Beating the Four-Percent Rule: A Data Driven Neural Network Approach to DC Pension Decumulation for Canadian Retirees[a]

Peter A.I. Forsyth Jr. [b] Dominik Luszczynski [c] Yiqing Huang [d] Yuying Li [e]

June 30, 2026

Abstract

We develop a data-driven neural network framework for generating optimal decumulation strategies for Canadian defined-contribution (DC) pension retirees. The optimal decumulation strategy solves a stochastic optimal control problem based on an objective function consistent with the revealed preferences of retail investors.

We employ a global-in-time neural network architecture within a data-driven framework. We show mathematically that general linear constraints on the control variables can be incorporated by reformulating the constrained optimization problem as an unconstrained one, via suitably designed constraint-feasible activation functions for a feedforward neural network. We establish mathematically that neural networks with these constraint-feasible output activation functions always yield feasible controls and can approximate any feasible control to arbitrary accuracy.

We compute optimal decumulation and allocation strategies using block bootstrap resampling of CPI-adjusted CRSP and CFMRC data during the period 1950:1 to 2024:12, in terms of Canadian dollars. Using these data, we assess and compare optimal decumulation strategies using up to six US and Canadian equity and bond indexes. Our results indicate that allowing variable withdrawals (within bound constraints) has a significantly larger effect on the optimal strategy compared with allowing variable asset allocation. The Canadian equal-weighted equity index and the U.S. capitalization-weighted equity index emerge as the most influential assets in optimal decumulation strategies.

Keywords: pension decumulation; stochastic control; neural networks; asset allocation; equity exposure constraints; data-driven methods

JEL Classification. G11 (Portfolio Choice); G17 (Financial Forecasting and Simulation); C61 (Optimization Techniques); C63 (Computational Techniques); D91 (Intertemporal Choice)

AMS Classification. 93E20 (Optimal stochastic control); 68T07 (Artificial neural networks and deep learning); 62P05(Applications of statistics to actuarial sciences and financial mathematics); 65C05 (Monte Carlo methods / resampling methods)

1 Introduction

Globally, retirees with Defined Contribution (DC) pension plans, which are increasingly replacing traditional Defined Benefit (DB) plans, face the challenge of making complex withdrawal and investment allocation decisions over extended time horizons. These decisions are crucial for managing longevity risk and achieving a desirable retirement lifestyle.

[a] available at `https://cmsweb.utsc.utoronto.ca/huang/research/multi_asset_nn.pdf`

[b] David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, `pa2forsyt@uwaterloo.ca`

[c] Department of Computer and Mathematical Sciences, University of Toronto Scarborough, Toronto, ON, Canada M1C 1A4, `dominik.luszczynski@mail.utoronto.ca`

[d] Department of Computer and Mathematical Sciences, University of Toronto Scarborough, Toronto, ON, Canada M1C 1A4, `iy.huang@utoronto.ca`

[e] David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, `yuying@uwaterloo.ca`

This global trend is evident in Canada, where DC pension plans are increasingly replacing DB plans (Thinking Ahead Institute, 2024; OSFI, 2022). In a DC plan, employers contribute to retirement accounts whose management becomes the responsibility of individual retirees. Consequently, many retirees must navigate the complex challenge of optimally managing their retirement savings throughout their post-working years.

A decumulation strategy consists of both withdrawal and investment allocation decisions made over a long time horizon, e.g., 30 years. Bengen (1994) is the seminal work on data-driven analysis of the DC-plan decumulation problem. Bengen considered a retiree with access to two assets: an equity index and a bond index, and studied static strategies in which the retirement portfolio is rebalanced to maintain a fixed proportion of equity, with the same amount of cash withdrawn each period. This work led to the popular and ubiquitous four-percent rule, which suggests withdrawing four percent of the initial capital each year, adjusted for inflation, over a thirty year retirement. In this paper, we consider the same scenario posed in Bengen (1994). Spending rules (Anarkulova et al., 2025), including the ubiquitous four-percent rule (Bengen, 1994), have long been the standard recommendation for retirees. However, the effectiveness of these simple rules depends heavily on the data used to evaluate them. We believe that a principled stochastic optimal control approach, with an objective function and constraints that align with retirees' goals, should yield better strategies. Decumulation can be formulated as a discrete-time optimal stochastic control problem, with annual asset allocations at each rebalancing date and annual withdrawal amounts as controls. Formulating decumulation as an optimal stochastic control problem makes the retirement goals and implementation constraints explicit.

Naturally, retirees have two goals: first, to maximize total withdrawals over the retirement planning period, and second, to minimize the risk of running out of savings before the end of the retirement period. To account for a retiree's idiosyncratic risk tolerance or their views on the future performance of available assets, constraints can be imposed on the permitted asset allocations. For example, minimum and maximum annual withdrawal amounts can be imposed, in addition to constraints on allocations such as bounds on total equity and leverage exposure.

Chen et al. (2025) apply neural network approximation methods to determine optimal decumulation strategies for DC pension plans, focusing on U.S. investors with access to an equity index and a treasury index. The authors examine a retiree operating under a no-leverage, no-shorting investment constraint. To validate their methodology, they benchmark the neural network results against solutions derived from Hamilton-Jacobi-Bellman partial differential equation (HJB PDE) solvers within a parametric market framework. They also apply the neural network approach to compute optimal investment strategies directly from bootstrap-resampled historical market data.

In Forsyth and Labahn (2026), assuming two investable assets, a stock index and a bond index, the discrete optimal decumulation control problem is solved by alternately solving a linear partial-integro differential equation (PIDE) at each rebalancing date and performing an optimization step. That work allows the use of leverage (i.e., borrowing to invest in stocks), subject to a lower bound on bond holdings. When evaluating the performance of the optimal strategy using block bootstrap resampling of historical data, an approach also used in Anarkulova et al. (2025) to assess standard spending rules, the authors find that restricting the portfolio's equity fraction to a maximum of 50% does not noticeably reduce portfolio efficiency. This is an attractive feature of the optimal decumulation strategy, since retirees are generally risk-averse. The most significant effect of the 50% equity restriction, relative to allowing a maximum stock fraction of up to 130% (i.e., leverage), is that the time for median withdrawals to reach their maximum is delayed by one year.

In this paper, we examine the optimal decumulation strategy for DC pension plans from a Canadian retiree's perspective, considering both the no-shorting, no-leverage case and the case of bounded leverage exposure. It is commonplace to assume that U.S. markets are superior to other markets. The MSCI world index (https://www.msci.com/documents/10199/255599/msci-world-index.pdf) is a capitalization weighted developed market benchmark index. This index has about 3.4% Canadian and 72.5% US equities by cap weight. It is commonly suggested that a diversified Canadian investor should have a relatively small weight in Canadian stocks. However, Figure 1.1 shows that there have been two decades, out of the last 75 years, where a Canadian investor was poorly served by US investments, which suggests that this may be a misconception. It indicates that, over the long term, equal-weight portfolios can outperform cap-weight portfolios, although equal-weight portfolios have underperformed in recent years (Tljaard and Mare, 2021).

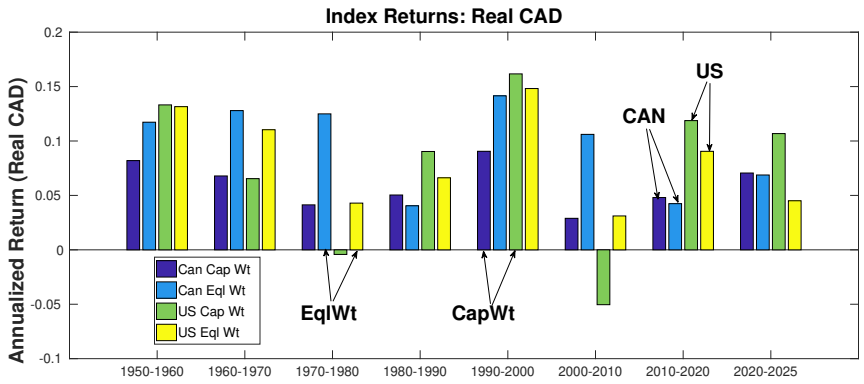

FIGURE 1.1: *U.S. and Canadian asset return comparison: Canadian data from CFMRC (`https://www.cfmrc.org/`), and CRSP (`https://www.crsp.org/`). These are total return indexes, inflation adjusted, in CAD*

The long-term superior returns generated by equal-weight portfolios are also discussed in (DeMiguel et al., 2009; Tljaard and Mare, 2021; Plyakha et al., 2021; Forsyth, 2022a). Plyakha et al. (2021) note that a large part of the enhanced return of the equal-weight strategy is due to rebalancing, not merely the small-cap effect.

Farago and Hjalmarsson (2023) point out that, because the return distribution of a buy-and-hold portfolio is highly skewed (owing to the approximate log-normality of cumulative returns), small, equally weighted portfolios can outperform the buy-and-hold (i.e., capitalization-weighted) portfolio with high probability. The most surprising result concerns small subportfolios. Using CRSP data for 1979–2016 to generate blocks of 30-year returns, a portfolio that picks 50 stocks at random and rebalances to equal weighting each month beats the market (the capitalization-weighted CRSP index) 96% of the time over 30 years.

Farago and Hjalmarsson (2023) emphasize that this non-intuitive result is largely driven by the equal-weight rebalancing effect, and use simulations based on simple assumptions to explain it. Motivated by these findings, we consider an investment universe comprising six assets: U.S. capitalization-weighted and equal-weighted indexes (in CAD), Canadian capitalization-weighted and equal-weighted indexes, and CAD-denominated short- and long-term bonds (all inflation-adjusted). Another central question for our investigation is whether retirees can achieve comparable risk-return tradeoffs while reducing equity exposure; specifically, we examine the impact of varying equity exposure on portfolio performance and sustainability.

Given the computational infeasibility of traditional PDE-based approaches with six investment assets, we employ a data-driven, global-in-time neural network methodology to derive optimal stochastic control policies. Because the stochastic optimal control problem incorporates stochastic linear constraints, solving it via stochastic gradient descent requires specially designed activation functions that transform the constrained optimization into an equivalent unconstrained formulation. We design the requisite custom activation functions for the decumulation problem and develop a general methodology for handling linear constraints based on the classical Minkowski–Weyl theorem of convex optimization duality.

We employ block bootstrap resampling of historical CRSP market data to model the stochasticity of traded asset returns, and our data-driven neural network approach then computes optimal decumulation strategies directly from these block bootstrap samples.

The main contributions of our work are as follows:

- Modeling framework: We develop a stochastic optimal control formulation for computing optimal decumulation strategies tailored to Canadian DC pension decumulation plans, and propose a data-driven neural network approach suitable for high-dimensional problems.
- Methodological innovation: We show how control constraints can be explicitly satisfied by designing constraint-feasible output activation functions for feed forward neural networks, and present a general framework for constructing such functions for arbitrary linear constraints, applied here to allocation constraints. For the withdrawal constraint, we design a custom constraint-feasible activation function

based on the specific stochastic linear constraints governing withdrawals. We prove mathematically that feed forward neural networks with these constraint-feasible output activations produce explicitly feasible controls and can achieve arbitrarily accurate approximations of any feasible control.

- Empirical analysis and key findings: Using historical market data spanning 1950:1–2024:12 from the Center for Research in Security Prices (CRSP) and the Canadian Financial Markets Research Center (CFMRC), we compute optimal constrained decumulation strategies across a six-asset portfolio. Our analysis reveals three key insights:

  (a) expanding the asset universe enables retirees to achieve superior pension decumulation performance;

  (b) retirees can effectively limit equity market exposure with only modest performance costs; and

  (c) the Canadian equal-weighted and U.S. capitalization-weighted equity indices emerge as the most significant portfolio components.

Next we describe the optimal stochastic control formulation for the pension decumulation in §2. The data driven NN approach to compute the optimal decumulation strategy in presented in §3. In §4, we propose a constraint-feasible activation function when the permissible control value set is described by linear constraints, e.g. for allocation controls and a constraint-feasible activation function for the withdrawal amount prescribed by adaptive stochastic bounds. We mathematically establish that feedforward NN with these constraint-feasible output activation functions can approximate any feasible control with arbitrary accuracy. The empirical performance and analysis of the computed optimal decumulation strategy is reported in §5. Concluding remarks are given in §6.

2 Problem Formulation

We adopt the decumulation setup in the standard spending rule literature (Anarkulova et al., 2025; Bengen, 1994) and consider a typical 65-year-old retiree at the onset of retirement with a 30-year planning horizon. We mathematically formulate the decumulation problem as an optimal stochastic control problem with a decumulation strategy as the control. Our primary objective is to identify optimal stochastic pension decumulation strategies that exceed the performance of existing heuristic spending rules, particularly by incorporating a broader range of asset classes, e.g., including equity indices and bonds from both U.S. and Canadian markets.

We consider the primary objectives for retirement planning for a typical 65-year-old retiree at the onset of retirement $t_0 = 0$. According to Statistics Canada,[1] the probability that a 65-year-old Canadian female attains age 95 is approximately 19%. Consequently, a 30-year planning horizon represents a prudent choice for our framework.

In our approach, we do not apply mortality weighting to future cash flows, as is standard practice in annuity pricing. Mortality weighting may be less meaningful for individual investors, since it reflects population-level averages rather than the specific circumstances of an individual retiree. This is also consistent with the Bengen scenario.

Let T denote the retirement horizon, e.g. $T = 30$. The main goal for the retiree is to obtain secure annual cashflow. We achieve this by imposing a lower bound on the withdrawal control and simultaneously minimizing the risk of running out of cash. We also cap the withdrawal to facilitate continual investment across time.

Additionally, we assume that the retiree withdraws funds and rebalance the portfolio annually. Since a minimum cash withdrawal is required each year, it is possible to exhaust the retirement fund.[2] We suppose that the retiree possesses real estate, which can generate a reverse mortgage if required, i.e., real estate can be used as collateral (Pfeiffer et al., 2013). This presumes that the retiree categorizes real estate separately from financial assets, which is consistent with behavioral finance (Shefrin and Thaler, 1988). Real estate in this case has a dual role: if market returns are exceptionally strong or the retiree passes away earlier

[1] `https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1310011401`.

[2] Ignoring the trivial and unlikely case where investing in riskless assets can fund 30-year minimum withdrawals.

than expected, this property can serve as a bequest. If market returns are poor or in the case of extreme longevity, real estate can be used to fund expenses if the retirement account is exhausted. Consequently, in our optimal control problem formulation, withdrawals continue even when wealth becomes negative. In this case, debt accumulates at the borrowing rate. This implementation is reasonable, considering that retirees typically possess other assets, such as real estate, that can serve as a hedge of last resort.

Assume that we have M discrete withdrawal/rebalancing times in [0,T], i.e.,

$$\mathcal{T} = \{t_i \mid i = 0, \ldots, M, \quad 0 = t_0 < t_1 < \ldots < t_M\}$$

For notational simplicity, we assume $t_i - t_{i-1} = \Delta t = T/M$.

At each rebalancing time t_i, $i = 0, 1, \ldots, M-1$, the retiree decides both the amount of cash $\mathfrak{q}_i$ to withdraw from the investment portfolio and the new asset percentage allocation $\mathfrak{p}_i$. At $t_M = T$, the portfolio is liquidated and no cash flow occurs, which is represented subsequently by setting $\mathfrak{q}_M = 0$.

We assume that the retiree has the initial wealth $W(t_0)$ and the retiree wants to invest in N assets.

We do not address tax considerations in this work. It is worth noting that defined contribution savings are held in tax-advantaged accounts, where rebalancing does not trigger tax liabilities. Furthermore, we also disregard transaction costs, as we anticipate their impact to be negligible with annual rebalancing. Computationally, incorporating transaction costs is straightforward (van Staden et al., 2018).

2.1 Control and wealth

Let $\boldsymbol{X}(t)$ represent the state or information a retiree incorporates into their asset allocation and withdrawal decision at time $t \in \mathcal{T}$. For each $i = 0, \ldots, M-1$, let $\mathfrak{p}_i(\cdot) = \mathfrak{p}_i\left(\boldsymbol{X}(t_i)\right) \in \mathbb{R}^N$ denote the vector of proportions of the retiree's wealth $W\left(t_i\right)$ at time t_i invested across a universe of N assets. The withdrawal control $\mathfrak{q}_i(\cdot)$ represents the amount withdrawn at time t_i.

Given the set of rebalancing times, the retiree's strategy is defined as

$$\mathcal{P} = \left\{\left(\mathfrak{q}(\boldsymbol{X}(t_i)), \mathfrak{p}_i\left(\boldsymbol{X}(t_i)\right)\right),\, i = 0, \ldots, M-1\right\}. \tag{2.1}$$

We also define $\mathfrak{q}_M = 0$ and $\mathfrak{p}_M = 0$ (no decisions are needed at T).

Let $t_i^+ = \lim_{\epsilon \to 0^+}(t_i + \epsilon)$, $t_i^- = \lim_{\epsilon \to 0^+}(t_i - \epsilon)$. $W(t_i^-)$ thus denotes the total wealth value the instant before withdrawals and rebalancing at t_i.

Then the wealth $W_{\mathcal{P}}(t_i^+)$ immediately after the withdrawal at t_i, under the control $\mathcal{P}$, is given by

$$W_{\mathcal{P}}(t_i^+) = W_{\mathcal{P}}(t_i^-) - \mathfrak{q}_i, \;\; i = 0,1 \ldots, M-1. \tag{2.2}$$

Without loss of generality, we assume here that withdrawal occurs prior to allocation. Consequently, the optimal control $\mathfrak{q}_i(\cdot)$ actually depends on the wealth immediately before withdrawal, while $\mathfrak{p}_i(\cdot)$ is a function of wealth immediately after withdrawal.

Let $\boldsymbol{R}(t_i) \in \mathbb{R}^N$ denote the random log asset returns in [t_i,t_{i+1}]. Then the wealth associated with a given control decision strategy is as follows: for $i = 0,1, \ldots, M-1$,

$$\begin{aligned} W_{\mathcal{P}}(t_i^+) &= W_{\mathcal{P}}(t_i^-) - \mathfrak{q}_i(\cdot) \\ \mathfrak{p}_i(\cdot) &= \mathfrak{p}_i(W(t_i^+)) \\ W_{\mathcal{P}}(t_{i+1}^-) &= \mathfrak{p}_i(\cdot)^T e^{\boldsymbol{R}(t_i)} W_{\mathcal{P}}(t_i^+) \end{aligned} \tag{2.3}$$

2.2 Constraints and feasibility

At time t_i, the control $\mathcal{P}$ of the decumulation problem is the pair $(\mathfrak{q}_i(W_{\mathcal{P}}(t_i^-)), \mathfrak{p}_i(W_{\mathcal{P}}(t_i^+)))$. We describe the constraints on the control by prescribing admissible control values: let $\mathcal{Z}(\cdot)$ represent the set of admissible values for the controls $(\mathfrak{q}_i(\cdot), \mathfrak{p}_i(\cdot))$. In this paper, we consider two types of allocation constraints: (i) no-shorting and no-leverage, and (ii) bounded leverage constraints (assuming solvency). For withdrawals, we impose maximum and minimum values. We further impose the constraint that in the event of insolvency, $W(t_i^+) < 0$, arising from withdrawals or leverage, allocation trading ceases and debt (negative wealth) accrues at the appropriate borrowing rate of return (i.e., a spread over the bond rate). We also assume that equity assets are liquidated at $t = t_M$.

More precisely, let W_i^+ denote the wealth immediately after the cash withdrawal, and W_i^- the total wealth immediately before withdrawal.

Assume that the N-th asset denote bond. We define

$$\mathcal{Z}_{\mathfrak{q}}(W_i^-,t_i) = \begin{cases} [\mathfrak{q}_{\min},\mathfrak{q}_{\max}] & \text{when } W_i^- \geq \mathfrak{q}_{\max},\ i=0,1,\ldots,M-1 \\ [\mathfrak{q}_{\min},\max(\mathfrak{q}_{\min},W_i^-)] & \text{when } W_i^- < \mathfrak{q}_{\max}, i=0,1,\ldots,M-1 \end{cases}, \tag{2.4}$$

$$\mathcal{Z}_{\mathfrak{p}}(W_i^+,t_i) = \begin{cases} \mathcal{Z}_{\mathfrak{p}}^+(W_i^+,t_i) & \text{when } W_i^+ > 0, i=0,1,\ldots,M-1 \\ \{[0;\ldots,0;1]\} & \text{when } W_i^+ \leq 0, i=0,1,\ldots,M-1 \end{cases}, \tag{2.5}$$

where $\mathcal{Z}_{\mathfrak{p}}^+(W_i^+,t_i)$ denotes the set of values of the feasible asset allocations when wealth is positive. Consequently, (2.5) implies that, when $W_i^+ \leq 0$, 100% is allocated to the bond. The constraints (2.4) state that if there is sufficient fund value, i.e., wealth is greater than $\mathfrak{q}_{\max}$, withdrawal can be any amount in $[\mathfrak{q}_{\min},\mathfrak{q}_{\max}]$. But if the current wealth is below $\mathfrak{q}_{\max}$, then the withdrawal amount is capped by the maximum of the current wealth and $\mathfrak{q}_{\min}$. In the event of insolvency, withdrawal must be $\mathfrak{q}_{\min}$.

The set of admissible values for $(\mathfrak{q}_i,\mathfrak{p}_i), t_i \in \mathcal{T}$, can then be written as

$$\mathcal{Z}(W_i^-,W_i^+,t_i) = \mathcal{Z}_{\mathfrak{q}}(W_i^-,t_i) \times \mathcal{Z}_{\mathfrak{p}}(W_i^+,t_i)\ . \tag{2.6}$$

For implementation purposes, we have written equation (2.6) in terms of wealth before and after cash withdrawal. However, we remind the reader that since $W_i^+ = W_i^- - \mathfrak{q}_i$, the controls are formally a function of the state $X(t_i^-)$ before the control is applied.

The admissible control set $\mathcal{A}$ can then be written as

$$\mathcal{A} = \left\{ (\mathfrak{q}_i,\mathfrak{p}_i)_{0\leq i\leq M-1} : (\mathfrak{p}_i,\mathfrak{q}_i) \in \mathcal{Z}(W_i^-,W_i^+,t_i) \right\}. \tag{2.7}$$

An admissible control $\mathcal{P} \in \mathcal{A}$ can be written as

$$\mathcal{P} = \{(\mathfrak{q}_i(\cdot),\mathfrak{p}_i(\cdot))\ :\ i=0,\ldots,M-1\} \in \mathcal{A}. \tag{2.8}$$

Using aforementioned specifications, we determine the optimal decumulation strategy as a solution to the following optimal stochastic control problem:

$$\min_{\mathcal{P}\in\mathcal{A}}\ E^{t_0,w_0}\left(F(W_{\mathcal{P}}(T),\{\mathcal{P}(t_i)\}_{i=0}^{M-1}) \right)$$

where the objective is to balance reward and risk. We break down the components of this problem:

- Sequence of Decisions $\{\mathcal{P}(t_i)\}_{i=0}^{M-1}$: This refers to the control decisions made at discrete time points $t_0,t_1,\ldots,t_{M-1}$ within the investment horizon $[t_0,T]$. These controls inherently contribute to the reward and risk profile (e.g., aggressive asset allocation leads to higher risk).

- Feasible Control $\mathcal{P}$: The control being optimized is $\mathcal{P}$, which represents a "feasible control." This $\mathcal{P}$ is chosen from the feasible set $\mathcal{A}$, which defines all permissible control strategies, i.e., a sequence of withdrawal and investment decisions over time.

- Conditional Expectation $E^{t_0,w_0}(\cdot)$: The $E^{t_0,w_0}(\cdot)$ operator denotes the expectation taken conditional on the initial state of the control. Specifically, it means the expectation is calculated given that the process starts at time $t_0 = 0$ with an initial wealth of $W(t_0) = w_0$. This implies that the risk and reward are evaluated considering the stochastic nature of future outcomes, starting from a known initial point.

- Risk-Reward Tradeoff Function $F(\cdot)$: This function, $F(W_{\mathcal{P}}(T),\{\mathcal{P}(t_i)\}_{i=0}^{M-1})$, quantifies the balance between the accumulated reward and the incurred risk for a given control strategy $\mathcal{P}$. In this paper, we assume that this tradeoff explicitly depends on two main components: Terminal wealth $W_{\mathcal{P}}(T)$ resulting from applying the control strategy $\mathcal{P}$. The reward depends on the withdrawals of the control sequence while risk depends on the terminal wealth. This assumption can also be generalized to capture path dependent risk or nonlinear reward.

- Objective Function: The primary goal is to minimize the conditional expectation of a risk-reward tradeoff function, denoted as $F(\cdot)$. This means we are looking for a control strategy $\mathcal{P}$ that yields the lowest expected value of this trade-off.

In summary: the problem aims to find a control strategy $\mathcal{P}$ that, when applied from an initial wealth w_0 at $t_0 = 0$, minimizes the expected value of a function F that weighs a risk measure based on the terminal wealth $W_{\mathcal{P}}(T)$ against a reward determined by the control strategy itself $\{\mathcal{P}(t_i)\}_{i=0}^{M-1}$ over the investment period. The specific form of F defines how the reward is measured (e.g., total withdrawals) and how the risk is quantified (e.g., variance of wealth, downside deviation, cost of control implementation). Next, we precisely specify the reward and risk measures used in $F(\cdot)$ in this work.

2.3 Expected Total Withdrawals (EW)

The primary goal in the DC-plan decumulation problem is to ensure regular and sufficient cashflows during the retirement period. This is achieved by maximizing expected total withdrawals while achieving a guaranteed minimum amount in each period.

Define $X_0^+ = X(t_0^+), X_0^- = X(t_0^-)$. Let $E^{X_0^-,t_0^-}(\cdot)$ denote the expectation conditional on the initial state. More precisely, below we define the expected total withdrawals $\mathrm{EW}^{X_0^-,t_0^-}$, viewed at the initial time:

$$\mathrm{EW}^{X_0^-,t_0^-} = E^{X_0^+,t_0^+}\left[\sum_{i=0}^{M-1} \mathfrak{q}_i\right], \tag{2.9}$$

where we assume that the retiree survives for the entire decumulation period, consistent with the Bengen (1994) scenario. We consider inflation adjusted (real) withdrawals, for which we do not discount withdrawals to the present time. Effectively we are discounting future cash flows with a real discount rate of zero, which is a conservative approach. This is also consistent with (Bengen, 1994). See Forsyth (2022b) for further discussion on why this is reasonable.

It is commonplace in actuarial applications to mortality-weight cash flows. While this is clearly appropriate for annuity providers, it does not seem to be very informative for an individual retiree. Consider the perspective of a 65-year-old male with median life expectancy of about 87. The standard mortality-weighting approach would weight the minimum cash flows at 22 years after retirement by one half. However, if the retiree is planning to live rather than planning to die (Pfau, 2018), he needs the entire minimum cash flow at age 87, not half of it.

The assumption of a thirty year retirement is consistent with the Bengen (1994) scenario, which is the prototypical spending rule. Anarkulova et al. (2025) note that *...spending practices demonstrate a revealed preference for spending rules over annuitization...Obtaining reliable, quantitative evidence on the 4% rule and alternative withdrawal rates is of critical importance given their widespread use.*

2.4 Expected Shortfall (ES) Risk

In addition to maximizing the expected total withdrawal, the retiree is concerned with the risk of severely depleting their savings. While alternative measures for risk can be used, see, e.g., (Forsyth et al., 2020; Forsyth and Li, 2026), we believe that expected shortfall $\mathrm{ES}_\alpha^{X_0^-,t_0^-}$, where $\alpha > 0$ is a small given probability level, intuitively matches well with the retiree's concern of depleting savings.

Specifically, this risk can be measured by the average of the worst terminal wealth outcomes at a given probability level α, typically $\alpha \in \{.01, .05\}$.

Let $W_{\mathcal{P}}(T)$ be the random terminal wealth generated by a decumulation strategy $\mathcal{P}$. Let $g(w)$ be the probability density function of $W_{\mathcal{P}}(T)$. Let W_α^* be the Value at Risk (VAR) of $W_{\mathcal{P}}(T)$ at probability level α, i.e., the α-quantile of $W_{\mathcal{P}}(T)$. For example, if $\alpha = .05$, then 95% of the outcomes have $W_{\mathcal{P}}(T) > W_\alpha^*$. Mathematically, W_α^* satisfies

$$\int_{-\infty}^{W_\alpha^*} g(w)\, dw = \alpha, \tag{2.10}$$

i.e. $Pr[W_{\mathcal{P}}(T) > W^*_\alpha] = 1-\alpha$. If W^*_α is sufficiently large, this suggests very low risk of running out of savings.[3]

Given the probability level α, the Expected Shortfall ES_α is the average of the worst α fraction of terminal wealth, i.e.,

$$\mathrm{ES}_\alpha^{X_0^-,t_0^-} = \frac{\int_{-\infty}^{W^*_\alpha} W_T\ g(w)\ dw}{\alpha}\ , \tag{2.11}$$

Note that here we define VaR in terms of the wealth value, not loss. Hence a larger value of ES (i.e., a larger value of average worst case terminal wealth) is desired. The standard definition of Conditional Value at Risk (CVAR) corresponds to the negative of ES in (2.11).

For each given control $\mathcal{P}$, assume that $W_{\mathcal{P}}(T)$ is the corresponding terminal wealth. Under some regularity conditions, ES_α is equivalent to

$$\mathrm{ES}_\alpha^{X_0^-,t_0^-} = \sup_{W'} E^{X_0^-,t_0^-}\left[W' + \frac{1}{\alpha}\min(W_{\mathcal{P}}(T) - W', 0)\right], \tag{2.12}$$

see Rockafellar and Uryasev (2000). The admissible set for W' in equation (2.12) is the set of possible values for $W_{\mathcal{P}}(T)$, and the maximizer W^*_α is the VaR at the probability level α.

2.5 EW-ES Stochastic Optimal Control Decumulation

Since a retiree has two conflicting objectives, maximizing expected total withdrawals and expected shortfall, we use a scalarization technique to find the Pareto points for this multi-objective optimization problem.

Specifically, for any given $\kappa > 0$, we solve:

$$\begin{aligned}
&\sup_{\mathcal{P}\in\mathcal{A}}\sup_{W'}\left\{ E^{X_0^-,t_0^-}\left[\sum_{i=0}^{M-1}\mathfrak{q}_i + \kappa\left(W' + \frac{1}{\alpha}\min(W_{\mathcal{P}}(T) - W', 0)\right)\quad\right]\right\}\\
&\text{subject to}\ \begin{cases}
\text{for}\quad i = 0,\ldots,M-1\\
\qquad \mathfrak{p}_i(\cdot) = \mathfrak{p}_i(W_{\mathcal{P}}(t_i^+)), \mathfrak{q}_i(\cdot) = \mathfrak{q}_i(W_{\mathcal{P}}(t_i^+))\\
\qquad (\mathfrak{q}_i(\cdot), \mathfrak{p}_i(\cdot)) \in \mathcal{Z}(W_{\mathcal{P}}(t_i^-), W_{\mathcal{P}}(t_i^+), t_i),\\
\qquad W_{\mathcal{P}}(t_i^+) = W_{\mathcal{P}}(t_i^-) - \mathfrak{q}_i(\cdot),\\
\qquad W_{\mathcal{P}}(t_{i+1}^-) = \mathfrak{p}_i(\cdot)^T \boldsymbol{R}(t_i) W_{\mathcal{P}}(t_i^+)
\end{cases}
\end{aligned} \tag{2.13}$$

where $\boldsymbol{R}(t_i)$ is the random log asset return in $[t_i, t_{i+1}]$.

We observe that when $W_{\mathcal{P}}(t) \gg W'$ as $t \to T$, a feasible control has little effect on either the expected total withdrawal or the expected shortfall values. This leads to some ill-posedness in optimal control in the region of near-term and high wealth. Thus, we add an additional term $E^{X_0^-,t_0^-}[\epsilon W_{\mathcal{P}}(T)]$ to the objective function of (2.13) to regularize this ill-posedness. In this fortunate state for the retiree, we can break investment policy ties either by setting $\epsilon < 0$, which will force allocations in bonds, or by setting $\epsilon > 0$, which will force investments into stocks. Choosing $|\epsilon| \ll 1$ ensures that this term only has an effect if $W_t \gg W^*$ and $t \to T$. See Forsyth (2022b) for more discussion of this.

Interchanging the $\sup\sup(\cdot)$, [4] (2.13) can be equivalently written as

$$\begin{aligned}
&\sup_{W'}\sup_{\mathcal{P}\in\mathcal{A}}\left\{ E^{X_0^-,t_0^-}\left[\sum_{i=0}^{M-1}\mathfrak{q}_i + \kappa\left(W' + \frac{1}{\alpha}\min(W_{\mathcal{P}}(T) - W', 0)\right)\quad + \epsilon W_{\mathcal{P}}(T)\right]\right\}\\
&\text{subject to}\ \begin{cases}
\text{for}\quad i = 0,\ldots,M-1\\
\qquad \mathfrak{p}_i(\cdot) = \mathfrak{p}_i(W_{\mathcal{P}}(t_i^+)), \mathfrak{q}_i(\cdot) = \mathfrak{q}_i(W_{\mathcal{P}}(t_i^+))\\
\qquad (\mathfrak{q}_i(\cdot), \mathfrak{p}_i(\cdot)) \in \mathcal{Z}(W_{\mathcal{P}}(t_i^-), W_{\mathcal{P}}(t_i^+), t_i),\\
\qquad W_{\mathcal{P}}(t_i^+) = W_{\mathcal{P}}(t_i^-) - \mathfrak{q}_i(\cdot),\\
\qquad W_{\mathcal{P}}(t_{i+1}^-) = \mathfrak{p}_i(\cdot)^T e^{\boldsymbol{R}(t_i)} W_{\mathcal{P}}(t_i^+)
\end{cases}
\end{aligned} \tag{2.14}$$

[3] In practice, the negative of W^*_α is often the reported VAR.

[4] Let $F = \sup_{(a,b)\in A\times B} f(a,b)$, then $\forall \epsilon > 0$, $\exists (a^*,b^*) \in A \times B$, s.t. $f(a^*,b^*) > F - \epsilon$. Then $F \geq \sup_{a\in A}\sup_{b\in B} f(a,b) \geq \sup_{b\in b} f(a^*,b) \geq f(a^*,b^*) > F - \epsilon$. Hence, $\epsilon \to 0$ implies $\sup_{a\in A}\sup_{b\in B} f(a,b) = F$. Similarly $\sup_{b\in B}\sup_{a\in A} f(a,b) = F$.

We emphasize that the expected shortfall $\mathrm{ES}_\alpha^{X_0^-,t_0^-}$ is conditional on the given initial state (X_0^-, t_0^-). Consequently (2.13) (or equivalently (2.14)) is not time consistent. and its optimal solution is sometimes regarded as *non-implementable*, since the retiree can be inclined to deviate at future $t > 0$ from the optimal strategy computed at time $t = 0$.

Let $(W^*, \mathcal{P}_0^*)$ denote the optimal solution to (2.13). Then the optimal strategy $\mathcal{P}_0^*$ clearly solves a single layer optimization problem in $\mathcal{P}_0^*$ only,

$$\begin{aligned}
&\sup_{\mathcal{P}\in\mathcal{A}} \left\{ E^{X_0^-,t_0^-}\left[\sum_{i=0}^{M-1} \mathfrak{q}_i + \kappa\left(W^* + \frac{1}{\alpha}\min(W_{\mathcal{P}}(T) - W^*, 0)\right) \quad + \epsilon W_T \right] \right\} \\
&\text{subject to } \begin{cases} \text{for} \quad i = 0,\ldots,M-1 \\ \quad \mathfrak{p}_i(\cdot) = \mathfrak{p}_i(W_{\mathcal{P}}(t_i^+)), \mathfrak{q}_i(\cdot) = \mathfrak{q}_i(W_{\mathcal{P}}(t_i^+)) \\ \quad (\mathfrak{q}_i(\cdot), \mathfrak{p}_i(\cdot)) \in \mathcal{Z}(W_{\mathcal{P}}(t_i^-), W_{\mathcal{P}}(t_i^+), t_i), \\ \quad W_{\mathcal{P}}(t_i^+) = W_{\mathcal{P}}(t_i^-) - \mathfrak{q}_i(\cdot), \\ \quad W_{\mathcal{P}}(t_{i+1}^-) = \mathfrak{p}_i(\cdot)^T e^{\boldsymbol{R}(t_i)} W_{\mathcal{P}}(t_i^+) \end{cases}
\end{aligned} \tag{2.15}$$

Fortunately, the linear shortfall formulation (2.15), in which W^* is a constant, is time consistent . Therefore, the retiree can rest assured that the optimal strategy $\mathcal{P}_0^*$ is also a time consistent strategy, when viewed from the perspective of the linear shortfall formulation (2.15). For further discussion of the relationship between time consistent and pre-commitment strategies, see (Vigna, 2014; Menoncin and Vigna, 2017; Vigna, 2017; Strub et al., 2019; Forsyth, 2020; Bjork et al., 2021; Cui et al., 2022).

3 Data Driven Global-in-Time Neural Network Solutions

When the number of assets N is less than 3 and asset dynamics are specified parametrically, the constrained stochastic optimal control problem (2.13) or (2.14) can be solved by converting it to a Hamilton-Jacobi-Bellman equation, derived from the dynamic programming principle, e.g., (Chen et al., 2025; Forsyth and Labahn, 2026).

To numerically compute optimal multi-asset decumulation strategies, we have developed a global-in-time neural network approach for solving constrained stochastic optimal control problems, see, e.g., (Chen et al., 2025). In this approach, we assume that controls are represented by a neural network (NN) with time as one of the inputs, i.e., a single NN represents the controls across all rebalancing times $t_0, \ldots, t_{M-1}$. Specifically, for the decumulation problem, we have $\mathfrak{q}_i\left(X(t_i^-)\right) = \mathfrak{q}_{\theta_\mathfrak{q}}\left(X(t_i^-), t_i\right), \mathfrak{p}_i\left(X(t_i^-)\right) = \mathfrak{p}_{\theta_\mathfrak{p}}\left(X(t_i^-), t_i\right)$, for $i = 0,\ldots,M-1$, where each of the functions $\mathfrak{q}_{\theta_\mathfrak{q}}$ and $\mathfrak{p}_{\theta_\mathfrak{p}}$ is represented by an NN. This is in contrast to a stacked NN approach in which a separate NN is used for the control at each rebalancing time, see e.g., (Han et al., 2018).

Specifically, we assume that a sufficient number of i.i.d. samples $\{\boldsymbol{R}^j(t_i), i = 0,1,\ldots,M-1, j = 1,\ldots L\}$ of the random log return $\boldsymbol{R}$ are given, drawn from either a nonparametric model (e.g., bootstrap block resampling) or a parametric stochastic model (e.g., a jump diffusion model Kou (2002)).

A sample-based NN approximation of (2.14) is given below:

$$\begin{aligned}
&\sup_{W'} \sup_{\mathcal{P}=(\mathfrak{p}_{\theta_\mathfrak{p}}, \mathfrak{q}_{\theta_\mathfrak{q}})\in\mathcal{A}} \left\{ \frac{1}{L}\sum_{j=1}^{L}\left[\sum_{i=0}^{M-1} \mathfrak{q}_{\theta_\mathfrak{q}}(X^j(t_i^-),t_i) + \kappa\left(W' + \frac{1}{\alpha}\min(W^j_{\mathcal{P}_\theta}(T) - W', 0)\right) \quad + \epsilon W^j_{\mathcal{P}_\theta}(T) \right] \right\} \\
&\text{subject to } \begin{cases} \text{for} \quad i = 0,\ldots,M-1, j = 1,\ldots,L \\ \quad (\mathfrak{q}_{\theta_\mathfrak{q}}(X^j(t_i^-),t_i), \mathfrak{p}_{\theta_\mathfrak{p}}(X^j(t_i^-),t_i)) \in \mathcal{Z}(W^j_{\mathcal{P}_\theta}(t_i^-), W^j_{\mathcal{P}_\theta}(t_i^+), t_i), \\ \quad W^j_{\mathcal{P}_\theta}(t_i^+) = W^j_{\mathcal{P}_\theta}(t_i^-) - \mathfrak{q}_{\theta_\mathfrak{q}}(X^j(t_i^-),t_i), \\ \quad W^j_{\mathcal{P}_\theta}(t_{i+1}^-) = \mathfrak{p}_{\theta_\mathfrak{p}}(X^j(t_i^-),t_i)^T e^{\boldsymbol{R}^j(t_i)} W^j_{\mathcal{P}_\theta}(t_i^+) \end{cases}
\end{aligned} \tag{3.1}$$

In the subsequent computational investigation, we analyze the optimal decumulation strategies computed in two different investment settings: first, a two-asset case ($N = 2$) consisting of an equity index and a bond index; and second, a six-asset case consisting of US cap- and equal-weighted indexes (in CAD), Canadian cap- and equal-weighted indexes, and CAD short- and long-term bonds (all CPI adjusted).

4 Constraint-Feasible Activation Functions

Solving the sample-based constrained NN optimization problem (3.1) is computationally challenging since there can be many assets, scenarios, and time periods. Furthermore, a critical computational challenge is that the solution to the constrained optimization problem (4.1) is computed using stochastic gradient descent, which is applicable only to unconstrained optimization. Recall also that we use a single global-in-time neural network to approximate the control $\mathcal{P} = \{\mathfrak{q}_i,\mathfrak{p}_i\}_{i=0}^{M-1}$ at all rebalancing times $\mathcal{T}$, using t_i as an input feature to the neural network approximation at time t_i, $i = 0,\ldots,M-1$.

In order to achieve feasibility, our approach is to design a custom NN constraint-feasible activation function for a feedforward NN so that the final output of the NN automatically yields a feasible control, and so that any feasible control can be approximated by an NN control arbitrarily well. This requires that, for any $\theta = (\theta_\mathfrak{p},\theta_\mathfrak{q}) \in \mathbb{R}^{\dim(\theta)}$, $(\mathfrak{q}_{\theta_\mathfrak{q}},\mathfrak{p}_{\theta_\mathfrak{p}}) \in \mathcal{A}$. Assuming existence of constraint-feasible activation functions, we then solve unconstrained minimization below

$$\sup_{W'} \sup_{\theta=(\theta_\mathfrak{p},\theta_\mathfrak{q})\in\mathbb{R}^{\dim(\theta)}} \left\{ \frac{1}{L}\sum_{j=1}^{L}\left[\sum_{i=0}^{M-1} \mathfrak{q}_{\theta_\mathfrak{q}}(X^j(t_i^-),t_i) + \kappa\left(W' + \frac{1}{\alpha}\min(W^j_{\mathcal{P}_\theta}(T) - W', 0)\right) \quad + \epsilon W^j_{\mathcal{P}_\theta}(T)\right]\right\}$$

$$\text{subject to} \begin{cases} \text{for} \quad i = 0,\ldots,M-1, j = 1,\ldots,L \\ \quad (\mathfrak{q}_{\theta_\mathfrak{q}}(X^j(t_i^-),t_i),\mathfrak{p}_{\theta_\mathfrak{p}}(X^j(t_i^-),t_i)) \in \mathcal{Z}(W^j_{\mathcal{P}_\theta}(t_i^-), W^j_{\mathcal{P}_\theta}(t_i^+),t_i), \quad \text{(ensured by activation function)} \\ \quad W^j_{\mathcal{P}_\theta}(t_i^+) = W^j_{\mathcal{P}_\theta}(t_i^-) - \mathfrak{q}_{\theta_\mathfrak{q}}(X^j(t_i^-),t_i), \\ \quad W^j_{\mathcal{P}_\theta}(t_{i+1}^-) = \mathfrak{p}_{\theta_\mathfrak{p}}(X^j(t_i^-),t_i)^T e^{\boldsymbol{R}^j(t_i)} W^j_{\mathcal{P}_\theta}(t_i^+) \end{cases} \tag{4.1}$$

More explicitly, we want to design a constraint-feasible NN such that

$$(\mathfrak{q}_{\theta_\mathfrak{q}}(X^j(t_i^-),t_i),\mathfrak{p}_{\theta_\mathfrak{p}}(X^j(t_i^-),t_i)) \in \mathcal{Z}(W^j_{\mathcal{P}_\theta}(t_i^-), W^j_{\mathcal{P}_\theta}(t_i^+),t_i)$$

This means that

- The constraints specified in (2.4) and (2.5) are satisfied for every possible value of the weights $\theta_\mathfrak{q}$ and $\theta_\mathfrak{p}$ of the NN.
- Simultaneously, any feasible control in the original stochastic optimal control problem can be approximated to arbitrary accuracy by the neural network controls.

Previous neural network approaches to discrete time financial stochastic control problems have achieved control feasibility by incorporating custom activation functions designed to satisfy problem constraints based on individual constraint specifications. For allocation problems in which shorting and leverage are not permitted, the softmax function is generally used as the final neural network activation function. This approach is effective because the set of feasible allocations of N assets consists of the N-simplex, and because the image of softmax is the interior of the N-simplex. See (Chen et al., 2025) and (Li and Forsyth, 2019) for example. Additionally, Ni et al. (2024) and van Staden et al. (2024) devise *leverage-feasible neural networks*, whose output activation functions are custom-designed composite functions, capable of achieving specified leverage-feasibility and arbitrary solution accuracy.

In this paper, we develop an global-in-time neural network approach for discrete-time financial optimal control problems for which the feasible control value set at each rebalancing time can be described by linear constraints. Specifically, we assume that the feasible set of control values

$$\begin{aligned} &\mathcal{Z}(W_i^-, W_i^+,t_i) = \mathcal{Z}_\mathfrak{q}(W_i^-, t_i) \times \mathcal{Z}_\mathfrak{p}(W_i^+,t_i), \\ &\mathcal{Z}_\mathfrak{p}(W_i^+,t_i) = \begin{cases} \mathcal{Z}_\mathfrak{p}^+(W_i^+,t_i) & \text{when } W_i^+ > 0, i = 0,1,\ldots M-1 \\ \{0\} & \text{when } W_i^+ \le 0, i = 0,1,\ldots M-1 \end{cases}, \end{aligned} \tag{4.2}$$

and both feasible sets $\text{ri}(\mathcal{Z}_\mathfrak{q})$ and $\text{ri}(\mathcal{Z}_\mathfrak{p})$ can be described by general linear constraints, where $\text{ri}(\mathcal{Z}_\mathfrak{p})$ and $\text{ri}(\mathcal{Z}_\mathfrak{q})$ denote the relative interior of $\mathcal{Z}_\mathfrak{p}$ and $\mathcal{Z}_\mathfrak{q}$ respectively. In our decumulation formulation, the feasible withdrawal control $\mathfrak{q} \in \mathcal{Z}_\mathfrak{q}$ is described by simple bounds. At any specified time t_i, for withdrawal control,

we ensure minimum wealth withdrawal and specify the following constraints: let $W = W_{\mathcal{P}}(t_i)$,

$$\mathcal{Z}_{\mathfrak{q}}(W) = \left\{ \mathfrak{q} : \left\{ \begin{array}{ll} \mathfrak{q} \in [q_{\min}, q_{\max}] & \text{if } W \geq q_{\max} \\ \mathfrak{q} \in [q_{\min}, W] & \text{if } q_{\min} < W < q_{\max} \\ \mathfrak{q} = q_{\min} & \text{if } W \leq q_{\min} \end{array} \right\} \right. \tag{4.3}$$

The feasible value set of allocation control in (4.4) is often time invariant and can be described as $\mathcal{Z}_{\mathfrak{p}}^{+} = \{\mathfrak{p} : A_{\mathfrak{p}} z \leq b_{\mathfrak{p}}\}$, where $A_{\mathfrak{p}}, b_{\mathfrak{p}}$ are constant matrices and vectors.

For allocation control $\mathfrak{p}$, we consider the base case of no shorting and no leverage, as well as the case in which bounded leverage is permitted. The allocation value set can be expressed via constraints which can be formulated as:

$$\left\{ \begin{array}{l} \sum_{j=1}^{N} z_j = 1 \\ \sum_{j=1}^{N_s} z_j \leq p_{s,\max} \\ z_j \geq 0 \text{ for } j \in \{1, \ldots, N-1\} \\ z_N \geq 1 - \max(1, p_{s,\max}) \end{array} \right. \tag{4.4}$$

where $\mathfrak{p}_i, i = 1, \ldots, N_s$, denotes risky assets, $N > N_s$, and $p_{s,\max}$ denotes the maximum total fraction that is permitted to the risky asset class.

Remark 4.1. *Typically, we can assume that there is only a single riskless asset, i.e., $N_S = N - 1$. Assume that shorting is not permitted,* (4.4) *can be written as*

$$\left\{ \begin{array}{l} \sum_{j=1}^{N} z_j = 1 \\ \sum_{j=1}^{N-1} z_j \leq p_{s,max} \\ z_j \geq 0 \text{ for } j \in \{1, \ldots, N-1\} \end{array} \right. \tag{4.5}$$

If $p_{s,max} \leq 1$, (4.4) *corresponds to the base case when both shorting and leverage are not allowed.*

If $p_{s,max} > 1$, it denotes the upper bound on the leverage. If the asset N represents a short-term bond, leverage is achieved via $\mathfrak{p}_N < 0$.

Remark 4.2. *For a Canadian retiree's perspectives, their equity exposure can further be decomposed into U.S. exposure and Canadian exposure. To model this, we can consider decompose the index set of all equity assets* {1,2, $\ldots, N-1$} *into $I_{U.S.}$ and Canadian I_{CAD} and requires exquity exposure not exceed* 70%*, while U.S. equity fraction is less than* 50% *of the total equity exposure. In other words, we formulate the following constraints:*

$$\left\{ \begin{array}{l} \sum_{j=1}^{N} z_j = 1 \\ \sum_{j=1}^{N-1} z_j \leq .7 \\ \sum_{j \in I_{U.S.}} z_j \leq 0.5 \sum_{j=1}^{N-1} z_j \\ z_j \geq 0 \text{ for } j \in \{1, \ldots, N-1\} \end{array} \right. \tag{4.6}$$

Linear constraints can always be equivalently expressed as linear inequality of the form $Az \leq b$. For example, if allocation prohibits shorting and leveraging, then

$$\mathcal{Z}_{\mathfrak{p}}^{+} = \{\mathfrak{p} : e^T z = 1, z \geq 0\} \tag{4.7}$$

where e is the column vector of ones, with the same dimension as $\mathfrak{p}$. Therefore, feasible region of constraints (4.7) can be expressed as $\{\mathfrak{p} : A_{\mathfrak{p}}\mathfrak{p} \leq b_{\mathfrak{p}}\}$, where $A_{\mathfrak{p}} = [e^T; -e^T; -I_{N,N}] \in \mathbb{R}^{(N+2)\times N}$ and $b_{\mathfrak{p}} = [1; -1; 0] \in \mathbb{R}^{N+2}$. Suitable $A_{\mathfrak{p}}$ and $b_{\mathfrak{p}}$ can be defined analogously for the constraints in (4.4).

The core idea of our approach here is to design the constraint-feasible activation for a feedforward neural network such that its final output is always feasible, while capable of approximating any feasible control arbitrarily well, provided the control lies within the feasible region defined by linear constraints.

Firstly, we define a feedforward neural network in Definition 4.1.

Definition 4.1. (Fully Connected Feedforward Neural Network). *A fully connected feedforward neural network (FNN) $\tilde{f}_{\boldsymbol{\theta}}$ maps an input vector $\boldsymbol{x} \in \mathbb{R}^{d_0}$ to an output vector $\boldsymbol{h} \in \mathbb{R}^{d_{K+1}}$, where FNN contains K hidden layers of sizes $d_1, \cdots, d_K$. The neural network is parameterized by the weight matrices $\boldsymbol{\theta}^{(k)} \in \mathbb{R}^{d_{k-1} \times d_k}$ and bias vectors $\boldsymbol{\theta}_b^{(k)} \in \mathbb{R}^{d_k}$, for $k = 1, \cdots, K+1$. Then, the output $\boldsymbol{h}$ is derived from the input $\boldsymbol{x}$ iteratively as follows.*

$$\begin{cases} \boldsymbol{x}^{(0)} = \boldsymbol{x}, \\ \boldsymbol{x}^{(k)} = \sigma\Big(\big(\boldsymbol{\theta}^{(k)}\big)^\top \cdot \boldsymbol{x}^{(k-1)} + \boldsymbol{\theta}_b^{(k)}\Big), \quad 1 \le k \le K, \\ \boldsymbol{h} = \big(\boldsymbol{\theta}^{(K+1)}\big)^\top \cdot \boldsymbol{x}^{(K)} + \boldsymbol{\theta}_b^{(K+1)}. \end{cases} \tag{4.8}$$

Here σ is the pointwise sigmoid activation function, i.e., for any vector $\boldsymbol{z}$, $[\sigma(\boldsymbol{z})]_i = \sigma(\boldsymbol{z}_i)$. For notational simplicity, we flatten and assemble all weight matrices and bias vectors into a single parameter vector $\boldsymbol{\theta} = (\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}_b^{(1)}, \cdots, \boldsymbol{\theta}^{(K+1)}, \boldsymbol{\theta}_b^{(K+1)})^\top \in \mathbb{R}^{N_{\boldsymbol{\theta}}}$, where $N_{\boldsymbol{\theta}} = \sum_{k=1}^{K+1}(d_{k-1} \cdot d_k + d_k)$. Furthermore, we use the 2-tuple $\Big(K, (d_1, \cdots, d_K)^\top\Big)$ to denote the hyperparameters, i.e., the number of hidden layers and the sizes of each hidden layer.

Let $\mathrm{Im}(\rho)$ denote the image of $\rho(\cdot)$. Below, we motivate the properties of the activation function required so that the unconstrained problem (4.1) can replace the constrained problem (3.1), following the NN universal approximation theorem. Specifically, applying the metric-space specialization of Theorem 3.3 of (Kratsios and Bilokopytov, 2020) (originally stated for general topological spaces), we obtain Lemma 4.1.

Lemma 4.1. *Let $\mathcal{X} \subset \mathbb{R}^l$ be a compact set, and $\mathcal{Y} \subset \mathbb{R}^m$. Let $\rho(\cdot) : \mathbb{R}^n \mapsto \mathcal{Y}$ satisfy the following:*

1. *ρ is continuous and has a right inverse on $\mathrm{Im}(\rho)$, i.e. $\exists \overrightarrow{\rho} : \mathrm{Im}(\rho) \mapsto \mathbb{R}^n$, s.t. $\rho(\overrightarrow{\rho}(z)) = z, \forall z \in \mathrm{Im}(\rho)$.*
2. *$\mathrm{Im}(\rho)$ is dense in $\mathcal{Y}$.*

Then, for any continuous $g : \mathcal{X} \mapsto \mathcal{Y}$, and any $\epsilon > 0$, there exists a choice of hyperparameters $\big(K, (d_1, \cdots, d_K)^\top\big)$ and parameter $\boldsymbol{\theta}$, such that the corresponding FNN $\tilde{f}_{\boldsymbol{\theta}} : \mathcal{X} \mapsto \mathbb{R}^n$ described in Definition 4.1 satisfies

$$\sup_{\boldsymbol{x} \in \mathcal{X}} \|\rho(\tilde{f}_{\boldsymbol{\theta}}(\boldsymbol{x})) - g(\boldsymbol{x})\| < \epsilon, \forall \boldsymbol{x} \in \mathcal{X}. \tag{4.9}$$

Here $\|\cdot\|$ denotes the vector norm.

Consider the allocation control $\mathfrak{p}$ with a total of N assets, for example. We design a NN with the output

$$\phi(X) = T_{\mathfrak{p}} \circ \tilde{f}_{\boldsymbol{\theta}}(X) \tag{4.10}$$

where $T_{\mathfrak{p}}(\cdot) : \mathbb{R}^{d_{K+1}} \mapsto \mathbb{R}^N$ is a constraint-feasible activation and $\tilde{f}_{\boldsymbol{\theta}}(\cdot) : \mathbb{R}^{d_0} \mapsto \mathbb{R}^{d_{K+1}}$ is defined in Definition 4.1. The "constraint-feasible activation function" $T_{\mathfrak{p}}$ is applied to the output of the feedforward NN.

Using Lemma 4.1, the goal is to choose constraint-feasible activation functions that are continuous and have a right inverse, with respect to the feasible regions $\mathcal{Z}_{\mathfrak{p}}$ and $\mathcal{Z}_{\mathfrak{q}}$, respectively[5].

We consider two different ways to construct an NN and a constraint-feasible output activation that together ensure feasibility.The first construction depends on specific forms of constraints, while the second method applies to any linear constraints. Subsequently, we illustrate both methods for the decumulation problem.

In this work, for the allocation control, we construct the constraint-feasible activation function using a technique which is applicable to the general linear constraint.

Assume that the feasibility of the allocation control can be expressed as time-invariant linear constraints $\mathcal{Z}_{\mathfrak{p}} = \{\mathfrak{p} : A_{\mathfrak{p}}\mathfrak{p} \le b_{\mathfrak{p}}\}$.

The basic idea follows from the classical Minkowski-Weyl theorem in convex analysis, e.g., Fukuda (2025), that any point in a polytope, i.e., bounded polyhedron, which can be expressed by $\mathcal{Z}_{\mathfrak{p}} = \{\mathfrak{p} : A_{\mathfrak{p}}\mathfrak{p} \le b_{\mathfrak{p}}\}$, is in the convex hull of its extreme points.

[5]Technically, Lemma 4.1 assumes a compact domain $\mathcal{X}$, while our domain of interest allows unbounded wealth W and so is non-compact. In practice, the behavior of the control at arbitrarily large wealth is not of interest, so we can bound W without any significant effect.

Suppose that $\mathfrak{n}_v$ denotes the total number of extreme points of $\mathcal{Z}_\mathfrak{p} = \{\mathfrak{p} : A_\mathfrak{p}\mathfrak{p} \leq b_\mathfrak{p}\}$.

Consider a feed-forward NN $\tilde{f}_{\boldsymbol{\theta}_\mathfrak{p}}$ with $d_{K+1} = \mathfrak{n}_v$ and let $\boldsymbol{h}$ denote the output of this FNN. Recall the softmax function for $\boldsymbol{h}$ below:

$$\text{softmax}(\boldsymbol{h})_i := \frac{\exp(\boldsymbol{h}_i)}{\sum_{j=1}^{\mathfrak{n}_v} \exp(\boldsymbol{h}_j)}, \quad i = 1, \ldots, \mathfrak{n}_v.$$

Let $\mathcal{Z}_\mathfrak{p} = \text{conv}\{\mathfrak{v}_1, \ldots, \mathfrak{v}_{\mathfrak{n}_v}\}$ be the set of extreme points for $\mathcal{Z}_\mathfrak{p} = \{\mathfrak{p} : A_\mathfrak{p}\mathfrak{p} \leq b_\mathfrak{p}\}$. These extreme points can be computed by standard polyhedral software packages like `cddlib` (Fukuda, 2024). `cddlib` uses the Double Description Method (Motzkin et al., 1953) to find extreme points of a polytope. This method starts by considering the simple polyhedron Z_0 defined by a subset of the inequalities of interest. This simple polyhedron is chosen so that it easy to describe as a sum of extreme points and rays:

$$Z_0 = \text{conv}\, V_0 + \text{cone}\, R_0. \tag{4.11}$$

In each subsequent step, the algorithm adds an additional inequality to the polyhedron of the previous step Z_i to generate a new polyhedron Z_{i+1}. It updates the vertices and rays so that:

$$Z_{i+1} = \text{conv}\, V_{i+1} + \text{cone}\, R_{i+1}. \tag{4.12}$$

Repeating this until all inequalities have been included will result in a description of the full polytope as the convex hull of its extreme points.

Given the extreme points $\{\mathfrak{v}_1, \ldots, \mathfrak{v}_{\mathfrak{n}_v}\}$ that correspond to a given set of feasible allocations $\mathcal{Z}_\mathfrak{p}$, the constraint-feasible activation function is defined by:

$$T_\mathfrak{p}(\boldsymbol{h}) = \sum_{i=1}^{\mathfrak{n}_v} \text{softmax}(\boldsymbol{h})_i \mathfrak{v}_i. \tag{4.13}$$

We can establish Lemma 4.2 below for the allocation constraint-feasible activation function $T_\mathfrak{p}(\cdot)$.

Lemma 4.2. *The activation function $T_\mathfrak{p}(\cdot)$ defined in (4.13) is a map from $\mathbb{R}^{\mathfrak{n}_v} \mapsto \text{ri}(\mathcal{Z}_\mathfrak{p}^+)$, is continuous, and has a right inverse. Furthermore, $\text{Im}(T_\mathfrak{p})$ is dense in $\mathcal{Z}_\mathfrak{p}^+$.*

Proof. It is clear that $T_\mathfrak{p}$ is continuous. It is also clear that

$$\text{Im}(T_\mathfrak{p}) \subset \text{ri}(\mathcal{Z}_\mathfrak{p}^+).$$

Let $\mathfrak{p} \in \text{ri}(\mathcal{Z}_\mathfrak{p})$. By basic convex analysis[6], there exists $0 < \lambda_i < 1, i = 1, \ldots, \mathfrak{n}_v$, such that $\sum_{i=1}^{\mathfrak{n}_v} \lambda_i = 1$ and

$$\mathfrak{p} = \sum_{i=1}^{\mathfrak{n}_v} \lambda_i \mathfrak{v}_i.$$

Define $\boldsymbol{h} \in \mathbb{R}^{\mathfrak{n}_v}$ by

$$\boldsymbol{h}_i = \log(\lambda_i), \ i = 1, \ldots, \mathfrak{n}_v$$

Clearly $\text{softmax}(\boldsymbol{h}) = \lambda$ and thus $T_\mathfrak{p}(h) = \mathfrak{p}$. In addition, define

$$\vec{T}_\mathfrak{p}(\mathfrak{p}) = \boldsymbol{h}.$$

Then

$$T_\mathfrak{p}(\vec{T}_\mathfrak{p}(\mathfrak{p})) = \mathfrak{p}.$$

This shows that $T_\mathfrak{p}$ has a right inverse, and that $\text{Im}(T_\mathfrak{p}) = \text{ri}(\mathcal{Z}_\mathfrak{p}^+)$. Since $\text{ri}(\mathcal{Z}_\mathfrak{p}^+)$ is dense in $\mathcal{Z}_\mathfrak{p}^+$, we also have that $\text{Im}(T_\mathfrak{p})$ is dense in $\mathcal{Z}_\mathfrak{p}^+$.

□

[6] See Hiriart-Urruty and Lemaréchal (2013), for example.

Using Lemma 4.2 and Lemma 4.1, it can be concluded that any feasible continuous allocation control can be approximated by $T_{\mathfrak{p}} \circ \tilde{f}_{\boldsymbol{\theta}_{\mathfrak{p}}}(X)$ arbitrarily well.

Note that we have focused on approximating the control in the $W > 0$ case above. When $W \leq 0$, (4.2) tells us that the set of feasible controls is trivial. In practice, we just implement this as a special case.

Application of Lemma 4.1 to the withdrawal control is complicated by the fact that the withdrawal feasible set depends on the wealth W. To address this issue, we first consider a feedforward NN with output $\boldsymbol{h} \in \mathbb{R}^2$, We design the activation function $\hat{T}_{\mathfrak{q}}$ with $\mathrm{Im}(\hat{T}_{\mathfrak{q}}) \subset \mathbb{R}^2$ explicitly as follows: let $\boldsymbol{h} \in \mathbb{R}^2$

$$\begin{cases} \left(\hat{T}_{\mathfrak{q}}(\boldsymbol{h})\right)_1 = \boldsymbol{h}_1 \\ \left(\hat{T}_{\mathfrak{q}}(\boldsymbol{h})\right)_2 = q_{\min} + \left(\min(\max(q_{\min},\boldsymbol{h}_1),q_{\max}) - q_{\min}\right)\left(\frac{e^{\boldsymbol{h}_2}}{1+e^{\boldsymbol{h}_2}}\right) \end{cases} \tag{4.14}$$

where $\boldsymbol{h} = \tilde{f}_{\boldsymbol{\theta}_{\mathfrak{q}}}$ is the output of a feedforward NN with $d_{K+1} = 2$ and $\boldsymbol{h}_1 = W$, the current wealth.

We also define the graph of corresponding to $\hat{T}_{\mathfrak{q}}(\cdot)$ as

$$\mathcal{Y}_{\mathfrak{q}} = \bigcup_{\boldsymbol{h}_1 > 0} \{(\boldsymbol{h}_1, \mathfrak{q}) : \mathfrak{q} \in \mathcal{Z}_{\mathfrak{q}}(\boldsymbol{h}_1)\}. \tag{4.15}$$

With these definitions in hand, we can prove Lemma 4.3, which demonstrates that $\hat{T}_{\mathfrak{q}}(\cdot)$ satisfies the conditions required in Lemma 4.1.

Lemma 4.3. *The constraint-feasible activation function $\hat{T}_{\mathfrak{q}}(\cdot)$ defined in (4.14) is a map from $\mathbb{R}^2 \mapsto \mathrm{ri}(\mathcal{Y}_{\mathfrak{q}}) = \bigcup_{\boldsymbol{h}_1>0}\{(\boldsymbol{h}_1, \mathfrak{q}) : \mathfrak{q} \in \mathrm{ri}(\mathcal{Z}_{\mathfrak{q}}(\boldsymbol{h}_1))\}$. In addition, $\hat{T}_{\mathfrak{q}}(\cdot)$ has a right-inverse and $\mathrm{Im}(\hat{T}_{\mathfrak{q}})$ is dense in $\mathcal{Y}_{\mathfrak{q}}$.*

Proof. Consider any $\boldsymbol{h}_1 \in \mathbb{R}$.

Assume that $\hat{\mathfrak{q}} = \hat{T}_{\mathfrak{q}}(\boldsymbol{h})$ for some $\boldsymbol{h}_2 \in \mathbb{R}$. Since

$$0 < \frac{e^{\boldsymbol{h}_2}}{1 + e^{\boldsymbol{h}_2}} < 1,$$

we have

$$q_{\min} < \hat{\mathfrak{q}}_2 < q_{\min} + \left(\min(\max(q_{\min},W),q_{\max}) - q_{\min}\right) = \min(q_{\max},W)$$

where $W = \boldsymbol{h}_1$. Hence $\hat{\mathfrak{q}} \in \mathrm{ri}(\mathcal{Y}_{\mathfrak{q}})$.

Conversely, let $\hat{\mathfrak{q}} \in \mathrm{ri}(\mathcal{Y}_{\mathfrak{q}})$. This implies that, for some $W = \boldsymbol{h}_1 \in \mathbb{R}$,

$$q_{\min} < \hat{\mathfrak{q}}_2 < q_{\min} + \left(\min(\max(q_{\min},W),q_{\max}) - q_{\min}\right) = \min(q_{\max},W).$$

Denote $\beta = \left(\min(\max(q_{\min},W),q_{\max}) - q_{\min}\right)$. Let $(\vec{\hat{T}}_{\mathfrak{q}})_1 = \hat{\mathfrak{q}}_1$ and

$$(\vec{\hat{T}}_{\mathfrak{q}})_2 = \log\left(\frac{\hat{\mathfrak{q}}_2 - q_{\min}}{\beta - (\hat{\mathfrak{q}}_2 - q_{\min})}\right)$$

Clearly, we have

$$\hat{T}_{\mathfrak{q}}(\vec{\hat{T}}_{\mathfrak{q}}(\hat{\mathfrak{q}})) = \hat{\mathfrak{q}}.$$

Hence $\mathrm{Im}(\hat{T}_{\mathfrak{q}})$ is dense in $\mathcal{Y}_{\mathfrak{q}}$, $T_{\mathfrak{q}}$ is continuous, and has right-inverse in $\mathbb{R}^2$. □

When applying Lemma 4.1 to the withdrawal control, we define as $g(X(t)) := (W(t), \mathfrak{q}(X(t))$, where $X(t) = (W(t),t)$ denotes the control state and $\mathfrak{q}(\cdot)$ denotes a withdrawal control function. The combination of Lemma 4.3 and Lemma 4.1 with this choice of g shows that we can approximate any continuous feasible withdrawal control via the combination of a feedforward neural network with $d_{K+1} = 2$ and the activation function $\hat{T}_{\mathfrak{q}}(\cdot)$ with any desired accuracy. However, if $\boldsymbol{h}_1$ only approximates W, where W denotes the current wealth, the approximating neural network can return an infeasible but sufficiently accurate approximation

to a feasible control. Fortunately, by explicitly replacing $h_1 = W$, i.e. using W as a passing through variable direatly to the activation function $\hat{T}_{\mathfrak{q}}$, we show that, in fact, there exists a FNN with $d_{K+1} = 1$ and (direct) activation function

$$T_{\mathfrak{q}}(\boldsymbol{h}) = \left\{ \quad \big(T_{\mathfrak{q}}(\boldsymbol{h})\big) = q_{\min} + \left(\min(\max(q_{\min}, W, q_{\max}) - q_{\min}\right)\left(\tfrac{e^{\boldsymbol{h}}}{1+e^{\boldsymbol{h}}}\right) \right. \tag{4.16}$$

which generates a feasible sufficiently accurate approximation for any continuous and feasible withdrawal control.

Note also that in our definition of $\mathcal{Y}_{\mathfrak{q}}$ we restricted ourselves to the case $W > 0$. When $W \leq 0$, the set of feasible controls is trivial, so we do not require a neural network approximation.

For the decumulation problem, the optimal withdrawal control is discontinuous (bang-bang)Nevertheless, for any specified tolerance, there exists a continuous function that approximates it in L_2 to within the required tolerance . Therefore, using Lemma 4.3 and Lemma 4.1, it can be concluded that any feasible withdrawal control (and in particular the optimal withdrawal control) can be approximated by $T_{\mathfrak{q}} \circ \tilde{f}_{\boldsymbol{\theta}_{\mathfrak{q}}}(X)$ arbitrarily well.

Using the constraint-feasible activation functions $T_{\mathfrak{p}}$ and $T_{\mathfrak{q}}$, the NN training problem becomes an unconstrained optimization problem, i.e., $\mathcal{A}_\theta = \mathbb{R}^{\dim(\theta)}$.

For a more rigorous proof, in the general case, that this strategy can approximate any control in a constrained optimal allocation problem, see Appendix C.

5 Empirically Discovered Characteristics of Optimal Strategies

Using the global-in-time neural network (NN) approach, we analyze and compare the performance of decumulation strategies from the perspective of a Canadian retiree, considering portfolios with both two and six assets. We derive insights into the relative importance of individual assets within the resulting optimal strategies. Additionally, we examine the effects of leverage, as well as constraints on withdrawal and allocation decisions. Finally, we evaluate the robustness of the computed optimal strategies.

5.1 Data

For our computational studies, we focus on an asset universe consisting of six assets listed in Table 5.1. Table 5.2 shows the correlation structure and summary statistics of these six assets.

Asset short name	Monthly data series	index name
Ctbill	Real Canadian 30-day Treasury bill index	
Cltbond	Real Canadian long-term bond index	
Cvw	Real Canadian capitalization-weighted equity index	CFMRC Value Weighted Index [7]
Cew	Real Canadian equal-weighted equity index	CFMRC EW
USvw	Real US capitalization-weighted equity index	VWRETD
USew	Real US equal-weighted equity index	EWRETD

TABLE 5.1: *Six assets under study. Historical monthly data range 1950-2025. US data source is from the Center for Research in Security Prices (CRSP), and the Canadian data source is from the Canadian Financial Markets Research Center (CFMRC). All prices are converted to Canadian dollars using the contemporaneous exchange rate and deflated using the Canadian CPI index.*

Our historical data covers the period 1950-2025, and is sourced from the Center For Research in Security Prices (CRSP, US data), and the Canadian Financial Markets Research Center (CFMRC, Canadian data). It consists of monthly prices, which we convert to monthly returns.

Since we adopt the perspective of a Canadian retiree, we convert all prices to Canadian dollars using the contemporaneous exchange rate. Furthermore, as we assume that retirees are chiefly concerned with the purchasing power of their wealth, we deflate all asset prices using the Canadian CPI index.

[7] In the common terminology this would be capitalization weighted

	USvw	USew	Cvw	Cew	Cltbond	Ctbill
USvw	1.0000	0.8511	0.7389	0.5990	0.2171	0.1426
USew	0.8511	1.0000	0.7181	0.7197	0.1256	0.1158
Cvw	0.7389	0.7181	1.0000	0.8710	0.2149	0.0913
Cew	0.5990	0.7197	0.8710	1.0000	0.1210	0.0685
Cltbond	0.2171	0.1256	0.2149	0.1210	1.0000	0.3270
Ctbill	0.1426	0.1158	0.0913	0.0685	0.3270	1.0000

(a) *Correlation matrix of the monthly returns of all six assets.*

	USvw	USew	Cvw	Cew	Cltbond	Ctbill
mean	0.0071	0.0075	0.0059	0.0086	0.0024	0.0011
std	0.0399	0.0497	0.0427	0.0508	0.0253	0.0046
min	-0.2244	-0.2719	-0.2319	-0.2835	-0.1068	-0.0222
5%	-0.0613	-0.0715	-0.0638	-0.0706	-0.0366	-0.0066
25%	-0.0163	-0.0186	-0.0180	-0.0170	-0.0121	-0.0017
50%	0.0108	0.0074	0.0090	0.0117	0.0017	0.0013
75%	0.0329	0.0359	0.0325	0.0370	0.0168	0.0038
95%	0.0646	0.0827	0.0664	0.0767	0.0445	0.0084
max	0.1515	0.3094	0.1832	0.2825	0.1527	0.0169

(b) *Summary statstics of the monthly returns of all six assets.*

TABLE 5.2: *Correlation structure and summary statistics describing the monthly returns of all six assets in the period from 1950:12 to 2025:1. A yearly transaction cost of 100 basis points has been incorporated into the equal-weight index returns.*

Unlike a capitalization-weighted index, an equal-weighted index requires more rebalancing because as stock prices drift, the weights drift away from being equal. To restore balance, funds constantly have to sell winners and buy losers, generating many transactions. To account for possible implementation costs, we impose a hundred-basis-point per-year penalty on Cew and USew, the two equal-weighted indices.

We use the historical data described in Section 5.1 to generate sample paths via the stationary block bootstrap procedure (Politis and Romano, 1994). Compared to fitting a parametric model to historical data, this approach has the advantage of retaining granular features of the real data such as autocorrelation. Compared with considering only the single realized historical path, as is done in the older spending rule literature (Bengen, 1994), the stationary bootstrap approach reduces risk in overfitting. We experimented with different block sizes of three and six months to generate the paths for training.

5.2 Optimal Decumulation when shorting and leverage are not permitted

Investment horizon T (years)	30
Assets	See Table 5.1
Initial portfolio value W_0	1000
Cash withdrawal times	$t = 0, 1, \ldots, 29$
Withdrawal range	$[q_{\min}, q_{\max}]$
Equity fraction range	$[0, 1]$
Borrowing spread μ_c^b	0.03
Rebalancing interval (years)	1

TABLE 5.3: *DC-plan decumulation base case setup. Monetary units: CAD$ in thousands.*

Table 5.3 summarizes the base parameter specification of the decumulation problem under study. We assume the initial wealth is 1,000,000 and the decumulation time horizon is 30 years. The portfolio is

rebalanced annually. The annual withdrawal amount is limited to an amount between 30,000 and 60,000 inclusively. In this section, shorting or leverage is not permitted and thus we assume that $p_{s,max} = 1$ in (4.4). We compute and assess the performance of the optimal decumulation strategies for two-assets and six-asset cases respectively, from perspectives of a Canadian retiree.

Previous work using neural networks to approximate the optimal control in the DC plan decumulation problem has mostly focused on the two-asset case (e.g. Forsyth and Labahn (2026)). We consider two separate two-asset settings. In the first setting, the retiree has access to Canadian short-term bonds (Ctbill) and the Canadian capitalization-weighted index (Cvw). In the second setting, the retiree instead invests in the US capitalization-weighted index (USvw). This second scenario accounts for the view popular among Canadian retail investors, that investing exclusively in Canadian equity is risky due to its high concentration in financials and natural resources stocks (Carrick, 2011).

Figure 5.1 shows that the EW/ES Pareto frontier induced by the six-asset setting dominates the Pareto frontier from the two-asset setting: for every level of tail-tail risk, a retiree is able to withdraw more funds in the six-asset setting than in either two-asset setting. For all three settings, the decummuation strategy outperforms the Begen's rule [8] significantly.

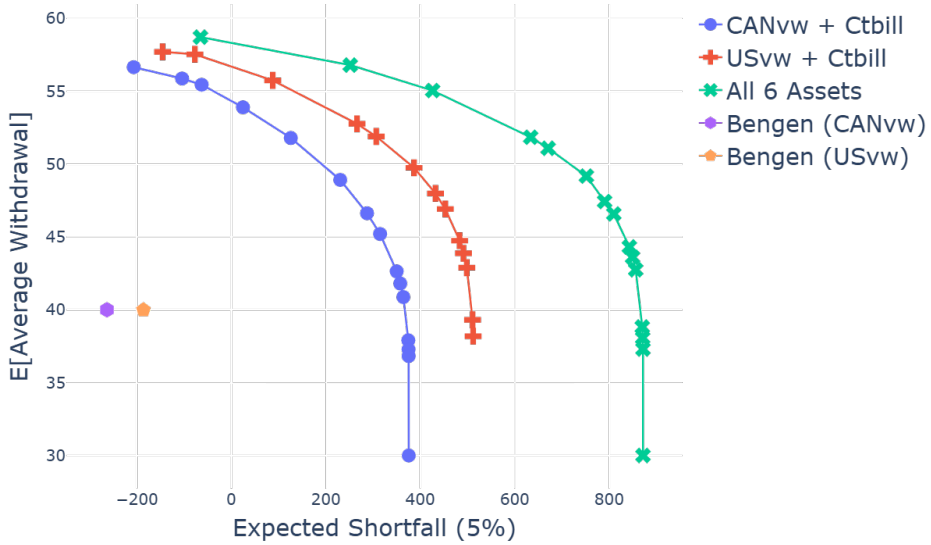

FIGURE 5.1: *EW/ES efficient frontier of DC-plan decumulation problem of the Canadian retiree. Comparing six asset settings with two separate two-asset settings. Markers on curves are κ values in range of* $[0.1,10^{10}]$. *The two markers on the left are the results of applying Bengen's rule to two different two-asset 50/50 stock-bond portfolios.*

Figure 5.2 compares the withdrawal control as a function of time and wealth in the two- and six-asset cases. It shows that the control is approximately bang-bang. Plots in Figure 5.2 indicate that withdrawing the maximum starts at a lower wealth cutoff level for the six-asset case, compared to both two-asset cases.

To analyze the allocation control as a function of time and wealth, we examine ES around 300. This is a conservative risk-averse use case to provide a robust longevity insurance buffer to ensure that in the worst 5% of market outcomes, a retiree has 300,000 in real wealth to fund maximum withdrawal amount between ages 96 and 100. Figure 5.3 is the allocation control as a function of time and wealth in a two-asset case when $\kappa = 0.683$. This corresponds to ES of 300.47, and EW of 52.06. For a given time, the higher the wealth, the less allocation goes into the risky asset (USvw) and the more allocation goes into the risk-free asset (Ctbill).

[8] We experiment with Bengen's rule using two 50/50 stock-bond portfolios, comprising an equity index and the Canadian short-term Treasury bond. The equity index is the US capitalization-weighted index in one portfolio and the Canadian capitalization-weighted index in the other portfolio.

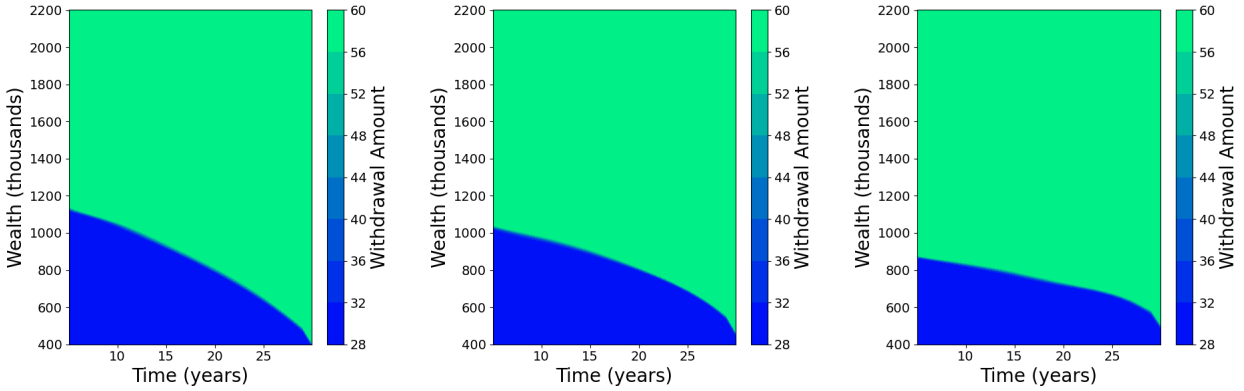

(a) *Two Asset Case: Cvw and Ctbill*

(b) *Two Asset Case: USvw and Ctbill*

(c) *Six Asset Case*

FIGURE 5.2: *Comparison of withdrawal control as a function of time and wealth in the two- and six-asset cases when* $\kappa = 1.530$, 0.683, *and* 0.292 *respectively. This corresponds to ES of* $301.39, 300.47$, *and* 299.17, *with AW of* $45.94, 52.06$, *and* 56.27 *for (a), (b), and (c) respectively. All three subplots show that the control is approximately bang-bang.*

Figure 5.4 shows asset allocation as a function of time and wealth in the six-asset case when $\kappa = 0.292$. It corresponds to ES of 299.17 and EW of 56.27. We also observe the similar allocation strategy that more allocation goes into risk-free asset (Ctbill) when we have more wealth for a given time. It is also interesting to observe that the US capitalization-weighted, Canadian equal-weighted, and the Canadian short-term T-bill are the major assets in the allocation. Both US equal-weighted and Canadian capitalization-weighted assets are close to zero allocations. Further analysis of key assets is in Section 5.5. Figure 5.5 shows percentiles of the wealth distribution as a function of time.

5.3 Impact of risky asset exposure

Let $\mathfrak{p}_1, \ldots, \mathfrak{p}_{N_s}$ denote the allocations to the risky assets (USvw, USew, Cvw, Cew, Cltbond), $p_{\text{s,max}}$is the maximum fraction that can be allocated to risky assets, and $\mathfrak{p}_N$ is the fraction allocated to the short-term bond asset (Ctbill). The constraints (4.4) model a cautious investor who wants to maintain a minimum fraction of their wealth in relatively safe bond assets. We implement these constraints using the strategy described in Section 4. We experiment with different risky asset exposure and withdrawal range settings to analyze the EW/ES efficient frontier sensitivity with respect to these constraints.

We study how risky asset exposure affects the EW/ES efficient frontier of DC-plan decumulation with the default withdrawal range of $[30, 60]$. Figure 5.6 compares the efficient frontiers of the six-asset cases with $p_{\text{s,max}} = 0.5, 1.0$, and 1.3, representing a moderate conservative baseline case, a high risk case, and a leveraged aggressive high risk case, respectively. Figure 5.6(a) shows that efficient frontiers are clustered in the region of $ES \in [0, 300]$, which is a prudent region to provide a longevity buffer (see Table B.5 for detailed numerical results). Figure 5.6(b) zooms into this region, where the gaps between the frontier of $p_{\text{s,max}} = 1.3$ and the frontier of $p_{\text{s,max}} = 1.0$ is very small, indicating a diminishing marginal utility of the leverage. It is worth mentioning that the noticeable gap between the efficient frontiers of $p_{\text{s,max}} = 0.5$ and $p_{\text{s,max}} = 1.0$ is no more than 2. Given the significant reduction in the risky asset exposure from 1.0 to 0.5, this gap is insignificant.

To analyze the risky asset exposure effect quantitatively, we conduct a sensitivity analysis of three metrics that measure global portfolio utility, global lifestyle utility, and local lifestyle utility (see Table 5.4) of a leverage-restricted decumulation strategy. The elasticity (η) of a metric M with respect to $p_{\text{s,max}}$measures how sensitive the metric is to the change of $p_{\text{s,max}}$and is defined as:

$$\eta_M = \frac{\partial \ln M}{\partial \ln p_{\text{s,max}}} \approx \frac{\%\Delta M}{\%\Delta p_{\text{s,max}}} = \frac{\Delta M/M}{\Delta p_{\text{s,max}}/p_{\text{s,max}}} \tag{5.1}$$

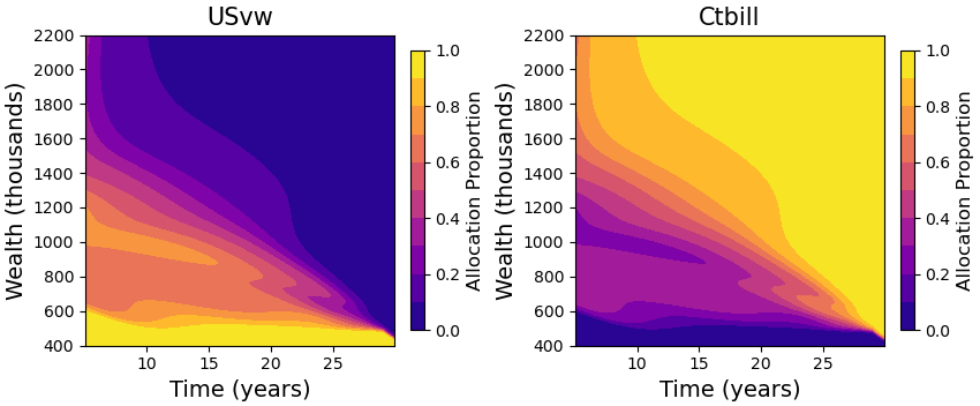

FIGURE 5.3: *The allocation control as a function of time and wealth in the two-asset case when* $\kappa = 0.683$. *This corresponds to ES of 300.47, and EW of 52.06.*

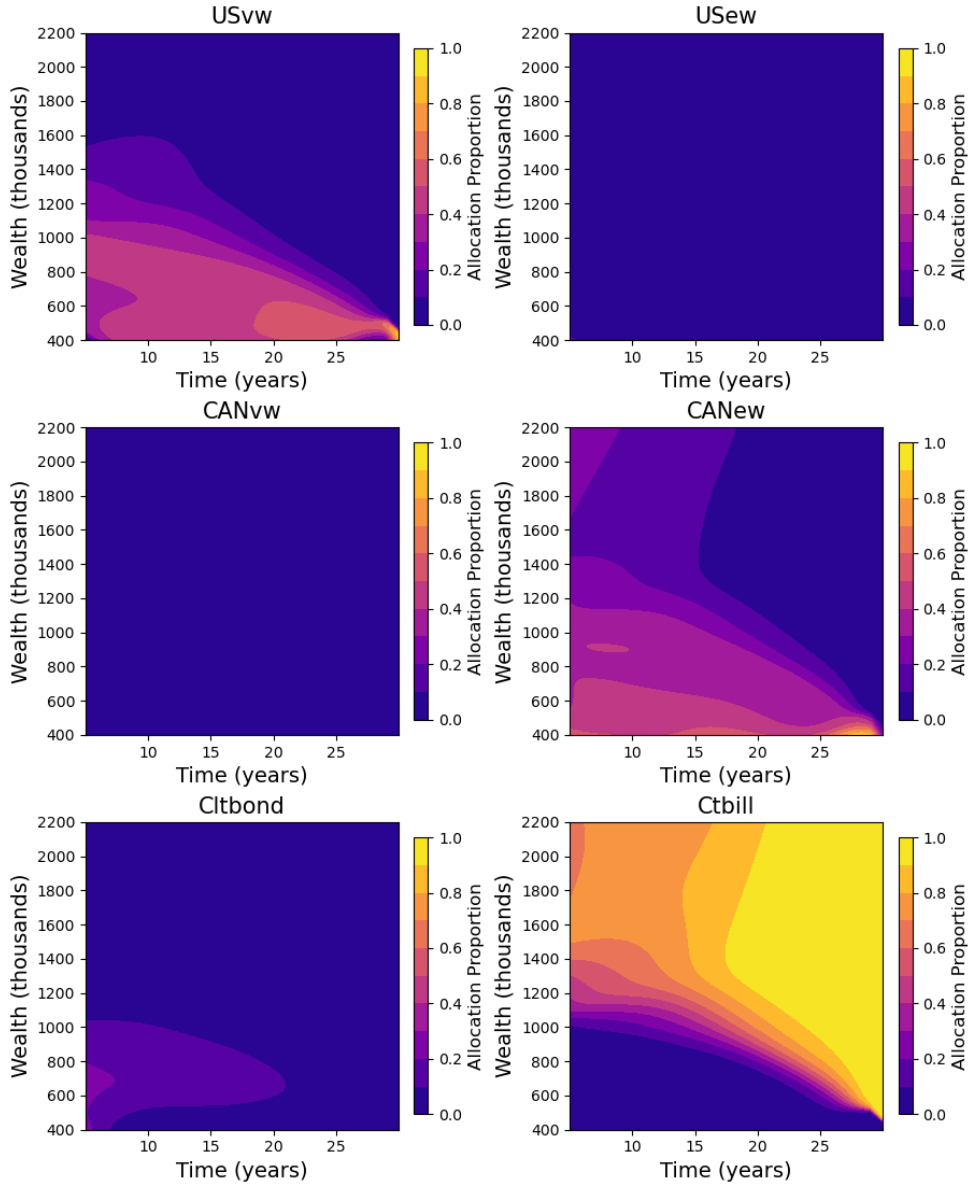

FIGURE 5.4: *The allocation control as a function of time and wealth in the six-asset case when* $\kappa = 0.292$. *This corresponds to ES of 299.17 and EW of 56.27.*

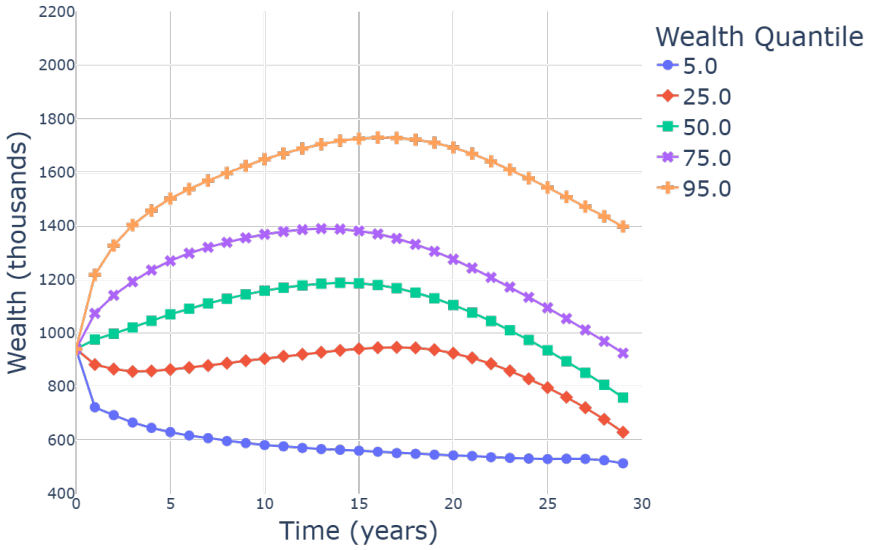

FIGURE 5.5: *The 5th, 25th, 50th, 75th, and 95th percentiles of the wealth distribution over time in the six asset base case when* $\kappa = 0.292$. *This corresponds to ES of 299.17 and EW of 56.27.*

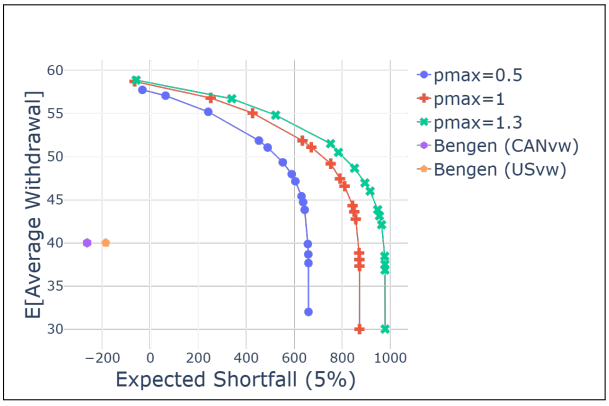

(a) *Comparison with Bengen's rule.*

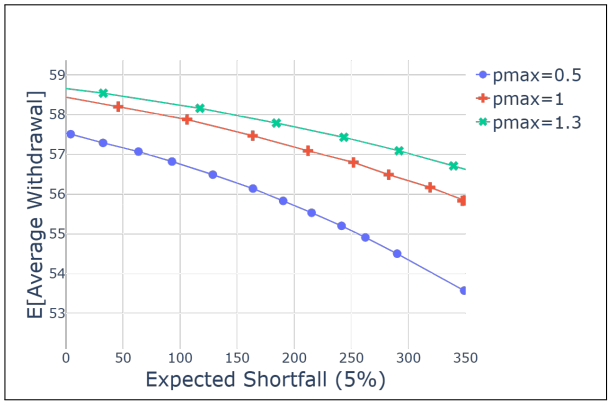

(b) *Case when* $ES \in [0, 300]$

FIGURE 5.6: *EW/ES efficient frontier of decumulation problem. Comparison of different risky asset allocation upper bound* $p_{s,max} \in \{0.5, 1.0, 1.3\}$. *(a) shows a wider range of EW and ES, and compares with Bengen's reference points on the plot. (b) shows the efficient frontier when* $ES \in [0, 300]$.

Metric Name	**Symbol**	**Economic Interpretation**	**Primary Role**
Raw Restricted AUC	A_R	Global Portfolio Utility	Global Efficiency
Surplus Restricted AUC	A_S	Global Lifestyle Utility	Global Lifestyle Efficiency
Surplus Withdrawal	EW_S	Local Lifestyle Utility	Local Lifestyle Efficiency

TABLE 5.4: *Leverage Efficiency Metrics*

5.3.1 Restricted Areas Under the Curve

ES and EW represent the risk and reward on the efficient frontier. To evaluate the decumulation strategy's global capacity to balance competing objectives of EW and ES holistically, we compute the Restricted Area Under the Curve (AUC) of the efficient frontier. General AUC measures the total reward-risk mass of the

strategy across the entire spectrum, which includes extreme risk scenarios that are economically less relevant to a retiree. Restricted AUC focuses on the restricted ES region of $\Omega = [ES_{\min}, ES_{\max}]$, which captures the realistic risk-reward trade-offs where a strategy is most likely to be deployed. Within this restricted domain, we further distinguish between the Raw Restricted AUC (A_R) and Surplus Restricted AUC (A_S).

The Raw Restricted AUC (A_R) measures the total mass of the frontier across region Ω, including the subsistence withdrawal floor EW_{ref} (e.g., $q_{\min}$). It describes the global utility available:

$$A_R(p_{\text{s,max}}) = \int_{\Omega} EW_{p_{\text{s,max}}}(ES)\, d(ES) \tag{5.2}$$

The Surplus Expected Withdrawal (EW_S) subtracts the subsistence floor (EW_{ref}) to isolate the utility of discretionary spending or "lifestyle" reward of a decumulation strategy. It is defined as

$$EW_S = EW - EW_{ref}\ . \tag{5.3}$$

The Surplus Restricted AUC (A_S) measures the total volume of discretionary utility available across region Ω generated by a decumulation strategy. We define A_S as the integral of the SEW across region Ω.

$$A_S(p_{\text{s,max}}) = \int_{\Omega} [EW_{p_{\text{s,max}}}(ES) - EW_{ref}]\, d(ES) \tag{5.4}$$

We focus our analysis on the prudent domain $\Omega_p = [0, 300]$ and the subsistence floor $EW_{ref} = q_{\min} = 30$. Table 5.5 compares both Raw and Surplus Restricted AUCs for three different levels of risky asset exposures and their elasticities.

When we increase the $p_{\text{s,max}}$from 0.5 to 1.0, which is an 100% increase, the gains in A_R and A_S are only 2.4% and 5.14% respectively. The elasticity of A_R and A_S are ultra low at 2.4% and 5.14% respectively, showing low marginal utility of risky asset exposure. Adding leverage by increasing the $p_{\text{s,max}}$from 1.0 to 1.3 only increases A_R and A_S by 0.75% and 1.56% respectively, showing diminishing marginal utility of leverage. The elasticity of A_R and A_S remains very low at 2.49% and 5.20%. This suggests that the optimal allocation policy is stochastically saturated at lower $p_{\text{s,max}}$levels. Higher equity exposure does not help much for retirees to achieve the desired lifestyle.

$p_{\text{s,max}}$	$\%\Delta p_{\text{s,max}}$	A_R	$\%\Delta A_R$	η_{A_R}	A_S	$\%\Delta A_S$	η_{A_S}
0.5	-	16848.39	-	-	7848.39	-	-
1.0	100	17252.17	2.40	2.40	8252.17	5.14	5.14
1.3	30	17380.97	0.75	2.49	8380.97	1.56	5.20

TABLE 5.5: *Raw and Surplus Restricted AUCs of the Efficient Frontier Comparisons.*

5.3.2 Surplus Expected Withdrawal

While the global metrics of Restricted AUC validate the strategy's overall efficiency, the retiree's utility is best measured by Surplus Expected Withdrawal (EW_S) at specific risk levels, which is the local efficiency of the decumulation strategy.

The elasticity of EW_S with respect to leverage ($p_{\text{s,max}}$) at a given risk level ES_{ref} is computed by

$$\eta_{EW_S}(p_{\max}; ES_{ref}) \approx \frac{\%\Delta EW_S}{\%\Delta p_{\text{s,max}}}\bigg|_{ES=ES_{ref}} \tag{5.5}$$

Table 5.6 presents the elasticity of EW_S when ES_{ref} is fixed at values that are equally spaced between the prudent region Ω_p = [0,300] where spacing is 60. The elasticity of the surplus expected withdrawal is between 0.03 and 0.09. This means that when a retiree increases $p_{\text{s,max}}$by one unit, the gain in EW_S is only between 3% - 9%, which confirms the diminishing marginal utility of leverage locally as well.

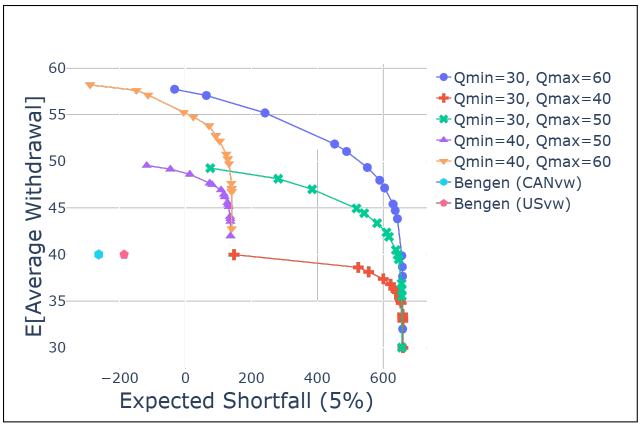

(a) *Six Asset Case:* $p_{s,max} = 0.5$

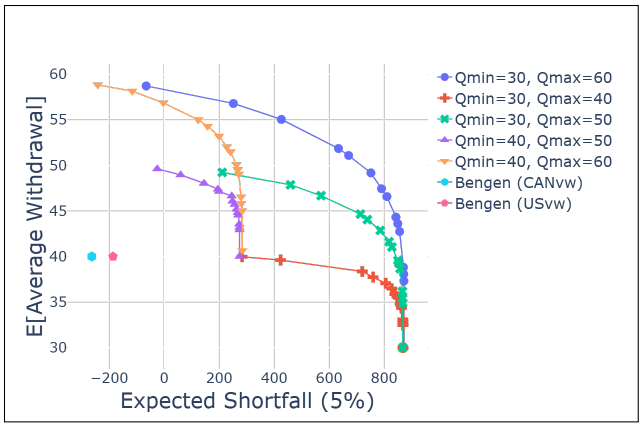

(b) *Six Asset Case:* $p_{s,max} = 1.0$

FIGURE 5.7: *Comparison of withdrawal boundary effect on efficient frontiers when* $p_{s,max} = 0.5$ *and* $p_{s,max} = 1$. *The plot shows that decreasing* q_{max} *lowers the EW without bringing too much benefit of increasing the ES. Increasing* q_{min} *shifts the efficient frontier significantly to the left (hence decreasing ES) without a significant gain in the EW.*

ES_{ref}	0		60		120		180		240		300	
$p_{\text{s,max}}$	EW	η_{EW_S}	EW	η_{EW_S}	EW	η_{EW_S}	EW	η_{EW_S}	EW	η_{EW_S}	EW	η_{EW_S}
0.5	57.54	-	57.09	-	56.57	-	55.95	-	55.21	-	54.34	-
1.0	58.44	0.03	58.12	0.04	57.78	0.05	57.34	0.05	56.89	0.07	56.34	0.08
1.3	58.66	0.03	58.41	0.03	58.15	0.04	57.81	0.06	57.45	0.07	57.02	0.09

TABLE 5.6: *Detailed comparison of* $p_{s,max}$= 0.5, $p_{s,max} = 1$ *and* $p_{s,max} = 1.3$ *for different* ES_{ref} *values in the six-asset case. Linear interpolation is used to obtain the EW values at* ES_{ref} *values equally spaced between* $[0, 300]$ *with spacing interval of 60.*

5.4 The withdrawal range bounds experiment results

We conducted a sensitivity analysis of the efficient frontier with respect to the withdrawal bounds. We fixed $p_{\text{s,max}}$to 1.0 and 0.5, and with each fixed value, we vary the withdrawal lower and upper bounds, and show how the efficient frontier moves when the bounds change. Figure 5.7 shows the results for the six-asset case for the two different $p_{\text{s,max}}$values. We use the default withdrawal bounds of $[30, 60]$ as the base case in the comparison and the main findings are as follows.

- Increasing q_{min} significantly shifts the efficient frontier inward and reduces ES, increasing portfolio longevity risk without a meaningful gain in terms of average expected withdrawal amount.
- Decreasing q_{max} significantly reduces the EW without a meaningful gain in ES, making reducing q_{max} irrational.
- The efficient frontier is more sensitive to the bounds of withdrawal rather than the risky asset exposure (i.e. $p_{\text{s,max}}$). The efficient frontier ES values clusters on the right side for a given q_{min} value, which indicates q_{min} bounds the longevity of the portfolio. The q_{max} bounds the maximum EW. When one choses a high q_{min}, reducing q_{max} only lowers the average EW without reducing the portfolio longevity risk. Merely reducing the q_{max} is an irrational decision.

5.5 Analysis of key assets

We have observed that compared to the two-asset base case, a retiree with access to the six assets can expect to do better: on average they will withdraw more from the retirement portfolio while being exposed to less

tail risk.

It is of interest to better understand the source of this improved performance. To test which assets are key to portfolio performance, we remove assets one at a time and evaluate our performance metrics.

Table 5.7 shows the average effect of dropping each of the non-short-term-bond assets on our performance metrics. From this table, it is clear that removal of the US capitalization-weighted Index or Canadian Equal-Weighted Index have the most significant impact on performance. The correlation structure and summary statistics of the six assets in Table 5.2 also confirm that the US capitalization-weighted index and Canadian equal-weighted index are the least correlated of all the equity assets. This suggests an explanation for the importance of these assets: incorporating these two assets into a DC-plan portfolio allows the retiree to benefit from high returns while mitigating risk through diversification.

We remove the US equal-weighted and Canadian capitalization-weighted assets from the portfolio. Figure 5.8 compares the efficient frontier obtained by using the four assets with the original six asset case when $p_{\text{s,max}}$ is set to 0.5 and 0.1. It is not surprising to see the two efficient frontiers are almost overlapping for both cases.

	USvw	USew	Cvw	Cew	Cltbond
Objective	1930.18	2007.53	2007.41	1859.66	1978.87
Avg Withdrawals	53.06	53.20	53.23	53.08	53.15
ES(5%)	249.45	347.77	345.12	168.58	314.51

TABLE 5.7: *Each column reports average objective function value, withdrawals, and expected shortfall at 5% when one asset is excluded. The reported values are averages across the efficient frontier. We consider* $\kappa \in \{0.01,0.1,0.25,0.4,0.6,0.7,1,1.4,175\}$ *and* $p_{s,max} \in \{0.5,0.6,0.7,0.8,0.9,1.0\}$.

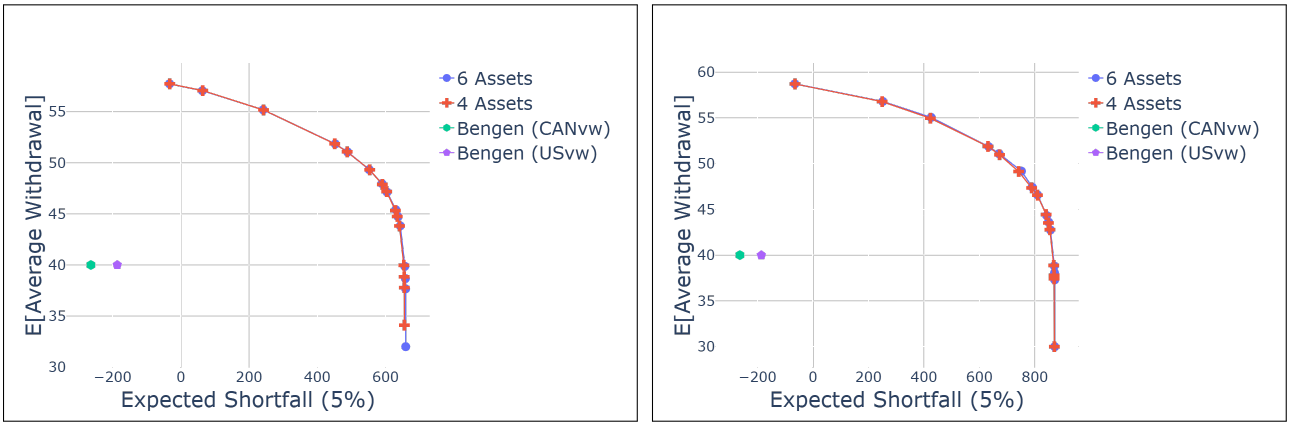

(a) $p_{s,max} = 0.5$ (b) $p_{s,max} = 1.0$

FIGURE 5.8: *Comparison of four assets with six assets when* $p_{s,max}$*is set to 0.5 and 1.0. The USew and Cvw are removed in the four asset case.The efficient frontiers are overlapping.*

5.6 Robustness Testing

Overfitting is a persistent problem in neural network literature. It is especially challenging for financial data, since return series have a low Signal-to-Noise ratio. Standard train/validation splits are not sufficient evidence that a learned strategy reflects market structure rather than artifacts of the training sample. We therefore subject the model to four additional checks: stability under perturbation, out-of-sample performance, out-of-distribution behavior, and sensitivity to the training distribution. These tests cannot confirm that the model has learned genuine market dynamics, but failure on any of them would be strong evidence that it has not.

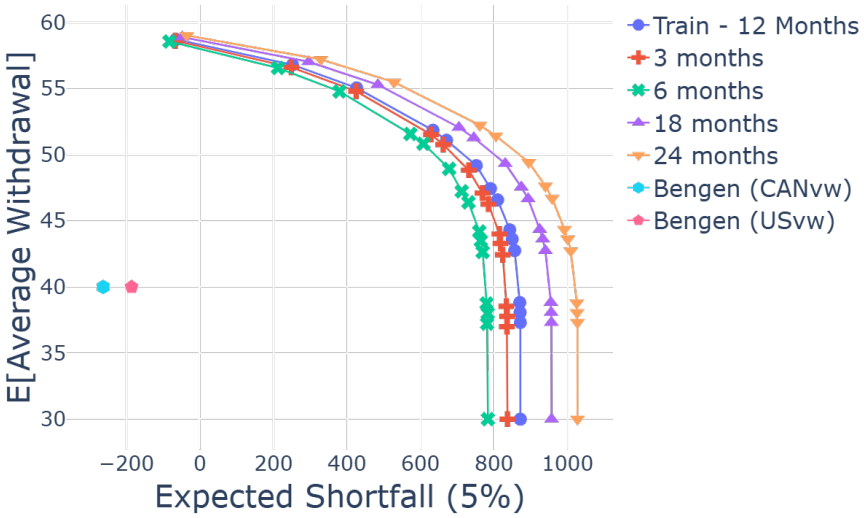

FIGURE 5.9: *Out of Distribution Test. Training block size is 12 months, and we vary the test block size from 3 months to 24 months.*

5.6.1 Stochastic Stability and Convergence Test

To assess the stability of the neural network optimization, we trained two models using two sets of paths generated by two different random seeds, both produced by stationary block bootstrapping of the same historical period. The results show the two efficient frontiers are almost identical cross both training runs (see Figure A.1(a)). The convergence of these two solutsions is an evidence that the learned policy reflects a global optimum rather than artifacts of a particular path realization. This corss-seed consistency is a practical check on solution quality.

5.6.2 Out-of-Sample Test

An out-of-sample test evaluates a model on new data that comes from the same underlying distribution as the training data. We used different random seeds to generate training paths and testing paths (see Figure A.1(b)). The out-of-sample efficient frontiers from the different random seeds are almost overlapping with the original efficient frontier generated by the training set, demonstrating the robustness of the model.

We further extended the stability test to a full out-of-sample test. We generated the set of test paths using a different random seed from the ones we used in the stability test and tested the two models obtained in the stability test using this test set. Figure A.1(c) shows the computed controls are barely changed. This is an evidence to show the learned model is robust.

5.6.3 Out of Distribution Test

We trained the model with a block size of 12 months and then tested the model using block sizes from 3 to 24 months (see Figure 5.9). Larger block size in test data shifts the efficient frontier up and right slightly. Smaller block size in test data shifts the efficient frontier down and left slightly. For the ES=300 use case, the efficient frontier shift is almost negligible. This shows the model is very robust to the out of distribution test.

5.6.4 Training Distribution Sensitivity Test

We trained the model with block sizes of 3 and 6 months and then tested the model using a block size of 12 months. Figure 5.10 shows the efficient frontier shift is small. For the typical case of ES=300, the shift is negligible.

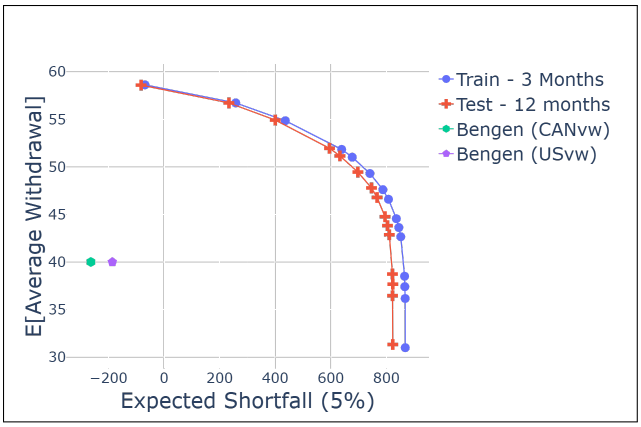

(a) *Six Asset Case: Training block size of 3 months.*

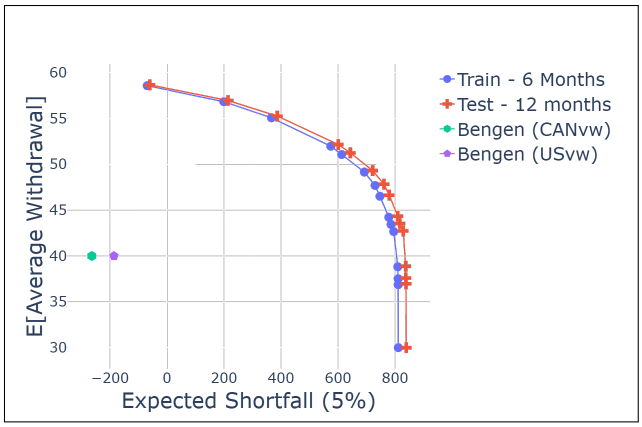

(b) *Six Asset Case: Training block size of 6 months.*

FIGURE 5.10: *Training Distribution Sensitivity Tests. The model is trained using block sizes of 3 and 6 months. Test block size is 12 months.*

5.7 Expanding Window Walkforward Validation Test

In previous sections, we use the stationary block bootstrap resampling to generate different training paths and validation testing paths from historical data in the entire time horizon of 1950 - 2025. In this section, we perform train/test split sequentially using an expanding window walkforward test. This test ensures there is no lookahead data bias.

At timestep t_i, where $i = 0, \ldots, 29$, we will generate training paths by using all the historical data available before t_i to train the withdrawal and allocation neural networks. We then apply the learned withdrawal and allocation strategies at $t_i, \ldots, t_{i+h}$, where h is the number of years we will not retrain the neural networks. At timestep of t_{i+h+1}, we will retrain the neural networks using block bootstrapped paths based on all the historical data available up to and including t_{i+h} and then apply the updated withdrawal and allocation strategies in $t_{i+h+1}, \ldots, t_{i+2h}$ and so on.

We selected the testing period of 1995 - 2024 where there are several major financial events, including the 1997 Asian financial crisis, the 2000 dot-com bubble burst, the 2008-2009 global financial crisis, the 2010-2012 European debt crisis, the 2020 COVID-19 market crash, and the 2021-2023 global inflation surge. We retrain the decmulation strategy neural networks every five years (i.e. $h = 5$). We used the block bootstrap resampling to generate 30 years of paths using data between 1995 and 2024 to study the decumulation stochastic characteristics. We also selected a single path, which is the actual historical real return path to test the strategy's actual performance. We experimented with different $p_{\mathrm{s,max}}$values. Figure 5.11 shows the experiment results. Each row in the plot is for a walkforward test result with a specific $p_{\mathrm{s,max}}$value. The left subplot in each row shows the result of the efficient frontier using the expanding window walkforward test in orange color. To compare it with the previous results, the blue curve on the plot is the efficient frontier generated by using the bootstrapped data of 1950 - 2025. The green curve is the efficient frontier of the walkforward validation set evaluated on the model trained using bootstrapped data from 1950 - 2025. Retirees will be interested in how the strategy works on real data. The right subplots show the efficient frontiers generated by evaluating the original model and the walkforward validation model using a single realization of actual asset returns.

The left plots show the statistical characteristics of each strategy. The walkforward validation curves are all under the curves generated by the model trained using the full return data from 1995-2024. This indicates the lookahead data bias gives an advantage to the model learned using both current and future observataion. The right side plots show that the model learned by using the walkforward validation outperforms the original model learned by using the entire dataset. The test curve using the walkforward model is under the test curve using the original model when the wealth is moderate. Both outperform the original model learned by using the entire dataset.

In all cases, the efficient frontiers are above that from the Bengen's reference portfolios. This clearly shows that the optimal stochastic dynamic decumulation strategy outperforms the Bengen's rule. It is worth

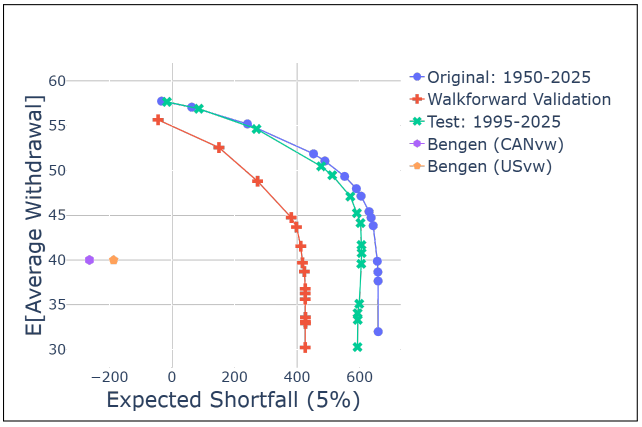

(a) $p_{s,max} = 0.5$, *bootstrapped test data for 30 years*

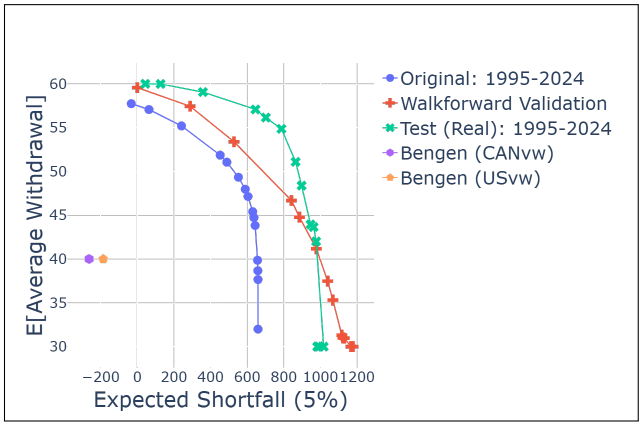

(b) $p_{s,max} = 0.5$, *single path realization*

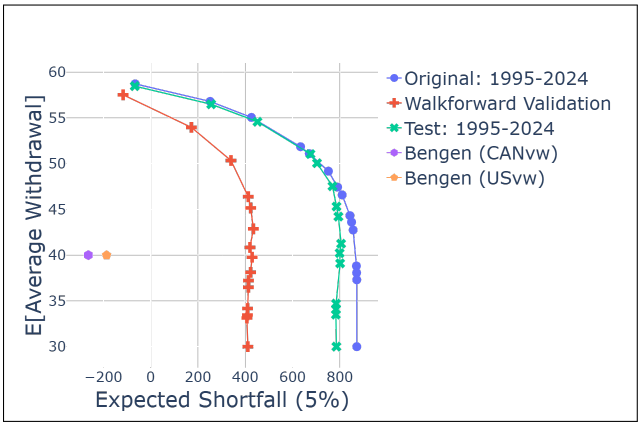

(c) $p_{s,max} = 1$, *bootstrapped test data for 30 years*

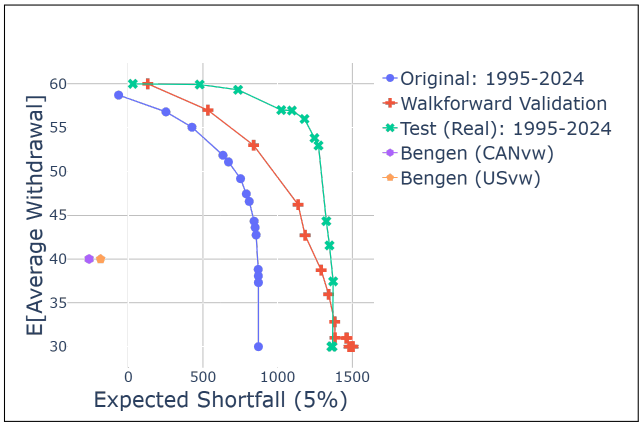

(d) $p_{s,max} = 1$, *single path realization*

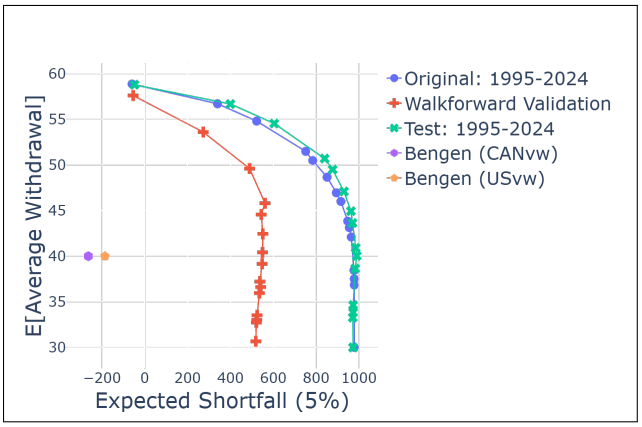

(e) $p_{s,max} = 1.3$, *bootstrapped test data for 30 years*

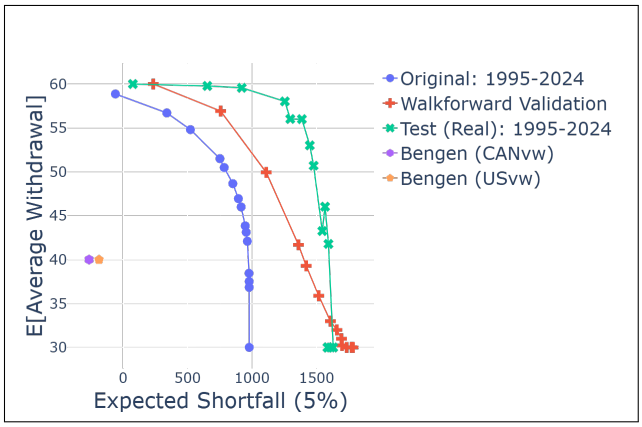

(f) $p_{s,max} = 1.3$, *single path realization*

FIGURE 5.11: *Walkforward Validation Tests. Models are retained every 5 years. Comparison of tests on 256K bootstrapped paths and a single pth realization of the actual returns*

noting that we also have experimented with retraining the models annually and the results are very close to the five-year retraining results using bootstrapped resampling data (see Figure A.2), which shows the model is robust.

6 Conclusion

In this paper, we developed optimal stochastic dynamic decumulation strategies for Canadian retirees, balancing the tradeoff between expected withdrawal (EW) and expected shortfall (ES) of terminal wealth. By using feedforward neural networks with time and current portfolio value as input features, we showed that the withdrawal and asset allocation control functions can be learned without succumbing to the curse of dimensionality. This comes from the fact that we are able to solve the optimal decumulation problem using a data driven NN approach without dynamic programming. In addition, we explicitly exploit the fact that the state variables governing control do not scale with the number of assets. This enabled us to solve, for the first time in this setting, an optimal stochastic six-asset decumulation portfolio.

A central methodological contribution is a data driven NN approach for solving constrained stochastic optimal control using the constraint-feasible output activation function, which transforms a constrained optimization problem into an unconstrained one while guaranteeing feasibility of the resulting controls. We proved that feedforward neural networks equipped with these activation functions can approximate any feasible control to arbitrary accuracy, giving the approach broad applicability beyond this specific decumulation setting.

Our sensitivity analysis, conducted using the newly introduced Restricted Area Under the Curve (RAUC) and Surplus Expected Withdrawal ($\mathrm{EW_S}$) metrics, leads to several practical implications for retirees and plan managers. Leverage offers negligible utility gains relative to the risk it introduces, while capping risky asset exposure at 0.5 costs little in terms of retiree utility. Our results and analysis suggest that conservative exposure limits are largely costless. By contrast, the withdrawal range exerts a far stronger influence on outcomes than the exposure cap, indicating that withdrawal policy design deserves more attention than asset allocation constraints alone. The Canadian equal-weighted and U.S. capitalization-weighted equity indices emerge as the most influential assets, and diversification across six asset classes meaningfully outperforms simpler two-asset portfolios. Our strategies also significantly outperform Bengen's rule, and the trained network remains robust even when retrained only every five years rather than annually.

More broadly, the proposed data-driven neural network framework is not specific to decumulation and can be extended to general optimal stochastic financial decision problems, a direction we intend to pursue in future work.

Appendices

A Additional plots

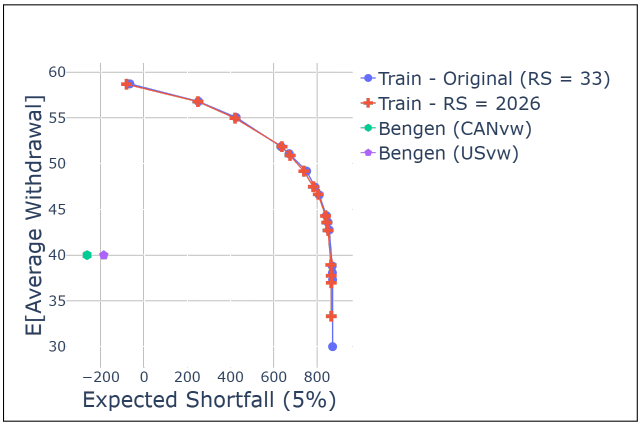

(a) *Stability Test*

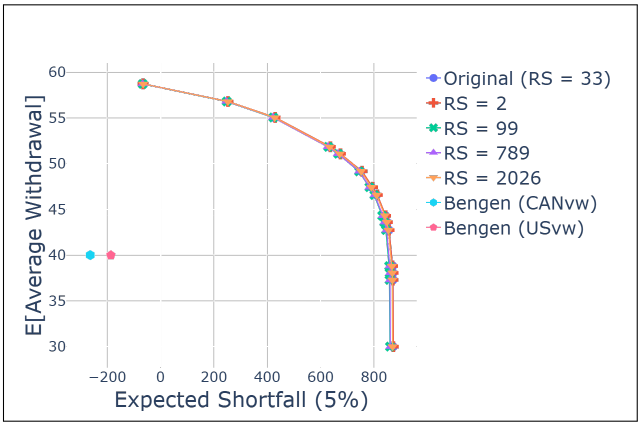

(b) *Out-of-Sample Test*

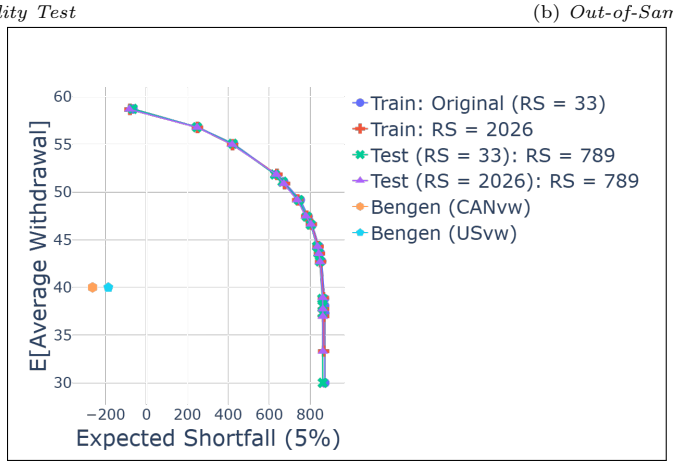

(c) *Full Out-of-Sample Test*

FIGURE A.1: *Various robustness evaluations. Subplot (a) is a Stability Test, where two models are trained using two sets of paths generated by different random seeds. Subplot (b) is the Out-of-Sample test, which involves using different random seeds used to generate training and testing paths. Finally, subplot (c) represents the full Out-of-Sample Test, where different random seeds were used to generate two training sets to train separate models. Both models where then validated against a test set generated by a different random seed. For each subplot, the training block size is 12 months.*

B Additional Tables

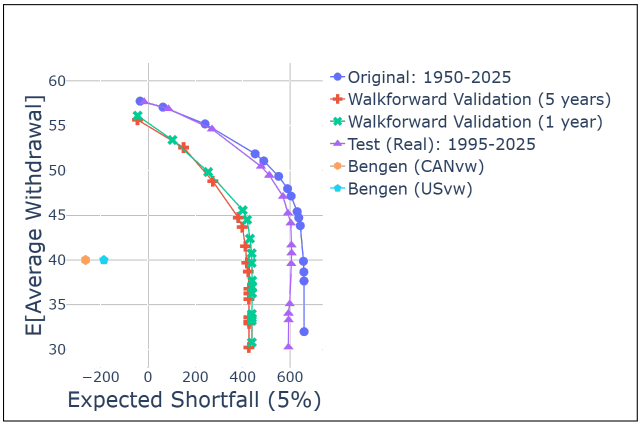

(a) $p_{s,max} = 0.5$, *bootstrapped test data for 30 years*

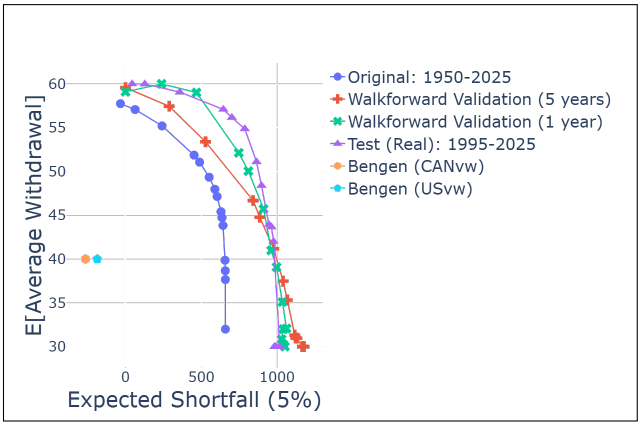

(b) $p_{s,max} = 0.5$, *single path realization*

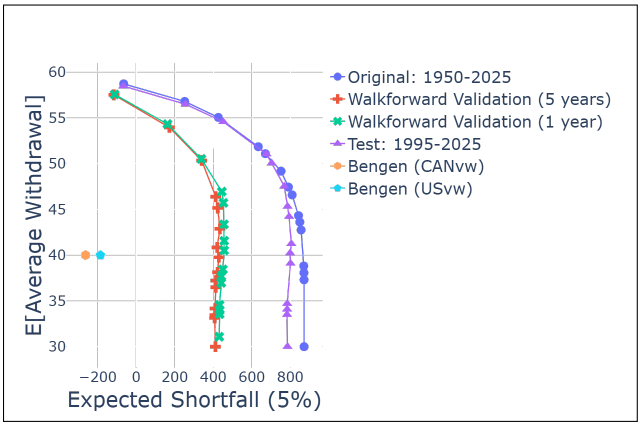

(c) $p_{s,max} = 1$, *bootstrapped test data for 30 years*

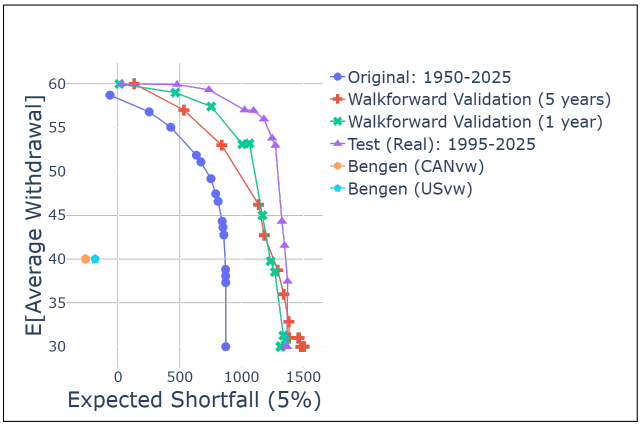

(d) $p_{s,max} = 1$, *single path realization*

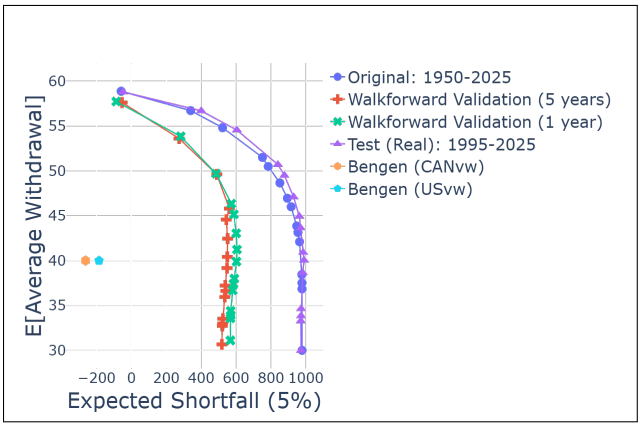

(e) $p_{s,max} = 1.3$, *bootstrapped test data for 30 years*

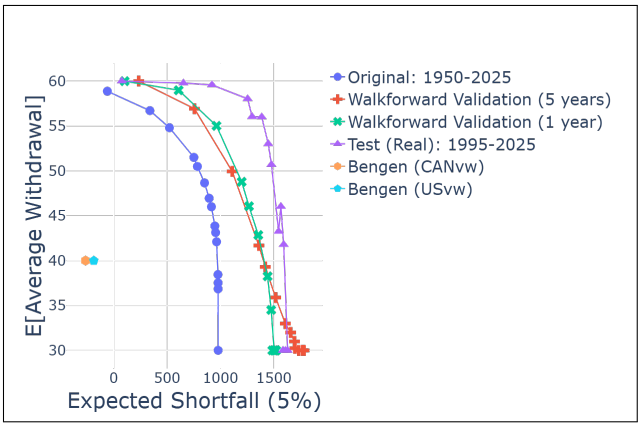

(f) $p_{s,max} = 1.3$, *single path realization*

FIGURE A.2: *Walkforward Validation Tests. Comparing models retained every 5 years and annually. Tests use 256K bootstrapped paths and a single pth realization of the actual returns.*

Expected Shortfall (5%)	USvw, Ctbill	6-asset	diff	relative diff
-100	57.641	58.736	1.095	0.019
-50	57.213	58.631	1.418	0.025
0	56.686	58.325	1.639	0.029
50	56.159	58.018	1.859	0.033
100	55.481	57.711	2.230	0.040
150	54.607	57.404	2.797	0.051
200	53.729	57.098	3.368	0.063
250	52.824	56.791	3.967	0.075
300	51.859	56.301	4.443	0.086

TABLE B.1: *Comparison of the average withdrawals of the 6-asset and USvw, Ctbill frontiers at a list of fixed expected shortfall values. Computed via linear interpolation*

Average Withdrawals	USvw, Ctbill	all	diff
48	429.863	772.433	342.570
49	401.800	747.607	345.807
50	370.194	713.303	343.109
51	332.426	672.903	340.477
52	294.658	626.488	331.830

TABLE B.2: *Comparison of the expected shortfall values of the 6-asset and USvw, Ctbill frontiers at a list of fixed average withdrawal values. Computed via linear interpolation*

Expected Shortfall (5%)	0.5	1	diff	relative diff
-100	57.828	58.736	0.909	0.016
-50	57.777	58.631	0.854	0.015
0	57.533	58.325	0.792	0.014
50	57.189	58.018	0.828	0.014
100	56.712	57.711	0.999	0.018
150	56.179	57.404	1.225	0.022
200	55.647	57.098	1.451	0.026
250	55.083	56.791	1.708	0.031
300	54.325	56.301	1.976	0.036

TABLE B.3: *Comparison of the expected shortfall values of the $p_{max} \in \{1,0.5\}$ frontiers at a list of fixed average withdrawal values. Computed via linear interpolation*

Average Withdrawals	0.5	1	diff
48	588.553	772.433	183.880
49	561.163	747.607	186.444
50	526.900	713.303	186.403
51	489.830	672.903	183.074
52	443.092	626.488	183.396

TABLE B.4: *Comparison of the average withdrawal values of the $p_{max} \in \{1,0.5\}$ frontiers at a list of expected shortfall values. Computed via linear interpolation .*

κ	$p_{max,s}=0.5$ AW	$p_{max,s}=1$ AW	$p_{max,s}=1.3$ AW	$p_{max,s}=0.5$ ES	$p_{max,s}=1$ ES	$p_{max,s}=1.3$ ES
0.01	58.18	58.91	59.09	-449.19	-452.83	-444.34
0.10	57.74	58.72	58.87	-32.93	-65.37	-59.16
0.25	57.07	56.80	56.71	63.46	251.78	339.53
0.40	55.20	55.05	54.81	241.41	426.28	522.43
0.50	53.57	52.85	52.60	348.84	585.75	690.81
0.60	51.86	51.85	51.50	452.64	633.90	751.25
0.65	51.44	51.34	51.11	471.28	660.75	773.27
1.00	49.34	49.18	48.66	552.03	751.79	851.12
1.40	47.97	47.44	46.95	589.60	790.42	894.28
1.75	47.14	46.58	46.00	604.41	809.85	915.85
3.00	45.29	44.40	43.85	630.92	840.48	947.39
3.75	42.69	43.63	43.13	637.36	847.47	954.84
5.00	43.73	42.68	42.10	643.18	853.91	964.04
25.0	39.82	38.94	38.45	656.37	863.31	976.87
50.0	38.95	37.88	37.52	657.69	863.73	977.90
100	30.00	30.00	36.85	656.63	863.22	978.15

TABLE B.5: *Detailed comparison of $p_{max,s}=0.5$, $p_{max,s}=1$ and $p_{max,s}=1.3$ cases with all 6 assets available.*

C Neural approximators of optimal allocation functions

C.1 Preliminaries: Convex Analysis

Definition C.1. *A* polyhedron $P \subset \mathbb{R}^N$ *is a set of the form*

$$P = \{x \in \mathbb{R}^N | Ax \leq b\} \tag{C.1}$$

for some $A \in \mathbb{R}^{M\times N}$ and $b \in \mathbb{R}^M$. A polytope *is a bounded polyhedron.*

We use the following classic convex analysis theorem:

Theorem 1 (Minkowski-Weyl)**.** *(See (Fukuda, 2025, 2.11)) For any $P \subset \mathbb{R}^N$ the following two statements are equivalent:*

1. *P is a polytope.*
2. *P is the convex hull of finitely many extreme points. i.e. $P = \text{conv}\{p_1, \ldots, p_r\}$ for $p_i \in \mathbb{R}^N$.*

Let $\mathcal{P}_N$ denote the set of polytopes in $\mathbb{R}^N$.

Definition C.2. *Given any polytope $P \in \mathcal{P}_N$, let $E(P)$ denote the corresponding finite set of extreme points.*

Definition C.3. *(See (Hiriart-Urruty and Lemaréchal, 2013, pg. 6)) For any $r > 0$, define the r-simplex as:*

$$\Delta^r := \{x \in \mathbb{R}^r : x \geq 0, \sum_i x_i = 1\}. \tag{C.2}$$

Definition C.4. *The affine hull of a set $S \subset \mathbb{R}^n$ is given by*

$$\text{aff } S := \{\sum_{i=1}^{p} \alpha_i x_i : p \in \mathbb{Z}_+, x_i \in S, \sum_i \alpha_i = 1\}. \tag{C.3}$$

Definition C.5. *The relative interior of a convex set $C \subset \mathbb{R}^n$ is the interior in the topology induced by its affine hull. That is:*

$$\operatorname{ri} C := \{x \in C : \exists r > 0 \;\; B(x,r) \cap \operatorname{aff}(C) \subset C\} \tag{C.4}$$

where $B(x,r)$ denotes the unit ball of radius r centered at x.

Theorem 2. *(See (Hiriart-Urruty and Lemaréchal, 2013, pg. 105)) For any convex set C,*

$$\operatorname{cl} C = \operatorname{cl} \operatorname{ri} C \tag{C.5}$$

where cl *denotes the closure of a set.*

Theorem 3. *(See (Hiriart-Urruty and Lemaréchal, 2013, pg. 104)) Let P be a polytope with r extreme points $\{e_1, \ldots, e_r\}$. Then*

$$\operatorname{ri} P := \{\sum_{i=1}^{r} \alpha_i e_i : \alpha_i > 0 \;\; \forall i, \sum_{i=1}^{r} \alpha_i = 1\}. \tag{C.6}$$

C.2 Preliminaries: Functional Analysis

Definition C.6. *(See (Folland, 1999, pg. 43)) Given two measure spaces $(X,\mathcal{M},\mu)$ and $(Y,\mathcal{N},\nu)$, a function $f : X \to Y$ is measurable if $f^{-1}(E) \in \mathcal{M}$ for all $E \in \mathcal{N}$. Call the set of such measurable functions $L(X,Y)$.*

Definition C.7. *Let $(X,\mathcal{M},\mu)$ and $(Y,\mathcal{N},\nu)$ be measure spaces. Additionally assume that Y is a subset of a separable Banach space. If $f : X \to Y$ is measurable, then define*

$$\|f\|_{L_1} = \int \|f\| d\mu, \tag{C.7}$$

where the integral is a Lebesgue integral. Furthermore, define

$$L_1(X,Y) := \{f \in L(X,Y) : \|f\|_{L_1} < \infty\}. \tag{C.8}$$

If we then identify functions that are equal almost everywhere, $L_1(X,Y)$ then becomes a Banach space under the norm $\|\cdot\|_{L_1}$ whenever Y is itself a separable Banach space.

Definition C.8. *A function $f \in L_1(X,Y)$ is called* simple *if it can be written as a a weighted sum of finitely many indicator functions of measurable sets. That is:*

$$f = \sum_{i=1}^{q} y_i I_{E_i} \tag{C.9}$$

for some $E_i \in \mathcal{M}$ and some $y_i \in Y$. Denote the set of such simple functions by $S(X,Y)$.

Theorem 4. *(See (Folland, 1999, pg. 159)) If Y is a separable Banach space, then $S(X,Y)$ is dense in $L_1(X,Y)$.*

Theorem 5. *Suppose C is a closed convex subset of $\mathbb{R}^n$. Then $S(X,C)$ is dense in $L_1(X,C)$.*

Proof. Given $f \in L_1(X,C)$. Apply Theorem 4 to find a sequence $f_n \xrightarrow{L_1} f$ with $\sum_i y_i^n I_{E_i^n} = f_n \in S(X,\mathbb{R}^n)$. If necessary, rewrite the f_n so the E_i^n are disjoint. Let P_C be the projection operator onto C (See (Hiriart-Urruty and Lemaréchal, 2013, pg. 118)). Consider the sequence $\{g_n\}$ with $g_n := P_c \circ f_n = \sum_i P_c(y_i^n) I_{E_i^n}$. Clearly $g_n \in S(X,C)$. We have:

$$\begin{aligned}
\|f - g_n\|_{L_1} &= \int \|f - g_n\|_1 && \text{(C.10)}\\
&= \int \|P_C \circ f - P_c \circ f_n\|_1 && \text{(C.11)}\\
&\leq \int C\|f - f_n\|_1 \;\; P_C \text{ is Lipshitz by Hiriart-Urruty and Lemaréchal (2013)} && \text{(C.12)}\\
&\to 0 && \text{(C.13)}
\end{aligned}$$

as desired.

□

Remark C.1. *We used the 1-norm on $\mathbb{R}^n$ above, but since all norms on Euclidian space produce the same topology, we could have chosen any norm.*

Definition C.9. *Let $C(X,Y)$ denote the continuous functions from X to Y.*

Theorem 6. *(See (Folland, 1999, pg. 245)) $C(X,\mathbb{R})$ is dense in $L_1(X,\mathbb{R})$.*

Theorem 7. *$C(X,\mathbb{R}^n)$ is dense in $L_1(X,\mathbb{R}^n)$.*

Proof. Follows from Theorem 6 via the decomposition of the vector L_1 norm into the sum of scalar L_1 norms. □

C.3 Preliminaries: Softmax

Definition C.10. *For any $r > 0$, define the function* $\text{softmax} : \mathbb{R}^r \to \text{ri}\,\Delta^r$ *by*

$$\text{softmax}(x)_i := \frac{\exp(x_i)}{\sum_{j=1}^{r} \exp(x_j)}. \tag{C.14}$$

Theorem 8. softmax^r *is a surjective map onto* $\text{ri}\,\Delta^r$.

Proof. It suffices to show that for each $p \in \text{ri}\,\Delta^r$, there exists a $z \in \mathbb{R}^r$ such that $\text{softmax}^r(z) = p$. We can achieve this by choosing

$$z_i = \begin{cases} \log\left(\frac{p_i}{p_n}\right) & \text{when } i < n \\ 0 & \text{when } i = n \end{cases}. \tag{C.15}$$

□

Theorem 9. *(See Gao and Pavel (2017))* softmax^r *is Lipshitz continuous.*

C.4 Preliminaries: Neural Networks

Definition C.11. *Let $\sigma : \mathbb{R} \to \mathbb{R}$. Let $m,n > 0$. A* neural-network layer *with activation function σ is a function $f : \mathbb{R}^n \to \mathbb{R}^m$ of the form*

$$f(x)_i = \sigma(\langle x,a_i \rangle + \theta_i) \tag{C.16}$$

for some $\{a_i\} \subset \mathbb{R}^n, \{\theta_i\} \subset \mathbb{R}$. Denote the set of neural network layers from $\mathbb{R}^n$ to $\mathbb{R}^m$ with activation function σ by $\mathbf{NNL}(n,m,\sigma)$

Definition C.12. *An l-layer* neural network *is a composition of l neural-network layers.*

Definition C.13. *Define*

$$\mathcal{M}(n,m,\sigma) := \{f \circ g | g \in \mathbf{NNL}(n,r,\sigma), f \in \mathbf{NNL}(r,m,I) \text{ for some } r > 0\} \tag{C.17}$$

where I denotes the identity function. Call this the set of shallow neural networks.

Theorem 10. *(See Pinkus (1999)[Thm 3.1]) Assume that $\sigma \in C(\mathbb{R},\mathbb{R})$. Then $\mathcal{M}(n,1,\sigma)$ is dense in the set of continuous functions on $\mathbb{R}^n$ in the topology of uniform convergence on compact sets if and only if σ is not a polynomial.*

Theorem 11. *Assume that $\sigma \in C(\mathbb{R},\mathbb{R})$. Then $\mathcal{M}(n,m,\sigma)$ is dense in $C(\mathbb{R}^m,\mathbb{R}^m,\sigma)$ in the topology of uniform convergence on compact sets.*

Proof. Given $f \in C(\mathbb{R}^m,\mathbb{R}^m,\sigma)$, we can find sequences $\{f_i^n\} \subset \mathcal{M}(m,1,\sigma)$ for each coordinate function $f_i \in C(\mathbb{R}^m,\mathbb{R})$ such that $f_i^n \to f_i$ uniformly on compact sets. But then $\{f^n\} \subset \mathcal{M}(n,m,\sigma)$ converges to f uniformly on compact sets as desired. □

C.5 Main Results

Definition C.14. *Let $P \in \mathcal{P}_N$, and let compact $X \subset \mathbb{R}^s$ be compact. Let $r := |E(P)|$, the number of extreme points of the polytope P. Let $p_1, \ldots, p_r$ be these extreme points. Define the map $T : \mathbb{R}^r \to \operatorname{ri} P$ by*

$$T(x) = \sum_{i=1}^{r} \operatorname{softmax}^r(x)_i p_i. \tag{C.18}$$

Note that as the composition of $\operatorname{softmax}^r$ *(which is Lipshitz by Theorem 9) and a linear map, T is Lipshitz. Let C denote its Lipshitz constant.*

1. *Let $V^1 := L_1(X,P)$ where L_1 is defined in Definition C.7.*
2. *Let $V^2 := S(X, P)$ where S is defined in Definition C.8.*
3. *Let $V^3 := S(X, \operatorname{ri} P)$ where* ri *is defined in Definition C.5.*
4. *Let $V^4 := \{T \circ g : g \in S(X,\mathbb{R}^r)\}$.*
5. *Let $V^5 := \{T \circ g : g \in C(X,\mathbb{R}^r)\}$.*
6. *Let $V^6 := \{T \circ g : g \in \mathcal{M}(s,r,\sigma)\}$ where $\mathcal{M}$ is defined in Definition C.13, and where σ is a fixed but arbitrary non-polynomial continuous function.*

Definition C.15. *In the following, when we say $B \subset L_1$ is approximated by $A \subset L_1$, we mean $\operatorname{cl} A \supset B$ in the topology induced by $\|\cdot\|_{L_1}$.*

Theorem 12 (Main Result)**.** *We have that:*

1. *V_1 is approximated by V_2.*
2. *V_2 is approximated by V_3.*
3. *V_3 is approximated by V_4.*
4. *V_4 is approximated by V_5.*
5. *V_5 is approximated by V_6.*

Proof. We have:

1. This follows directly from Theorem 5, which states that the simple functions mapping to P are dense in $L_1(X,P)$.
2. Given a simple function $f = \sum_{i=1}^{q} y_i I_{E_i}$ with each $y_i \in P$, approximate it by $f' = y'_i \sum_{i=1}^{q} I_{E_i}$ with $y'_i \in \operatorname{ri} P$. We have $\|f - f'\|_1 = \sum_{i=1}^{q} \|y_i - y'_i\| \mu(E_i)$. By Theorem 2, we can chose the y'_i to make each $\|y_i - y'_i\|$ arbitrarily small, which proves the result.
3. Given $\sum_{i=1}^{s} y_i I_{E_i} \in V_3$ with each $y_i \in \operatorname{ri} P$, we can first apply Theorem 3 to find $\{\alpha_j^i\}$ such that $\sum_j \alpha_j^i = 1$, $\alpha_j^i > 0$ such that $\sum_j \alpha_j^i p_j = y_i$. We can then apply Theorem 8 to find z_i such that $\operatorname{softmax}^r(z_i) = \alpha_i$. Thus combining these two facts, we have $T(z_i) = y_i$. Then let $g = \sum_{i=1}^{s} z_i I_{E_i} \in S(X,\mathbb{R}^r)$. We have $f = T \circ g$, proving the result.
4. Consider an arbitrary $f \in V_4$. We can write $f = T \circ g$ for some $g \in S(X,\mathbb{R}^r)$. By 6, we can find a sequence $\{g_n\}$ such that $g_n \in C(X,\mathbb{R}^r)$ and $g_n \xrightarrow{L_1} g$. Consider the sequence $\{h_n\}$ with $h_n := T \circ g_n$.

Clearly $\{h_n\} \subset V_5$. We have

$$\int \|f - h_n\| = \int \|T \circ g - h_n\| \tag{C.19}$$

$$= \int \|T \circ g - T \circ g_n\| \tag{C.20}$$

$$\leq C \int \|g - g_n\| \qquad T \text{ is } C\text{-Lipshitz} \tag{C.21}$$

$$\to 0 \tag{C.22}$$

as desired.

5. Consider an arbitrary $f \in V_5$. We can write $f = T \circ g$ for some $g \in C(X,\mathbb{R}^r)$. By Theorem 11 we can find a sequence $\{g_n\}$ such that $g_n \in \mathcal{M}(s,r,\sigma)$ and $g_n \xrightarrow{u} g$. Consider the sequence $\{h_n\}$ with $h_n := T \circ g_n$. Clearly $\{h_n\} \subset V_6$. We have

$$\int \|f - h_n\| = \int \|\text{softmax}^r \circ g - h_n\| \tag{C.23}$$

$$= \int \|T \circ g - T \circ g_n\| \tag{C.24}$$

$$\leq C \int \|g - g_n\| \qquad T \text{ is } C\text{-Lipshitz} \tag{C.25}$$

$$\leq C\mu(X) \sup_{x \in X} \|g(x) - g_n(x)\| \tag{C.26}$$

$$\to 0 \qquad \text{By uniform convergence} \tag{C.27}$$

as desired.

□

Remark C.2. *Consider a dynamic optimal asset allocation problem with compact state space X, where the valid asset allocations are specified by the polytope P. The set $V_1 = L_1(X,P)$ describes all valid controls, including highly discontinuous bang-bang-style controls.*

By Theorem 12 all such valid controls can be approximated arbitrarily well in the L_1 sense by combining softmax *with a shallow neural network. Furthermore, since $V_6 \subset V_1$, each approximator is itself a valid control.*

References

Anarkulova, A., S. Cederburg, M. S. O'Doherty, and R. W. Sias (2025). The safe withdrawal rate: Evidence from a broad sample of developed markets. *Journal of Pension Economics and Finance 24(3)*, 464–500.

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

Bjork, T., M. Khapko, and A. Murgoci (2021). *Time inconsistent control theory with finance applications*. New York: Springer Finance.

Carrick, R. (2011). The case for embracing global markets. *The Globe and Mail*.

Chen, M., M. Shirazi, P. A. Forsyth, and Y. Li (2025). Machine learning and hamilton-jacobi-bellman equation for optimal decumulation: a comparison study. *Journal of Computational Finance 29:1*, 77–118. Working paper, Cheriton School of Computer Science, `https://cs.uwaterloo.ca/~paforsyt/benchmarkNNpaper.pdf`.

Cui, X., D. Li, X. Qiao, and M. Strub (2022). Risk and potential: An asset allocation framework with applications to robo-advising. *Journal of the Operations Research Society of China 10*, 529–558.

DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio? *Review of Financial Studies 22*(5), 1915–1953.

Farago, A. and E. Hjalmarsson (2023). Small rebalanced portfolios often beat the market over long295 horizons. *The Review of Asset Pricing Studies 13*, 307–342.

Folland, G. B. (1999). *Real analysis: modern techniques and their applications*. John Wiley & Sons.

Forsyth, P. and G. Labahn (2026). Numerical methods for optimal decumulation of a dc pension plan. In B. Koraen and D. Roose (Eds.), *Lecture Notes in Computational Science and Engineering (LNCSE)*. Springer.

Forsyth, P. A. (2020). Multi-period mean CVAR asset allocation: Is it advantageous to be time consistent? *SIAM Journal on Financial Mathematics 11:2*, 358–384.

Forsyth, P. A. (2022a). Equal weight vs. capitalization weight indexes. White paper, Cheriton School of Computer Science, University of Waterloo.

Forsyth, P. A. (2022b). 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. and Y. Li (2026). Risk measures for defined contribution (dc) pension plan decumulation. *North American Actuarial Journal*, 1–32.

Forsyth, P. A., K. R. Vetzal, and G. Westmacott (2020). Optimal asset allocation for a DC pension decumulation with a variable spending rule. *ASTIN Bulletin 50:2*, 419–447.

Fukuda, K. (2024). cddlib.

Fukuda, K. (2025). Frequently asked questions in polyhedral computation.

Gao, B. and L. Pavel (2017). On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*.

Han, J., A. Jentzen, and E. Weinan (2018). Solving high-dimensional partial differential equations using deep learning. *PNAS 115*(34), 8505–8510.

Hiriart-Urruty, J.-B. and C. Lemaréchal (2013). *Convex analysis and minimization algorithms I: Fundamentals*, Volume 305. Springer science & business media.

Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science 48*, 1086–1101.

Kratsios, A. and I. Bilokopytov (2020). Non-Euclidean universal approximation. *Advances in Neural Information Processing Systems 33*, 10635–10646.

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

Menoncin, F. and E. Vigna (2017). Mean-variance target based optimisation for defined contribution pension schemes in a stochastic framework. *Insurance: Mathematics and Economics 76*, 172–184.

Motzkin, T., G. Thompson, and R. Thrall (1953). The double description method. *Contributions to the Theory of Games* (24), 51.

Ni, C., Y. Li, and P. Forsyth (2024). Neural network approach to portfolio optimization with leverage constraints: a case study on high inflation investment. *Quantitative Finance 24*(6), 753–777.

OSFI (2022). Canadian Office of the Superintendent of Financial Institutions (OSFI), Registered pension plans (rpp) and other typesof saving plans coverage in Canada, https://www.osfi-bsif.gc.ca/en/oca/oca-factsheets-other-reports/registered-pension-plans-rpp-other-types-savings-plans-coverage-canada.

Pfau, W. (2018). An overview of retirement income planning. *Journal of Financial Counseling and Planning 29(1)*, 114–120. Doi: 10.1891/1052-3073.29.1.114.

Pfeiffer, S., J. R. Salter, and H. E. Evensky (2013). Increasing the sustainable withdrawal rate using the standby reverse mortgage. *Journal of Financial Planning 26:12*, 55–62.

Pinkus, A. (1999). Approximation theory of the mlp model in neural networks. *Acta numerica 8*, 143–195.

Plyakha, Y., R. Uppal, and G. Vilkov (2021). Equal or value weighting? implications for asset pricing tests. In C. Zopounidis, R. Benkraiem, and I. Kalaitzoglou (Eds.), *Financial Risk Management and Modeling*, Chapter 9, pp. 295–347. Springer International.

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

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

Shefrin, H. M. and R. H. Thaler (1988). The behavioral life-cycle hypothesis. *Economic Inquiry 26*, 609–643.

Strub, M., D. Li, X. Cui, and J. Gao (2019). Discrete-time mean-CVaR portfolio selection and time-consistency induced term structure of the CVaR. *Journal of Economic Dynamics and Control 108*. Article 103751 (electronic).

Thinking Ahead Institute (2024). Global pension assets study 2024. https://www.thinkingaheadinstitute.org/content/uploads/2024/02/GPAS-2024.pdf.

Tljaard, B. H. and E. Mare (2021). Why has the equal weight portfolio underperformed and what306 can we do about it? *Quantitative Finance 21:11*, 1855–1868.

van Staden, P., P. Forsyth, and Y. Li (2024). Smart leverage? rethinking the role of leveraged exchange traded funds in constructing portfolios to beat a benchmark. *arXiv preprint arXiv:2412.05431*.

van Staden, P. M., D.-M. Dang, and P. Forsyth (2018). Time-consistent mean-variance portfolio optimization: a numerical impulse control approach. *Insurance: Mathematics and Economics 83*, 9–28.

Vigna, E. (2014). On efficiency of mean-variance based portfolio selection in defined contribution pension schemes. *Quantitative Finance 14*, 237–258.

Vigna, E. (2017). Tail optimality and preferences consistency for intertemporal optimization problems. Working paper no. 502 , Collegio Carlo Alberto, Università Degli Studi di Torino.