# Optimal Management of a Retirement Account with a Tontine Overlay: A Neural Network Approach

M. M. Shirazi* P. A. Forsyth† K. R. Vetzal‡

March 30, 2026

**Note:** The views expressed in this work are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

**Abstract**

Individual tontine retirement accounts can reduce longevity risk by transferring assets from deceased to surviving tontine pool members. Forsyth et al. (2024) consider an example where retirees participate in a tontine that lets them choose annual withdrawals within a specified range and also select their preferred asset allocations. Retirees are assumed to want to maximize total withdrawals and to minimize a measure of the left tail risk of their portfolio values at the end of a 30-year horizon. Forsyth, Vetzal, and Westmacott (2024) determine the optimal withdrawals each year and allocations to equity and fixed income investments using dynamic programming and partial differential equation methods. The tontine performs very well, but the approach used works for only a small number of risky factors, e.g. two investment assets. Here we consider the same context, but we use a neural network approach to determine the optimal withdrawals and asset allocations. This lets us explore a more realistic investment environment with additional assets and alternative assumptions about mortality risk. Examples show that the tontine is fairly robust to random changes in mortality. Having more investment assets gives even better performance, allowing high annual withdrawals with low risk of portfolio depletion.

**Keywords:** Tontine; decumulation; neural network approach; optimal stochastic control.
**JEL codes:** G11, G22

## 1 Introduction

Access to traditional defined benefit (DB) pension plans continues to trend downwards for employees. In many cases, employees have been transitioned to some form of defined contribution (DC) plan. By 2024, just 15% of private sector workers in the U.S. had access to a DB plan, while 70% had a DC plan available. The corresponding figures for 2010 were 20% and 59% for DB and DC plans respectively.[1] Broadly similar trends have been seen in many other countries as well. DB plans are often viewed by sponsoring organizations as both expensive and risky, due to factors such as lengthy periods of low interest rates requiring riskier investment strategies and perhaps increased contributions to fund the specified benefits as well as mortality improvements increasing liabilities due to to benefits being paid out over longer average retirement periods.

A DC plan does not magically erase these risk exposures, but transfers them from sponsors to plan participants. Employees need to invest contributions appropriately to accumulate savings for retirement,

*Current address: Bank of Canada, Ottawa ON Canada K1A 0G9, Mshirazi@bank-banque-canada.ca. While contributing to this research project, Shirazi was a graduate student at the Cheriton School of Computer Science, University of Waterloo, Waterloo ON Canada N2L 3G1.

†Cheriton School of Computer Science, University of Waterloo, Waterloo ON Canada N2L 3G1, paforsyt@uwaterloo.ca.

‡School of Accounting and Finance, University of Waterloo, Waterloo ON Canada N2L 3G1, kvetzal@uwaterloo.ca.

[1]Source: (U.S. Bureau of Labor Statistics, 2024). If the sample is expanded to civilian workers (i.e. private sector plus state and local government workers), the corresponding 2024 numbers are 24% for DB plans (down from 31% in 2010) and 66% for DC plans (up from 54% in 2010).

while in addition to managing their investment portfolios retirees must determine appropriate withdrawals to pay for living expenses rather than simply receiving their defined pension benefits (assuming DB plan solvency). While some participants may wish to tackle these challenges themselves, others may benefit from professional financial planning advice. However, most traditional DC plans expose participants to higher longevity risk because there is no pooling of mortality risk. In DB plans, the benefits paid to participants who live longer are at least partially funded by the contributions of those who die relatively young. This *risk pooling* is an essential characteristic of insurance, and it does not happen in most traditional DC plans. The accumulated plan assets of participants who pass away go to their estate beneficiaries, rather than being available to benefit longer-lived plan members. This means that DC plan members are much more exposed to the risk of outliving their savings.

A modern tontine is a structure which potentially offers a partial remedy for the lack of longevity risk pooling in traditional DC plans. While many designs are possible, in general participants in a tontine combine their assets into a shared pool which re-distributes the pool assets of deceased members to surviving members as "mortality credits". This enhances the investment accounts of survivors, supporting them into very advanced ages. This possibility has resulted in numerous recent research studies into various types of pooled mortality structures (see, e.g. Milevsky and Salisbury, 2015, 2016; Bräughtigam et al., 2017; Milevsky et al., 2018; Fullmer, 2019; Bernhardt and Donnelly, 2019; Winter and Planchet, 2022; Gemmo et al., 2020; Weinert and Gründl, 2021; Chen et al., 2021; Milevsky, 2022; Chen and Rach, 2023; Vetzal et al., 2023; Forsyth et al., 2024). Mortality pooling for DC type plans historically has been quite rare. Two examples are TIAA's College Retirement Equities Fund, which dates to 1952, and the University of British Columbia faculty's Variable Life Pension Annuity, which dates to 1967. However, recently some additional products have debuted, including Australia's QSuper Lifetime Pension (in 2021) and two Canadian structures: the Purpose Longevity Pension Fund (2021), and the GuardPath Modern Tontine mutual fund (2022).[2]

It was noted above that tontine structures may offer partial protection against longevity risk. The reason why this insurance is incomplete lies in an important distinction between two types of longevity risk:

- *Idiosyncratic longevity risk* is the risk associated with each individual's lifespan, assuming that the overall mortality of the tontine pool members is consistent with expectations.

- *Systematic longevity risk* is the risk connected to the entire pool's average lifespan. For example, unanticipated advancements in healthcare may mean that the entire pool lives on average longer than expected.

A tontine pool can offer diversification benefits that mitigate idiosyncratic longevity risk. Mortality credits earned by pool survivors due to asset transfers from deceased pool members support higher levels of annual withdrawals for survivors during retirement. However, if the average lifespan of pool participants increases, mortality credits will be delayed, and pool members who live relatively long periods will receive lower benefits.

An obvious alternative to a tontine structure is annuitization, which offers complete longevity risk protection assuming solvency of the annuity provider. However, annuities have proven to be quite unpopular. There are numerous possible reasons for this (see MacDonald et al., 2013, for an extensive discussion), but it is worth emphasizing here that annuities are typically quite expensive, particularly in a relatively low interest rate environment. In addition, virtually all annuities provide long-term fixed nominal periodic payments, so annuity purchasers who experience lengthy retirements are exposed to considerable inflation risk.

By providing partial protection (i.e. protection against idiosyncratic but not systematic longevity risk), modern tontines can offer a form of cheaper insurance compared to annuities. This is a principal motivation for studying tontines. Forsyth et al. (2024) consider a perpetual tontine pool with annual real (i.e. inflation-adjusted) withdrawals that are subject to upper and lower bounds. Forsyth et al. compare this structure with two simpler alternatives which lack tontine mortality credits over a 30-year retirement period. The first alternative is based on the well-known 4% rule (Bengen, 1994), which assumes that the retiree constantly withdraws an inflation-adjusted 4% of the initial lump sum savings available at the inception of retirement and follows a strategy with constant portfolio weights between equities and fixed income investments. The second alternative allows annual withdrawals within a specified range and flexible portfolio weights over time. Optimal strategies for the withdrawals and portfolio weights are determined using a numerical optimal

[2]Note, however, that the GuardPath product was terminated in January 2025. For additional details about the other recent offerings, see qsuper.qld.gov.auand www.retirewithlongevity.com.

control approach (Forsyth, 2022). Forsyth et al. (2024) show that allowing flexible withdrawals and portfolio weights gives a marked improvement over a scenario with fixed portfolio weights and fixed withdrawals, but this flexible approach is in turn dramatically improved by incorporating a tontine overlay.

However, the approach used by Forsyth et al. (2024) to determine withdrawal and asset allocation strategies is based on numerical solution techniques which have strengths but also some significant weaknesses. The main strength is that the numerical approach used by Forsyth et al. provably converges to the correct optimal solution. However, this approach requires specifying and estimating a parametric model of asset returns and computational complexity considerations limit its applicability to a small number of risk factors.

In this article we consider a neural network (NN) approach for determining optimal asset allocation and withdrawal strategies in the same tontine setup as Forsyth et al. (2024). Our NN approach is able to handle a higher number of risk factors. As an illustration, we explore allocation over six different investment assets (four equity indices and two fixed income indices), whereas Forsyth et al. are limited to considering just two investment assets (one equity index and one fixed income index). In addition, the NN approach does not have to rely on a specified parametric model of asset returns: it can use such a model, but it is also possible to use bootstrap resampling of historical data instead. Put differently, the NN approach just requires having a set of asset return paths, but it doesn't matter whether these paths are produced by Monte Carlo simulation of a parametric model or by a resampling procedure. Overall, our NN approach offers the ability to handle more risk factors and avoids necessarily having to specify and estimate a parametric model of asset returns. While it does not provably converge to the correct optimal strategies, we show that it comes quite close to the strategies generated by the numerical approach of Forsyth et al. when restricted to just two investment assets and in the context of a given parametric model for asset returns.

The remainder of this article proceeds as follows. Section 2 describes the economic environment. Section 3 presents the NN formulation. Section 4 provides some initial validation tests, and Sections 5 and 6 explore some alternative assumptions and broader investment opportunities respectively. Section 7 summarizes and concludes.

# 2 Problem Description

This section describes the investment scenario in detail. Note that some parts of this material are quite similar to that provided in Forsyth et al. (2024), which is to be expected since the purpose of this paper is to provide an alternative and more general NN-based approach for solving the same type of investment problem.

## 2.1 Overview

Following Forsyth et al. (2024), we consider an investor who has just retired. To be conservative about longevity, we assume that this investor plans to withdraw funds from a retirement portfolio over a horizon of $T = 30$ years. Withdrawals and portfolio rebalancing both occur at an annual frequency. On rebalancing dates, the investor first withdraws funds and then reallocates the assets of the portfolio. The amount of the withdrawal is flexible, within a pre-specified range. The investment portfolio can be a mixture of broad stock index funds and constant maturity bond index funds. The investor cannot use leverage or short sales. The investor also has the choice of participating in a tontine pool consisting of individual tontine accounts (Fullmer and Sabin, 2019).

Due to the requirement of withdrawing at least a minimum amount each year, it is possible for the investor's portfolio value to become negative. If this happens, the investor can no longer invest in equities and subsequent withdrawals are restricted to the minimum amount. The investor is presumed to finance this by taking on debt, which will grow at the borrowing rate. It is assumed that the investor has access to other assets outside of the investment portfolio such as mortgage-free real estate, which can be used to fund accumulated debt.

The investor faces a tradeoff between risk and reward. In particular, the investor would like to maximize total withdrawals throughout the horizon (the reward), but higher withdrawals run the risk of depleting the portfolio at some point during the horizon. This risk is measured by *expected shortfall* (ES), which is the mean of the worst 5% of outcomes of the portfolio value at the end of the retirement period. While ES is

measured at this terminal point, lower (more negative) values of ES indicate that the investor ran out of money and started accumulating debt at an earlier stage. The investor is assumed to try to maximize a weighted sum of total withdrawals and ES, where the weights depend on risk-aversion. Strategies that run out of money earlier are penalized by both lower ES and reduced withdrawals.

## 2.2 Tontine Pool

We allow the investor to participate in a perpetual pool of individual tontine accounts. This pool is not restricted to new members and is managed by a financial organization in return for a fee. The assets of pool participants who die are redistributed to surviving pool members.

This redistribution is assumed to be actuarially fair. Following Fullmer (2019), the expected value of participating in a tontine is then the same as the expected value of staying out of the pool. In other words, participation is a fair game with expected losses balanced by expected gains. Consider a member $j$ of the pool at time $t_{i-1}$ with a portfolio worth $v^j_{i-1}$. If this investor dies during the period $(t_{i-1},t_i)$, his investment value $v^j_i$ is redistributed at the end of the period. If this investor survives, he receives a mortality credit (or *tontine gain*) $c^j_i$ as a redistribution from the accounts of other pool members who pass away during the period. Let the probability of death for investor $j$ during $(t_{i-1},t_i)$ be $\mathfrak{q}^j_{i-1}$. The expected loss for investor $j$ due to dying in the period $(t_{i-1},t_i)$ is then $-\mathfrak{q}^j_{i-1}v^j_i$, and his expected gain is his survival probability $(1-\mathfrak{q}^j_{i-1})$ multiplied by the expected mortality credit, denoted by $c^j_i$. Actuarial fairness requires then requires that

$$-\mathfrak{q}^j_{i-1}v^j_i + (1-\mathfrak{q}^j_{i-1})c^j_i = 0, \tag{1}$$

which can be rearranged to

$$c^j_i = \frac{\mathfrak{q}^j_{i-1}}{(1-\mathfrak{q}^j_{i-1})}v^j_i. \tag{2}$$

The tontine gain rate for investor $j$ is defined to be

$$\left(\mathbb{T}^g_i\right)^j = \frac{\mathfrak{q}^j_{i-1}}{(1-\mathfrak{q}^j_{i-1})}.$$

Since our retirement problem is considered from the standpoint of a single retiree, we can drop the superscript $j$ and define

$$\mathbb{T}^g_i = \mathfrak{q}_{i-1}/(1-\mathfrak{q}_{i-1}). \tag{3}$$

Note that the tontine gain is independent of the allocation strategies of the other participants in the tontine pool. This counterintuitive result can be justified on the basis of the "small bias condition" (Sabin and Forman, 2016). Essentially, this condition requires that the expected total amount forfeited during any time period is large compared to any single member's expected gain.

The probability of death can be estimated from actuarial mortality tables, but the actual number of deaths is random and can differ from the expected value. Since mortality credits will depend on actual deaths, we need to adjust for this. Following Fullmer (2019) and Forsyth et al. (2024), suppose there are $m$ investors in the pool at $t_{i-1}$ and define the indicator function

$$\mathbf{1}^j_i = \begin{cases} 1 & \text{if investor } j \text{ is alive at } t_{i-1} \text{ and at } t_i \\ 0 & \text{if investor } j \text{ is alive at } t_{i-1} \text{ and dead at } t_i \end{cases}. \tag{4}$$

Total realized gains and forfeitures in the pool during the period $(t_{i-1},t_i)$ are then

$$\text{Total gains} = \sum_{j=1}^{m} \mathbf{1}^j_i c^j_i = \sum_{j=1}^{m} \mathbf{1}\left(\frac{\mathfrak{q}^j_{i-1}}{1-\mathfrak{q}^j_{i-1}}\right) v^j_i, \tag{5}$$

$$\text{Total forfeitures} = \sum_{j=1}^{m} (1-\mathbf{1}^j_i) v^j_i. \tag{6}$$

To adjust for the difference between actual and expected deaths, we follow Fullmer and Sabin (2019) by introducing the *group gain*

$$G_i = \frac{\sum_{j=1}^{m}(1-\mathbf{1}_i^j)v_i^j}{\sum_{j=1}^{m}\mathbf{1}_i^j c_i^j} \tag{7}$$

where $\mathbf{1}_i^j$ is the actual state of member $j$. We modify the mortality credits earned by investor $j$ to

$$\hat{c}_i^j = G_i\left(\frac{\mathrm{q}_{i-1}^j}{1-\mathrm{q}_{i-1}^j}\right)v_i^j, \tag{8}$$

so as to ensure that total mortality credits paid out at $t_i$ equals total forfeitures during the preceding period. This transforms the tontine gain rate equation (3) into a stochastic process where

$$\mathbb{T}_i^g = G_i\left(\frac{\mathrm{q}_{i-1}}{1-\mathrm{q}_{i-1}}\right) \tag{9}$$

and $G_i$ is a random number.

To get a sense of representative values of $G_i$, Fullmer and Sabin (2019) run Monte Carlo simulations of a tontine pool which takes in 1,000 new members each year. Pool members have randomly assigned ages, random initial investment contributions, and random years of death generated in accordance with an actuarial mortality table. They are also randomly allocated to a fixed proportion investment strategy (all stocks, all bonds, or a 50/50 mixture of stocks and bonds). Based on 10,000 simulations, Fullmer and Sabin show that the mean group gain is very close to 1 with a standard deviation that declines as the pool size increases from about 0.4 during the early years of the simulation to around 0.1 after about 15 years.

## 2.3 Notational Conventions and Description of Controls

Let the number of assets available to the investor be $N_a$ and the real (inflation-adjusted) amount invested in asset $k$ at time $t$ be $A_k(t)$, so that the total value of the investor's portfolio at time $t$ is

$$W_t = \sum_{k=1}^{N_a} A_k(t). \tag{10}$$

For simplicity, we will henceforth refer to the total portfolio value as the investor's wealth, with the understanding that this does not include the value of any assets that the investor owns such as real estate that could be used to fund debt in the event of portfolio depletion.

Let $\mathcal{T}$ be a set of evenly-spaced discrete times at which rebalancing and withdrawals are permitted,

$$\mathcal{T} = \{t_0 = 0 < t_1 < t_2 < \cdots < t_M = T\}. \tag{11}$$

The start of the investment period is $t_0 = 0$, the end of it is $t_M = T$, and the length of the period between these discrete times is $t_i - t_i - 1 = \Delta t = T/M$. At each time $t_i \in \mathcal{T}$, the following order of events is assumed: (i) if the investor is participating in a tontine pool, mortality credits are received from the pool, as reflected in equation (9); (ii) the investor withdraws an amount of cash $q_i$ from the portfolio; and (iii) the investor rebalances the portfolio. The portfolio is liquidated at the end of the horizon $T$. Consistent with the fact that DC plan assets are usually owned in tax-advantaged accounts which do not result in taxes being owed upon rebalancing, we ignore such taxes. We also ignore transaction costs, as these can be expected to be minimal since we assume rebalancing dates are relatively infrequent (e.g. annually). However, we assume that a constant fee $\mathbb{T}^f$ is charged to participate in the tontine.

For clarity involving the sequence of events at $t_i \in \mathcal{T}$, for any function $f(t)$ we denote $f(t_i^+) = \lim_{\epsilon\to 0^+} f(t_i+\epsilon)$ and $f(t_i^-) = \lim_{\epsilon\to 0^+} f(t_i-\epsilon)$. We also denote the controlled underlying process by $X(t) = \{A_k(t) : k = 1,2,\ldots,N_a\}$, $t \in [0,T]$. This is determined by the *allocation control* $p$ (more details below).

Portfolio withdrawals are assumed to depend on the value of the portfolio after the receipt of tontine gains $\mathbb{T}^g$ and the payment of the tontine fee $\mathbb{T}^f$, but prior to rebalancing, as well as time. We denote this value as

$$W(t_i^-) = \left(\sum_{k=1}^{N_a} A_k(t_i^-)\right)(1+\mathbb{T}_i^g)\exp(-(\Delta t)\mathbb{T}^f). \tag{12}$$

The amount withdrawn is determined by the *withdrawal control* $q$. At time $t_i \in \mathcal{T}$, we have $q_i(\cdot) = q(X(t_i^-),t_i)$, i.e. the amount withdrawn depends on the portfolio value immediately prior to the withdrawal event. This control determines the evolution of the portfolio value at $t_i$:

$$W(t_i^+) = W(t_i^-) - q_i, \quad t_i \in \mathcal{T}. \tag{13}$$

The withdrawal and allocation controls are both formally functions of the state before withdrawal $X(t_i^-)$. However, as described above rebalances occur after withdrawals, we can consider the allocation control as depending on $W(t_i^+)$. More specifically, let $p_k(\cdot)$ denote the fraction of wealth invested in the $k$-th asset after rebalancing, so that

$$A_k(t_i^+) = p_k(X(t_i^+),t_i^+)W(t_i^+). \tag{14}$$

The full underlying allocation control is then

$$\boldsymbol{p}(X(t_i^+),t_i^+) = \left\{p_k(X(t_i^+),t_i^+) \ : \ k = 1,2,\ldots,N_a\right\}. \tag{15}$$

## 2.4 Constraints

Both the allocation control and the withdrawal control are subject to constaints. The constraints on the allocation control are a prohibition of short positions and no use of leverage, except in the case of insolvency. The withdrawal control must specify that the amount withdrawn lie within a lower bound $q_{\min}$ and an upper bound $q_{\max}$. Due to minimum withdrawals at each rebalancing time, insolvency can occur. In this case, (i) no stock holdings are allowed; (ii) debt will accumulate at a borrowing rate; and (iii) subsequent withdrawals are restricted to the minimum amounts. In addition, all equity positions are assumed to be liquidated into risk-free bonds at the end of the investment horizon $T$. For ease of representation, let $W_i^- = W(t_i^-)$ be the pre-withdrawal wealth at $t_i$ and let $W_i^+ = W(t_i^+)$ be the post-withdrawal wealth at $t_i$. In addition, assume that $p_1$ represents the proportion of wealth invested in risk-free bonds, which serve as the benchmark for the borrowing rate in the event of insolvency. The overall control at time $t_i$ is described by a pair $(q_i(\cdot),\boldsymbol{p}(\cdot))$ from the admissible set $\mathcal{Z}(W_i^-,W_i^+,t_i)$. Mathematically, we can then represent the constraints as

$$\mathcal{Z}_q(W_i^-,t_i) = \begin{cases} [q_{\min},q_{\max}] & t_i \in \mathcal{T} \ ; \ W_i^- > q_{\max} \\ [q_{\min},W_i^-] & t_i \in \mathcal{T} \ ; \ q_{\min} < W_i^- < q_{\max} \\ q_{\min} & t_i \in \mathcal{T} \ ; \ W_i^- < q_{\min} \end{cases}, \tag{16}$$

$$\mathcal{Z}_p(W_i^+,t_i) = \left\{ \boldsymbol{p} \in \mathbb{R}^{N_a} \;\middle|\; \begin{cases} \sum_{k=1}^{N_a} p_k = 1 & t_i \neq t_M \\ p_k \in [0,1] \ k = \{1,\ldots,N_a\} & W_i^+ > 0 \ ; \ t_i \in \mathcal{T} \ ; \ t_i \neq t_M \\ \boldsymbol{p} = \{1,0,\ldots,0\} & W_i^+ \leq 0 \ ; \ t_i \in \mathcal{T} \ ; \ t_i \neq t_M \\ \boldsymbol{p} = \{1,0,\ldots,0\} & t_i = t_M \end{cases} \right. , \tag{17}$$

$$\mathcal{Z}(W_i^-,W_i^+,t_i) = \mathcal{Z}_q(W_i^-,t_i) \times \mathcal{Z}_p(W_i^+,t_i). \tag{18}$$

Denote the admissible control set

$$\mathcal{A} = \left\{(q_i(\cdot,t_i),\boldsymbol{p}(\cdot,t_i))_{0\leq i\leq M} : (q_i(\cdot,t_i),\boldsymbol{p}(\cdot,t_i)) \in \mathcal{Z}(W_i^-,W_i^+,t_i)\right\}, \tag{19}$$

so that an admissible control $\mathcal{P} \in \mathcal{A}$ may be expressed as

$$\mathcal{P} = \{(q_i(\cdot,t_i),\boldsymbol{p}(\cdot,t_i)) \ : \ i = 0,\ldots,M\}. \tag{20}$$

Finally, define $\mathcal{P}_n$ as $\mathcal{P}_{t_n} \subset \mathcal{P}$ as the remainder of the set of controls from $t_n$ on, i.e.

$$\mathcal{P}_n = \{(q_n(\cdot,t_n),\boldsymbol{p}(\cdot,t_n)),(q_{n+1}(\cdot,t_{n+1}),\boldsymbol{p}(\cdot,t_{n+1})),\ldots,(q_M(\cdot,t_M),\boldsymbol{p}(\cdot,t_M))\}. \tag{21}$$

## 2.5 Objective Function

Following Forsyth et al. (2024), we assume that the investor seeks the best tradeoff between reward, defined as expected total withdrawals, and risk, measured in terms of expected shortfall. More specifically, reward is expected withdrawals

$$\mathrm{EW}(X_0^-, t_0^-) = E_{\mathcal{P}_0}^{X_0^-, t_0^-}\left[\sum_{i=0}^{M} q_i\right], \tag{22}$$

and expected shortfall is

$$\mathrm{ES}_\alpha = \frac{\int_{-\infty}^{W'_\alpha} W_T \mathcal{G}(W_T) dW_T}{\alpha}, \tag{23}$$

where $\mathcal{G}(W_T)$ is the probability density of terminal wealth $W_T$ and $W'_\alpha$ is the "Value-at-Risk" of level $\alpha$, i.e.

$$\int_{-\infty}^{W'_\alpha} \mathcal{G}(W_T) dW_T = \alpha. \tag{24}$$

We use $\alpha = .05$, which implies that ES is the mean of the 5% left tail of the distribution of $W_T$. Following Rockafellar and Uryasev (2000), ES can be expressed in the form

$$\mathrm{ES}_\alpha = \sup_{W'} E\left[W' + \frac{\min(W_T - W', 0)}{\alpha}\right]. \tag{25}$$

Under a control $\mathcal{P}$ and given the initial state $X_0$, this becomes

$$\mathrm{ES}_\alpha(X_0^-, t_0^-) = \sup_{W'} E_{\mathcal{P}_0}^{X_0^-, t_0^-}\left[W' + \frac{\min(W_T - W', 0)}{\alpha}\right], \tag{26}$$

where the candidate values of $W'$ are the set of possible values for $W_T$. Note that this expectation is set at $t_0$, so $W'$ is fixed throughout the investment horizon. For readers who are concerned about time-consistency, note that we are effectively considering an induced time-consistent strategy, not a time-inconsistent version of an expected shortfall strategy.[3]

The investor's overall objective is to find a control strategy $\mathcal{P}_0$ subject to the constraints (16) and (17) that maximizes a weighted sum of EW and ES, plus a stabilization term $\epsilon W_T$. Denoting the overall objective function by $F$, we have

$$F = \sup_{W'} E_{\mathcal{P}_0}^{x_0^-, t_0^-}\left[\overbrace{\sum_{i=0}^{M} q_i}^{\text{EW}} + \overbrace{\left\{\kappa\left(W' + \frac{\min(W_T - W', 0)}{\alpha}\right)\right\}}^{\text{ES}} + \overbrace{\epsilon W_T}^{\text{stabilization}}\right] \tag{27}$$

In equation (27), $\kappa$ is effectively a risk-aversion parameter (larger values of $\kappa$ place more weight on the expected shortfall of terminal wealth), and the stabilization term mitigates ill-posedness in the problem when $W_t \gg W'$ as $t \to T$. The issue here is that if wealth is relatively high near the end of the horizon, withdrawals will be the maximum permitted and the allocation control will have almost no effect since the risk measure only depends on the extreme left tail of the distribution. Setting $\epsilon$ to a small positive value will effectively push the asset allocation into the riskiest asset, while setting $\epsilon$ to a small negative value will result in the portfolio being concentrated in the least risky investment. In addition to seeking the optimal control strategy $\mathcal{P}_0$, formulation (27) makes clear that we need to also optimize with respect to $W'$ for specified values of $\kappa$ and $\alpha$.

[3] See Appendix A of Forsyth et al. (2024) for a detailed discussion.

# 3 Neural Network Formulation

Forsyth et al. (2024) construct a dynamic programming solution to the optimization problem that relies on (i) specifying a particular parametric model for asset returns between rebalancing dates; and (ii) the use of numerical methods for solving partial integro differential equations. This approach is provably convergent, but it suffers from the "curse of dimensionality", i.e. it is restricted to a small number of risk factors due to complexity considerations. As a result, Forsyth et al. (2024) consider just two investment assets, a broad-based equity index and a constant maturity short-term government bond index under the asssumption that the returns of each of these investments can be described by double exponential jump diffusion processes.

Following Li and Forsyth (2019) and Chen et al. (2023), an alternative approach to solving the optimization problem uses NN methods to approximate the solution. This type of approach offers a couple of potential advantages. First, while it requires having a set of asset returns between rebalancing dates, the source of these returns is immaterial. Any desired parametric model can be used, or returns can be directly sampled from observed historical data without specifying any model. Second, it is easier to incorporate additional risk factors, so we can consider a broader set of investment assets. On the other hand, this type of approach is relatively slow (compared to alternatives when they are feasible) and is not provably convergent. Although a sufficently large NN can approximate a continuous function to any desired accuracy, the implicit requirement is that the optimal solution of a non-convex optimization problem can be determined. This assumption is, of course, questionable.

The basic idea is to approximate the control $\mathcal{P}$ through two feed-forward fully-connected NNs. The first NN gives an approximation $\hat{q}_i$ for the withdrawal control $q_i$ given NN weights and biases $\theta_q$. Since withdrawals depend on wealth prior to rebalancing,

$$\hat{q}\left(W_i^-, t_i^-, \theta_q\right) \simeq q_i(W_i^-), \quad i = 0, \ldots, M. \tag{28}$$

The second NN generates an approximation $\hat{p}_i$ for the allocation control $p_i$ using NN weights and biases $\theta_p$, depending on post-withdrawal wealth:

$$\hat{p}\left(W_i^+, t_i^+, \theta_p\right) \simeq p_i(W_i^-), \quad i = 0, \ldots, M. \tag{29}$$

These two NNs give the approximation $\hat{\mathcal{P}} = \{(\hat{q}(\cdot), \hat{p}(\cdot))\} \simeq \mathcal{P}$.[4] In the taxonomy of Powell (2023), this approach is termed policy function approximation. This technique was described in Han and Weinan (2016), Buehler et al. (2019), and Li and Forsyth (2019). Observe that equations (28) and (29) are global-in-time approximations (Hu and Laurière, 2023). This technique is based on directly approximating the control, and does not use dynamic programming (van Staden et al., 2025). This contrasts with typical dynamic programming based reinforcement learning algorithms. Dynamic programming based methods require approximation of high dimensional conditional expectations (van Staden et al., 2023) in order to determine a low dimensional control function, which is undesirable.

As indicated above, the inputs to these NNs both contain time and prevailing wealth, but the wealth is measured prior to the withdrawal for $\hat{q}$ and after the withdrawal for $\hat{p}$. The NNs can have different numbers of hidden layers and different numbers of nodes per layer, and activation functions at the output layer. As in Li and Forsyth (2019), a standard softmax activation function is used for the output layer of the allocation NN. This automatically generates values between zero and one for each node of this output layer, which are then readily interpreted as portfolio weights consistent with the no-shorting and no-leverage constraints (17).

Following Chen et al. (2023), a modified sigmoid activation function is used for the withdrawal NN which gives output consistent with the withdrawal constraints (16). Let

$$Q = \begin{cases} q_{\max} - q_{\min} & \text{if } W_i^- > q_{\max} \\ W_i^- - q_{\min} & \text{if } q_{\min} < W_i^- < q_{\max} \\ 0 & \text{if } W_i^- < q_{\min} \end{cases}, \tag{30}$$

[4] In some cases we allow for the group gain $G$ to be random. We can formally add $G$ as an additional input to our NNs, but in some sense this is not necessary as $G$ can be viewed as just another random variable affecting wealth, with a more significant impact later on in the investment horizon.

which can be expressed more concisely as

$$Q = \max\left[\left((\min(q_{\max}, W_i^-) - q_{\min}\right), 0\right]. \tag{31}$$

The output layer of the withdrawal NN contains a single node, which is used to determine the approximate withdrawal amount $\hat{q}$ as follows. A sigmoid activation function is used to produce for this output layer, which produces a number $\beta \in [0,1]$. The withdrawal amount is then specified as

$$\hat{q} = q_{\min} + Q \times \beta. \tag{32}$$

To summarize, the input layer to the withdrawal NN contains time $t_i$ and prevailing wealth measured prior to the withdrawal, $W_i^-$. The output layer contains a single node, which is used to determine the withdrawal amount $\hat{q}_i$. Prevailing wealth post-withdrawal is $W_i^+ - \hat{q}_i$, which along with time $t_i$ is used for the input layer for the allocation NN. The output layer of this NN contains $N_a$ nodes which are used to determine the asset allocations $\hat{p}_i$.

Within this general framework, there is a lot of flexibility in terms of designing the withdrawal and allocation NNs in terms of the number of hidden layers, the number of nodes per layer, activation functions for the hidden layers, etc. A simple example is shown in Figure 1, which depicts a case where the withdrawal NN has two hidden layers with five nodes each and the allocation NN has three hidden layers with four nodes each. In addition, the figure implicitly assumes that there are $N_a = 3$ assets, so the output layer has three nodes, with the allocation to asset $k$ at time $t_i$ designated as $\hat{p}_{i_k}$, $k \in \{1,2,3\}$.

Next, we describe the NN training optimization procedure. Suppose that a set of $N$ joint asset return paths is available. These random paths are assumed to be independent, although the returns of different assets on a path can be correlated. We approximate the expectation in (27) using $N$ sample paths as follows:

$$\begin{aligned}
&\hat{V}(\theta_q, \theta_p, W', X(t_0^-), t_0^-) = \\
&\quad \frac{1}{N}\sum_{j=1}^{N}\left[\sum_{i=0}^{M}\hat{q}((W_i^-)^j, t_i; \theta_q) + \kappa\left(W' + \frac{\min\left((W_T)^j - W', 0\right)}{\alpha}\right) + \epsilon(W_T)^j | X(t_0^-) = (W_0^-, t_0^-)\right],
\end{aligned} \tag{33}$$

subject to the evolution of asset values between rebalancing dates determined by the given joint return paths, and

$$\begin{aligned}
(W_i^+)^j &= (W_i^-)^j - \hat{q}_i\left((W_i^-)^j, t_i, \theta_q\right) \quad (X_i^+)^j = (t_i^+, W_i^+)^j \\
(W_i^-)^j &= \left(\sum_{k=1}^{N_a}(A_k(t_i^-))^j\right)\left(1 + \mathbb{T}_i^g\right)\exp\left(-(\Delta t)\mathbb{T}^f\right) \\
(A_k(t_i^+))^j &= (W_i^+)^j \hat{p}_{k_i}\left((W_i^+)^j, t_i, \theta_p\right)\big) \\
(\hat{q}_i(\cdot), \hat{\boldsymbol{p}}_i(\cdot)) &\in \mathcal{Z}\left((W_i^-)^j, (W_i^+)^j, t_i\right)
\end{aligned} \tag{34}$$

for $i = 0, \ldots, M$, $t_i \in \mathcal{T}$, the superscript $j$ represents the $j$-th path of joint returns, and $\theta_q$ and $\theta_p$ represent the parameters of the two controls. The optimal parameters found by training the NNs are used to generate the control functions $\hat{q}^*(\cdot) := \hat{q}(\cdot; \theta_q^*)$ and $\hat{p}^*(\cdot) := \hat{p}(\cdot; \theta_p^*)$ respectively.

The NN approach sketched above relies on a set of joint asset return paths being available, but it is agnostic about the source of these paths. One possibility is to specify a parametric model for returns, estimate it based on historical data, and then run Monte Carlo simulations of it to generate the paths. Another possibility is to use a bootstrap procedure such as the stationary block bootstrap to directly sample return paths from the historical data (Politis and Romano, 1994; Politis and White, 2004; Patton et al., 2009). This method involves repeatedly drawing blocks of return data of random size (with replacement) from observed historical return series. The random block size follows a geometric distribution with a specified expected block size. Return blocks are drawn simultaneously for all assets considered, thus preserving the joint correlation structure within each block. A return path for the entire investment horizon is formed by pasting enough blocks together, and this is repeated many times to generate a sufficiently large number of paths. Although we primarily rely here on this bootstrap approach, some of tests discussed below do rely on simulated returns from a parametric model. Of course there are other methods for generating data from a single historical path such as Generative Adversial Netorks (GANs) (Yoon et al., 2019; van Staden et al., 2025). However, we will not pursue these alternative approaches here.

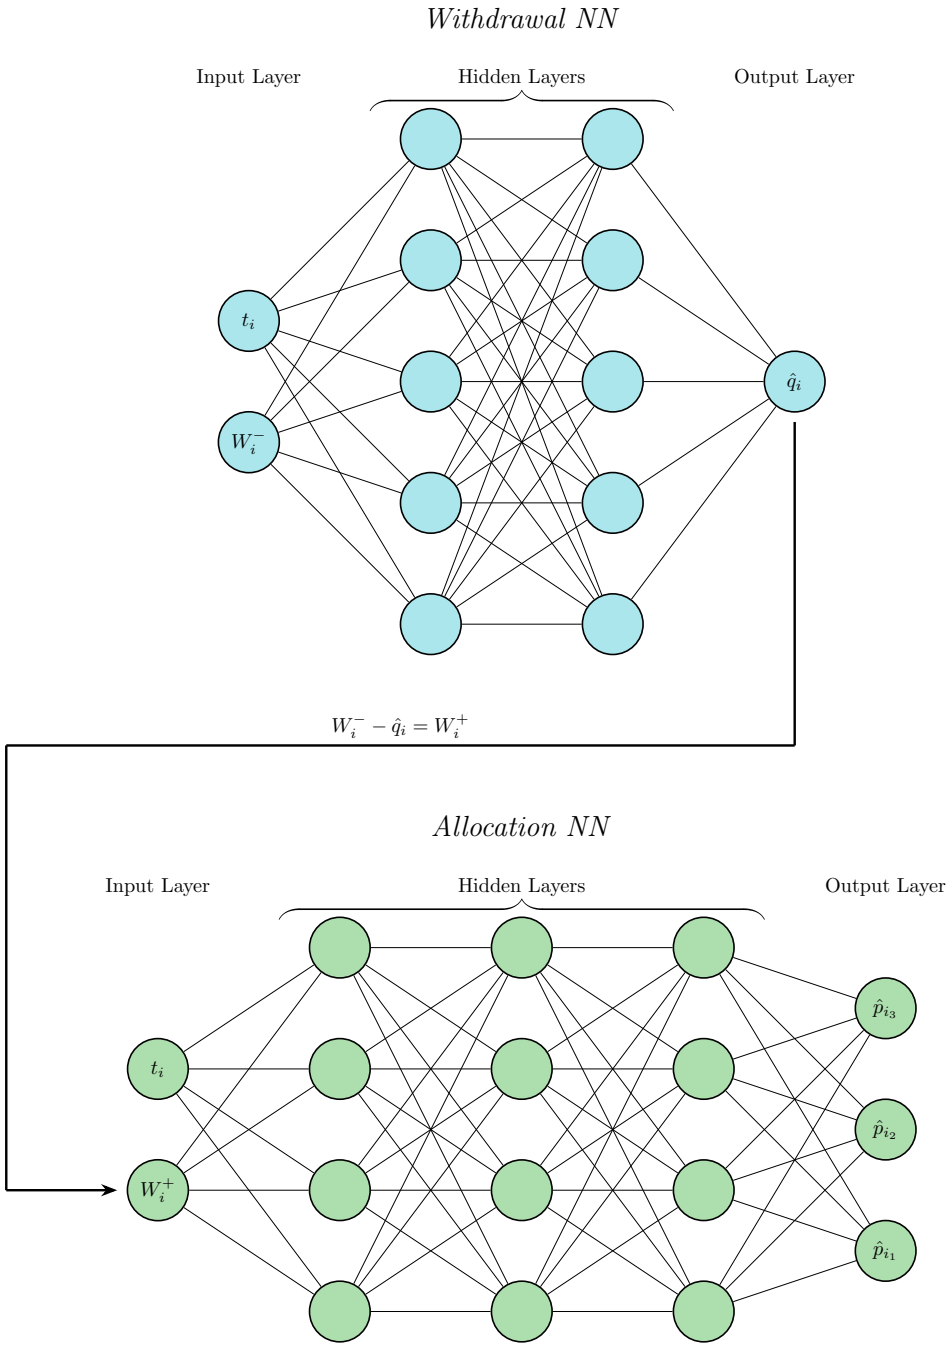


FIGURE 1: *An illustration of possible withdrawal and allocation NNs. The input layer for the withdrawal NN contains the rebalancing time $t_i$ and pre-withdrawal wealth $W_i^-$. The output layer has a single node which is used to determine the withdrawal amount $\hat{q}_i$ at $t_i$ via a modified sigmoid function. The post-withdrawal wealth $W_i^+ = W_i^- - \hat{q}_i$ is one of the two components of the input layer for the allocation NN, along with $t_i$. This illustration assumes there are three assets, so the output layer for the allocation NN has three nodes. The asset allocations $\hat{p}_i$ are determined through a softmax activation function.*

| Base Case Investment Scenario Parameters | |
|---|---|
| Retiree | 65 year old Canadian male |
| Mortality table | CPM2014 |
| Tontine group gain | $G = 1$ |
| Tontine gain rate | $\mathbb{T}_i^g = \mathfrak{q}_{i-1}/(1-\mathfrak{q}_{i-1})$ |
| Tontine fees | $\mathbb{T}^f = 50$ basis points per annum |
| Investment horizon | $T = 30$ years |
| Initial portfolio value | $W_0 = 1000$ |
| Minimum annual withdrawal | $q_{\min} = 40$ |
| Maximum annual withdrawal | $q_{\max} = 80$ |
| Withdrawal/rebalancing times | $t = 0,1,\ldots,29$ |
| Borrowing spread | 200 basis points per annum |
| Stabilization term (see equation (27)) | $\epsilon = -10^{-4}$ |

TABLE 1: *Base case investment scenario parameters. All monetary units are in thousands of real USD.*

# 4 Initial Tests

We begin by considering the same scenario and environment analyzed by Forsyth et al. (2024). This involves a case where there are just two assets, a bond index formed by rolling over investments in 30-day U.S. Treasury bills and an equity market index, represented by the CRSP value-weighted U.S. stock market index.[5] Both of these indexes are considered in real terms (i.e. adjusted for inflation). Forsyth et al. (2024) estimate a double exponential jump-diffusion model for the returns of these two indexes using monthly data from 1926 through 2020. Since this parametric model will only be used in our initial test, its exact details are not of paramount importance here. Readers who want more information can find further description of the returns model and data in Sections 4 and 9 of Forsyth et al. (2024). Parameter estimates are reported in Table 1 of that paper.

The investment scenario considered involves a just-retired 65 year old male with initial wealth of 1,000 (in thousands of USD). This retiree participates in a tontine pool with gain rate defined as in equation (3). The relevant mortality table is the CPM2014 table from the Canadian Institute of Actuaries.[6] The group gain defined in equation (7) is assumed for the purposes of these tests to be constant and equal to one, so that the tontine gain rate (9) simplifies to $\mathbb{T}_i^g = \mathfrak{q}_{i-1}/(1-\mathfrak{q}_{i-1})$. The retiree's investment horizon is 30 years. The tontine fee $\mathbb{T}^f$ is set to 50 basis points per annum. At the start of each year, the retiree withdraws an amount from the invested portfolio of at least 40 and at most 80. The portfolio is rebalanced annually. In case of insolvency, the investor is forced to continue with minimum annual withdrawals and to accumulate debt. The relevant real interest rate is the realized annual return on 30 day Treasury bills plus a borrowing spread of 2%. All of these details are summarized in Table 1.

## 4.1 Validation of NN Approach in Base Case

Our first test is an attempt to replicate the base case results reported in Forsyth et al. (2024). The approach of Forsyth et al. uses numerical partial differential equation (PDE) methods to calculate the optimal withdrawal and allocation controls under an assumed parametric model. As in Shirazi (2024), we use allocation and withdrawal NNs to approximate the same controls. Our implementation uses the PyTorch library (Paszke et al., 2019) with the Adam optimizer (Kingma and Ba, 2017). Further computational details for the NNs are provided in Appendix A.

As indicated by equation (27), the investor trades off higher expected withdrawals during retirement against expected shortfall, a measure of the risk of depleting the investor's portfolio. This tradeoff can

[5]Specifically, the results we present were calculated based on data from Historical Indexes, ©2020 Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business. Wharton Research Data Services (WRDS) was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

[6]This mortality table is available at www.cia-ica.ca.

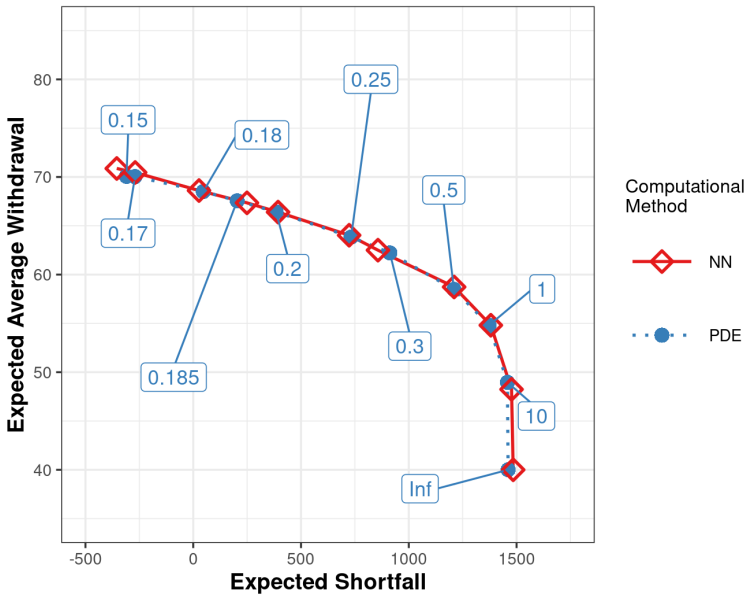


FIGURE 2: *Initial test: comparison of performance of NN approach (solid line with diamond symbols) with that of PDE solution of Forsyth et al. (2024) (dotted line with circles). Values of $\kappa$ are shown in boxes attached to the performance of the PDE solution at various points. Red diamonds indicate the NN solution performance for values of $\kappa$ closest to those shown for the PDE solution. Scenario parameters given in Table 1. Return paths generated by Monte Carlo simulation of the parametric double exponential jump diffusion model of Forsyth et al. (2024). Results shown based on $2.56 \times 10^6$ simulated return paths. The investment portfolio is allocated between the value-weighted CRSP U.S. equity index and the CRSP 30-day U.S. T-bill index.*

be seen in a plot of expected average withdrawals (expected total withdrawals divided by the number of withdrawals, to get the reward in annual terms) vs. expected shortfall. The data underlying this plot for the PDE approach of Forsyth et al. (2024) can be found by using numerical PDE methods to solve for the optimal withdrawal and allocation controls across a range of values of the risk-aversion parameter $\kappa$, and then using Monte Carlo simulations of paths of asset returns under the assumed parametric model. Similar data can be generated for our NN approach by training the allocation and withdrawal NNs on a set of Monte Carlo return paths, and then evaluating the training performance of the model on these paths.

Figure 2 shows the results for a range of values of $\kappa$ that are shown in the boxes attached to the results for the PDE approach. The horizontal axis shows the expected shortfall, while the vertical axis is the expected average withdrawal.[7] Setting $\kappa$ to a very high value (designated in this plot as "Inf") places so much emphasis on maximizing the expected shortfall that the withdrawals are held to the minimum value of 40 each year. As $\kappa$ is reduced, withdrawals increase and the expected shortfall is reduced. The main point here is that there is generally a reasonably close correspondence between the NN approach described in this paper with the earlier work of Forsyth et al. (2024). However, it is also worth noting that for relatively low values of $\kappa$ the retiree can withdraw an average of almost 70 each year in real terms without incurring negative expected shortfall. This is close to 7% of the portfolio value at the inception of retirement, far exceeding the standard 4% rule (Bengen, 1994) which would correspond to annual withdrawals of 40.

As an additional test, we train the NN using returns produced by block bootstrap resampling of the CRSP data from 1926:1 to 2022:12. Figure 3 shows the training performance of our NN approach for several expected block sizes, ranging from 1 month to two years. Relatively large differences are seen towards the right side of the plot, corresponding to high values of $\kappa$. However, cases such as these with expected shortfall in excess of 1,000 are quite unrealistic: recall that these cases are average values for the worst 5% of outcomes, and they exceed the initial portfolio value of 1,000. Moreover, these residual values after 30 years

[7] The plot shows the expected average withdrawal rather than expected total withdrawals as in the objective function (27). This gives the interpretation of expected average annual income, in thousands of USD.

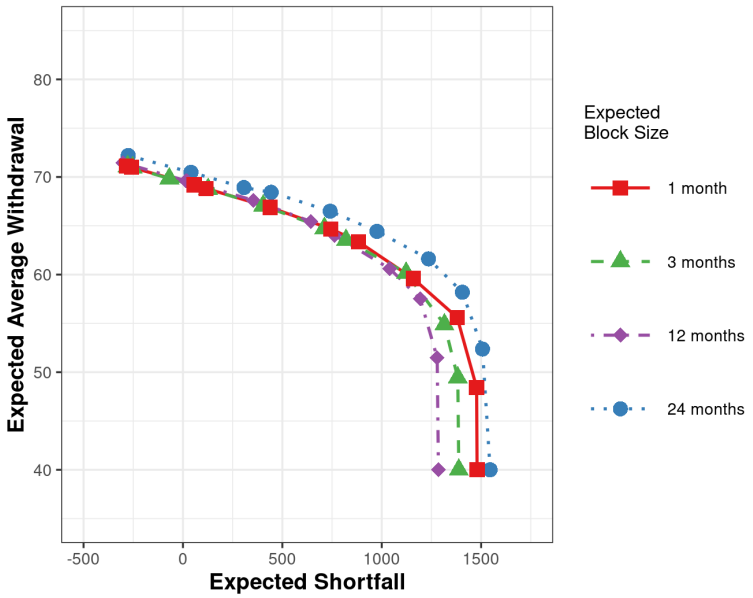


FIGURE 3: *Expected block size test: comparison of NN training performance for resampled returns with various expected block sizes. Scenario parameters are given in Table 1. Training data contains* $2.56 \times 10^5$ *paths of bootstrapped return data, with expected block sizes indicated in the figure legend. The investment portfolio is allocated between the value-weighted CRSP U.S. equity index and the CRSP 30-day U.S. Treasury bill index, both expressed in real terms. The same values of* $\kappa$ *are used as in Figure 2.*

of retirement do not subsequently go to the retiree's estate, but rather are reallocated to other members of the tontine pool. As such, it is of more practical interest to trade off lower expected shortfall in exchange for higher average annual withdrawals. For example, an expected shortfall somewhere in the range of zero to 500 would correspond to a very low probability of portfolio depletion having to be financed by debt, while also allowing relatively high average withdrawals and not being likely to result in large transfers to other tontine pool participants.[8] With the caveat in mind that results could be somewhat different, particularly for unrealistically high values of $\kappa$, henceforth we will restrict attention to the case of an expected block size of 24 months.[9]

A common concern in the NN context is overfitting to the training data. We evaluate this by conducting an out-of-sample test, comparing performance for training data with that for another set of bootstrapped return paths generated with a different random seed ("testing data"). Figure 4 shows a very close correspondence between the two, with the largest discernible differences being quite small and for high values of $\kappa$ which lead to high expected shortfall. Based on this, it does not appear that overfitting to the training data is much of a concern in our context.

# 5 Some Alternative Assumptions

So far we have assumed that mortality for the tontine pool follows the CPM2014 mortality table and that the group gain (7) is constant and equal to one. We now investigate the effects of altering these assumptions, beginning with changes to the assumed mortality table while leaving the group gain unaltered. First, we consider what we refer to as *unexpected* improvements in mortality. This means that there is an improvement in mortality (i.e. a reduction in mortality rates), but that it is not factored into training the NN. Put differently, we train the model based on the CPM2014 table, but then test it on another data set with

[8]Since the expected shortfall measure is an average across the worst 5% of cases, it is of course possible for individual paths to result in either much higher or much lower terminal wealth than this average. However, such cases would be quite rare.

[9]This is roughly equal to the average value of the optimal expected block sizes for the CRSP equity and 30-day Treasury bill indexes reported in Table 2 of Forsyth et al. (2024).

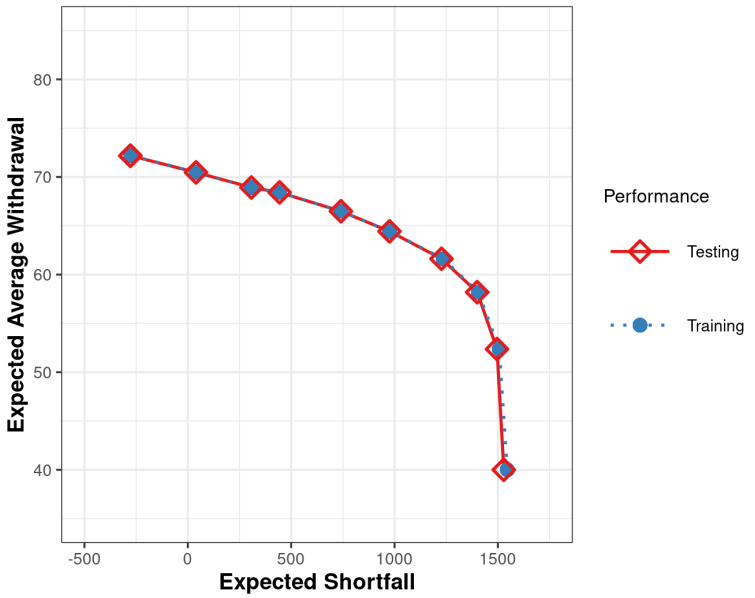


FIGURE 4: *Out-of-sample test: comparison of NN training performance results (solid line with diamond symbols) with out-of-sample test (dotted) line with circles. Scenario parameters given in Table 1. Both training and testing data consist of* $2.56 \times 10^5$ *paths of bootstrapped return data, generated with different random seeds and an expected block size of 24 months. The investment portfolio is allocated between the value-weighted CRSP U.S. equity index and the CRSP 30-day U.S. Treasury bill index, both expressed in real terms. The same values of* $\kappa$ *are used as in Figure 2.*

mortality rates that are uniformly reduced by 10%. Second, we consider *expected* mortality improvements. Note that the CPM2014 table already reflects forecasted improvements, so we are effectively just considering improvements beyond those already built into the table. We again consider a uniform reduction by 10%, but now these advances are incorporated into the NN model training.

From the perspective of our retiree who is assumed to have a 30-year horizon, any mortality improvements should result worse performance due to receiving lower mortality credits. This observation is borne out in Figure 5. This is particularly clear for high values of $\kappa$ (i.e. cases with relatively low average annual withdrawals and high expected shortfall). As argued above, however, such cases are likely to be of limited practical interest. In the more relevant range of smaller but positive expected shortfall (e.g. in the range of 100 to 500), the decline in average annual withdrawals is noticeable but quite minor. It is also clear that incorporating the mortality improvement into training the NN leads to better outcomes, partly offsetting the performance degradation due to reduced mortality credits.

As noted in Section 2.2, the group gain $G$ is a multiplicative adjustment factor that adjusts for discrepancies between actual and expected mortality credits. So far we have simply assumed $G = 1$, but in reality $G$ will vary somewhat from year to year, with larger deviations expected to be seen for tontine pools with fewer members. In a given year $G$ would be determined by a complex mixture of factors, including how many pool members die, the ages and initial investments in the pool of those who die, and the investment and withdrawal strategies that were followed by those deceased members. The simulations of Fullmer and Sabin (2019) give a mean for $G$ of just about one with a standard deviation of 0.1 once the pool has reached sufficient size (approximately 10,000 participants). Forsyth et al. (2024) apply the control from the base case (constant $G$) to Monte Carlo simulations with randomized group gain, assuming a normal distribution with mean 1 and standard deviation 0.1, and find negligible differences in performance. However, Fullmer and Sabin (2019) report a higher standard deviation of about 0.4 with a smaller pool size consisting of about 1,000 members. It is therefore of interest to investigate the effects of this increased randomness, particularly since it can be easily accommodated in training the NNs.[10]

[10]By contrast, as group gain randomness introduces an extra state variable, this would increase the complexity of the PDE

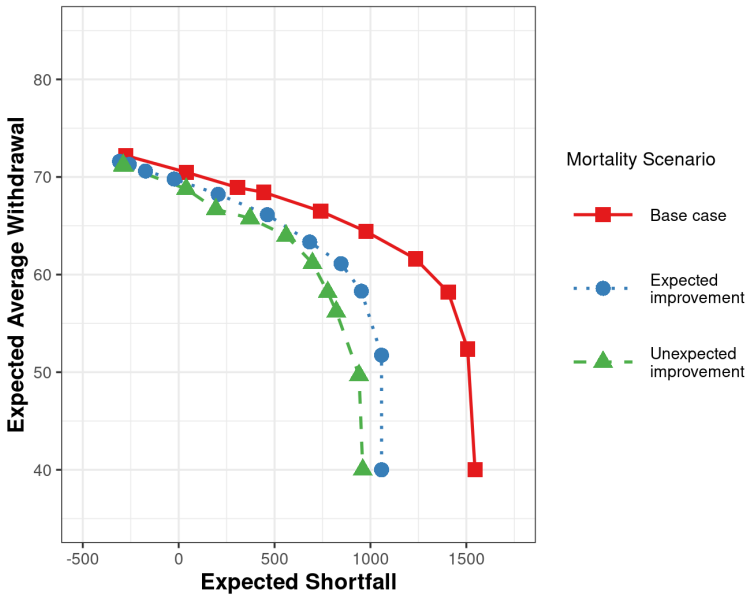


FIGURE 5: *Mortality test: comparison of base case with mortality improvement scenarios. Expected improvement is a situation where mortality rates are 10% lower than in the base case which is based on the CPM2014 mortality table, and this longevity gain is factored into training the NN. Unexpected improvement is similar, but the improvement is not incorporated in training. Remaining scenario parameters given in Table 1. Training data contains* $2.56 \times 10^5$ *paths of bootstrapped return data with an expected block size of 24 months. The investment portfolio is allocated between the value-weighted CRSP U.S. equity index and the CRSP 30-day U.S. Treasury bill index, both expressed in real terms. The same values of* $\kappa$ *are used as in Figure 2.*

We consider cases with a mean of $G$ equal to one and standard deviations of $\sigma \in \{0.1, 0.2, 0.3, 0.4\}$ (in addition to the base case where $\sigma = 0$). To generate maximum impact, we assume a uniform distribution with these characteristics. Figure 6 shows that increased uncertainty in the group gain leads to lower performance, but the effects are fairly minor and again largest for cases with high expected shorfall (in excess of the original 1,000 value of the investment portfolio at the outset of retirement). Comparing with Figure 5, we see that the degradation in performance for low expected average withdrawals is much less severe with random group gain as opposed to improved mortality. In addition, a further comparison of this plot with that in Figure 1 of Forsyth et al. (2024) suggests that even in the face of high group gain uncertainty, being in the tontine offers dramatically better performance compared to an optimized investment strategy without tontine participation or to a simple strategy with minimum withdrawals and constant proportional investment allocation. The two plots are not completely comparable since that of Forsyth et al. is based on a parametric model, while Figure 6 uses bootstrap resampling, but the differences between the cases with and without the tontine in Forsyth et al. (2024) are vastly larger than the range of the various cases of random $G$ considered here. For example, for average annual withdrawals of 60, the range here for expected shortfall is approximately 1,050 to 1,230. On the other hand, in Forsyth et al. (2024) for this level of average withdrawals the expected shortfall is a bit over 1,000 with the tontine vs. around $-250$ without it.

# 6 Enhanced Investment Opportunities

As noted above, an advantage of the NN approach is that it is relatively easy to incorporate additional risk factors, compared to PDE-type methods. This makes it possible to explore scenarios with a larger set of potential investments. The examples explored to here have all involved just two investment assets, the CRSP value-weighted U.S. equity index and the CRSP 30-day U.S. Treasury bill index, both denominated in

framework of Forsyth et al. (2024).

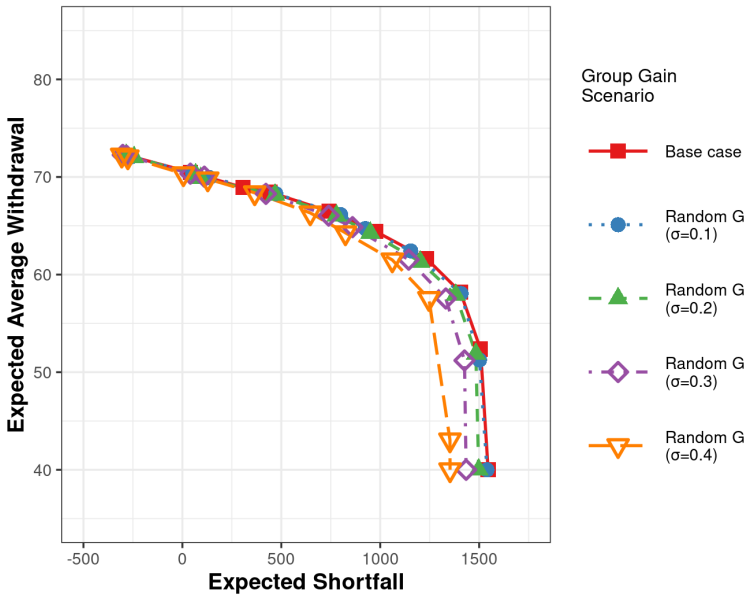


FIGURE 6: *Comparison of the effects of randomized group gain G with base case scenario. Random G cases assume a uniform distribution with a mean of one and standard deviation as indicated. Remaining scenario parameters given in Table 1. Training data contains* $2.56 \times 10^5$ *paths of bootstrapped return data with an expected block size of 24 months. The investment portfolio is allocated between the value-weighted CRSP U.S. equity index and the CRSP 30-day U.S. Treasury bill index, both expressed in real terms. The same values of* $\kappa$ *are used as in Figure 2.*

real USD. In this section we switch to working in real Canadian dollars (CAD), with the following possible investments. First, on the equity side, we allow the investor to allocate his portfolio between the U.S. CRSP value-weighted index, the U.S. CRSP equal-weighted index, the Canadian Financial Markets Research Centre (CFMRC) value-weighted Canadian equity index, and the CFMRC equal-weighted Canadian equity index. To reiterate, these four potential equity investments are all denominated in real CAD.[11] On the fixed income side, the investor can choose between the Canadian 30-day Treasury bill index and the Canadian long-term bond index, both of which are obtained from CFMRC and expressed in real CAD. One unfortunate aspect of the Canadian data is that it is only available for all of these Canadian indexes as of November 1950, giving us 866 total observations through the end of 2022. Aside from the change in investment opportunities and the switch to real CAD (in thousands), the scenario considered here is the same as outlined above in Table 1.

Table 2 provides simple descriptive statistics for real monthly log returns in CAD of the six investment assets over the sample period. Unsurprisingly, the two fixed income alternatives (30-Day Treasury bill index and long-term bond index) have the lowest mean and median, but also the lowest standard deviation. They also have the smallest range from the minimum to maximum observed values, and from the 5th to 95th percentiles. Between the two fixed income investments, the Treasury bill index exhibits lower risk (standard deviation) and lower average returns. Turning to the four equity indexes, the lowest average return is for the Canadian value-weighted equity index while the highest average return is for the Canadian equal-weighted equity index. Mean returns for the two U.S. indexes lie in the middle, with the equal-weighted option offering the higher average. In addition to having higher mean returns, the two equal-weighted indexes have higher standard deviations compared to their value-weighted counterparts. The two Canadian equity indexes have the highest correlation of returns, but not by much over that of the two U.S. equity indexes. As a final observation, while the differences in average monthly returns may not seem large, compounding them over a long-term investment horizon of several decades leads to dramatically different total returns at the end of the horizon. A potential concern is that equal-weighted indexes are more expensive to maintain than value-

[11] Nominal returns for the U.S. indexes are converted to CAD using historical exchange rate data from the Bank of Canada, and then adjusted for inflation using Statistics Canada's consumer price index.

| Descriptive Statistics<br>Monthly Real Log Returns, 1950:11 to 2022:12 ($N = 866$) | | | | | | |
|---|---|---|---|---|---|---|
| Asset 1: Canadian 30-Day Treasury Bill Index<br>Asset 2: Canadian Long-Term Bond Index<br>Asset 3: Canadian Value-Weighted Equity Index<br>Asset 4: Canadian Equal-Weighted Equity Index<br>Asset 5: U.S. Value-Weighted Equity Index<br>Asset 6: U.S. Equal-Weighted Equity Index | | | | | | |
| | Asset 1 | Asset 2 | Asset 3 | Asset 4 | Asset 5 | Asset 6 |
| Mean | 0.0011 | 0.0022 | 0.0048 | 0.0081 | 0.0059 | 0.0071 |
| Std. Deviation | 0.0047 | 0.0249 | 0.0436 | 0.0518 | 0.0403 | 0.0500 |
| Minimum | -0.0224 | -0.1130 | -0.2638 | -0.3326 | -0.2541 | -0.3165 |
| 5th percentile | -0.0067 | -0.0365 | -0.0664 | -0.0732 | -0.0640 | -0.0750 |
| Median | 0.0014 | 0.0017 | 0.0087 | 0.0126 | 0.0104 | 0.0080 |
| 95th percentile | 0.0084 | 0.0429 | 0.0644 | 0.0754 | 0.0621 | 0.0807 |
| Maximum | 0.0168 | 0.1421 | 0.1682 | 0.2497 | 0.1411 | 0.2704 |
| | | | Correlation Matrix | | | |
| | Asset 1 | Asset 2 | Asset 3 | Asset 4 | Asset 5 | Asset 6 |
| Asset 1 | 1.0000 | 0.3306 | 0.0867 | 0.0614 | 0.1416 | 0.1148 |
| Asset 2 | 0.3306 | 1.0000 | 0.1988 | 0.1097 | 0.2016 | 0.1095 |
| Asset 3 | 0.0867 | 0.1988 | 1.0000 | 0.8783 | 0.7414 | 0.7248 |
| Asset 4 | 0.0614 | 0.1097 | 0.8783 | 1.0000 | 0.6075 | 0.7225 |
| Asset 5 | 0.1416 | 0.2016 | 0.7414 | 0.6075 | 1.0000 | 0.8616 |
| Asset 6 | 0.1148 | 0.1095 | 0.7248 | 0.7225 | 0.8616 | 1.0000 |

TABLE 2: *Descriptive statistics for monthly real log returns from 1950:11 to 2022:12. All return series are measured in real CAD.*

weighted indexes, due to higher rebalancing costs. To partially address this, in an admittedly *ad hoc* way, we have reduced the historical real returns for both equal-weighted indexes by 1% per year (compounded monthly) in our analysis below.

Figure 7 compares the performance when the investor has access to the set of assets described above with a situation where the investor is restricted to two investment assets (in this case, the two assets are the U.S. value-weighted equity index and the Canadian 30-day Treasury bill index). In order to generate the curves shown, different values of $\kappa$ were used for the two cases, so here we do not plot any symbols at points with various $\kappa$ values.[12] As we would expect, performance is clearly improved when more investment assets are available, although once again the largest differences are for cases with high expected shortfall which are less relevant. It is also worth observing that the two asset case shows much better performance here compared to earlier two asset cases shown in Figure 2. The highest expected shortfall is about 2,000 here compared to 1,500 in the earlier plot. Expected shortfall is still positive here with expected average withdrawals in the low 70s, compared to the high 60s earlier. There are two reasons for these differences. First, the investment assets are different (here we have a Canadian Treasury bill index instead of a U.S. Treasury bill index) and we are using real CAD instead of real USD. Second, as mentioned above the Canadian data is only available beginning in November 1950, while the U.S. data goes back to 1926. Consequently Figure 7 excludes highly volatile periods associated with the Great Depression and the Second World War, which are represented in our earlier plots.

In addition to the two scenarios described above, Figure 7 shows the single point that results from adopting the simple policy of consistently maintaining equal weights across all six investment assets and withdrawing the minimum of 40 every year, without being in a tontine pool. This results in slightly negative ES (approximately $-45$). Of course, any higher withdrawals would lead to lower ES. The enhanced efficiency

[12] In particular, in order to generate instances with relatively low expected shortfall in the multi-asset case, we need to include lower values of $\kappa$ than with the two asset case. The lowest case included in Figure 7 is $\kappa = 0.02$.

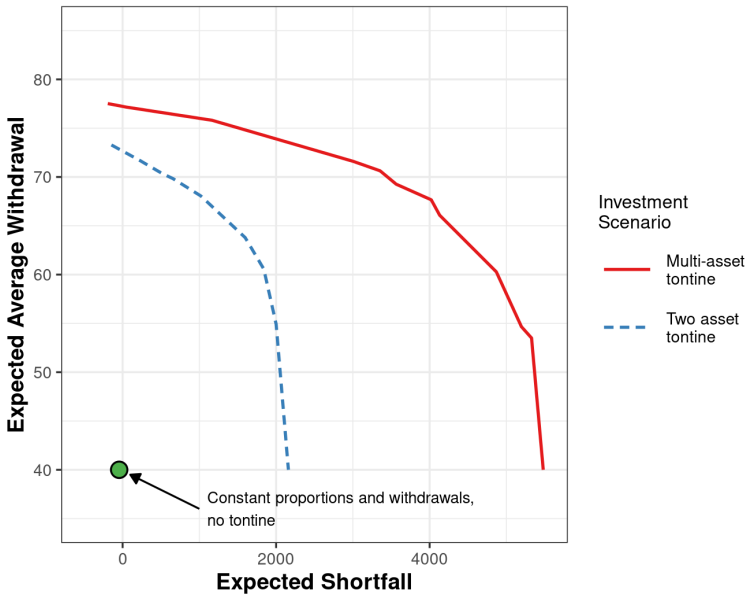


FIGURE 7: *Comparison of multi-asset tontine with two asset tontine. In the multi-asset case, the investor has access to four equity market indexes: the CRSP value-weighted and equal-weighted U.S. market indexes and the CFMRC value-weighted and equal-weighted Canadian market indexes. The investor can also allocate funds to either the CFMRC Canadian 30-day Treasury bill index and the CFMRC Canadian long-term bond index. All of these indexes are expressed in real CAD. In the two asset case, the investor is constrained to investing the CRSP value-weighted U.S. equity market index and the CFMRC Canadian 30-day Treasury bill index. Remaining scenario parameters given in Table 1. Training data contains* $2.56 \times 10^5$ *paths of bootstrapped return data with an expected block size of 24 months. Different values of* $\kappa$ *are used for the two cases, roughly similar to those used in Figure 2. Also shown as a single point is a policy with constant proportions (equal weights across the six investment assets) and constant withdrawals of the annual minimum, with no tontine.*

offered by the tontine pool and a dynamic investment allocation strategy is clear, even with just two assets: the same ES can be attained with expected average withdrawals of more than 70; alternatively, an ES of over 2,000 can be achieved with constant withdrawals of 40 per year. Expanding the set of possible investments leads to further potential improvements.

We take a more detailed look by considering a single point on the multi-asset curve in Figure 7. The point we consider is generated by $\kappa = 0.03$. It produces expected shortfall of 46.25 and expected average withdrawal of 77.16. Given the relatively high average withdrawal and low (but positive) expected shortfall, cases similar to this are likely to be of most interest to the retiree. We begin by examining the withdrawals in more detail. Figure 8 displays a heatmap documenting the dependence of annual withdrawals on (pre-withdrawal) real wealth and time.[13] Outside of a very narrow range, the annual withdrawals are all either at the maximum of 80 or the minimum of 40. Of course, withdrawals are reduced when available real wealth is lower.

While Figure 8 indicates what the model implies withdrawals should be given time and prevailing wealth, it does not say anything about what withdrawals are likely to be since there is no information about the distribution of wealth over time. The left plot of Figure 9 shows the 5th, 50th, and 95th percentiles of this distribution over time. The median level of wealth stays fairly flat over the first decade, declines steadily over the second decade, and subsequently rises quite dramatically, reflecting high mortality credits from the tontine towards the end of the horizon. The 95th percentile rises rapidly over the first few years, has a fairly flat period lasting about 10 years, and then rises rapidly. The 5th percentile declines quickly in the first few

[13]This plot and other heatmaps shown below are basically a sequence of 30 vertical bars that have been pasted together. The represented decision point is at the left edge of the bar (i.e. the first bar spans the period $t = [0,1]$ and is for the decision at $t = 0$, the last one spans $t = [29,30]$ and provides the decision at $t = 29$).

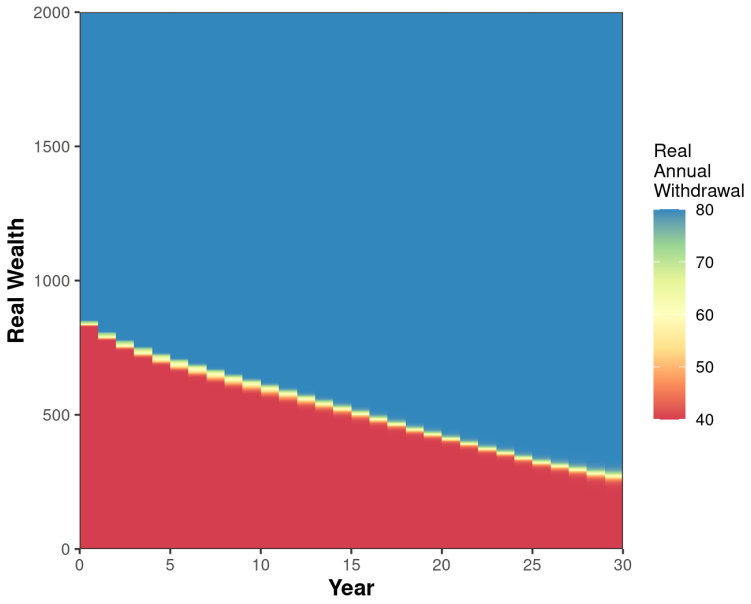


FIGURE 8: *Heatmap of real annual withdrawals as a function of pre-withdrawal real wealth and time ($t = 0,1,\ldots,29$).*

years, then more slowly but steadily for the next couple of decades, and then modestly recovers over the last few years. It might seem odd that in this case the wealth level does not recover much in comparison to, for example the median. The reason for this becomes clear from the right plot of Figure 9, which shows the same percentiles of the annual withdrawals distribution throughout the horizon (i.e. at $t = 0,1,\ldots,29$). The 50th and 95th percentiles are constant at the maximum permitted value of 80 throughout. By contrast, the 5th percentile starts at 80, but drops quickly to the minimum amount of 40. It remains there for a couple of decades, before recovering to the maximum level in the last few years. In this case high mortality credits towards the end lead to a recovery in withdrawals, more so than an increase in wealth. Of course, in the other cases this does not happen. Since withdrawals are already at the upper bound, the mortality credits go entirely towards building wealth.

Continuing with the multi-asset tontine case for $\kappa = 0.03$, we now explore the model's implied asset allocation. Although a total of six assets are available (two U.S. equity indexes, two Canadian equity indexes, and two Canadian fixed income indexes), the allocation control suggests a portfolio that is primarily dominated by just three assets: Canadian Treasury bills, the Canadian equal-weighted equity index, and the U.S. value-weighted equity index. Figure 10 shows the 5th, 50th, and 95th percentiles of the percentage invested in each of the six assets. The initial investment at $t = 0$ is a rough even split between the Canadian equal-weighted equity index and the U.S. value-weighted equity index. The median investment fraction in both of these equity indexes trends downward over time, but more slowly for the Canadian index. The 95th percentiles remain consistently around 50% before dropping off in the last few years. The median percentage invested in Canadian Treasury bills climbs steadily over the entire horizon, to the point where the portfolio is almost entirely in this asset at the end. The 5th percentile is flat at basically zero over the first two decades, but increases rapidly to about 50% at the end. The other fixed income asset, long-term bonds, is barely used as even the 95th percentile of the investment fraction peaks at about 10%. The choice of short-term Treasury bills over long-term bonds is perhaps due to superior inflation hedging (Spierdijk and Umar, 2015). The Canadian value-weighted equity index is not used much: similar to long-term bonds, even the 95th percentile of the percentage invested tops out at close to 10%. The preference for the domestic equal-weighted index over its value-weighted counterpart makes sense in light of Table 2 as the equal-weighted case offers signficantly higher returns on average with what appears to be modestly elevated risk.[14] In most

[14]Recall that our investment results here are based on reducing the equal-weighted real returns by 1% per year compared to

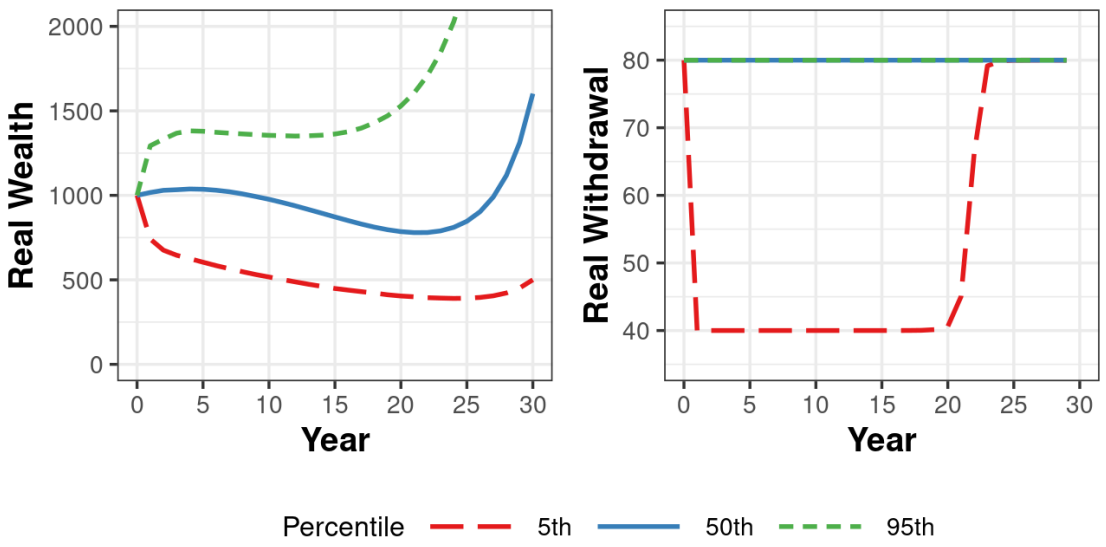


FIGURE 9: *Percentiles of pre-withdrawal real wealth (left) and real withdrawals (right) for the multi-asset tontine with* $\kappa = 0.03$.

cases, the U.S. equal-weighted index is not used much: the median allocation is basically zero throughout. However, the 95th percentile rises rapidly from zero to over 50% in the first few years, before declining gradually back to zero over the ensuing decade, and remaining at that level thereafter. Outside of these relatively rare cases for the U.S. equal-weighted index, the model generally suggests investing mostly in the Canadian equal-weighted and U.S. value-weighted equity indexes early on, and gradually shifting into short-term Treasury bills over time for this particular choice of $\kappa$.

More detailed insight into the model's recommended allocation can be found in Figure 11, which provides heatmaps of the allocation to the three main assets (Canadian Treasury bill index and equal-weighted equity index, and U.S. value-weighted equity index) as a function of real wealth over time.[15] At relatively high levels of real wealth, the NN suggests investing almost entirely in Canadian Treasury bills, especially near the end of the 30-year horizon. Preferring a safer fixed income investment over riskier equity alternatives near the end of the horizon for high wealth makes sense since withdrawals will be at the upper bound and there is no need to take on much risk (and, as noted above, mortality credits would be high, further reducing the incentive to take on risk). At very low levels of wealth near the end of the horizon, the investment portfolio is dominated by the Canadian equal-weighted equity index. Of course, it is not surprising that there would be a tendency to take a riskier approach here since a positive outcome may lead to some recovery, both in terms of enhanced withdrawals and higher expected shortfall. It is interesting, however, that in such cases the U.S. value-weighted index is avoided. Although at the initial real wealth of 1,000 roughly half of the portfolio is allocated to this index, the NN control reduces this to basically zero outside of a fairly narrow range of wealth that generally declines over time, except during the last 5 years.

The choice of the Canadian equal-weighted equity index over its value-weighted counterpart reflects its strongly superior returns over the sample period from November 1950 through December 2022 (see Table 2). On the other hand, the U.S. equal-weighted equity index also outperformed its value-weighted counterpart during this period. However, in the U.S. case the NN typically suggests investing in the lower average return index. One potential reason may be better diversification, as the correlation between real CAD returns for the Canadian equal-weighted equity index and the U.S. equal-weighted index was about 0.72, and the corresponding correlation for the U.S. value-weighted index was about 0.61. It is also worth noting that the statistics reported in Table 2 for the equal-weighted indexes are prior to our *ad hoc* reduction by 1%

those used in Table 2 in an effort to adjust for the higher cost of maintaining equal weights.

[15] Note that these heatmaps depend on post-withdrawal wealth, unlike Figure 8 which depends on pre-withdrawal wealth.

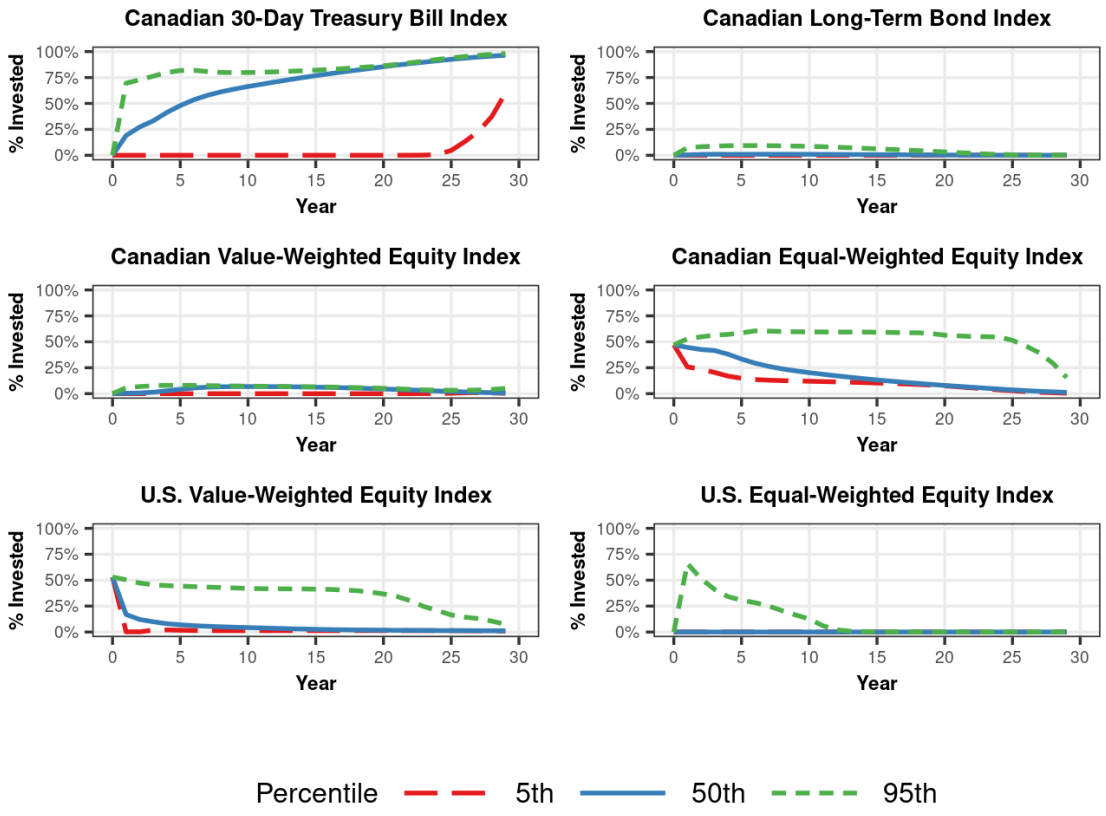


FIGURE 10: *Percentiles of pre-withdrawal real wealth (left) and real withdrawals (right) for the multi-asset tontine with* $\kappa = 0.03$.

annually. This reduction shrinks the gap between average returns for the U.S. equal-weighted and value-weighted indexes to a very small amount, but the corresponding gap for the Canadian indexes is quite a bit larger. Another consideration from Table 2 is that median (not average) monthly real returns were higher for the equal-weighted equity index compared to its value-weighted counterpart in Canada but not in the U.S.[16]

As a final experiment, we split the historical data roughly into halves. We use the first part (1950:11 to 1985:12) to generate the training data, and the second part (1986:1 to 2022:12) to generate the testing data. Figure 12 shows the results. Clearly the performance degrades somewhat for the testing data, but the differences are generally not huge. This provides further confidence in the out-of-sample performance of the NN solution approach.

[16]As noted above in the discussion accompanying Figure 10, there are some rare cases where the model suggests investing a significant fraction in the U.S. equal-weighted equity index during the first few years. We can deduce from the heatmaps in Figure 11 what circumstances would lead to this switch from the U.S. value-weighted index to its equal-weighted counterpart. This would happen if the level of real wealth declined substantially during the first few years. Suppose, for example, that after 5 years real wealth has declined by 50% from its initial 1,000 to 500. The heatmaps in Figure 11 show basically zero allocation to Treasury bills and the U.S. value-weighted equity index, and roughly a 60% allocation to the Canadian equal-weighted equity index. The remaining 40% would then be mostly invested in the U.S. equal-weighted equity index since the only other two assets (Canadian long-term bond index and Canadian value-weighted equity index) are barely used, as shown in Figure 10.

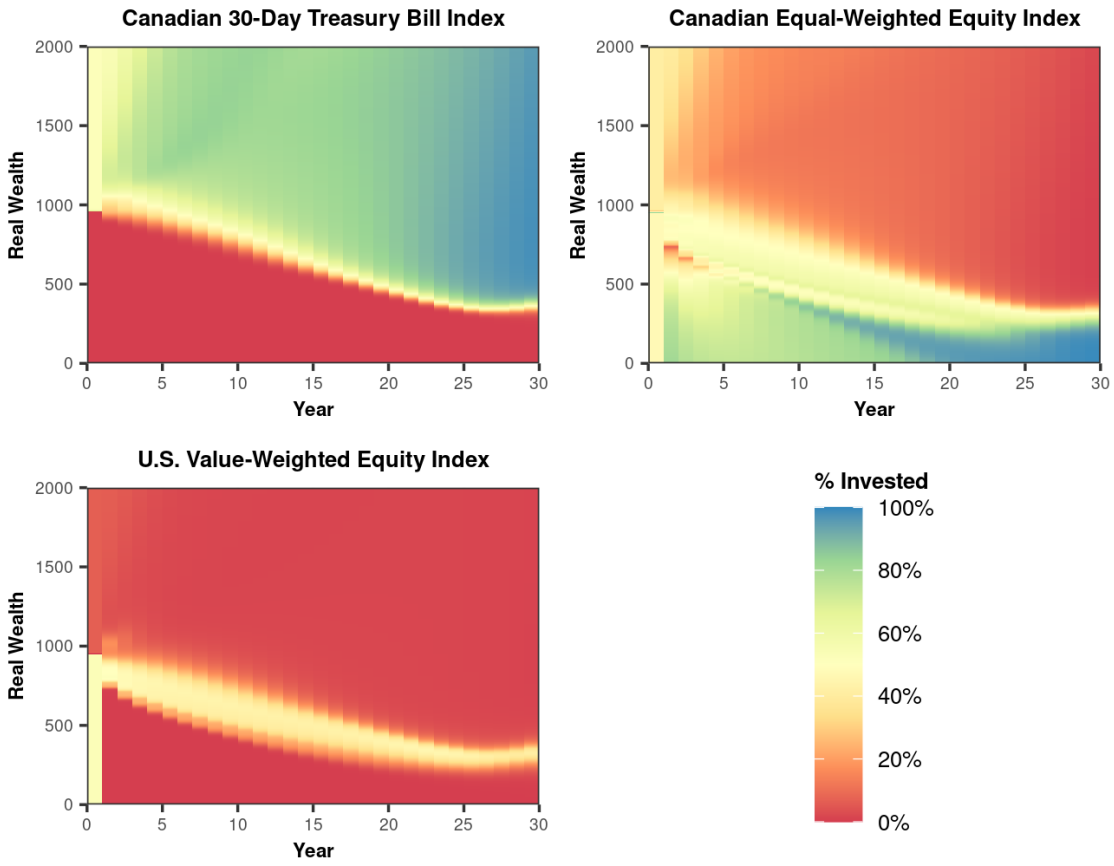


FIGURE 11: *Heatmaps of the investment allocation as a function of time and post-withdrawal wealth for the Canadian 30-day Treasury bill index (upper left), the Canadian equal-weighted equity index (upper right), and the U.S. value-weighted equity index (lower left) for the multi-asset tontine with* $\kappa = 0.03$.

# 7 Summary and Conclusion

Retirees in DC plans are confronted with two important decisions: how much to withdraw from their portfolios, and how to allocate the investment of their portfolios across various investments. The NN approach developed here can potentially be useful in optimizing these decisions. Compared to previous dynamic programming approaches, it is far easier to extend to scenarios with multiple investment options as well as alternative assumptions about mortality risk. It also offers the flexibility of being able to determine optimal policies using any type of parametric model for asset returns, or simply from available historical return data.

Examples given above have shown that results using our approach correspond quite closely with those produced using dynamic programming in a setting with a small number of investment assets based on a specified parametric model for returns. Most of our work here, however, uses boostrapped historical return data rather than a parametric model. Results provided here show that the expected block size used in the bootstrapping approach has relatively small impact over the range of greatest interest (i.e. high average annual withdrawals and low but positive expected shortfall). Similar results were seen with random fluctuations in tontine group gains and, to a lesser extent, with unanticipated systematic mortality improvement. Adding more investment assets lead to improved outcomes in the region of primary interest (i.e. even higher withdrawals with low positive expected shortfall), and, of less importance, much higher expected shortfall for cases with low withdrawals.

A possible topic for future research is investigating the use of more international investment return data.

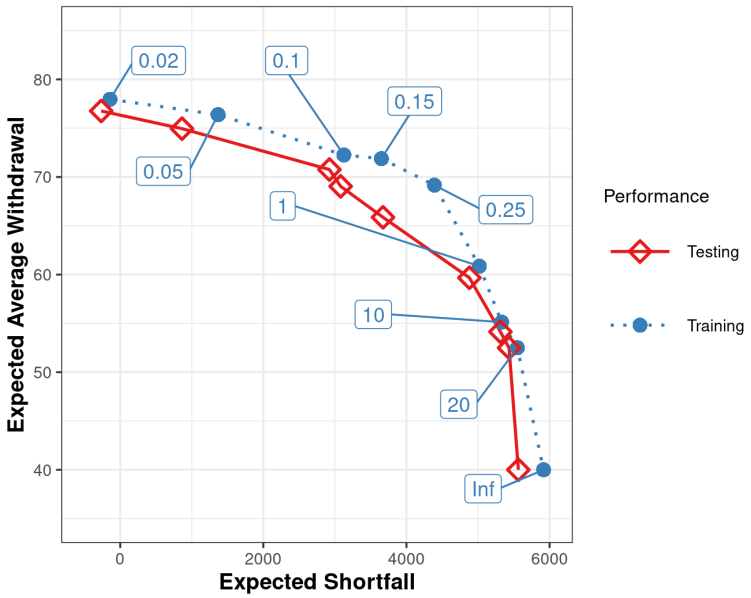


FIGURE 12: *Comparison of training and testing performance for the multi-asset tontine. The training data contains* $2.56 \times 10^5$ *paths of bootstrapped return data from 1950:11 to 1985:12, and the testing data contains* $2.56 \times 10^5$ *paths of bootstrapped return data from 1986:1 to 2022:12. In each case the expected block size is 24 months. Values of* $\kappa$ *are shown in boxes attached to the training performance at various points.*

Enhanced diversification opportunities may allow for a somewhat higher upper limit for annual withdrawals. While in this work we have used the basic reward measure of expected total withdrawals and the tail risk measure of expected shortfall (formulated as an induced time-consistent strategy), the NN approach can easily be adapted to alternative measures of risk and reward. It may be worth exploring the implications of different objective functions. Finally, we have only considered very simple *ad hoc* scenarios for mortality risk. More sophisticated approaches have been developed in the actuarial literature (see, e.g. Dahl, 2004; Blackburn and Sherris, 2013; Kessy et al., 2022), and it may well be worth investigating their impact in the context of individual tontine retirement accounts.

**Acknowledgements:** The Natural Sciences and Engineering Research Counicl of Canada (NSERC) supported Forsyth's work through grant RGPIN-2017-03760. The Canadian Securities Institute Research Foundation supported Vetzal's work through a grant awarded in 2021-2022.

**Declaration:** The authors have no conflicts of interest to report.

**Data Availability Statement:** Data used in this study are not available due to third party licencing restrictions. Computer code is not available due to University of Waterloo intellectual property policy.

# References

Bengen, W. (1994). Determining withdrawal rates using historical data. *Journal of Financial Planning 7*, 171–180.

Bernhardt, T. and C. Donnelly (2019). Modern tontine with bequest: Innovation in pooled annuity products. *Insurance: Mathematics and Economics 86*, 168–188.

Blackburn, C. and M. Sherris (2013). Consistent dynamic affine mortality models for longevity risk applications. *Insurance: Mathematics and Economics 53*, 54–73.

Bräughtigam, M., M. Guillén, and J. P. Nielsen (2017). Facing up to longevity with old actuarial methods: A comparison of pooled funds and income tontines. *Geneva Papers on Risk and Insurance 42*, 406–422.

Buehler, H., L. Gonon, J. Teichmann, and B. Wood (2019). Deep hedging. *Quantitative Finance 19*(8), 1271–1291.

Chen, A., M. Guillen, and M. Rach (2021). Fees in tontines. *Insurance: Mathematics and Economics 100*, 89–106.

Chen, A. and M. Rach (2023). The tontine puzzle. SSRN 4106903.

Chen, M., M. Shirazi, P. A. Forsyth, and Y. Li (2023). Machine learning and Hamilton-Jacobi-Bellman equation for optimal decumulation: A comparison study. Working paper, Cheriton School of Computer Science, University of Waterloo.

Dahl, M. (2004). Stochastic mortality in life insurance: Market reserves and mortality-linked insurance contracts. *Insurance: Mathematics and Economics 35*, 113–136.

Forsyth, P. A. (2022). A stochastic control approach to defined contribution plan decumulation: "the nastiest, hardest problem in finance". *North American Actuarial Journal 26*(2), 227–252.

Forsyth, P. A., K. R. Vetzal, and G. Westmacott (2024). Optimal performance of a tontine overlay subject to withdrawal constraints. *ASTIN Bulletin 54*, 94–128.

Fullmer, R. K. (2019). Tontines: A practitioner's guide to mortality-pooled investments. CFA Institute Research Foundation Brief.

Fullmer, R. K. and M. J. Sabin (2019). Individual tontine accounts. *Journal of Accounting and Finance 19*(8), 31–61.

Gemmo, I., R. Rogalla, and J.-H. Weinart (2020). Optimal portfolio choice with tontines under systematic longevity risk. *Annals of Actuarial Science 14*, 302–315.

Han, J. and E. Weinan (2016). Deep learning approximation for stochastic control problems. Working paper, https://doi.org/10.48550/arXiv.1611.07422.

Hu, R. and M. Laurière (2023). Recent developments in machine learning methods for stochastic control and games. Working paper, https://doi.org/10.48550/arXiv.2303.10257.

Kessy, S. R., M. Sherris, A. M. Villegas, and J. Ziveyi (2022). Mortality forecasting using stacked regression ensembles. *Scandinavian Actuarial Journal 2022*(7), 591–626.

Kingma, D. P. and J. Ba (2017). Adam: A method for stochastic optimization. arXiv:1412.6908v9.

Li, Y. and P. A. Forsyth (2019). A data driven neural network approah to optimal asset allocation for target based defined contribution pension plans. *Insurance: Mathematics and Economics 86*, 189–204.

MacDonald, B.-J., B. Jones, R. J. Morrison, R. L. Brown, and M. Hardy (2013). Research and reality: A literature review on drawing down retirement financial savings. *North American Actuarial Journal 17*, 181–215.

Milevsky, M. A. (2022). *How to Build A Modern Tontine: Algorithms, Scripts and Tips*. Springer.

Milevsky, M. A. and T. S. Salisbury (2015). Optimal retirement income tontines. *Insurance: Mathematics and Economics 64*, 91–105.

Milevsky, M. A. and T. S. Salisbury (2016). Equitable retirement income tontines: Mixing cohorts without discriminating. *ASTIN Bulletin 46*, 571–604.

Milevsky, M. A., T. S. Salisbury, G. Gonzalez, and H. Jankowski (2018). Annuities versus tontines in the 21st century: A Canadian case study. Society of Actuaries.

Paszke, A., S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc.

Patton, A., D. Politis, and H. White (2009). Correction to: Automatic block-length selection for the dependent bootstrap. *Econometric Reviews 28*, 372–375.

Politis, D. and J. Romano (1994). The stationary bootstrap. *Journal of the American Statistical Association 89*, 1303–1313.

Politis, D. and H. White (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews 23*, 53–70.

Powell, W. B. (2023). A universal framework for sequential decision problems. *OR/MS Today 50*(1). https://tinyurl.com/PowellORMSFeature/.

Rockafellar, R. T. and S. Uryasev (2000). Optimization of conditional-value-at-risk. *Journal of Risk 2*, 21–42.

Sabin, M. J. and J. B. Forman (2016). The analytics of a single period tontine. Working paper, SSRN 2874160.

Shirazi, M. (2024). Optimal decumulation for retirees using tontines: A dynamic neural network approach. Master's Thesis, University of Waterloo.

Spierdijk, L. and Z. Umar (2015). Stocks, bonds, T-bills, and inflation hedging: From great moderation to great recession. *Journal of Economics and Business 79*, 1–37.

U.S. Bureau of Labor Statistics (2024). Bls data finder 1.1. https://data.bls.gov/dataQuery/find?fq=survey:%5Bnb%5D&s=popularity:D.

van Staden, P. M., P. A. Forsyth, and Y. Li (2023). Beating a benchmark: Dynamic programming may not be the right numerical approach. *SIAM Journal on Financial Mathematics 14*(2), 407–451.

van Staden, P. M., P. A. Forsyth, and Y. Li (2025). A global-in-time neural network approach to dynamic portfolio optimization. *Applied Mathematical Finance*, forthcoming.

Vetzal, K. R., G. Westmacott, and P. A. Forsyth (2023). The modern tontine: Potential and practice. Research report, University of Waterloo and Richardson Wealth, available at https://web.richardsonwealth.com/documents/1833563/1833619/paper_VWF.pdf.

Weinert, J.-H. and H. Gründl (2021). The modern tontine: An innovative instrument for longevity risk management in an aging society. *European Actuarial Journal 11*, 49–86.

Winter, P. and F. Planchet (2022). Modern tontines as a pension solution: A practical overview. *European Actuarial Journal 12*, 3–32.

Yoon, J., D. Jarrett, and M. van der Schaar (2019). Time-series generative adversial networks. *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*, https://proceedings.neurips.cc/paper/2019/file/c9efe5f26cd17ba6216bbe2a7d26d490-Paper.pdf.

| | |
|---|---|
| Hidden layers per network | 2 |
| Nodes per hidden layer | 10 for cases with two assets,<br>14 otherwise |
| Nodes have biases | True |
| No. of iterations (#itn) | 50,000 |
| Stochastic gradient descent mini-batch size | 1,000 |
| No. of training paths | $2.56 \times 10^5$ |
| Optimizer | Adam |
| Initial Adam learning rate for $\theta_q$ and $\theta_p$ | 0.05 |
| Initial Adam learning rate for $W'$ | 0.04 |
| Adam learning rate decay schedule | $[0.70 \times \#\text{itn}, 0.97 \times \#\text{itn}]$, $\gamma = 0.20$ |
| Adam $\beta_1$ | 0.9 |
| Adam $\beta_2$ | 0.998 |
| Adam weight decay ($\ell_2$ penalty) | 0.0001 |

TABLE 3: *Details of NN framework hyper-parameters.*

# Appendix A Computational Details

This appendix provides details of the NN methods used to produce the results reported in this paper. For the cases with two assets, the withdrawal and allocation networks were each implemented with two hidden layers having 10 nodes each with biases. For cases with more assets, the number of nodes per layer was increased to 14. The NNs were trained using stochastic gradient descent with the Adaptive Momentum (Adam) optimization algorithm (Kingma and Ba, 2017). The parameters for the allocation and withdrawal NNs and the auxiliary parameter $W'$ were trained with different initial learning rates. The same learning rate schedule and decay parameters were employed. Weight decay ($\ell_2$ penalty) was used to enhance training stability. These and other training hyper-parameters are summarized in Table 3.

As training proceeds, the minimum loss function value is tracked. The model selected is the one that had the minimum loss value based on the entire training dataset by the end of the specified number of training epochs. The optimization problem is more difficult for higher values of the risk-aversion parameter $\kappa$, since this places more weight on the ES component of the objective function that is only affected by sample paths in the lowest 5% of the terminal wealth distribution. Due to the sparsity of such paths, we use transfer learning to mitigate these issues. We start by training the model for the lowest value of $\kappa$ under consideration from a random initialization, and then initiate the training for higher values of $\kappa$ based on the optimized model for the preceding value of $\kappa$.

Learning is also improved by standardizing the input wealth based on the first two moments of wealth samples computed using a reference strategy. In particular, we consider a constant withdrawal of 40 each year and a constant proportion investment strategy (equally-weighted across the included assets) and calculate results for $2.56 \times 10^5$ simulated return paths. Denote the associated wealth vector at time $t$ by $W_t^b$, and its mean and standard deviation across the paths by $\bar{W}_t^b$ and $\sigma(W_t^b)$ respectively. The feature input $W_t$ to the allocation and withdrawal NNs is then normalized to $\tilde{W}_t$ according to

$$\tilde{W}_t = \frac{W_t - W_t^b}{\sigma(W_t^b)}.$$