, the starting-rate translation is economically meaningful only for $0 \le f_g < 1$. 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 valuation functionals in (7.3) are approximated from K independent sample paths, each representing a full set of exogenous drivers for index dynamics, mortality inputs, and CPI. The resulting Monte Carlo estimators of the exact expectation and upper-tail CVaR are denoted by $\widehat{\mathbb{E}}[Z_g]$ and $\widehat{\mathrm{CVaR}}_{\alpha_g}^{+}(Z_g)$, respectively. These sample paths are indexed by $k = 1, \ldots, K$, and all path-dependent quantities carry the superscript (k).

The pricing procedure uses the learned NN-parameterized controls obtained from Section 6. These controls are held fixed during the MBG valuation and applied across all sample paths. For each path $k = 1, \ldots, K$, the wealth evolution follows the dynamics generated by the learned NN-parameterized policy $\widehat{\mathcal{U}}_0^*$. We write $\{X_t^{(k)}\}_{t \in [0,T]}$ for the resulting state trajectory on path k; all withdrawals and allocations at the rebalancing times $t_m \in \mathcal{T}$ are the learned network outputs $\widehat{q}_m^*(W_{m^-}^{(k)}, t_m)$ and $\widehat{\boldsymbol{p}}_m^*(W_{m^+}^{(k)}, t_m)$ evaluated with the path-specific inputs.

As summarized in Remark 3.1, 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

$g_m^{(k)} = \delta_{m-1}^{(k)}/\big(1-\delta_{m-1}^{(k)}\big)$ when computing mortality credits via (4.4)–(4.7). In addition, they drive the simulation of death times and MBG payouts along each path as outlined in Algorithm 7.1 below. Let $\{\mathrm{CPI}_m^{(k)}\}_{m=0}^{M}$ denote the CPI values used on path k at the decision times. The mortality inputs determine the timing of the MBG payout, while the CPI values enter the nominal withdrawal accumulation and the nominal–real conversion used to compute the pathwise real (inflation-adjusted) MBG payout.

The simulation-based pricing details are described in Algorithm 7.1. For each path $k = 1, \ldots, K$, we denote by $Z_g^{(k)}$ the simulated real (inflation-adjusted) dollar MBG payout, obtained by applying (7.1) with the path-specific withdrawal history and death time. Algorithm 7.1 simulates the account under

Algorithm 7.1 Simulation-based MBG pricing under learned NN-parameterized controls

1: initialize the active index list $\mathcal{K} \leftarrow \{1, \ldots, K\}$;
2: initialize MBG payouts $Z_g^{(k)} \leftarrow 0$, $k = 1, \ldots, K$;
3: initialize cumulative nominal withdrawals $\mathcal{C}^{(k)} \leftarrow 0$, $k = 1, \ldots, K$;
4: initialize $X_{0^-}^{(k)} = x_0$ and compute $W_{0^-}^{(k)}$ for all $k \in \mathcal{K}$;
5: **for** $m = 0$ **to** $M-1$ **do**
6:   **for each** $k \in \mathcal{K}$ **do**
7:     compute pathwise $g_m^{(k)}$ via (4.4) (see Remark 3.1) and $W_{m^-}^{(k)}$ via (4.7); {mortality credit}
8:     $q_m^{(k)} \leftarrow \widehat{q}_m^{*}\big(W_{m^-}^{(k)}, t_m\big)$; {real (inflation-adjusted) dollars}
9:     $\mathcal{C}^{(k)} \leftarrow \mathcal{C}^{(k)} + q_m^{(k)} \dfrac{\mathrm{CPI}_m^{(k)}}{\mathrm{CPI}_0}$; {nominal withdrawal accumulation}
10:     $W_{m^+}^{(k)} \leftarrow W_{m^-}^{(k)} - q_m^{(k)}$;
11:     obtain $X_{m^+}^{(k)}$ via (4.10) using $\widehat{\boldsymbol{p}}_m^{*}\big(W_{m^+}^{(k)}, t_m\big)$; {rebalancing}
12:     compute $X_{(m+1)^-}^{(k)}$ via (4.2); {simulation over $[t_m^+, t_{m+1}^-]$}
13:     draw $B_m^{(k)} \sim \mathrm{Bernoulli}(\delta_m^{(k)})$; {death-time}
14:     **if** $B_m^{(k)} = 1$ **then**
15:       $Z_g^{(k)} \leftarrow \dfrac{\mathrm{CPI}_0}{\mathrm{CPI}_{m+1}^{(k)}} \max\Big(L_0 - \mathcal{C}^{(k)},\, 0\Big)$;
16:       remove path k from the active list $\mathcal{K}$;
17:     **end if**
18:   **end for**
19:   **if** $\mathcal{K} = \emptyset$ **then**
20:     break; {all deaths processed}
21:   **end if**
22: **end for**
23: compute the sample mean: $\widehat{\mathbb{E}}[Z_g] \;=\; \dfrac{1}{K}\displaystyle\sum_{k=1}^{K} Z_g^{(k)}$;
24: sort the simulated payouts in descending order to obtain $Z_{(1)} \geq Z_{(2)} \geq \cdots \geq Z_{(K)}$;
25: compute the empirical upper-tail CVaR estimator: $\widehat{\mathrm{CVaR}}_{\alpha_g}^{+}(Z_g) \;=\; \dfrac{1}{\lceil \alpha_g K \rceil}\displaystyle\sum_{j=1}^{\lceil \alpha_g K \rceil} Z_{(j)}$;
26: return $\big(\widehat{\mathbb{E}}[Z_g], \widehat{\mathrm{CVaR}}_{\alpha_g}^{+}(Z_g)\big)$;

the learned NN-parameterized policy until death or the terminal horizon.[5] On Line 13, the death event over $(t_m^+, t_{m+1}^-]$ is determined by a Bernoulli experiment based on the pathwise death probability $\delta_m^{(k)}$. If death occurs, the algorithm computes the real (inflation-adjusted) MBG payout on Line 15, records it as

[5] The estimator of $\mathrm{CVaR}_{\alpha_g}^{+}(Z_g)$ in Algorithm 7.1 is an order-statistic upper-tail average, computed after sorting the simulated payouts in descending order. When $\alpha_g K$ is an integer, it coincides with the empirical CVaR associated with the infimum representation in (7.3); otherwise, it is implemented as the average of the largest $\lceil \alpha_g K \rceil$ simulated MBG payouts.

$Z_g^{(k)}$, and removes the path from the active list. Paths with no death before the horizon retain $Z_g^{(k)} = 0$.

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

$$\widehat{f}_g \;:=\; (L_0)^{-1}\left(\widehat{\mathbb{E}}[Z_g] + \lambda\, \widehat{\mathrm{CVaR}}^+_{\alpha_g}(Z_g)\right). \tag{7.5}$$

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

Remark 7.2 (Inflation treatment in MBG pricing)**.** *In the numerical experiments, wealth is simulated in real (inflation-adjusted) units using bootstrapped real (inflation-adjusted) 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 (inflation-adjusted) t_0 dollars.*

To obtain pathwise inflation adjustments consistent with the bootstrap, CPI is treated as an additional series in the resampling procedure: monthly CPI changes are bootstrapped jointly with the asset-return blocks, 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 $\mathrm{CPI}_m^{(k)}$ *on Line 9 and* $\mathrm{CPI}_{m+1}^{(k)}$ *on Line 15.*

8 Contract pooling and tail risk

The preceding section evaluates the representative contract on a stand-alone basis. In practice, an issuer would hold a portfolio of broadly homogeneous MBG contracts, so part of the payout risk can diversify across contracts. Let J denote the number of contracts over which MBG costs are aggregated; this is distinct from the tontine-pool size underlying the homogeneous large-pool mortality-credit approximation. Let $Z_{g,1}, \ldots, Z_{g,J}$ denote MBG payouts on these contracts, each with the same marginal distribution as the representative-contract payout Z_g. Define the aggregate and pooled per-contract payouts by

$$S_{g,J} := \sum_{j=1}^{J} Z_{g,j}, \qquad \overline{Z}_{g,J} := \frac{S_{g,J}}{J}. \tag{8.1}$$

Linearity gives $\mathbb{E}[\overline{Z}_{g,J}] = \mathbb{E}[Z_g]$. More importantly, subadditivity and positive homogeneity of CVaR give the diversification inequality

$$\mathrm{CVaR}^+_{\alpha_g}(\overline{Z}_{g,J}) = \frac{1}{J}\,\mathrm{CVaR}^+_{\alpha_g}(S_{g,J}) \le \frac{1}{J}\sum_{j=1}^{J}\mathrm{CVaR}^+_{\alpha_g}(Z_{g,j}) = \mathrm{CVaR}^+_{\alpha_g}(Z_g). \tag{8.2}$$

Thus, contract pooling cannot increase the upper-tail CVaR per contract.

To obtain a practical analytical approximation, decompose the representative-contract CVaR as

$$\mathrm{CVaR}^+_{\alpha_g}(Z_g) = \mathbb{E}[Z_g] + \underbrace{\left(\mathrm{CVaR}^+_{\alpha_g}(Z_g) - \mathbb{E}[Z_g]\right)}_{\text{excess-tail risk}}. \tag{8.3}$$

Suppose that contract-specific mortality risk is independent across contracts and is the dominant source of diversifiable excess-tail variation, while common-factor dependence is abstracted from in this approximation. For sufficiently large J, a normal approximation suggests that the excess-tail component of the aggregate CVaR scales approximately as $\sqrt{J}$, and hence as $J^{-1/2}$ per contract. Approximating the coefficient of this $J^{-1/2}$ term by the representative-contract excess-tail risk in (8.3) gives the heuristic

$$\mathrm{CVaR}^+_{\alpha_g}(S_{g,J}) \approx J\,\mathbb{E}[Z_g] + \sqrt{J}\left(\mathrm{CVaR}^+_{\alpha_g}(Z_g) - \mathbb{E}[Z_g]\right), \tag{8.4}$$

and hence

$$\mathrm{CVaR}^+_{\alpha_g}(\overline{Z}_{g,J}) \approx \mathbb{E}[Z_g] + J^{-1/2}\big(\mathrm{CVaR}^+_{\alpha_g}(Z_g) - \mathbb{E}[Z_g]\big). \tag{8.5}$$

Under this approximation, $\mathrm{CVaR}^+_{\alpha_g}(\overline{Z}_{g,J}) \to \mathbb{E}[Z_g]$ as $J \to \infty$.

Using (8.5), define the corresponding approximate pooled-contract load $f^{\mathrm{a}}_{g,J}$ by

$$f^{\mathrm{a}}_{g,J} := (L_0)^{-1} \left(\mathbb{E}[Z_g] + \lambda \left[\mathbb{E}[Z_g] + J^{-1/2} \big(\mathrm{CVaR}^{+}_{\alpha_g}(Z_g) - \mathbb{E}[Z_g] \big) \right] \right). \tag{8.6}$$

Consequently, as $J \to \infty$, $f^{\mathrm{a}}_{g,J} \longrightarrow \frac{(1+\lambda)\mathbb{E}[Z_g]}{L_0}$.

For numerical implementation, the exact expectation and upper-tail CVaR in (8.5)–(8.6) are replaced by the Monte Carlo estimators $\widehat{\mathbb{E}}[Z_g]$ and $\widehat{\mathrm{CVaR}}^{+}_{\alpha_g}(Z_g)$ from Algorithm 7.1. The resulting estimated approximate quantities are defined by

$$\begin{aligned} \widehat{\mathrm{CVaR}}^{+,\mathrm{a}}_{\alpha_g}(\overline{Z}_{g,J}) &:= \widehat{\mathbb{E}}[Z_g] + J^{-1/2} \left(\widehat{\mathrm{CVaR}}^{+}_{\alpha_g}(Z_g) - \widehat{\mathbb{E}}[Z_g] \right), \\ \widehat{f}^{\mathrm{a}}_{g,J} &:= (L_0)^{-1} \left(\widehat{\mathbb{E}}[Z_g] + \lambda\, \widehat{\mathrm{CVaR}}^{+,\mathrm{a}}_{\alpha_g}(\overline{Z}_{g,J}) \right). \end{aligned} \tag{8.7}$$

These are the quantities evaluated in the contract-pooling numerical results.

The approximation abstracts from common market, inflation, and systematic-mortality dependence, which may leave residual tail risk. A direct portfolio-level CVaR calculation preserving these dependencies would require joint simulation of MBG liabilities across contracts and is left for future work.

9 Numerical results: retirement outcomes and MBG pricing

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. Unless otherwise noted, all return series in this section are measured at a monthly frequency in real AUD terms, where "real" means inflation-adjusted. Monetary quantities in this section are reported in thousands of real AUD.

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. Detailed construction of the asset-return, AUD/USD exchange rate, and CPI series is reported in Appendix C.1.

Historical market and bootstrap sampling. Future index paths in the internationally diversified experiments are generated nonparametrically using the stationary block bootstrap [15, 38, 43, 44]. 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 [38] 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.

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

Asset	Mean (ann.)	Geo. mean (ann.)	Vol (ann.)	$\mathrm{VaR}_{0.05}$ (mo.)	$\mathrm{CVaR}_{0.05}$ (mo.)
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

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

As indicated in Tables 9.1–9.2, U.S. equity has the highest historical average real return in this sample and is only weakly correlated with Australian equity (correlation ≈ 0.33, well below 1)[6]. 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 learned 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 3.2 and Remark 3.1.

For the deterministic specification, we use the 2021 column of the Australian male 1×1 period life table from the Human Mortality Database (HMD, [26]); the full file `mltper_1x1.txt` covers calendar years 1921–2021.[7] We hold the period-table year fixed at $y_0 = 2021$ over the retirement horizon. Thus, with $a_0 = 65$, the deterministic death-probability sequence is $\delta_{m-1} = q_{65+(m-1),\,2021}$, $m = 1, \ldots, M$, consistently with (3.4). In a homogeneous pool we set $\delta^{\ell}_{m-1} \equiv \delta_{m-1}$ for all members ℓ, and the associated tontine gain rates g_m entering the wealth recursion follow from (3.3).

For stochastic mortality, annual Australian male deaths and central exposures for ages 55–95 and calendar years 1980–2021 are used. The central exposures are converted to initial exposures, and both the LC and CBD models are fitted using a Binomial likelihood with a logit link in R using StMoMo [59]. We simulate 256,000 mortality sequences and extract each sequence diagonally from the projected mortality surface, beginning with age 65 in $y_{\text{proj}} = 2022$. The data construction, calibration, and simulation details are reported in Appendix C.2.

9.3 Experimental setup and scenarios

For the internationally diversified experiments, we retain the same decumulation framework used in the benchmark validation (Table B.1), but now from the perspective of a 65–year–old Australian male retiree investing in the four indices described in Subsection 9.1. 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 $L_0 = 1000$ (thousand AUD at inception) for the tontine product with the embedded MBG overlay described in Section 2. The planning horizon is $T = 30$ years with annual decision times $t_m = m$, $m = 0, \ldots, M$, 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 [46], we use an annual

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

[7]Available at `https://www.mortality.org/File/GetDocument/hmd.v6/AUS/STATS/mltper_1x1.txt`.

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 the lower-tail $\text{CVaR}_{0.05}(W_T)$. For a given scalarization parameter $\gamma > 0$, we first train the NNs to obtain the learned policy $\widehat{\mathcal{U}}_0^*(\gamma)$ 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 $\widehat{\mathcal{U}}_0^*(\gamma)$ fixed, using the actuarial pricing rule (7.2) and the simulation procedure in Section 7.

9.4 EW–CVaR 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 [22], we include a constant–real–withdrawal, constant–weight 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 AUD) 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 maximizes $\text{CVaR}_{0.05}$ of terminal wealth.

In the two–asset case, the search is over equity fractions $p_s^d \in \{0, 0.1, \ldots, 1\}$ with $p_b^d = 1 - p_s^d$; the best constant–weight benchmark has a 50% domestic equity allocation and $\text{CVaR}_{0.05}(W_T) \approx -617.1$. 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 $\text{CVaR}_{0.05}(W_T) \approx -446.1$. These two constant–strategy benchmarks appear as single points in the efficient–frontier figures for the two–asset and four–asset experiments, respectively.

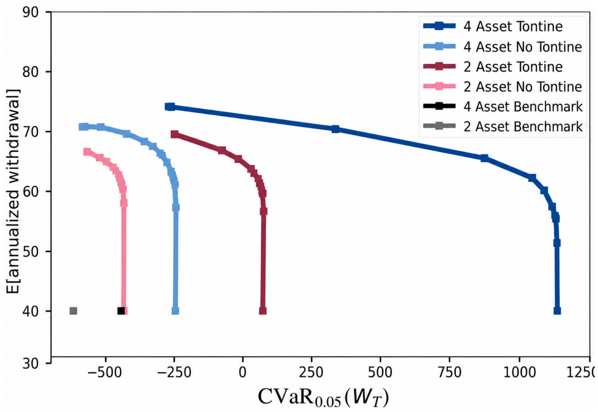

FIGURE 9.1: *EW–CVaR efficient frontiers for domestic–only (two–asset) and internationally diversified (four–asset) portfolios, with and without a tontine overlay. Constant–strategy benchmarks are shown as single points.*

Figure 9.1 shows the EW–CVaR efficient frontiers traced out by the learned policies in the two–asset (domestic–only) and four–asset (internationally diversified) cases, both with and without a tontine overlay, together with the corresponding constant–strategy benchmarks.

We make the following key observations about Figure 9.1:

- The internationally diversified (four–asset) constant–strategy benchmark lies essentially on the two–asset no–tontine frontier. Thus, at the 4%–rule withdrawal level, the internationally diversified constant–strategy benchmark achieves an EW–CVaR trade–off comparable to that generated by the domestic–only learned EW–CVaR policy without a tontine overlay. Moreover, relative to the two–asset constant–strategy benchmark, international diversification produces a substantial improvement in CVaR at the same withdrawal level. These comparisons provide evidence that international diversification can materially improve the risk–reward trade–off even before longevity pooling is introduced.

- In both the domestic–only and internationally diversified settings, the efficient frontiers traced out by the learned policies lie well above and to the right of their corresponding constant–strategy benchmark points. Within the calibrated setting, this shows that 4%–rule–style strategies with constant real withdrawals and constant portfolio weights are substantially less efficient than the learned NN 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, corresponding to a substantial lower–tail terminal–wealth buffer, the four–asset tontine frontier supports expected annual withdrawals of about 70 (thousand real AUD), or roughly 7% of initial wealth, compared with the fixed 40 (thousand real AUD, 4%) annual withdrawal under the constant–strategy 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 $\text{CVaR}_{0.05}(W_T)$. Allowing the retiree to invest in both Australian and U.S. assets therefore supports higher expected annualized withdrawals for a comparable, or even improved, level of downside risk relative to 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. Thus, 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. Across a broad range of scalarization parameters γ, the four–asset tontine strategies deliver both higher expected annual withdrawals and higher $\text{CVaR}_{0.05}(W_T)$ than the corresponding four–asset no–tontine strategies. Together with the preceding comparison, this illustrates the substantial combined benefit of international diversification and longevity pooling.

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

9.5 EW–CVaR frontiers: effects of stochastic mortality

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.

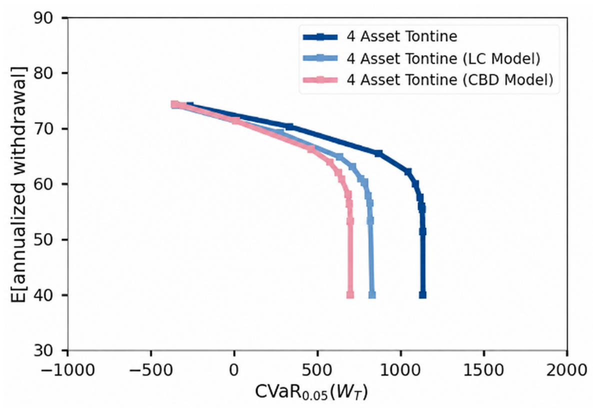

FIGURE 9.2: *EW–CVaR efficient frontiers for the four–asset tontine portfolio under deterministic table mortality and stochastic mortality (LC and CBD models).*

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.

A key point for interpretation is that the horizon corresponds to a very advanced age (age 95 in the baseline calibration with a 30–year horizon). Consequently, frontier points with very high terminal–wealth CVaR, for example, $\text{CVaR}_{0.05}(W_T) \approx 1000$ (thousand real AUD), correspond to policies that place substantial weight on preserving residual wealth at age 95. Values around $\text{CVaR}_{0.05}(W_T) \approx 500$ already represent a sizeable terminal–wealth buffer relative to the annual withdrawal amounts considered. Therefore, for economic interpretation, we focus on the region of Figure 9.2 at or below roughly

$\mathrm{CVaR}_{0.05}(W_T) = 500$, where the controls correspond to higher expected withdrawals rather than preserving large terminal balances.

Over this 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 and downward shift. This shift is consistent with systematic longevity improvement in the stochastic mortality models: when death probabilities fall relative to the deterministic baseline table, survivors receive smaller mortality credits, which lowers terminal wealth in the adverse outcomes that determine $\mathrm{CVaR}_{0.05}(W_T)$. Although systematic longevity improvement under the LC and CBD models reduces the mortality-credit support available to survivors relative to the deterministic baseline, the four–asset tontine continues to offer an attractive combination of high expected withdrawals and improved $\mathrm{CVaR}_{0.05}(W_T)$ relative to the domestic–only and constant–weight strategies considered 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 generally lying slightly to the left of the LC curve over this region. Conditional on modelling systematic longevity risk, the specific choice between the LC and CBD specifications therefore has only a relatively minor effect on the EW–CVaR trade–off.

As a further robustness check, we evaluate the learned controls trained under LC mortality on out-of-sample CBD mortality realizations, without retraining, and compare them with controls trained and evaluated under CBD mortality. The resulting frontiers are nearly coincident; for $\gamma = 1.5$ and $\lambda = 0.05$, the corresponding representative-contract equivalent MBG loads are $\widehat{f}_g = 0.114$ and $\widehat{f}_g = 0.113$, respectively. Full results are reported in Appendix D.

The learned controls provide a wealth- and time-dependent interpretation of the frontier gains from international diversification. The complete control heatmaps and the comparison with the domestic-only allocation are reported in Appendix E.

9.6 Representative-contract MBG pricing

We now report simulation–based representative-contract MBG pricing results using the actuarial pricing formula (7.2) and the Monte Carlo procedure in Algorithm 7.1. As noted in Section 2, the deceased member's tontine account is forfeited to the pool and supports mortality credits for surviving members; it is therefore unavailable to finance the MBG death benefit, which must be funded separately. The effect of aggregating MBG costs across contracts is discussed in Subsection 9.7. In all cases, the MBG is priced *ex post* under the fixed learned policy as discussed previously.

Unless otherwise stated, the MBG pricing experiments use the common base–case parameter set reported in Table 9.3. The choice $\gamma = 1.5$ selects a representative point on the EW–CVaR frontier, L_0 is measured in thousands of AUD at inception, and β_0 is the reference payment rate device described in Remark 7.1.

TABLE 9.3: *Base–case parameters for the MBG pricing experiments.*

γ	L_0	α_g	λ	β_0
1.5	1000	5%	0.05	0.05

The prudential-buffer coefficient $\lambda = 0.05$ is an illustrative loading choice. Sensitivity to the prudential-buffer coefficient λ, the upper-tail probability α_g, and the retiree's scalarization parameter γ, is reported in Appendix F.1.

Under this base case, Table 9.4 compares MBG pricing across four settings: two–asset (domestic–only) and four–asset (internationally diversified) portfolios, each under deterministic table mortality and stochastic LC mortality. The final column reports the implied post–load notional starting payment rate $(1 - \widehat{f}_g)\beta_0$ (see Remark 7.1). Monte Carlo standard errors are reported in Appendix F.2 (Table F.2).

The central finding from Table 9.4 is that the material improvement in the EW–CVaR frontier from international diversification translates into only a modest change in MBG pricing. Across the four settings, the expected payouts and upper-tail CVaR estimates remain close, yielding equivalent loads between 0.102 and 0.109.

Table 9.4: *MBG pricing across asset and mortality settings under the base–case parameters in Table 9.3.*

Setting	$\widehat{\mathbb{E}}[Z_g]$	$\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g)$	$\widehat{f}_g$	$(1-\widehat{f}_g)\beta_0$
2 assets, table mortality	65.63	736.62	0.102	4.49%
2 assets, LC mortality	66.68	755.48	0.104	4.48%
4 assets, table mortality	69.33	736.82	0.106	4.47%
4 assets, LC mortality	70.69	758.28	0.109	4.46%

To examine this pricing stability further, we use the fact that the MBG payout $Z_g \geq 0$ to write $\widehat{\mathbb{E}}[Z_g] = \widehat{\mathbb{P}}(Z_g > 0)\widehat{\mathbb{E}}[Z_g \mid Z_g > 0]$, which decomposes the expected MBG payout into the frequency and conditional severity of positive payouts. Table 9.5 reports these components across the four settings.

Table 9.5: *Decomposition of $\widehat{\mathbb{E}}[Z_g]$ from Table 9.4.*

Setting	$\widehat{\mathbb{P}}(Z_g > 0)$	$\widehat{\mathbb{E}}[Z_g \mid Z_g > 0]$
2 assets, table mortality	17.59%	373.2
2 assets, LC mortality	16.78%	397.3
4 assets, table mortality	19.30%	359.2
4 assets, LC mortality	18.33%	385.7

As shown in Table 9.5, moving from two to four assets slightly raises the probability of a positive MBG payout but lowers its conditional mean, while replacing table mortality with LC mortality has the opposite effect. These offsetting frequency and severity effects keep expected payouts close across the four settings. The corresponding CVaR estimates in Table 9.4 also vary only modestly.

While this decomposition explains the modest variation across the asset and mortality specifications, the large upper tail remains economically important for the equivalent MBG load. The actuarially fair expected–cost component $\widehat{\mathbb{E}}[Z_g]$ is modest (about 6.6%–7.1% of L_0). Thus, under actuarially fair pricing ($\lambda = 0$), the implied equivalent load would be only $\widehat{f}_g \approx \widehat{\mathbb{E}}[Z_g]/L_0 \approx$ 6%–7%, which translates, for $\beta_0 = 5\%$, into an illustrative starting-rate reduction of roughly 33–35 bps. By contrast, the tail measure $\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g)$ is large (about 0.74–0.76 of L_0, roughly an order of magnitude larger than $\widehat{\mathbb{E}}[Z_g]$). With $\lambda = 0.05$, the resulting equivalent load is $\widehat{f}_g \approx 0.102$–$0.109$, corresponding to a post–load notional starting payment rate of about 4.46%–4.49%.

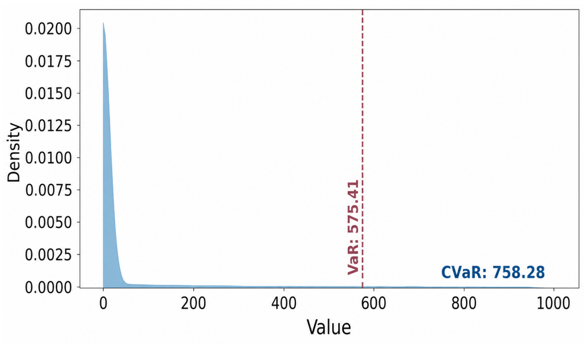

(a) Unconditional empirical distribution of Z_g

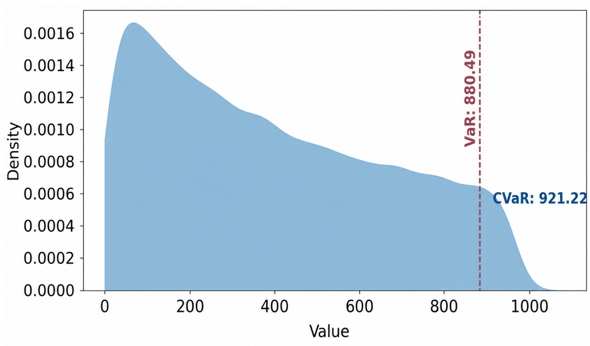

(b) Empirical distribution conditional on $Z_g > 0$

Figure 9.3: *Empirical distributions of the representative-contract MBG payout Z_g (in real dollars at inception): the four–asset tontine with stochastic mortality (LC) under the base–case parameters from Table 9.3.*

Figure 9.3 illustrates the distributional source of this large tail contribution. It 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 upper-tail probability is $\alpha_g = 5\%$. Figure 9.3 (a) shows the unconditional empirical distribution of Z_g, which has a substantial point mass at zero (many scenarios in which the guarantee is out of the money) and a long right tail. Figure 9.3 (b) shows the empirical distribution conditional on $Z_g > 0$, highlighting that when the MBG is triggered the payout can still be large and widely spread. In Figure 9.3 (a), the 5% upper-tail VaR threshold is approximately 575 (thousand real AUD) and the corresponding upper-tail mean is $\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g) \approx 758$, consistent with Table 9.4.

Taken together, the representative-contract results show that international diversification materially improves retirement outcomes while changing the MBG load only modestly; the large upper-tail exposure nevertheless materially raises the load above its expected-cost level.

9.7 Contract pooling results

Section 8 develops the contract-book approximation and defines the estimated approximate quantities in (8.7). We now evaluate these quantities for the four–asset LC base case. Table 9.4 reports the representative-contract estimates for this case: $\widehat{\mathbb{E}}[Z_g] = 70.69$, $\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g) = 758.28$, and $\widehat{f}_g = 0.109$ for $\lambda = 0.05$. Table 9.6 reports the resulting estimated approximate pooled per-contract CVaR, equivalent load, and implied post–load notional starting payment rate $(1 - \widehat{f}^{\mathrm{a}}_{g,J})\beta_0$, with Monte Carlo standard errors reported in Appendix G.

TABLE 9.6: *Approximate contract-pooling results: the four–asset tontine with stochastic mortality (LC), based on the estimators in* (8.7) *with* $\lambda = 0.05$ *and* $\beta_0 = 0.05$.

J	$\widehat{\mathrm{CVaR}}^{+,\mathrm{a}}_{0.05}(\overline{Z}_{g,J})$	$\widehat{f}^{\mathrm{a}}_{g,J}$	$(1 - \widehat{f}^{\mathrm{a}}_{g,J})\beta_0$
1	758.28	0.109	4.46%
10	288.13	0.085	4.57%
50	167.93	0.079	4.60%
100	139.45	0.078	4.61%
500	101.44	0.076	4.62%
1000	92.44	0.075	4.62%
10,000	77.57	0.075	4.63%
20,000	75.56	0.074	4.63%
$J \to \infty$	70.69	0.074	4.63%

Table 9.6 shows that the estimated approximate pooled per-contract upper-tail CVaR $\widehat{\mathrm{CVaR}}^{+,\mathrm{a}}_{0.05}(\overline{Z}_{g,J})$ declines substantially as J increases. Specifically, relative to the representative-contract $\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g) = 758.28$ in Table 9.4, the estimated approximate pooled per-contract counterpart $\widehat{\mathrm{CVaR}}^{+,\mathrm{a}}_{0.05}(\overline{Z}_{g,J})$ falls to 92.44 at $J = 1000$ and to 75.56 at $J = 20{,}000$.

At $\lambda = 0.05$, Table 9.6 shows that the estimated approximate equivalent load $\widehat{f}^{\mathrm{a}}_{g,J}$ falls from 0.109 at $J = 1$ to 0.075 at $J = 1000$ and 0.074 at $J = 20{,}000$, while the corresponding implied post–load starting rate $(1 - \widehat{f}^{\mathrm{a}}_{g,J})\beta_0$ rises from 4.46% to 4.62% and 4.63%, respectively. As $J \to \infty$, evaluating the limit implied by (8.6) at the representative-contract estimates gives $\widehat{f}^{\mathrm{a}}_{g,J} \to 0.074$, with a corresponding implied post–load starting rate of 4.63%. This limiting load is close to the representative-contract expected-cost load $\widehat{\mathbb{E}}[Z_g]/L_0 = 70.69/1000 = 0.071$, and materially below the representative-contract load $\widehat{f}_g = 0.109$ reported in Table 9.4. The limiting load remains slightly above the expected-cost load because the prudential buffer is applied to the entire pooled per-contract CVaR. As $J \to \infty$, the approximate pooled CVaR converges to $\mathbb{E}[Z_g]$, so the buffer converges to $\lambda\mathbb{E}[Z_g]$ rather than to zero.

10 Conclusion and future work

This paper develops a pricing framework for a return-of-purchase-price MBG overlay on an individual-account tontine. The representative-member EW–CVaR optimization problem incorporates dynamic withdrawals and rebalancing across domestic and foreign assets, with mortality credits determined under either deterministic or stochastic mortality. We use NN parameterizations to compute the wealth- and time-dependent withdrawal and rebalancing controls. Under the fixed-horizon plan-to-live convention, the decumulation policy is determined first and the MBG is then valued ex post through an implementation-agnostic equivalent up-front load combining the expected payout under the physical measure with an upper-tail CVaR prudential buffer. The resulting representative-contract payout distribution is then used to approximate the effect of contract pooling on per-contract tail risk and pricing.

The central MBG result is that expected payouts are modest, while the representative-contract payout has a severe upper tail. Under the pooling approximation, diversification of contract-specific death-timing risk substantially reduces the excess-tail component and lowers the approximate price per contract. International diversification materially improves the EW–CVaR retirement-income trade-off, with the largest frontier gains when combined with the tontine overlay. Stochastic mortality has only a modest effect on the efficient frontier, while both international diversification and stochastic mortality have limited effects on representative-contract MBG pricing.

Natural extensions include modelling death as a stopping time with explicit bequest preferences; allowing the MBG terms and funding mechanism to feed back into the decumulation problem; developing a direct portfolio-level CVaR calculation accounting for cross-contract dependence; incorporating regulatory constraints and heterogeneous pool members; and testing broader asset, currency-hedging, return, and mortality specifications.

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] René Carmona and Mathieu Laurière. Deep learning for mean field games and mean field control with applications to finance. *arXiv preprint arXiv:2107.04568*, 2021.

[9] Chang Chen and Duy-Minh Dang. Convergence of neural network policies for risk–reward optimization. *arXiv preprint arXiv:2603.06563*, 2026.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] Michel Denuit, Jan Dhaene, Marc Goovaerts, and Rob Kaas. *Actuarial Theory for Dependent Risks: Measures, Orders and Models*. John Wiley & Sons, 2005.

[15] 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.

[16] Catherine Donnelly. Actuarial fairness and solidarity in pooled annuity funds. *ASTIN Bulletin: The Journal of the IAA*, 45(1):49–74, 2015.

[17] Catherine Donnelly, Montserrat Guillén, and Jens Perch Nielsen. Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, 56:14–27, 2014.

[18] 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.

[19] 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.

[20] 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.

[21] Peter A. Forsyth and George Labahn. Numerical methods for optimal decumulation of a defined contribution pension plan. *Working paper*, 2026.

[22] 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.

[23] Robert K. Fullmer. Tontines: A practitioner's guide to mortality-pooled investments, 2019. Accessed: January 2026.

[24] Jonathan T. Guyton and William J. Klinger. Decision rules and maximum initial withdrawal rates. *Journal of Financial Planning*, 19(3):48–58, 2006.

[25] Jiequn Han et al. Deep learning approximation for stochastic control problems. *arXiv preprint arXiv:1611.07422*, 2016.

[26] 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 accessed on 22 June 2026. DOI: 10.4054/HMD.Countries.20260615.

[27] 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.

[28] 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.

[29] S.G. Kou. A jump diffusion model for option pricing. *Management Science*, 48:1086–1101, August 2002.

[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] 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.

[33] Thomas Mathews. A history of australian equities. Research Discussion Paper RDP 2019-04, Reserve Bank of Australia, 2019.

[34] Moshe A Milevsky and Thomas S Salisbury. Optimal retirement income tontines. *Insurance: Mathematics and economics*, 64:91–105, 2015.

[35] 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.

[36] 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.

[37] OECD. Pension markets in focus, 2019. Accessed: January 2026.

[38] 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.

[39] 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.

[40] Wade D Pfau. An overview of retirement income planning. *Journal of Financial Counseling and Planning*, 29(1):114–120, 2018.

[41] 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.

[42] Huyên Pham and Xavier Warin. Mean-field neural networks: Learning mappings on Wasserstein space. *Neural Networks*, 168:380–393, 2023.

[43] Dimitris N Politis and Joseph P Romano. The stationary bootstrap. *Journal of the American Statistical association*, 89(428):1303–1313, 1994.

[44] Dimitris N Politis and Halbert White. Automatic block-length selection for the dependent bootstrap. *Econometric reviews*, 23(1):53–70, 2004.

[45] W Powell. A universal framework for sequential decision problems. *OR/MS Today February. https://tinyurl.com/PowellORMSfeature*, 2023.

[46] QSuper. QSuper Product Disclosure Statement for Income Account and Lifetime Pension. `https://qsuper.qld.gov.au/pds`, July 2025. Issue date: 1 July 2025.

[47] A Max Reppen, H Mete Soner, and Valentin Tissot-Daguette. Deep stochastic optimization in finance. *Digital Finance*, 5(1):91–111, 2023.

[48] Anders Max Reppen and Halil Mete Soner. Deep empirical risk minimization in finance: Looking into the future. *Mathematical Finance*, 33(1):116–145, 2023.

[49] R.T. Rockafellar and J.O. Royset. Random variables, monotone relations, and convex analysis. *Mathematical Programming*, 148:297–331, 2014.

[50] R.T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.

[51] Michael J Sabin and Jonathan Barry Forman. The analytics of a single-period tontine. *Available at SSRN 2874160*, 2016.

[52] Hersh M Shefrin and Richard H Thaler. The behavioral life-cycle hypothesis. *Economic inquiry*, 26(4):609–643, 1988.

[53] Bruno H Solnik. Why not diversify internationally rather than domestically? *Financial analysts journal*, 30(4):48–54, 1974.

[54] 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.

[55] Thinking Ahead Institute. Global pension assets study 2023. `https://www.thinkingaheadinstitute.org/research-papers/global-pension-assets-study-2023/`, 2023. Accessed: January 2026.

[56] 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.

[57] United Nations, Department of Economic and Social Affairs, Population Division. World population prospects 2024, data sources. UN DESA/POP/2024/DC/NO. 11, 2024.

[58] 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.

[59] 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.

[60] Shaun S. Wang. A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67(1):15–36, 2000.

Appendices

A NN implementation and transfer learning

Separate fully connected feedforward networks parameterize the withdrawal and rebalancing controls. The networks use the same hidden-layer architecture throughout the numerical experiments; only the dimension of the rebalancing output changes with the number of assets. The principal implementation settings are summarized in Table A.1.

Transfer learning across the γ-grid. The efficient frontiers are computed over the ordered grid $\gamma \in \{0.05, 0.10, 0.15, 0.18, 0.20, 0.50, 0.80, 1, 1.5, 2, 5, 10, 50, 10^{11}\}$. At smaller values of γ, the objective places relatively greater weight on expected withdrawals. As γ increases, greater weight is placed on lower-tail CVaR, whose empirical contribution is determined by only the lowest α fraction of terminal-wealth outcomes. Training the more tail-focused problems directly from a cold start can therefore be less stable.

Following the sequential transfer-learning procedure in [10], the initial low-γ grid points are trained from cold starts. Thereafter, the trained weights and biases at one value of γ are used to initialize the optimization at the next value. In the reported four–asset experiments, the first four grid points use cold starts, and transfer learning is applied subsequently. At every grid point, all NN parameters and the CVaR threshold W are re-optimized;

Table A.1: *Common NN architecture and principal training parameters.*

Specification	**Value**
Inputs	Standardized wealth and time-to-go
Hidden architecture	2 layers, 10 sigmoid units per layer
Output maps	Scaled sigmoid (withdrawal); softmax (rebalancing)
Optimization	Adam
Mini-batch size	1000
Initial learning rates	0.04–0.05 for the NN parameters and CVaR threshold W
Multistep schedule	Yes
Maximum iterations	30,000
Retained iterate	Best empirical objective attained during training
Training / independent test / pricing paths	$N = 256{,}000$ / 256,000 / $K = 256{,}000$

transfer learning therefore provides only the initialization and does not constrain the solution at the new value of γ. The independent out-of-sample test set is generated using a different random seed from the training sample.

B Benchmark validation

This appendix provides a benchmark validation of the NN implementation by reproducing the synthetic-market efficient frontiers of [22]. That paper solves the two-asset, deterministic-mortality EW–CVaR decumulation problem, with and without a tontine overlay, using a dynamic-programming/PIDE method. The purpose of this benchmark exercise is to verify that the NN-parameterized policy approach in Section 6 reproduces the reference frontiers in a setting where a PIDE-based benchmark is available.

Accordingly, we adopt the benchmark specification of [22]: a broad equity index and a short-term government bond index, modelled by correlated double-exponential jump-diffusion processes, together with deterministic mortality from the 2014 Canadian Pensioner Mortality Table (CPM2014). In the notation of Section 3.2, CPM2014 supplies a fixed deterministic sequence of age-specific one-year conditional death probabilities $\{\delta_{m-1}\}_{m=1}^{M}$; hence $\delta_{m-1}^{(n)} \equiv \delta_{m-1}$ on every training path. The full synthetic-market dynamics and calibration are reported below for reproducibility.

B.1 Benchmark scenario

The base-case retirement scenario mirrors the specification in Section 11 and Table 3 of [22]. Unless otherwise stated, monetary quantities in this appendix are expressed in thousands of real dollars, where "real" means inflation-adjusted. A concise summary is given in Table B.1.

In this setting, the investor controls annual withdrawals $q_m \in [q_{\min}, q_{\max}]$ and the equity fraction $p_m \in [0, 1]$. If required withdrawals exhaust the account, the portfolio is liquidated and subsequent withdrawals are financed as debt growing at the bond rate plus the borrowing spread μ_b^c. Mortality affects the account only through deterministic mortality credits based on CPM2014, so all synthetic-market training paths share the same mortality profile.

B.2 Synthetic-market asset dynamics

We restrict attention to the two domestic assets used in [22]: a broad domestic equity index and a constant-maturity domestic short-term government bond index. The benchmark uses the same correlated double-exponential jump-diffusion model as [22, Section 4].

Let $\{S_t\}_{0\le t\le T}$ denote the real (inflation-adjusted) value invested in the domestic stock index at time t, and let $\{B_t\}_{0\le t\le T}$ denote the real (inflation-adjusted) value invested in the domestic bond index. In the absence of investor actions, both assets evolve according to correlated double-exponential jump-diffusion processes.

Let ξ_s be the stock-index jump multiplier, so that a jump at time t changes the index value according to $S_t = \xi_s S_{t^-}$. We assume $\ln(\xi_s)$ has the asymmetric double-exponential density

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

[8] This matches the continuously charged 0.5% per annum tontine fee used in [22] under annual decision times.

TABLE B.1: *Base-case retirement scenario used for benchmark validation, matching the specification in [22]. 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	$t = 0, 1, \ldots, 29$ (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
Mortality-credit rule	Homogeneous large-pool representative-member rule in (3.2)–(3.3) (Section 3.1)
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	Table B.2

Between rebalancing times, and in the absence of active control, the stock index process evolves as

$$\frac{dS_t}{S_{t^-}} = (\mu_s - \lambda_s \, \kappa_s) \, dt + \sigma_s \, dZ_s(t) + d\left(\sum_{i=1}^{\pi_s(t)} (\xi_{s,i} - 1)\right), \quad t \in [t_{m-1}^+, t_m^-], \quad t_{m-1} \in \mathcal{T}. \tag{B.2}$$

Here, μ_s and σ_s are the real (inflation-adjusted) drift and instantaneous volatility, $\{Z_s(t)\}_{t\in[0,T]}$ is a standard Brownian motion, and $\{\pi_s(t)\}_{0\le t\le T}$ is a Poisson process with intensity $\lambda_s > 0$. The jump amplitudes $\{\xi_{s,i}\}_{i=1}^{\infty}$ are i.i.d. with the same distribution as ξ_s. The compensated drift adjustment is

$$\kappa_s = \mathbb{E}[\xi_s - 1] = \frac{\zeta_s \, \eta_{s,1}}{\eta_{s,1} - 1} + \frac{(1 - \zeta_s) \, \eta_{s,2}}{\eta_{s,2} + 1} - 1. \tag{B.3}$$

The Brownian motion, Poisson process, and jump multipliers are mutually independent.

Similarly, between rebalancing times, the bond index process evolves as

$$\frac{dB_t}{B_{t^-}} = \left(\mu_b - \lambda_b \kappa_b + \mu_b^c \mathbf{1}_{\{B_{t^-}<0\}}\right) dt + \sigma_b \, dZ_b(t) + d\Bigg(\sum_{i=1}^{\pi_b(t)} (\xi_{b,i} - 1)\Bigg), \quad t \in [t_{m-1}^+, t_m^-], \quad t_{m-1} \in \mathcal{T}. \tag{B.4}$$

The parameters in (B.4) are defined analogously to those in (B.2). The term $\mu_b^c \mathbf{1}_{\{B_{t^-}<0\}}$ applies the borrowing spread $\mu_b^c \ge 0$ when the bond amount is negative. The bond jump multiplier ξ_b satisfies

$$\varphi_b(y) = \zeta_b \, \eta_{b,1} \, e^{-\eta_{b,1} y} \, \mathbf{1}_{y\ge 0} + (1 - \zeta_b) \, \eta_{b,2} \, e^{\eta_{b,2} y} \, \mathbf{1}_{y<0}, \quad \zeta_b \in [0, 1], \quad \eta_{b,1} > 1, \quad \eta_{b,2} > 0. \tag{B.5}$$

The diffusion components are correlated, $dZ_s(t) \, dZ_b(t) = \rho_{sb} \, dt$, while the jump processes and jump sizes are mutually independent and independent of the Brownian motions.

B.3 Synthetic-market calibration

Following [22], the synthetic market is calibrated to monthly real (inflation-adjusted) total-return series for the CRSP value-weighted equity index and the CRSP US 30-day T-bill index over 1926:1–2020:12, with both series deflated by the US CPI. We do not re-estimate these parameters. Instead, the benchmark validation uses the annualized real (inflation-adjusted) parameter values reported in Table 1 of [22], reproduced in Table B.2.

TABLE B.2: *Annualized real (inflation-adjusted) parameter values for the double-exponential jump-diffusion model, taken from [22]. 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

B.4 Validation results

We validate our NN implementation by reproducing the synthetic-market EW–CVaR efficient frontiers reported in [22], both with and without a tontine overlay. Figure B.1 compares the PIDE-based reference frontiers from [22] with the NN frontiers obtained from our training procedure, together with the constant-withdrawal/constant-allocation benchmark.

For $\alpha = 0.05$, the horizontal axis reports the lower-tail $\text{CVaR}_{0.05}(W_T)$, while the vertical axis reports the expected annualized withdrawal $\mathbb{E}[\sum_{m=0}^{M} q_m]/T$; larger values on both axes are preferred. The constant-policy benchmark uses an annual withdrawal of 40 and a constant equity allocation of 10%, with the remaining wealth invested in bonds and with no tontine overlay.

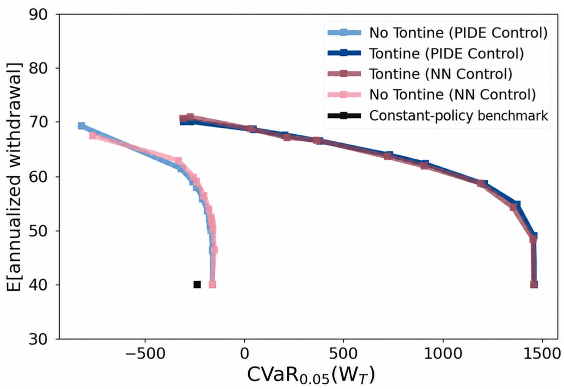

FIGURE B.1: *Validation of the EW–CVaR efficient frontiers in the synthetic market. Reference frontiers with and without a tontine overlay, labelled respectively as "No Tontine (PIDE Control)" and "Tontine (PIDE Control)", are compared with the corresponding NN frontiers, together with the constant withdrawal/constant allocation benchmark. Units: thousands of real dollars.*

Overall, the NN approach produces EW–CVaR efficient frontiers in close agreement with the PIDE-based reference frontiers and preserves their qualitative structure. Across the range of scalarization parameters γ used to trace the risk–reward trade-off, the NN frontiers closely track the reference curves in both the tontine and no-tontine cases.

Beyond pointwise agreement, the NN frontiers preserve the expected qualitative behavior as γ varies: as risk aversion increases, the expected annualized withdrawal decreases, while the lower-tail $\text{CVaR}_{0.05}(W_T)$ improves. The substantial improvement in the risk–reward trade-off generated by the tontine overlay is likewise preserved.

C Subsections 9.1 and 9.2: supplementary details

This appendix provides supplementary details on data construction, mortality calibration, and NN control comparisons for the internationally diversified portfolio experiments in Section 9.

C.1 Subsection 9.1: Construction of the asset-return and CPI series

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 [33] 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 [28], which is based on micro–level Commonwealth bond data targeting a 10–year maturity [11]. 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 B.3, 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.

C.2 Subsection 9.2: Mortality data and calibration details

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 $D_{a,y}$ and central exposures-to-risk $E^c_{a,y}$ for ages $a = 55, \dots, 95$ and calendar years $y = 1980, \dots, 2021$ from the HMD tables `Deaths_1x1.txt` and `Exposures_1x1.txt`.[9]

Following [59], the central exposures are converted to initial exposures using $E^0_{a,y} \approx E^c_{a,y} + 0.5D_{a,y}$, with empirical one–year death probabilities given by $q_{a,y} = D_{a,y}/E^0_{a,y}$. Both LC and CBD are then estimated under the Binomial–logit specification: $D_{a,y} \sim \text{Binomial}\big(E^0_{a,y}, q_{a,y}\big)$, with $\operatorname{logit} q_{a,y} = \eta_{a,y}$. Consequently, the fitted and simulated model outputs are one–year death probabilities directly; no conversion from central death rates to death probabilities is applied.

These historical mortality series are used to calibrate the LC and CBD models from the generalized age–period–cohort (GAPC) family, following Subsection 3.2. Calibration and forecasting are implemented in R using the StMoMo package [59], 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 [59] for details.

For each model, the projected period factors are simulated using the random-walk-with-drift specifications in (3.7) and (3.9). We generate 256,000 stochastic mortality sequences using random seed 3. Each sequence is extracted diagonally from the projected age–calendar-year surface, beginning with $q_{65,2022}$ and continuing with $q_{66,2023}, q_{67,2024}, \dots$. The simulation output contains 31 probabilities corresponding to ages 65–95; the first $M = 30$ probabilities, through $q_{94,2051}$, enter the mortality-credit and death-time calculations over the 30-year horizon.

D Subsection 9.5: Mortality-model robustness

For each scalarization parameter γ, one set of learned NN controls is trained under LC mortality and another under CBD mortality. Both sets are then evaluated using the same out-of-sample asset, CPI, and CBD mortality paths, with the LC-trained controls applied without retraining. This common evaluation specification isolates sensitivity to the mortality model used during training. For the representative-contract MBG comparison, we use $\gamma = 1.5$, $L_0 = 1000$, $\alpha_g = 5\%$, $\lambda = 0.05$, and $\beta_0 = 0.05$.

Figure D.1 shows that the two frontiers are nearly coincident. Table D.1 similarly shows differences of less than one thousand real AUD in both the expected MBG payout and its upper-tail CVaR. Because the same evaluation paths are used, the differences can also be assessed on a paired basis. The paired LC-minus-CBD differences are 0.786 (0.012) for $\widehat{\mathbb{E}}[Z_g]$, 0.810 (0.050) for $\widehat{\text{CVaR}}^{+}_{0.05}(Z_g)$, 8.26×10^{-4} (1.31×10^{-5}) for $\widehat{f}_g$, and -0.413 (0.007) basis points for the implied starting payment rate, where Monte Carlo standard errors are shown in parentheses. Thus, although the small differences are estimated precisely, they are economically negligible. These results support

[9] Available 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`.

robustness of both the retirement-income frontiers and the representative-contract MBG valuation to the choice between LC and CBD as the training mortality model under the common CBD evaluation.

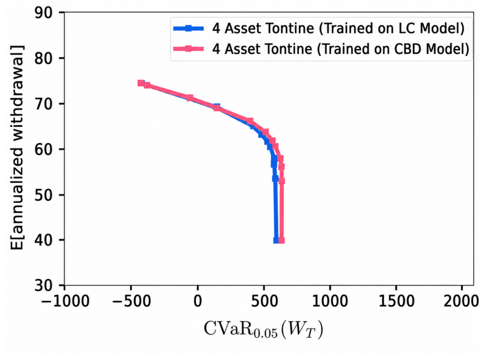

FIGURE D.1: *Mortality-model robustness of the four–asset tontine EW–CVaR frontier under out-of-sample CBD evaluation.*

TABLE D.1: *Representative-contract MBG pricing under out-of-sample CBD evaluation at* $\gamma = 1.5$.

Training model	$\widehat{\mathbb{E}}[Z_g]$	$\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g)$	$\widehat{f}_g$	$(1-\widehat{f}_g)\beta_0$
LC	75.49	760.79	0.114	4.43%
CBD	74.71	759.98	0.113	4.44%

E Subsection 9.5: Learned control heatmaps

To interpret the retirement-outcome gains from international diversification, Figure E.1 plots the learned 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 learned policy. Control heatmaps under deterministic table mortality and the CBD mortality model are qualitatively very similar and are therefore omitted for brevity.

We make several qualitative observations from Figure E.1:

- 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 E.1 (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 E.1 (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 learned policy.

  Along the 95th–percentile path the equity allocations (Figures E.1 (a) and (b)) are gradually reduced, 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 E.1 (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 learned control assigns a high weight to U.S. equity, 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 state-dependent high-growth allocation in a low-probability tail region, rather than a typical allocation across the state space. When wealth is low, the learned control allocates more to growth assets in order to improve the expected withdrawal profile, with foreign equity providing the main source of additional growth exposure.

  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 source of growth exposure 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 [21].
- Comparing Figures E.1 (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 learned control uses domestic

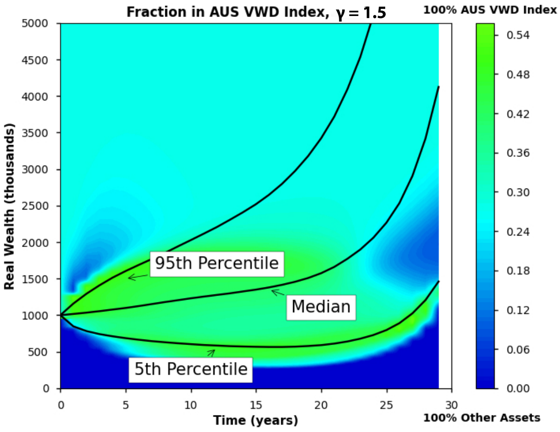

(a) Australian equity index (domestic)

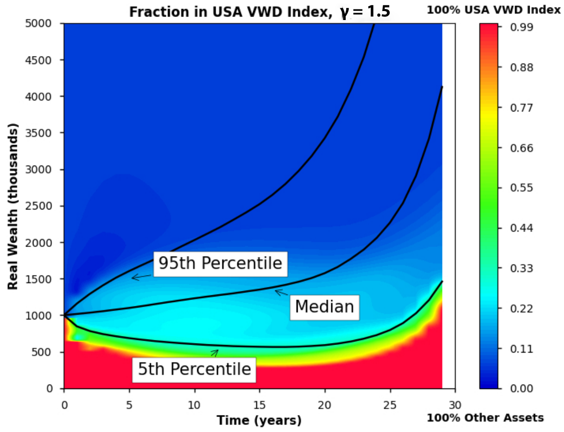

(b) U.S. equity index

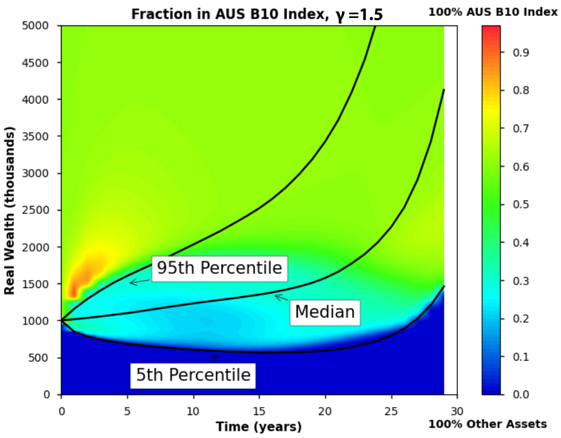

(c) Australian 10–year government bond index (domestic)

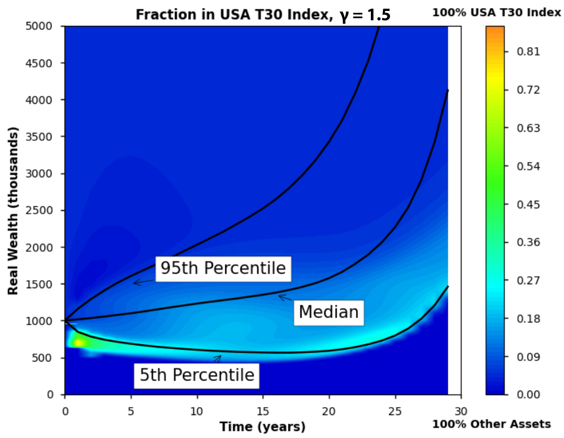

(d) U.S. 30–day T–bill index

FIGURE E.1: *Learned rebalancing controls: the four–asset tontine with stochastic mortality (LC) for $\gamma = 1.5$. Colours show the fraction of wealth invested in each asset as a function of time and real wealth.*

equity as the core growth asset, while foreign equity provides a more concentrated source of growth exposure in low–wealth regions of the state space.

Taken together, the four–asset heatmaps show that international diversification changes how equity exposure is allocated across assets and wealth states, rather than producing a uniformly higher total equity allocation. Australian equity remains the core growth asset at moderate wealth, while U.S. equity provides a more concentrated source of growth exposure in low–wealth regions.

Two– versus four–asset equity allocation. To complement the four–asset controls in Figure E.1, Figure E.2 reports the learned domestic-equity allocation for the corresponding two–asset domestic-only tontine under stochastic LC mortality.

Figure E.2 shows that, in the domestic-only case, the equity allocation is primarily wealth-dependent rather than strongly time-dependent. Along the median wealth trajectory, the learned control remains in a moderate equity range for most of the horizon, approximately 25%–45%. Near and below the 5th–percentile trajectory, however, the allocation approaches 100%, since domestic equity is the only available growth asset.

However, as shown in Figure E.1, once international diversification is allowed, the growth role performed solely by domestic equity in the two–asset setting is divided between Australian and U.S. equities. Australian equity remains part of the core growth allocation at moderate wealth, whereas U.S. equity becomes more prominent in low–wealth regions. This control comparison indicates that the four–asset frontier improvement reflects the state-dependent allocation of equity exposure across domestic and foreign markets, rather than a uniformly higher total equity exposure.

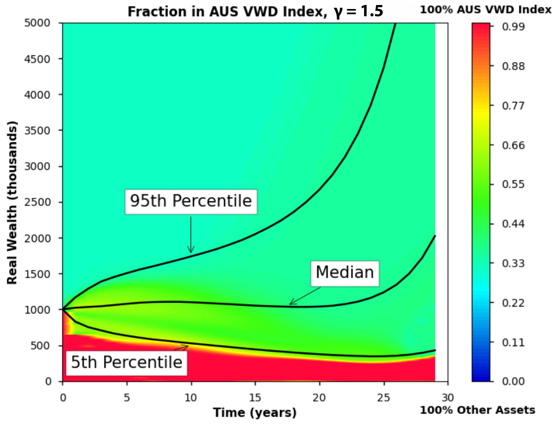

FIGURE E.2: *Learned fraction invested in domestic equity in the two–asset domestic-only tontine under stochastic mortality (LC model), for the representative frontier point $\gamma = 1.5$. Colours show the allocation as a function of time and real wealth.*

F Subsection 9.6: supplementary results

F.1 Representative-contract MBG load sensitivity

We study the sensitivity of the representative-contract 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 F.1 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 $\lambda \in \{0, 0.05, 0.5\}$. Monte Carlo standard errors are reported in Table F.3. For fixed (γ, α_g), $\widehat{f}_g$ depends linearly on λ, with slope $\widehat{\text{CVaR}}^+_{\alpha_g}(Z_g)/L_0$.

TABLE F.1: *Representative-contract MBG load sensitivity: the four–asset tontine with stochastic mortality (LC).*

(a) $\alpha_g = 5\%$

γ	$\widehat{\mathbb{E}}[Z_g]$	$\widehat{\text{CVaR}}^+_{0.05}(Z_g)$	$\widehat{f}_g$ $(\lambda = 0)$	$\widehat{f}_g$ $(\lambda = 0.05)$	$\widehat{f}_g$ $(\lambda = 0.5)$
1.5	70.69	758.28	0.071	0.109	0.450
0.5	58.13	729.37	0.058	0.095	0.423
0.2	47.51	664.39	0.048	0.081	0.380

(b) $\alpha_g = 1\%$

γ	$\widehat{\mathbb{E}}[Z_g]$	$\widehat{\text{CVaR}}^+_{0.01}(Z_g)$	$\widehat{f}_g$ $(\lambda = 0)$	$\widehat{f}_g$ $(\lambda = 0.05)$	$\widehat{f}_g$ $(\lambda = 0.5)$
1.5	70.69	917.56	0.071	0.117	0.529
0.5	58.13	915.19	0.058	0.104	0.516
0.2	47.51	871.95	0.048	0.091	0.483

Table F.1 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 $\widehat{f}_g \approx 5\%$–$7\%$ across the reported γ values (about 24–35 bps off a notional $\beta_0 = 5\%$ starting rate). At $\alpha_g = 5\%$, the illustrative value $\lambda = 0.05$ gives loads of approximately 0.081–0.109, while the stress value $\lambda = 0.5$ gives loads of approximately 0.380–0.450. Reducing the upper-tail probability from $\alpha_g = 5\%$ to $\alpha_g = 1\%$ further increases the loads to approximately 0.091–0.117 for $\lambda = 0.05$ and 0.483–0.529 for $\lambda = 0.5$.

Varying γ affects the load in a secondary, policy–induced way: changing γ changes the learned policy along the efficient frontier and thereby shifts the distribution of Z_g through the timing and magnitude of withdrawals, which determine the cumulative nominal withdrawals made before death. Over the ranges considered here, this effect is modest compared with the direct impact of (λ, α_g) in (7.5).

The sensitivity patterns in Table F.1 are not specific to the four–asset LC setting. At $\gamma = 1.5$, the two– and four–asset portfolios under both table and LC mortality produce similar upper-tail CVaR estimates and, by the linear dependence of $\widehat{f}_g$ on λ noted above, similar sensitivity to λ. Moreover, reducing α_g from 5% to 1% increases $\widehat{f}_g$ by approximately 0.008–0.009 at $\lambda = 0.05$ and by approximately 0.08–0.09 at $\lambda = 0.5$ in every setting. Detailed cross-setting results are reported in Subsection F.3.

Finally, because (λ, α_g) represent prudential choices that are generally not observable from public product disclosures, these sensitivity results should be viewed as pricing stress tests rather than as estimates of any particular provider's loading rule.

F.2 Tables 9.4-9.5 and Table F.1

All MBG pricing experiments use $K = 256{,}000$ Monte Carlo observations. Tables F.2 and F.3 report standard errors (SE) for the estimates in Tables 9.4–9.5 and Table F.1, respectively. For the illustrative post–load starting rate in Table 9.4, its standard error in percentage points is $5\,\mathrm{SE}(\widehat{f}_g)$ when $\beta_0 = 0.05$.

TABLE F.2: *Monte Carlo standard errors (SE) corresponding to estimates in Tables 9.4-9.5.*

Setting	$\mathrm{SE}(\widehat{\mathbb{E}}[Z_g])$	$\mathrm{SE}(\widehat{\mathrm{CVaR}}^{+}_{0.05})$	$\mathrm{SE}(\widehat{f}_g)$	$\mathrm{SE}(\widehat{\mathbb{P}}(Z_g > 0))$	$\mathrm{SE}(\widehat{\mathbb{E}}[Z_g \mid Z_g > 0])$
2 assets, table mortality	0.36	2.01	4.4×10^{-4}	0.08	1.3
2 assets, LC mortality	0.37	1.96	4.5×10^{-4}	0.07	1.3
4 assets, table mortality	0.36	1.96	4.4×10^{-4}	0.08	1.2
4 assets, LC mortality	0.38	1.86	4.5×10^{-4}	0.08	1.3

TABLE F.3: *Monte Carlo standard errors (SE) corresponding to estimates in Table F.1.*

(a) $\alpha_g = 5\%$

γ	$\mathrm{SE}(\widehat{\mathbb{E}}[Z_g])$	$\mathrm{SE}(\widehat{\mathrm{CVaR}}^{+}_{0.05})$	$\mathrm{SE}(\widehat{f}_g)$ $(\lambda = 0)$	$\mathrm{SE}(\widehat{f}_g)$ $(\lambda = 0.05)$	$\mathrm{SE}(\widehat{f}_g)$ $(\lambda = 0.5)$
1.5	0.38	1.86	3.8×10^{-4}	4.5×10^{-4}	1.2×10^{-3}
0.5	0.35	2.18	3.5×10^{-4}	4.4×10^{-4}	1.4×10^{-3}
0.2	0.31	2.40	3.1×10^{-4}	4.2×10^{-4}	1.5×10^{-3}

(b) $\alpha_g = 1\%$

γ	$\mathrm{SE}(\widehat{\mathbb{E}}[Z_g])$	$\mathrm{SE}(\widehat{\mathrm{CVaR}}^{+}_{0.01})$	$\mathrm{SE}(\widehat{f}_g)$ $(\lambda = 0)$	$\mathrm{SE}(\widehat{f}_g)$ $(\lambda = 0.05)$	$\mathrm{SE}(\widehat{f}_g)$ $(\lambda = 0.5)$
1.5	0.38	0.99	3.8×10^{-4}	4.0×10^{-4}	7.3×10^{-4}
0.5	0.35	1.11	3.5×10^{-4}	3.8×10^{-4}	7.7×10^{-4}
0.2	0.31	1.32	3.1×10^{-4}	3.5×10^{-4}	8.5×10^{-4}

F.3 Sensitivity across asset and mortality specifications

This subsection reports the representative-contract cross-setting sensitivity results summarized in Subsection F.1. Holding $\gamma = 1.5$ fixed, Table F.4 compares the upper-tail CVaR estimates and equivalent MBG loads for $\alpha_g = 5\%$ and $\alpha_g = 1\%$ across the two– and four–asset portfolios under both table and LC mortality, using the illustrative loading coefficient $\lambda = 0.05$ and the stress value $\lambda = 0.5$. Standard errors are shown in parentheses.

Reducing α_g from 5% to 1% increases the equivalent load by approximately 0.008–0.009 when $\lambda = 0.05$ and by approximately 0.080–0.087 when $\lambda = 0.5$ across the four settings. The similar upper-tail CVaR estimates also imply similar sensitivity to λ. Thus, international diversification changes the level of the MBG load only modestly and does not materially alter its sensitivity to either prudential-buffer parameter.

G Monte Carlo uncertainty for contract-pooling results

The estimated quantities in Table 9.6 are obtained from the pooling approximation and are functions of the representative-contract estimators $\widehat{\mathbb{E}}[Z_g]$ and $\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g)$. Using the $K = 256{,}000$ pathwise payouts for the four–asset LC base case, we estimate the variances and covariance of these two estimators from the same simulated sample. Since the quantities in (8.7) are linear functions of these estimators, the variances of the pooled quantities are computed directly from the estimated variances and covariance, thereby accounting for the dependence between the two estimators. Table G.1 reports the resulting standard errors for $\lambda = 0.05$ and $L_0 = 1000$. These Monte Carlo standard errors quantify the sampling uncertainty inherited from the representative-contract mean and CVaR estimates. They do not capture approximation error associated with the $J^{-1/2}$ pooling rule or uncertainty in the learned policy and calibrated model inputs.

TABLE F.4: *Sensitivity of representative-contract MBG pricing to the upper-tail probability across asset and mortality specifications;* $\gamma = 1.5$. *Standard errors are shown in parentheses.*

(a) $\alpha_g = 5\%$

Setting	$\widehat{\mathrm{CVaR}}^{+}_{0.05}(Z_g)$	$\widehat{f}_g$ ($\lambda = 0.05$)	$\widehat{f}_g$ ($\lambda = 0.5$)
2 assets, table mortality	736.62 (2.01)	0.102 (4.4×10^{-4})	0.434 (1.3×10^{-3})
2 assets, LC mortality	755.48 (1.96)	0.104 (4.5×10^{-4})	0.444 (1.3×10^{-3})
4 assets, table mortality	736.82 (1.96)	0.106 (4.4×10^{-4})	0.438 (1.3×10^{-3})
4 assets, LC mortality	758.28 (1.86)	0.109 (4.5×10^{-4})	0.450 (1.2×10^{-3})

(b) $\alpha_g = 1\%$

Setting	$\widehat{\mathrm{CVaR}}^{+}_{0.01}(Z_g)$	$\widehat{f}_g$ ($\lambda = 0.05$)	$\widehat{f}_g$ ($\lambda = 0.5$)
2 assets, table mortality	910.94 (1.07)	0.111 (3.8×10^{-4})	0.521 (7.6×10^{-4})
2 assets, LC mortality	916.91 (0.96)	0.113 (3.9×10^{-4})	0.525 (7.1×10^{-4})
4 assets, table mortality	911.07 (1.12)	0.115 (3.9×10^{-4})	0.525 (7.8×10^{-4})
4 assets, LC mortality	917.56 (0.99)	0.117 (4.0×10^{-4})	0.529 (7.3×10^{-4})

TABLE G.1: *Monte Carlo standard errors corresponding to the estimated approximate quantities in Table 9.6. Monetary quantities are in thousands of real AUD.*

J	$\mathrm{SE}\left(\widehat{\mathrm{CVaR}}^{+,\mathrm{a}}_{0.05}(\overline{Z}_{g,J})\right)$	$\mathrm{SE}(\widehat{P}^{\mathrm{a}}_{g,J})$	$\mathrm{SE}(\widehat{f}^{\mathrm{a}}_{g,J})$
1	1.857	0.451	4.51×10^{-4}
10	0.803	0.411	4.11×10^{-4}
50	0.551	0.402	4.02×10^{-4}
100	0.496	0.399	3.99×10^{-4}
500	0.426	0.396	3.96×10^{-4}
1000	0.410	0.396	3.96×10^{-4}
10,000	0.386	0.394	3.94×10^{-4}
20,000	0.383	0.394	3.94×10^{-4}
$J \to \infty$	0.375	0.394	3.94×10^{-4}