Optimal Asset Allocation for DC Pension Decumulation with a Variable Spending Rule

Peter A. Forsyth[a] Kenneth R. Vetzal[b] Graham Westmacott[c]

Revised: March 12, 2020

Abstract

We determine the optimal asset allocation to bonds and stocks using an Annually Recalculated Virtual Annuity (ARVA) spending rule for DC pension plan decumulation. Our objective function minimizes downside withdrawal variability for a given fixed value of total expected withdrawals. The optimal asset allocation is found using optimal stochastic control methods. We formulate the strategy as a solution to a Hamilton Jacobi Bellman (HJB) Partial Integro Differential Equation (PIDE). We impose realistic constraints on the controls (no shorting, no leverage, discrete rebalancing), and solve the HJB PIDEs numerically. Compared to a fixed weight strategy which has the same expected total withdrawals, the optimal strategy has a much smaller average allocation to stocks, and tends to de-risk rapidly over time. This conclusion holds in the case of a parametric model based on historical data, and also in a bootstrapped market based on the historical data.

Keywords: DC pension plans, decumulation, optimal control, HJB equation, annually recalculated virtual annuity

1 Introduction

Throughout the developed economies, defined benefit (DB) pension plans are disappearing. In an effort to reduce corporate risk exposures, firms are moving employees to defined contribution (DC) plans. In a typical DC plan, the employee and employer contribute a fraction of the employee's annual salary to a tax-advantaged fund. The employee then can choose to invest the funds in a variety of investment vehicles. These usually include bond and equity index funds. A common default choice in the US is a Target Date Fund (TDF). In a TDF, the fraction invested in stocks declines over time in a prescribed manner, according to the anticipated retirement date.

However, once employees reach retirement age, they are faced with constructing a decumulation strategy. In other words, pensioners then face the difficult problems of (i) calculating how much to withdraw from their investment portfolios each year; and (ii) determining their asset allocation strategies.

[a]David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, paforsyt@uwaterloo.ca, +1 519 888 4567 ext. 34415.

[b]School of Accounting and Finance, University of Waterloo, Waterloo ON, Canada N2L 3G1, kvetzal@uwaterloo.ca, +1 519 888 4567 ext. 36518.

[c]PWL Capital, 20 Erb Street W., Suite 506, Waterloo, ON, Canada N2L 1T2, gwestmacott@pwlcapital.com, +1 519 880 0888.

The shift to a DC plan transfers income and longevity risk from the firm to the plan member. The plan member is now exposed to investment risk during both the accumulation and decumulation phases, as well as longevity risk during decumulation. It might seem that use of an annuity would be a good strategy for decumulation, since this eliminates both of these risks during that phase. However, it is well known that when given a choice, few investors annuitize (Peijnenburg et al., 2016). While this has been traced to the investor's desire to retain control over their portfolio, there are many other potential factors in play as well. For example, MacDonald et al. (2013) list more than three dozen reasons to avoid annuitization. Among these are the general lack of true inflation protection and the fact that in many cases, annuities are poorly priced in practice.

A well-known decumulation strategy relies on the *4% rule*. This can be traced to the work of Bengen (1994), which was based on historical backtests of a portfolio that had 50% invested in stocks and 50% invested in bonds, rebalanced annually. The final recommendation in Bengen (1994) was that withdrawing initially at a rate of 4% of the initial portfolio value and with successive withdrawals adjusted for inflation was a safe strategy, in that the investor would have never run out of funds over any rolling 30-year historical period considered. More generally, there is a large academic literature on decumulation strategies. A small but representative sample includes Blake et al. (2003), Gerrard et al. (2004; 2006), Smith and Gould (2007), Milevsky and Young (2007), Freedman (2008), and Liang and Young (2018). For overviews of various strategies for decumulation, see MacDonald et al. (2013) and Bernhardt and Donnelly (2018).

We emphasize that throughout this work we focus on real cash flows and real investment returns. This follows the tradition of sources such as Bengen (1994) who concentrated on inflation-adjusted withdrawals. The investment period we consider is sufficiently long that, even in a sustained environment of low inflation, the cumulative effects of inflation on purchasing power in the long run should not be ignored.

Rather than specifying a strategy which withdraws a fixed real amount each year, we consider a strategy that responds to the actual investment experience. A recent suggestion is based on an *Annually Recalculated Virtual Annuity* (ARVA) (Waring and Siegel, 2015; Westmacott and Daley, 2015). An ARVA rule can summarized as follows:

> *"Each year, one should spend (at most) the amount that a freshly purchased annuity–with purchase price equal to the then-current portfolio and priced at the current interest rates and number of years of required cash flows remaining–would pay out in that year."* (Waring and Siegel, 2015, p. 91)

The trade-off here is that real withdrawals will fluctuate in order to prevent the possibility of running out of cash. Of course, the downside is that in the event of very poor portfolio investment performance, the withdrawals may become minuscule (though the value of the ARVA portfolio will not drop to zero). Waring and Siegel (2015) justify this strategy by contending that DC plan decumulation is basically an annuitization problem. This does not require an actual annuity (hence the use of *virtual* in the ARVA designation), but does require *annuity thinking*. Note that longevity is taken into account in an approximate way by setting the current cash flow horizon to be the date at which 80% of the investors (at that time) have passed away. This is discussed in more detail later in this article.

Waring and Siegel (2015) focus exclusively on the withdrawal rule, as opposed to the asset allocation strategy. Our objective here is to optimize the asset allocation strategy for an ARVA type decumulation rule. We consider two measures of performance:

- The expected total real (i.e. inflation-adjusted) withdrawals over the lifetime of the strategy.
- The expected downside variability of the withdrawals from one year to the next.

This results in a multi-objective optimization problem, which we solve using a scalarization approach. In other words, for a fixed value of the expected total withdrawal, we find the strategy which gives us the smallest possible withdrawal variability. We formulate this problem as a Hamilton-Jacobi-Bellman (HJB) Partial Integro Differential Equation (PIDE), which we solve numerically. This allows us to impose realistic constraints on our strategy, i.e. no short-selling or leverage and infrequent (yearly) rebalancing.

Of course, another reason to avoid annuitization is the desire to leave a bequest. This issue was discussed in Forsyth et al. (2019), where the decumulation problem was formulated based on fixed withdrawal amounts and the risk measure was based on the final wealth distribution. However, in principle, the ARVA spending rule is designed to spend down the investor's wealth by the end of the cash flow horizon, so adding a risk measure based on terminal wealth conflicts with the ARVA philosophy.

Using parameters of a stochastic model estimated from 90 years of market data, we compute and store the optimal strategy from the numerical solution of the HJB equation. We then use this strategy in Monte Carlo simulations to generate statistics of interest. One set of simulations is based on the same parametric stochastic model used to determine the investment strategy. We label this type of simulation the *synthetic market*. As a robustness check, we also use the stored strategy in historical bootstrap backtests. We refer to the market based on bootstrap resampling as the *historical market*. For comparison purposes, in addition to our optimal strategies we also test fixed weight strategies with ARVA withdrawals in both the synthetic and historical markets.

The outline of this paper is as follows. Section 2 formulates the assumptions concerning the underlying stochastic processes. Section 3 defines the ARVA spending rule, and Section 4 gives a mathematical description of the optimal asset allocation problem. Section 5 describes the numerical method used to solve for the optimal asset allocation. Section 6 describes the data and calibration methods. A convergence test is given in Section 7. Illustrative results are provided in Sections 8 and 9 for the synthetic and bootstrapped historical markets. Section 10 provides some general conclusions.

2 Formulation

For simplicity we assume that there are only two assets available in the financial market, namely a risky asset and a risk-free asset. In practice, the risky asset would be a broad equity market index fund. We assume that the investor's allocation to these two assets is rebalanced periodically, not continuously.

The investment horizon (i.e. the total time of the decumulation phase) is T. S_t and B_t respectively denote the *amounts* invested in the risky and risk-free assets at time t, $t \in [0, T]$. The investor's total wealth at time t is defined as

$$\text{Total wealth} \equiv W_t = S_t + B_t. \tag{2.1}$$

To reduce subscript clutter, we will occasionally use the notation $S_t \equiv S(t)$, $B_t \equiv B(t)$ and $W_t \equiv W(t)$. Since we focus on real cash flows, S_t and B_t should be seen as being in real terms.

In general, S_t and B_t will depend on the investor's strategy over time, as well as changes in the real unit prices of these assets. Absent a control determined by the investor (i.e. withdrawing funds or rebalancing the portfolio), S_t and B_t will only change as a result of movements in real asset prices. In this case (absence of control), we assume that S_t follows a jump diffusion process under the objective measure (i.e. using real probabilities, not risk-neutral probabilities) of the form

$$\frac{dS_t}{S_{t^-}} = (\mu - \zeta_\xi E[\xi - 1])\, dt + \sigma\, d\mathcal{Z} + d\left(\sum_{i=1}^{\pi_t} (\xi_i - 1)\right), \tag{2.2}$$

where $S_{t^-} = \lim_{\epsilon \to 0^+} S_{t-\epsilon}$, μ is the uncompensated drift rate, σ is the diffusive volatility, $\mathcal{Z}$ is a standard Brownian motion, π_t is a Poisson process with positive intensity parameter ζ_ξ, and ξ_i are i.i.d. positive random variables. We assume that ξ_i, π_t, and $\mathcal{Z}$ are all mutually independent. Equation (2.2) augments standard geometric Brownian motion with occasional discontinuous jumps. Adding in jumps allows us to incorporate the effects of large market movements (e.g. market crashes) into our analysis. As we only consider cases where the portfolio is discretely rebalanced, the jump process models the cumulative effects of substantial changes in the real price of the risky asset between rebalancing times.

When a jump occurs, $S_t = \xi S_{t^-}$. We assume that $\log(\xi)$ follows a double exponential distribution (Kou, 2002). Conditional on a jump occurring, p_u is the probability of an upward jump, while $p_d = 1 - p_u$ is the chance of a downward jump. The density function for $y = \log(\xi)$ is

$$f(y) = p_u\, \eta_1 e^{-\eta_1 y} \mathbf{1}_{y \geq 0} + p_d\, \eta_2 e^{\eta_2 y} \mathbf{1}_{y < 0}. \tag{2.3}$$

Given (2.3), the expected value of the jump multiplier ξ is

$$E[\xi] = \frac{p_u\, \eta_1}{\eta_1 - 1} + \frac{p_d\, \eta_2}{\eta_2 + 1}. \tag{2.4}$$

Letting $\kappa = E[\xi - 1]$, we can write equation (2.2) more informally as

$$\frac{dS_t}{S_{t^-}} = (\mu - \zeta_\xi \kappa)\, dt + \sigma\, d\mathcal{Z} + (\xi - 1)\, d\mathbb{Q} \tag{2.5}$$

where $d\mathbb{Q} = 1$ with probability $\zeta_\xi\, dt$ and $d\mathbb{Q} = 0$ with probability $1 - \zeta_\xi\, dt$.

In the absence of control, we assume that the dynamics of the amount B_t invested in the risk-free asset are simply

$$dB_t = r B_t\, dt, \tag{2.6}$$

where r is the constant real risk-free rate. This is obviously a gross simplification of the actual bond market, but it allows us to compute a relatively simple asset allocation strategy.

Summarizing the model, we assume that there is a constant risk-free interest rate so that the amount invested in the risk-free asset is described by equation (2.6), apart from times when the investor either rebalances the portfolio or withdraws cash from it. Similarly excluding rebalances and withdrawals, the amount invested in the risky asset is described by the jump diffusion model (2.2). Equations (2.2) and (2.6) represent what we referred to earlier as the synthetic market. These equations are highly simplified models of actual stock and bond index processes. We use them to determine the asset allocation strategy in the synthetic market. We first evaluate the performance of the strategy using Monte Carlo simulations driven by the same underlying assumptions (i.e. in the synthetic market), but we subsequently assess the performance of the same strategy in the historical market, i.e. using bootstrapped historical bond and stock market returns. This second set of tests subjects the strategy to complications such as stochastic interest rates, correlations in stock and bond indices, and random changes in asset volatilities. As will be seen below, the performance of the strategy in the historical market is very similar to that found in the synthetic market, indicating that the parsimonious representation (2.2)-(2.6) suffices to determine an asset allocation strategy in our context.

3 ARVA Spending Rule

The ARVA spending rule is based on the following idea. Each year, a virtual annuity is constructed based on the current portfolio value, the current number of remaining years of required cash flows, and the prevailing real interest rate. The withdrawal amount in a given year is then based on the fixed virtual annuity payment per year. The virtual annuity is recomputed each year, so the annual payments will fluctuate in response to the investment experience. In contrast to the ubiquitous 4% rule, ARVA is efficient, in the sense that all accumulated assets are withdrawn at the end of the period of required cash flows.

Note that the ARVA rule is based on determining the number of years of remaining cash flows required, at each withdrawal time. As discussed by Waring and Siegel (2015) and Westmacott (2017), taking into account mortality can front load spending into periods when retirees are more active. Waring and Siegel (2015) observe that simply basing the virtual annuity on the remaining life expectancy results in very high front load spending with a rather large drop in spending in later years. Waring and Siegel end up suggesting a blend of current expected life expectancy and the maximum possible lifespan.

We use the methods suggested in Westmacott (2017) to add a mortality boost to the ARVA spending rule. To avoid being overly conservative and assuming a maximum possible lifespan (117 is the oldest recorded Canadian) we assume that retirees are merely in the top 20% as measured by longevity. Suppose that a retiree is x years old at $t = 0$. Assuming that the $x + t$ year old retiree is alive at time t, let $T_x^*(t)$ be the time at which 80% of the cohort of $x + t$ year olds have passed away, given that all members of the cohort were alive at time t. In other words, conditional on an investor being alive at time t, $T_x^*(t)$ is the time at which there is just a 20% chance that this investor will still be alive.

This rule provides some front end spending, while not reducing spending too precipitously during later years. The relative size of front load spending to back end spending can be adjusted by varying the fraction of the cohort assumed to have passed away. We emphasize that these spending rules are always based on assuming an annuity which pays out for the entire remaining years of required cash flows. This is, of course, not the same withdrawal amount as a currently purchased lifetime annuity (in general).

Given the real interest rate r, the present value of a continuously paid annuity, which pays at a rate of one dollar per year, for $(T_x^*(t) - t)$ years, is denoted by the annuity factor $a(t)$ where

$$a(t) \quad = \quad \frac{1 - \exp[-r(T_x^*(t) - t)]}{r}. \tag{3.1}$$

Consequently, the continuous real annuity payment for $T_x^*(t) - t$ years that can be purchased at time t with wealth $W(t)$ is $W(t)/a(t)$. Consider a set of withdrawal times $\mathcal{T}$

$$\mathcal{T} \equiv \{t_0 = 0 < t_1 < \cdots < t_M = T\}, \tag{3.2}$$

where $t = 0$ denotes the time that the x year old retiree begins to withdraw money from the DC plan. We specify that any two points of $\mathcal{T}$ are equidistant with $t_i - t_{i-1} = \Delta t = T/M$, $i = 1, \ldots, M$. In the following we will let $\Delta t = 1$ year. If we restrict ourselves to annual payments at times t_i, we can convert the continuous payment above into a lump sum received in advance of the interval $[t_i, t_{i+1}]$. The lump sum (i.e. the withdrawal at t_i) is $W(t_i)A(t_i)$, where the ARVA multiplier $A(t_i)$ is given by

$$\begin{aligned} A(t_i) &= \int_{t_i}^{t_{i+1}} \frac{e^{-r(t'-t_i)}}{a(t)} \, dt' \ , \\ &\text{withdrawal at } t_i = A(t_i)W(t_i) \ . \end{aligned} \tag{3.3}$$

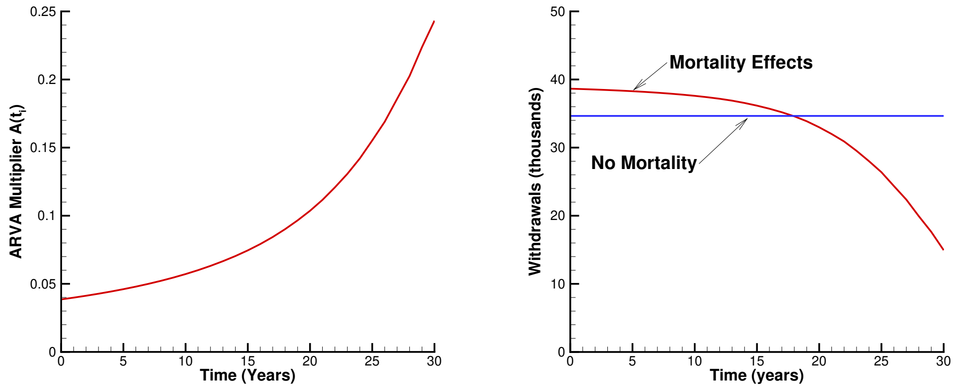

(a) *ARVA multiplier $A(t_i)$.* (b) *ARVA annual spending.*

FIGURE 3.1: *ARVA multiplier and spending amounts per year. The real interest rate $r = .0048$. CPM 2014 mortality tables are used. The investor is assumed to be a 65 year old male at $t = 0$. In panel (b), it is assumed that the initial portfolio value at $t = 0$ is $W(0) =$ \$1,000,000 and that the portfolio is invested entirely in the risk-free asset. Mortality Effects: assumes equation (3.3) used. No Mortality: fixed payments which exhaust all wealth after 30 years.*

We use the CPM 2014 mortality tables (male) from the Canadian Institute of Actuaries[1] to compute $T_x^*(t)$ with $x = 65$.

Figure 3.1(a) shows the ARVA multiplier $A(t_i)$ (assuming lump sum annual payments) for a Canadian male who begins withdrawing at age 65 at $t = 0$. As time passes $T_x^*(t) - t$ shrinks, and so the ARVA multiplier becomes larger. In other words, given a fixed amount of wealth, the annuity that can be received based on 80% of ones cohort passing away becomes larger. Assuming an initial wealth of $W(0) =$ \$1,000,000, Figure 3.1(b) shows the withdrawal amounts per year based on the ARVA spending rule. The portfolio is entirely invested in the risk-free asset with an assumed real return of $r = .004835$, which was the average real return of one month US T-bills over the period 1926:1-2016:12 (see Section 6). Although $A(t_i)$ increases with t_i, $W(t_i)$ decreases due to withdrawals. The net effect is that withdrawal amounts $A(t_i)W(t_i)$ decrease with time. If mortality effects are ignored, then the real fixed lump sum yearly payment that precisely exhausts the initial wealth after 30 years would be about \$34,650 per year.[2]

We can see from Figure 3.1(b) that the ARVA rule with a mortality boost shifts spending to earlier years, but this comes at the cost of reduced spending compared to a fixed term annuity after about year 20, i.e. age 85 (assuming, of course, that the investor has not passed away). As points of reference, the CPM 2014 tables indicate that the probabilities that a 65 year old Canadian male attains the ages of 85, 95, or 100 are .58, .13, and .02 respectively.

However, this pattern of reducing spending in the latter stages of retirement is perhaps not unreasonable. Studies show a decline in spending by about 1% per year after age 70, followed by a 2% decline per year after age 80 (Vettese, 2018). More precisely, in the US total consumption stays fairly flat during retirement but the allocation to healthcare increases significantly to be the second

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

[2] In contrast, assuming an initial capital of \$1,000,000 a fairly priced real life annuity would generate about \$49,960 per year. However, in the Canadian context real annuities are essentially unavailable. As of February 2019, online posted rates for a life annuity (no guarantee) for a 65-year old Canadian male were in the range of \$58,000 to \$65,160 per year (nominal). A 2% annual inflation rate would reduce the real value of a payment of \$60,000 to about \$33,000 after 30 years.

largest expense for those aged 91 and older.[3] In Canada total consumption declines slightly on a per-adult basis by about 5% from early sixties to early seventies. The percentage of health-related spending does increase but from a low base of 3% to 6%.[4] In summary, the data suggests that overall consumption declines from age 70, but the allocation to health and housing costs increases. The allocation to health care is a larger proportion of income in the US compared to Canada. In addition, in the Canadian context, deferring government benefits in terms of the Canada Pension Plan from age 65 to age 70 results in a 48% increase in annuity income. This deferred government annuity strategy can be used to offset the declining ARVA payments.

Our objective is to improve these spending patterns with a high probability by investing the portfolio in a combination of risky and risk-free assets. The comparison with the fixed term annuity is not quite fair, since the ARVA with mortality effects will not exhaust the portfolio at $t = 30$. In fact, at $t = 30$, $A(t = 30) = .24$. This means that after the payment at age 95, which is 0.24 $W(30)$, then 0.76 $W(30)$ remains in the investment portfolio. If this portfolio is then invested in risk-free assets, three additional equal payments of 0.253 $W(30)$ can be made at ages 96, 97, and 98. This will fund the retiree through to his 99^{th} birthday.

Remark 3.1. *[ARVA and ruin] Note that the ARVA strategy has no possibility of ruin. However, the withdrawal amounts may be become very small. The use of $T_x^*(t)$ to specify the remaining time of required cash flows allows us to front-end load spending in the early years of retirement, while allowing a reasonable buffer against longevity (i.e. the retiree has only a 20% chance of being alive at the end of the current annuity horizon). Note that this is a conditional probability, so that $T_x^*(t)$ increases with t. In the original paper (Waring and Siegel, 2015), the authors suggest various possibilities for $T_x^*(t)$, including the current life expectancy and the maximum possible lifespan. Waring and Siegel (2015) observe that use of current life expectancy results in a large front end spending, with a very rapid drop in spending in later years (provided the retiree is still alive). Westmacott (2017) suggests using the 20% rule as a way to shift spending to early years, while not causing too rapid a drop in spending in later years. Asset allocation strategies with fixed withdrawals and focusing on risk of ruin are analyzed in Forsyth et al. (2019).*

4 Optimization Problem

Let $\mathcal{T}$ be the set of withdrawal/rebalancing times. At each $t_i \in \mathcal{T}$, the investor (i) withdraws an amount of cash Q_i from the portfolio and then (ii) rebalances the portfolio. If the (unconstrained) ARVA rule is followed, then $Q_i = A(t_i)W(t_i)$, where $A(t_i)$ is defined in equation (3.3). As noted in the introduction, we consider a multi-objective problem which involves attempting to maximize reward while minimizing risk. This is in the spirit of mean-variance optimization, albeit using different measures of reward and risk. This general type of approach has been previously used by authors such as Freedman (2008) and Menoncin and Vigna (2017), among others. The obvious alternative would be to maximize some form of utility function. There are many possibilities here, including constant relative risk aversion (e.g. Milevsky and Young, 2007), constant absolute risk aversion (e.g. Liang and Young, 2018), recursive preferences (e.g. Blake et al., 2014), or habit formation (e.g. de Jong and Zhou, 2014). We do not pursue utility functions here, for several reasons. First, in some cases they would require additional parameters (e.g. separate parameters

[3] `https://www.vanguardcanada.ca/advisors/en/article/markets-economy/a-look-at-graying-populations`
[4] `https://www150.statcan.gc.ca/n1/en/pub/11f0027m/11f0027m2011067-eng.pdf?st=wqDBAi5o`

for risk-aversion and intertemporal substitution, in the case of recursive preferences). Second, a subjective discount rate often needs to be specified. This would have a similar effect as the mortality boost described above, i.e. it would tend to increase spending in earlier periods. Morever, such a parameter would be difficult to estimate. Third, some utility specifications are incompatible with the ARVA framework. For example, in the case of habit formation, "the habit serves as a floor in the required benefit level" (de Jong and Zhou, 2014, p. 37), and any benefit less than this level effectively results in an infinite loss of utility. On the other hand, as pointed out in Remark 3.1 above, the ARVA framework effectively removes the probability of ruin by taking on the risk of extremely small withdrawals. Taken together, these various considerations lead us to avoid the use of utility functions in this work.

As a measure of reward, we consider

$$E\left[\sum_{i=0}^{i=M} Q_i\right]. \tag{4.1}$$

Equation (4.1) is the expected sum of all real withdrawals. This captures the simple intuition that the investor seeks to maximize total real withdrawals. We do not explicitly consider the present value of the withdrawals because it is not clear what discount rate should be used given the risk of the withdrawals. The withdrawals are determined in part by the performance of an investment strategy having portfolio weights that change randomly over time, in response to realized past returns. As such, it is not possible to specify an appropriate discount rate.

Our measure of risk is

$$E\left[\sum_{i=1}^{i=M} \left((Q_i - Q_{i-1})^-\right)^2\right], \tag{4.2}$$

where $(Q_i - Q_{i-1})^- \equiv \min(Q_i - Q_{i-1}, 0)$. This is a measure of the downside variability in withdrawals, and reflects the idea that the retiree generally wants to avoid year-to-year declines in withdrawals. However, keep in mind that applying the mortality boost as described in Section 3 will tend to reduce withdrawals over time. These reductions will be reflected in the risk measure given in equation 4.2, even though they are in a sense a deliberate choice (from the mortality boost), not a consequence of poor investment allocation or performance. Overall, the investor wants to maximize reward (4.1) while minimizing risk (4.2). These are clearly conflicting goals, and we search for Pareto optimal strategies using a scalarization approach.

Given a time dependent function $h(t)$, we use the shorthand notation

$$h_i^+ = h(t_i^+) \equiv \lim_{\epsilon \to 0^+} h(t_i + \epsilon) \quad ; \quad h_i^- = h(t_i^-) \equiv \lim_{\epsilon \to 0^+} h(t_i - \epsilon) \; t_i \in \mathcal{T} \; .$$

For $t_i \in \mathcal{T}$, let $S_i^- = S_{t_i^-}$, $S_i^+ = S_{t_i^+}$, $B_i^- = B_{t_i^-}$, $B_i^+ = B_{t_i^+}$ Similarly, define total wealth as $W_i^- = S_i^- + B_i^-$ and $W_i^+ = S_i^+ + B_i^+$.

Remark 4.1 (Relation to jump process (2.5))**.** *Note that the jump process (2.5) is considered to apply only at times $t \notin \mathcal{T}$, while the notation W_i^-, W_i^+ refers to the times the instant before and after the rebalancing times $t \in \mathcal{T}$. In other words, we suppose that the risky asset follows a jump diffusion process* between *rebalancing times (in the absence of control). Rebalancing is assumed to occur instantaneously, so that the probability of a jump ocurring in (t_i^-, t_i^+) is zero. Informally, we suppose that during the interval (t_i^-, t_i^+), the risky asset value is frozen. As a concrete example, this would be the case if the investor (i) liquidated the risky asset and invested in riskless bonds just before the rebalancing time, and then (ii) purchased the desired amount of the risky asset just after the rebalancing time.*

Define the state variable $Q(t)$ for $t \in (t_{i-1}, t_i)$ as $Q(t) = Q_{i-1}$ for $t \in (t_{i-1}, t_i)$. In other words, for any time between withdrawal dates, $Q(t)$ represents the withdrawal amount at the previous withdrawal time. Finally, denote by $X(t) = (S(t), B(t), Q(t))$, $t \in [0, T]$, the multi-dimensional controlled underlying process, and let $x = (s, b, q)$ be the realized state of the system.

The control for our problem is the fraction allocated to equities at t_i^+, $p_i = p_i(X_i^-, t_i^-)$, where $X_i^- = (S_i^-, B_i^-, Q_i^-)$. Our optimization problem is then

$$\max_{\{p_0,\ldots,p_{M-1}\}} \left\{ E\left[\sum_{i=0}^{i=M} Q_i\right] - \lambda E\left[\sum_{i=1}^{i=M} \left((Q_i - Q_{i-1})^-\right)^2\right]\right\}$$
$$\text{subject to} \begin{cases} (S_t, B_t) \text{ follow processes (2.2)-(2.6)}; \quad t \notin \mathcal{T} \\ Q_i = \min(A_i W_i^-, Q_{\max}) \\ W_i^+ = W_i^- - Q_i\,;\; S_i^+ = p_i W_i^+\,;\; B_i^+ = (1-p_i) W_i^+\,; \quad t \in \mathcal{T} \\ p_i = p_i(X_i^-, t_i^-)\,;\; p_i \in Z\,;\; Z = [0,1] \end{cases}, \tag{4.3}$$

where $\lambda > 0$ is the scalarization parameter, $A_i \equiv A(t_i)$ is defined in equation (3.3) and Z is the admissible set. In problem (4.3), we impose the constraints that no-shorting and no-leverage are permitted (i.e. $p_i \in Z = [0,1]$). We also restrict the withdrawal amount to be at most $Q_{\max}$ in order to minimize the effects of large low probability withdrawals. We solve problem (4.3) using dynamic programming, working backwards from the investment horizon $t = T$ to $t = 0$. Note that, given the control p_i, then S_i^+, B_i^+ are entirely determined by quantities at t_i^-,

$$\begin{aligned} S_i^+ &= p_i(W_i^- - \min(A_i W_i^-, Q_{\max})) \\ B_i^+ &= (1-p_i)(W_i^- - \min(A_i W_i^-, Q_{\max}))\,. \end{aligned} \tag{4.4}$$

In the interval (t_i, t_{i+1}), we define the value function $V(s, b, q, t)$ as

$$V(s,b,q,t) = \max_{\hat{p}_{i+1}} E\left[\sum_{k=i+1}^{M} Q_k - \lambda \sum_{k=i+1}^{M} \left((Q_k - Q_{k-1})^-\right)^2 \middle| S(t) = s, B(t) = b, Q(t) = q\right] \tag{4.5}$$

where $\hat{p}_{i+1} = \{p_{i+1}, \ldots, p_{M-1}\}$. For $t \in (t_i, t_{i+1})$, there are no external cash flows or controls applied, as well as no discounting (all quantities are real). Thus the tower property gives for $h < (t_{i+1} - t_i)$

$$\begin{aligned} V(s,b,q,t) = E&\left[V(S(t+h), B(t+h), Q(t+h), t+h) \middle| S(t) = s, B(t) = b, Q(t) = q\right] \\ &t \in (t_i, t_{i+1} - h)\,. \end{aligned} \tag{4.6}$$

Assuming (S_t, B_t) follow the processes (2.2)-(2.6) and noting that $Q(t)$ is constant in (t_i, t_{i+1}), Itô's Lemma (for a jump diffusion) with $h \to 0$ gives the PIDE for $V(s, b, q, t)$ in the interval (t_i, t_{i+1}):

$$V_t + \frac{\sigma^2 s^2}{2} V_{ss} + (\mu - \zeta_\xi \kappa) V_s - \zeta_\xi V + rbV_b + \int_{-\infty}^{+\infty} V(e^y s, b, q, t) f(y)\, dy = 0. \tag{4.7}$$

Across the rebalancing/withdrawal time (t_i^-, t_i^+), the value function satisfies

$$
\begin{aligned}
V(s,b,q,t_i^-) &= \max_{p' \in Z} \left\{ V\left(p'\, w^+, (1-p')\, w^+, q_i^+, t_i^+\right) \right\} + q_i^+ - \lambda \left((q_i^+ - q)^- \right)^2 \\
w^- &= s + b \\
q_i^+ &= \min\left(A_i w^-, Q_{\max}\right) \\
w^+ &= w^- - q_i^+ \ .
\end{aligned} \tag{4.8}
$$

Equations (4.8) can be simplified for implementation purposes. Define

$$
\begin{aligned}
q_i^+\left(w^-\right) &= \min\left(A_i w^-, Q_{\max}\right) \\
w^+\left(w^-\right) &= w^- - q_i^+\left(w^-\right) \\
p_i^*\left(w^-\right) &= \operatorname*{arg\,max}_{p' \in Z} V\left(p'\, w^+\left(w^-\right), (1-p')\, w^+\left(w^-\right), q_i^+\left(w^-\right), t_i^+\right) \\
\hat{V}_i\left(w^-\right) &= V\left(p_i^*\left(w^-\right)\ w^+\left(w^-\right), \left(1 - p_i^*\left(w^-\right)\right)\ w^+\left(w^-\right), q_i^+\left(w^-\right), t_i^+\right)
\end{aligned} \tag{4.9}
$$

so that across the rebalancing time (t_i^-, t_i^+) we have

$$
V(s,b,q,t_i^-) = \hat{V}_i\left(w^-\right) + q_i^+\left(w^-\right) - \lambda \left(\left(q_i^+\left(w^-\right) - q \right)^- \right)^2 . \tag{4.10}
$$

Equation (4.10) shows that the optimal rebalancing fraction $p_i\left(s,b,q,t_i^-\right) = p^*\left(w^-, t_i^-\right)$ is a function of only $w^- = (s+b)$ and time. Note that this contrasts with typical glide path strategies in *to and through* target date funds, where the fraction invested in equities is a function of time only (Forsyth et al., 2019).

5 Numerical Method

We use dynamic programming to solve the optimization problem (4.3) on the computational domain $\Omega = (s,b,q,t) \in [s_{\min}, s_{\max}] \times [0, b_{\max}] \times [0, q_{\max}] \times [0, T]$. At $t = T$ we have

$$
V(s,b,q,T^+) = 0; \qquad (s,b,q,T) \in \Omega. \tag{5.1}
$$

We use equation (4.8) to advance the solution (backwards in time) from $t_i^+ \rightarrow t_i^-$. Then we use equation (4.7) to advance the solution (backwards in time) from $t_i^- \rightarrow t_{i-1}^+$.

We discretize the intervals $[0, b_{\max}]$ and $[0, q_{\max}]$ using an unequally spaced grid having $n_b \times n_q$ nodes. The (constrained) ARVA spending rule means that $0 \le q \le q_{\max}$. Setting $q_{\max} = Q_{\max}$, then no boundary conditions are required at $q = 0, q_{\max}$. We simply solve the PIDE (4.7) along the planes $q = 0$ and $q = q_{\max}$. Similarly, the ARVA spending rule and the no-leverage constraint imply that $b \ge 0$. In addition, we artificially set the interest payments to zero at $b = b_{\max}$. We then solve PIDE (4.7) along the $b = 0$ plane. We solve PIDE (4.7) (setting the term rbV_b to zero) along the $b = b_{\max}$ plane. We use the Fourier-based method described in Forsyth and Labahn (2019) with an equally spaced $x = \log s$ grid in the s direction with n_x nodes. To avoid wrap-around pollution, we use a buffer zone where we extend the solution by constant values for $s < s_{\min}, s > s_{\max}$, as described in Forsyth and Labahn (2019). More precisely, $V(s < s_{\min}, b, q, t) = V(s_{\min}, b, q, t)$ and $V(s > s_{\max}, b, q, t) = V(s_{\max}, b, q, t)$, which is imposed at the end of each timestep. The local maximization problem in equation (4.8) is solved using exhaustive search by discretizing the admissible range of p using an equally spaced grid with n_b nodes. Linear interpolation is used to evaluate $V(\cdot)$ at off-grid points. For further details, see Forsyth and Labahn (2019) and Dang and Forsyth (2014). Choosing $s_{\max}, b_{\max}$ sufficiently large will result in the effect of the artificial boundary conditions being small in regions of interest. We will verify this in our numerical experiments.

In order to determine

$$E\left[\sum_{i=0}^{i=M} Q_i\right]; \qquad E\left[\sum_{i=1}^{i=M} \left((Q_i - Q_{i-1})^-\right)^2\right] \tag{5.2}$$

separately, we solve an additional PIDE for $U(s,b,q,t)$ defined by

$$U(s,b,q,t) = E_{\{p^*_{i+1},\ldots,p^*_{M-1}\}}\left[\sum_{k=i+1}^{M} Q_k \middle| S(t)=s, B(t)=b, Q(t)=q\right]; \quad t \in \left(t_i^+, t_{i+1}^-\right). \tag{5.3}$$

where $p^*_i(s{+}b, t_i^-)$ are the optimal controls determined from equation (4.10). Across the rebalancing times (t_i^-, t_i^+) we have

$$\begin{aligned} U(s,b,q,t_i^-) &= U\left(p^*_i w^+, (1-p^*_i)\, W^+, q_i^+\left(w^-\right), t_i^+\right) + q_i^+\left(w^-\right) \\ w^- = (s+b); &\quad w^+ = w^- - q_i^+\left(w^-\right). \end{aligned} \tag{5.4}$$

At $t = T$, we have the initial condition

$$U(s,b,q,T^+) = 0; \quad (s,b,q,T) \in \Omega. \tag{5.5}$$

From $t_i^- \to t_{i-1}^+$, we have

$$U_t + \frac{\sigma^2 s^2}{2} U_{ss} + (\mu - \zeta_\xi \kappa)\, U_s - \zeta_\xi U + rbU_b + \int_{-\infty}^{+\infty} U(e^y s, b, q, t) f(y)\, dy = 0. \tag{5.6}$$

Given an initial portfolio value W_0 along with $V(0, W_0, 0, 0^-)$ and $U(0, W_0, 0, 0^-)$, it is straightforward to determine the quantities of interest in equation (5.2).

6 Data and Parameters

The data we use was obtained from Dimensional Returns 2.0 under licence from Dimensional Fund Advisors Canada. In particular, we use the Center for Research in Security Prices Deciles (1-10) index. This is a total return value-weighted index of US stocks. We also use one month Treasury bill returns for the risk-free asset. Both the equity returns and the Treasury bill returns are in nominal terms, so we adjust them for inflation by using the US CPI index. All of the data used was at the monthly frequency, with a sample period of 1926:1 to 2016:12.

To avoid known problems with other approaches, we use the method described in Dang and Forsyth (2016) and Forsyth and Vetzal (2017) based on the thresholding technique of Mancini (2009) and Cont and Mancini (2011). A tuning parameter α is required which, in intuitive terms, identifies a jump if the absolute value of the detrended log return is more than $\alpha\sigma\sqrt{\Delta t}$, where σ is the annualized diffusive volatility and Δt is the time interval (measured in years) between observations of the data series. Table 6.1 shows the estimated parameters of process (2.2) for the real stock return index, with $\alpha = 3$. The (uncompensated) drift parameter μ is a bit below 9%, the diffusive volatility σ is around 15%, and jumps are expected occur about once every $1/\zeta_\xi \approx 3$

Real CRSP Value-Weighted Index					
μ	σ	ζ_ξ	p_u	η_1	η_2
0.08753	0.14801	0.34065	0.25806	4.67877	5.60389
Real 1-Month Treasury Bill Index					
Mean return		Volatility		Correlation	
0.004835		0.018920		0.06662	

TABLE 6.1: *Annualized parameter estimates for jump diffusion model (see equation* (2.2)*) of the real CRSP value-weighted equity index and mean annualized real rate of return for 1-month US Treasury bills (*$\log[B(T)/B(0)]/T$*). Also reported are the annualized volatility of the real rate of return for Treasury bills and the correlation between real returns for the Treasury bill and value-weighted equity indexes. Sample period 1926:1 to 2016:12. Data obtained from Dimensional Returns 2.0 under licence from Dimensional Fund Advisors Canada.*

years. Downward jumps are almost three times as likely to occur as upward jumps, and the average magnitude of an upward jump ($1/\eta_1$) is a bit higher than the average magnitude of a downward jump ($1/\eta_2$). Table 6.1 also shows that the estimated value of r for the bond process (2.6) (i.e. the average annual return) is 0.4835%. For information purposes, we also provide the volatility of the real Treasury bill index return as well as its correlation with the equity market index. The volatility is quite low (less than 2%), and the two return series have slightly positive correlation over the sample period from 1926:1 to 2016:12.

7 Convergence Test

We begin by conducting a convergence test of our numerical method. We consider the scenario documented in Table 7.1. Monetary units in the table are in thousands of dollars, so that the initial portfolio value $W_0 = 1{,}000$ implies an initial wealth of \$1 million. The investor withdraws cash immediately and at the end of each of the next 30 years, and is not permitted to withdraw more than \$100,000 per year. The portfolio is rebalanced annually, at the cash withdrawal times. As indicated in Table 7.1, the market parameters used are from Table 6.1. The summary statistics provided here are based on the average expected withdrawal $\bar{\mathcal{Q}}$ and the average withdrawal variability $\bar{\mathcal{V}}_q$. These two quantities are defined as follows:

$$
\begin{aligned}
\text{Average expected withdrawal} = \bar{\mathcal{Q}} &= \frac{1}{M+1}\, E\left[\sum_{i=0}^{i=M} Q_i\right] \\
\text{Average withdrawal variability} = \bar{\mathcal{V}}_q &= \sqrt{E\left[\frac{1}{M}\sum_{i=1}^{i=M} \left((Q_i - Q_{i-1})^-\right)^2\right]} \\
(Q_i - Q_{i-1})^- &= \min\left(Q_i - Q_{i-1}, 0\right).
\end{aligned} \tag{7.1}
$$

We take the square root in equation (7.1) so that $\bar{\mathcal{V}}_q$ and $\bar{\mathcal{Q}}$ have the same units.

We solve the optimization problem (4.3) by solving the PIDEs (4.5)-(4.8) and equations (5.3)-(5.6). We discretize the problem in the (s, b, q) directions using $s_{\min} = .04, s_{\max} = b_{\max} = 10^6$, and $q_{\max} = Q_{\max}$. Increasing $s_{\max}, b_{\max}$ (by a factor of ten) and decreasing $s_{\min}$ (dividing by ten) resulted in no change to the solution to six figures. We use the Fourier method described in Forsyth

Investment horizon T (years)	30
Equity market index	CRSP value-weighted index (real)
Risk-free asset index	1-month Treasury bill index (real)
Initial portfolio value W_0	1000
Cash withdrawal times	$t = 0,1,\ldots,30$
$Q_{\max}$	100
Rebalancing interval (years)	1
Market parameters	See Table 6.1

TABLE 7.1: *Input data for examples. Monetary units: thousands of dollars.*

	PIDE			Monte Carlo	
Grid (n_x, n_b, n_q)	Value Function	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$
$\lambda = 2.0$					
$256 \times 153 \times 157$	1257.23	61.4008	3.2284	61.3628	3.1231
$512 \times 305 \times 313$	1290.10	61.2590	3.1338	61.2232	3.1050
$1024 \times 609 \times 625$	1297.95	61.2060	3.1094	61.1830	3.1006
$\lambda = 1.0$					
$256 \times 153 \times 157$	1647.22	67.8729	3.8388	67.8621	3.7580
$512 \times 305 \times 313$	1661.53	67.8883	3.7803	67.8786	3.7559
$1024 \times 609 \times 625$	1665.24	67.8931	3.7651	67.8808	3.7552

TABLE 7.2: *Convergence test of the solution of equations* (4.5)-(4.8) *and equations* (5.3)-(5.6) *used to compute the optimal asset allocation. Monte Carlo results based on 640,000 simulated paths in the synthetic market, with controls computed from the PIDE at the indicated grid size. Input data provided in Table 7.1. Monetary units: thousands of dollars.*

and Labahn (2019), which requires that $s_{\min} > 0$. There is no timestepping error for the Fourier method between rebalancing times. Table 7.2 provides a convergence study in which we compute various quantities of interest for a sequence of grid sizes. The value function is $V(0, W_0, 0, 0^-)$ where $V(\cdot)$ is defined in equation (4.5). We compute and store the optimal asset allocations from the PIDE solver, and then carry out Monte Carlo simulations to verify the solution. If we are in the asymptotic convergence region, then convergence of the PIDE Fourier method should be monotonic, and we can estimate a rate of convergence. This appears to be true from Table 7.2, and the rate of convergence appears to be between first and second order for the PIDE solution. Making the pessimistic assumption that convergence is first order, then for $\lambda = 2.0$ we can assume that $\bar{\mathcal{Q}} = 61.206 \pm .05$, $\bar{\mathcal{V}}_q = 3.1094 \pm .025$, while for $\lambda = 1.0$, we can estimate that $\bar{\mathcal{Q}} = 67.8931 \pm .005$, and $\bar{\mathcal{V}}_q = 3.7651 \pm .01$. The computational cost of our method is dominated by the cost of the PIDE solve. Suppose we have a grid with $n_x \times n_b \times n_q$ nodes. If we double the number of nodes in each direction, then this will cost $8(1 + 1/\log_2(n_x))$ more computational time (the $\log_2$ term comes from the FFT algorithm). Results reported in the remainder of the paper use the control from the finest PIDE grid.

p	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$	p	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$	p	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$
0.00	33.0	1.11	0.45	55.7	3.45	0.75	66.1	4.56
0.20	41.9	1.88	0.50	58.0	3.68	0.80	67.0	4.73
0.25	44.6	2.21	0.55	60.1	3.88	0.85	67.9	4.90
0.30	47.4	2.56	0.60	62.0	4.06	0.90	68.5	5.07
0.35	50.3	2.90	0.65	63.6	4.24	0.95	69.1	5.24
0.40	53.1	3.88	0.70	64.9	4.40	1.00	69.6	5.42

TABLE 8.1: *Results for ARVA spending rule when the portfolio is rebalanced at each rebalancing date to a fixed risky asset weight p. Monetary units: thousands of dollars. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

8 Synthetic Market Examples

We now explore some examples based on the input data given in Table 7.1. This section presents results in the synthetic market. Recall that this means that we compute the control using the parameters from Table 6.1 and then assess performance by Monte Carlo simulation assuming exactly the same parameters. More specifically, we use the following steps:

1. We solve problem (4.3) to determine the optimal asset allocation strategy. This assumes that the value of the risky equity market and risk-free bond indexes evolve according to equations (2.2) and (2.6) respectively, with the parameters provided in Table 6.1. We store the generated optimal controls.

2. We generate Monte Carlo simulated paths of the two indexes over the investment horizon, calculating values at each rebalancing date according to processes (2.2) and (2.6) with the parameters in Table 6.1.

3. We then apply the stored controls to each path, calculating statistics such as the average withdrawal and withdrawal variability. We then compute averages and percentiles of the relevant path statistics across the simulated paths.

For purposes of comparison, we also evaluate the performance of fixed weight strategies. In these cases Step 1 above is skipped, and in Step 3 we just rebalance to constant specified portfolio weights.

As a first case, we consider the ARVA spending rule with a fixed (constant) equity allocation at each rebalancing time. Table 8.1 shows the results for this constant weight asset allocation strategy. As the fixed equity weight is increased, both the average expected withdrawal $\bar{\mathcal{Q}}$ and the average withdrawal variability $\bar{\mathcal{V}}_q$ rise monotonically. Note that even for $p = 0$ (all money invested in bonds), $\bar{\mathcal{V}}_q > 0$. This is because the ARVA spending rule results in declining payments over time, due to the front end loading of the mortality boost as indicated in Figure 3.1(b).

We next solve problem (4.3) to determine the optimal strategy according to our criteria. Optimal asset allocation results are shown for various values of the scalarization parameter λ in Table 8.2. Higher values of λ correspond to higher risk-aversion since more weight is placed on the risk term in the objective function. As the table shows, reducing λ leads to monotonically increasing reward $\bar{\mathcal{Q}}$ and risk $\bar{\mathcal{V}}_q$.

Table 8.2 shows that the average expected withdrawal $\bar{\mathcal{Q}}$ is 67.9 when $\lambda = 1.0$. From Table 8.1, the same average expected withdrawal is obtained for a fixed weight strategy with $p = 0.85$. However, the fixed weight strategy has higher average withdrawal variability $\bar{\mathcal{V}}_q$ of 4.73 compared

λ	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$
5.0	36.4	1.34
4.0	39.1	1.54
3.0	53.1	2.55
2.0	61.2	3.10
1.0	67.9	3.76
0.5	70.2	4.29

TABLE 8.2: *Results for ARVA spending rule when the portfolio is rebalanced according to the optimal asset allocation, defined as the solution to problem (4.3). Monetary units: thousands of dollars. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

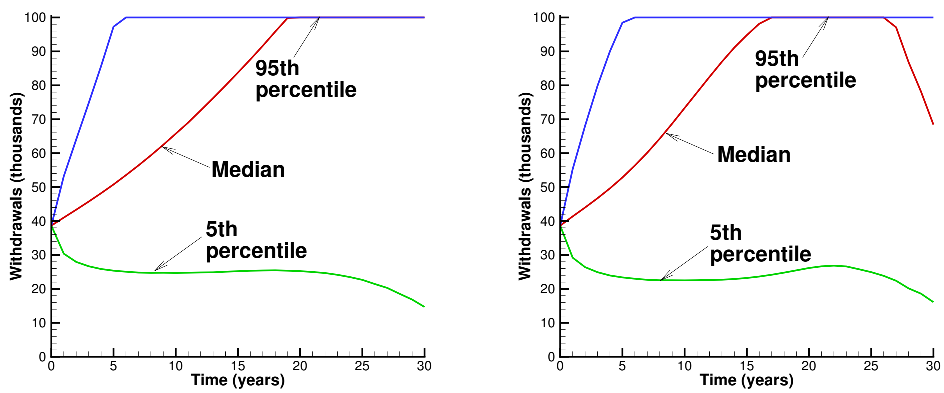

(a) *Fixed weight strategy* $p = 0.85$. (b) *Optimal strategy* $\lambda = 1.0$.

FIGURE 8.1: *Percentiles of withdrawal amounts over time for fixed weight strategy with* $p = 0.85$ *and optimal strategy with* $\lambda = 1.0$*. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

to 3.76 for the optimal asset allocation. An overall indication of the general pattern of withdrawals over time is provided in Figure 8.1, which shows the 5^{th}, 50^{th}, and 95^{th} percentiles of the distribution of withdrawals for the fixed weight strategy in panel (a) and the optimal strategy in panel (b). At a broad level, the two cases appear to be quite similar. The initial withdrawal is around \$40,000. The 95^{th} percentile of withdrawals rises rapidly over about the first 5 years to the maximum specified amount of \$100,000 and remains there throughout the horizon for each case. Conversely, the 5^{th} percentile of withdrawals quickly drops below \$30,000 and remains there over most of the horizon, before tailing off a bit further in the final few years. The median withdrawal rises more slowly than the 95^{th} percentile, but does attain the allowed maximum in each case. This happens slightly faster under the optimal strategy compared to the fixed weight strategy. The median withdrawal remains constant at the maximum amount throughout the horizon for the fixed weight strategy, but it drops off during the last 3 years of the horizon for the optimal strategy.

Figure 8.2 depicts the 5^{th}, 50^{th}, and 95^{th} percentiles of the fraction of the portfolio allocated to equities over time for the optimal strategy with $\lambda = 1$. Keep in mind that this strategy produces the same average expected withdrawal as a fixed weight strategy that annually rebalances to having

85% invested in the risky equity market index. The optimal strategy starts out with all funds in the risky asset. In the 5^{th} percentile case, the portion of the portfolio in the risky asset drops very quickly, down to about 30% after 5 years and reaching zero after about 15 years. The median fraction has $p = 1$ for the first 7 years. This declines to around 10-15% for years 20-25, and thereafter increases back to about 30% at the end of the horizon. In the 95^{th} percentile case, the portfolio is entirely invested in the risky asset for almost 20 years, and then falls off to being about 30% at risk at the end of the horizon. Overall, an investor who follows the optimal strategy will need to initially put all of his funds in the risky asset, but he will likely to be able to reduce his equity market risk exposure substantially over time. Reaching the same average expected withdrawal with a fixed weight strategy requires keeping a consistently high equity weighting throughout the horizon. Of course, this leads to higher withdrawal variability, as measured by $\bar{\mathcal{V}}_q$.

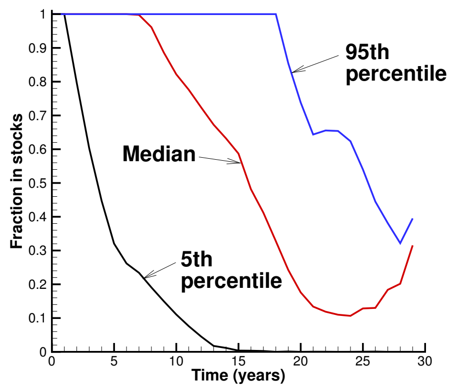

Figure 8.2: *Percentiles of control p over time for optimal strategy with $\lambda = 1.0$. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

Similar results are shown for $\lambda = 2.0$ in Figures 8.3 and 8.4. From Table 8.2, the average expected withdrawal in this case is $\bar{\mathcal{Q}} = 61.2$. The constant weight strategy which gives the same value of $\bar{\mathcal{Q}}$ is $p^* = .58$, found by interpolating the results reported in Table 8.1. This constant weight strategy has $\bar{\mathcal{V}}_q = 3.99$, compared to the optimal strategy which has $\bar{\mathcal{V}}_q = 3.10$. Unlike in Figure 8.1 above, here (Figure 8.3) the median withdrawal amount rises over the first several years and subsequently falls, but it never comes close to the maximum allowed withdrawal. Comparing Figures 8.2 and 8.4, we see that in the more risk-averse case ($\lambda = 2$), the fraction optimally put at risk declines from the initial value of $p = 1$ much earlier.

Table 8.3 shows some statistics about the distributions of final wealth for the optimal strategies with $\lambda = \{1.0, 2.0\}$ and the constant weight strategies which generate the same average expected withdrawals $\bar{\mathcal{Q}}$. The final wealth is at $t = 30$ years. Recall that there is enough cash remaining to fund 3 years of payments (after the payment at $t = 30$). As a result, this takes the retiree through to his 99^{th} birthday. The final wealth values at the 5^{th} percentiles are comparable with the fixed weight strategies with the same $\bar{\mathcal{Q}}$. However, the fixed weight strategies have much higher median and 95^{th} percentile terminal portfolio values, which is to be expected due to the higher average allocation to equities. These large values of final wealth are due to low probability very favourable investment results, coupled with the cap on withdrawals of $100,000 per year.

As another point of comparison between the fixed weight strategies and the optimal strategy, we

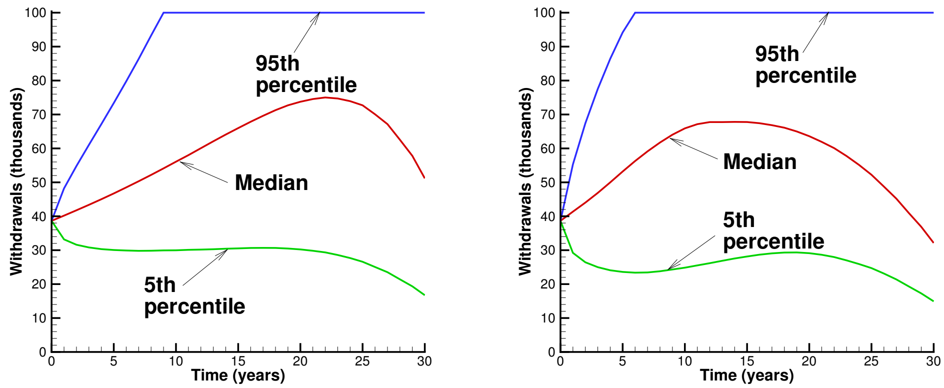

(a) *Fixed weight strategy $p = .58$.* (b) *Optimal strategy $\lambda = 2.0$.*

FIGURE 8.3: *Percentiles of withdrawal amounts over time for fixed weight strategy with $p = 0.58$ and optimal strategy with $\lambda = 2.0$. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

Strategy	Median W_T	5^{th} percentile W_T	95^{th} percentile W_T
		$\bar{\mathcal{Q}} = 61.2$	
Optimal $\lambda = 1.0$	179	47.9	610
Fixed $p = .85$	352	46.8	10800
		$\bar{\mathcal{Q}} = 67.9$	
Optimal $\lambda = 2.0$	100	46.6	443
Fixed $p = .58$	161	52.5	2612

TABLE 8.3: *Statistics of final wealth W_T after withdrawal at $T = 30$ years. Monetary units: thousands of dollars. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

consider the time averaged median fraction in the risky asset. If $Median[p_i]$ is the median fraction invested in the risky asset at time t_i, then we define the time averaged median fraction as

$$\text{time averaged median fraction in stocks} \quad = \quad \frac{1}{M+1} \sum_{i=0}^{i=M} Median[p_i] \ . \tag{8.1}$$

Table 8.4 shows the time averaged median fraction invested in the risky asset for the cases of $\lambda = \{1.0, 2.0\}$ compared with the fixed weight strategies which give the same $\bar{\mathcal{Q}}$. When $\lambda = 1.0$, the time averaged median value of p is .57, versus the fixed weight of .85. Similarly, for the case with $\lambda = 2.0$, the time averaged median of p is .40, compared to the fixed weight of $p = .58$.

From Figure 8.4 we can see that the optimal strategy ($\lambda = 2$) has a median fraction in stocks of 1.0 during the early years of retirement, which then drops rapidly. This is contrary to the usual advice given to retirees. However, from Table 8.4 we can see that time averaged fraction in stocks for this strategy is 0.40. In order to generate the same average expected withdrawal, a fixed weight strategy requires $p = .58$, with considerably greater withdrawal variability. In other words, although the optimal strategy has a maximum equity fraction larger than the fixed weight strategy

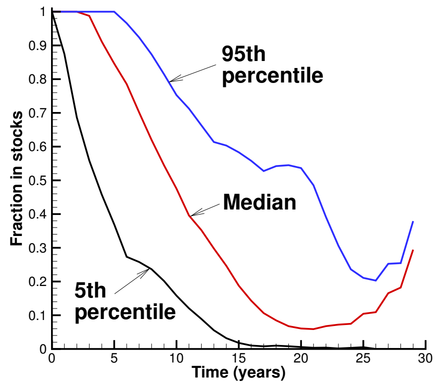

FIGURE 8.4: *Percentiles of control p over time for optimal strategy with $\lambda = 2.0$. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

Optimal				Fixed weight		
λ	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$	Time Averaged Median of p	p	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$
1.0	67.9	3.76	.57	.85	67.9	4.90
2.0	61.2	3.10	.40	.58	61.2	3.99

TABLE 8.4: *Results for optimal strategy as in problem* (4.3) *and fixed weight strategies having the same expected average withdrawal $\bar{\mathcal{Q}}$. Monetary units: thousands of dollars. Results computed using Monte Carlo simulations in the synthetic market with 640,000 paths. Input data provided in Table 7.1.*

with the same expected withdrawal, it is the time averaged equity fraction which contributes to overall risk. We discuss this in greater detail in the next section.

8.1 Analysis of the Objective Function

We now provide a heuristic analysis as to why the optimal strategy for objective function (4.3) tends to reduce the average amount in the risky asset. Consider the risk term in problem (4.3):

$$E\left[\sum_{i=1}^{i=M}\left((Q_i - Q_{i-1})^-\right)^2\right] = \sum_{i=1}^{i=M} E\left[\left((Q_i - Q_{i-1})^-\right)^2\right]. \tag{8.2}$$

For ease of exposition, assume that $t_i - t_{i-1} = \Delta t$ is small. If $t_i \ll T, T_x^*(t_i)$, then we can also assume that $A(t_i) = O(\Delta t); A(t_i) = A(t_{i-1}) + O(\Delta t)^2$. Examining one term of the sum in equation (8.2) gives

$$\begin{aligned}\left((Q_i - Q_{i-1})^-\right)^2 &\le \left((Q_i - Q_{i-1})\right)^2\\ &= A_i^2\left(W_i^- - W_{i-1}^- + O(\Delta t)^2\right)^2\end{aligned}$$

$$\simeq A_i^2 \left(W_i^- - \frac{W_{i-1}^+}{1 - A_i} + O(\Delta t)^2 \right)^2$$
$$= A_i^2 \left(\Delta W_{i-1}^+ + O(\Delta t) \right)^2, \tag{8.3}$$

where $\Delta W_{i-1}^+ = \left(W_i^- - W_{i-1}^+ \right)$. If we continuously rebalance to a fixed weight p_{i-1} for $t \in (t_{i-1}, t_i)$, then from equations (2.5) and (2.6), and Itô's Lemma for jump processes, we obtain

$$\frac{dW}{W} = [p_{i-1}(\mu - r) + r]\ dt - \zeta_\xi \kappa\, dt + p_{i-1}\sigma\, dZ + p_{i-1}(\xi - 1)\, d\mathbb{Q}. \tag{8.4}$$

If we assume that $\Delta t = t_i - t_{i-1} \simeq dt$ and that $dW \simeq (W_i^- - W_{i-1}^+) \equiv \Delta W_{i-1}^+$, then equation (8.4) becomes

$$\frac{\Delta W_{i-1}^+}{W_{i-1}^+} = [p_{i-1}(\mu - r) + r]\ dt - \zeta_\xi \kappa\, dt + p_{i-1}\sigma\, dZ + p_{i-1}(\xi - 1)\, d\mathbb{Q}. \tag{8.5}$$

Substituting equation (8.5) into equation (8.3) and taking expectations gives

$$E\left[((Q_i - Q_{i-1})^-)^2 \right] \leq E\left[(Q_i - Q_{i-1})^2 \right]$$
$$= \left(W_{i-1}^+ \right)^2 A_i^2 \left[p_{i-1}^2 \sigma^2\, dt + p_{i-1}^2 E\left[(\xi - 1)^2 \right] \zeta_\xi\, dt \right] + o(dt)$$
$$\simeq \left(W_{i-1}^+ \right)^2 A_i^2 \left[p_{i-1}^2 \sigma^2\, \Delta t + p_{i-1}^2 E\left[(\xi - 1)^2 \right] \zeta_\xi\, \Delta t \right] + o(\Delta t). \tag{8.6}$$

Substituting equation (8.6) into equation (8.2) and ignoring terms of $o(\Delta t)$ gives

$$E\left[\sum_{i=1}^{i=M} ((Q_i - Q_{i-1})^-)^2 \right] \leq \sum_{i=1}^{i=M} p_{i-1}^2 \left(W_{i-1}^+ \right)^2 A_i^2 \left[\sigma^2 + E\left[(\xi - 1)^2 \right] \zeta_\xi \right] \Delta t. \tag{8.7}$$

From equation (8.7) we can see that reducing the weighted average fraction invested in the risky asset (p_{i-1}) will also reduce the upper bound on the risk term (8.2), which is consistent with the numerical results. As well, we can see that when W_{i-1}^+ becomes small (for large times) the weight multiplying the risky asset fraction in equation (8.7) becomes small, hence maximizing problem (4.3) would focus on maximizing the expected total withdrawals, which would tend to increase the fraction invested in the risky asset at later times. This effect can be seen in Figures 8.2 and 8.4. The non-smooth percentile curves in these plots for larger times arise because with little time remaining the control has a small influence on maximizing the expected total withdrawals.

Finally, note that in the limit as $t_i - t_{i-1} \to 0$, the risk term on the right hand side of (8.7) becomes a weighted portfolio quadratic variation. This has previously been suggested as a standalone risk measure in sources such as Brugiere (1996), Forsyth et al. (2012), and van Staden et al. (2019).

9 Bootstrap Tests

The results reported above have all been in the synthetic market, following the 3 step procedure outlined at the start of Section 8. We now replace the second step involving Monte Carlo simulation by bootstrap resampling of the historical data to generate simulated paths of the values of the risk and risk-free assets, in the absence of control. In this *historical market*, the other two steps remain as before. Although we still compute the optimal asset allocation strategy by solving problem (4.3), assuming as before that S_t and B_t follow processes (2.2) and (2.6) respectively, the performance tests themselves make no assumptions regarding the stochastic processes followed by the value of the equity and bond market indexes.

To construct a single bootstrap resampled path for asset returns, we use the stationary block bootstrap to account for possible serial dependence (see, e.g. Politis and White, 2004; Patton et al., 2009). We start at a random month in the 1926:1 to 2016:12 sample period. We draw a block of data starting in that month (we simultaneously sample both the bond and the stock indexes). The length of the block is determined by drawing a random value from a geometric distribution having mean (i.e. expected blocksize) $\hat{b}$. We continue to draw blocks of data in this way and paste them together until we have a path that covers the entire horizon of $T = 30$ years. This procedure is repeated many times to generate a large number of resampled paths. Note that we draw the blocks of data with replacement, so it is possible for us to use a historical period more than once in a single path. We wrap the data around so that if the size of a particular block extends past the end of the sample period in 2016:12, values for the remaining duration of that block are taken from the start of the sample period, beginning in 1926:1. See Forsyth and Vetzal (2019) for a detailed description of the bootstrap algorithm.

In principle, it is possible to estimate the optimal expected blocksize $\hat{b}$. However, if we apply the algorithm described in Patton et al. (2009) to our data, we find very different estimates for the two indexes: the value for the equity market index is about 3.5 months, while the value for the bond market index is around 57 months. This poses a problem since we sample simultaneously from both indexes. Consequently, we give results for several expected blocksizes.

One final point should be noted about our procedures. In our bootstrap tests, the bond and stock returns are computed using the actual historical returns. The ARVA annuity factor (3.3) is determined using the long term average real T-bill rate (recall that this is $r = .0048$). Since this rate is very low, this is a conservative approach, which essentially means that fluctuations in withdrawals are primarily driven by the actual observed asset returns, instead of projections about future real interest rates. Alternatively, it would be possible to use the most recently observed historical short rate in the ARVA annuity factor computation. However, this can cause volatility in the withdrawals solely due to the bootstrapping procedure, even when the portfolio returns are not volatile.

Table 9.1 shows the results. We also provide comparable results for the fixed weight strategy which gives the same value of $\bar{\mathcal{Q}}$ in the synthetic market. For any given expected blocksize, the optimal strategy has a much smaller average allocation to the risky asset, while having a very similar average total withdrawal. We also observe that the results in Table 9.1 are relatively insensitive to expected blocksize, which suggests that the strategies are quite robust. Comparing Table 9.1 with the earlier Table 8.4 from the synthetic market, we observe that the results for the expected average withdrawal $\bar{\mathcal{Q}}$ are quite similar. For example, with $\lambda = 1$ we had $\bar{\mathcal{Q}} = 67.9$ and $\bar{\mathcal{V}}_q = 3.76$ in the synthetic market (Table 8.4) for the optimal strategy. The corresponding bootstrap resampled values for $\bar{\mathcal{Q}}$ in Table 9.1 range from 67.6 to 69.3 as the expected blocksize increases from 6 months to 5 years, and the corresponding resampled values for $\bar{\mathcal{V}}_q$ range from 3.81 to 3.97. Overall, for the optimal strategy the average expected withdrawal $\bar{\mathcal{Q}}$ in the historical market is quite close to that for the idealized synthetic market. The average withdrawal variability $\bar{\mathcal{V}}_q$ is slightly higher in the historical market, but this is not surprising since the resampled paths will have stochastic interest rates and randomly changing volatility, neither of which are features of our synthetic market. The historical market results for the fixed weight strategies in Table 9.1 are also quite close to their synthetic market counterparts in Table 8.4. For example, with $p = .85$ the average expected withdrawal ranges from 68.0 to 69.4 in the historical market, compared to 67.9 in the synthetic market. Using this fixed weight gives average withdrawal variability that ranges

	Optimal			Fixed Weight		
λ	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$	Average Median $[p]$	p	$\bar{\mathcal{Q}}$	$\bar{\mathcal{V}}_q$
Expected blocksize = 0.5 years						
1.0	67.6	3.81	.57	.85	68.0	4.76
2.0	60.6	3.22	.40	.58	60.7	3.95
Expected blocksize = 1.0 years						
1.0	68.0	3.90	.57	.85	68.3	4.84
2.0	60.9	3.28	.40	.58	61.0	4.01
Expected blocksize = 2.0 years						
1.0	68.1	3.97	.57	.85	68.6	5.00
2.0	60.9	3.36	.40	.58	61.0	4.14
Expected blocksize = 5.0 years						
1.0	69.3	3.96	.56	.85	69.4	5.12
2.0	61.6	3.37	.40	.58	61.3	4.24

TABLE 9.1: *Results for optimal strategy as in problem* (4.3) *and fixed weight strategy with the same expected average withdrawal $\bar{\mathcal{Q}}$ in the synthetic market. Monetary units: thousands of dollars. Results computed in the historical market with 100,000 bootstrap resampled paths. Input data provided in Table 7.1. Controls computed in the synthetic market using parameters from Table 6.1 and stored, then applied to resampled historical data.*

from 4.76 to 5.12 in the historical market, versus 4.90 in the synthetic market.

Figure 9.1 shows the percentiles of withdrawals over time for the fixed weight strategy with $p = .85$ and the optimal strategy with $\lambda = 1$ in the historical market with an expected blocksize of two years. The two panels here are quite similar to the corresponding synthetic market plots from Figure 8.1, another signal that the synthetic market strategy is robust when tested on historical market data. Figure 9.2 shows the percentiles of the fraction of the investment portfolio allocated to equities over time, based on bootstrap resampling with an expected blocksize of two years. These results are also quite similar to the corresponding synthetic market results shown in Figure 8.2.

10 Conclusion

An ARVA spending rule results in variable withdrawals, which eliminates the possibility of ruin over the specified horizon. The risk of ruin is effectively replaced by the risk of withdrawal variability. The main positive feature of an ARVA rule is the fact that withdrawals reflect the investment experience. In addition, a mortality boost can be used to front end load the withdrawals. On the other hand, compared to an annuity, there is some possibility of very low withdrawals later on in life. Combining an ARVA rule with investing in a portfolio of risky assets and risk-free assets leads to a higher average expected withdrawal compared to a fairly priced annuity. Under an ARVA rule the investor retains full control over their portfolio, unlike for an annuity.

We compared two possible approaches to managing the investment portfolio under an ARVA

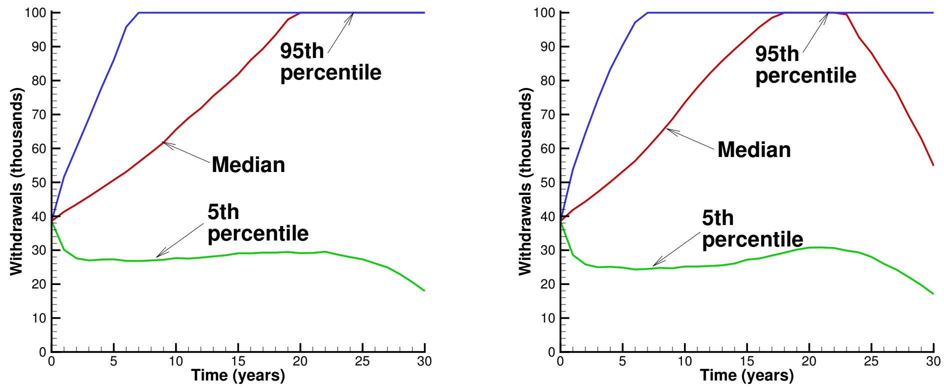

(a) *Fixed weight strategy* $p = .85$.

(b) *Optimal strategy* $\lambda = 1.0$.

FIGURE 9.1: *Percentiles of withdrawal amounts over time for fixed weight strategy with* $p = .85$ *and optimal strategy with* $\lambda = 1.0$. *Results computed in the historical market with 100,000 bootstrap resampled paths and expected blocksize of 2 years. Input data provided in Table 7.1. Control computed in the synthetic market using parameters from Table 6.1 and stored, then applied to resampled historical data.*

spending rule: a fixed weight strategy, and a strategy based on optimal control. The optimal control strategy minimized a downside measure of withdrawal variability, for a given expected average withdrawal. For the same expected average withdrawal, the optimal strategy has smaller withdrawal variability, smaller average investment over time in the risky asset, and similar final wealth at the 5^{th} percentile, compared to a fixed weight strategy. However, the fixed weight strategy has a higher median terminal wealth compared to the optimal strategy. This is to be expected due to the higher average weight in risky assets (for the same expected average withdrawal) compared to the optimal strategy and the cap imposed on withdrawals. These results hold for both a parametric model based on historical time series, as well as bootstrap resampled backtests.

The synthetic market results (parametric model) and the bootstrapped historical market results are very similar for either the optimal strategy or the fixed weight strategies. This suggests that an ARVA spending rule which adapts withdrawals to investment experience results in a very robust strategy, i.e. insensitive to market parameter misspecification.

Overall, a combination of an ARVA spending rule and an optimal control approach to reduce withdrawal variability, result in a decumulation strategy which has a high probability of achieving desirable outcomes. This does, however, come at the cost of high median equity fractions for short periods of time. Nevertheless, the time averaged (median) equity fraction is much smaller than the equivalent constant weight strategy, which we argue is the appropriate risk measure in this case.

A possible avenue for future research is to impose both maximum and minimum withdrawal amounts under an ARVA spending rule. This would ameliorate withdrawal variability, but now there is a risk of ruin. In this case, a possible objective function would maximize the expected total withdrawals, and minimize a risk measure such as probability of ruin or CVAR.

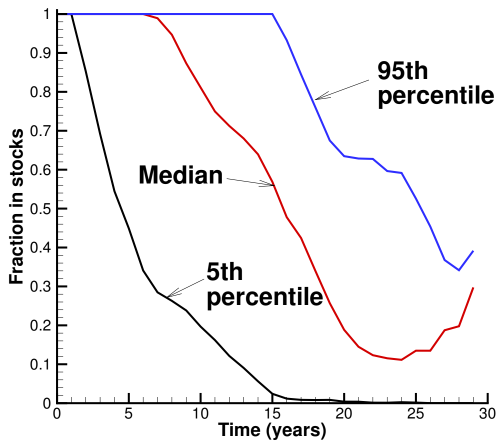

FIGURE 9.2: *Percentiles of control p over time for optimal strategy with $\lambda = 1.0$. Results computed in the historical market with 100,000 bootstrap resampled paths and expected blocksize of 2 years. Input data provided in Table 7.1. Control computed in the synthetic market using parameters from Table 6.1 and stored, then applied to resampled historical data.*

Acknowledgements

Peter Forsyth acknowledges support from the Natural Sciences and Engineering Research Council of Canada (NSERC), under grant RGPIN-2017-03760. Peter Forsyth and Kenneth Vetzal acknowledge support from the University of Waterloo.

Declarations of Interest

The authors have no conflicts of interest to declare.

References

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

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

Blake, D., A. J. G. Cairns, and K. Dowd (2003). Pensionmetrics 2: Stochastic pension plan design during the distribution phase. *Insurance: Mathematics and Economics 33*, 29–47.

Blake, D., D. Wright, and Y. Zhang (2014). Age-dependent investing: Optimal funding and investment strategies in defined contribution pension plans when members are rational life cycle financial planners. *Journal of Economic Dynamics and Control 38*, 105–124.

Brugiere, P. (1996). Optimal portfolio and optimal trading in a dynamic continuous time framework. AFIR Colloquium, Nurenberg, Germany.

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

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.

de Jong, F. and Y. Zhou (2014). Habit formation: Implications for pension plans. Design Paper 32, Netspar.

Forsyth, P., J. Kennedy, S. Tse, and H. Windcliff (2012). Optimal trade execution: a mean-quadratic variation approach. *Journal of Economic Dynamics and Control 36*, 1971–1991.

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

Forsyth, P. A. and K. R. Vetzal (2017). Robust asset allocation for long-term target-based investing. *International Journal of Theoretical and Applied Finance 20*(3). 1750017 (electronic).

Forsyth, P. A., K. R. Vetzal, and G. Westmacott (2019). Management of portfolio depletion risk through optimal life cycle asset allocation. *North American Actuarial Journal 23:3*, 447–468.

Freedman, B. (2008). Efficient post-retirement asset allocation. *North American Actuarial Journal 12*, 228–241.

Gerrard, R., S. Haberman, and E. Vigna (2004). Optimal investment choices post-retirement in a defined contribution pension scheme. *Insurance: Mathematics and Economics 35*, 321–342.

Gerrard, R., S. Haberman, and E. Vigna (2006). The management of decumulation risk in a defined contribution pension plan. *North American Actuarial Journal 10*, 84–110.

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

Liang, X. and V. R. Young (2018). Annuitization and asset allocation under exponential utility. *Insurance: Mathematics and Economics 79*, 167–183.

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.

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

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

Milevsky, M. A. and V. R. Young (2007). Annuitization and asset allocation. *Journal of Economic Dynamics and Control 31*, 3138–3177.

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.

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

Smith, G. and D. P. Gould (2007, Spring). Measuring and controlling shortfall risk in retirement. *Journal of Investing 16*, 82–95.

van Staden, P. M., D.-M. Dang, and P. A. Forsyth (2019). Mean-quadratic variation portfolio optimization: A desirable alternative to time-consistent mean-variance optimization? *SIAM Journal on Financial Mathematics 10:3*, 815–856.

Vettese, F. (2018). *Retirement Income for Life: Getting More Without Saving More*. Toronto: Milner.

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

Westmacott, G. (2017). The retiree's dilemma: the Deckards. PWL Capital White Paper, `http://www.pwlcapital.com/retirees-dilemmma-deckards/`.

Westmacott, G. and S. Daley (2015). The design and depletion of retirement portfolios. PWL Capital White Paper, `http://www.pwlcapital.com/design-depletion-of-retirement-portfolios/`.