

Spouse-Protected Tontines: Household Decumulation via Neural-Network Optimization

Duy-Minh Dang*

Yukan Perumal†

September 30, 2026

Abstract

We develop a spouse-protected tontine in which a first death changes the household state but generates no pool transfer. The same account remains attached to the household contract until extinction and funds a spouse-only continuation phase if the retiree dies first. We derive contract-level actuarial-fairness conditions and a finite-pool mortality-credit allocation rule with exact ex post budget balance. Under a homogeneous large-pool approximation, we formulate a multidimensional decumulation problem for a representative household, with withdrawal and rebalancing controls while the retiree is alive. The objective balances expected cumulative real household payments against the Conditional Value-at-Risk (CVaR) of terminal shortfalls relative to household reserve targets, without conditioning on survival to the horizon. We develop a numerical solution method for this problem based on a global-in-time neural-network parameterization of admissible control policies. We quantify policy-induced spouse-continuation costs using expected-cost and risk-loaded payment-scale loads at both representative-contract and book levels. We characterize the large-book limit of average per-contract continuation cost as the expected representative-contract cost conditional on common market information.

Numerical experiments calibrated to Australia show that the spouse-only phase is a first-order and persistent household event. In these experiments, book-level diversification removes most of the representative-contract excess upper-tail cost of spouse continuation in the large-book limit, without changing expected per-contract continuation cost. A cross-country mortality comparison indicates that the prevalence and persistence of the spouse-only phase are not specific to Australia.

Keywords: spouse-protected tontines; household decumulation; longevity pooling; mortality credits; spouse-continuation costs; neural-network policies; conditional value-at-risk.

AMS Subject Classification: 91G05, 91G10, 91G60, 91G70, 93E20, 68T07

1 Introduction

Tontines pool idiosyncratic longevity risk by redistributing the account balances of contracts that exit the pool to contracts that remain active. The resulting mortality credits can enhance retirement income while retaining greater investment and withdrawal flexibility than a conventional life annuity [11, 12, 16, 17, 26]. Many modern tontine optimization models nevertheless take the pooled unit to be a single retiree who purchases the contract: the contract exits when the retiree dies, and no household payment is made thereafter.

Spouse protection addresses a different retirement-income need. Commercial lifetime-income products in Australia and other markets, such as Canada, the UK, and the US, can offer joint-life options under which payments continue for a surviving spouse or partner [1, 23, 27, 33, 34, 41]. QSuper’s Lifetime

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

†School of Mathematics and Physics, The University of Queensland, St Lucia, Brisbane 4072, Australia, email: y.perumal@student.uq.edu.au

Pension provides a pooled example: its spouse-protection option continues payments for the life of the last survivor, bases the initial payment rate on the younger spouse’s age, and applies a lower payment rate than the single option [33]. These product terms specify the last-survivor income benefit but leave open, for a tontine design, how a spouse-protected contract should participate in the pool, when its balance should be forfeited, and how mortality-credit gains should depend on household survival status.

This paper draws on two related but largely separate literatures. On the one hand, modern tontine and pooled-annuity research develops actuarially fair pooling rules and optimal retirement-income designs, predominantly for single-life participants [3, 11, 16, 26, 28]. On the other hand, research on retired couples and joint-life products recognizes that consumption, portfolio choice, survivor income, and annuity values depend on both lifetimes [20, 38, 44]. Related work develops longevity-credit mechanisms for joint-life benefits [13] and extends longevity pooling to multi-life benefits, with large-pool convergence results for survivor-only credits [14]. To our knowledge, existing studies do not jointly address last-survivor tontine pooling, controlled household withdrawals and portfolio choice under investment risk, and the resulting policy-induced spouse-continuation cost and its book-level tail risk.

We develop a spouse-protected tontine in which each pooled unit is a household contract that remains active until household extinction. Under this design, pool-exit probabilities and mortality-credit gain rates are household-state dependent. While both members are alive, the contract exits during a period only if both members die during that period; if the retiree dies first and the spouse survives, the contract enters a spouse-only continuation phase and remains active in the pool. Crucially, a first death changes the household state but generates no forfeiture or other pool transfer; the existing account remains attached to the household contract until extinction. To isolate this household-state mechanism, we consider a homogeneous pool in which every contract includes spouse protection; mixed pools containing both single-life and spouse-protected contracts are outside the present scope. Throughout, *spouse protection* denotes the overall product-design feature, whereas *spouse continuation* denotes the spouse-only payment phase or stream that may follow the retiree’s death.

The household contract has four survival states: both members alive, retiree only, spouse only, and neither alive. We derive the state-dependent pool-exit probabilities, conditions for individual-contract actuarial fairness, exact ex post budget balance at the finite-pool level, and the corresponding large-pool tontine gain rates. Withdrawals and portfolio rebalancing are controlled while the retiree remains alive; if the retiree dies first, the surviving spouse receives a capped, rule-based spouse-continuation payment and the portfolio follows a passive rebalancing rule. Household reward is expected cumulative real household payments, while risk is the CVaR of a terminal shortfall measured against household-state-dependent reserve targets. Both criteria are evaluated without conditioning on survival to the horizon.

Under the homogeneous large-pool approximation, the resulting decumulation problem for a representative household is a multidimensional stochastic control problem. It combines a four-asset investment universe comprising domestic and foreign equity and bond indices with discrete household mortality states, household-state-dependent mortality-credit gain rates and payment rules, insolvency, and a nonlinear tail-risk objective. We develop a numerical solution method for this problem based on a global-in-time neural-network (NN) parameterization of admissible withdrawal and asset-allocation control policies. Admissibility is enforced through the network output transformations. This approach follows the NN policy literature for multi-period portfolio and decumulation problems [7, 19, 24, 28]; recent convergence theory provides convergence-in-probability results for NN policy approximations in constrained risk–reward control problems [6].

We quantify the spouse-continuation payment obligation induced by the learned household policy through expected-cost and risk-loaded payment-scale loads. We then distinguish representative-contract tail risk from contract-book risk. In a finite book, contracts share market returns but have conditionally independent household mortality histories given common market information. We characterize the large-book limit of average per-contract continuation cost as the expected representative-contract cost

conditional on this information. Household-specific mortality risk therefore diversifies across contracts, while common-market risk remains. Expected per-contract continuation cost is unchanged by pooling, whereas the prudential tail buffer must reflect book-level risk.

During the preparation of this paper, we became aware of the contemporaneous work of [4], who develop a fully funded mutual couple tontine with dependent partner lifetimes and actuarially fair mortality and widowing credits. Withdrawals are prescribed to achieve a target average income together with random credits, rather than optimized jointly with portfolio choice under investment risk. First death releases reserves for redistribution as widowing credits to couples joint-alive at the beginning of the period.

We study a different same-account continuation design: first death generates no forfeiture or other pool transfer, and the existing account remains attached to the household contract until extinction. This convention preserves the link between retiree-alive withdrawal and investment decisions and the subsequent spouse-continuation stream. Our focus is therefore on how a household should decumulate and invest when the same account may later support a surviving spouse, how large the resulting policy-induced spouse-continuation cost is, and how its tail risk changes with book size. This book-level analysis concerns cumulative continuation payments under common market risk, rather than the large-pool behavior of individual longevity credits analyzed in [14].

The contributions of the paper are threefold.

- *Household-contract tontine design.* We develop a spouse-protected tontine in which a first death changes the household state but generates no forfeiture or other pool transfer, and the same account remains attached to the household until household extinction. We derive the associated household-state pool-exit probabilities, contract-level actuarial-fairness conditions for mortality-credit allocation, and a finite-pool allocation rule with exact ex post budget balance.
- *Household reward–risk optimization.* We formulate a decumulation problem under investment risk that distinguishes controlled retiree-alive withdrawals from rule-based spouse-continuation payments in the spouse-only state and solve it using admissible global-in-time NN policies.
- *Continuation-cost measurement and book-level diversification.* We define representative-contract and book-level expected-cost and risk-loaded spouse-continuation loads, characterize the conditional large-book limit, and quantify the reduction in excess upper-tail cost through household-specific mortality diversification.

The numerical experiments use the 2021 Australian male and female period life tables from the Human Mortality Database [21]. The baseline household consists at inception of a male retiree aged 65 and a female spouse aged 60 and is followed for 35 years, until the spouse reaches age 95. The initial contract balance is AUD 1 million. The Australian calibration yields three main findings:

- The spouse-only continuation phase is a first-order event: approximately 71% of contracts enter this state, which lasts about 12.25 years on average conditional on entry.
- Under capped continuation, expected representative-contract continuation payments remain substantial and relatively stable as the evaluated policies vary along the reward–risk frontier. The expected-cost spouse-continuation load nevertheless increases as policies become more conservative.
- Contract-book diversification leaves expected continuation cost unchanged but sharply reduces its excess upper tail. Under the baseline capped design, the conditional large-book limit reduces the representative-contract excess upper-tail cost by approximately 97% for the policies evaluated along the reward–risk frontier.

The Australian household-state pattern is not exceptional. Under the same household ages, 35-year horizon, and independent-mortality specification, the corresponding 2021 male and female period life tables for Canada, the United States, Japan, and the Netherlands imply that, across the five countries, the probability of entering the spouse-only phase ranges from approximately 68.1% to 77.0%. Its unconditional mean duration ranges from 8.70 to 10.23 years, while its mean duration conditional on

entry ranges from 12.25 to 13.34 years. The spouse-only continuation phase is therefore a first-order and persistent household event across these mortality environments.

The remainder of the paper is organized as follows. Section 2 develops the household-state tontine and mortality-credit mechanics. Section 3 specifies the household payment, portfolio, and insolvency framework. Section 4 formulates the household reward–risk problem, and Section 5 presents the neural-network policy approximation. Section 6 defines the representative-contract and book-level spouse-continuation loads and develops the conditional large-book diversification framework. Numerical experiments and their implications for spouse-protected tontine design are reported in Section 7. Section 8 concludes the paper and outlines some possible research directions. Appendix D reports the cross-country household-state comparison. Appendix G examines the effects of removing the continuation cap on expected continuation costs and tail risk.

2 Tontine modeling with spouse-protection

We build on and generalize the individual-account tontine framework of [16] to a spouse-protected household-contract setting. At inception, each pooled unit is a *spouse-protected household contract* comprising a retiree and spouse. There is no mixed entry pool: all pooled units enter as spouse-protected household contracts, rather than as a mixture of spouse-protected and single-life contracts. As retiree and spouse deaths occur, a contract may move from the two-member state to a one-survivor state, but it remains the same contract until it exits the pool. Hereafter, these pooled units are simply referred to as contracts unless otherwise noted.

The central modeling distinction is that mortality pooling is based on contract exit from the pool, not on the retiree’s death alone. If the retiree dies while the spouse remains alive, the contract continues in the pool and household payments may continue under the spouse-protection rule. Mortality credits are therefore determined by the household survival state of each contract and are generated only when a contract exits the pool. In this section, we focus on the mortality-pooling mechanics of the spouse-protected tontine.

2.1 Preliminaries

Unlike the “plan-to-live” convention in retirement planning [16, 28, 30], we do not condition on household survival to the horizon. Mortality is incorporated through the household-state process introduced below. The finite horizon is therefore a computational truncation horizon rather than an assumption that either the retiree or the spouse survives to that date. Once neither household member remains alive, household payments and mortality credits cease.

Let a_R and a_S denote the retiree’s and spouse’s ages at inception. Since spouse continuation may remain relevant after the retiree’s death, choosing the horizon solely from the retiree’s age can mechanically truncate the spouse-only phase when the spouse is younger. To avoid this, we select a terminal reference age x^* for the younger of the retiree and spouse and set

$$T = \lceil x^* - \min\{a_R, a_S\} \rceil.$$

This convention does not condition on survival of either household member to the reference age x^* .

We consider the finite planning horizon $[0, T]$ with annual decision times

$$\mathcal{T} = \{t_m \mid t_m = m, m = 0, \dots, M\}, \quad M = T,$$

at which mortality credits are allocated, fees are deducted, household payments are made, and portfolio positions are rebalanced.¹ Here, $t_0 = 0$ is the inception date. Household payments and rebalancing decisions are made at t_m , $m = 0, \dots, M - 1$, while mortality credits are allocated at t_m^- , $m = 1, \dots, M$.

¹Annual spacing is assumed for simplicity. In practice, review and rebalancing dates are typically fixed in advance, e.g. annually or semi-annually.

For notational convenience, let $t^- = t - \varepsilon$ and $t^+ = t + \varepsilon$ denote the instants immediately before and after time t , with $\varepsilon \rightarrow 0^+$. For a generic time-dependent quantity $f(t)$ and any $t_m \in \mathcal{T}$, we write

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

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

The tontine pool consists of L contracts, indexed by $\ell = 1, \dots, L$. A contract is called *active* at time t if at least one household member is alive at t .

During the inter-decision period $[t_{m-1}^+, t_m^-]$, asset prices evolve and mortality outcomes are realized. At t_m^- , these outcomes are observed, after which the period- m decision and implementation sequence is

$$t_m^- \longrightarrow t_m \longrightarrow t_m^+.$$

For contracts active at the start of the period, the events are ordered as follows.

- At t_m^- , each contract's portfolio balance is observed and its household survival state is updated. Any contract for which neither member remains alive exits the pool and forfeits its balance.

The forfeited balances are redistributed immediately as mortality credits to the contracts that remain active after the mortality update, subject to the pool-closing convention in Subsection 2.5.

- At t_m , $m = 0, \dots, M - 1$, the period- m household payment is made. When the retiree is alive, this payment is the controlled withdrawal; when only the spouse is alive, it is the spouse-protection continuation payment, determined by the continuation rule rather than by a new withdrawal decision.

At t_m^+ , the post-payment portfolio is then updated: contracts for which the retiree remains alive and post-payment wealth is strictly positive may rebalance; if post-payment wealth is nonpositive, insolvency is triggered; and spouse-only contracts follow the passive rebalancing rule.

At the terminal date $t_M = T$, the portfolio is liquidated and no further household payment or rebalancing decision is made.

To formalize mortality credits, we denote by $\{H_t^\ell\}_{0 \leq t \leq T}$ the household survival-state process for contract ℓ , where H_t^ℓ records at time t whether the retiree and spouse are alive. In the period-by-period recursion, the relevant quantities are the pre- and post-decision states

$$H_{m-}^\ell := H_{t_m^-}^\ell, \quad H_{m+}^\ell := H_{t_m^+}^\ell, \quad \ell = 1, \dots, L,$$

with values in

$$\mathcal{H} := \{11, 10, 01, 00\}.$$

The two digits indicate the survival status of the retiree and spouse, respectively:

$$11 : \text{both alive}, \quad 10 : \text{retiree only}, \quad 01 : \text{spouse only}, \quad 00 : \text{neither alive}.$$

Since mortality is updated at t_m^- and no further mortality event occurs between t_m^- and t_m^+ , we have $H_{m+}^\ell = H_{m-}^\ell$. Contract ℓ remains active in the pool after the mortality update at t_m^- if and only if $H_{m-}^\ell \neq 00$; if $H_{m-}^\ell = 00$, it exits the pool at t_m^- .

2.2 Spouse protection

Spouse protection changes the pooled unit from a single retiree to a household contract. The retiree's death alone does not cause pool exit: the contract remains active in states 11, 10, and 01, and exits only when $H_{m-}^\ell = 00$. Consequently, spouse protection changes the household-state-dependent pool-exit probabilities and hence the corresponding mortality-credit gain rates. In particular, while both household members are alive, pool exit is governed by their joint-death probability.

The payment convention distinguishes the retiree-alive control states from the spouse-only continuation state. In states 11 and 10, the household payment is a controlled withdrawal and the portfolio may be rebalanced. In state 01, no new withdrawal or rebalancing control is chosen; the surviving spouse instead receives a rule-based *spouse-continuation payment*, and the portfolio follows the passive

update convention. The precise payment, continuation-payment cap, admissibility, and insolvency rules are specified in Section 3.

The amount of the spouse-continuation payment does not enter the current-period mortality-credit gain rate directly; it affects post-payment wealth and therefore the subsequent contract evolution. The resulting spouse-only payment stream is quantified later through the policy-induced spouse-continuation load.

2.3 Household-state transitions and contract exit

Let $\delta_{R,m-1}^\ell$ and $\delta_{S,m-1}^\ell$ denote the conditional one-year death probabilities of the retiree and spouse in contract ℓ over the period $[t_{m-1}^+, t_m^-]$, given survival at t_{m-1}^+ . Throughout this subsection, we assume conditional independence of retiree and spouse mortality within each period. The resulting one-period household-state transitions are

$$\begin{aligned}\mathbb{P}\left(H_{m-}^\ell = 11 \mid H_{(m-1)+}^\ell = 11\right) &= (1 - \delta_{R,m-1}^\ell)(1 - \delta_{S,m-1}^\ell), \\ \mathbb{P}\left(H_{m-}^\ell = 10 \mid H_{(m-1)+}^\ell = 11\right) &= (1 - \delta_{R,m-1}^\ell)\delta_{S,m-1}^\ell, \\ \mathbb{P}\left(H_{m-}^\ell = 01 \mid H_{(m-1)+}^\ell = 11\right) &= \delta_{R,m-1}^\ell(1 - \delta_{S,m-1}^\ell), \\ \mathbb{P}\left(H_{m-}^\ell = 00 \mid H_{(m-1)+}^\ell = 11\right) &= \delta_{R,m-1}^\ell\delta_{S,m-1}^\ell,\end{aligned}\tag{2.1}$$

together with

$$\mathbb{P}\left(H_{m-}^\ell = 10 \mid H_{(m-1)+}^\ell = 10\right) = 1 - \delta_{R,m-1}^\ell, \quad \mathbb{P}\left(H_{m-}^\ell = 00 \mid H_{(m-1)+}^\ell = 10\right) = \delta_{R,m-1}^\ell, \tag{2.2}$$

$$\mathbb{P}\left(H_{m-}^\ell = 01 \mid H_{(m-1)+}^\ell = 01\right) = 1 - \delta_{S,m-1}^\ell, \quad \mathbb{P}\left(H_{m-}^\ell = 00 \mid H_{(m-1)+}^\ell = 01\right) = \delta_{S,m-1}^\ell. \tag{2.3}$$

State 00 is absorbing and corresponds to permanent exit of the contract from the pool.

Let

$$\mathcal{L}_{m-1} := \left\{ \ell \in \{1, \dots, L\} : H_{(m-1)+}^\ell \neq 00 \right\}$$

denote the set of contracts active at time t_{m-1}^+ . For each $\ell \in \mathcal{L}_{m-1}$, the one-period pool-exit probability over $[t_{m-1}^+, t_m^-]$ is given by

$$\bar{\delta}_{m-1}^\ell := \mathbb{P}\left(H_{m-}^\ell = 00 \mid H_{(m-1)+}^\ell\right) = \begin{cases} \delta_{R,m-1}^\ell \delta_{S,m-1}^\ell, & H_{(m-1)+}^\ell = 11, \\ \delta_{R,m-1}^\ell, & H_{(m-1)+}^\ell = 10, \\ \delta_{S,m-1}^\ell, & H_{(m-1)+}^\ell = 01. \end{cases} \tag{2.4}$$

Thus, a contract in state 11 exits the pool during the period only if both members die before t_m^- . A contract already in state 10 or 01 exits during the period if the surviving household member dies before t_m^- .

For each $\ell \in \mathcal{L}_{m-1}$, let $\mathcal{W}_{m-}^\ell$ denote the real (inflation-adjusted) account balance of contract ℓ at time t_{m-}^- , before mortality credits are distributed; hereafter, we also refer to this account balance as the wealth of contract ℓ . The mortality-credit allocation in this section is formulated for solvent tontine accounts, so we assume $\mathcal{W}_{m-}^\ell > 0$ for all $\ell \in \mathcal{L}_{m-1}$. Once an account is depleted, it receives no further mortality credits, and subsequent minimum household payments are financed from other assets or sources of finance under the insolvency convention in Subsection 3.6. We define the post-mortality-update contract-active indicator at t_{m-}^- and its conditional expectation by

$$\mathbf{1}_m^\ell = \begin{cases} 1, & H_{m-}^\ell \neq 00, \\ 0, & H_{m-}^\ell = 00, \end{cases} \quad \mathbb{E}_{m-1}[\mathbf{1}_m^\ell] = 1 - \bar{\delta}_{m-1}^\ell, \quad \ell \in \mathcal{L}_{m-1}. \tag{2.5}$$

Here, $\mathbb{E}_{m-1}[\cdot]$ denotes expectation conditional on all information available at t_{m-1}^+ (the start of $[t_{m-1}^+, t_m^-]$). If $\mathbf{1}_m^\ell = 0$, the contract exits the pool and its account balance is forfeited. If $\mathbf{1}_m^\ell = 1$, the contract remains active in the pool and receives a mortality credit c_m^ℓ at time t_m^- .

Remark 2.1 (Pool exit and mortality credits). *For spouse-protected contracts, the relevant event for mortality pooling is contract exit from the pool, not the retiree's death alone. The cases in (2.4) therefore encode the relevant one-period pool-exit probability for each contract state. Consequently, the conditional probability entering the actuarially fair mortality credit formulas is the pool-exit probability $\bar{\delta}_{m-1}^\ell$ for contract ℓ in (2.4), rather than the retiree death probability alone.*

We next derive the mortality-credit allocation associated with this pool-exit probability, proceeding from individual-contract fairness, through the pool-level budget constraint and large-pool approximation, to the contract-level tontine gain rate used in the wealth recursion. Throughout this derivation, we abstract from administration, investment-management, and transaction costs; fees are introduced separately in the next section.

2.4 Individual-contract fairness condition

The fairness condition is imposed at the contract level. During $[t_{m-1}^+, t_m^-]$, a contract that exits the pool forfeits its account balance, while a contract that remains active is eligible to receive a mortality credit. Let Ω_m be the sigma-field generated by the information available at t_{m-1}^+ and the realized balances $\{\mathcal{W}_{m-}^k\}_{k \in \mathcal{L}_{m-1}}$, before forfeiture and mortality-credit allocation. We assume that current-period mortality outcomes and asset returns are conditionally independent given the start-of-period information, so $\mathbb{P}(\mathbf{1}_m^\ell = 1 \mid \Omega_m) = 1 - \bar{\delta}_{m-1}^\ell$. For $\bar{\delta}_{m-1}^\ell < 1$, individual-contract fairness balances expected forfeitures against expected credit receipts conditional on Ω_m :

$$\bar{\delta}_{m-1}^\ell \mathcal{W}_{m-}^\ell = (1 - \bar{\delta}_{m-1}^\ell) \mathbb{E}[c_m^\ell \mid \Omega_m, \mathbf{1}_m^\ell = 1], \quad (2.6)$$

where the expectation conditions on contract ℓ remaining active, but not on the current-period mortality outcomes of the other contracts. Solving (2.6) gives the actuarial-fairness benchmark

$$\mathbb{E}[c_m^\ell \mid \Omega_m, \mathbf{1}_m^\ell = 1] = \frac{\bar{\delta}_{m-1}^\ell}{1 - \bar{\delta}_{m-1}^\ell} \mathcal{W}_{m-}^\ell. \quad (2.7)$$

Thus, the expected mortality credit is proportional to the contract's own wealth, with the proportionality factor determined by the contract's pool-exit probability. Substituting the pool-exit probabilities in (2.4) into the general fairness condition (2.7) gives the state-specific fair mortality-credit factor using $\delta_{R,m-1}^\ell \delta_{S,m-1}^\ell$ in state 11, $\delta_{R,m-1}^\ell$ in state 10, and $\delta_{S,m-1}^\ell$ in state 01. In particular, since the fair-credit multiplier $p/(1-p)$ increases with the pool-exit probability p , the state-11 factor is no greater than either corresponding one-survivor factor.

Remark 2.2 (Finite-pool budget balance and actuarial fairness). *Equation (2.7) is a contract-level ex ante actuarial-fairness condition, not an ex post budget identity. In a finite pool, actual mortality credits must be funded by realized forfeitures and therefore depend on the realized pool-exit outcomes and account balances of all contracts. This distinction is especially important under spouse protection: a first death produces a one-survivor state without a forfeiture, while spouse-continuation payments may reduce the account balance before household extinction. Hence realized forfeitures may be insufficient to pay all unscaled fair mortality-credit amounts in (2.7).*

The group-gain factor introduced below scales these unscaled fair mortality-credit amounts to match aggregate realized forfeitures, thereby enforcing exact ex post pool-level budget balance while surviving recipients remain.

2.5 Pool-level budget constraint and large-pool approximation

For the finite-pool allocation below, assume $0 < \bar{\delta}_{m-1}^\ell < 1$ for all $\ell \in \mathcal{L}_{m-1}$. When at least one contract remains active after the mortality update, realized mortality credits must match realized forfeitures:

$$\sum_{\ell \in \mathcal{L}_{m-1}} \mathbf{1}_m^\ell c_m^\ell = \sum_{\ell \in \mathcal{L}_{m-1}} (1 - \mathbf{1}_m^\ell) \mathcal{W}_{m-}^\ell. \quad (2.8)$$

The left-hand side is the total mortality credit paid to contracts that remain active after the mortality update, and the right-hand side is the total forfeited account balance of contracts that exit the pool during the period. This identity is an ex post pool-level budget constraint.

As in [37], introduce the realized group-gain factor

$$\Gamma_m = \frac{\sum_{k \in \mathcal{L}_{m-1}} (1 - \mathbf{1}_m^k) \mathcal{W}_{m-}^k}{\sum_{k \in \mathcal{L}_{m-1}} \mathbf{1}_m^k \frac{\bar{\delta}_{m-1}^k}{1 - \bar{\delta}_{m-1}^k} \mathcal{W}_{m-}^k}. \quad (2.9)$$

The factor Γ_m is a random variable ex ante: it is determined by the realized pool-exit outcomes and account balances observed at t_m^- . Its numerator is the total realized forfeited account balance available for redistribution, whereas its denominator is the aggregate actuarially fair mortality-credit amount for contracts that remain active after the mortality update. Thus, on the event that the denominator is positive, $\Gamma_m > 1$ means realized forfeitures exceed this aggregate amount, $\Gamma_m < 1$ means realized forfeitures are insufficient, and $\Gamma_m = 0$ occurs when no account balance is forfeited.

Scaling the fair-credit benchmark on the right-hand side of (2.7) by Γ_m defines the finite-pool mortality-credit allocation rule

$$c_m^\ell = \frac{\bar{\delta}_{m-1}^\ell}{1 - \bar{\delta}_{m-1}^\ell} \mathcal{W}_{m-}^\ell \Gamma_m, \quad \text{for all } \ell \in \mathcal{L}_{m-1} \text{ with } \mathbf{1}_m^\ell = 1. \quad (2.10)$$

This allocation satisfies the pool-level budget identity (2.8) exactly whenever at least one contract remains active: the total realized mortality credits paid to contracts that remain active equal the total realized forfeited account balance collected from contracts that exit the pool. Since $\Gamma_m \geq 0$ whenever account balances are nonnegative, the mortality credits paid under (2.10) are also nonnegative.

Following [16, Remark 3.2], if all contracts in $\mathcal{L}_{m-1}$ exit during the period, we set $\Gamma_m = 0$, close the pool, and pay the remaining account balances to the respective estates. This exceptional settlement replaces the mortality-credit allocation in (2.8)–(2.10).

The group-gain factor enforces exact ex post pool-level budget balance, but in a finite heterogeneous pool the allocation rule can introduce a bias in expected mortality credits relative to the individual-contract actuarial-fairness condition (2.7). Following the finite-pool analysis in [37] and Condition 3.1 in [16], this bias is negligible when the pool is sufficiently large and each contract's unscaled actuarially fair mortality-credit amount is small relative to aggregate expected forfeitures:

$$\frac{\bar{\delta}_{m-1}^\ell}{1 - \bar{\delta}_{m-1}^\ell} \mathcal{W}_{m-}^\ell \ll \sum_{k \in \mathcal{L}_{m-1}} \bar{\delta}_{m-1}^k \mathcal{W}_{m-}^k, \quad \text{for all } \ell \in \mathcal{L}_{m-1}. \quad (2.11)$$

Remark 2.3 (Spouse protection and the small-bias condition). *In a spouse-protected pool, a large number of contracts does not by itself ensure (2.11). All contracts begin in state 11, where pool exit requires both household members to die; a first death instead creates a one-survivor state without forfeiture, and spouse-continuation payments may reduce the account balance before eventual exit. The condition therefore also requires sufficiently large aggregate expected forfeitures to support the mortality-credit allocation.*

When (2.11) holds, and the pool is large enough that realized pool exits and realized forfeited account balances are close to their expected aggregate levels, the random group-gain factor satisfies $\mathbb{E}[\Gamma_m] \simeq 1$ with small variance; see [16, 37]. Accordingly, for the remainder of the paper, we adopt the large-pool approximation $\Gamma_m \equiv 1$. With this approximation, the mortality credit for contract ℓ that remains active after the mortality update at time t_m^- becomes

$$c_m^\ell = \frac{\bar{\delta}_{m-1}^\ell}{1 - \bar{\delta}_{m-1}^\ell} \mathcal{W}_{m-}^\ell, \quad \text{for all } \ell \in \mathcal{L}_{m-1} \text{ with } \mathbf{1}_m^\ell = 1. \quad (2.12)$$

Under the large-pool and small-bias assumptions in (2.11), the rule (2.12) preserves the key properties of the finite-pool allocation with high accuracy: (i) individual-contract actuarial fairness, as expressed by (2.7); (ii) the pool-level budget constraint (2.8), exactly when the realized group-gain factor is retained in (2.10) while surviving recipients remain and approximately under the large-pool approximation $\Gamma_m \equiv 1$; and (iii) non-negativity of mortality credits whenever account balances are nonnegative.

2.6 Household-state tontine gain rate

For the homogeneous spouse-protected pool studied in this paper, we suppress the superscript ℓ and consider a representative contract conditional on the contract remaining active at t_m^- and having a positive account balance. Combining the large-pool mortality-credit rule (2.12) with the household-state pool-exit probabilities in (2.4), the mortality credit can be written in tontine-gain form as

$$c_m = g_m(H_{(m-1)+}) \mathcal{W}_{m-}, \quad (2.13)$$

where the state-dependent tontine gain rate is

$$g_m(H_{(m-1)+}) = \begin{cases} \frac{\delta_{R,m-1} \delta_{S,m-1}}{1 - \delta_{R,m-1} \delta_{S,m-1}}, & H_{(m-1)+} = 11, \\ \frac{\delta_{R,m-1}}{1 - \delta_{R,m-1}}, & H_{(m-1)+} = 10, \\ \frac{\delta_{S,m-1}}{1 - \delta_{S,m-1}}, & H_{(m-1)+} = 01. \end{cases} \quad (2.14)$$

For each start-of-period household state, the gain rate has the usual tontine-uplift form: the numerator is the corresponding one-period pool-exit probability, and the denominator is the corresponding probability of remaining active in the pool. Thus, the mortality credit remains proportional to the contract's own account balance, while the proportionality factor depends on the household state. Equations (2.13)–(2.14) define the mortality-credit mechanism used in the household control framework below.

3 Household control framework

Following the decision-time chronology and household-state conventions specified in Section 2, we formulate the controlled wealth recursion for a representative spouse-protected contract. Mortality credits enter through the household-state tontine gain rate $g_m(H_{(m-1)+})$ in (2.14). We denote the retiree-alive states by

$$\mathcal{H}_R := \{11, 10\}.$$

Withdrawal and rebalancing controls are available only in $\mathcal{H}_R$; state 01 follows the spouse-continuation payment and passive portfolio-update conventions, while state 00 is absorbing.

3.1 Index positions and contract wealth

We consider the investment portfolio held in a spouse-protected contract, with access to four real (inflation-adjusted) asset classes: (i) a domestic stock index fund, (ii) a domestic bond index fund, (iii) a foreign stock index fund, converted to domestic currency, and (iv) a foreign bond index fund, converted to domestic currency. A subscript $\iota \in \{d, f\}$ distinguishes domestic and foreign asset classes, and a superscript s or b denotes stock and bond index positions, respectively.

Let $S^d(t)$, $B^d(t)$, $S^f(t)$, and $B^f(t)$ respectively 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 write $S_t^d \equiv S^d(t)$, $B_t^d \equiv B^d(t)$, $S_t^f \equiv S^f(t)$, and $B_t^f \equiv B^f(t)$.

We denote by $\{X_t\}_{0 \leq t \leq T}$ the controlled index-position process, where

$$X_t = (S_t^d, B_t^d, S_t^f, B_t^f)^\top \in \mathbb{R}^4. \quad (3.1)$$

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

$$X_t = \mathcal{M}_{t_m, t}(X_{m+}, \varepsilon_{m+1}), \quad t \in [t_m^+, t_{m+1}^-], \quad (3.2)$$

where $\mathcal{M}_{t_m, t}$ is a transition operator and ε_{m+1} collects the exogenous return drivers 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.

We define by $\{\mathcal{W}_t\}_{0 \leq t \leq T}$ the contract wealth process before applying mortality credits and the tontine management fee, where

$$\mathcal{W}_t = S_t^d + B_t^d + S_t^f + B_t^f \quad (3.3)$$

In particular, at decision time t_m^- we write $\mathcal{W}_{m-} := \mathcal{W}_{t_m^-}$. The quantity $\mathcal{W}_t$ is allowed to become non-positive; under the insolvency convention, negative wealth is interpreted as accumulated debt.

3.2 Mortality credits and tontine fee

To rule out mortality-credit distribution at t_0 and after contract termination or insolvency, we set the mortality credit to zero at t_0 and apply it for $m \geq 1$ only when the contract remains active and satisfies the solvency condition $\mathcal{W}_{m-} > 0$. Accordingly, define the modified state-dependent tontine gain rate by

$$\hat{g}_m = \begin{cases} g_m(H_{(m-1)^+}), & m = 1, \dots, M, H_{m-} \neq 00, \mathcal{W}_{m-} > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (3.4)$$

Here, $g_m(\cdot)$ is defined in (2.14). The resulting contract wealth after mortality-credit distribution at time t_m^- , but before any fee, household payment, or portfolio allocation, is

$$\widetilde{W}_{m-} = (1 + \hat{g}_m) \mathcal{W}_{m-}. \quad (3.5)$$

We allow for a baseline tontine management fee, deducted once per year at times t_m , $m = 1, \dots, M$, and modeled as a proportional charge on the account balance. Let $\varrho \in (0, 1)$ denote the yearly fee rate, and define

$$\varrho_m = \varrho \mathbf{1}_{\{m \geq 1, \widetilde{W}_{m-} > 0, H_{m-} \neq 00\}}, \quad m = 0, \dots, M. \quad (3.6)$$

The contract wealth after fee deduction is

$$W_{m-} = (1 - \varrho_m) \widetilde{W}_{m-}, \quad \widetilde{W}_{m-} \text{ given in (3.5)}. \quad (3.7)$$

In the remainder of the paper, W_{m-} denotes contract wealth at time t_m^- after mortality credits and tontine management fee, but before the period- m household payment and any portfolio allocation. If $W_{m-} \leq 0$, the contract is insolvent at the pre-payment stage; scheduled household payments may still be made, but no further trading or mortality credits are applied (see Subsection 3.6).

3.3 Household payments

This subsection defines the household payment process used in the control problem. The withdrawal control is applied only in retiree-alive states 11 and 10; in the spouse-only state 01, the automatic spouse-protection continuation payment applies. We first specify the withdrawal bands and admissible withdrawal control, then combine controlled withdrawals and spouse-protection continuation payments into a unified payment rule.

3.3.1 Household withdrawal bands

We specify withdrawal bounds directly at the household level. Let $[\underline{q}_{11}, \bar{q}_{11}]$ denote the admissible real (inflation-adjusted) withdrawal band when both household members are alive, i.e. when $H_{m-} = 11$, and let $[\underline{q}_1, \bar{q}_1]$ denote the corresponding one-survivor real (inflation-adjusted) withdrawal band. We impose

$$0 < \underline{q}_1 \leq \underline{q}_{11}, \quad \underline{q}_{11} \leq \bar{q}_{11}, \quad \underline{q}_1 \leq \bar{q}_1, \quad \bar{q}_1 \leq \bar{q}_{11}. \quad (3.8)$$

The band $[\underline{q}_{11}, \bar{q}_{11}]$ applies in state 11 and corresponds to a two-person household. The band $[\underline{q}_1, \bar{q}_1]$ is the one-survivor withdrawal band and applies when the retiree is the sole survivor, i.e. in state 10. We emphasize that state 01 is also a one-survivor household, but it is not a control state: once only

the spouse remains alive, payments are determined by the automatic spouse-protection continuation rule rather than by a controlled withdrawal. The ordering in (3.8) reflects the assumption that a one-survivor household requires no more real (inflation-adjusted) spending than a two-person household.

For subsequent use, define the state-dependent lower and upper controlled withdrawal bounds by

$$\underline{q}(H) := \begin{cases} \underline{q}_{11}, & H = 11, \\ \underline{q}_1, & H = 10, \end{cases} \quad \bar{q}(H) := \begin{cases} \bar{q}_{11}, & H = 11, \\ \bar{q}_1, & H = 10. \end{cases} \quad (3.9)$$

3.3.2 Withdrawal control

For $m = 0, \dots, M-1$, we represent the withdrawal control by a feedback function that maps pre-payment wealth, time, and household state into a withdrawal amount:

$$q_m(\cdot) : (W_{m-}, t_m, H_{m-}) \mapsto q_m = q(W_{m-}, t_m, H_{m-}),$$

where W_{m-} is given in (3.7). The withdrawal control is available only in retiree-alive states, i.e. when $H_{m-} \in \mathcal{H}_R$. Its admissible set is household-state dependent and is given by $\mathcal{Z}_q(W_{m-}, t_m, H_{m-})$, where

$$\mathcal{Z}_q(\cdot) = \begin{cases} [\underline{q}(H_{m-}), \bar{q}(H_{m-})], & H_{m-} \in \mathcal{H}_R, W_{m-} \geq \bar{q}(H_{m-}), & (3.10a) \\ [\underline{q}(H_{m-}), \max\{\underline{q}(H_{m-}), W_{m-}\}], & H_{m-} \in \mathcal{H}_R, 0 < W_{m-} < \bar{q}(H_{m-}), & (3.10b) \\ \{\underline{q}(H_{m-})\}, & H_{m-} \in \mathcal{H}_R, W_{m-} \leq 0, & (3.10c) \\ \{0\}, & H_{m-} \notin \mathcal{H}_R. & (3.10d) \end{cases}$$

Here, $\underline{q}(\cdot)$ and $\bar{q}(\cdot)$ are defined in (3.9). Cases (3.10a) and (3.10b) apply in retiree-alive states when the contract is solvent at t_m^- . In case (3.10a), pre-payment wealth is large enough that the full household-state-dependent withdrawal band is admissible. In case (3.10b), the contract remains solvent but wealth has fallen below the relevant upper bound, so the upper endpoint is tightened while the relevant lower bound is preserved. Case (3.10c) applies when the contract is already insolvent at t_m^- ; the admissible set then collapses to the relevant lower bound, so no withdrawal control is applied. Case (3.10d) records that the withdrawal control is inactive outside retiree-alive states. No withdrawal control is introduced at t_M .

This formulation is also consistent with evidence that retired households adjust spending in response to cash-flow changes, including categories often perceived as relatively fixed; see the related discussion in [28]. It is therefore natural to model withdrawals as a flexible control within state-dependent bounds rather than as a rigid deterministic rule.

This constraint (3.10) admits a natural economic interpretation. In state 11, the upper bound $\bar{q}_{11}$ represents the household's desired annual real spending level while both members are alive, while $\underline{q}_{11}$ is a contingency floor the household is willing to adopt to reduce depletion risk. In state 10, the corresponding quantities $\bar{q}_1$ and $\underline{q}_1$ represent the desired and minimum real spending levels for a one-survivor retiree household. The optimal control uses the flexibility in (3.10) to preserve higher withdrawals when the account is well funded, but cuts withdrawals toward the relevant lower bound in adverse market or wealth states.

3.3.3 Spouse-protection continuation payment

We now define the automatic spouse-protection continuation payment. Let

$$\tau_R := \inf\{m \in \{1, \dots, M\} : H_{m-} = 01\}, \quad (3.11)$$

with the convention $\inf \emptyset = \infty$. The event $\{\tau_R \leq m\}$ is determined by the observed household-state history up to time t_m^- . On the event $\{\tau_R < \infty\}$, the retiree dies during $[t_{\tau_R-1}^+, t_{\tau_R}^-]$ and the spouse is alive at $t_{\tau_R}^-$.

Since the last controlled withdrawal before the retiree's death may have been chosen under the larger two-person band, we cap the spouse-protection continuation payment at the one-survivor upper bound. Define the continuation-base variable recursively by

$$B_{-1} := 0, \quad B_m = \begin{cases} \min\{q_m, \bar{q}_1\}, & H_{m-} \in \mathcal{H}_R, \\ B_{m-1}, & H_{m-} \in \{01, 00\}. \end{cases} \quad (3.12)$$

The process $\{B_m\}_{m=-1}^{M-1}$ stores the capped one-survivor continuation base associated with the most recent controlled withdrawal while the retiree was alive. If the household first enters the spouse-only state at time τ_R , then $B_{\tau_R-1} = \min\{q_{\tau_R-1}, \bar{q}_1\}$ is the automatic continuation amount inherited by the surviving spouse.

3.3.4 Unified household payment rule

We now unify the withdrawal control in retiree-alive states 11 and 10 with automatic spouse-protection continuation payments in the spouse-only state 01. Specifically, for $m = 0, \dots, M-1$, the household payment at time t_m is

$$C_m = \begin{cases} q_m, & H_{m-} \in \{11, 10\}, W_{m-} > 0, \\ \underline{q}(H_{m-}), & H_{m-} \in \{11, 10\}, W_{m-} \leq 0, \\ B_{m-1}, & H_{m-} = 01, W_{m-} > 0, \\ \underline{q}_1, & H_{m-} = 01, W_{m-} \leq 0, \\ 0, & H_{m-} = 00. \end{cases} \quad \begin{aligned} (3.13a) \\ (3.13b) \\ (3.13c) \\ (3.13d) \\ (3.13e) \end{aligned}$$

Cases (3.13a) and (3.13c) cover positive pre-payment wealth. In retiree-alive states 11 and 10, the payment is the controlled withdrawal q_m ; in the spouse-only state 01, it is the automatic spouse-protection continuation amount B_{m-1} . Cases (3.13b) and (3.13d) apply when the contract is already insolvent, with the relevant minimum payment imposed mechanically. Case (3.13e) records that no household payment is made after pool exit.

After the household payment C_m is made, post-payment wealth is

$$W_{m+} = \begin{cases} W_{m-} - C_m, & H_{m-} \in \{11, 10, 01\}, \\ 0, & H_{m-} = 00, \end{cases} \quad m = 0, \dots, M-1. \quad (3.14)$$

At the terminal date $t_M = T$, no further household payment or rebalancing is made. We define terminal wealth by

$$W_T = \begin{cases} W_{M-}, & H_{M-} \in \{11, 10, 01\}, \\ 0, & H_{M-} = 00. \end{cases} \quad (3.15)$$

Thus, terminal wealth is recorded only if at least one household member remains alive at the horizon.

Remark 3.1 (Minimum payments and insolvency). *At first sight it may seem natural to cease household payments once the account is depleted, i.e. when $W_{m-} \leq 0$. However, the lower bounds $\underline{q}_{11}$ and $\underline{q}_1$ represent minimum real spending levels for the two-person state 11 and the one-survivor states 10 and 01. Accordingly, in retiree-alive states 11 and 10, (3.13b) sets $C_m = \underline{q}(H_{m-})$ even when the contract is already insolvent at t_m^- . Hence $W_{m+} = W_{m-} - C_m < 0$, which is equivalent to borrowing the shortfall.*

The same principle applies to the spouse-only state 01. When $W_{m-} > 0$, (3.13c) pays the spouse-protection continuation amount $C_m = B_{m-1}$; if this amount exceeds available wealth, the shortfall is recorded as debt. When the contract is already insolvent at t_m^- , (3.13d) applies the one-survivor minimum $C_m = \underline{q}_1$ mechanically. Any resulting debt then evolves according to the insolvency convention in Subsection 3.6. Accumulated debt increases terminal shortfall if at least one household member remains alive at the horizon.

3.4 Rebalancing control

For $m = 0, \dots, M - 1$ and household states $H_{m-} \in \mathcal{H}_R$, we denote the rebalancing control by $\mathbf{p}_m(\cdot)$, a vector of proportions of post-payment wealth allocated to the asset indices. The control depends on the post-payment wealth W_{m+} (defined in (3.14)), the decision time t_m , and the household state H_{m-} :

$$\mathbf{p}_m(\cdot) : (W_{m+}, t_m, H_{m-}) \mapsto \mathbf{p}_m = \mathbf{p}(W_{m+}, t_m, H_{m-}), \quad H_{m-} \in \mathcal{H}_R,$$

where

$$\mathbf{p}_m = (p_{s,m}^d, p_{b,m}^d, p_{s,m}^f). \quad (3.16)$$

We enforce no leverage and no short-selling whenever the contract is solvent. To define the admissible set of rebalancing controls, introduce

$$\Delta^{(3)} := \left\{ (p_1, p_2, p_3) \in [0, 1]^3 \mid \sum_{i=1}^3 p_i \leq 1 \right\}, \quad \mathbf{b}^{(3)} := (0, 1, 0).$$

The three explicitly parameterized proportions correspond to domestic stock, domestic bond, and foreign stock, while the residual proportion is allocated to the foreign bond.

For retiree-alive states $H_{m-} \in \mathcal{H}_R$, the admissible rebalancing control set is

$$\mathcal{Z}_p(W_{m+}, t_m) = \begin{cases} \Delta^{(3)}, & W_{m+} > 0, \\ \{\mathbf{b}^{(3)}\}, & W_{m+} \leq 0. \end{cases} \quad (3.17)$$

The singleton $\{\mathbf{b}^{(3)}\}$ is a bookkeeping placeholder for the control variable after insolvency and represents full allocation to the domestic bond/debt component; it is not a genuine rebalancing decision. Outside retiree-alive states, the admissible set (3.17) is not used. No rebalancing control is introduced at t_M .

For $m = 0, \dots, M - 1$, if $H_{m-} \in \mathcal{H}_R$ and $W_{m+} > 0$, the post-rebalancing state X_{m+} is

$$X_{m+} = W_{m+} (p_{s,m}^d, p_{b,m}^d, p_{s,m}^f, p_{b,m}^f)^\top, \quad \text{where } p_{b,m}^f := 1 - p_{s,m}^d - p_{b,m}^d - p_{s,m}^f. \quad (3.18)$$

If the household payment results in $W_{m+} \leq 0$, no rebalancing control is applied; instead, the insolvency convention in Subsection 3.6 applies. If $H_{m-} = 01$ and $W_{m+} > 0$, no rebalancing control is applied; the portfolio is updated according to the passive rebalancing rule in Subsection 3.5. If $H_{m-} = 00$, the state is absorbing and no further portfolio update occurs.

3.5 Passive rebalancing and contract termination

If $H_{m-} = 01$, no new withdrawal or rebalancing control is chosen. The spouse-protection continuation payment is determined by (3.13). To describe the passive portfolio update in this state, recall that $\mathcal{W}_{m-}$, defined in (3.3), is the contract wealth before mortality credits and fees, while W_{m-} , defined in (3.7), is the wealth after mortality credits and the tontine management fee, but before the household payment.

When $\mathcal{W}_{m-} > 0$ and $W_{m-} > 0$, define the pre-payment asset-position vector by

$$X_{m-}^{\text{pay}} := \frac{W_{m-}}{\mathcal{W}_{m-}} X_{m-}. \quad (3.19)$$

The scaling in (3.19) updates the asset positions to reflect mortality credits and the tontine fee while preserving the pre-payment portfolio proportions. If either $\mathcal{W}_{m-} \leq 0$ or $W_{m-} \leq 0$, the contract is already insolvent at the pre-payment stage, and (3.19) is not used.

If $H_{m-} = 01$ and the spouse-protection continuation payment leaves the contract solvent, i.e. $W_{m+} > 0$, the passive rebalancing rule is

$$X_{m+} = \frac{W_{m+}}{W_{m-}} X_{m-}^{\text{pay}}, \quad \text{when } H_{m-} = 01, \mathcal{W}_{m-} > 0, W_{m-} > 0, W_{m+} > 0. \quad (3.20)$$

Equation (3.20) preserves the pre-payment portfolio proportions after the spouse-protection continuation payment is made. Equivalently, the payment is funded by reducing each asset position proportionally;

no new portfolio allocation is chosen in state 01. If $W_{m+} \leq 0$, the passive rebalancing rule is not used and the insolvency convention applies.

If $H_{m-} = 00$, the contract has exited the pool and neither household member remains alive. We then set

$$C_m = 0, \quad W_{m+} = 0, \quad X_{m+} = (0, 0, 0, 0)^\top. \quad (3.21)$$

State 00 is absorbing: once the contract has exited the pool, no further household payments, mortality credits, or portfolio updates occur. The reset in (3.21) ends the household wealth recursion; settlement of external debt after extinction is outside the model.

3.6 Insolvency

For an active contract, insolvency occurs at t_m^+ whenever the post-payment wealth in (3.14) satisfies

$$H_{m-} \neq 00, \quad W_{m+} \leq 0.$$

Thus, insolvency is triggered after the household payment has been made and before any rebalancing control or passive rebalancing rule is applied. If $W_{m-} \leq 0$, the contract has carried debt forward from an earlier period.

Once the contract is insolvent, no further discretionary withdrawal or rebalancing control is applied and no further mortality credits are distributed. The applicable scheduled minimum payments under (3.13) continue while the contract remains active, with any resulting shortfall carried as debt in the domestic bond/debt component. Accordingly, the portfolio of an insolvent active contract is set to

$$X_{m+} = W_{m+}(0, 1, 0, 0)^\top, \quad \text{when } H_{m-} \neq 00 \text{ and } W_{m+} \leq 0. \quad (3.22)$$

In the baseline implementation, the debt position accrues at the domestic bond-index return plus a borrowing spread of 2% per annum.

3.7 Admissible policy set

A household control policy is a sequence of non-anticipative feedback rules

$$\nu = \{(q_m, \mathbf{p}_m)\}_{m=0}^{M-1},$$

where q_m is the withdrawal control and $\mathbf{p}_m$ is the rebalancing control. The withdrawal control is applied using (W_{m-}, t_m, H_{m-}) and is constrained by $\mathcal{Z}_q(W_{m-}, t_m, H_{m-})$ defined in (3.10). The rebalancing control is defined only on retiree-alive states $H_{m-} \in \mathcal{H}_R$, and is constrained by $\mathcal{Z}_p(W_{m+}, t_m)$ defined in (3.17). The admissible household control policy set is

$$\mathcal{A} = \left\{ \nu = \{(q_m, \mathbf{p}_m)\}_{m=0}^{M-1} \left| \begin{array}{l} q_m(W_{m-}, t_m, H_{m-}) \in \mathcal{Z}_q(W_{m-}, t_m, H_{m-}), \\ \mathbf{p}_m(W_{m+}, t_m, H_{m-}) \in \mathcal{Z}_p(W_{m+}, t_m), \quad \text{whenever } H_{m-} \in \mathcal{H}_R \end{array} \right. \right\}. \quad (3.23)$$

Outside retiree-alive states, no rebalancing control is defined: state 01 uses the passive rebalancing rule, while state 00 is absorbing. The solvency restrictions are already embedded in $\mathcal{Z}_q$ and $\mathcal{Z}_p$, defined in (3.10) and (3.17), respectively.

4 Household reward-risk optimization

We now specify the household-level reward and risk criteria used in the household optimization problem. Throughout, all household payments and contract wealth are measured in real (inflation-adjusted) terms. The formulation reflects two central features of the spouse-protection setting: household payments depend on the household state, and the terminal reserve targets depend on the terminal household state.

Let $x_0 = X_{0-}$ and $h_0 = H_{0-}$ denote the initial index-position vector and household state, respectively. For any admissible policy ν , we write $\mathbb{E}_{x_0, h_0}^\nu[\cdot] := \mathbb{E}^\nu[\cdot | X_{0-} = x_0, H_{0-} = h_0]$, where the expectation is taken over the joint law of the exogenous asset-return and household-state processes, with contract wealth and household payments induced by the policy ν .

4.1 Reward

The reward criterion is the expected cumulative real (inflation-adjusted) household payments over the planning horizon. Given $\nu \in \mathcal{A}$, let C_m^ν denote the period- m household payment generated by the unified payment rule (3.13). This rule already distinguishes controlled withdrawals in retiree-alive states from automatic spouse-protection continuation payments in the spouse-only state, and incorporates the lower-bound payment convention when the contract is already insolvent. We therefore define the household reward functional by

$$\mathcal{R}(\nu) := \mathbb{E}_{x_0, h_0}^\nu \left[\sum_{m=0}^{M-1} C_m^\nu \right]. \quad (4.1)$$

Remark 4.1 (Discounting and household-state averaging). *All payments and wealth are expressed in real dollars, and the baseline reward functional (4.1) sets the real discount factor equal to one at each decision time. More generally, prescribed real discount factors $\{d_m\}_{m=0}^{M-1}$ could be introduced by replacing $\sum_{m=0}^{M-1} C_m^\nu$ with $\sum_{m=0}^{M-1} d_m C_m^\nu$.*

Mortality is incorporated through the stochastic household-state process $\{H_t\}_{0 \leq t \leq T}$. Accordingly, the expectation in (4.1) averages the state-contingent household payments over the joint law of asset returns and household mortality, so no additional survival-probability weighting is required.

4.2 Risk

On the risk side, we do not use terminal contract wealth itself as the risk variable, nor do we penalize low terminal wealth uniformly across all terminal household states. In a single-retiree formulation with a fixed planning horizon, terminal wealth can often be used directly as a proxy for terminal adequacy. In a spouse-protected contract, however, the same terminal wealth has different implications depending on the household state: it may be inadequate for a two-person household, sufficient for a one-survivor household, or irrelevant for continuing household reserve needs once neither household member remains alive. We therefore measure terminal risk through a state-contingent shortfall variable, applying CVaR to the shortfall from state-dependent terminal reserve targets rather than directly to terminal wealth.

Let W_T^ν denote the contract wealth at time T under policy $\nu \in \mathcal{A}$, and write $H_T := H_{M-}$ for the terminal household state after the final mortality update. To measure downside risk under policy ν , we define the associated state-contingent terminal shortfall, denoted by L_T^ν .

Specifically, let w_{11}^* , w_{10}^* , and w_{01}^* be nonnegative terminal reserve targets, corresponding respectively to the terminal household states 11, 10, and 01. These reserve targets are model inputs and may be specified externally using household retirement-spending benchmarks. A natural ordering is

$$w_{11}^* \geq \max\{w_{10}^*, w_{01}^*\},$$

since state 11 corresponds to a two-person household. The one-survivor targets w_{10}^* and w_{01}^* may be chosen equal or specified differently, depending on whether retiree-only and spouse-only terminal reserve needs are treated as symmetric.

The state-contingent terminal shortfall random variable L_T^ν is

$$L_T^\nu := \mathbf{1}_{\{H_T=11\}}(w_{11}^* - W_T^\nu)_+ + \mathbf{1}_{\{H_T=10\}}(w_{10}^* - W_T^\nu)_+ + \mathbf{1}_{\{H_T=01\}}(w_{01}^* - W_T^\nu)_+, \quad (4.2)$$

where $(x)_+ := \max\{x, 0\}$. No shortfall term is included when $H_T = 00$, since no household member remains alive at the horizon.

To focus on adverse household outcomes, we apply CVaR directly to the state-contingent terminal shortfall. For a confidence level $\alpha \in (0, 1)$, define the household risk functional under policy ν by

$$\mathcal{Q}_\alpha(\nu) := \text{CVaR}_\alpha(L_T^\nu) := \inf_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{1-\alpha} \mathbb{E}_{x_0, h_0}^\nu [(L_T^\nu - \eta)_+] \right\}. \quad (4.3)$$

For example, $\text{CVaR}_{0.95}(L_T^\nu)$ corresponds to the upper 5% tail of the household terminal shortfall. This specification penalizes inadequate terminal reserves only in household states where at least one household member remains alive, with the shortfall target determined by the terminal household state.

4.3 Household value function

The household optimization problem trades off two competing criteria: increasing expected cumulative household payments and reducing terminal reserve shortfall risk. We combine these criteria using a scalarization parameter $\lambda > 0$, which converts the reward–risk trade-off into a single scalar objective. Since $\mathcal{Q}_\alpha(\nu)$ is a risk functional applied to a loss variable, it enters the objective with a negative sign.

The household value function associated with this scalarized reward–risk criterion is

$$V(x_0, h_0, t_0^-) := \sup_{\nu \in \mathcal{A}} \{ \mathcal{R}(\nu) - \lambda \mathcal{Q}_\alpha(\nu) \}. \quad (4.4)$$

For any admissible household control policy $\nu \in \mathcal{A}$ and auxiliary CVaR threshold $\eta \in \mathbb{R}$, define the household objective associated with (ν, η) by

$$V(x_0, h_0, t_0^-; \nu, \eta) := \mathbb{E}_{x_0, h_0}^\nu \left[\sum_{m=0}^{M-1} C_m^\nu - \lambda \left(\eta + \frac{1}{1-\alpha} (L_T^\nu - \eta)_+ \right) \right], \quad (4.5)$$

subject to the following system evolution and control-admissibility conditions:

$$\left\{ \begin{array}{ll} \text{(initial state)} & X_{0-} = x_0, \quad H_{0-} = h_0, \\ \text{(index-position dynamics)} & \{X_t\} \text{ evolves via (3.2), } \quad t \notin \mathcal{T}, \\ \text{(household state)} & H_{m-} \text{ evolves according to (2.1)–(2.3),} \\ \text{(tontine and fee)} & \mathcal{W}_{m-}, \widetilde{\mathcal{W}}_{m-}, W_{m-} \text{ are computed via (3.4)–(3.7),} \\ \text{(withdrawal control)} & q_m \in \mathcal{Z}_q(W_{m-}, t_m, H_{m-}) \text{ as defined in (3.10),} \\ \text{(household payment)} & B_m, C_m, W_{m+} \text{ are computed via (3.12)–(3.14),} \\ \text{(rebalancing control)} & \mathbf{p}_m \in \mathcal{Z}_p(W_{m+}, t_m) \text{ when } H_{m-} \in \mathcal{H}_R, \text{ as defined in (3.17),} \\ \text{(asset update)} & X_{m+} \text{ is computed via (3.18), (3.19)–(3.20), (3.21), or (3.22),} \end{array} \right. \quad (4.6)$$

Using the explicit CVaR representation (4.3), and since ν and η range over the independent feasible sets $\mathcal{A}$ and $\mathbb{R}$, respectively, the iterated suprema commute. Hence, subject to the system evolution and control-admissibility conditions in (4.6), the household value function $V(x_0, h_0, t_0^-)$ can equivalently be written as

$$V(x_0, h_0, t_0^-) = \sup_{\nu \in \mathcal{A}} \sup_{\eta \in \mathbb{R}} V(x_0, h_0, t_0^-; \nu, \eta) = \sup_{\eta \in \mathbb{R}} \sup_{\nu \in \mathcal{A}} V(x_0, h_0, t_0^-; \nu, \eta). \quad (4.7)$$

Remark 4.2 (Pre-commitment and induced time-consistent interpretation). *The scalarized problem (4.4) is evaluated at inception and therefore defines a pre-commitment objective. Since CVaR is generally not time-consistent in multi-period settings, the reward-risk problem is formally time-inconsistent: after observing future wealth and household states, the household may prefer to solve a different continuation problem. Nevertheless, the pre-commitment formulation is useful for identifying the inception-time reward-risk trade-off. If $(\nu_\lambda^*, \eta_\lambda^*)$ attains the suprema in (4.7), then fixing η_λ^* gives the induced expected-value control problem*

$$\sup_{\nu \in \mathcal{A}} \mathbb{E}_{x_0, h_0}^\nu \left[\sum_{m=0}^{M-1} C_m^\nu - \frac{\lambda}{1-\alpha} (L_T^\nu - \eta_\lambda^*)_+ \right],$$

up to the constant term $-\lambda \eta_\lambda^$. With the threshold fixed, the objective is an additive expected-value criterion with a terminal shortfall penalty and therefore has the usual dynamic-programming interpretation. We interpret the computed policy ν_λ^* as the induced time-consistent strategy for the fixed-threshold objective associated with $(\alpha, \lambda, \eta_\lambda^*)$; see, for example, [9, 15, 39].*

For a fixed confidence level α , let ν_λ^* denote an optimizer of (4.4) for a given $\lambda > 0$. The household reward-risk frontier is the set

$$\mathcal{S}(\alpha) := \{ (\mathcal{Q}_\alpha(\nu_\lambda^*), \mathcal{R}(\nu_\lambda^*)) : \lambda > 0 \}. \quad (4.8)$$

As $\lambda > 0$ varies, the corresponding optimizers of (4.4) generate this set. We report the frontier with $\mathcal{Q}_\alpha(\nu_\lambda^*)$ on the horizontal axis and $\mathcal{R}(\nu_\lambda^*)$ on the vertical axis.

5 Neural network formulation

To solve the household reward–risk stochastic control problem numerically, we adopt a neural-network policy-approximation approach in the spirit of the rapidly developing literature on direct NN-based control methods [5, 24, 28, 35, 36, 43]. These methods avoid classical dynamic-programming discretizations by parameterizing the control itself and training it end-to-end. They are often described as “global-in-time” machine-learning approaches to stochastic control [19], in contrast to stacked-NN schemes that use a separate network at each rebalancing step [18, 42].

5.1 Approximation of admissible controls

The central idea of our NN formulation is to approximate an admissible household control policy $\nu \in \mathcal{A}$ by parameterizing its two feedback components with separate feedforward, fully connected neural networks: a withdrawal network and a rebalancing network. Specifically, given trainable network parameters θ_q and θ_p , define the NN-parameterized withdrawal and rebalancing controls by

$$\begin{aligned}\widehat{q}_m(\cdot; \theta_q) &:= \widehat{q}(W_{m-}, t_m, H_{m-}; \theta_q) \simeq q_m(W_{m-}, t_m, H_{m-}), \\ \widehat{p}_m(\cdot; \theta_p) &:= \widehat{p}(W_{m+}, t_m, H_{m-}; \theta_p) \simeq p_m(W_{m+}, t_m, H_{m-}).\end{aligned}\tag{5.1}$$

The NN-parameterized household control policy is the sequence of paired controls

$$\widehat{\nu} = \{(\widehat{q}_m(\cdot; \theta_q), \widehat{p}_m(\cdot; \theta_p))\}_{m=0}^{M-1},\tag{5.2}$$

where $\widehat{q}_m(\cdot; \theta_q)$ and $\widehat{p}_m(\cdot; \theta_p)$ are defined in (5.1). The NN-parameterized controls in (5.2) are applied according to the admissibility conventions in Section 3: the withdrawal control is applied only when $H_{m-} \in \mathcal{H}_R$ and $W_{m-} > 0$, while the rebalancing control is evaluated only when $H_{m-} \in \mathcal{H}_R$ and $W_{m+} > 0$. Outside these conditions, the corresponding network output is not used; the household-payment convention, passive rebalancing update, termination convention, or insolvency convention applies instead.

Let

$$\widehat{\mathcal{A}} := \{\widehat{\nu} \in \mathcal{A} \mid \widehat{\nu} \text{ has the form (5.2)}\}$$

denote the class of admissible NN-parameterized household control policies. We now approximate ν in the household objective $V(x_0, h_0, t_0^-; \nu, \eta)$ by restricting attention to the NN-parameterized policy class $\widehat{\mathcal{A}}$. Specifically, for fixed $\widehat{\nu} \in \widehat{\mathcal{A}}$ and $\eta \in \mathbb{R}$, define the NN-induced household objective by

$$V_{NN}(x_0, h_0, t_0^-; \widehat{\nu}, \eta) := \mathbb{E}_{x_0, h_0}^{\widehat{\nu}} \left[\sum_{m=0}^{M-1} C_m^{\widehat{\nu}} - \lambda \left(\eta + \frac{1}{1-\alpha} (L_T^{\widehat{\nu}} - \eta)_+ \right) \right].\tag{5.3}$$

Here, the expectation in (5.3) is taken under the NN-induced system evolution obtained from (4.6), with the controls $q_m(\cdot)$ and $p_m(\cdot)$ replaced by their NN parameterizations $\widehat{q}_m(\cdot)$ and $\widehat{p}_m(\cdot)$ in (5.1). The household value function $V(x_0, h_0, t_0^-)$, defined in (4.7), is approximated by restricting the policy search to $\widehat{\mathcal{A}}$:

$$V(x_0, h_0, t_0^-) \simeq V_{NN}(x_0, h_0, t_0^-) := \sup_{\widehat{\nu} \in \widehat{\mathcal{A}}} \sup_{\eta \in \mathbb{R}} V_{NN}(x_0, h_0, t_0^-; \widehat{\nu}, \eta).\tag{5.4}$$

5.2 Network architecture for controls

The network architecture is designed so that admissibility is enforced by construction. We use one feedforward network for the withdrawal control and one feedforward network for the rebalancing control, with time, wealth, and household state used as inputs. The output layers are chosen to map unconstrained network outputs into the admissible sets (3.10) and (3.17).

Withdrawal control. Let $z_q \in \mathbb{R}$ denote the pre-activation output of the withdrawal network’s final layer. For $H_{m-} \in \mathcal{H}_R$ and $W_{m-} > 0$, the state-dependent lower and upper withdrawal bounds are

$\underline{q}(H_{m-})$ and $\bar{q}(H_{m-})$, as defined in (3.9). We enforce the admissibility condition (3.10) using the scaled sigmoid transformation

$$\hat{q}(W_{m-}, t_m, H_{m-}; \theta_q) = \underline{q}(H_{m-}) + \max(\min\{\bar{q}(H_{m-}), W_{m-}\} - \underline{q}(H_{m-}), 0) \frac{1}{1 + e^{-z_q}}. \quad (5.5)$$

Since the sigmoid function maps into $(0, 1)$, this ensures that $\hat{q}_m(\cdot) \in \mathcal{Z}_q(W_{m-}, t_m, H_{m-})$ whenever $H_{m-} \in \mathcal{H}_R$ and $W_{m-} > 0$. If $H_{m-} \notin \mathcal{H}_R$ or $W_{m-} \leq 0$, the withdrawal network output is not used.

Rebalancing control. The rebalancing network outputs four logits $(z_1, z_2, z_3, z_4) \in \mathbb{R}^4$. Applying a four-way softmax gives

$$\hat{p}_m^{(i)} = \frac{e^{z_i}}{\sum_{\ell=1}^4 e^{z_\ell}}, \quad i = 1, \dots, 4. \quad (5.6)$$

This guarantees $\hat{p}_m^{(i)} \in [0, 1]$ and $\sum_{i=1}^4 \hat{p}_m^{(i)} = 1$. The first three components define the explicitly parameterized rebalancing control, $\hat{\mathbf{p}}_m(\cdot) = (\hat{p}_m^{(1)}, \hat{p}_m^{(2)}, \hat{p}_m^{(3)})$, corresponding to domestic stock, domestic bond, and foreign stock, while $\hat{p}_m^{(4)}$ is the residual foreign-bond allocation.

For $m = 0, \dots, M-1$, the softmax transformation enforces the simplex constraint in $\mathcal{Z}_p(W_{m+}, t_m)$ automatically whenever $H_{m-} \in \mathcal{H}_R$ and $W_{m+} > 0$. If $H_{m-} \notin \mathcal{H}_R$ or $W_{m+} \leq 0$, the rebalancing network output is not used. If the contract becomes insolvent after the household payment, the insolvency convention applies before any rebalancing is carried out. In state 01, the passive rebalancing rule applies; in state 00, the contract is absorbing. No terminal rebalancing control is introduced at t_M .

With these output transformations in place, an NN-parameterized policy $\hat{\nu}$ is fully determined by its trainable network parameters (θ_q, θ_p) . Accordingly, we write $V_{NN}(\cdot; \hat{\nu}, \eta)$ as $V_{NN}(\cdot; \theta_q, \theta_p, \eta)$ to make this dependence explicit. The optimization problem (5.4) then becomes the unconstrained optimization problem

$$V_{NN}(x_0, h_0, t_0^-) = \sup_{\theta_q, \theta_p, \eta} V_{NN}(x_0, h_0, t_0^-; \theta_q, \theta_p, \eta). \quad (5.7)$$

Here, the supremum is taken over all trainable weights and biases in the withdrawal and rebalancing networks, together with the scalar CVaR threshold η . We denote the optimal trainable parameters and CVaR threshold by θ_q^* , θ_p^* , and η^* , respectively.

We emphasize that, although the original household control problem (4.7) is constrained, the NN-parameterized training problem in (5.7) is unconstrained in the trainable variables. The withdrawal and rebalancing admissibility constraints are enforced through the output transformations in (5.5)–(5.6). This permits the use of standard gradient-based optimizers such as Adam.

5.3 Empirical objective

To evaluate the NN-induced objective in (5.7), we approximate its expectation using a finite training sample of N independent realizations of the exogenous asset-return and household-state processes. These realizations are indexed by $n = 1, \dots, N$, and all sample-dependent quantities carry the superscript “ (n) ”. For example, $\{X_t^{(n)}\}_{0 \leq t \leq T}$ is the sample- n realization of the controlled index-position process, $\{H_{m-}^{(n)}\}_{m=0}^M$ is the corresponding household-state realization, and $W_T^{(n)}$ is the terminal contract wealth. Unless otherwise noted, the dependence of controlled pathwise quantities, such as $X_t^{(n)}$ and $W_T^{(n)}$, on the current network parameters (θ_q, θ_p) is suppressed to keep notation light.

For any (θ_q, θ_p) , let $C_m^{(n)}$ denote the period- m household payment generated along the n -th training path under the current network parameters (θ_q, θ_p) . The corresponding sample-wise household terminal shortfall is

$$L_T^{(n)} = \mathbf{1}_{\{H_T^{(n)}=11\}}(w_{11}^* - W_T^{(n)})_+ + \mathbf{1}_{\{H_T^{(n)}=10\}}(w_{10}^* - W_T^{(n)})_+ + \mathbf{1}_{\{H_T^{(n)}=01\}}(w_{01}^* - W_T^{(n)})_+. \quad (5.8)$$

where $H_T^{(n)} := H_{M-}^{(n)}$. The associated NN-induced objective $V_{NN}(x_0, h_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, \eta)$ is approximated by the empirical objective

$$\begin{aligned} V_{NN}(x_0, h_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, \eta) &\approx \widehat{V}_{NN}(x_0, h_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, \eta) \\ &:= \frac{1}{N} \sum_{n=1}^N \sum_{m=0}^{M-1} C_m^{(n)} - \lambda \left[\eta + \frac{1}{(1-\alpha)N} \sum_{n=1}^N (L_T^{(n)} - \eta)_+ \right]. \end{aligned} \quad (5.9)$$

We train the NN policy by maximizing the empirical objective (5.9) over the network parameters and the CVaR threshold. Thus,

$$(\boldsymbol{\theta}_q^*, \boldsymbol{\theta}_p^*, \eta^*) := \arg \max_{\boldsymbol{\theta}_q, \boldsymbol{\theta}_p, \eta} \widehat{V}_{NN}(x_0, h_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, \eta). \quad (5.10)$$

Equivalently, one may minimize the empirical loss

$$\mathcal{L}(\boldsymbol{\theta}_q, \boldsymbol{\theta}_p, \eta) := -\widehat{V}_{NN}(x_0, h_0, t_0^-; \boldsymbol{\theta}_q, \boldsymbol{\theta}_p, \eta).$$

The resulting learned withdrawal and rebalancing controls are the sequences

$$\widehat{q}^* := \{\widehat{q}_m^*(\cdot)\}_{m=0}^{M-1}, \quad \widehat{p}^* := \{\widehat{p}_m^*(\cdot)\}_{m=0}^{M-1},$$

where, for each $m = 0, \dots, M-1$,

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

Thus, the learned NN household control policy is

$$\widehat{\nu}^* = (\widehat{q}^*, \widehat{p}^*) = \{(\widehat{q}_m^*(\cdot), \widehat{p}_m^*(\cdot))\}_{m=0}^{M-1}.$$

The network architecture, training hyperparameters, sequential transfer-learning protocol, and out-of-sample evaluation design used in the numerical experiments are reported in Appendix A.

6 Spouse-continuation load and book-level diversification

This section quantifies the spouse-continuation cost induced by an optimal household control policy. We first consider the representative-contract expected-cost load and tail risk, then book-level diversification, and finally use the excess upper-tail cost after diversification to construct the risk-loaded continuation cost and payment-scale load. These are reporting quantities and do not enter the household optimization.

Let ν^* denote an optimal household control policy for the reward–risk problem (4.4). Under ν^* , define the total payments in the retiree-alive states and in the spouse-only phase respectively by

$$Y_r^{\nu^*} := \sum_{m=0}^{M-1} C_m^{\nu^*} \mathbf{1}_{\{H_{m-} \in \mathcal{H}_R\}}, \quad Z_c^{\nu^*} := \sum_{m=0}^{M-1} C_m^{\nu^*} \mathbf{1}_{\{H_{m-} = 01\}}. \quad (6.1)$$

Here, $C_m^{\nu^*}$ denotes the household payment at t_m generated by the unified payment rule (3.13) under ν^* . All payment quantities are expressed in real dollars, and no additional discount factor is applied.

For the remainder of this section, we suppress the policy superscript and write $Y_r := Y_r^{\nu^*}$ and $Z_c := Z_c^{\nu^*}$. We refer to Z_c as the representative-contract spouse-continuation cost. Here, “cost” refers to cumulative spouse-continuation payments over the planning horizon. Under the withdrawal and capped spouse-continuation payment rules, and using $\bar{q}_{11} \geq \bar{q}_1$, the finite horizon gives

$$0 \leq Y_r \leq M\bar{q}_{11}, \quad 0 \leq Z_c \leq M\bar{q}_1. \quad (6.2)$$

6.1 Representative-contract expected-cost load and tail risk

For brevity, write $\mathbb{E}^{\nu^*}[\cdot] := \mathbb{E}_{x_0, h_0}^{\nu^*}[\cdot]$ for expectation under the fixed policy and initial state. The expected-cost spouse-continuation load, expressed as the share of expected total household payments attributable to the spouse-continuation cost, is

$$f_c(\nu^*) := \frac{\mathbb{E}^{\nu^*}[Z_c]}{\mathbb{E}^{\nu^*}[Y_r] + \mathbb{E}^{\nu^*}[Z_c]}. \quad (6.3)$$

Let $\alpha_c \in (0, 1)$ denote an upper-tail probability. For an integrable nonnegative cost Z generated under ν^* , define

$$\text{CVaR}_{\alpha_c}^+(Z) := \inf_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{\alpha_c} \mathbb{E}^{\nu^*}[(Z - \eta)_+] \right\}, \quad (6.4)$$

and define the excess upper-tail cost above the mean by

$$\Delta_{\alpha_c}^+(Z) := \text{CVaR}_{\alpha_c}^+(Z) - \mathbb{E}^{\nu^*}[Z]. \quad (6.5)$$

For example, $\alpha_c = 0.05$ corresponds to the upper 5% tail. In particular, $\Delta_{\alpha_c}^+(Z_c)$ measures the amount by which the representative-contract upper-tail spouse-continuation cost exceeds its mean. We retain $\text{CVaR}_{\alpha_c}^+(Z_c)$ as a representative-contract tail-cost benchmark. The tail probability α_c is used for risk loading and need not coincide with the confidence level in the household reward-risk objective.

6.2 Book-level diversification

Let $\mathcal{G}$ denote the sigma-field generated by the common market environment, namely, the asset returns over the planning horizon to which all contracts in the book are exposed. Consider a homogeneous book of spouse-protected contracts following the same policy ν^* , with household mortality histories conditionally independent and identically distributed given $\mathcal{G}$. We use J to denote the number of contracts over which spouse-continuation costs are aggregated; it is distinct from the tontine-pool size underlying the homogeneous large-pool gain rates. We first examine the average per-contract cost for a finite book and then its large-book limit.

6.2.1 Finite-book average per-contract cost

Consider a book of J spouse-protected contracts, and let $Z_{c,j}$ denote the continuation cost of contract j , where $j = 1, \dots, J$. Define the finite-book average per-contract spouse-continuation cost by

$$\bar{Z}_{c,J} := \frac{1}{J} \sum_{j=1}^J Z_{c,j}. \quad (6.6)$$

By homogeneity, linearity of expectation, and the positive homogeneity and subadditivity of CVaR,

$$\mathbb{E}^{\nu^*}[\bar{Z}_{c,J}] = \mathbb{E}^{\nu^*}[Z_c], \quad \text{CVaR}_{\alpha_c}^+(\bar{Z}_{c,J}) \leq \text{CVaR}_{\alpha_c}^+(Z_c), \quad \Delta_{\alpha_c}^+(\bar{Z}_{c,J}) \leq \Delta_{\alpha_c}^+(Z_c). \quad (6.7)$$

6.2.2 Large-book limiting average per-contract cost

By (6.2), Z_c is bounded. Hence, under the conditional-i.i.d. assumption, the conditional law of large numbers gives

$$\bar{Z}_{c,J} \longrightarrow \bar{Z}_c := \mathbb{E}^{\nu^*}[Z_c | \mathcal{G}], \quad J \rightarrow \infty, \quad \text{almost surely.} \quad (6.8)$$

We refer to $\bar{Z}_c$ as the large-book limiting average per-contract spouse-continuation cost. Equation (6.8) shows that idiosyncratic household-mortality risk in the continuation cost is diversified away in the large-book limit, whereas common-market risk remains. Furthermore, the law of iterated expectations gives

$$\mathbb{E}^{\nu^*}[\bar{Z}_c] = \mathbb{E}^{\nu^*}[\mathbb{E}^{\nu^*}[Z_c | \mathcal{G}]] = \mathbb{E}^{\nu^*}[Z_c]. \quad (6.9)$$

Thus, book-level diversification leaves the expected cost per contract unchanged.

Finite-book convergence rate. Define the average conditional idiosyncratic variance by

$$\sigma_{c,\text{id}}^2(\nu^*) := \mathbb{E}^{\nu^*} \left[\text{Var}^{\nu^*}(Z_c \mid \mathcal{G}) \right] \leq \frac{M^2 \bar{q}_1^2}{4} < \infty, \quad (6.10)$$

where the inequality follows from (6.2) and Popoviciu's inequality. In particular, the bound is independent of the book size J .

The conditional i.i.d. assumption and the Cauchy-Schwarz inequality give

$$\mathbb{E}^{\nu^*} \left[(\bar{Z}_{c,J} - \bar{Z}_c)^2 \right] = \frac{\sigma_{c,\text{id}}^2(\nu^*)}{J}, \quad \mathbb{E}^{\nu^*} [|\bar{Z}_{c,J} - \bar{Z}_c|] \leq \frac{\sigma_{c,\text{id}}(\nu^*)}{\sqrt{J}}. \quad (6.11)$$

The variational representation (6.4) implies that upper-tail CVaR is $1/\alpha_c$ -Lipschitz in L^1 . Hence

$$|\text{CVaR}_{\alpha_c}^+(\bar{Z}_{c,J}) - \text{CVaR}_{\alpha_c}^+(\bar{Z}_c)| \leq \frac{\sigma_{c,\text{id}}(\nu^*)}{\alpha_c \sqrt{J}}, \quad |\Delta_{\alpha_c}^+(\bar{Z}_{c,J}) - \Delta_{\alpha_c}^+(\bar{Z}_c)| \leq \frac{\sigma_{c,\text{id}}(\nu^*)}{\alpha_c \sqrt{J}}, \quad (6.12)$$

because $\bar{Z}_{c,J}$ and $\bar{Z}_c$ have the same mean by (6.7) and (6.9). Thus, $\Delta_{\alpha_c}^+(\bar{Z}_{c,J})$ converges to $\Delta_{\alpha_c}^+(\bar{Z}_c)$ with an $\mathcal{O}(J^{-1/2})$ error bound.

Remark 6.1 (Square-root book-size approximation). *If the common market environment is degenerate, i.e. the shared asset returns are deterministic, then $\bar{Z}_c = \mathbb{E}^{\nu^*}[Z_c]$ is deterministic and $\Delta_{\alpha_c}^+(\bar{Z}_c) = 0$. A normal approximation to the finite-book average $\bar{Z}_{c,J}$ suggests that its excess upper-tail cost decreases approximately as $J^{-1/2}$, motivating the heuristic $\Delta_{\alpha_c}^+(\bar{Z}_{c,J}) \approx \frac{\Delta_{\alpha_c}^+(Z_c)}{\sqrt{J}}$.*

The bound (6.12) establishes only an $\mathcal{O}(J^{-1/2})$ bound on the difference between $\Delta_{\alpha_c}^+(\bar{Z}_{c,J})$ and $\Delta_{\alpha_c}^+(\bar{Z}_c)$; it does not identify the coefficient in the heuristic. If instead the common market environment is non-degenerate, $\Delta_{\alpha_c}^+(\bar{Z}_{c,J})$ converges to $\Delta_{\alpha_c}^+(\bar{Z}_c)$, which is generally nonzero. The numerical experiments therefore estimate both quantities directly rather than impose the heuristic.

Under the independent-mortality specification of Section 2, retiree and spouse mortality are independent, and household mortality is independent of the shared asset returns. Consequently, $\bar{Z}_c$ has a probability-weighted representation. Specifically, for $\beta = 1, \dots, M-1$, define

$$\begin{aligned} a_\beta &:= \mathbb{P}(H_{(\beta-1)+} = 11, H_{\beta-} = 01) = \mathbb{P}(H_{(\beta-1)+} = 11) \delta_{R,\beta-1} (1 - \delta_{S,\beta-1}), \\ s_{\beta,m} &:= \mathbb{P}(H_{m-} = 01 \mid H_{(\beta-1)+} = 11, H_{\beta-} = 01) = \prod_{\ell=\beta}^{m-1} (1 - \delta_{S,\ell}), \quad m = \beta, \dots, M-1, \end{aligned} \quad (6.13)$$

with an empty product equal to one. Let $C_m^{(\beta)}$, $m = \beta, \dots, M-1$, denote the spouse-continuation payment at t_m for a household that enters state 01 at t_β^- and whose spouse remains alive through t_m^- . Then

$$\bar{Z}_c = \sum_{\beta=1}^{M-1} a_\beta \sum_{m=\beta}^{M-1} s_{\beta,m} C_m^{(\beta)}. \quad (6.14)$$

Equation (6.14) averages over all possible entry times and durations of the spouse-only phase while retaining the uncertainty in the shared asset returns. Its derivation is given in Appendix B.

For reporting purposes, when $\Delta_{\alpha_c}^+(Z_c) > 0$, define the large-book diversification ratio by

$$D_c(\nu^*) := 1 - \Delta_{\alpha_c}^+(\bar{Z}_c) / \Delta_{\alpha_c}^+(Z_c). \quad (6.15)$$

Conditional Jensen's inequality gives $(\mathbb{E}^{\nu^*}[Z_c \mid \mathcal{G}] - \eta)^+ \leq \mathbb{E}^{\nu^*}[(Z_c - \eta)^+ \mid \mathcal{G}]$ for each $\eta \in \mathbb{R}$. Hence, by (6.4) and (6.9), $D_c(\nu^*) \in [0, 1]$ and measures the fraction of the representative-contract excess upper-tail cost diversified away in the large-book limit.

6.3 Book-level risk-loaded cost and payment-scale load

For notational convenience, set $\bar{Z}_{c,\infty} := \bar{Z}_c$. For $J \in \mathbb{N} \cup \{\infty\}$ and prudential-buffer coefficient $\lambda_c \geq 0$, define

$$\Pi_{c,J}(\nu^*) := \mathbb{E}^{\nu^*}[\bar{Z}_{c,J}] + \lambda_c \Delta_{\alpha_c}^+(\bar{Z}_{c,J}), \quad (6.16)$$

$$f_{c,J}^{\text{RL}}(\nu^*) := \frac{\Pi_{c,J}(\nu^*)}{\mathbb{E}^{\nu^*}[Y_r] + \Pi_{c,J}(\nu^*)}. \quad (6.17)$$

Here, $\Pi_{c,J}(\nu^*)$ is the risk-loaded average per-contract continuation cost, including a prudential loading on the excess upper-tail cost remaining after book-level diversification, while $f_{c,J}^{\text{RL}}(\nu^*)$ is the corresponding normalized continuation-cost share under the fixed policy. For the representative-contract case $J = 1$, write $\Pi_c(\nu^*) := \Pi_{c,1}(\nu^*)$ and $f_c^{\text{RL}}(\nu^*) := f_{c,1}^{\text{RL}}(\nu^*)$.

When $\lambda_c = 0$, (6.17) reduces to the expected-cost load $f_c(\nu^*)$ in (6.3). Since $\mathbb{E}^{\nu^*}[\bar{Z}_{c,J}] = \mathbb{E}^{\nu^*}[Z_c]$ for every $J \in \mathbb{N} \cup \{\infty\}$, when $\lambda_c = 1$, (6.16) reduces to $\Pi_{c,J}(\nu^*) = \text{CVaR}_{\alpha_c}^+(\bar{Z}_{c,J})$. Moreover, (6.12) implies

$$|\Pi_{c,J}(\nu^*) - \Pi_{c,\infty}(\nu^*)| \leq \frac{\lambda_c \sigma_{c,\text{id}}(\nu^*)}{\alpha_c \sqrt{J}},$$

so the finite-book risk-loaded cost $\Pi_{c,J}(\nu^*)$ converges to its large-book limit $\Pi_{c,\infty}(\nu^*)$ with an $\mathcal{O}(J^{-1/2})$ error bound. If $\bar{Z}_{c,J}$ is deterministic, its excess upper-tail cost vanishes and $\Pi_{c,J}(\nu^*) = \mathbb{E}^{\nu^*}[Z_c]$.

6.4 Simulation-based numerical method

In numerical work, the exact optimal household control policy ν^* is replaced by the learned NN-parameterized household control policy $\hat{\nu}^*$ from Section 5. The learned policy is held fixed during load estimation, and the notation above is used with expectations taken under $\hat{\nu}^*$.

Representative-contract quantities are estimated from out-of-sample realizations of the joint asset-return and household-state processes. For each selected finite book size J and each Monte Carlo replication, one realization of the shared asset returns is common to all contracts in the book, while J conditionally independent household mortality histories are generated. The resulting finite-book averages $\bar{Z}_{c,J}$ are used to estimate $\text{CVaR}_{\alpha_c}^+(\bar{Z}_{c,J})$, $\Delta_{\alpha_c}^+(\bar{Z}_{c,J})$, $\Pi_{c,J}(\hat{\nu}^*)$, and $f_{c,J}^{\text{RL}}(\hat{\nu}^*)$, and to assess convergence to the large-book limit.

Separately, the large-book limiting variable $\bar{Z}_c$ is evaluated using (6.14), which integrates over household mortality conditional on each realization of the shared asset returns. These values estimate $\text{CVaR}_{\alpha_c}^+(\bar{Z}_c)$, $D_c(\hat{\nu}^*)$, $\Pi_{c,\infty}(\hat{\nu}^*)$, and $f_{c,\infty}^{\text{RL}}(\hat{\nu}^*)$. As a consistency check, the sample means of Z_c , $\bar{Z}_{c,J}$, and $\bar{Z}_c$ should agree up to Monte Carlo error.

The complete numerical algorithm and empirical estimators are given in Appendix C.

7 Numerical experiments

This section quantifies the economic and actuarial effects of spouse protection in the tontine. After specifying the baseline household, mortality inputs, financial data, and return-generation procedure, we examine household-state incidence, the household reward–risk frontier, the representative-contract spouse-continuation load, finite-book diversification, and the conditional large-book limit. We conclude with their implications for product design. The main analysis distinguishes the continuation-cost tail risk that is diversified across contracts from the common-market component that remains at the book level. Supplementary analyses of the continuation cap and learned household controls are reported in the appendices; Monte Carlo evaluation uncertainty is reported in Appendix H.

7.1 Baseline household scenario and mortality inputs

The baseline household consists of a male retiree aged $a_R = 65$ and a female spouse aged $a_S = 60$ at inception. With terminal reference age 95 for the younger protected life and annual decision times, the horizon is

$$T = M = \lceil 95 - \min\{65, 60\} \rceil = 35.$$

Unless otherwise noted, all monetary quantities are reported in thousands of real Australian dollars. The household starts with initial wealth $W_0 = 1000$ and, while solvent, is subject to the no-shorting and no-leverage constraints in Section 3. The annual tontine management fee is set to $\varrho = 0.11\%$, consistent with the fee structure of QSuper’s Lifetime Pension pool [33]. Debt positions accrue at the domestic bond-index return plus a borrowing spread of 2% per annum.

We calibrate the household withdrawal bands using the modest and comfortable spending estimates in the December 2025 quarter of the ASFA Retirement Standard [2], and set

$$q_{11} = 51.299, \quad \bar{q}_{11} = 77.375, \quad q_1 = 35.503, \quad \bar{q}_1 = 54.840.$$

The first pair applies to the two-person state 11. The second is the one-survivor band used for controlled withdrawals in state 10 and for the spouse-continuation rule in state 01.

The terminal targets represent capital reserves for spending beyond T , not annual payments within the horizon. Since the younger protected life reaches age 95 at T , we use the ASFA over-85 homeowner spending figures [2], rather than the age-65–84 figures used for the withdrawal bands. We choose a fixed ten-year spending multiple as the baseline, giving

$$w_{11}^* = 10 \times 72.093 = 720.93, \quad w_{10}^* = w_{01}^* = 10 \times 52.235 = 522.35.$$

For deterministic mortality, we use the 2021 single-year Australian male and female period life tables from the Human Mortality Database [21]. We hold these age-specific period schedules fixed over the horizon, so that only age advances. With $t_m = m$, the retiree and spouse one-year death probabilities are

$$\delta_{R,m-1} = q_{a_R+(m-1),2021}^M, \quad \delta_{S,m-1} = q_{a_S+(m-1),2021}^F, \quad m = 1, \dots, M.$$

The homogeneous baseline pool uses the same mortality sequences for every contract:

$$\delta_{R,m-1}^\ell \equiv \delta_{R,m-1}, \quad \delta_{S,m-1}^\ell \equiv \delta_{S,m-1}.$$

These sequences determine the household-state transitions (2.1)–(2.3), the contract-exit probabilities (2.4), and the state-dependent tontine gain rates (2.14).

7.2 Financial data and return generation

Having specified the household and mortality inputs, we now describe the market data used to train and evaluate the learned policies. We use a four-asset monthly return panel comprising Australian equities, Australian government bonds, U.S. equities converted to AUD, and U.S. 30-day Treasury bills converted to AUD. Following [28], the four series are aligned over the common sample period 1935:1–2022:12 and converted to real AUD returns using Australian CPI. Detailed data construction and descriptive statistics are provided in Appendix E.

Future return paths are generated nonparametrically by applying the stationary block bootstrap jointly to the four-dimensional monthly return vector [10, 29, 31, 32]. We use an expected block length of 24 months, thereby retaining serial dependence within each series and contemporaneous dependence across assets. The resulting real AUD return paths are used for NN training and out-of-sample policy evaluation. Network specifications and the training and evaluation samples are reported in Appendix A; Monte Carlo evaluation uncertainty is assessed in Appendix H.

7.3 Household-state probabilities and spouse-only duration

We first assess the incidence and persistence of spouse continuation under the baseline household and fixed-period mortality schedules in Subsection 7.1. Let

$$\pi_h(m) := \mathbb{P}(H_{m-} = h), \quad h \in \{11, 10, 01, 00\}, \quad m = 0, \dots, M,$$

denote the probability that the household is in state h at time t_m^- . The initial condition is $\pi_{11}(0) = 1$ and $\pi_{10}(0) = \pi_{01}(0) = \pi_{00}(0) = 0$. The analytic household-state probabilities are obtained directly from the

household-state transition law (2.1)–(2.3). Under the deterministic one-year death probabilities $\delta_{R,m-1}$ and $\delta_{S,m-1}$,

$$\begin{aligned}\pi_{11}(m) &= \pi_{11}(m-1)(1 - \delta_{R,m-1})(1 - \delta_{S,m-1}), \\ \pi_{10}(m) &= \pi_{10}(m-1)(1 - \delta_{R,m-1}) + \pi_{11}(m-1)(1 - \delta_{R,m-1})\delta_{S,m-1}, \\ \pi_{01}(m) &= \pi_{01}(m-1)(1 - \delta_{S,m-1}) + \pi_{11}(m-1)\delta_{R,m-1}(1 - \delta_{S,m-1}), \\ \pi_{00}(m) &= 1 - \pi_{11}(m) - \pi_{10}(m) - \pi_{01}(m).\end{aligned}$$


Thus, the full household-state probability profile is available analytically, year by year, from the mortality inputs alone. A particularly useful summary statistic is the probability that the household ever enters the spouse-only state. Since state 01 can only be entered from state 11, this probability is

$$\mathbb{P}(\exists m \in \{1, \dots, M\} : H_{m-} = 01) = \sum_{j=0}^{M-1} \left(\prod_{\ell=0}^{j-1} (1 - \delta_{R,\ell})(1 - \delta_{S,\ell}) \right) \delta_{R,j}(1 - \delta_{S,j}). \quad (7.1)$$

For the baseline Australian household, this gives

$$\mathbb{P}(\exists m \leq M : H_{m-} = 01) \approx 0.7100.$$

Hence, roughly 71% of baseline households are observed in the spouse-only state at least once by the terminal date.



(a) Analytic household-state probabilities over time.

(b) Analytic distribution of spouse-only duration D_{01} .

FIGURE 7.1: *Household-state probabilities and spouse-only duration for the baseline Australian household.*

Figure 7.1(a) reports the analytic household-state probabilities $\pi_{11}(m), \pi_{10}(m), \pi_{01}(m), \pi_{00}(m)$ over the full horizon. As expected, the probability of state 11 declines monotonically, while the probability of state 00 rises monotonically. More importantly, the spouse-only state 01 is materially more likely than the retiree-only state 10 over the economically relevant part of the horizon, reflecting both the younger spouse age in the baseline scenario and lower female mortality in the Australian life tables.

To quantify the length of the spouse-only phase, define $D_{01} := \sum_{m=1}^M \mathbf{1}_{\{H_{m-}=01\}}$, an annual state-count measure that includes the terminal observation. Its distribution is obtained analytically by combining the probability of entering state 01 at each date with the probability of remaining there for a given number of subsequent years before either the spouse dies or the horizon is reached. Figure 7.1(b) reports the resulting distribution. The unconditional mean spouse-only duration is approximately 8.70 years, and conditional on ever entering state 01 it is approximately 12.25 years. These state-duration statistics differ from the expected number of spouse-only payment dates: payments stop at $M-1$, giving $\sum_{m=1}^{M-1} \pi_{01}(m) \approx 8.5058$, consistent with the constant-payment diagnostic in Appendix H.

Overall, these results show that under realistic Australian mortality inputs, spouse continuation is both likely to arise and likely to persist for a substantial period. This provides direct motivation

for modeling spouse protection explicitly in the household decumulation problem. The corresponding international comparison in Appendix D shows that this qualitative pattern is not specific to Australia. A separate Monte Carlo validation based on 10^7 household paths, not reported here, closely matches the analytic formulas.

Because every contract enters the pool in state 11, spouse protection also weakens mortality pooling at the start of retirement. While both household members are alive, pool exit requires both deaths within the period, so the tontine gain is based on the joint-death probability $\delta_{R,m-1}\delta_{S,m-1}$ rather than an individual death probability. Spouse protection therefore delays forfeitures and attenuates mortality-credit gains while the contract remains in state 11.

7.4 Household reward–risk frontier

Having established that spouse continuation is both likely and persistent, we next examine the trade-off between expected cumulative real household payments and state-contingent terminal-shortfall risk. For each scalarization parameter λ in a non-uniform grid over $[0.15, 100]$, we train an NN-parameterized household control policy $\hat{\nu}_\lambda^*$ using the NN parameterization in Section 5 and the sequential transfer-learning protocol in Appendix A.

For later reference, define the ever-insolvent probability under the learned policy $\hat{\nu}_\lambda^*$ by

$$p_{\text{ins}}^{\text{ever}}(\hat{\nu}_\lambda^*) := \mathbb{P}_{x_0, h_0}^{\hat{\nu}_\lambda^*} \left(\exists m \in \{0, \dots, M-1\} : H_{m-} \neq 00 \text{ and } W_{m+}^{\hat{\nu}_\lambda^*} \leq 0 \right). \quad (7.2)$$

Thus, $p_{\text{ins}}^{\text{ever}}(\hat{\nu}_\lambda^*)$ is the probability that the contract becomes insolvent at least once while it remains active during the planning horizon.

From the computed frontier, we select the policies $\hat{\nu}_{0.3}^*$, $\hat{\nu}_1^*$, and $\hat{\nu}_{15}^*$, which we call *payment-seeking*, *intermediate*, and *conservative*, respectively. These policies span materially different reward–risk positions and are carried through the subsequent continuation-cost experiments. Table 7.1 reports their out-of-sample evaluation statistics.

TABLE 7.1: *Selected positions on the household reward–risk frontier and their associated learned policies.*

Frontier position	λ	$\mathcal{R}(\hat{\nu}_\lambda^*)$	$\text{CVaR}_{0.95}(L_T^{\hat{\nu}_\lambda^*})$	$p_{\text{ins}}^{\text{ever}}(\hat{\nu}_\lambda^*)$ (%)
Payment-seeking	0.3	1860.4	1218.0	26.7
Intermediate	1	1704.8	840.5	12.6
Conservative	15	1584.8	793.3	12.0

Table 7.1 shows that the payment-seeking policy generates the highest expected household payments but also the largest terminal-shortfall tail and insolvency probability. Moving to the intermediate policy reduces $\text{CVaR}_{0.95}(L_T)$ by approximately 31.0% and $p_{\text{ins}}^{\text{ever}}$ by 14.1 percentage points, while $\mathcal{R}$ decreases by approximately 8.4%. Moving from the intermediate to the conservative policy reduces terminal-shortfall CVaR by a further 5.6% and insolvency by only 0.6 percentage points, while reducing $\mathcal{R}$ by another 7.0%. The intermediate policy therefore lies near the visible knee of the frontier and is used as the reference policy for the learned household withdrawal and allocation controls presented in Appendix F.

For each trained policy, the corresponding frontier point is

$$(\mathcal{Q}_{0.95}(\hat{\nu}_\lambda^*), \mathcal{R}(\hat{\nu}_\lambda^*)) = \left(\text{CVaR}_{0.95}(L_T^{\hat{\nu}_\lambda^*}), \mathcal{R}(\hat{\nu}_\lambda^*) \right),$$

consistent with (4.8). The resulting frontier is shown in Figure 7.2.

Figure 7.2 shows the expected reward–risk trade-off: increasing λ lowers both cumulative household payments and terminal-shortfall risk, with diminishing marginal risk reduction beyond the intermediate position. The three selected policies therefore provide distinct reward–risk environments for assessing whether greater policy conservatism materially reduces the spouse-continuation obligation.




FIGURE 7.2: *Household reward-risk frontier under the baseline spouse-protected tontine specification. The square markers identify the three representative positions in Table 7.1.*

7.5 Representative-contract spouse-continuation load

We first evaluate the spouse-continuation obligation for one representative contract ($J = 1$) under the three selected learned policies. Throughout this subsection, the learned household policy is held fixed and the baseline continuation rule is $B_m = \min\{q_m, \bar{q}_1\}$. The retiree-alive and spouse-only payment streams $Y_r^{\hat{p}^*}$ and $Z_c^{\hat{p}^*}$ are defined in (6.1). We set $\alpha_c = 0.05$ and, unless otherwise stated, take the prudential-buffer coefficient to be $\lambda_c = 1$. Detailed capped-versus-uncapped sensitivity results under the same fixed learned policies are reported in Appendix G.

TABLE 7.2: *Representative-contract spouse-continuation load under the selected learned policies. The last column reports the baseline risk-loaded payment-scale load $f_{c,1}^{\text{RL}}$ at $\alpha_c = 0.05$ and $\lambda_c = 1$.*

Frontier position	λ	$\mathbb{E}[Y_r]$	$\mathbb{E}[Z_c]$	$\Delta_{0.05}^+(Z_c)$	f_c (%)	$f_{c,1}^{\text{RL}}$ (%)
Payment-seeking	0.3	1422.2	438.1	1096.2	23.5	51.9
Intermediate	1	1264.6	440.1	1043.8	25.8	54.0
Conservative	15	1151.7	433.1	1027.5	27.3	55.9

Table 7.2 shows that expected spouse-continuation payments remain in the narrow range $[433.1, 440.1]$, whereas expected retiree-alive payments fall by approximately 270.5 from the payment-seeking to the conservative policy. Consequently, f_c rises from 23.5% to 27.3% even though the absolute expected continuation cost changes by less than 2%. Across these policies, the increase in f_c is driven mainly by the lower retiree-alive payment scale, not by higher expected continuation cost. The same table reports a large representative-contract excess upper-tail cost, with $\Delta_{0.05}^+(Z_c)$ declining only modestly from 1096.2 to 1027.5. Under the baseline capped rule, relatively stable expected continuation costs therefore coexist with substantial representative-contract tail risk. This motivates the finite-book analysis in the next subsection.

Figure 7.3 shows how the risk-loaded payment-scale load varies with $\lambda_c \in [0, 1]$. At $\lambda_c = 1$, Table 7.2 reports loads of 51.9%–55.9%, substantially above the expected-cost loads.




FIGURE 7.3: *Representative-contract risk-loaded spouse-continuation load under the three learned policies.*

7.6 Finite-book spouse-continuation load

We next evaluate finite-book spouse-continuation costs for $J \in \{1, 10, 100, 1000\}$ under capped continuation. In each of the $K = 256,000$ replications, all contracts in a book share the same market-return path, while their household mortality histories are conditionally independent given that path. Smaller books use the first J histories from the same pool of 1,000 mortality draws for each market path; this shared sampling design is described in Appendix H.

TABLE 7.3: *Finite-book average per-contract spouse-continuation cost under capped continuation. The excess upper-tail cost is evaluated at $\alpha_c = 0.05$, and $f_{c,J}^{\text{RL}}$ is reported at $\lambda_c = 1$.*

Frontier position	J	$\mathbb{E}[\bar{Z}_{c,J}]$	$\Delta_{0.05}^+(\bar{Z}_{c,J})$	$f_{c,J}^{\text{RL}}$ (%)
Payment-seeking	1	438.1	1096.2	51.9
	10	437.7	327.6	35.0
	100	437.4	116.1	28.0
	1000	437.4	54.8	25.7
Intermediate	1	440.1	1043.8	54.0
	10	439.7	314.6	37.4
	100	439.5	107.2	30.2
	1000	439.4	46.9	27.8
Conservative	1	433.1	1027.5	55.9
	10	432.7	307.3	39.1
	100	432.5	103.3	31.7
	1000	432.4	44.8	29.3

Table 7.3 shows that, for each policy, the estimated mean continuation cost varies by less than 0.2% over $J \in \{1, 10, 100, 1000\}$, consistent with the mean identity in (6.7). The excess upper-tail cost, by contrast, falls by about 70% at $J = 10$, 89%–90% at $J = 100$, and 95%–96% at $J = 1000$, relative to the representative-contract value. At $\lambda_c = 1$, the risk-loaded load consequently falls from 51.9%–55.9% for $J = 1$ to 25.7%–29.3% for $J = 1000$. Book-level diversification therefore reduces the prudential tail buffer, not the expected continuation cost.

7.7 Large-book limiting spouse-continuation load

We next evaluate the conditional large-book limit $\bar{Z}_c = \mathbb{E}[Z_c \mid \mathcal{G}]$ in (6.8) under capped continuation. Conditional on each common market-return path, household mortality is integrated out, so the resulting distribution isolates the common-market component of continuation-cost risk. The probability-weighted evaluation is specified in Algorithm C.2 in Appendix C.

TABLE 7.4: *Representative-contract excess upper-tail cost and conditional large-book limiting average per-contract spouse-continuation cost under capped continuation. The excess upper-tail costs and diversification ratio are evaluated at $\alpha_c = 0.05$, and $f_{c,\infty}^{\text{RL}}$ is reported at $\lambda_c = 1$.*

Frontier position	$\Delta_{0.05}^+(Z_c)$	$\mathbb{E}[\bar{Z}_c]$	$\Delta_{0.05}^+(\bar{Z}_c)$	D_c (%)	$f_{c,\infty}^{\text{RL}}$ (%)
Payment-seeking	1096.2	437.4	29.0	97.35	24.7
Intermediate	1043.8	439.4	24.9	97.62	26.9
Conservative	1027.5	432.4	26.9	97.38	28.5

The limiting means in Table 7.4 differ from the representative-contract means in Table 7.2 by less than 0.2%, consistent with the tower-property identity (6.9). Book-level diversification therefore leaves the expected per-contract continuation cost unchanged. By contrast, the limiting excess upper-tail cost is only 24.9–29.0, compared with 1027.5–1096.2 for a representative contract. Accordingly, $D_c = 97.35\%$ –97.62%: diversification of household mortality removes almost all of the representative-contract excess

upper-tail cost under capped continuation. The smaller but nonzero limiting tail is the common-market component that remains even in an arbitrarily large book.

At $\lambda_c = 1$, Table 7.4 reports limiting risk-loaded payment-scale loads of 24.7%–28.5%, only about one percentage point above the corresponding expected-cost loads in Table 7.2. Moreover, the $J = 1000$ loads in Table 7.3 are only 0.8–1.0 percentage points above their large-book limits. Thus, large books remove most of the prudential loading associated with household-specific mortality risk, but not the underlying expected continuation cost or the residual common-market tail. The corresponding capped-versus-uncapped large-book comparison is reported in Appendix G.2.

To compare finite books with this limit, define, for $J > 1$,

$$R_{c,J} := \frac{\Delta_{0.05}^+(\bar{Z}_{c,J}) - \Delta_{0.05}^+(\bar{Z}_c)}{\Delta_{0.05}^+(Z_c) - \Delta_{0.05}^+(\bar{Z}_c)},$$

whenever the denominator is positive. Figure 7.4 compares the estimates with the $J^{-1/2}$ benchmark from Remark 6.1.




FIGURE 7.4: *Finite-book convergence toward the conditional large-book limiting excess upper-tail spouse-continuation cost under capped continuation. The plotted statistic is $R_{c,J}$; the reference curve is $J^{-1/2}$.*

Across the three policies, $R_{c,J}$ is approximately 0.280–0.284 at $J = 10$, 0.076–0.082 at $J = 100$, and 0.018–0.024 at $J = 1000$. These values are broadly consistent with, but generally below, the square-root benchmarks 0.316, 0.100, and 0.0316. The experiment supports the predicted rapid approach to the conditional large-book limit; it does not establish an empirical convergence rate.

7.8 Implications for spouse-protected tontine design

Household risk control and continuation funding must be assessed together. Tables 7.1 and 7.2 show that greater policy conservatism reduces terminal-shortfall risk but raises the expected-cost continuation load, chiefly by lowering retiree-alive payments. Lower terminal-shortfall risk therefore need not imply a lower continuation load.

Book size should affect the prudential buffer, not the expected per-contract obligation. Using the representative-contract excess tail to set a prudential buffer for a diversified book would fail to reflect the diversification benefit shown in Tables 7.3 and 7.4. Scaling it toward zero ignores residual common-market risk. The load should reflect the book’s remaining tail risk, with expected cost retained separately. This is a model-based funding assessment, not a market price or regulatory capital requirement.

The continuation cap and book diversification are not substitutes. The fixed-policy comparisons in Appendix G show that the cap reduces expected cost and the residual common-market tail, neither of which disappears with book size. This funding benefit must be weighed against the restriction on the payment inherited by the spouse. The comparison quantifies the effect of the rule; it does not establish an optimal cap or compare designs after policy reoptimization.

8 Conclusion

This paper develops a tontine with spouse protection at the household-contract level. Pool exit occurs at household extinction rather than at the retiree’s death, leading to household-state-dependent mortality

credits and exact ex post budget balance at the finite-pool level. The household chooses withdrawals and portfolio rebalancing while the retiree is alive, whereas the spouse-only phase follows a rule-based continuation payment and passive rebalancing convention. Admissible global-in-time neural-network policies approximate the solution of the resulting multidimensional reward–risk problem. The associated policy-dependent continuation obligation is evaluated at both the representative-contract and book levels, with the conditional large-book limit isolating the common-market risk that remains after household-mortality diversification.

The numerical results show that spouse continuation is frequent and persistent, with a substantial expected cost that varies little across the selected learned policies under capped continuation. In the baseline capped experiments, book-level diversification reduces the representative-contract excess upper-tail cost by approximately 97% in the conditional large-book limit across the selected learned policies. Expected per-contract continuation cost is unchanged, and a nonzero common-market tail remains. The fixed-policy capped-versus-uncapped comparisons further show that the continuation rule affects both the expected obligation and the residual market-sensitive tail. The cross-country mortality comparison indicates that the prevalence and persistence of the spouse-only phase are not specific to Australia.

The framework therefore links household-extinction pooling, endogenous decumulation, and the funding of spouse continuation within a single model. Its central implication is that continuation funding must distinguish the expected obligation from a prudential tail buffer assessed at book level: neither representative-contract risk nor its mechanical scaling toward zero captures the residual exposure of a large book. Future work could incorporate dependent spousal mortality, heterogeneous household types, stochastic mortality, and separately optimized alternative continuation designs.

Appendices

A Neural-network training and evaluation specifications

The reported policies use two separate global-in-time, fully connected feedforward networks: one for the withdrawal control and one for the rebalancing control. At each decision time, both networks receive three inputs: time-to-go, standardized contract wealth, and the current household state. Table A.1 summarizes the architecture and the principal training specifications used in the production run.

The scalarization parameters are trained sequentially over the ordered grid

$$\Lambda = \{0.15, 0.16, 0.17, 0.18, 0.185, 0.2, 0.25, 0.3, 0.5, 0.7, 1, 1.5, 2, 2.5, 3, 5, 7.5, 10, 15, 20, 50, 100\}.$$

As λ increases, the objective places greater weight on $\text{CVaR}_{0.95}(L_T)$. The gradient contribution from this risk term is driven primarily by realizations in the upper tail of the terminal shortfall distribution and can therefore be sparse within a training mini-batch. This makes the more risk-sensitive policies harder to train from an independent random initialization.

Following the sequential transfer-learning approach of [7] and the broader transfer-learning literature [40], the policy at $\lambda = 0.15$ is initialized randomly (a cold start). For each subsequent λ , the withdrawal-network parameters, rebalancing-network parameters, and CVaR threshold are initialized from the immediately preceding trained solution. The ordering of Λ is therefore part of the numerical protocol. In related decumulation experiments, this continuation strategy was found to improve training stability and reduce the incidence of convergence to inferior local solutions at more risk-sensitive frontier points [7].

During each training run, running-best checkpointing retains the parameter set that attains the best observed empirical objective, rather than necessarily using the final iterate. After training, the learned policies are held fixed and evaluated on fresh joint asset-return and household-state realizations. For the selected policies $\lambda = 0.3, 1$, and 15 , Monte Carlo uncertainty is assessed using 100 independent evaluation batches of 2,560 realizations each. Common batch realizations are used across the selected

TABLE A.1: *Neural-network architecture and training specifications.*

Specification	Withdrawal network	Rebalancing network
Inputs	Time-to-go, standardized wealth, household state	
Hidden layers	2	2
Nodes per hidden layer	10	10
Hidden-layer activation	Logistic sigmoid	Logistic sigmoid
Output dimension	1	4
Output transformation	Scaled sigmoid, (5.5)	Four-way softmax, (5.6)
Bias terms	Yes	Yes
Batch/layer normalization	No	No
Dropout	No	No
Training realizations	256,000	
Training horizon	$T = M = 35$	
Stationary-bootstrap mean block length	24 months	
Training mini-batch size	1,000	
Number of training iterations	30,000	
CVaR confidence level	$\alpha = 0.95$	
Initial CVaR threshold	$\eta_0 = 100$	
Policy-network optimizer	Adam	
Initial policy-network learning rate	0.04	
Policy-network Adam coefficients	$(\beta_1, \beta_2) = (0.9, 0.998)$	
Policy-network weight decay	10^{-4}	
CVaR-threshold optimizer	Adam	
Initial CVaR-threshold learning rate	0.05	
CVaR-threshold regularization	None	
Learning-rate schedule	Multistep	
Random seed	5	
Sequential transfer learning	Yes	
Running-best checkpointing	Yes	

policies and in the paired capped–uncapped continuation comparisons. The shared-path evaluation design and batch-based uncertainty estimates are detailed in Appendix H.

B Derivation of (6.14)

Recall that $Z_c = \sum_{m=0}^{M-1} C_m \mathbf{1}_{\{H_{m-}=01\}}$ and $\bar{Z}_c = \mathbb{E}^{\nu^*}[Z_c \mid \mathcal{G}]$. Since state 01 can be entered at most once, through a transition $11 \rightarrow 01$,

$$Z_c = \sum_{\beta=1}^{M-1} \sum_{m=\beta}^{M-1} C_m^{(\beta)} \mathbf{1}_{\{H_{(\beta-1)+}=11, H_{\beta-}=01, H_{m-}=01\}},$$

where, on the event in the indicator, $C_m = C_m^{(\beta)}$, and $C_m^{(\beta)}$ is $\mathcal{G}$ -measurable. Hence,

$$\begin{aligned} \bar{Z}_c &= \sum_{\beta=1}^{M-1} \sum_{m=\beta}^{M-1} C_m^{(\beta)} \mathbb{P}(H_{(\beta-1)+}=11, H_{\beta-}=01, H_{m-}=01 \mid \mathcal{G}) \\ &= \sum_{\beta=1}^{M-1} \sum_{m=\beta}^{M-1} C_m^{(\beta)} \mathbb{P}(H_{(\beta-1)+}=11, H_{\beta-}=01, H_{m-}=01) = \sum_{\beta=1}^{M-1} \sum_{m=\beta}^{M-1} C_m^{(\beta)} a_{\beta} s_{\beta,m} = \sum_{\beta=1}^{M-1} a_{\beta} \sum_{m=\beta}^{M-1} s_{\beta,m} C_m^{(\beta)}, \end{aligned}$$

where the second equality uses independence of household mortality and the shared asset returns. This gives (6.14).

C Simulation algorithm for spouse-continuation loads

The two algorithms below implement the finite-book and large-book spouse-continuation calculations in Section 6. Algorithm C.1 estimates the finite-book quantities for a selected book size J , with $J = 1$ corresponding to the representative-contract case. Algorithm C.2 evaluates the large-book limiting quantities by integrating over household mortality using (6.14). In the algorithms below, the “continuation rule” refers to the choice between capped and uncapped continuation: $B_m = \min\{q_m, \bar{q}_1\}$ under capped continuation and $B_m = q_m$ under uncapped continuation.

C.1 Representative-contract or finite-book estimates

Let K be the number of Monte Carlo replications and let $J \geq 1$ be the book size. Within each replication, all J contracts share the same asset returns, while their household mortality histories are conditionally independent.

ALGORITHM C.1: *Representative-contract or finite-book estimates*

- 1: choose the number of evaluation realizations K , book size J , prudential-buffer coefficient $\lambda_c \geq 0$, upper-tail probability $\alpha_c \in (0, 1)$, and continuation rule;
- 2: initialize arrays $\bar{\mathcal{Y}}_J \leftarrow [0, \dots, 0]$ and $\bar{\mathcal{Z}}_J \leftarrow [0, \dots, 0]$, each of length K ;
- 3: **for** $k = 1, \dots, K$ **do**
- 4: generate one realization $\{\varepsilon_{m+1}^{(k)}\}_{m=0}^{M-1}$ of the exogenous return drivers in (3.2);
- 5: **for** $j = 1, \dots, J$ **do**
- 6: initialize $X_{0-}^{(k,j)} \leftarrow x_0$, $H_{0-}^{(k,j)} \leftarrow 11$, $B_{-1}^{(k,j)} \leftarrow 0$, $Y_r^{(k,j)} \leftarrow 0$, and $Z_c^{(k,j)} \leftarrow 0$;
- 7: **for** $m = 0, \dots, M - 1$ **do**
- 8: compute $\mathcal{W}_{m-}^{(k,j)}$, $\widetilde{W}_{m-}^{(k,j)}$, and $W_{m-}^{(k,j)}$ using (3.4)–(3.7);
- 9: **if** $H_{m-}^{(k,j)} \in \mathcal{H}_R$ **then**
- 10: **if** $W_{m-}^{(k,j)} > 0$ **then**
- 11: set $q_m^{(k,j)} \leftarrow \hat{q}_m^*(W_{m-}^{(k,j)}, t_m, H_{m-}^{(k,j)})$;
- 12: **else**
- 13: set $q_m^{(k,j)} \leftarrow \underline{q}(H_{m-}^{(k,j)})$;
- 14: **end if**
- 15: set $C_m^{(k,j)} \leftarrow q_m^{(k,j)}$;
- 16: set $B_m^{(k,j)} \leftarrow \min\{q_m^{(k,j)}, \bar{q}_1\}$ under capped continuation, and $B_m^{(k,j)} \leftarrow q_m^{(k,j)}$ under uncapped continuation;
- 17: update $Y_r^{(k,j)} \leftarrow Y_r^{(k,j)} + C_m^{(k,j)}$ and $W_{m+}^{(k,j)} \leftarrow W_{m-}^{(k,j)} - C_m^{(k,j)}$;
- 18: **if** $W_{m+}^{(k,j)} > 0$ **then**
- 19: compute $\hat{p}_m^*(W_{m+}^{(k,j)}, t_m, H_{m-}^{(k,j)})$ and update $X_{m+}^{(k,j)}$ using (3.18);
- 20: **else**
- 21: update $X_{m+}^{(k,j)}$ using (3.22);
- 22: **end if**
- 23: **else if** $H_{m-}^{(k,j)} = 01$ **then**
- 24: set $B_m^{(k,j)} \leftarrow B_{m-1}^{(k,j)}$;
- 25: set $C_m^{(k,j)} \leftarrow B_{m-1}^{(k,j)}$ if $W_{m-}^{(k,j)} > 0$, and $C_m^{(k,j)} \leftarrow \underline{q}_1$ otherwise;
- 26: update $Z_c^{(k,j)} \leftarrow Z_c^{(k,j)} + C_m^{(k,j)}$ and $W_{m+}^{(k,j)} \leftarrow W_{m-}^{(k,j)} - C_m^{(k,j)}$;
- 27: **if** $\mathcal{W}_{m-}^{(k,j)} > 0$, $W_{m-}^{(k,j)} > 0$, and $W_{m+}^{(k,j)} > 0$ **then**
- 28: update $X_{m+}^{(k,j)}$ passively using (3.19)–(3.20);
- 29: **else**
- 30: update $X_{m+}^{(k,j)}$ using (3.22);

```

31:         end if
32:     else
33:         set  $B_m^{(k,j)} \leftarrow 0$ ,  $C_m^{(k,j)} \leftarrow 0$ ,  $W_{m^+}^{(k,j)} \leftarrow 0$ , and  $X_{m^+}^{(k,j)} \leftarrow (0, 0, 0, 0)^\top$ ;
34:     end if
35:     propagate  $X_{(m+1)^-}^{(k,j)}$  from  $X_{m^+}^{(k,j)}$  using  $\varepsilon_{m+1}^{(k)}$  via (3.2);
36:     simulate  $H_{(m+1)^-}^{(k,j)}$  from  $H_{m^+}^{(k,j)} = H_{m^-}^{(k,j)}$  using (2.1)–(2.3);
37: end for
38: end for
39: set  $\bar{Y}_{r,J}^{(k)} \leftarrow \frac{1}{J} \sum_{j=1}^J Y_r^{(k,j)}$  and  $\bar{Z}_{c,J}^{(k)} \leftarrow \frac{1}{J} \sum_{j=1}^J Z_c^{(k,j)}$ ;
40: set  $\bar{\mathcal{Y}}_J^{(k)} \leftarrow \bar{Y}_{r,J}^{(k)}$  and  $\bar{\mathcal{Z}}_J^{(k)} \leftarrow \bar{Z}_{c,J}^{(k)}$ ;
41: end for
42: compute  $\hat{\mathbb{E}}[Y_r] \leftarrow K^{-1} \sum_{k=1}^K \bar{\mathcal{Y}}_J^{(k)}$  and  $\hat{\mathbb{E}}[Z_c] \leftarrow K^{-1} \sum_{k=1}^K \bar{\mathcal{Z}}_J^{(k)}$ ;
43: sort  $\bar{\mathcal{Z}}_J$  in descending order as  $\bar{Z}_{c,J,(1)} \geq \dots \geq \bar{Z}_{c,J,(K)}$ ;
44: compute  $\widehat{\text{CVaR}}_{\alpha_c}^+(\bar{Z}_{c,J}) \leftarrow ([\alpha_c K])^{-1} \sum_{k=1}^{[\alpha_c K]} \bar{Z}_{c,J,(k)}$ ;
45: compute  $\hat{\Delta}_{\alpha_c}^+(\bar{Z}_{c,J}) \leftarrow \widehat{\text{CVaR}}_{\alpha_c}^+(\bar{Z}_{c,J}) - \hat{\mathbb{E}}[Z_c]$ ;
46: compute  $\hat{f}_c \leftarrow \frac{\hat{\mathbb{E}}[Z_c]}{\hat{\mathbb{E}}[Y_r] + \hat{\mathbb{E}}[Z_c]}$ ;
47: compute  $\hat{\Pi}_{c,J} \leftarrow \hat{\mathbb{E}}[Z_c] + \lambda_c \hat{\Delta}_{\alpha_c}^+(\bar{Z}_{c,J})$  and  $\hat{f}_{c,J}^{\text{RL}} \leftarrow \frac{\hat{\Pi}_{c,J}}{\hat{\mathbb{E}}[Y_r] + \hat{\Pi}_{c,J}}$ ;
48: return  $\hat{\mathbb{E}}[Y_r]$ ,  $\hat{\mathbb{E}}[Z_c]$ ,  $\widehat{\text{CVaR}}_{\alpha_c}^+(\bar{Z}_{c,J})$ ,  $\hat{\Delta}_{\alpha_c}^+(\bar{Z}_{c,J})$ ,  $\hat{f}_c$ ,  $\hat{\Pi}_{c,J}$ , and  $\hat{f}_{c,J}^{\text{RL}}$ .

```

When $J = 1$, $\bar{Y}_{r,1}^{(k)} = Y_r^{(k,1)}$ and $\bar{Z}_{c,1}^{(k)} = Z_c^{(k,1)}$, giving the representative-contract estimates; $J > 1$ gives the corresponding finite-book estimates.

C.2 Large-book limiting estimates

The large-book limiting average per-contract cost is evaluated using (6.14). For each possible entry time t_β^- , household mortality is not simulated: state 11 is imposed before t_β^- , state 01 is imposed from t_β^- onward, and the weights $s_{\beta,m}$ account for spouse survival through subsequent payment dates.

ALGORITHM C.2: *Large-book limiting estimates*

```

1: generate  $K$  realizations  $\{\varepsilon_{m+1}^{(k)}\}_{m=0}^{M-1}$ ,  $k = 1, \dots, K$ , of the exogenous return drivers in (3.2);
2: use the same  $\alpha_c$ ,  $\lambda_c$ , and the continuation rule, as in Algorithm C.1 with  $J = 1$ ;
   input the resulting representative-contract estimates  $\hat{\mathbb{E}}[Y_r]$  and  $\hat{\Delta}_{\alpha_c}^+(Z_c)$ ;
3: input the weights  $\{a_\beta, s_{\beta,m}\}$  from (6.13);
4: initialize  $\bar{\mathcal{Z}} \leftarrow [0, \dots, 0]$  of length  $K$ ;
5: for  $k = 1, \dots, K$  do
6:     for  $\beta = 1, \dots, M - 1$  do
7:         initialize  $X_{0^-}^{(k,\beta)} \leftarrow x_0$ ,  $B_{-1}^{(k,\beta)} \leftarrow 0$ ,  $Y_r^{(k,\beta)} \leftarrow 0$ , and  $Z_c^{(k,\beta)} \leftarrow 0$ ;
8:         impose  $H_{m^-}^{(k,\beta)} = H_{m^+}^{(k,\beta)} = 11$  for  $m < \beta$  and  $H_{m^-}^{(k,\beta)} = H_{m^+}^{(k,\beta)} = 01$  for  $m \geq \beta$ ;
9:         when computing the period- $\beta$  tontine gain, use  $H_{(\beta-1)^+}^{(k,\beta)} = 11$ ;
10:        for  $m = 0, \dots, M - 1$  do
11:            apply lines 8–35 of Algorithm C.1, with  $j$  replaced by  $\beta$ , using  $\varepsilon_{m+1}^{(k)}$  and the imposed household states;

```

```

12:   end for
13:   retain  $C_m^{(k,\beta)}$ ,  $m = \beta, \dots, M-1$ ;
14:   end for
15:   set  $\bar{Z}_c^{(k)} \leftarrow \sum_{\beta=1}^{M-1} a_\beta \sum_{m=\beta}^{M-1} s_{\beta,m} C_m^{(k,\beta)}$  and  $\bar{Z}^{(k)} \leftarrow \bar{Z}_c^{(k)}$ ;
16: end for
17: compute  $\hat{\mathbb{E}}[\bar{Z}_c] \leftarrow K^{-1} \sum_{k=1}^K \bar{Z}^{(k)}$ ;
18: sort  $\bar{Z}$  in descending order as  $\bar{Z}_{c,(1)} \geq \dots \geq \bar{Z}_{c,(K)}$ ;
19: compute  $\widehat{\text{CVaR}}_{\alpha_c}^+(\bar{Z}_c) \leftarrow ([\alpha_c K])^{-1} \sum_{k=1}^{[\alpha_c K]} \bar{Z}_{c,(k)}$ ;
20: compute  $\hat{\Delta}_{\alpha_c}^+(\bar{Z}_c) \leftarrow \widehat{\text{CVaR}}_{\alpha_c}^+(\bar{Z}_c) - \hat{\mathbb{E}}[\bar{Z}_c]$ ;
21: if  $\hat{\Delta}_{\alpha_c}^+(\bar{Z}_c) > 0$ , compute  $\hat{D}_c \leftarrow 1 - \frac{\hat{\Delta}_{\alpha_c}^+(\bar{Z}_c)}{\hat{\Delta}_{\alpha_c}^+(\bar{Z}_c)}$ ;
22: compute  $\hat{\Pi}_{c,\infty} \leftarrow \hat{\mathbb{E}}[\bar{Z}_c] + \lambda_c \hat{\Delta}_{\alpha_c}^+(\bar{Z}_c)$ ;
23: compute  $\hat{f}_{c,\infty}^{\text{RL}} \leftarrow \frac{\hat{\Pi}_{c,\infty}}{\hat{\mathbb{E}}[Y_r] + \hat{\Pi}_{c,\infty}}$ ;
24: return  $\hat{\mathbb{E}}[\bar{Z}_c]$ ,  $\widehat{\text{CVaR}}_{\alpha_c}^+(\bar{Z}_c)$ ,  $\hat{\Delta}_{\alpha_c}^+(\bar{Z}_c)$ ,  $\hat{D}_c$  when defined,  $\hat{\Pi}_{c,\infty}$ , and  $\hat{f}_{c,\infty}^{\text{RL}}$ .

```

D International household-state sensitivity

To assess whether the importance of spouse continuation is specific to the Australian mortality calibration, we repeat the analytic household-state calculation using the 2021 male and female period life tables from the Human Mortality Database for Australia, Canada, the United States, Japan, and the Netherlands [21]. In every country, the household consists at inception of a male retiree aged 65 and a female spouse aged 60 and is followed for 35 years, until the spouse reaches age 95. We hold the 2021 age-specific mortality schedule fixed over the horizon and retain the baseline assumption that the two household lifetimes are independent. Durations use the same annual state count D_{01} , including the terminal observation, as in Subsection 7.3.

TABLE D.1: *Household-state probabilities and spouse-only duration under selected 2021 HMD period life tables. The household consists of a male retiree aged 65 and a female spouse aged 60 and is followed over a 35-year horizon.*

Country	Probability spouse-only phase occurs (%)	Mean spouse-only duration (years)	Conditional mean duration given entry (years)
Australia	71.0	8.70	12.25
Canada	70.2	9.03	12.85
United States	68.1	9.09	13.34
Japan	77.0	10.23	13.28
Netherlands	71.0	8.84	12.44

Table D.1 shows that the probability of entering the spouse-only phase ranges from approximately 68.1% in the United States to 77.0% in Japan. Its unconditional mean duration ranges from 8.70 to 10.23 years, while the conditional mean given entry lies between 12.25 and 13.34 years. Canada and the Netherlands are close to the Australian baseline; Japan has the highest entry probability and unconditional duration, while the United States has the lowest entry probability but the longest conditional duration. Thus, the conclusion that spouse continuation is both likely and long-lasting is robust across the selected mortality environments. This comparison isolates mortality effects; it does not recalibrate financial returns, spending bands, or product rules for each country.

E Financial data construction and return characteristics

The four-asset investment universe consists of Australian equities, Australian government bonds, U.S. equities, and U.S. 30-day Treasury bills. The underlying monthly data are aligned over the common sample period 1935:1–2022:12 and expressed in real Australian dollars.

The Australian equity series is a long-horizon capitalization-weighted total-return index constructed from Bloomberg All Ordinaries data together with historical dividend-yield information from [25] and the Reserve Bank of Australia. The Australian bond series is a long-run government-bond total-return index formed by combining the historical government-bond return series of [22] and [8] with a recent Bloomberg Australian government-bond total-return index.

TABLE E.1: *Summary statistics of monthly real returns, 1935:1–2022:12. The first three return measures are annualized. VaR_{0.05} and CVaR_{0.05} refer to the lower 5% tail of monthly returns; returns are reported as decimals.*

Asset	Mean (ann.)	Geometric mean (ann.)	Volatility (ann.)	VaR _{0.05}	CVaR _{0.05}
U.S. 30-day Treasury 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
Australian 10-year government bond	0.012	0.012	0.036	-0.015	-0.025
Australian equity index	0.069	0.059	0.151	-0.065	-0.100

TABLE E.2: *Correlation matrix of monthly real AUD returns, 1935:1–2022:12.*

	U.S. 30-day Treasury bill	U.S. equity	Australian bond	Australian equity
U.S. 30-day Treasury bill	1.00	0.34	0.17	-0.22
U.S. equity	0.34	1.00	0.10	0.33
Australian bond	0.17	0.10	1.00	0.11
Australian equity	-0.22	0.33	0.11	1.00

The U.S. asset block consists of a broad equity total-return series and a 30-day Treasury-bill total-return series. Both are converted from USD to AUD using an end-of-month AUD/USD exchange-rate series compiled from Federal Reserve archives and Reserve Bank of Australia data. All four nominal return series are converted to real AUD returns using Australian CPI. Further details on data construction, splicing, currency conversion, and preprocessing are given in [28].


Tables E.1 and E.2 report the main return characteristics and dependence structure of the resulting monthly real AUD panel.

U.S. equity has the highest historical average real return in the sample and is only moderately correlated with Australian equity. The full-sample correlation is approximately 0.33, increasing to 0.48 over 1990:1–2022:12. By contrast, after conversion to real AUD, the U.S. Treasury bill has a near-zero average return and substantial exchange-rate-induced volatility. The lower historical volatility of Australian government bonds helps explain their defensive role in the learned allocations reported in Appendix F.


F Learned household withdrawal and allocation controls

This appendix presents the learned withdrawal and four-asset allocation controls under the intermediate policy $\hat{\nu}_1^*$ selected in Subsection 7.4. Controls are defined only in the retiree-alive states $\mathcal{H}_R = \{11, 10\}$; state 01 follows the continuation-payment and passive-rebalancing rules. We examine withdrawal behavior first, followed by the allocation pattern across funding levels and household states.

Figure F.1 shows that the learned withdrawals are close to the relevant lower or upper band over most of the wealth–time domain. The wealth threshold for the maximum withdrawal is generally higher in state 11 than in state 10, leaving a larger minimum-withdrawal region when both members are alive.



(a) Household state 11.




(b) Household state 10.


FIGURE F.1: *Learned withdrawal controls under the intermediate policy $\hat{v}_1^*$. Colors report annual controlled withdrawals as functions of time and pre-payment wealth.*

This is consistent with the larger two-person payment band and reserve target, together with weaker mortality-credit support under the joint-death pool-exit rule.


Figures F.2 and F.3 together report the complete four-asset allocation under the same learned policy, with paired panels for states 11 and 10.




(a) Australian bond, state 11.



(b) Australian bond, state 10.




(c) U.S. equity, state 11.




(d) U.S. equity, state 10.


FIGURE F.2: *Learned Australian government-bond allocations (top row) and U.S. equity allocations (bottom row) under the intermediate policy $\hat{v}_1^*$. Colors denote fractions of post-payment wealth.*




(a) Australian equity, state 11.



(b) Australian equity, state 10.



(c) U.S. Treasury bill, state 11.



(d) U.S. Treasury bill, state 10.

FIGURE F.3: *Learned Australian equity allocations (top row) and U.S. Treasury-bill allocations (bottom row) under the intermediate policy $\hat{v}_1^*$. Colors denote fractions of post-payment wealth.*

The allocation surfaces exhibit a growth-to-preservation pattern. U.S. equity exposure is largest in poorly funded regions, particularly early in the horizon (Figure F.2), while Australian equity is used mainly at low and intermediate funding levels (Figure F.3). Once wealth is sufficiently high, the Australian government-bond allocation rises sharply; the U.S. Treasury-bill component remains small over most of the plotted domain. The learned policy therefore shifts from growth seeking when underfunded to domestic-bond wealth preservation when the terminal reserve target is well funded.

The shift toward Australian government bonds generally occurs at a higher wealth level in state 11 (Figure F.2). This mirrors the higher withdrawal threshold in Figure F.1 and is consistent with the greater spending and reserve requirements and weaker mortality-credit support while both household members remain alive.

G Continuation-cap sensitivity

This appendix compares capped and uncapped spouse continuation under the three selected learned policies. Each policy is held fixed, with no retraining under the uncapped rule. We first examine representative-contract continuation costs and household downside risk, then consider the cap's effect on the conditional large-book limit.

G.1 Representative-contract comparison

For a representative contract, we evaluate each learned policy under

$$B_m^{\text{cap}} = \min\{q_m, \bar{q}_1\}, \quad B_m^{\text{uncap}} = q_m.$$

The asset-return and household-mortality paths are identical in each paired comparison, so only the continuation-base convention changes. For τ_R defined in (3.11), define the unconditional cap-binding probability and mean cap excess by

$$p_{\text{cap}}(\hat{\nu}_\lambda^*) := \mathbb{P}_{x_0, h_0}^{\hat{\nu}_\lambda^*} \left(q_{\tau_R-1}^{\hat{\nu}_\lambda^*} > \bar{q}_1, \tau_R < \infty \right), \quad (\text{G.1})$$

$$e_{\text{cap}}(\hat{\nu}_\lambda^*) := \mathbb{E}_{x_0, h_0}^{\hat{\nu}_\lambda^*} \left[\left(q_{\tau_R-1}^{\hat{\nu}_\lambda^*} - \bar{q}_1 \right)_+ \mathbf{1}_{\{\tau_R < \infty\}} \right]. \quad (\text{G.2})$$

TABLE G.1: *Representative-contract cap incidence, expected continuation costs, and expected-cost loads for the three selected learned policies.*

Frontier position	λ	p_{cap} (%)	e_{cap}	$\mathbb{E}[Z_c]$ capped	$\mathbb{E}[Z_c]$ uncapped	f_c capped (%)	f_c uncapped (%)
Payment-seeking	0.3	50.3	10.9	438.1	555.6	23.5	28.1
Intermediate	1	38.7	8.4	440.1	528.0	25.8	29.5
Conservative	15	22.7	5.0	433.1	476.3	27.3	29.3

Table G.1 shows that the cap-binding probability and unconditional mean cap excess decrease across the selected policies as conservatism increases. Removing the cap raises both expected continuation cost and the expected-cost load under every policy, with the largest increases under the payment-seeking policy. Expected continuation costs remain close across policies under capped continuation but are substantially more dispersed under the uncapped rule. Thus, limiting the inherited continuation base to the one-survivor upper band $\bar{q}_1$ stabilizes the expected continuation obligation across the selected policies.




FIGURE G.1: *Representative-contract risk-loaded spouse-continuation loads under capped and uncapped continuation. The policy is held fixed, $\alpha_c = 0.05$, and the excess-CVaR loading is used.*

Figure G.1 compares the risk-loaded payment-scale loads as the prudential-buffer coefficient λ_c varies. Unlike the expected-cost loads in Table G.1, these loads incorporate the continuation-cost excess tail. Their values at $\lambda_c = 1$, together with the reported changes in terminal-shortfall risk and insolvency, are collected in Table G.2.

Table G.2 shows that removing the cap raises the risk-loaded continuation load and worsens both household downside-risk measures under the fixed learned policies. The increases are largest under the payment-seeking policy and smallest under the conservative policy. Together with Table G.1, these results show that the cap has its strongest effect near the high-payment end of the frontier, where larger withdrawals can otherwise be inherited by the surviving spouse.

G.2 Conditional large-book comparison

The representative-contract comparison includes both household-specific mortality risk and common-market risk. We now compare the two continuation rules in the conditional large-book limit, where

TABLE G.2: *Representative-contract risk-loaded loads and reported changes in household downside risk when the continuation cap is removed. Loads are evaluated at $\alpha_c = 0.05$ and $\lambda_c = 1$. Changes are uncapped minus capped; terminal-shortfall CVaR changes are in thousands of real Australian dollars, and insolvency-probability changes are in percentage points.*

Frontier position	$f_{c,1}^{\text{RL}}$ capped (%)	$f_{c,1}^{\text{RL}}$ uncapped (%)	Change in $\text{CVaR}_{0.95}(L_T)$	Change in $p_{\text{ins}}^{\text{ever}}$
Payment-seeking	51.9	58.8	35.4	5.7
Intermediate	54.0	58.2	31.6	0.4
Conservative	55.9	57.2	1.1	< 0.1

household mortality is integrated out. This separates the cap’s effect on expected continuation cost from its effect on the residual common-market tail. For each rule, D_c compares the limiting excess tail with the representative-contract excess tail under that same rule.

TABLE G.3: *Conditional large-book continuation-cap sensitivity. The upper-tail probability is $\alpha_c = 0.05$; $f_{c,\infty}^{\text{RL}}$ is reported at $\lambda_c = 1$.*

Frontier position	Rule	$\mathbb{E}[\bar{Z}_c]$	$\text{CVaR}_{0.05}^+(\bar{Z}_c)$	$\Delta_{0.05}^+(\bar{Z}_c)$	D_c (%)	$f_{c,\infty}^{\text{RL}}$ (%)
Payment-seeking	Capped	437.4	466.5	29.0	97.35	24.7
	Uncapped	554.5	658.1	103.6	92.96	31.6
Intermediate	Capped	439.4	464.3	24.9	97.62	26.9
	Uncapped	526.9	641.2	114.4	90.72	33.6
Conservative	Capped	432.4	459.4	26.9	97.38	28.5
	Uncapped	475.6	604.0	128.4	87.91	34.4

Table G.3 shows that removing the cap raises both the limiting expected continuation cost and the residual excess tail under every policy. Across the three policies, the residual excess tail is 24.9–29.0 under capped continuation, compared with 103.6–128.4 under uncapped continuation. The lower uncapped diversification ratios indicate that a larger share of the corresponding representative-contract excess tail survives in the large-book limit. Thus, the cap reduces not only the expected continuation-payment level but also the market-sensitive tail that remains after household-specific mortality risk has diversified.




FIGURE G.2: *Conditional large-book risk-loaded spouse-continuation loads under capped and uncapped continuation. The upper-tail probability is $\alpha_c = 0.05$.*

Figure G.2 shows how the capped and uncapped large-book loads depend on λ_c . At $\lambda_c = 1$, Table G.3 reports capped loads of 24.7%–28.5% and uncapped loads of 31.6%–34.4%. The capped loads remain close to the corresponding expected-cost loads in Table G.1, whereas uncapped continuation carries a larger residual tail buffer. This difference cannot be captured by scaling the representative-contract excess tail mechanically by $J^{-1/2}$, because the rule-dependent common-market component does not diversify away.

H Numerical robustness

The learned policies are held fixed and evaluated on $K = 256,000$ common market-return paths, organized into $B = 100$ non-overlapping batches of 2,560 paths. The same market paths are retained across the three selected policies, book sizes, and paired capped–uncapped comparisons. This analysis therefore measures Monte Carlo evaluation uncertainty conditional on the learned policies; it does not include variability from retraining the neural networks.

For each market path, 1,000 independent household-mortality histories are generated under the fixed 2021 Australian age-specific period schedules. The first 1, 10, 100, and 1,000 histories define the nested books with $J = 1, 10, 100$, and 1,000, respectively. Contract outcomes are averaged within each book before any book-level statistic is computed. In the conditional large-book calculation, household mortality is instead integrated out using Algorithm C.2 and (6.14). Setting each spouse-only payment to one gives 8.5058442 expected payment dates through $M - 1$, the same quantity $\sum_{m=1}^{M-1} \pi_{01}(m)$ identified in Subsection 7.3. Debt positions accrue at the domestic bond-index return plus a borrowing spread of 2% per annum.

The principal mean and CVaR estimates are calculated from the complete K -vector. In particular, with $\alpha_c = 0.05$, upper-tail CVaR is the arithmetic mean of the largest 12,800 observations. Batches are used only to estimate Monte Carlo uncertainty. If $\hat{S}_b$ denotes the statistic calculated within batch b , the batch-based standard-error estimate is

$$\widehat{\text{se}}(\hat{S}) = \left[\frac{1}{B(B-1)} \sum_{b=1}^B \left(\hat{S}_b - \frac{1}{B} \sum_{r=1}^B \hat{S}_r \right)^2 \right]^{1/2}.$$

For ratios and capped–uncapped comparisons, the relevant statistic or difference is formed within each paired common-path batch before applying this formula. The full-sample point estimate is not replaced by the average of the batch estimates.

TABLE H.1: *Representative-contract evaluation estimates and Monte Carlo standard errors for the selected learned policies. Entries are reported as estimate (standard error). Probability and load standard errors are in percentage points. The continuation quantities use capped continuation, $\alpha_c = 0.05$, and $\lambda_c = 1$.*

Frontier position	λ	$\mathcal{R}$	$\text{CVaR}_{0.95}(L_T)$	$p_{\text{ins}}^{\text{ever}}$ (%)	$\mathbb{E}[Z_c]$	f_c (%)	$f_{c,1}^{\text{RL}}$ (%) $\lambda_c = 1$
Payment-seeking	0.3	1860.4 (0.7)	1218.0 (5.1)	26.71 (0.08)	438.1 (0.9)	23.55 (0.05)	51.90 (0.04)
Intermediate	1	1704.8 (0.7)	840.5 (9.9)	12.59 (0.07)	440.1 (0.8)	25.82 (0.05)	53.99 (0.04)
Conservative	15	1584.8 (0.6)	793.3 (9.5)	11.95 (0.07)	433.1 (0.8)	27.33 (0.05)	55.91 (0.04)

For the capped finite-book results, the Monte Carlo standard error of $\Delta_{0.05}^+(\bar{Z}_{c,J})$ is at most 0.75 for $J = 10$, 0.22 for $J = 100$, and 0.11 for $J = 1000$; the corresponding standard errors of $f_{c,J}^{\text{RL}}$ are at most 0.03, 0.01, and 0.01 percentage points. The standard errors of the large-book diversification ratios are below 0.01 percentage points. All estimated expected continuation costs for $J = 10, 100, 1000$ and the conditional large-book limit agree with the representative-contract mean within paired 95% Monte Carlo intervals. The estimated tail costs also satisfy, for every selected policy and continuation rule,

$$\text{CVaR}_{0.05}^+(\bar{Z}_c) \leq \text{CVaR}_{0.05}^+(\bar{Z}_{c,J}) \leq \text{CVaR}_{0.05}^+(Z_c).$$

The numerical conclusions in Section 7 are therefore not driven by evaluation-sample variation at the reported simulation size.

Acknowledgments

This work was supported by computational resources provided by The University of Queensland via the Friday supercomputer.

References

- [1] Allianz Australia Life Insurance Limited. Allianz Guaranteed Income for Life (AGILE): Product Disclosure Statement. https://www.allianzretireplus.com.au/content/dam/onemarketing/azau/allianzretireplus_com_au/pds-2026/AGILE_PDS_19%20Jan_2026_F1.pdf, January 2026. Issued 19 January 2026.
- [2] Association of Superannuation Funds of Australia. ASFA Retirement Standard: Super balances needed for comfortable retirement reach all-time high. <https://www.superannuation.asn.au/media-release/asfa-retirement-standard-super-balances-needed-for-comfortable-retirement-reach-all-time-high/>, February 2026. Published 24 February 2026.
- [3] Thomas Bernhardt and Catherine Donnelly. Modern tontine with bequest: innovation in pooled annuity products. *Insurance: Mathematics and Economics*, 86:168–188, 2019.
- [4] Marie-Fleur Borac and Peter Hieber. Couple tontines: Mutual joint-life and last-survivor insurance products. Available at <https://ssrn.com/abstract=6760536>, April 2026. SSRN preprint, first posted April 2026.
- [5] Hans Buehler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- [6] Chang Chen and Duy-Minh Dang. Convergence of neural network policies for risk–reward optimization. *arXiv preprint arXiv:2603.06563*, 2026.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] Catherine Donnelly. Actuarial fairness and solidarity in pooled annuity funds. *ASTIN Bulletin: The Journal of the IAA*, 45(1):49–74, 2015.
- [12] Catherine Donnelly, Montserrat Guillén, and Jens Perch Nielsen. Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, 56:14–27, 2014.
- [13] Catherine Donnelly and Xin Zhang. Joint life income in a pooled annuity fund. Technical report, Institute and Faculty of Actuaries, January 2023. Commissioned report.
- [14] Catherine Donnelly and Xin Zhang. Longevity risk sharing for multi-life benefits in pooled annuity funds: State-transition and Möbius decomposition frameworks. SSRN preprint 6979498, June 2026. Version dated 22 June 2026.
- [15] 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.
- [16] 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.
- [17] Richard K. Fullmer. Tontines: A practitioner’s guide to mortality-pooled investments, 2019. Accessed: January 2026.
- [18] Jiequn Han and Weinan E. Deep learning approximation for stochastic control problems. *arXiv preprint arXiv:1611.07422*, 2016.
- [19] 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.
- [20] Andreas Hübener, Raimond Maurer, and Ralph Rogalla. Optimal portfolio choice with annuities and life insurance for retired couples. *Review of Finance*, 18(1):147–188, 2014.
- [21] Human Mortality Database. Human mortality database. Max Planck Institute for Demographic Research (Germany), University of California, Berkeley (USA), and French Institute for Demographic Studies (France). Available at <https://www.mortality.org/>, 2026. Data downloaded on 29 July 2026. Database version 20260615, DOI: 10.4054/HMD.Countries.20260615.
- [22] 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.
- [23] Legal & General. What is a joint life annuity? <https://www.legalandgeneral.com/retirement/pension-annuity/guides/joint-life-annuity/>, June 2025. Accessed 8 August 2026.

- [24] 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.
- [25] Thomas Mathews. A history of australian equities. Research Discussion Paper RDP 2019-04, Reserve Bank of Australia, 2019.
- [26] Moshe A Milevsky and Thomas S Salisbury. Optimal retirement income tontines. *Insurance: Mathematics and economics*, 64:91–105, 2015.
- [27] N. M. Superannuation Proprietary Limited. MyNorth Super and Pension: Product Disclosure Statement—Part A. <https://www.northonline.com.au/content/dam/north-online/documents/pds/insights/my-north-super-pension-pds.pdf>, September 2025. Issue no. 14, issued 30 September 2025; read with the supplementary Product Disclosure Statement dated 1 July 2026.
- [28] German Nova Orozco, Duy-Minh Dang, and Peter A. Forsyth. Money-back tontines for retirement decumulation: Neural-network optimization under systematic longevity risk, 2026. <https://arxiv.org/abs/2602.16212>.
- [29] 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.
- [30] Wade D Pfau. An overview of retirement income planning. *Journal of Financial Counseling and Planning*, 29(1):114–120, 2018.
- [31] Dimitris N Politis and Joseph P Romano. The stationary bootstrap. *Journal of the American Statistical association*, 89(428):1303–1313, 1994.
- [32] Dimitris N Politis and Halbert White. Automatic block-length selection for the dependent bootstrap. *Econometric reviews*, 23(1):53–70, 2004.
- [33] QSuper. QSuper Product Disclosure Statement for Income Account and Lifetime Pension. <https://qsuper.qld.gov.au/pds>, July 2025. Issue date: 1 July 2025.
- [34] RBC Insurance. Payout annuities. <https://www.rbcinsurance.com/en-ca/annuities/>. Accessed 8 August 2026.
- [35] A Max Reppen, H Mete Soner, and Valentin Tissot-Daguette. Deep stochastic optimization in finance. *Digital Finance*, 5(1):91–111, 2023.
- [36] Anders Max Reppen and Halil Mete Soner. Deep empirical risk minimization in finance: Looking into the future. *Mathematical Finance*, 33(1):116–145, 2023.
- [37] Michael J Sabin and Jonathan Barry Forman. The analytics of a single-period tontine. *Available at SSRN 2874160*, 2016.
- [38] Lisanne Sanders and Bertrand Melenberg. Estimating the joint survival probabilities of married individuals. *Insurance: Mathematics and Economics*, 67:88–106, 2016.
- [39] 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.
- [40] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. *arXiv preprint arXiv:1808.01974*, 2018.
- [41] TIAA. How does the tiaa traditional annuity work? <https://www.tiaa.org/public/learn/retirement-planning-and-beyond/how-do-traditional-annuities-work>. Accessed 8 August 2026.
- [42] 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.
- [43] 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.
- [44] Manuel Ventura-Marco, Carlos Vidal-Meliá, and Juan Manuel Pérez-Salamero González. Joint life care annuities to help retired couples to finance the cost of long-term care. *Insurance: Mathematics and Economics*, 113:122–139, 2023.