Risk Measures for DC Pension Plan Decumulation

Peter A. Forsyth[a] Yuying Li[b]

November 12, 2025

Abstract

Optimal decumulation of a Defined Contribution (DC) pension plan can be viewed as a problem in optimal stochastic control, which requires specification of an objective function, a combination of reward and risk. An intuitive specification of reward is the sum of withdrawals over the retirement period. This paper investigates three tail risk measures for running out of savings in a DC plan decumulation strategy, which includes (i) expected shortfall (ii) linear shortfall and (iii) probability of shortfall. From the perspective of all optimal solutions, we establish that, under suitable regularity assumptions, the set of optimal controls corresponding to all expected reward expected shortfall Pareto efficient frontier curves is identical to the set of optimal controls associated with all expected reward and linear shortfall Pareto efficient frontier curves. To better understand the impact of a chosen risk measure, we compare its optimal controls across all three combinations of risk measures and reward performance criteria, while fixing the values of the risk aversion and wealth/probability level parameters at reasonable levels. This comparison reveals a clear preference for the linear shortfall risk measure, which yields more desirable optimal strategies. From a practical point of view, we show that allowing variable withdrawals has a large effect on reducing risk, compared to dynamic asset allocation.

Keywords: decumulation, stochastic control, risk
JEL codes: G11, G22
AMS codes: 91G, 65N06, 65N12, 35Q93

1 Introduction

Internationally, there is a growing movement to replace Defined Benefit (DB) pension plans with Defined Contribution (DC) plans. A study of the P7 countries[1] reveals that in terms of fraction of total pension assets, DC plans have increased from 37% in 2003 to 58% in 2023 (Thinking Ahead Institute, 2024). In terms of individual countries, Australia has 88% of pension assets in DC plans, while Japan has only 5% of pension assets in DC plans. In the Netherlands, all DB plans will transition to collective DC plans by 2028.[2] The trend towards DC plans seems inevitable, since

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

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

[1]Australia, Canada, Japan, Netherlands, Switzerland, UK, US

[2]"The End of the Dutch Defined Benefit Model A Steeper Euro Swap Curve Ahead," https://www.pimco.com/eu/en/insights/the-end-of-the-dutch-defined-benefit-model-a-steeper-euro-swap-curve-ahead

corporations and governments no longer desire to take on the risk of providing the guarantees implicit in DB plans.

During the accumulation phase of a DC plan, the burden of deciding on an asset allocation usually is relegated to the investor. However, upon retirement, the DC plan holder is faced with an even bigger challenge. During the decumulation stage of a DC plan, the retiree must decide on a withdrawal schedule and an asset allocation. Surveys have revealed that retirees fear running out of savings more than death (Hill, 2016). Consequently, it seems clear that the retiree wants to withdraw as much as possible, but avoid ruin. The decumulation problem has been termed "the nastiest, hardest problem in finance," by William Sharpe (Ritholz, 2017).

While it is often suggested that retirees purchase annuities to reduce the risk of depletion of savings, annuities are not popular with DC plan holders (Peijnenburg et al., 2016). Milevsky and Young (2007a) note that a survey of US retirees suggested that less than 2% of respondents chose to annuitize. In the Canadian context, Carrick (2020) reports that only 4% of retirees with DC plans at a major Canadian insurer planned to annuitize. In fact, MacDonald et al. (2013) suggest that avoidance of annuities may be entirely rational.[3] For example, in the North American context, true inflation protected annuities are virtually unobtainable.

Variable annuities, where the retiree invests in risky assets, but with minimum cash flow guarantees, have been suggested as hedge against longevity risk. These guarantees come in several forms, perhaps the most popular include Guaranteed Minimum Withdrawals (GMWB) which guarantee minimum cash flows over a fixed term, (Bauer et al., 2008; Dai et al., 2008; Chen et al., 2008; Chen and Forsyth, 2008; Monig and Bauer, 2016) and Guaranteed Lifelong Withdrawal Benefits (GLWB), which provide minimum cash flows for life (Piscopo and Haberman, 2011; Azimzadeh and Forsyth, 2015; Horneff et al., 2015; Feng and Yi, 2019; Landriault et al., 2021; Bacinello et al., 2024). However, the fees charged for these products are controversial (Monig and Bauer, 2016; Forsyth and Vetzal, 2014). Before the financial crisis of 2008 (GFC), many academics argued that insurance companies were undercharging for these guarantees, in the sense that the fees could not cover the cost of hedging. This turned out to be the case, since some insurance companies took massive writedowns on their portfolio of variable annuities after the GFC, see Forsyth (2016). More recently, it appears that the total fees charged by insurance companies (including fees charged on the underlying assets) are much larger than the no-arbitrage value, indicating that perhaps variable annuities are not so attractive to retirees (Forsyth, 2016).

Another product which incorporates a longevity hedge is a modern tontine (Donnelly et al., 2014; Milevsky and Salisbury, 2016; Chen et al., 2019; Fullmer and Sabin, 2019; Chen and Rach, 2022; Forsyth et al., 2024). Existing DC plans which incorporate modern tontine features include the variable annuity funds offered by TIAA[4], the University of British Columbia Pension Plan[5], and the Australian QSuper fund[6].

The modern tontine which is most closely related to the work in this paper is the individual tontine account (Fullmer and Sabin, 2019). This product works as follows: the retiree makes an irrevocable investment in a pooled fund. If the investor passes away, their portfolio is redistributed among the surviving members of the fund. For the survivors, they receive mortality credits as well as the returns from their own investment portfolio. The retiree maintains full control over how their portfolio is invested.

Unlike an annuity, a tontine offers no guaranteed payouts, since returns depend on both market performance and the survival of other participants. This added risk, has the following advantage:

[3] See also *"When do you need insurance?"* `https://donezra.com/217-when-do-you-need-insurance/`

[4] `https://www.tiaa.org/public/`

[5] `https://faculty.pensions.ubc.ca/`

[6] `https://qsuper.qld.gov.au/`

for the same initial investment, the expected payouts from a tontine are typically higher than those of a traditional annuity.

From a mathematical point of view, an individual tontine account can be modelled as a traditional DC decumulation account with a tontine overlay (Forsyth et al., 2024). However, these products are still in their infancy [7]. True, actuarially fair, individual tontine accounts do not seem to be available at this time.

An extensive study of decumulation strategies can be found in Bernhardt and Donnelly (2018). As an example, typical wealth management advice given to retirees is usually some variant of the ubiquitous *4% rule* (Bengen, 1994). This rule suggests that retirees should (i) invest in a portfolio of 50% bonds and 50% equities, rebalanced annually and (ii) withdraw 4% of the initial capital each year (adjusted for inflation). We can consider that this advice is given to a 65-year old retiree, who wants to be sure that they do not run out of savings if he/she lives to age 95.[8]

This advice is justified on the basis of historical rolling 30 year periods, using US data. A retiree following this advice would never have run out of savings over any of these rolling thirty year periods. Various adjustments to this rule have been suggested many times, see e.g. Guyton and Klinger (2006). However, both the advice and historical tests can be criticized. Rolling thirty-year periods obviously have very high correlations. Use of constant weight stock allocation is somewhat simplistic, as is use of a constant (in real terms) withdrawal rate. In fact Irlam (2014) used dynamic programming methods to conclude that deterministic (i.e. glide path) allocation strategies are sub-optimal.[9] More recently, Anarkulova et al. (2025) suggest that the safe withdrawal rate might be much lower than the the 4% rule. In contrast to rolling historical periods, Anarkulova et al. (2025) use block bootstrap resampling to test withdrawal strategies. We will also use bootstrap resampling to test our results in this paper.

Nevertheless, the *four per cent rule* has seen wide adoption since the original publication over thirty years ago, and can be regarded as the default advice.

Contrary to commonly held beliefs, it appears that retirees are somewhat flexible in annual spending. A survey in Bannerje (2021) indicates that retirees actually adjust their lifestyle (i.e. what are perceived as fixed expenses) to match their cash flows.

In fact, recent surveys indicate that, if anything, many retirees underspend on the basis of their financial assets (Rappaport, 2019; Ackerly et al., 2021). Browning et al. (2016) suggests that these assets are being held as a reserve against unexpected medical expenses. However, Canada has a comprehensive public health care system, yet Hamilton (2001) finds that senior Canadian couples 85 and older either save or give away about 25% of their income.

All these facts indicate that we should allow some flexibility in withdrawals from pension savings, in order to ameliorate return risk across time. Various rules have been proposed to allow variable withdrawals. Bengen (2001) updates his original work on the four per cent rule, and now suggests fixed percentage withdrawals, with a floor and ceiling. Bengen (2001) concludes that these withdrawal rules can increase the safe withdrawal rate compared to the original *four per cent* rule. Variable withdrawals are also advocated in Pfau (2015). Guyton and Klinger (2006) explore various heuristics governing portfolio allocation, withdrawals, and caps on inflation adjustments. Another thread of the literature focuses on withdrawal rules which depend on the current state of the investment account and and the current life expectancy of the retiree (Dus et al., 2005; Horneff et al., 2008; Waring and Siegel, 2015; Forsyth et al., 2022).

[7] Saskatchewan Co-operative Superannuation Society Variable Payment Life Annuity (VPLA): `https://www.benefitsandpensionsmonitor.com/pensions/retirement-planning/a-new-annuity-product-for-dc-plan-members/386168`

[8] The probability of a 65-year old Canadian male living to age 95 is about 0.13.

[9] A constant weight strategy is trivially deterministic.

Perhaps the most rigorous approach to the decumulation problem is to formulate this as a problem in optimal stochastic control. The controls in this case, are (i) the asset allocation, i.e. the stock/bond split and (ii) the withdrawal amounts (real) per year, subject to maximum and minimum constraints.

We note that the vast majority of the literature on decumulation fundamentally uses a variety of heuristic rules, which are then tested using Monte Carlo simulation (based on assumed parametric market models) or on historical data. However, once we decide to use an optimal control approach, it is unnecessary to consider heuristic rules. The optimal control, by definition is the best possible strategy which optimizes the objective function (which encapsulates the risk-reward tradeoff). While this is tautologically true, there may be other characteristics of the optimal control, which may not be appealing to retirees. A key consequence of this approach is that it compels the retiree (or advisor) to explicitly formulate the decumulation strategy in terms of a well-defined objective function and a set of appropriate constraints. It is therefore crucial to understand how the choice of risk measure and constraints influences the resulting optimal control, which constitutes one of the main goals of this work. Such understanding can help avoid reliance on suboptimal heuristic strategies and lead to more principled and effective decision-making in retirement planning.

One possible objective function in the decumulation problem is a utility function, combining the withdrawals and final portfolio value. However, it seems clear (from the popularity of the four per cent rule), that investors prefer to delineate the trade off between risk (running out of savings) and reward (maximizing withdrawals).

Since we essentially consider the same problem as formulated by Bengen (1994), the obvious measure of reward is the expected sum of the withdrawals (EW), inflation adjusted, over a 30 year period. However, the choice of risk measure is not so clear. Clearly retirees are primarily concerned with running out of savings, which implies that we should be focused on left tail measures of risk.

One main goal of this paper is to carry out, in the context of decumulation, a thorough investigation of the following reasonable tail risk measures in terms of portfolio value at year 30:

- Expected shortfall ES, i.e., the mean of the worst α fraction of the outcomes. Typically $\alpha = .05$.
- Linear shortfall LS, i.e. weighting negative portfolio values linearly. This is also known as the first partial moment, or mean excess loss (Dus et al., 2005).
- Probability of final portfolio value being negative (PS).

We first formalize an equivalence, from the perspective of all optimal solutions, between expected withdrawal reward and expected shortfall risk (EW-ES) and expected withdrawal reward and linear shortfall risk (EW-LS) efficient frontiers. Our main theoretical result is that, under some regularity conditions, the set of optimal controls associated with all expected reward EW and expected shortfall ES Pareto efficient frontier curves is identical to the set of optimal controls for all expected reward EW and linear shortfall LS Pareto efficient frontier curves. Note a main difference between EW-ES and EW-LS is in the parameter which specifies tail-risk level. In the EW–LS formulation, this parameter specifies an explicit target for terminal wealth, while in the EW–ES formulation, it defines a probability level associated with the expected shortfall constraint. Hence the equivalence theoretical result implies that the choice between EW-ES and EW-LS is essentially a preference in specifying probability level or wealth level when measuring tail risk. We further compare the efficient frontiers generated using all three risk measures above when a reasonable parameter for probability/wealth level.is given. We investigate the impact of the risk measure by examining performance of its optimal controls in all three reward and risk measure combinations. We calibrate a parametric stochastic model for stocks and bonds based on almost a century of data. We solve the optimal control problem via dynamic programming using the parametric model. The controls are tested out-of-sample, using block bootstrap resampling of historical data (Politis and Romano, 1994; Cogneau and Zakalmouline, 2013; Dichtl et al., 2016; Anarkulova et al., 2022; 2025).

Our computational investigation leads to the following practically relevant conclusions:

- We recommend use of expected total withdrawals (as a measure of reward) and linear shortfall LS as a measure of risk when computing optimal decumulation strategies. Linear shortfall is an excellent practical measure of tail risk. Linear shortfall LS is (i) trivially time consistent (ii) weights shortfall[10] (iii) is close to optimal in terms of expected shortfall and probability of shortfall (iv) has an intuitive interpretation and (v) has robust performance in out of sample bootstrap resampling tests.

- Using Linear Shortfall as a risk measure, our example computations suggest that that the variable withdrawal strategy is more important than the asset allocation strategy, in achieving optimal results. Consequently, the exact solution to the optimal control problem can be approximated by a set of simple rules, which are easy to implement, without an appreciable loss of efficiency.[11]

- The optimal controls based on the LS risk measure are robust, as determined by out-of-sample block bootstrap resampling of historical data. This result can be explained by the low sensitivity to the asset allocation strategy, and the *bang-bang* control nature of the optimal withdrawals.

2 Problem Setting

Spending rules (such as the four per cent rule) are clearly popular with retirees. It is interesting to note the following quotation from (Anarkulova et al., 2025)

> *"Current retirement spending practices demonstrate a revealed preference for spending rules over annuitization, such that the efficacy of spending rules is an important issue. ... Obtaining reliable, quantitative evidence on the 4% rule and alternative withdrawal rates is of critical importance given their widespread use."*

Due to its wide acceptance in wealth management, we consider the scenario discussed in (Bengen, 1994). We consider a 65-year old retiree who desires fixed minimum annual (real) cash flows over a 30 year time horizon. We also impose a cap on maximum withdrawals in any year. From the CPM2014 table from the Canadian Institute of Actuaries[12], the probability that a 65-year old Canadian male attains the age of 95 is about 0.13. However, use of a 30 year time horizon is considered a prudent test for having a low probability of running out of savings. In addition, note that we will not mortality weight future cash flows, as is done when averaging over a population for pricing annuities. Note that the probability that a 65 year old Canadian male will attain the age of 87 is about 50%. This has the consequence that the mortality weighted cash flows of an 87 year old (ignoring discounting) would be reduced by one half. However, being alive is a binary state, i.e.

[10] Being short \$100,000 is worse than being short \$1.

[11] This is primarily due to the fact that the optimal withdrawal policy is a bang-bang control (Forsyth, 2022).

[12] `www.cia-ica.ca/docs/default-source/2014/214013e.pdf`.

either alive or dead. If alive, the retiree needs the full amount of the minimum required cash flows, not one half. Hence, a conservative approach is to assume living to the age of 95. This is consistent with the Bengen (1994; 2001) scenario, which is popular with retirees. Since we are motivated by the problem of decumulation strategies for individual retirees, we will not mortality weight future cash flows, as is done when averaging over a population for pricing annuities.

Note that this approach (assumption of a relatively long retirement) is consistent with the advice in Pfau (2018)

> *"Play the long game. A retirement income plan should be based on planning to live, not planning to die. A long life will be expensive to support, and it should take precedence over death planning..."*

Since we allow investing in risky assets, with a minimum cash withdrawal each year, it is possible to exhaust savings. One possibility to address this issue would be to then cease withdrawing from the portfolio when the portfolio value becomes negative. However, this approach would not penalize running out of cash sooner rather than later.

Consequently when the portfolio value is negative, we continue to withdraw cash from the portfolio, which is equivalent to borrowing cash. This debt accumulates at the borrowing rate. The fact that the portfolio can become negative, and the required cash flows can add to debt, means that any tail risk measure will penalize these states. Hence, the optimal stochastic control will find strategies which make these states as unlikely as possible.

We can view the accumulation of debt, in the event that the portfolio has negative wealth, purely as a mathematical device, which penalizes insolvent states.

However, this also has a very simple practical interpretation. Essentially, we are assuming that the investor has other assets, e.g. real estate, which can be used as a hedge of last resort. In practice, accumulated debt due to exhausting savings could be funded using a reverse mortgage, with real estate as collateral (Pfeiffer et al., 2013).

Note that real estate is not fungible with financial assets, except in extreme circumstances. This mental bucketing of assets is a common tenet of behavioral finance (Shefrin and Thaler, 1988). With respect to real estate, if investment performance is strong or the retiree passes away early, the property can serve as part of the bequest.

Consequently, the 65 year old retiree plans to live to the age of 95, which will require minimum cash flows each year. The decumulation plan induces a quantifiable measure of risk. As long as the risk can be hedged with an existing real estate asset, then the plan is considered reasonable. Since there is only a 0.13 probability that the 65 year old Canadian male will attain the age of 95, then, it is likely that the retirement assets (and the real estate asset) can serve as a bequest.

What about the case of extreme longevity? One idea would be to simply extend the retirement planning horizon to the largest possible age in a mortality table. However, this would mean that the minimum cash flows would have to be reduced significantly (compared to the standard 30 year horizon). Since extreme longevity has a small probability of occurring, it would seem sensible to purchase insurance against this risk, since this should be inexpensive during the early years of retirement.

Insurance products which hedge longevity risk include variable annuities and modern tontines, as discussed in the introduction. However, currently, the total fees being charged for variable annuities are not attractive (Forsyth, 2016), and true, fair, individual tontine accounts (Fullmer and Sabin, 2019) do not seem to be available at the current time.

Another possibility is to consider a two stage approach, whereby the portfolio horizon is based on a high probability of survival, i.e. 15-20 years. At the end of this horizon, the retiree may switch to an annuity, or continue managing the decumulation process (Dus et al., 2005; Milevsky and Young, 2007b; Habib et al., 2020; Forsyth, 2021).

However, our main focus in this paper is to compare risk measures for decumulation, assuming that the retiree does not wish to annuitize, for a fixed time horizon. Optimal switching to an annuity is beyond the scope of this paper.

2.1 Notation, Formulation

The investor has access to two funds: a stock index and a constant maturity bond index. At any instant in time t, let the *amount* invested in the stock index fund be denoted by $S_t \equiv S(t)$, and similarly the amount invested in the bond index is denoted by $B_t \equiv B(t)$. These amounts are real, i.e. inflation adjusted. The total (real) value of the portfolio W_t is then

$$W_t = S_t + B_t \ . \tag{2.1}$$

For any time dependent function $g(t)$, we use the notation

$$g(t^+) \equiv \lim_{\epsilon \to 0^+} g(t+\epsilon) \ \ ; \ \ g(t^-) \equiv \lim_{\epsilon \to 0^+} g(t-\epsilon) \ \ . \tag{2.2}$$

Consider a set of discrete withdrawal/rebalancing times $\mathcal{T}$,

$$\mathcal{T} = \{t_0 = 0 < t_1 < t_2 < \ldots < t_M = T\}, \tag{2.3}$$

where T is the investment horizon. For ease of notation, we assume that $t_i - t_{i-1} = \Delta t = T/M$ is constant.

At each rebalancing time $t_i, i = 0, \ldots, M-1$, the investor first (i) withdraws an amount of cash $\mathfrak{q}_i$ from the portfolio and then (ii) rebalances the portfolio. More precisely

$$W(t_i^+) \ \ = \ \ W(t_i^-) - \mathfrak{q}_i \ . \tag{2.4}$$

Denote the state of the system at each time by $\mathcal{X}(t), t \in [0,T]$. Informally, the state can be regarded as the information necessary to model the control from time t onwards (Powell, 2025).

Let the rebalancing control $\mathfrak{p}(\mathcal{X}(t_i^-))$ be the fraction in stocks after withdrawals, then,

$$\begin{aligned} S(t_i^+) \ &= \ \mathfrak{p}_i(\mathcal{X}(t_i^-)) W(t_i^+) \\ & \quad \ \mathfrak{p}_i(\mathcal{X}(t_i^-)) \equiv \mathfrak{p}(\mathcal{X}(t_i^-), t_i) \\ B(t_i^+) \ &= \ W(t_i^+) - S(t_i^+) \ . \end{aligned} \tag{2.5}$$

We can regard the amount withdrawn $\mathfrak{q}_i(\cdot)$ as an additional control i.e. $\mathfrak{q}_i(\mathcal{X}(t_i^-)) = \mathfrak{q}(\mathcal{X}(t_i^-), t_i)$. Note we make the implicit assumption that the optimal controls are of feedback form, i.e. only a function of the state and time.

Based on the parametric SDE model for (S_t,B_t) in Appendix A and Forsyth (2022), we will assume in the following that $\mathcal{X}(t) = (S(t), B(t)), t \in [0,T]$, with the realized state of the system denoted by $x = (s,b)$. More generally, of course, it may be necessary to include other variables to define the state (e.g. *lifting the state space* to include path dependent variables).[13]

In the special case that there are no transaction costs $\mathfrak{q}_i(\cdot) = \mathfrak{q}_i(W_i^-)$ and $\mathfrak{p}_i(\cdot) = \mathfrak{p}(W_i^+)$, i.e. the amount withdrawn is only a function of total wealth before withdrawals, and the rebalancing

[13] A classic example is the pricing of an Asian option, which depends of the observed average stock price A_t. If the stock price S_t follows GBM, then the state space for an Asian option is lifted to (S_t, A_t).

fraction is only a function of wealth after withdrawals. Note that it is straightforward to include transaction costs, but if typical costs for ETFs are included, this has a very small impact on the controls (Dang and Forsyth, 2014).

The control at time t_i is given by $(\mathfrak{q}_i(\cdot), \mathfrak{p}_i(\cdot))$, where $(\cdot)$ denotes the control as a function of its state. We specify feasibility of control by prescribing the set of admissible *values* of the controls by $\mathcal{Z}$, i.e.,

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

where

$$\mathcal{Z}_{\mathfrak{q}}(W_i^-, t_i) = \begin{cases} [\mathfrak{q}_{\min}, \mathfrak{q}_{\max}] & t_i \in \mathcal{T} \ ; \ t_i \neq t_M \ ; \ W_i^- \geq \mathfrak{q}_{\max} \\ [\mathfrak{q}_{\min}, \max(\mathfrak{q}_{\min}, W_i^-)] & t_i \in \mathcal{T} \ ; \ t \neq t_M \ ; \ W_i^- < \mathfrak{q}_{\max} \\ \{0\} & t_i = t_M \end{cases} , \tag{2.7}$$

$$\mathcal{Z}_{\mathfrak{p}}(W_i^+, t_i) = \begin{cases} [\mathfrak{p}_{\min}, \mathfrak{p}_{\max}] & W_i^+ > 0 \ ; \ t_i \in \mathcal{T} \ ; \ t_i \neq t_M \\ \{0\} & W_i^+ \leq 0 \ ; \ t_i \in \mathcal{T} \ ; \ t_i \neq t_M \\ \{0\} & t_i = t_M \end{cases} . \tag{2.8}$$

These expressions encapsulate the following constraints:

- Setting $\mathfrak{p}_{\min} = 0$, $\mathfrak{p}_{\max} = 1$ (our base case) enforces no shorting, no leverage (assuming solvency, i.e., when $W_i^+ > 0$),
- Maximum $\mathfrak{q}_{\max}$ and minimum $\mathfrak{q}_{\min}$ withdrawal constraints,
- In the case of insolvency $W_i^+ < 0$, trading ceases and debt accumulates at the borrowing rate,
- At $t = t_M$, all stocks are liquidated no withdrawals $\mathfrak{q}_M = 0$,
- If $W_i^- < \mathfrak{q}_{\max}$, the investor attempts to avoid insolvency, but always withdraws at least $\mathfrak{q}_{\min}$.

Recall that we assume that the retiree can finance the debt using other assets, e.g. a real estate hedge of last resort. At first sight it might seem appropriate to simply cease withdrawals if insolvent. However, by assumption, the retiree needs a minimum cash flow of $\mathfrak{q}_{\min}$ each year. Therefore, we penalize any set of controls which causes the retiree to exhaust his savings (and access the assumed real estate hedge) in order to fund the minimum cash flows. Note that allowing debt to accumulate also penalizes early insolvency compared to late insolvency.

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

$$\mathcal{A} = \left\{ (\mathfrak{q}_i, \mathfrak{p}_i)_{0 \leq i \leq M} : (\mathfrak{p}_i, \mathfrak{q}_i) \in \mathcal{Z}(W_i^-, W_i^+, t_i) \right\} . \tag{2.9}$$

For notational simplicity, we denote a dynamic control by $\mathcal{P}$, and an admissible control $\mathcal{P} \in \mathcal{A}$ can be written as

$$\mathcal{P} = \{(\mathfrak{q}_i(\cdot), \mathfrak{p}_i(\cdot)) \ : \ i = 0, \ldots, M\} \ . \tag{2.10}$$

We also define $\mathcal{P}_n \equiv \mathcal{P}_{t_n} \subset \mathcal{P}$ as the tail of the set of controls in $[t_n, t_{n+1}, \ldots, t_M]$, i.e.

$$\mathcal{P}_n = \{(\mathfrak{q}_n(\cdot), \mathfrak{p}_n(\cdot)), \ldots, (\mathfrak{q}_M(\cdot), \mathfrak{p}_M(\cdot))\} \ . \tag{2.11}$$

For notational completeness, we also define the tail of the admissible control set $\mathcal{A}_n$ as

$$\mathcal{A}_n = \left\{ (\mathfrak{q}_i, \mathfrak{p}_i)_{n \leq i \leq M} : (\mathfrak{q}_i, \mathfrak{p}_i) \in \mathcal{Z}(W_i^-, W_i^+, t_i) \right\} , \tag{2.12}$$

so that $\mathcal{P}_n \in \mathcal{A}_n$.

3 Risk and reward

We now define the specific risk and reward measures used in the stochastic optimal control formulations.

3.1 Reward

Define $E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}[\cdot]$ as the expectation conditional on the observation at time t_0^-, state $\mathcal{X}_0^-$, under control $\mathcal{P}_0$. We then define reward as

$$\mathrm{EW}(\mathcal{X}_0^-, t_0^-) \;=\; E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}\left[\sum_{i=0}^{M} \mathfrak{q}_i\right] , \tag{3.1}$$

which is the total expected withdrawals in $[0,T]$, with $t_M = T$. We will use EW as the reward measure in all cases. Note that $\mathfrak{q}_i$ is inflation adjusted and that we do not discount the future cash flows. We view this as a conservative approach and is consistent with the Bengen (1994) scenario.

3.2 Risk

Although the selection of a reward measure is relatively clear, identifying a suitable measure of risk is more subtle. We examine the following reasonable candidates.

PS We define PS risk as the probability of shortfall w.r.t. a terminal wealth level $\mathbb{W}$,

$$\mathrm{PS}(\mathcal{X}_0^-, t_0^-) = Prob[W_T < \mathbb{W}] = E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}[\mathbf{1}_{W_T < \mathbb{W}}] . \tag{3.2}$$

Usually, $\mathbb{W}$ is zero, i.e., we are concerned with running out of cash. We want to *minimize* PS risk.

LS Linear shortfall

$$\mathrm{LS}(\mathcal{X}_0^-, t_0^-) \;=\; E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}[\min(W_T - \mathbb{W}, 0)] . \tag{3.3}$$

Note that PS risk does not differentiate bad outcomes. Clearly, being short 1\$ is not as bad as being short 1000\$. LS weights the bad outcomes. Since LS is defined in terms of final wealth, not losses, we want to *maximize* LS risk measure.

ES ES is the mean of the worst α fraction of outcomes. A common choice is $\alpha = .05$. More precisely, let W_T be the wealth associated with $\mathcal{P}_0{}^{\mathcal{X}_0^-, t_0^-}$

$$\begin{aligned} \mathrm{ES}(\mathcal{X}_0^-, t_0^-) &= E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}\left[\frac{W_T \mathbf{1}_{W_T < \mathbb{W}}}{\alpha}\right] \\ &\text{subject to } \left\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}[\mathbf{1}_{W_T < \mathbb{W}}] = \alpha. \right. \end{aligned} \tag{3.4}$$

We want to *maximize* ES risk measure.

One of the main goals of this paper is to compare and contrast these different reward-risk combinations, both mathematically and computationally.

3.3 Summary of Acronyms

For future reference, Table 3.1 lists the acronyms used in this paper.

Acronym	Description
EW (expected withdrawals)	$E[\sum_{i=0}^{M} \mathfrak{q}_i]$
PS (probability of shortfall)	$E[\mathbf{1}_{W_T < \mathbb{W}}]$
LS (linear shortfall)	$E[\min(W_T - \mathbb{W}, 0)]$
ES (expected shortfall)	$E\left[\frac{W_T \mathbf{1}_{W_T < \mathbb{W}}}{\alpha}\right]$ s.t. $E[\mathbf{1}_{W_T < \mathbb{W}}] = \alpha$

TABLE 3.1: *Definition of acronyms.*

4 Pareto points

We will use a scalarization technique to determine Pareto optimal points for the multi-objective problems balancing risk and reward. As an example consider problem EW-PS. Informally, given an scalarization parameter $\kappa > 0$, we seek the optimal control $\mathcal{P}_0$ that maximizes

$$\mathrm{EW}(\mathcal{X}_0^-, t_0^-) - \kappa \ \mathrm{PS}(\mathcal{X}_0^-, t_0^-) \ . \tag{4.1}$$

Varying κ traces out an efficient frontier in the (EW, PS) plane. For any fixed value of PS, the corresponding point on the efficient frontier is the largest possible value of EW.

4.1 PS, LS

We solve optimal control problem for weighted reward and risk combinations, e.g., EW-PS, EW-LS. To be precise, for each reward and risk parameter pair, we define the function $G(W_T, \mathbb{W})$ below,

$$\mathrm{PS} \ : \ G_{PS}(W_T, \mathbb{W}) = -\mathbf{1}_{W_T < \mathbb{W}} \tag{4.2}$$

$$\mathrm{LS} \ : \ G_{LS}(W_T, \mathbb{W}) = \min(W_T - \mathbb{W}, 0) \ , \tag{4.3}$$

where $\mathbb{W}$ is a specified wealth level. Assuming a risk aversion scaling parameter κ, the general problem for EW-xS, ($x = \{$P,L$\}$) can be written as

$$\begin{aligned} \text{EW-xS}_{t_0}(\mathbb{W}, \kappa) : \sup_{\mathcal{P}_0 \in \mathcal{A}} \Bigg\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^+, t_0^+} \Bigg[\sum_{i=0}^{M} \mathfrak{q}_i + \kappa G_{\mathrm{xS}}(W_T, \mathbb{W}) \\ + \epsilon W_T \Bigg| \mathcal{X}(t_0^-) = (s,b) \Bigg] \Bigg\} \end{aligned} \tag{4.4}$$

$$\text{subject to} \begin{cases} (S_t, B_t) \text{ follow processes (A.3) and (A.4);} \quad t \notin \mathcal{T} \\ W_\ell^+ = W_\ell^- - \mathfrak{q}_\ell \, ; \ \mathcal{X}_\ell^+ = (S_\ell^+, B_\ell^+) \\ W_\ell^- = S(t_i^-) + B(t_i^-) \\ S_\ell^+ = \mathfrak{p}_\ell(\cdot) W_\ell^+ \, ; \ B_\ell^+ = (1 - \mathfrak{p}_\ell(\cdot)) W_\ell^+ \\ (\mathfrak{q}_\ell(\cdot), \mathfrak{p}_\ell(\cdot)) \in \mathcal{Z}(W_\ell^-, W_\ell^+, t_\ell) \\ \ell = 0, \ldots, M \ ; \ t_\ell \in \mathcal{T} \end{cases} . \tag{4.5}$$

Note that we have added a stabilization term ϵW_T to the objective function in equation (4.4). In our examples, we use $|\epsilon| \ll 1$, with the result that the summary statistics are virtually the same as with $\epsilon = 0$. However, without the stabilization term, plots of the heat maps of the control are not well determined for large values of the wealth, as $t \to T$. This problem can be traced to the *Warren Buffet* effect.[14] Essentially, Warren Buffet can choose any investment strategy, and will never run out of savings.

More formally, the control problem is ill-posed when $t \to T, W_t \gg \mathbb{W}$. In this case, due to the maximum withdrawal constraint, and since the $Prob[W_t < \mathbb{W}] \simeq 0$, then the control has almost no effect on the objective function. Addition of the stabilization term regularizes the problem (see e.g. Chen et al. (2025)). We will discuss this further in later sections, and Appendix C.

4.2 Expected Shortfall ES

We are interested in the relationship between the above reward-risk formulations with the same reward but ES risk, i.e.,

$$\mathrm{EW}(\mathcal{X}_0^-, t_0^-) + \kappa\ \mathrm{ES}(\mathcal{X}_0^-, t_0^-)\ . \tag{4.6}$$

We formulate the EW-ES optimal control problem using the technique in Rockafellar and Uryasev (2000), more precisely ($0 < \alpha < 1$)

$$\begin{aligned}
\text{EW-ES}_{t_0}(\alpha, \kappa) :\ & \sup_{\mathcal{P}_0 \in \mathcal{A}} \Bigg\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-} \Bigg[\sum_{i=0}^{M} q_i \ + \ \kappa \sup_{W'} \Bigg(W' + \frac{1}{\alpha} \min(W_T - W', 0) \Bigg) \\
& \qquad\qquad + \epsilon W_T \Bigg| (s,b) \Bigg] \Bigg\} \\
\text{subject to } & \Big\{ \text{Conditions (4.5)} \quad .
\end{aligned} \tag{4.7}$$

Interchanging the order in *sup sup*$\{\cdot\}$ in problem (4.7), we equivalently have

$$\begin{aligned}
\text{EW-ES}_{t_0}(\alpha, \kappa) :\ & \sup_{W'} \sup_{\mathcal{P}_0 \in \mathcal{A}} \Bigg\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-} \Bigg[\sum_{i=0}^{M} q_i \ + \ \kappa \Bigg(W' + \frac{1}{\alpha} \min(W_T - W', 0) \Bigg) \\
& \qquad\qquad + \epsilon W_T \Bigg| \mathcal{X}_0^- = (s,b) \Bigg] \Bigg\} \\
\text{subject to } & \Big\{ \text{Conditions (4.5)} \quad .
\end{aligned} \tag{4.8}$$

Note that, as for the EW-xS problems, we have added a stabilization term ϵW_T to the objective function.

Remark 4.1 (Pre-commitment policy)**.** *Note that the optimal control for problem (4.8) is formally a pre-commitment policy (Forsyth, 2020a). We delay further discussion concerning this issue to Section 4.4.*

4.3 Properties of optimal solution of EW-ES formulation (4.8)

Let $\mathcal{P}_0$ be any permissible control for problem (4.8) and W_T be the terminal wealth corresponding to $\mathcal{P}_0$. Consider the wealth level associated with wealth distribution W_T, which is the maximizer below:[15]

[14] Warren Buffet is 95 years old, with a net worth of 150 billion USD.

[15] The $\arg\max$ is well defined since $\sup_{\mathcal{P}}\{\cdot\}$ is a continuous function of W'.

$$\mathbb{W} \;=\; \arg\max_{W'} \Bigg\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-} \Bigg[\sum_{i=0}^{M} q_i \;+\; \kappa \Big(W' + \frac{1}{\alpha} \min(W_T - W', 0) \Big) + \epsilon W_T \Bigg| \mathcal{X}_0^- = (s,b) \Bigg] \Bigg\} . \tag{4.9}$$

Following Rockafellar and Uryasev (2000), (4.9) is equivalent to the probability constraint below, under the assumption of continuity in distribution of W_T,

$$E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-} [\mathbf{1}_{W_T < \mathbb{W}}] \;=\; \alpha \; . \tag{4.10}$$

In other words, $\mathbb{W}$ is the α-VaR of W_T. This result can be easily seen from the following argument. Let $E_{\mathcal{P}_0}$ denote $E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-}$ for notational simplicity. Consider

$$\begin{aligned} & E_{\mathcal{P}_0} \Big(\mathbb{W} + \frac{1}{\alpha} \min(W_T - \mathbb{W}, 0) \Big) \\ & \text{subject to} \\ & \Big\{ E_{\mathcal{P}_0} [\mathbf{1}_{W_T \le \mathbb{W}}] = \alpha \quad . \end{aligned} \tag{4.11}$$

Let $g_{\mathcal{P}_0}(W_T)$ be the density of W_T under control $\mathcal{P}_0$. Then, write equation (4.11) as

$$\int_{-\infty}^{+\infty} \mathbb{W} \; g_{\mathcal{P}_0}(W_T) \; dW_T + \frac{1}{\alpha} \int_{-\infty}^{\mathbb{W}} (W_T - \mathbb{W}) \; g_{\mathcal{P}_0}(W_T) \; dW_T \tag{4.12}$$

subject to

$$\Big\{ \textstyle\int_{-\infty}^{\mathbb{W}} g_{\mathcal{P}_0}(W_T) \; dW_T = \alpha \quad . \tag{4.13}$$

We can write (4.12) as

$$\mathbb{W} \int_{-\infty}^{+\infty} g_{\mathcal{P}_0}(W_T) \; dW_T - \frac{\mathbb{W}}{\alpha} \int_{-\infty}^{\mathbb{W}} G_{\mathcal{P}_0}(W_T) \; dW_T + \frac{1}{\alpha} \int_{-\infty}^{\mathbb{W}} W_T \; g_{\mathcal{P}_0}(W_T) \; dW_T \; . \tag{4.14}$$

Using equation (4.13), this becomes

$$\mathbb{W} - \mathbb{W} + \frac{1}{\alpha} \int_{-\infty}^{\mathbb{W}} W_T \; g_{\mathcal{P}_0}(W_T) \; dW_T \;=\; E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-} \left[\frac{W_T \mathbf{1}_{W_T < \mathbb{W}}}{\alpha} \right] . \tag{4.15}$$

Thus, when (4.9) is satisfied, we have

$$\begin{aligned} E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-} [\mathbf{1}_{W_T < \mathbb{W}}] \;&=\; \alpha \\ E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-} \left[\frac{W_T \mathbf{1}_{W_T < \mathbb{W}}}{\alpha} \right] \;&=\; E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-} \left[\Big(\mathbb{W} + \frac{1}{\alpha} \min(W_T - \mathbb{W}, 0) \Big) \right] . \end{aligned} \tag{4.16}$$

Consequently, for the optimal $\mathbb{W}^*$ associated with W_T^* of the control $\mathcal{P}_0^*$ for EW-ES$_{t_0}(\alpha,\kappa)$, i.e.,

$$\mathbb{W}^* \;=\; \arg\max_{W'} \sup_{\mathcal{P}_0 \in \mathcal{A}} \Bigg\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^-,t_0^-} \Bigg[\sum_{i=0}^{M} q_i \;+\; \kappa \Big(W' + \frac{1}{\alpha} \min(W_T - W', 0) \Big) + \epsilon W_T \Bigg| \mathcal{X}_0^- = (s,b) \Bigg] \Bigg\} , \tag{4.17}$$

we have (4.16) also holds, or equivalently,

$$\begin{aligned} Prob[W_T^* < \mathbb{W}^*] &= \alpha \\ ES &= \text{mean of worst } \alpha \text{ fraction of outcomes} \\ &= E_{\mathcal{P}_0^*}^{\mathcal{X}_0^-, t_0^-}\left[\left(\mathbb{W}^* + \frac{1}{\alpha}\min(W_T^* - \mathbb{W}^*, 0)\right)\right] . \end{aligned} \tag{4.18}$$

From (4.18), we see immediately that $\mathbb{W}^*$ is the α-VaR (value at risk) of the terminal wealth W_T^* associated with the optimal control.

Risk and reward formulation EW-ES$_{t_0}$ assumes that an investor views risk in terms of a given probability level α, which implicitly defines a wealth level specifying the tail. Alternatively, an investor may think of risk with respect to a wealth level $\mathbb{W}$ explicitly, e.g., using $\mathbb{W} = 0$ represents the risk of negative wealth at $t = T$. For a fixed target wealth level $\mathbb{W}$, naturally we can consider the linear shortfall Pareto optimization:

$$\begin{aligned} \text{EW-LS}_{t_0}(\mathbb{W}, \widehat{\kappa}) : \qquad & \sup_{\mathcal{P}_0 \in \mathcal{A}} \Biggl\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}\Biggl[\sum_{i=0}^{M} q_i \; + \; \hat{\kappa} \; \min(W_T - \mathbb{W}, 0) + \epsilon W_T \\ & \qquad \Bigg| \mathcal{X}(t_0^-) = (s,b) \Biggr]\Biggr\} , \\ \text{subject to} \qquad & \Bigl\{ \text{Conditions (4.5)} \qquad , \end{aligned} \tag{4.19}$$

where $\hat{\kappa} > 0$ is a given weight parameter. Note that we notationally distinguish risk aversion parameters for EW-ES and EW-LS, in order to describe their precise connection subsequently.

The following Proposition 4.1 states that for any point on the EW-ES efficient frontier, there exists a pair of parameters $(\mathbb{W}, \widehat{\kappa})$ such that the EW-LS policy generates this exact same point (see also Forsyth (2020a)).

Proposition 4.1 (Optimal EW-ES strategy solves EW-LS)**.** *The pre-commitment strategy $\mathcal{P}^*$ which solves EW-ES$_{t_0}(\alpha,\kappa)$* (4.7) *is a solution to EW-LS$_{t_0}(\mathbb{W},\frac{\kappa}{\alpha})$* (4.19) *with the fixed wealth level $\mathbb{W} = \mathbb{W}^*$ defined in* (4.17).

Proof. Let $\mathcal{P}_0^*$ solve (4.7). Then it solves (4.8) due to equivalence between (4.8) and (4.7). Consequently $\mathcal{P}_0^*$ also solves the linear shortfall problem EW-LS$_{t_0}(\mathbb{W}^*,\frac{\kappa}{\alpha})$ in (4.19), i.e., $\mathcal{P}_0^*$ solves

$$\begin{aligned} \text{EW-LS}_{t_0}(\mathbb{W}, \hat{\kappa}) : \qquad & \sup_{\mathcal{P}_0 \in \mathcal{A}} \Biggl\{ E_{\mathcal{P}_0}^{\mathcal{X}_0^-, t_0^-}\Biggl[\sum_{i=0}^{M} q_i \; + \; \hat{\kappa}\min(W_T - \mathbb{W}, 0) \\ & \qquad + \epsilon W_T \Bigg| (s,b) \Biggr]\Biggr\} \\ \text{subject to} \; & \Bigl\{ \text{Conditions (4.5)} \qquad , \end{aligned}$$

with $\hat{\kappa} = \frac{\kappa}{\alpha}$, $\mathbb{W} = \mathbb{W}^*$, and $\mathbb{W}^*$ defined in (4.17). □

Let us consider sets of all permissible parameter pairs for EW-ES$_{t_0}$ and EW-LS$_{t_0}$ respectively:

$$\begin{aligned} \mathcal{D}_{ES} &= \{(\alpha, \kappa) \mid 0 < \alpha < 1 \; ; \; \kappa > 0\} \\ \mathcal{D}_{LS} &= \{(\mathbb{W}, \kappa) \mid \mathbb{W} \in \mathbb{R} \; ; \; \kappa > 0\} \; . \end{aligned} \tag{4.20}$$

We can understand the precise relationship between EW-ES$_{t_0}$ and EW-LS$_{t_0}$ in terms of optimal controls across all parameter pairs. Define sets of all optimal controls associated with these sets of parameter pairs:

$$
\begin{aligned}
\mathcal{H}^*_{\text{ES}} &= \{\mathcal{P}^*_0 : \ \mathcal{P}^*_0 \text{ solves EW-ES}_{t_0}(\alpha,\kappa)(4.7) \text{ for some } (\alpha,\kappa) \in \mathcal{D}_{ES}\} \\
\mathcal{H}^*_{\text{LS}} &= \{\mathcal{P}^*_0 : \ \mathcal{P}^*_0 \text{ solves EW-LS}_{t_0}(\mathbb{W},\hat{\kappa})(4.19) \text{ for some } (\mathbb{W},\hat{\kappa}) \in \mathcal{D}_{LS}\} \ .
\end{aligned}
\tag{4.21}
$$

We then have the following Corollary, which immediately follows from Proposition 4.1:

Corollary 4.1. *Let $\mathcal{H}^*_{ES}$ and $\mathcal{H}^*_{LS}$ be defined in* (4.21). *Then the set $\mathcal{H}^*_{ES}$ of optimal controls for Problem EW-ES$_{t_0}$ is a subset of the set $\mathcal{H}^*_{LS}$ of optimal controls for Problem EW-LS$_{t_0}$.*

4.4 Time inconsistent EW-ES and time consistent EW-LS

While Proposition 4.1 indicates that EW-ES$_{t_0}$ and EW-LS$_{t_0}$ share a common solution when the wealth level $\mathbb{W} = \mathbb{W}^*$, these two dynamic optimization formulations have different properties in terms of time consistency. To see this, we first recall the concept of time consistency and relate its relevance to the EW-ES$_{t_0}(\alpha,\kappa)$ problem, (4.8).

Consider the optimal control $\mathcal{P}^*_0 = (\mathcal{P}^*)^{t_0}$ computed at t_0 from problem (4.8) at all rebalancing times,

$$
(\mathcal{P}^*)^{t_0}(\mathcal{X}(t_i^-), t_i) \ , \ i = 0, \ldots, M \ , \tag{4.22}
$$

i.e., (4.22) denotes the optimal control $(\mathcal{P}^*)^{t_0}$ at any time $t_i \geq t_0$, as a function of the state variables $\mathcal{X}(t)$.

Next we solve the problem (4.8) starting at a later time $t_k, k > 0$, and denote the optimal control starting at t_k by:

$$
(\mathcal{P}^*)^{t_k}(\mathcal{X}(t_i^-), t_i) \ , \ i = k, \ldots, M \ . \tag{4.23}
$$

In general, the solution of (4.8) computed at t_k is not equivalent to the solution computed at t_0:

$$
(\mathcal{P}^*)^{t_k}(\mathcal{X}(t_i^-), t_i) \neq (\mathcal{P}^*)^{t_0}(\mathcal{X}(t_i^-), t_i) \ ; \ i \geq k > 0. \tag{4.24}
$$

This non-equivalence makes problem (4.8) *time inconsistent*, implying that the optimal control computed at t_k, $k > 0$, deviates from the control determined at time t_0. The optimal control $\mathcal{P}^*_0 = (\mathcal{P}^*)^{t_0}$ determined at the initial time is considered a *pre-commitment* control since the investor would need to commit to following the strategy at all times following t_0, even if the optimal control recomputed at future time becomes different. Some authors describe pre-commitment controls as non-implementable because of the incentive to deviate from the initial control.

Following Proposition 4.1, the pre-commitment control for EW-ES$_{t_0}$ (4.8), fortunately, is optimal for EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$, for which $\mathbb{W}$ is fixed at the optimal value (at time zero) in (4.17).

With a fixed $\mathbb{W}$, EW-LS$_{t_0}(\mathbb{W},\kappa/\alpha)$ uses a target-based linear shortfall as its measure of risk, and EW-LS$_{t_0}$ is trivially time consistent. Furthermore, $\mathbb{W}$ has the convenient interpretation of a disaster level of final wealth, as specified at time zero.

While the pre-commitment strategy $\mathcal{P}^*$ from EW-ES$_{t_0}(\alpha,\kappa)$, (4.8), is time inconsistent when viewed as a solution to EW-ES, this strategy is time consistent with respect to EW-LS$_{t_0}(\mathbb{W}^*,\frac{\kappa}{\alpha})$. In other words, conditional on information at t_n and fixed $\mathbb{W}^*$, the future decision $\{(\mathcal{P}^*)^{t_n}(\mathcal{X}(t_i^-), t_i)\ ;\ i = n, \ldots, M\}$ of the optimal pre-commitment EW-ES control, computed at t_0, solves

$$\begin{aligned}\text{EW-LS}_{t_n}\left(\mathbb{W}^*,\kappa/\alpha\right): \qquad \sup_{\mathcal{P}_n \in \mathcal{A}} \Bigg\{ & E_{\mathcal{P}_n}^{\mathcal{X}_n^-, t_n^-} \Bigg[\sum_{i=n}^{M} q_i \ + \ \frac{\kappa}{\alpha} \min(W_T - \mathbb{W}^*, 0) \\ & + \epsilon W_T \Bigg| \mathcal{X}(t_n^-) = (s, b) \Bigg] \Bigg\},\end{aligned} \tag{4.25}$$

for any given permissible stock and bond value pair (s,b).

Remark 4.2 (EW-ES $\rightarrow$ EW-LS). *Proposition 4.1 essentially tells us that any optimal control $\mathcal{P}^*$ from EW-ES problem (4.8), solves some EW-LS problem (4.25) with a fixed wealth level $\mathbb{W}^*$. Since EW-LS is time consistent, the EW-ES optimal control $\mathcal{P}^*$ is time consistent when Pareto optimality is assessed with EW-LS with this fixed wealth level $\mathbb{W}^*$.*

Since the optimal control $\mathcal{P}^*$ for $\text{EW-ES}_{t_0}(\alpha, \kappa)$ solves $\text{EW-LS}_{t_n}(\mathbb{W}^*, \kappa/\alpha)$ at any t_n, where $\mathbb{W}^*$ is the α-VaR of the conditional terminal wealth W_T^*, conditional on $W_0^* = s + b$, we can regard $\mathcal{P}^*$ as an induced time consistent strategy for $\text{EW-LS}_{t_n}(\mathbb{W}^*, \kappa/\alpha)$ (Strub et al., 2019). Consequently the investor has no incentive to deviate from the induced time consistent strategy, determined at time zero. Hence this policy is implementable.

For more detailed analysis concerning the subtle distinctions involved in pre-commitment, time consistent, and induced time consistent strategies, please consult Bjork and Murgoci (2010; 2014); Vigna (2014; 2022); Strub et al. (2019); Forsyth (2020a); Bjork et al. (2021); Landriault et al. (2021).

4.5 Equivalence between EW-ES and EW-LS problem

Problem $\text{EW-ES}_{t_0}(\alpha,\kappa)$ requires specification of the parameter pair (α,κ) while problem $\text{EW-LS}_{t_0}(\mathbb{W},\hat{\kappa})$ needs stipulation of parameter pair $(\mathbb{W},\hat{\kappa})$. We now consider an equivalence relationship between EW-ES_{t_0} and EW-LS_{t_0} by considering optimal controls from all parameter pairs. In other words, we consider the relationship between the $\text{EW-ES}_{t_0}(\alpha,\kappa)$ efficient surface, as (α,κ) varies, and the $\text{EW-LS}_{t_0}(\mathbb{W},\hat{\kappa})$ efficient surface, as $(\mathbb{W},\hat{\kappa})$ varies.

Proposition 4.1 in §4.3 establishes that the optimal control associated with any point on the EW-ES efficient surface generates some point on the EW-LS efficient surface. Now we establish that, under certain conditions, the converse also holds. In other words, given a point on an EW-LS efficient surface, then the associated EW-LS optimal control will generate some point on the EW-ES efficient surface.

Given $(\mathbb{W},\hat{\kappa})$, $\hat{\kappa} > 0$, let $\mathcal{P}_0^*$ be the optimal control at a point on the EW-LS_{t_0} efficient surface and W_T^* be the corresponding terminal wealth. To connect this optimal EW-LS solution $\mathcal{P}_0^*$ to some point on the EW-ES_{t_0} efficient surface precisely, we define

$$\alpha_{\hat{\kappa}}^*(\mathbb{W}) = \text{prob}(W_T^* < \mathbb{W}), \quad W_T^* \text{ is the terminal wealth of } \mathcal{P}_0^* \text{ which solves } \text{EW-LS}_{t_0}(\mathbb{W}, \hat{\kappa}) \tag{4.26}$$

Remark 4.3 (Construction of $\alpha_{\hat{\kappa}}^*(\mathbb{W})$). *Given $(\mathbb{W}, \hat{\kappa})$, and an optimal control $\mathcal{P}_0^*(\mathbb{W}, \hat{\kappa})$ which solves Problem EW-LS$_{t_0}(\mathbb{W}, \hat{\kappa})$, we determine $\alpha_{\hat{\kappa}}^*(\mathbb{W})$ from*

$$\alpha_{\hat{\kappa}}^*(\mathbb{W}) \quad = \quad E_{\mathcal{P}_0^*(\mathbb{W},\hat{\kappa})}[\mathbf{1}_{\{W_T^* < \mathbb{W}\}}] \ . \tag{4.27}$$

To ensure a proper correspondence to EW-ES_{t_0}, we consider the solution to EW-LS_{t_0} satisfying $0 < \alpha_{\hat{\kappa}}^*(\mathbb{W}) < 1$, i.e., we consider a restricted domain for EW-LS_{t_0} as:

$$\mathcal{D}_{LS}^+ \quad = \quad \{(\mathbb{W}, \hat{\kappa}) \mid \ 0 < \alpha_{\hat{\kappa}}^*(\mathbb{W}) < 1 \ \text{ and } \hat{\kappa} > 0\} \ . \tag{4.28}$$

Assumption 4.1 (invertibility of $\alpha^*_{\hat{\kappa}}(\mathbb{W})$)**.** *The function $\alpha^*_{\hat{\kappa}}(\mathbb{W})$ in* (4.26) *is well defined and is invertible at $(\mathbb{W},\hat{\kappa}) \in \mathcal{D}^+_{LS}$, i.e., for any $\mathbb{W}' \neq \mathbb{W}, (\mathbb{W}',\hat{\kappa}) \in \mathcal{D}^+_{LS}, \alpha^*_{\hat{\kappa}}(\mathbb{W}') \neq \alpha^*_{\hat{\kappa}}(\mathbb{W})$.*

Note that here we only assume that, for each given $(\mathbb{W},\hat{\kappa})$, EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$ yields a unique probability value $\alpha^*_{\hat{\kappa}}(\mathbb{W})$ but we do not assume uniqueness of the optimal controls for EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$.

Proposition 4.2 (Optimal EW-LS$_{t_0}$ strategy solves EW-ES$_{t_0}$)**.** *Suppose Assumption 4.1 holds at $(\mathbb{W},\hat{\kappa}) \in \mathcal{D}^+_{LS}$, then a solution to EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$ is a solution to EW-ES$_{t_0}(\alpha^*_{\hat{\kappa}}(\mathbb{W}),\alpha^*_{\hat{\kappa}}(\mathbb{W})\hat{\kappa})$, with $(\alpha^*_{\hat{\kappa}}(\mathbb{W}),\alpha^*_{\hat{\kappa}}(\mathbb{W})\hat{\kappa}) \in \mathcal{D}_{ES}$.*

Proof. Consider EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$ for a given $(\mathbb{W},\hat{\kappa}) \in \mathcal{D}^+_{LS}$. Let $\alpha^*_{\hat{\kappa}}(\mathbb{W})$ be defined in (4.26). Consider EW-ES$_{t_0}(\alpha^*,\alpha^*\hat{\kappa})$ where $\alpha^* = \alpha^*_{\hat{\kappa}}(\mathbb{W})$. Note that by definition of $\mathcal{D}^+_{LS}$, we must have $(\alpha^*,\alpha^*\hat{\kappa}) \in \mathcal{D}_{ES}$. Proposition 4.1 shows that a solution of EW-ES$_{t_0}(\alpha^*,\hat{\kappa}\alpha^*)$ is a solution to linear shortfall EW-LS$_{t_0}(\mathbb{W}^*,\hat{\kappa})$ problem for $\mathbb{W}^*$ defined in (4.17) with prob$(W^*_T < \mathbb{W}^*) = \alpha^*_{\hat{\kappa}}(\mathbb{W}^*) = \alpha^*$. Hence

$$\alpha^*_{\hat{\kappa}}(\mathbb{W}) = \alpha^*_{\hat{\kappa}}(\mathbb{W}^*) = \alpha^*.$$

Since $\alpha^*_{\hat{\kappa}}(\mathbb{W})$ is invertible, we have that $\mathbb{W} = \mathbb{W}^*$.

Assume that $\mathcal{P}^*_0$ solves EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$, then $\mathcal{P}^*_0$ solves EW-LS$_{t_0}(\mathbb{W}^*,\hat{\kappa})$ (4.19) , since $\mathbb{W} = \mathbb{W}^*$. Let $\hat{\kappa} = \frac{\kappa}{\alpha^*}$, where $\mathbb{W}^*$ defined in (4.17). Then $\mathcal{P}^*_0$ solves

$$\begin{aligned} &\sup_{\mathcal{P}_0 \in \mathcal{A}} \Bigg\{ E^{\mathcal{X}_0^-,t_0^-}_{\mathcal{P}_0} \Bigg[\sum_{i=0}^{M} q_i \ + \ \kappa(\mathbb{W}^* + \frac{1}{\alpha}\min(W_T - \mathbb{W}^*,0) \\ &\qquad +\epsilon W_T \Bigg| (s,b) \Bigg] \Bigg\} \\ &\text{subject to } \Big\{ \text{Conditions (4.5).} \quad , \end{aligned}$$

Since $\mathbb{W}^*$ is defined in (4.17), then $\mathcal{P}^*_0$ solves (4.8), and hence (4.7).

In other words, a solution $\mathcal{P}^*_0$ to EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$ solves EW-ES$_{t_0}(\alpha^*,\alpha^*\hat{\kappa})$, where $\alpha^* = \alpha^*_{\hat{\kappa}}(\mathbb{W})$. This completes the proof. □

Applying Corollary 4.1 and Proposition 4.2, we obtain the following Corollary 4.2.

Corollary 4.2. *Suppose Assumption 4.1 holds for any $(\mathbb{W},\hat{\kappa}) \in \mathcal{D}^+_{LS}$. Let $\mathcal{H}^*_{ES}$ and $\mathcal{H}^*_{LS}$ be defined in* (4.21)*. Then the set $\mathcal{H}^*_{ES}$ of optimal controls for Problem EW-ES$_{t_0}$ is identical to the set $\mathcal{H}^{*+}_{LS}$ of optimal controls for Problem EW-LS$_{t_0}$, where*

$$\mathcal{H}^{*+}_{LS} = \{\mathcal{P}^*_0(\mathbb{W},\hat{\kappa}) \in \mathcal{H}^*_{LS} \text{ and } (\mathbb{W},\hat{\kappa}) \in \mathcal{D}^+_{LS}\}.$$

Proof. From Proposition 4.1, any optimal solution $\mathcal{P}^*_0$ of EW-ES$_{t_0}(\alpha,\kappa)$, $(\alpha,\kappa) \in \mathcal{D}_{ES}$, solves EW-LS$_{t_0}(\mathbb{W}^*,\hat{\kappa})$ where $\hat{\kappa} = \frac{\kappa}{\alpha}$, and $0 < \alpha = \alpha^*_{\hat{\kappa}}(\mathbb{W}^*) < 1$, $\mathbb{W}^*$ defined in (4.17). Hence we have

$$\mathcal{H}^*_{\mathrm{ES}} \subset \mathcal{H}^{*+}_{\mathrm{LS}}.$$

Conversely, following Proposition 4.2, if $\mathcal{P}^*_0 \in \mathcal{H}^{*+}_{\mathrm{LS}}$, then $\mathcal{P}^*_0 \in \mathcal{H}^*_{\mathrm{ES}}$. Hence

$$\mathcal{H}^*_{\mathrm{ES}} \equiv \mathcal{H}^{*+}_{\mathrm{LS}}.$$

□

Corollary 4.2 essentially states that, under Assumption 4.1, the set of optimal controls associated with all points on the EW-ES$_{t_0}$ Pareto efficient surface, $(\alpha,\kappa) \in \mathcal{D}_{ES}$, is identical to the set of optimal controls for all points on the EW-LS$_{t_0}$ Pareto efficient surface with $(\mathbb{W},\hat{\kappa}) \in \mathcal{D}^+_{LS}$. It would be interesting to discover conditions on Problem EW-LS$_{t_0}$ which are required to guarantee that Assumption 4.1 holds. We leave this for future work.

5 Numerical Comparison of Different Risk-Reward Pairs

In 4.5, we have established an equivalence relationship between EW-ES_{t_0} and EW-LS_{t_0} across all parameter pairs. In practice, a retiree maybe only interested in the efficient frontier for a fixed probability level α for EW-ES_{t_0} or wealth level $\mathbb{W}$ for EW-LS_{t_0}. In this section we computationally assess and compare, based on market data, EW-ES_{t_0}, EW-LS_{t_0}, and EW-PS_{t_0} efficient frontiers when parameter α and $\mathbb{W}$ are fixed. Note that EW-PS_{t_0} does not have equivalence relationship with EW-ES_{t_0} or EW-LS_{t_0}, since the probability of shortfall risk ignores the magnitude of loss.

5.1 Data

We use data from the Center for Research in Security Prices (CRSP) on a monthly basis over the 1926:1-2023:12 period.[16] Our base case tests use the CRSP US 30 day T-bill for the bond asset and the CRSP value-weighted total return index for the stock asset. This latter index includes all distributions for all domestic stocks trading on major U.S. exchanges. All of these various indexes are in nominal terms, so we adjust them for inflation by using the U.S. CPI index, also supplied by CRSP. We use real indexes since investors funding retirement spending should be focused on real (not nominal) wealth goals.

We use the parametric model for the real stock index and real constant maturity bond index described in Appendix A.

Remark 5.1 (Choice of 30-day T-bill for the bond index). *It might be argued that the bond index should hold longer-dated bonds such as 10-year Treasuries since this would allow the investor to harvest the term premium. Long-term bonds enjoyed high real returns during 1990-2022. However, it is unlikely that this will continue to be true over the next 30 years. For example, during the period 1950-1983, long term bonds had negative real returns (Hatch and White, 1985), while short-term T-bills had positive real returns. If one imagines that the next 30 years will be a period with inflationary pressures, this suggests that the defensive asset should be short-term T-bills. Note that the historical real return of short-term T-bills over 1926:1-2023:12 is approximately zero. Hence our use of T-bills as the defensive asset is a conservative approach going forward.*

Remark 5.2 (Sensitivity to Calibrated Parameters). *Some readers might suggest that the stochastic processes (A.3-A.4) are simplistic, and perhaps inappropriate. However, we will test the optimal strategies (computed assuming processes (A.3-A.4)) with calibrated parameters in Table A.1 using bootstrap resampled historical data (see Section 5.2 below). The computed strategy seems surprisingly robust to model misspecification. Similar results have been noted for the case of multi-period mean-variance controls (van Staden et al., 2021). We conjecture that this robustness is due to the self-correcting nature of feedback controls.*

5.2 Historical Market

We compute and store the optimal controls using the parametric model (A.3) and (A.4), following the same procedure as in the synthetic market case. However, we compute statistical quantities using the stored controls, but using bootstrapped historical return (test) data directly. In this case, we make no assumptions concerning the stochastic processes followed by the stock and bond

[16] More specifically, results presented here were calculated based on data from Historical Indexes, ©2024 Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business. Wharton Research Data Services (WRDS) was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

Data series	Optimal expected block size $\hat{b}$ (months)
Real 30-day T-bill index	50.805
Real CRSP value-weighted index	3.17535

Table 5.1: *Optimal expected blocksize $\hat{b} = 1/v$ when the blocksize follows a geometric distribution $Pr(b = k) = (1 - v)^{k-1} v$. The algorithm in Patton et al. (2009) is used to determine $\hat{b}$. Historical data range 1926:1-2023:12.*

indices. We remind the reader that all returns are inflation-adjusted. We use the stationary block bootstrap method (Politis and Romano, 1994; Politis and White, 2004; Patton et al., 2009; Cogneau and Zakalmouline, 2013; Dichtl et al., 2016; Cavaglia et al., 2022; Simonian and Martirosyan, 2022; Anarkulova et al., 2022).

Sampling the data in blocks accounts for serial correlation in the data series. A key parameter is the expected blocksize. We use the algorithm in Patton et al. (2009) to determine the optimal blocksize for the bond and stock returns separately, see Table 5.1. Informally, the expected blocksize corresponds to the average memory length of the time series. Short term bond returns are primarily driven by policies of central banks, which generally favour gradual changes to interest rates. We can see that the expected blocksize for T-bills, from Table 5.1, is about 51 months, which is consistent with intuition. On the other hand, the expected blocksize for stocks, from Table 5.1, is about 3 months, indicating that, stocks show a very short memory length, which is also consistent with intuition.

In our application, we use a paired sampling approach to simultaneously draw returns from both time series. This paired sampling technique preserves historical correlations between the stock and bond indexes. However, there does not seem to be an existing algorithm which estimates the expected blocksize for a paired time series. In this case, a reasonable estimate for the blocksize for the resampling algorithm of a paired time series would be about 2.0 years, which is roughly the average of the two expected blocksize estimates for the separate time series for stocks and bonds. We will give results for a range of blocksizes as a check on the robustness of the bootstrap results. Detailed pseudo-code for block bootstrap resampling is given in Forsyth and Vetzal (2019).

5.3 Investment Scenario

Table 5.2 shows our base case investment scenario. We use thousands of dollars as our units of wealth. For example, a withdrawal of 40 per year corresponds to \$40,000 per year (all values are real, i.e. inflation-adjusted), with an initial wealth of 1000 (i.e. \$1,000,000). This would correspond to the use of the four per cent rule (Bengen, 1994). Recall that we assume that the investor has real estate, which is in a separate mental bucket (Shefrin and Thaler, 1988). Real estate is a hedge of last resort, used to fund required minimum cash flows (Pfeiffer et al., 2013). We assume that the retiree owns mortgage free real estate worth \$400,000, of which \$200,000 can be easily accessed using a reverse mortgage.

5.4 Numerical Results

We use the numerical method described in (Forsyth, 2022; Forsyth et al., 2024) to compute the optimal controls, which is based on solving a Partial Integro-Differential Equation (PIDE), combined

Investment horizon T (years)	30.0
Equity market index	CRSP Cap-weighted index (real)
Bond index	30-day T-bill (US) (real)
Initial portfolio value W_0	1000
Mortgage free real estate	400
Cash withdrawal/rebalancing times	$t = 0{,}1.0, 2.0, \ldots, 29.0$
Withdrawal range (per year)	$[\mathfrak{q}_{\min}, \mathfrak{q}_{\max}] = [30, 60]$
Equity fraction range	$[\mathfrak{p}_{\min}, \mathfrak{p}_{\max}] = [0{,}1]$
Borrowing spread μ_c^b	0.03
Rebalancing interval (years)	1.0
Target Wealth $\mathbb{W}$ (EW-xS, $x = \{P{,}L\}$)	0.0
α (EW-ES)	.05
Stabilization ϵ (see equation (4.4))	-10^{-4}
Market parameters	See Table A.1

TABLE 5.2: *Input data for examples. Monetary units: thousands of dollars.*

with discretizing the controls and finding the optimal values by exhaustive search. A brief overview is given in Appendix B.

We compute optimal strategies from EW-ES$_{t_0}(\alpha, \kappa)$ with $\alpha = 0.05$ and $\kappa > 0$. For EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$ and EW-PS$_{t_0}(\mathbb{W},\hat{\kappa})$, we compute optimal strategies with $\mathbb{W} = 0$ and $\hat{\kappa} > 0$.

5.5 Convergence

Figure 5.1 shows the convergence of the EW-LS frontier, as a function of the PIDE nodes. The curves for different numbers of nodes essentially overlap, indicating satisfactory convergence. Appendix B.1 shows convergence as the number of grid nodes increases, for a single point on the synthetic market EW-LS efficient frontier. Similar results were obtained for the other strategies.

5.6 Performance Frontiers: EW-LS, EW-PS, EW-ES

In this section, we compare efficient frontier curves computed from EW-ES$_{t_0}(\alpha,\kappa)$ with fixed $\alpha = 0.05$ and $\kappa > 0$, and EW-xS$_{t_0}(\mathbb{W},\hat{\kappa})$ ($x =$P,L) with $\mathbb{W} = 0$ and $\hat{\kappa} > 0$. We choose this comparison setting since it seems more immediately relevant from a retiree's perspective.

We note that it is relevant to compare performance of optimal strategies using risk measures which are not directly included in their respective objective functions. This helps an investor understand consequences of implementing a specific strategy in terms of different but relevant performance measures. When the performance of optimal controls from EW-$\mathcal{X}$ (with $\mathcal{X}$ = one of ES, LS, or PS) is plotted in the domain EW-$\mathcal{Y}$, where $\mathcal{X}$ can be different from $\mathcal{Y}$, we refer the curve as a performance frontier. The term efficient frontier refers to the case when $\mathcal{Y} = \mathcal{X}$. The efficient frontier must plot above all other performance frontiers. We assess the performance of these strategies in the performance domain (EW,LS), (EW,PS), and (EW,ES) in Figures 5.2, 5.3(a), and 5.3(b) respectively.

Since all three formulations share the same reward, the top left sides of performance frontiers from all strategies are expected to converge asymptotically (as risk aversion parameter goes to zero) in all performance domains (EW,LS), (EW,PS), and (EW,ES).

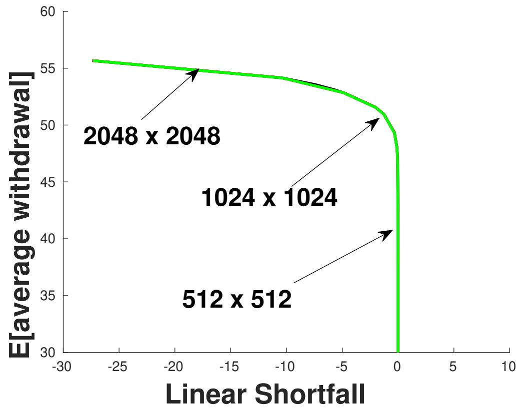

FIGURE 5.1: *EW-LS convergence test. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1. The optimal control is determined by solving the PIDEs as described in Appendix B. Grid refers to the grid used in the algorithm in Appendix B: $n_x \times n_b$, where n_x is the number of nodes in the $\log s$ direction, and n_b is the number of nodes in the $\log b$ direction. Units: thousands of dollars (real). The controls are stored, and then the final results are obtained using a Monte Carlo method, with 2.56×10^6 simulations. Target wealth $\mathbb{W} = 0.0$. Maximum fraction on equities $p_{\max} = 1.0$.*

Note that efficient frontier curves in either performance domain (EW,ES) or (EW,LS), from EW-ES$_{t_0}(\alpha,\kappa)$ and EW-LS$_{t_0}(\mathbb{W},\hat{\kappa})$, for fixed α and $\mathbb{W}$, are not expected to coincide, except possibly at single points.

As the risk aversion parameter increases, performance frontiers on the right side from different formulations are expected to deviate from each other more significantly. Overall, Figure 5.2, Figure 5.3(a), and 5.3(b) do demonstrate larger differences in performance frontier curves on the right side, see particularly EW-ES frontiers in Figure 5.3(b). This confirms that the choice of objective function is important in achieving risk control. Subsequently we compare and contrast performance frontiers in more detail.

5.6.1 EW-LS Performance

Figure 5.2 plots frontier curves in the (EW,LS) domain. We compute EW-LS, EW-ES and EW-PS optimal controls, but plot their (EW,LS) performance measures in the same figure. Naturally the efficient frontier curve of the EW-LS control must plot above all the other curves (since the objective function of EW-LS aligns with the specified measures). However, it is interesting to see that the EW-ES and EW-PS controls are not overly suboptimal, relative to EW-LS$_{t_0}(\mathbb{W} = 0,\hat{\kappa})$, using (EW,LS) criteria.

From Proposition 4.1, we expect that there is a point (with target wealth $\mathbb{W} = 0$) at which the EW-ES and EW-LS curve coincide. In Figure 5.2, we can see that this point occurs at $EW \simeq 52.3$. Figure 5.2 also shows the results for the Bengen strategy (Bengen, 1994).[17] We can see that the Bengen strategy is considerably suboptimal compared to any of the other strategies. However, it is only fair to point out that the Bengen strategy always withdraws 40 per year (units thousands), while the other strategies have minimum withdrawals of 30 per year.

[17] Recall that the recommended policy is to withdraw 4% of the initial capital per year, inflation adjusted, and to rebalance to a weight of 50% stocks annually. The 4% withdrawal would correspond to 40 per year for our example.

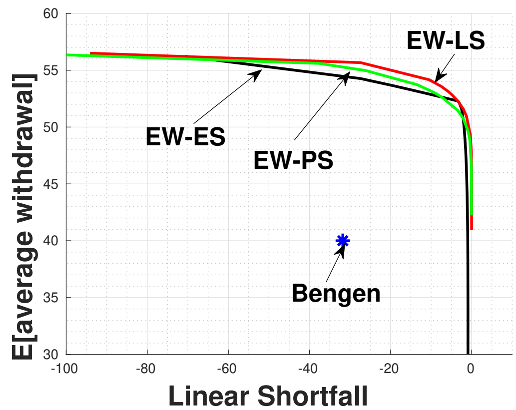

FIGURE 5.2: *EW-LS performance frontier. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1. Synthetic market. Controls computed using EW-PS and EW-ES, and results plotted in terms of EW-LS criteria. The EW-LS frontier plots above the the controls computed using EW-PS and EW-ES objective functions. The Bengen control withdraws 40 per year, and rebalances annually to 50% bonds and 50% stocks. The wealth target level* $\mathbb{W} = 0$ *for both EW-LS*$_{t_0}$ *and EW-PS*$_{t_0}$.

5.6.2 EW-PS Performance

Figure 5.3(a) plots the EW-PS performance frontier. Along the x-axis we plot $Prob[W_T > 0] = 1 - Prob[W_T < 0]$, to produce consistent shapes for the frontiers. As before, we compute the EW-ES and EW-LS efficient controls, but plot them using PS as a risk measure. As expected, the EW-PS frontier plots above the other curves (it is, after all, the efficient strategy according to the $Prob[W_T > 0]$ risk measure).

There is somewhat more variation in these curves compared to Figure 5.2. In particular, the EW-PS and EW-LS curves generate $Prob[W_T > 0] \simeq 0.9998$ for the largest values of κ (the right hand most point on the curves). In contrast, the EW-LS strategy never gets above $Prob[W_T > 0] \simeq 0.994$. We also see that the EW-LS curve flat-tops to the left of $EW \simeq 52$.[18]

5.6.3 EW-ES Performance

Figure 5.3(b) shows the EW-ES performance frontier. We also show (EW,ES) measures for optimal EW-PS$_{t_0}(0,\kappa)$ and EW-LS$_{t_0}(0,\kappa)$ strategies. The curves for EW-PS and EW-LS are very similar. However, for $ES > 0$, the EW-ES curve is dramatically different. This can be explained as follows.

The EW-ES strategy moves towards maximizing ES as κ becomes large. This comes at the expense of decreased $Prob[W_t > 0]$ (from Figure 5.3(a)). On the other hand, the EW-LS and EW-PS strategies have no risk if $W_T > 0$, so focus entirely on increasing EW, if $W_t > 0$ as $t \rightarrow T$. Effectively, this means that for the EW-LS and EW-PS strategies, it does not make sense to consider points in the performance frontier which are below the *knee* of the curves. To the left of the *knee* of the curves, EW-LS is very close to the EW-ES curve.

Recall that we have assumed that the retiree can access \$200K using a reverse mortgage with real estate as collateral. Consequently, as a rule of thumb, any point on any frontier which has $ES > -200$ is acceptable from a risk management point of view. In other words the mean of the

[18] If we extend the x-axis to the left, then, eventually, all three curves meet. However, these points have an uncomfortably large $Prob[W_T] < 0$, hence are not of practical interest.

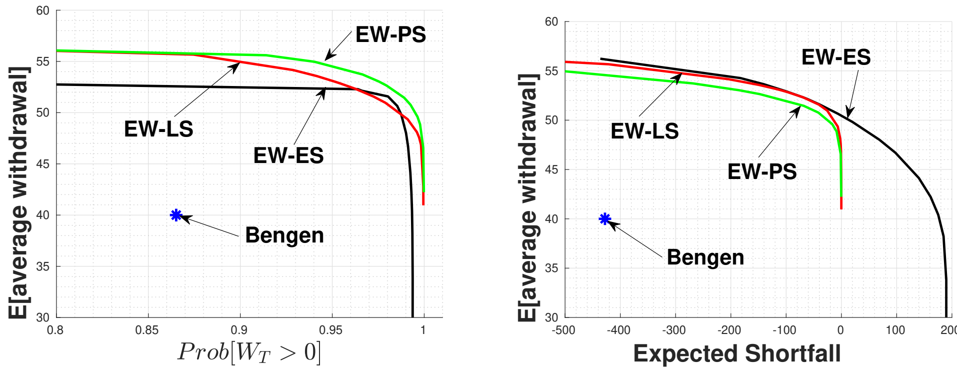

(a) *EW-PS performance frontier. Results for EW-LS and EW-ES also shown.*

(b) *EW-ES performance frontier. Results for EW-LS and EW-PS also shown.*

FIGURE 5.3: *EW-PS and EW-ES performance frontiers. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1. Synthetic market. The Bengen control withdraws 40 per year, and rebalances annually to 50% bonds and 50% stocks. Target wealth $\mathbb{W} = 0$ for EW-LS and EW-PS.*

worst 5% of the outcomes can be hedged using real estate. In particular, we can see that the Bengen strategy fails this risk management test.

5.6.4 Comparison of performance frontiers

Our first observation is that in all cases, whatever the strategy or risk measure, the Bengen strategy is significantly sub-optimal. We also make the following additional observations:

- Firstly recall that we use the same target wealth level $\mathbb{W} = 0$ for both EW-LS$_{t_0}$ and EW-PS$_{t_0}$. We observe that their frontier curves are close in all three measurement domains, (EW,LS), (EW,PS) and (EW,ES), even asymptotically as the risk aversion parameter $\kappa \to +\infty$. This suggests that choosing EW-LS$_{t_0}$ also leads to good performance in terms of EW-PS$_{t_0}$. Furthermore, since linear shortfall is an expectation of piecewise linear shortfall, i.e., $E(\max(W_T - \mathbb{W}, 0))$, while the probability function is the expectation of a discontinuous indicator function, i.e., $E(\mathbf{1}_{W_T < \mathbb{W}})$, consequently solving EW-LS$_{t_0}$ can be computationally preferable to solving EW-PS$_{t_0}$.

- Choosing a suitable risk measure as part of the objective function which aligns with the desired decumulation goals does matter. Different risk measures can lead to very different performing strategies. This is particularly important in decumulation. For example, Figure 5.3(a) shows that the optimal EW-ES$_{t_0}$ strategy is inefficient at minimizing the probability of negative terminal wealth. The smallest probability of negative wealth achieved is at the expense of steeply diminishing reward. Similarly Figure 5.3(b) shows that, with the wealth target $\mathbb{W} = 0$, the 5% ES risk associated with the optimal EW-LS$_{t_0}$ and EW-PS$_{t_0}$ strategies as the risk aversion parameter $\hat{\kappa} \to +\infty$ is far from the optimal 5%-ES risk achievable.

- While Figure 5.2 seems to indicate that all the strategies perform reasonably well in terms of the LS risk measure, we note that there is a similar steeper drop from the optimal EW-ES$_{t_0}$ strategy as the risk aversion parameter goes to $+\infty$. We further note that the scales of the horizontal axis in Figure 5.2 and Figure 5.3(a) are very different. Optimal EW-ES$_{t_0}$ strategies are unable to achieve the minimum LS risk and the smallest LS risk strategy is achieved at the expense of suboptimal rewards.

From Figure 5.3(a) we can see that, in terms of the PS risk measure, EW-LS plots a bit below the optimal EW-PS frontier. However, the EW-ES curve has a very unusual behaviour (in terms of PS risk). Any increase in EW above about 53 causes a large decrease in $Prob[W_T > 0]$.

Turning attention to Figure 5.3(b), in terms of ES risk, all strategies behave similarly to the left of $ES \simeq 0$. However, the EW-ES efficient frontier continues to generate positive ES for $EW < 50$. Essentially, this is because the EW-ES strategy focuses on maximizing ES, but at the expense of giving up increases in $Prob[W_T > 0]$. Recall from our previous discussion that the EW-PS and EW-LS strategies only make sense if we look at points to the left of the *knee* of the curves.

From a practical point of view, it is not clear that maximizing ES when it is positive is consistent with the retiree's view of risk (i.e. running out of savings).

5.7 Summary of comparisons of performance frontiers

In examining Figures 5.2-5.3, we should bear in mind the following

- Efficient frontiers typically have two well characterized regions. In the extreme right most portion of the curve, small decreases in ES, (1-PS), or LS (i.e. small increases in risk) result in large increases in the reward. In the extreme left hand portion of the curve, small increases in reward require large increases in risk. The *knee* of the curve separates these two regions. Consequently, most investors will choose a point in the knee of the efficient frontier, since this represents a good trade-off between risk and reward.
- Robustness can be measured by examining how the frontier EW-$\mathcal{X}$ behaves when plotted with $\mathcal{Y}$ as a risk measure.

In summary, Figures 5.2-5.3 show that

- EW-PS and EW-LS have similar *knee* regions in all the plots. EW-ES is somewhat of an outlier, with some plots having very small knee regions, and other plots showing very large knee regions.

This behaviour can be traced to the fact that EW-LS and EW-PS have an investor specified parameter $\mathbb{W} = 0$, which is the disaster level of final wealth. The efficient frontiers consequently are easy to interpret. In contrast, EW-ES determines parameter $\mathbb{W}$ as part of the optimization process, which sometimes leads to non-intuitive behaviour.

We will further differentiate these strategies, in Section 6, by examining the Cumulative Distribution Functions (CDF) of the final wealth based on bootstrap resampling data.

5.8 Tests for robustness: bootstrap resampling

We compute and store the optimal controls for EW-PS, EW-LS and EW-ES objective functions, based on the parametric market model described in Appendix A. We then test these controls using block bootstrap resampling of the market data in 1926:1-2023:12 (Politis and Romano, 1994; Politis and White, 2004; Patton et al., 2009; Cogneau and Zakalmouline, 2013; Dichtl et al., 2016; Cavaglia et al., 2022; Simonian and Martirosyan, 2022; Anarkulova et al., 2022).

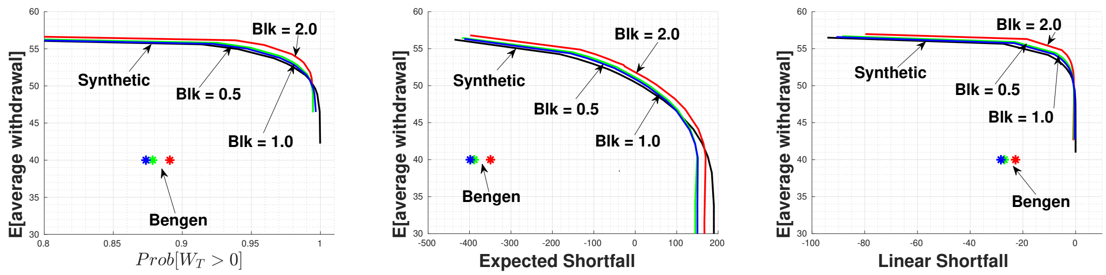

(a) *EW-PS performance frontier, bootstrapped simulations.*

(b) *EW-ES performance frontier, bootstrapped simulations.*

(c) *EW-LS performance frontier, bootstrapped simulations.*

FIGURE 5.4: *Optimal controls computed using the synthetic market model. These controls tested using bootstrapped historical data. Expected blocksizes (years) shown.* 10^6 *bootstrap resamples. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1. The Bengen control withdraws 40 per year, and rebalances annually to 50% bonds and 50% stocks. The Bengen results are also shown for expected blocksizes of* $0.5, 1.0, 2.0$ *years.*

Figure 5.4 compares the synthetic market results (test and train on the parametric market model) as well as testing this control on bootstrapped historical data, for all three objective functions: EW-LS, EW-PS and EW-ES. The bootstrapped tests are carried out for a range of expected blocksizes. In all cases, for all blocksizes, the performance frontiers are quite close, indicating that the controls computed using the parametric market model in Appendix A are relatively robust to model misspecification.

Further insight can be obtained by examining the summary statistics in Table 5.3 (synthetic market) and Table 5.4 (historical market). It would seem that the EW-LS strategy is a good compromise, having a relatively small $Prob[W_T < 0]$, and with an expected shortfall close to the optimal value from the EW-ES solution. Note that $\mathbb{W}$ is a byproduct of the optimization algorithm for the EW-ES problem. This may not correspond, intuitively, to the investor's preferences. For example, as we move rightward along the EW-ES performance frontier, $\mathbb{W}$ becomes a large positive value. Any value of W_T to the left of this point, is regarded (by the objective function) as a bad outcome, which probably does not correspond to most investor's concept of risk.

In contrast, $\mathbb{W}$ is an input parameter for EW-PS and EW-LS. In our case, since our main concern is running out of cash, setting $\mathbb{W} = 0$ is clearly a reasonable choice.

Note that Table 5.4 shows that the ES(5%) result for the EW-LS control (computed in the synthetic market) is actually better than for the EW-ES control (also computed in the synthetic market), when tested in the historical market. This suggests that the EW-LS control is more robust than the EW-ES control.

Another test of robustness is shown in Table 5.5. Here, we rank each strategy, in terms of performance, according to each risk criteria, in the historical market. All strategies have approximately the same $EW \simeq 53$. In this case, we can see that EW-LS is the clear winner.

5.9 Comparison with the Bengen strategy

Consider Figure 5.4(a), with EW $\simeq 50$. This translates to average withdrawals of 5% of initial wealth with $Prob[W_T < 0] < 1\%$. Contrast this with the bootstrapped results for the Bengen strategy, where the withdrawals are 40 per year (4% of initial wealth (real)), with a probability of failure $> 10\%$.

Strategy	κ	$E[\sum_i \mathfrak{q}_i]/M$	LS($\mathbb{W}=0$)	ES(5%)	$Prob[W_T < 0]$	$\mathbb{W}$
EW-ES	0.5925	52.97	-11.106	-102.36	.271	-31.15
EW-LS	9.3822	52.99	-5.3332	-106.66	.048	0.0
EW-PS	2670.9	53.04	-9.2747	-185.40	.027	0.0

TABLE 5.3: *Synthetic market, summary statistics for EW-PS, EW-LS, and EW-ES objective functions, EW $\simeq 53$ for all strategies. LS refers to $E[\min(W_T - \mathbb{W}, 0)]$, ES(5%) is the mean of the worst five per cent of the outcomes. $\mathbb{W}$ is specified for EW-PS and EW-LS, while it is an outcome of the EW-ES optimization. Scenario in Table 5.2. Parameters in Table A.1. Units: thousands of dollars (real). M is the total number of withdrawals (rebalancing dates).*

Strategy	$E[\sum_i \mathfrak{q}_i]/M$	LS($\mathbb{W}=0$)	ES(5%)	$Prob[W_T < 0]$
EW-ES	53.18	-5.4192	-46.21	0.250
EW-LS	52.93	-1.5448	-30.85	0.0226
EW-PS	53.12	-3.705	-74.02	0.0120

TABLE 5.4: *Block bootstrap resampling, summary statistics for EW-PS, EW-LS, and EW-ES objective functions, EW $\simeq 53$ for all strategies. Blocksize two years, 10^6 bootstrap resamples. Optimal controls computed in the synthetic market. LS refers to $E[\min(W_T - \mathbb{W}, 0)]$, ES(5%) is the mean of the worst five per cent of the outcomes. Scenario in Table 5.2. Parameters in Table A.1. Units: thousands of dollars (real). M is the total number of withdrawals (rebalancing dates).*

Strategy	Rank			
	LS($\mathbb{W}=0$)	ES(5%)	$Prob[W_T < 0]$	Total Score
EW-ES	3	2	3	8
EW-LS	1	1	2	4
EW-PS	2	3	1	6

TABLE 5.5: *Ranking of strategies, historical market. Each strategy is ranked (first, second or third). Optimal controls computed in the synthetic market. Total score is the sum of the rows, smaller is better. EW $\simeq 53$ for all strategies. Data is from Table 5.4. Block bootstrap resampling, summary statistics for EW-PS, EW-LS, and EW-ES objective functions. Blocksize two years, 10^6 bootstrap resamples. LS refers to $E[\min(W_T - \mathbb{W}, 0)]$, ES(5%) is the mean of the worst five per cent of the outcomes. Scenario in Table 5.2. Parameters in Table A.1. Units: thousands of dollars (real).*

Similarly, Figure 5.4(b) has (EW,ES) = (50,0) for the EW-ES optimal strategy, compared with (EW,ES) = $\simeq$ (40, -350) for the Bengen strategy.

Finally, Figure 5.4(c) gives (EW,LS) = (50,0) for the EW-LS optimal policy, compared with (EW,LS) = $\simeq$ (40, -20) for the Bengen strategy.

Of course, all these comparisons come with the caveat that the Bengen strategy withdraws a fixed amount per year, while the results for the optimal strategies are in terms of expected withdrawals.

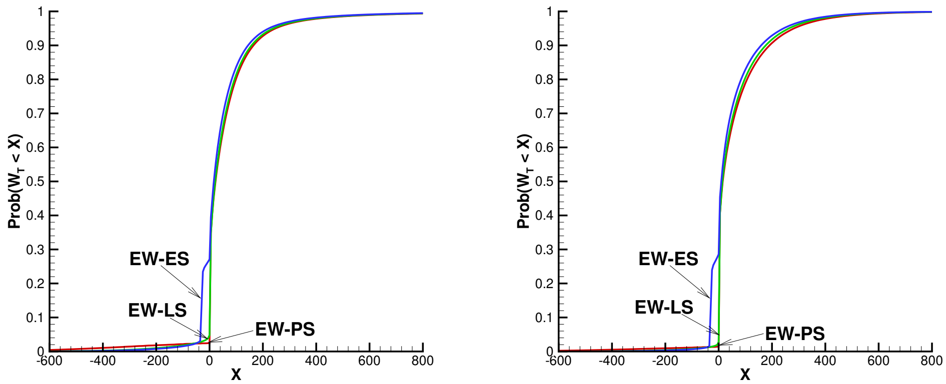

(a) *CDF curves, synthetic market*

(b) *CDF curves, bootstrapped historical market.*

FIGURE 6.1: *CDF curves, all strategies have the same average $EW \simeq 53$. Optimal controls computed using the synthetic market model. Tests in the synthetic market Figure 6.1(a) and the historical market, Figure 6.1(b) shown. Expected blocksize: two years. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1.*

6 CDFs of the optimal strategies

Figure 6.1 shows the CDF curves for the final wealth W_T for all three strategies. The results are shown for both the synthetic and historical market. For each strategy, the point on the efficient frontier was selected so that $EW \simeq 53$. It is interesting to observe that all strategies have similar CDFs for $X > 0$ ($Prob[W_T > 0]$), and rapidly increase to the right of this point. This indicates that all strategies are efficient in the sense that there is little unspent wealth at $t = 30$ years (age 95). This contrasts with the Bengen policy, which has a non-trivial probability of either running out of cash or ending up with large unspent wealth.

Figure 6.2(b) focuses on the area of the CDF curves near $X = 0$. Examining the synthetic market results, Figure 6.2(a), we can see that the EW-PS and EW-LS curves behave very similarly near $W_T = 0$, but there is a difference in the left tail, as might be expected. We can see that EW-PS does an excellent job of producing small $Prob[W_T < 0]$. However, this strategy does not do well in the left tail compared with EW-LS. The EW-ES strategy, on the other hand, has a fairly high probability that $W_T < 0$, compared with either EW-PS or EW-LS. However, this is a bit misleading, since the EW-ES CDF plots below the other strategies for $X < -40$. The historical market CDFs, Figure 6.2(b), are qualitatively similar to the synthetic market curves.

6.1 Summary of comparison of CDFs

Our conclusions from Figures 6.1-6.2 are

- The CDFs for all strategies are similar for $W_T > 0$.
- EW-PS performs poorly in the left tail compared to the other strategies.
- EW-ES performs poorly in terms of $Prob[W_T < 0]$ compared to the other strategies.

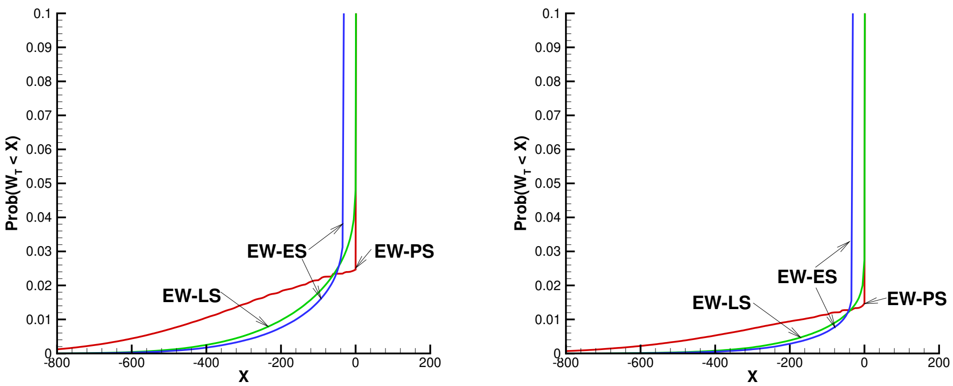

(a) *CDF curves, synthetic market*

(b) *CDF curves, bootstrapped historical market.* 10^6 *bootstrap resamples.*

FIGURE 6.2: *Zoomed plots, CDF curves, all strategies have the same average EW $\simeq$ 53. Optimal controls computed using the synthetic market model. Tests in the synthetic market Figure 6.1(a) and the historical market, Figure 6.1(b) shown. Expected blocksize: two years. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1.*

Note that the disaster level of final wealth $\mathbb{W} = 0$ is specified by the investor for EW-LS and EW-PS, while $\mathbb{W}$ is a byproduct of the optimization algorithm for the EW-ES strategy.

7 Comments: EW-LS, EW-PS and EW-ES strategies

The EW-PS optimal control, using PS risk (probability of running out of savings), seems at first sight to be an appealing intuitive strategy. However, the CDF of the final wealth shows that this strategy generates a very fat left tail. This is simply due to the fact that PS risk does not weight the amounts less than $\mathbb{W}$.

The EW-ES optimal control also has a simple intuitive interpretation. The ES (mean of the worst 5% of the outcomes) is a dollar amount that can be compared with, for example, the retiree's real estate hedge of last resort. However, in some cases, the ES can be large and positive, which does not correspond to what we would normally think of as risk. In addition, EW-ES is formally time inconsistent. There is, of course, an induced time consistent policy, which is simply the EW-LS control with suitable $\mathbb{W}$.

The EW-LS control is trivially time-consistent. The investor specified parameter $\mathbb{W}$ in the EW-LS objective function is easily interpreted as the disaster level of final wealth. The EW-LS controls also perform reasonably well using ES (expected shortfall) or PS (probability of shortfall) as risk measures. In particular, the knee of the EW-LS strategy,[19] when plotted using ES as a risk measure, has a similar ES risk as the optimal EW-ES strategy.

The EW-LS control is also more robust, when tested in the historical (bootstrapped) market, compared to the other strategies.

Consequently, we recommend use of the EW-LS control for decumulation.

[19] See Section 5.7 for a definition of the knee portion of the efficient frontier.

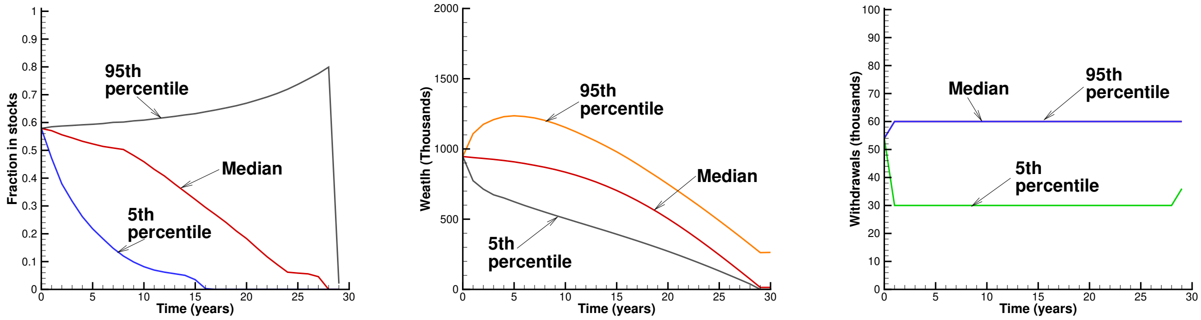

(a) *Percentiles fraction in stocks*

(b) *Percentiles wealth*

(c) *Percentiles withdrawals*

FIGURE 8.1: *Scenario in Table 5.2. EW-LS control computed from problem EW-LS Problem (4.4). Parameters based on the real CRSP index, and real 30-day T-bills (see Table A.1). Control computed and stored from the Problem (4.4) in the synthetic market. Control used in the historical market,* 10^6 *bootstrap samples.* $q_{min} = 30, q_{\max} = 60$ *(per year), EW* $\simeq 53.0$*. Units: thousands of dollars. Expected blocksize two years.*

8 Details of EW-LS strategies on historical market

Figure 8.1(a) shows the percentiles of the optimal fraction in stocks, versus time, in the historical market. Initially, the fraction in stocks is a bit less than 0.60. The median fraction drops smoothly down to zero near year 29. At the fifth percentile, complete de-risking occurs at about year 16. In the case of poor investment returns, the allocation to stocks is 0.60-0.80 at the 95th percentile.

Figure 8.1(b) shows the wealth percentiles in the historical market. We can see that W_T just approaches zero at the 5th percentile, at year 29. Again, we remind the reader that it is assumed that the retiree has real estate which can be used to fund a shortfall at less than the 5th percentile. The expected shortfall at the 5% level in this case is about -30.

Remark 8.1 (Shortfall hedge)**.** *Recall that we have assumed that the retiree has mortgage-free real estate, which is not considered to be part of the DC wealth. Assuming that a reverse mortgage can be obtained for one half the value of the real estate, this suggests that, in this case, real estate valued (in real terms)* $> 60{,}000$ *can manage a shortfall of* $30{,}000$*. Consequently, our advice to retirees is that they should select points on the efficient frontier where the ES in dollar terms,* [20] *is less than one-half the net value of their real estate.*

Finally, we can see from Figure 8.1(c) that the median withdrawal rapidly increases to the maximum withdrawal by year one.

The heat maps for the optimal fraction in stocks and the optimal withdrawals are shown in Figure 8.2. Figure 8.2(b) shows that the withdrawal control is approximately bang-bang, i.e. it is only ever optimal to withdraw $\mathfrak{q}_{\max}$ or $\mathfrak{q}_{\min}$ and nothing in between. For an explanation of this, see Forsyth (2022).

[20] More precisely here, we mean $|\min(ES{,}0)|$.

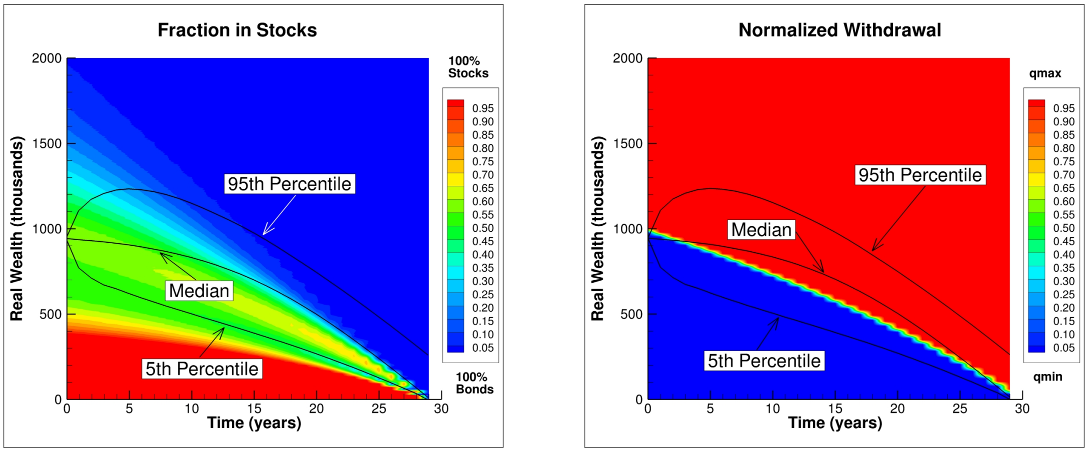

(a) *Fraction in stocks* (b) *Withdrawals*

FIGURE 8.2: *Optimal EW-LS. Heat map of controls: fraction in stocks and withdrawals, computed from Problem EW-LS (4.4). Real capitalization weighted CRSP index, and real 30-day T-bills. Scenario given in Table 5.2. Control computed and stored from the Problem 4.4 in the synthetic market.* $q_{min} = 30, q_{\max} = 60$ *(per year). EW* $\simeq 53.0$. *Percentiles from bootstrapped historical market. Normalized withdrawal* $(q - q_{\min})/(q_{\max} - q_{\min})$. *Units: thousands of dollars.*

8.1 Intuitive Explanation of the Optimal Controls

Superficially, these heat maps appear to be complex. However, in practice, these heat maps can be reduced to simple rules.

First, we draw attention to Figure 8.2(b). Due to the (approximate) bang-bang nature of the optimal control, the withdrawal control is conceptually simple. If the observed total wealth is in the blue region, withdraw at $q_{\min}$, otherwise, withdraw at $q_{\max}$. Basically, if the observed market returns are doing well, then withdraw at the maximum rate. If market returns are poor, withdraw at the minimum rate. It is only necessary to keep track of the location of the boundary separating the maximum withdrawal region from the minimum withdrawal region as a function of portfolio value and time.

Turning attention to the asset allocation heat map, Figure 8.2(a), we note the following. From the 5'th percentile to the median percentile, the allocation to stocks is in the range 50-60%. The red region occurs if the portfolio has done poorly. In this case, going to 100% stocks represents the *pulling the goalie* strategy.[21] On the other hand, note that there is a rapid reduction in allocation to stocks if the portfolio does well. There is a sharp transition between the green and the upper blue regions.

If we are willing to give up a bit in optimal performance, we can approximate the allocation heat map by the following rules. Set the stock allocation to 60%, unless you are well into the blue region in Figure 8.2(a). If you are fortunate, and your portfolio does well (blue region) then you can de-risk your portfolio completely.

Note that our allocation control approximation ignores the red region in Figure 8.2(a). As we shall see in Section 9, the efficient frontiers are not very sensitive to the asset allocation control, hence replacing the red region by a green approximation will not be largely sub-optimal.

To summarize, a financial advisor could use the following method of devising an optimal policy, and explaining the strategy a specific retiree.

[21] In ice hockey, if your team is down by one goal, with one minute left to play, then a common strategy is to pull your goalie, giving you an extra attacker. However, since your goal is empty this is a risky strategy.

- Determine the point on the LS frontier where the retiree is comfortable with the risk-reward tradeoff. This would require setting the minimum withdrawal requirements, and estimating the size of other assets (i.e. real estate) which could be used as hedge of last resort.

- Construct the asset allocation heat maps and the withdrawal heat maps.

- The withdrawal control is quite simple, withdraw the minimum in the blue region, and the maximum in the red region.

- An approximate, but slightly sub-optimal, policy for asset allocation would be to invest 60% in stocks unless in the blue region (asset allocation heat map). If the retiree is well into the blue region, then he can completely de-risk.

Summary advice to the retiree

- Allocate 60% to stocks, unless the market does really well, and then we can de-risk the portfolio. Most likely we will stick to 60% stocks for the first 10 years.

- Withdraw at the minimum rate for the first few years, and see how your portfolio does. If returns are good, we can withdraw at the maximum amount each year. Most likely, you will withdraw at the minimum rate for the first two-three years, and then withdraw at the maximum rate.

9 Sensitivity to Control Constraints

The base case in Table 5.2 has the minimum withdrawal $\mathfrak{q}_{\min} = 30$ and the maximum withdrawal $\mathfrak{q}_{\max} = 60$. Figure 9.1 shows the effect of varying these minimum and maximum withdrawal constraints, while keeping the stock allocation fraction in the range $[\mathfrak{p}_{\min}, \mathfrak{p}_{\max}] = [0,1]$. The optimal policy for each set $[\mathfrak{q}_{\min}, \mathfrak{q}_{\max}]$ is computed for the EW-LS objective function in the synthetic market. Each frontier is labelled with the range of $\mathfrak{q}_{\min} - \mathfrak{q}_{\max}$, instead of $[\mathfrak{q}_{\min},\mathfrak{q}_{\max}]$, for presentation clarity in these figures. The performance frontiers are shown in terms LS, ES and PS risk measures. Note that this is a true efficient frontier only in the (EW,LS) domain.

Figure 9.2 shows the base case in Table 5.2, using the EW-LS objective function, but varying the range of values of the fraction in stocks $[\mathfrak{p}_{\min}, \mathfrak{p}_{\max}]$, while keeping the withdrawal control range $[\mathfrak{q}_{\min}, \mathfrak{q}_{\max}] = [30,60]$. Each frontier is labelled with the permitted range $[\mathfrak{p}_{\min}, \mathfrak{p}_{\max}]$. The frontiers are shown in terms LS, ES and PS risk measures.

These results suggest that varying withdrawals has a much bigger effect on the efficient frontier (as a function of various risk measures) than variable asset allocation. In fact, even using a fairly narrow range of permitted fraction in stocks $[\mathfrak{p}_{\min}, \mathfrak{p}_{\max}] = [0.40, 0.60]$ does not lower the efficient frontier by much. This may be acceptable in the interests of simplicity.

Since the precise asset allocation does not have a large effect on the efficient frontier, this means that model misspecification, in terms of the asset allocation control, will not greatly affect the efficient frontier. This is consistent with our observations in Section 5.8, that the optimal controls were robust in out-of-sample bootstrap resampling tests.

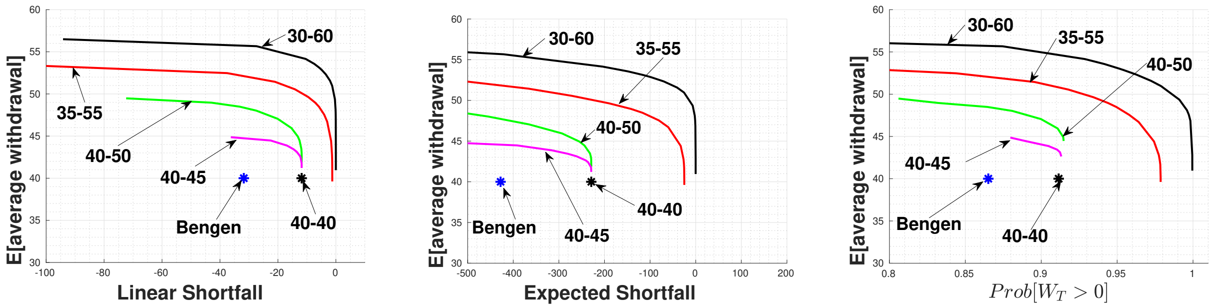

(a) *LS frontier, (EW,LS) domain.*

(b) *LS frontier, (EW,ES) domain.*

(c) *LS frontier, (EW,PS) domain.*

FIGURE 9.1: *Scenario in Table 5.2. EW-LS control computed from problem EW-LS Problem (4.4). Parameters based on the real CRSP index, and real 30-day T-bills (see Table A.1). Control computed from the Problem (4.4) in the synthetic market. Effect of varying* $[\mathfrak{q}_{min}, \mathfrak{q}_{\max}]$*, as labelled on the curves. Base case* $[\mathfrak{q}_{\min} = 30, \mathfrak{q}_{\max} = 60]$*. Stock allocation fraction* $[\mathfrak{p}_{\min}, \mathfrak{p}_{\max}] = [0{,}1]$*. Bengen (Bengen, 1994) strategy: rebalance to 50% stocks, constant withdrawal 40 per year. Units: thousands of dollars. All quantities inflation adjusted.*

10 Practical Implications

From a practical point of view, we suggest that EW-LS controls can be viewed as a decision rule that parallels, but improves upon, the four per cent rule.

How would this be applied in practice? An advisor would interview a prospective retiree, and elicit information such as (i) minimum real cash flows required (ii) maximum allocation to equities and (iii) value of other assets (which will usually be real estate). It will also be necessary to determine the disaster level ($\mathbb{W}$) of wealth at age 95. If there is no bequest motive, $\mathbb{W} = 0$ might be appropriate. If the retiree has a desire for a bequest, or if she is concerned with extreme longevity, then $\mathbb{W} > 0$ would be more in line with her concerns.

Then, the EW-LS criteria is used to generate an efficient frontier, satisfying the constraints. Each point on the frontier can be discussed in terms of (i) expected annual withdrawals (ii) probability of exhausting savings and (iii) expected shortfall. Some points on the frontier can be easily rejected, e.g, the expected shortfall is larger than the value of real estate, or the probability of ruin is too large. From the points that remain, the retiree will chose a point which is suitable, based on her risk preferences. Insurance products. which can hedge against extreme longevity, can also be discussed at this decision time.

Once a desirable point on the frontier is selected, then the heat maps for this point can be constructed. Each year, the retiree (or perhaps a wealth manager) determines the value of investment portfolio, and then consults the heat maps to determine the withdrawal amount, and the new asset allocation.

11 Conclusions

As noted in Anarkulova et al. (2025), retirees and wealth advisors demonstrate a revealed preference for spending rules for decumulation of DC pension plans. Almost all previous work on spending rules postulates heuristic strategies and tests these rules using historical data.

We follow a different methodology here. We determine the spending rules as the solution of an optimal stochastic control problem. The control problem is solved numerically, based on a

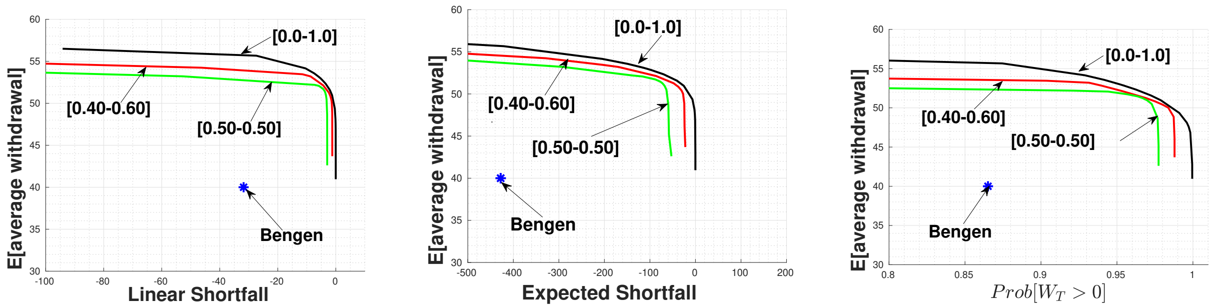

(a) *LS frontier, (EW,LS) domain.*

(b) *LS frontier, (EW,ES) domain.*

(c) *LS frontier, (EW,PS) domain.*

FIGURE 9.2: *Scenario in Table 5.2. EW-LS control computed from problem EW-LS Problem (4.4). Parameters based on the real CRSP index, and real 30-day T-bills (see Table A.1). Control computed from the Problem (4.4) in the synthetic market. Effect of varying fraction in stocks* $[\mathfrak{p}_{min}, \mathfrak{p}_{\max}]$*, as labelled on the curves. Base case* $[\mathfrak{p}_{\min}, \mathfrak{p}_{\max}] = [0,1]$*. Withdrawal range* $[\mathfrak{q}_{min}, \mathfrak{q}_{\max}] = [30,60]$*. Bengen (Bengen, 1994) strategy: rebalance to 50% stocks, constant withdrawal 40 per year. Units: thousands of dollars. All quantities inflation adjusted.*

parametric model of long term stock and bond returns.

For an optimal control problem, the first order of business is to specify the objective function, in terms of risk and reward. Since we allow variable withdrawals (subject to maximum and minimum constraints) we define reward as the sum of expected (real) withdrawals over a 30 year retirement (EW).

We focus on Expected Shortfall ES, Linear Shortfall LS, and Probability of Shortfall PS risk measures. We establish mathematically that, under certain assumptions, the set of optimal controls associated with all expected reward and expected shortfall EW-ES Pareto efficient frontier curves is identical to the set of optimal controls for all expected reward and linear shortfall EW-LS Pareto efficient frontier curves. This has the consequence that the set of optimal controls for EW-ES$_{t_0}$ are time consistent under the EW-LS$_{t_0}$ risk measure.

Fixing disaster level $\mathbb{W} = 0$ and probability level of $\alpha = 5\%$, based on our analysis and computational assessment of various risk measures, we conclude that risk as measured by linear shortfall LS, i.e. linearly weighting the final wealth below zero, is an appropriate risk measure.

As noted, the optimal EW-LS control is computed using a parametric market model. However, this control has been tested out-of-sample using block bootstrap resampling of historical data. These tests show that the optimal control is relatively robust to parameter misspecification.

Bootstrap resampling of historical data shows that the 4% rule (initial capital: one million, withdrawing 4% real of initial capital per year) has a probability of failure $> 10\%$, and expected shortfall ES(5%) $< -\$350{,}000$. In contrast, under bootstrap resampling tests, the EW-LS optimal control can withdraw 5% of initial wealth annually, on average, (adjusted for inflation) with a 98% probability of success, with an ES(5%) $\simeq -\$15{,}000$.

The EW-LS controls are dynamic. Both withdrawal amounts and stock allocation depend on the realized portfolio wealth (and time to go). However, the controls are summarized as easy to interpret heat maps, which makes implementation of these optimal controls straightforward.

Our numerical experiments also reveal that allowing flexible withdrawals has a much larger effect on the efficient frontier, compared to the asset allocation. For example, allowing flexible withdrawals, but constraining the fraction in equities to the range $[0.40, 0.60]$ results in strategies which are only slightly sub-optimal, compared to allowing the equity fraction be in the full range

[0.0, 1.0]. This may be an acceptable trade-off, i.e. giving up some efficiency in return for simplicity, and a cap on the maximum fraction allocated to stocks.

Finally, we note that the optimal controls can be computed directly from the bootstrapped resampled data, without specifying a parametric model of the underlying stock and bond processes. This requires use of machine learning techniques (Li and Forsyth, 2019; Ni et al., 2022; van Staden et al., 2023; 2024). These methods also allow use of more assets in terms of investment choices. We leave further study of machine learning techniques in the context of DC decumulation for future work.

12 Acknowledgements

Forsyth's work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2017-03760. Li's work was supported by a Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2020-04331.

13 Declaration

The authors have no conflicts of interest to report.

Appendix

A Parametric Model

We assume that the investor has access to two funds: a broad market stock index fund and a constant maturity bond index fund. The investment horizon is T. Let S_t and B_t respectively denote the real (inflation adjusted) *amounts* invested in the stock index and the bond index respectively. In general, these amounts will depend on the investor's strategy over time, as well as changes in the real unit prices of the assets. In the absence of an investor determined control (i.e. cash withdrawals or rebalancing), all changes in S_t and B_t result from changes in asset prices. We model the stock index as following a jump diffusion.

In addition, we follow the usual practitioner approach and directly model the returns of the constant maturity bond index as a stochastic process (see, e.g. Lin et al., 2015; MacMinn et al., 2014). As in MacMinn et al. (2014), we assume that the constant maturity bond index follows a jump diffusion process. Empirical justification for this can be found in Forsyth et al. (2022), Appendix A.

Let $S_{t^-} = S(t-\epsilon), \epsilon \to 0^+$, i.e. t^- is the instant of time before t, and let ξ^s be a random number representing a jump multiplier. When a jump occurs, $S_t = \xi^s S_{t^-}$. Allowing for jumps permits modelling of non-normal asset returns. We assume that $\log(\xi^s)$ follows a double exponential distribution (Kou, 2002; Kou and Wang, 2004). If a jump occurs, u^s is the probability of an upward jump, while $1-u^s$ is the chance of a downward jump. The density function for $y = \log(\xi^s)$ is

$$f^s(y) = u^s \eta_1^s e^{-\eta_1^s y} \mathbf{1}_{y \geq 0} + (1-u^s) \eta_2^s e^{\eta_2^s y} \mathbf{1}_{y<0} \ . \tag{A.1}$$

We also define

$$\gamma_\xi^s = E[\xi^s - 1] = \frac{u^s \eta_1^s}{\eta_1^s - 1} + \frac{(1-u^s)\eta_2^s}{\eta_2^s + 1} - 1 \ . \tag{A.2}$$

In the absence of control, S_t evolves according to

$$\frac{dS_t}{S_{t^-}} = \left(\mu^s - \lambda_\xi^s \gamma_\xi^s\right)\, dt + \sigma^s\, dZ^s + d\left(\sum_{i=1}^{\pi_t^s} (\xi_i^s - 1)\right), \tag{A.3}$$

where μ^s is the (uncompensated) drift rate, σ^s is the volatility, dZ^s is the increment of a Wiener process, π_t^s is a Poisson process with positive intensity parameter λ_ξ^s, and ξ_i^s are i.i.d. positive random variables having distribution (A.1). Moreover, ξ_i^s, π_t^s, and Z^s are assumed to all be mutually independent.

Similarly, let the amount in the bond index be $B_{t^-} = B(t-\epsilon), \epsilon \to 0^+$. In the absence of control, B_t evolves as

$$\frac{dB_t}{B_{t^-}} = \left(\mu^b - \lambda_\xi^b \gamma_\xi^b + \mu_c^b \mathbf{1}_{\{B_{t^-}<0\}}\right)\, dt + \sigma^b\, dZ^b + d\left(\sum_{i=1}^{\pi_t^b} (\xi_i^b - 1)\right), \tag{A.4}$$

where the terms in equation (A.4) are defined analogously to equation (A.3). In particular, π_t^b is a Poisson process with positive intensity parameter λ_ξ^b, and ξ_i^b has distribution

$$f^b(y = \log \xi^b) = u^b \eta_1^b e^{-\eta_1^b y} \mathbf{1}_{y \geq 0} + (1 - u^b) \eta_2^b e^{\eta_2^b y} \mathbf{1}_{y<0}\ , \tag{A.5}$$

and $\gamma_\xi^b = E[\xi^b - 1]$. ξ_i^b, π_t^b, and Z^b are assumed to all be mutually independent. The term $\mu_c^b \mathbf{1}_{\{B_{t^-}<0\}}$ in equation (A.4) represents the extra cost of borrowing (the spread).

The diffusion processes are correlated, i.e. $dZ^s \cdot dZ^b = \rho_{sb}\, dt$. The stock and bond jump processes are assumed mutually independent. See Forsyth (2020b) for justification of the assumption of stock-bond jump independence.

We use the threshold technique (Mancini, 2009; Cont and Mancini, 2011; Dang and Forsyth, 2016) to estimate the parameters for the parametric stochastic process models. Since the index data is in real terms, all parameters reflect real returns. Table A.1 shows the results of calibrating the models to the historical data. The correlation ρ_{sb} is computed by removing any returns which occur at times corresponding to jumps in either series, and then using the sample covariance. Further discussion of the validity of assuming that the stock and bond jumps are independent is given in Forsyth (2020b).

CRSP	μ^s	σ^s	λ^s	u^s	η_1^s	η_2^s	ρ_{sb}
	0.087323	0.147716	0.316326	0.225806	4.3591	5.53370	0.095933
30-day T-bill	μ^b	σ^b	λ^b	u^b	η_1^b	η_2^b	ρ_{sb}
	0.0032	0.0140	0.3878	0.3947	61.5350	53.4043	0.095933

TABLE A.1: *Parameters for parametric market models (A.3 and (A.4, fit to CRSP data (inflation adjusted) for 1926:1-2023:12.*

B Numerical Techniques

We solve problems (4.4) using the techniques described in detail in Forsyth and Labahn (2019); Forsyth (2020a; 2022). We give only a brief overview here.

We localize the infinite domain to $(s,b) \in [s_{\min}, s_{\max}] \times [b_{\min}, b_{\max}]$, and discretize $[b_{\min},b_{\max}]$ using an equally spaced $\log b$ grid, with n_b nodes. Similarly, we discretize $[s_{\min}, s_{\max}]$ on an equally spaced $\log s$ grid, with n_s nodes. For case $b < 0$, we define a reflected grid $b' = -b$, with the $n_b \times n_s$ nodes. This represents the insolvent case nodes. The PIDE for $b' > 0$ has the same form as for $b > 0$. This idea can be used more generally if leverage is permitted, which we do not explore in this work. Localization errors are minimized using the domain extension method in Forsyth and Labahn (2019).

At rebalancing dates, we solve the local optimization problem by discretizing $(\mathfrak{q}(\cdot), \mathfrak{p}(\cdot))$ and using exhaustive search. Between rebalancing dates, we solve a two dimensional partial integro-differential equation (PIDE) using Fourier methods (Forsyth and Labahn, 2019; Forsyth, 2022). Finally, in the case of EW-ES, the outer optimization over $\mathbb{W}$ is solved using a one-dimensional method.

We used the value $\epsilon = -10^{-4}$ in equation (4.4), which forces the investment strategy to be bond heavy if the remaining wealth in the investor's account is large, and $t \to T$. Using this small value gave the same results as $\epsilon = 0$ for the summary statistics, to four digits. This is simply because the states with very large wealth have low probability. However, this stabilization procedure produced smoother heat maps for large wealth values, without altering the summary statistics appreciably.

B.1 Convergence Test: Synthetic Market

Table B.1 shows a detailed convergence test for the base case problem given in Table 5.2, for the EW-ES problem. The results are given for a sequence of grid sizes, for the dynamic programming algorithm in (Forsyth, 2022) and Appendix B. The dynamic programming algorithm appears to converge at roughly a second order rate. The optimal control computed using dynamic programming is stored, and then used in Monte Carlo computations. The Monte Carlo results are in good agreement with the dynamic programming solution. For all the numerical examples, we will use the 2048×2048 grid, since this seems to be accurate enough for our purposes.

	Algorithm in (Forsyth, 2022) and Appendix B			Monte Carlo	
Grid	LS	$E[\sum_i \mathfrak{q}_i]/M$	Value Function	LS	$E[\sum_i \mathfrak{q}_i]/M$
512×512	-1.40884	50.9082	1484.981	-1.26443	50.938
1024×1024	-1.32050	50.9491	1488.864	-1.27396	50.953
2048×2048	-1.30148	50.9643	1489.880	-1.28189	50.963

TABLE B.1: *EW-LS convergence test. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1. The Monte Carlo method used* 2.56×10^6 *simulations. The MC method used the control from the solving the PIDEs as described in Appendix B.* $\kappa = 30, \mathbb{W} = 0.0$. *Grid refers to the grid used in the Algorithm in Appendix B:* $n_x \times n_b$, *where* n_x *is the number of nodes in the* $\log s$ *direction, and* n_b *is the number of nodes in the* $\log b$ *direction. Units: thousands of dollars (real).* M *is the total number of withdrawals (rebalancing dates).*

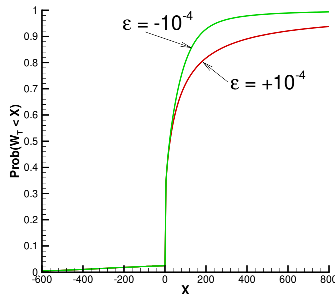

FIGURE C.1: *EW-PS CDF of the terminal wealth, with stabilization parameters shown. Point on curve where $EW \simeq 53.0$. Real stock index: deflated real capitalization weighted CRSP, real bond index: deflated 30 day T-bills. Scenario in Table 5.2. Parameters in Table A.1. Synthetic market.*

C Effect of Stabilization term

Recall that the optimization problem, for all objective functions, becomes ill-posed along any path where $W_t \gg \mathbb{W}$, $t \to T$. To remove this problem, the stabilization term

$$\epsilon E[W_T] \tag{C.1}$$

is added to each objective function. We set $|\epsilon| \ll 1$, to ensure that this term has little effect unless we are in the ill-posed region. Essentially, if $\epsilon < 0$, then this forces the portfolio to invest 100% in bonds. On the other hand, if $\epsilon > 0$, then the portfolio will invest 100% in stocks. We remark here that these choices are essentially arbitrary: by assumption, the 95-year old retiree has a short life expectancy, and has large wealth, so that even with maximum withdrawals, there is almost zero probability of running out of cash.

To verify that the choice of positive or negative ϵ has little effect near $W_t = \mathbb{W}$, Figure C.1 shows the CDF curves for the EW-PS strategy (EW $= 53.0$), for both positive and negative ϵ. We can see that both curves overlap for $W_T \leq 100$. Consequently, left tail risk measures will be identical for both cases, and there will be no differences in average withdrawals, since we will be constrained by the maximum withdrawal specification. To the right of $W_T = 100$, we can see that for $\epsilon > 0$, there is higher probability of obtaining larger W_T compared to the case $\epsilon < 0$. This is of course expected, since investing in all stocks, (when W_t is large) will have a larger expected portfolio value.

The CDF curves for $\pm\epsilon$ for EW-ES and EW-LS policies are similar.

References

Ackerly, A., N. N, and D. Nikles (2021). To spend or not to spend? BlackRock white paper, `https://www.blackrock.com/us/individual/literature/whitepaper/spending-retirement-assets-final-whitepaper.pdf`.

Anarkulova, A., S. Cederburg, and M. S. O'Doherty (2022). Stocks for the long run? Evidence from a broad sample of developed markets. *Journal of Financial Economics 143:1*, 409–433.

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*, 464–500.

Azimzadeh, P. and P. A. Forsyth (2015). The existence of optimal bang-bang controls for GMxB contracts. *SIAM Journal on Financial Mathematics 6*, 117–139.

Bacinello, A. R., R. Maggistro, and I. Zoccolan (2024). Risk-neutral valuation of glwb riders in variable annuities. *Insurance: Mathematics and Economics 114*, 1–14.

Bannerje, S. (2021). Decoding retiree spending. R. Rowe Price Insights on Retirement, `https://www.troweprice.com/institutional/us/en/insights/articles/2021/q1/decoding-retiree-spending-na.html`.

Bauer, D., A. Kling, and J. Russ (2008). A universal pricing framework for guaranteed minimum benefits in variable annuities. *ASTIN Bulletin 38*, 621–651.

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

Bengen, W. P. (2001). Conserving client portfolios during retirement, part IV. *Journal of Financial Planning 14:5*, 110–119.

Bernhardt, T. and C. Donnelly (2018). Pension decumulation strategies: A state of the art report. Technical Report, Risk Insight Lab, Heriot Watt University.

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

Bjork, T. and A. Murgoci (2010). A general theory of Markovian time inconsistent stochastic control problems. SSRN 1694759.

Bjork, T. and A. Murgoci (2014). A theory of Markovian time inconsisent stochastic control in discrete time. *Finance and Stochastics 18*, 545–592.

Browning, C., R. Guo, Y. Cheng, and M. S. Finke (2016). Spending in retirement: determining the consumption gap. *Journal of Financial Planning 26(2)*, 42–53.

Carrick, R. (2020). Have a defined contribution pension plan? here's how to make the most of it. *The Globe and Mail Feb. 20*.

Cavaglia, S., L. Scott, K. Blay, and S. Hixon (2022). Multi-asset class factor premia: A strategic asset allocation perspective. *The Journal of Portfolio Management 48:9*, 14–32.

Chen, A., P. Hieber, and J. K. Klein (2019). Tonuity: A novel individual-oriented retirement plan. *ASTIN Bulletin 49*, 5–30.

Chen, A. and M. Rach (2022). The tontine puzzle. SSRN 4106903.

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.

Chen, Z. and P. A. Forsyth (2008). A numerical scheme for the impulse control formulation for pricing variable annuities with a Guaranteed Minimum Withdrawal Benefit (GMWB). *Numericshe Mathematik 109*, 535–569.

Chen, Z., K. R. Vetzal, and P. A. Forsyth (2008). The effect of modelling parameters on the value of GMWB guarantees. *Insurance: Mathematics and Economics 43*, 165–173.

Cogneau, P. and V. Zakalmouline (2013). Block bootstrap methods and the choice of stocks for the long run. *Quantitative Finance 13*, 1443–1457.

Cont, R. and C. Mancini (2011). Nonparametric tests for pathwise properties of semimartingales. *Bernoulli 17*, 781–813.

Dai, M., K. Y. Kwok, and J. Zong (2008). Guaranteed minimum withdrawal benefit in variable annuity. *Mathematical Finance 18:4*, 595–611.

Dang, D.-M. and P. A. Forsyth (2014). Continuous time mean-variance optimal portfolio allocation under jump diffusion: a numerical impulse control approach. *Numerical Methods for Partial Differential Equations 30*, 664–698.

Dang, D.-M. and P. A. Forsyth (2016). Better than pre-commitment mean-variance portfolio allocation strategies: a semi-self-financing Hamilton-Jacobi-Bellman equation approach. *European Journal of Operational Research 250*, 827–841.

Dichtl, H., W. Drobetz, and M. Wambach (2016). Testing rebalancing strategies for stock-bond portfolios across different asset allocations. *Applied Economics 48*, 772–788.

Donnelly, C., M. Guillén, and J. B. Nielsen (2014). Bringing cost transparency to the life annuity market. *Insurance: Mathematica and Economics 56*, 14–27.

Dus, I., R. Maurer, and O. Mitchell (2005). Betting on death and capital markets in retirement: a shortfall analysis of life annuities versue phased withdrawal plans. *Financial Services Review 14*, 169–196.

Feng, R. and B. Yi (2019). Quantitative modeling of risk management strategies: Stochastic reserving and hedging of variable annuity guaranteed benefits. *Insurance: Mathematics and Economics 85(c)*, 60–73.

Forsyth, P. and G. Labahn (2019). $\epsilon-$Monotone Fourier methods for optimal stochastic control in finance. *Journal of Computational Finance 22:4*, 25–71.

Forsyth, P., K. Vetzal, and G. Westmacott (2024). Optimal performance of a tontine overlay subject to withdrawal constraints. *ASTIN Bulletin 54:1*, 94–128.

Forsyth, P. A. (2016). Variable annuities: Fees too high or too low? Global Risk Institute White Paper, `https://globalriskinstitute.org/publication/variable-annuities/`.

Forsyth, P. A. (2020a). 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. (2020b). Optimal dynamic asset allocation for DC plan accumulation/decumulation: Ambition-CVAR. *Insurance: Mathematics and Economics 93*, 230–245.

Forsyth, P. A. (2021). Two stage decumulation strategies for DC plan investors. *International Journal of Theoretical and Applied Finance 24:1*. Article 2150007, 1-31.

Forsyth, P. A. (2022). A stochastic control approach to defined contribution plan decumulation: The nastiest, hardest problem in finance. *North American Actuarial Journal 26:2*, 227–251.

Forsyth, P. A. and K. R. Vetzal (2014). An optimal stochastic control framework for determining the cost of hedging of variable annuities. *Journal of Economic Dynamics and Control 44*, 29–53.

Forsyth, P. A. and K. R. Vetzal (2019). Optimal asset allocation for retirement savings: deterministic vs. time consistent adaptive strategies. *Applied Mathematical Finance 26:1*, 1–37.

Forsyth, P. A., K. R. Vetzal, and G. Westmacott (2022). Optimal control of the decumulation of a retirement portfolio with variable spending and dynamic asset allocation. *ASTIN Bulletin 51:3*, 905–938.

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

Guyton, J. T. and W. J. Klinger (2006). Decision rules and maximum initial withdrawal rates. *Journal of Financial Planning 19:3*, 45–58.

Habib, F., H. Huang, A. Mauskopf, B. Nikolic, and T. S. Salisbury (2020). Optimal allocation to deferred income annuities. *Insurance: Mathematics and Economics 90*, 94–104.

Hamilton, M. (2001, December). The Financial Circumstances of Elderly Canadians and the Implications for the Design of Canada's Retirement Income System. In P. Grady and A. Sharpe (Eds.), *The State of Economics in Canada: Festschrift in Honour of David Slater*, The State of Economics in Canada: Festschrift in Honour of David Slater, pp. 225–253. Centre for the Study of Living Standards.

Hatch, J. E. and R. W. White (1985). Canadian stocks, bonds, bills and inflation: 1950-1983. Financial Analysts Research Foundation Mongraph Series, Number 19.

Hill, C. (2016). Older people fear this more than death. `https://www.marketwatch.com/story/older-people-fear-this-more-than-death-2016-07-18`.

Horneff, V., R. Maurer, O. S. Mitchell, and R. Rogalla (2015). Optimal life cycle portfolio choice with variable annuities offering liquidity and investment downside protection. *Insurance: Mathematics and Economics 63*, 91–107.

Horneff, W. J., R. Maurer, O. Mitchell, and I. Dus (2008). Following the rules: Integrating asset allocation and annuitization in retirement portfolios. *Insurance: Mathematica and Economics 42*, 396–408.

Irlam, G. (2014). Portfolio size matters. *Journal of Personal Finance 13*(2), 9–16.

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

Kou, S. G. and H. Wang (2004). Option pricing under a double exponential jump diffusion model. *Management Science 50*, 1178–1192.

Landriault, D., B. Li, D. Li, and Y. Wang (2021). High-water mark fee structure in variable annuities. *Journal of Risk and Insurance 88*, 1057–1094.

Li, Y. and P. A. 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.

Lin, Y., R. MacMinn, and R. Tian (2015). De-risking defined benefit plans. *Insurance: Mathematics and Economics 63*, 52–65.

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

MacMinn, R., P. Brockett, J. Wang, Y. Lin, and R. Tian (2014). The securitization of longevity risk and its implications for retirement security. In O. S. Mitchell, R. Maurer, and P. B. Hammond (Eds.), *Recreating Sustainable Retirement*, pp. 134–160. Oxford: Oxford University Press.

Mancini, C. (2009). Non-parametric threshold estimation models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics 36*, 270–296.

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

Milevsky, M. A. and V. R. Young (2007a). Annuitization and asset allocation. *Journal of Economic Dynamics and Control 31*, 3238–3177.

Milevsky, M. A. and V. R. Young (2007b). The timing of annuitization: investment dominance and mortality risk. *Insurance: Mathematics and Economics 40*, 135–144.

Monig, T. and D. Bauer (2016). Revisiting the risk-neutral approach to optimal policyholder behaviour: a study of withdrawal guarantees in variable annuities. *Review of Finance 20*, 759–794.

Ni, C., Y. Li, P. A. Forsyth, and R. Caroll (2022). Optimal asset allocation for outperforming a stochastic benchmark. *Quantitative Finance 22(9)*, 1595–1626.

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

Peijnenburg, K., T. Nijman, and B. J. Werker (2016). The annuity puzzle remains a puzzle. *Journal of Economic Dynamics and Control 70*, 18–35.

Pfau, W. D. (2015). Making sense out of variable spending strategies for retirees. *Journal of Financial Planning 28*(10), 42–51.

Pfau, W. D. (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.

Piscopo, G. and S. Haberman (2011). The valuation of guaranteed lifelong withdrawal benefit options in variable annuity contracts and the impact of mortality risk. *North American Actuarial Journal 15(1)*, 59–76.

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

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

Powell, W. (2025). On state variables. https://castle.princeton.edu/statevariables/.

Rappaport, A. (2019). Insights on spending and asset management in retirement. SOA report, https://www.soa.org/499c88/globalassets/assets/files/resources/research-report/2019/2019-spending-asset-management-report.pdf.

Ritholz, B. (2017). Tackling the 'nastiest, hardest problem in finance'. www.bloomberg.com/view/articles/2017-06-05/tackling-the-nastiest-hardest-problem-in-finance.

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.

Simonian, J. and A. Martirosyan (2022). Sharpe parity redux. *The Journal of Portfolio Management 48:9*, 183–193.

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.

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

van Staden, P. M., D.-M. Dang, and P. Forsyth (2021). The surprising robustness of dynamic mean-variance portfolio optimization to model misspecification errors. *European Journal of Operational Research 289:2*, 74–792.

van Staden, P. M., P. A. Forsyth, and Y. Li (2024). A global-in-time neural network approach to dynamic portfolio optimization. *Applied Mathematical Finance 31:3*, 131–163.

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

Vigna, E. (2022). Tail optimality and preferences consistency for intertemporal optimization problems. *SIAM Journal on Financial Mathematics 13:1*. Article 10.1137/21M1435422, electronic.

Waring, M. B. and L. B. Siegel (2015). The only spending rule you will ever need. *Financial Analysts Journal 71:1*, 91–107.