Management of Portfolio Depletion Risk Through Optimal Life Cycle Asset Allocation

Peter A. Forsyth* Kenneth R. Vetzal† Graham Westmacott‡

Draft version: December 14, 2018

Abstract

Members of defined contribution (DC) pension plans must take on additional responsibilities for their investments, compared to participants in defined benefit (DB) pension plans. The transition from DB to DC plans means that more employees are faced with these responsibilities. We explore the extent to which DC plan members can follow financial strategies that have a high chance of resulting in a retirement scenario that is fairly close to that provided by DB plans. Retirees in DC plans typically must fund spending from accumulated savings. This leads to the risk of depleting these savings, i.e. portfolio depletion risk. We analyze the management of this risk through life cycle optimal dynamic asset allocation, including the accumulation and decumulation phases. We pose the asset allocation strategy as an optimal stochastic control problem. Several objective functions are tested and compared. We focus on the risk of portfolio depletion at the terminal date, using such measures as conditional value at risk (CVAR) and probability of ruin. A secondary consideration is the median terminal portfolio value. The control problem is solved using a Hamilton-Jacobi-Bellman formulation, based on a parametric model of the financial market. Monte Carlo simulations which use the optimal controls are presented to evaluate the performance metrics. These simulations are based on both the parametric model and bootstrap resampling of 91 years of historical data. The resampling tests suggest that target-based approaches which seek to establish a safety margin of wealth at the end of the decumulation period appear to be superior to strategies which directly attempt to minimize risk measures such as the probability of portfolio depletion or CVAR. The target-based approaches result in a reasonably close approximation to the retirement spending available in a DB plan. There is a small risk of depleting the retiree's funds, but there is also a good chance of accumulating a buffer which can be used to manage unplanned longevity risk, or left as a bequest.

Keywords: Depletion risk, life cycle asset allocation, optimal control, decumulation

JEL Codes: G11, G22

1 Introduction

Nobel laureate William Sharpe has referred to decumulation (i.e. the use of savings to fund spending during retirement) as "the nastiest, hardest problem in finance" (Ritholz, 2017). Retirees are

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

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

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

confronted with portfolio depletion risk and longevity risk, as well as additional uncertainties associated with unexpected inflation, the level of other sources of income such as government benefits, and the changing utility of income over time. We focus on depletion risk, which is the chance of running out of money, even when the retirement period is specified, due to the demand for constant income from a volatile portfolio. Depletion risk can be assessed in a variety of ways, such as the probability of ruin (i.e. depleting savings to a level that is insufficient to fund planned withdrawals), or the magnitude of ruin if it occurs.

Depletion risk is clearly much less important for individuals in traditional DB pension plans, and is arguably only a concern if the plan is insolvent. However, solvent DB plans still leave retirees exposed to other risks such as possible reductions of government benefits, elevated health care expenses, etc. The long-term shift from DB to DC pension plans exposes more individuals to significant depletion risk, and has resulted in increased focus on how DC plan members should manage their investments both during the accumulation phase when saving for retirement and during the post-retirement decumulation period. In this article we explore the question of how investments can be managed by DC plan members (or other individuals who are saving on their own for retirement) so as to give a reasonably close approximation to the retirement spending available from a DB plan.

Numerous studies have explored issues related to the management of investments for retirement saving and spending. Many of them focus on either just the accumulation phase or the decumulation period, as opposed to the entire life cycle. Examples of studies that concentrate on accumulation include Cairns et al. (2006), Vigna (2014), Yao et al. (2014), Guan and Liang (2014), Wu and Zeng (2015), Chen and Delong (2015), Donnelly et al. (2015), Donnelly et al. (2017), Dahlquist et al. (2018), and Christiansen and Steffensen (2018). Representative papers that focus on various aspects of decumulation include 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 an overview of the various strategies for decumulation, we refer the reader to MacDonald et al. (2013). Among the articles that consider both accumulation and decumulation (i.e. the entire life cycle) are Dammon et al. (2004), Cocco et al. (2005), Blake et al. (2014), Horneff et al. (2015), Campanele et al. (2015), Fagereng et al. (2017), and Michaelides and Zhang (2017).

The papers cited above use a variety of approaches and address a diverse set of issues. It is standard in the financial economics literature to develop models based on maximizing some form of utility function, typically defined over intermediate consumption as well as a final bequest. In contrast, actuarial papers are often based on statistical criteria (e.g. mean-variance optimization, minimization of ruin probability, etc.), as well as utility maximization. Some studies assume that retirees will be forced to annuitize at a pre-determined age, others try to determine the best time to annuitize given the option to do so. Many articles incorporate the effects of stochastic labour income during the accumulation phase, or of different models for financial market returns (e.g. stochastic interest rates, stochastic volatility of equity market returns, regime-switching specifications, etc.).

Our focus is deliberately narrow, compared to many other studies. As noted previously, the fundamental issue we address is the extent to which an asset allocation scheme can be designed to lead to approximately the same outcome as that which would be experienced by a DB pension plan member. This leads us to avoid utility maximization in favour of objective functions based on statistical criteria, for the following reasons: (i) the design of a DB plan cannot take into account the individual preferences (e.g. risk-aversion) of plan members; (ii) standard utility functions which have infinite marginal utility at zero wealth cannot be applied in our setting because we can only minimize ruin probability, not completely eliminate it; and (iii) in our experience typical retirees are concerned with concrete issues such as the probability of portfolio depletion and the size of a possible bequest, so it is generally easier for practitioners to discuss these issues with their clients as opposed to trying to determine the parameters of a utility function.[1] We also ignore other factors such as stochastic labour income, other wealth that plan members may have (e.g. home ownership), government social programs, taxes, etc. At an individual level, such factors can clearly be quite important. However, it would not be feasible for a DB plan to incorporate them, so we exclude them here.

We consider a prototypical DC plan holder who is assumed to be 35 years old with stable employment. This individual plans to work until age 65, so our accumulation period is 30 years. During this time, the annual combined contribution to the holder's DC account by the employee and employer amounts to 20% of the employee's salary, which is assumed to be constant in real terms.[2]

During the decumulation phase, the retiree faces longevity risk and perhaps has a bequest motive. Due to risk pooling and the earning of *mortality credits*, it is often suggested that retirees should purchase annuities. However, it is well-known that most retirees are very reluctant to do so (MacDonald et al., 2013; Peijnenburg et al., 2016). In fact, MacDonald et al. (2013) list 39 reasons (behavioural and rational) to avoid annuitization. For example, most annuities do not provide true inflation protection, are poorly priced, and retirees desire to have access to capital for emergencies. In addition, in the current low interest rate environment, annuities generate very low cash flow streams. We refer the reader to MacDonald et al. (2013) for a thorough discussion of this issue. We therefore assume that our 35 year old DC plan holder has no intention to annuitize upon retirement, and so adopts an asset allocation strategy which will be operational *to and through* the retirement date. Recommended final salary replacement ratios (including government programs) are variously estimated from 40% to 70%. We assume constant real withdrawals of 40% of final salary (excluding any government benefits). Given possible increases in longevity, and having ruled out the use of annuities, it is prudent for our plan holder to allow for a lengthy decumulation period, which we assume to last 30 years. By using a fixed, lengthy time for fixed cash outflows, we sidestep the issue of longevity risk. We recognize that this is a weakness of our analysis, but it appears to be a reasonable approach in the absence of any desire to annuitize. Having ruled out annuitization, a conservative estimate of longevity seems prudent.

Of course, one of the challenges of planning an investment strategy for DC plan holders is that retirees as a group have diverse and changing requirements. We focus attention on a specific group with the following attributes:

- the DC plan investors know what real income is required from their investment capital during retirement;
- the primary goal of retirees is to sustain the specified income for a pre-defined period with minimal risk;
- retirees value a portfolio which de-risks rapidly, while satisfying their income requirements.

[1] Note that industry surveys suggest that retirees are extremely concerned about possibly exhausting their savings. For example, see `https://www.allianzlife.com/about/news-and-events/news-releases/Generations-Ahead-Study-2017`.

[2] As mentioned above, we ignore labour income risk. Many studies assume that real earnings are expected to follow a hump-shaped pattern, rising rapidly until about age 35, then more slowly until around age 45-50, and slowly declining thereafter (see, e.g. Cocco et al., 2005; Blake et al., 2014). It is common to add diffusive shocks to this trend, though Cocco et al. (2005) calculate that the utility costs of assuming labour income has no risk are not high, absent a very large negative shock to income, which would be highly unlikely in a diffusive model. It is also worth noting here that the hump-shaped pattern described above has been questioned recently by Rupert and Zanella (2015), who find that while wage rates do rise rapidly in the early years of a typical employee's career, they do not decline prior to retirement. Average income does fall on average during those years, but this is due to a reduction in hours worked by some employees transitioning into retirement.

We specifically exclude from our consideration those who consider other goals (e.g. a bequest) more important than sustaining a specified level of income. If the investor is fortunate to be able to secure their retirement income at an agreed level, we assume that in this circumstance, their preference is to leave any surplus as a bequest, rather than increasing spending.

For our working scenario, we postulate 30 years of accumulation at 20% of salary followed by 30 years of decumulation at 40% of final salary. The contributions and withdrawals are each assumed to be constant in real terms. Specifying a constant real annual withdrawal means that, as noted above, we are attempting so far as possible to create a defined benefit (DB) experience. We adjust the asset allocation throughout the life cycle to minimize the adverse consequences.

In the absence of annuitization, we emphasize that the ability to generate a specified real income with minimal risk is the best a DC plan holder can expect in the quest to obtain a DB plan experience. She retains longevity risk as well as the assumption of a finite (specified) investment horizon.

Target date funds (TDFs) are popular investment products which cater to the market for retirement saving. A standard TDF begins with a high allocation to equities and moves to a higher weighting in bonds as retirement approaches. The fraction invested in equities over time is called a *glide path*. Typically, these glide paths are *deterministic* strategies, i.e. the equity fraction is only a function of time to go. Total assets invested in US TDFs at the end of 2017 were over \$1.1 trillion.[3] The rationale for the high initial equity allocation to stocks is often based on human capital considerations, i.e. a young DC plan holder has many years of bond-like cash flows from employment, and can take on a large equity risk in the DC account. As retirement approaches, the future income from employment diminishes, and hence the holder should switch to bonds. However, recent work calls into question the effectiveness of the TDF type of approach (see, e.g. Arnott et al., 2013; Graf, 2017; Westmacott and Daley, 2015; Forsyth et al., 2017; Forsyth and Vetzal, 2017b). For example, Forsyth et al. (2017) and Forsyth and Vetzal (2017b) show that for a fixed value of target expected wealth at the end of the accumulation period, there is always a constant weight strategy that achieves the same target expected wealth as a deterministic glide path with a similar cumulative standard deviation. More recently, deterministic strategies have also been suggested for *to and through* funds, i.e. both the accumulation and the decumulation phases (O'Hara and Daverman, 2017).

We initially consider some deterministic strategies, for which the asset allocation is either constant or a deterministic function of time (i.e. a glide path). Our main focus, however, is on *adaptive* strategies, in which the asset allocation depends on realized wealth to date in addition to time. We specify adaptive strategies as optimal stochastic control problems. We test several candidate objective functions, and assess their suitability in terms of metrics of interest to retirees such as the probability of portfolio depletion (i.e. ruin) and the conditional value at risk (CVAR), which measures how severe ruin is likely to be if it does occur. Amongst objective functions which have similar risk statistics, we prefer the strategies which generate larger median values of terminal wealth (e.g. a potential bequest). In effect, we view the objective function strictly as a means to shape the probability density of the outcome of the investment process, not as an end in itself. The first objective function we consider is minimizing the probability of ruin, before the end of the decumulation phase. We then consider mean-CVAR strategies (Gao et al., 2017; Strub et al., 2017), as well as target-based approaches (Vigna, 2014; Menoncin and Vigna, 2017) that correspond to multi-period mean variance strategies (Li and Ng, 2000; Dang et al., 2017).

We assume that the investment account contains only a stock index and a bond index. We model the real (inflation-adjusted) stock index as following a jump diffusion model (Kou and Wang,

[3]Investment Company Fact Book (2018), available at `www.ici.org`.

2004). We fit the parameters of this model to monthly US data over the 1926:1-2016:12 period. We consider two markets in our simulation analysis. The *synthetic market* assumes that the stock and bond processes follow the models with constant parameters fit to the historical time series. Given an objective function, we determine optimal strategies by solving a Hamilton-Jacobi-Bellman equation in the synthetic market. We use a fully numerical approach, which allows us to impose realistic constraints: infrequent rebalancing (yearly) and no leverage/no-shorting constraints. The entire distribution function of the strategy is then determined by Monte Carlo simulations in the synthetic market. As a stress test, we apply these strategies to bootstrap resampling of the historical data, which we refer to as the *historical market*. The bootstrap tests make no assumptions about the actual processes followed by the stock and bond indexes. In some cases, we reject strategies which appear promising based on synthetic market results due to poor performance in the bootstrapped historical market. This highlights the importance of resampling to assess the robustness of recommended strategies, which has rarely been done in the prior actuarial literature on long-term asset allocation.[4]

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 market index fund. For example, many wealth managers have funds which have a fixed weight of domestic and foreign equity markets.

The investment horizon (over both the accumulation and decumulation phases) is T. S_t and B_t respectively denote the *amounts* invested in the risky and risk-free assets at time t, $t \in [0, T]$. In general, these amounts will depend on the investor's strategy over time, including contributions, withdrawals, and portfolio rebalances, as well as changes in the unit prices of the assets. Suppose for the moment that the investor does not take any action with respect to the controllable factors, so that any change in the value of the investor's portfolio is due to changes in asset prices. We refer to this as the absence of control. In this case, we assume that S_t follows a jump diffusion process.

Let $t^- = t - \epsilon, \epsilon \to 0^+$, i.e. t^- is the instant of time before t, and let ξ be a random number representing a jump multiplier. When a jump occurs, $S_t = \xi S_{t^-}$. Allowing discontinuous jumps lets us explore the effects of severe market crashes on the risky asset holding.

More precisely, in the absence of control, S_t evolves according to

$$\frac{dS_t}{S_{t^-}} = (\mu - \lambda E[\xi - 1])\ dt + \sigma\, d\mathcal{Z} + d\left(\sum_{i=1}^{\pi_t} (\xi_i - 1)\right), \tag{2.1}$$

where μ is the (uncompensated) drift rate, σ is the volatility, $d\mathcal{Z}$ is the increment of a Wiener process, π_t is a Poisson process with positive intensity parameter λ, and ξ_i are i.i.d. positive random variables having distribution given below. Moreover, ξ_i, π_t, and $\mathcal{Z}$ are assumed to all be mutually independent.

We assume that $\log \xi$ follows a double exponential distribution (Kou and Wang, 2004). If a jump occurs, p_{up} is the probability of an upward jump, while $1 - p_{up}$ is the chance of a downward jump. The mean upward and downward log jump sizes are $1/\eta_1$ and $-1/\eta_2$ respectively. The density function for $y = \log \xi$ is

$$f(y) = p_{up} \eta_1 e^{-\eta_1 y} \mathbf{1}_{y \geq 0} + (1 - p_{up}) \eta_2 e^{\eta_2 y} \mathbf{1}_{y < 0}. \tag{2.2}$$

[4] Donnelly et al. (2017) conduct some resampling experiments, but only for the equity market (not the bond market), and over a relatively short period of time.

We note that

$$E[y] = \frac{p_{up}}{\eta_1} - \frac{(1-p_{up})}{\eta_2} \; ; \qquad E[\xi] = \frac{p_{up}\eta_1}{\eta_1 - 1} + \frac{(1-p_{up})\eta_2}{\eta_2 + 1}. \tag{2.3}$$

We focus on jump diffusion models for long-term equity dynamics since sudden drops in the equity index can have a devastating impact on retirement portfolios, particularly during the decumulation phase. Since we consider discrete rebalancing, the jump process models the cumulative effects of large market movements between rebalancing times.[5]

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

$$dB_t = rB_t\, dt, \tag{2.4}$$

where r is the (constant) risk-free rate. This is obviously a simplification of the actual bond market. In any case, we will test our strategies in a bootstrapped historical market which introduces inflation shocks and stochastic interest rates.

We define the investor's total wealth at time t as

$$\text{Total wealth } \equiv W_t = S_t + B_t. \tag{2.5}$$

Since we specify the real withdrawals during decumulation, the objective functions which we consider below are all defined in terms of terminal wealth W_T. If the portfolio is solvent, we impose the constraints that shorting stock and using leverage (i.e. borrowing) are not permitted, which would be typical of a retirement savings account. In the event of portfolio depletion, withdrawals cause an accumulation of debt.

3 Data, synthetic market, and historical market

The data used in this work 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 (CRSP) Deciles (1-10) index. This is a total return value-weighted index of US stocks. We also use one month Treasury bill (T-bill) returns for the risk-free asset.[6] 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. We use real indexes since long-term retirement saving should be attempting to achieve real (not nominal) wealth goals. All of the data used was at the monthly frequency, with a sample period of 1926:1 to 2016:12.

In our tests, we consider a *synthetic* and an *historical* market. The synthetic market is generated by assuming processes (2.1) and (2.4). We fit the parameters to the historical data using the methods described in Appendix A. We then use these parameters to determine optimal strategies and carry out Monte Carlo computations. As a test of robustness, we also carry out tests using bootstrap resampling of the actual historical data, which we call the historical market. In this case, we make no assumptions about the underlying stochastic processes. We use the stationary block resampling method described in Appendix B. A crucial parameter for block bootstrap resampling is the expected blocksize. We carry out our tests using a range of expected blocksizes. Although the

[5] A possible extension would be to incorporate stochastic volatility. However, previous work has shown that stochastic volatility effects are small for the long-term investor (Ma and Forsyth, 2016). This can be traced to the fact that stochastic volatility models are mean-reverting, with typical mean-reversion times of less than one year.

[6] We have also carried out tests using a 10 year US treasury as the bond asset (Forsyth and Vetzal, 2017a). The results are qualitatively similar to those reported in this paper.

Investment horizon (years)	60
Equity market index	Value-weighted CRSP deciles 1-10 US market index
Risk-free asset index	1-month T-bill
Initial investment W_0	0.0
Real investment each year	20.0 ($0 \le t_i \le 30$), -40.0 ($31 \le t_i \le 60$)
Rebalancing interval (years)	1
Market parameters	See Appendix A

TABLE 4.1: *Input data for examples. Cash is invested at* $t_i = 0,1,\ldots,30$ *years, and withdrawn at* $t_i = 31,32,\ldots,60$ *years. Units for real investment: thousands of dollars.*

absolute performance of variance strategies is mildly sensitive to the choice of blocksize, the relative performance of the various strategies appears to be insensitive to blocksize. See Appendix B for more discussion.

4 Investment scenario

Let the inception time of the investment be $t_0 = 0$. We consider a set $\mathcal{T}$ of pre-determined *rebalancing times*,

$$\mathcal{T} \equiv \{t_0 = 0 < t_1 < \cdots < t_M = T\}. \tag{4.1}$$

For simplicity, we specify $\mathcal{T}$ to be equidistant with $t_i - t_{i-1} = \Delta t = T/M$, $i = 1,\ldots,M$. At each rebalancing time t_i, $i = 0,1,\ldots,M$, the investor (i) injects an amount of cash q_i into the portfolio, and then (ii) rebalances the portfolio. At $t_M = T$, the portfolio is liquidated. If $q_i < 0$, this corresponds to cash withdrawals. Let $t_i^- = t_i - \epsilon$ ($\epsilon \to 0^+$) be the instant before rebalancing time t_i, and $t_i^+ = t_i + \epsilon$ be the instant after t_i. Let $p(t_i^+, W_i^+) = p_i$ be the fraction in the risky asset at t_i^+.

Table 4.1 shows the parameters for our investment scenario. As discussed previously, this corresponds to an individual with a constant salary of \$100,000 per year (real) who saves 20% of her salary for 30 years, then withdraws 40% of her final real salary for 30 years in retirement. The target salary replacement level of 40% is at the lower end of the recommended range, but it is possible that government benefits could increase this to a more desirable level. We do not consider escalating the (real) contribution during the accumulation phase (which also impacts the desired replacement ratio), although this is arguably more realistic. Assuming flat contributions and withdrawals, we can interpret the above scenario as an investment strategy which allows real withdrawals of twice as much as real contributions. We shall see that this rather modest objective still entails significant risk. As indicated in Table 4.1, we assume yearly rebalancing.[7]

5 Constant weight strategies and linear glide paths

Let p denote the fraction of total wealth that is invested in the risky asset, i.e.

$$p = \frac{S_t}{S_t + B_t}. \tag{5.1}$$

[7]More frequent rebalancing has little effect for long-term (> 20 years) investors (Forsyth and Vetzal, 2017c).

Strategy	$Median[W_T]$	$Mean[W_T]$	$std[W_T]$	$Pr[W_T < 0]$	5% CVAR
Glide path	935	1385	1795	.15	-483
$p = .40$	992	1542	2093	.16	-482
$p = .60$	2922	5422	8882	.093	-516
$p = .80$	6051	14832	34644	.082	-592

TABLE 5.1: *Synthetic market results for deterministic strategies, assuming the scenario given in Table 4.1. W_T denotes real terminal wealth after* 60 *years, measured in thousands of dollars. Statistics based on* 6.4×10^5 *Monte Carlo simulation runs. The constant weight strategies have equity fraction p. The glide path is linear with $p_{\max} = .80$ and $p_{\min} = 0.0$.*

A *deterministic* glide path restricts the admissible strategies to those where $p = p(t)$, i.e. the strategy depends only on time and cannot take into account the actual value of W_t at any time. Clearly this is a very restrictive assumption, but it is commonly used in TDFs.

We consider two cases: $p(t) = const.$ and a linear glide path

$$p(t) = p_{\max} + (p_{\min} - p_{\max}) \frac{t}{T}. \tag{5.2}$$

Note that this is a *to and through* strategy, since $t = 0$ indicates the beginning of the accumulation phase, while $t = T$ represents the end of the decumulation phase.

Monte Carlo simulations were carried out for the scenario given in Table 4.1, using constant weight strategies and the linear glide path in equation (5.2). We run these simulations in the synthetic market, assuming processes (2.1) and (2.4), with parameters given in Appendix A. The results are shown in Table 5.1. Here, 5% CVAR refers to mean of the worst 5% of the outcomes, defined in terms of wealth, not losses.[8]

The results in Table 5.1 show the high risks associated with deterministic strategies. Note the very high dispersion of final wealth as indicated by the large standard deviations and the large differences between the means and medians. Consistent with the findings reported for the accumulation phase by Forsyth et al. (2017) and Forsyth and Vetzal (2017b), the results here for the entire life cycle for a linear glide path are similar to the results for a constant weight strategy having the same time-averaged weighting in stocks (i.e. $p = .40$ in this case). It is interesting to note that while the high constant weighting in equities ($p = 0.8$) has a much higher dispersion of final wealth compared to lower allocations, the $p = 0.8$ strategy has a smaller probability of ruin (i.e. $Pr[W_T < 0]$) and larger median value of terminal wealth compared to the lower equity allocation strategies. The downside for the $p = .8$ case compared to the $p = .6$ case is an increase in the tail risk (5% CVAR).

Table 5.2 shows the results for constant proportion strategies based on bootstrap resampling of the historical market, for a range of expected blocksizes.[9] Since we sample simultaneously from the stock and bond historical time series, the choice of blocksize is not obvious (see Appendix B). A reasonable choice would appear to be an expected blocksize of $\simeq 2$ years. Nevertheless, the ranking of the three constant weight strategies is preserved across all blocksizes, i.e. the higher allocation to equities is superior (in terms of $Pr[W_T < 0]$) compared to the smaller allocation to equities. Note that the historical backtests show that the probability of ruin for a typical suggested equity weighting of .6 is in the range $.05 - .09$ depending on the assumed expected blocksize.

[8]See Appendix C for a precise definition of CVAR as used in this work.

[9]Results for the linear glide path are again similar to the constant proportion case with $p = .40$ and have been excluded from Table 5.2 to save space.

Strategy	$\hat{b}$	$Median[W_T]$	$Mean[W_T]$	$std[W_T]$	$Pr[W_T < 0]$	5% CVAR
$p = .40$	0.5	900	1337	1683	.16	-490
$p = .60$	0.5	2767	4592	6251	.085	-488
$p = .80$	0.5	5893	12120	21278	.071	-540
$p = .40$	1.0	955	1367	1637	.16	-493
$p = .60$	1.0	2896	4614	5814	.081	-466
$p = .80$	1.0	6075	12028	18991	.068	-514
$p = .40$	2.0	961	1339	1530	.15	-461
$p = .60$	2.0	2931	4248	4955	.07	-389
$p = .80$	2.0	6151	10865	15023	.054	-411
$p = .40$	5.0	965	1306	1451	.14	-438
$p = .60$	5.0	2890	4068	4326	.051	-275
$p = .80$	5.0	5986	9768	12543	.034	-190

TABLE 5.2: *Historical market results for constant proportion strategies with equity fraction p, assuming the scenario given in Table 4.1. W_T denotes real terminal wealth after* 60 *years, measured in thousands of dollars. Statistics based on* 10,000 *stationary block bootstrap resamples of the historical data from 1926:1 to 2016:12. $\hat{b}$ is the expected blocksize, measured in years.*

6 Adaptive strategies: overview

We will attempt to improve on deterministic strategies by allowing the rebalancing strategy to now depend on the accumulated wealth, i.e. $p_i = p_i(W_i^+, t_i)$. We will specify an objective function, and compute the optimal controls in the synthetic market. This involves the numerical solution of a Hamilton-Jacobi-Bellman (HJB) equation to determine the controls. We use the numerical methods from Dang and Forsyth (2014; 2016) and Forsyth and Labahn (2018), and refer the reader to these sources for a detailed description of the HJB equation and solution techniques. We emphasize that, given an objective function, solving the HJB equation gives the provably optimal strategy in the constant parameter synthetic market. The following several sections consider various possible objective functions in this context. Primarily, we focus on downside risk measures, such as CVAR and probability of ruin. We regard the median terminal wealth to be of secondary importance.

7 Minimize probability of ruin

Many retirees place a premium on reducing the probability of ruin, i.e. portfolio depletion. Therefore, as a first attempt at defining a suitable objective function, we directly minimize probability of ruin. A similar objective function for the accumulation phase of DC plans has been suggested in Tretiakova and Yamada (2011). Consider a level of terminal wealth W^{min}. We wish to solve the following optimization problem:

$$\min_{\{(p_0,c_0),\,\ldots,\,(p_{M-1},c_{M-1})\}} Pr\left[W_T < W^{min}\right]$$

$$\text{subject to} \begin{cases} (S_t, B_t) \text{ follow processes (2.1)-(2.4)}; \quad t \notin \mathcal{T} \\ W_i^+ = W_i^- + q_i - c_i\,;\ S_i^+ = p_i W_i^+\,;\ B_i^+ = W_i^+ - S_i^+\,; \quad t \in \mathcal{T} \\ p_i = p_i(W_i^+, t_i)\,;\ 0 \le p_i \le 1 \\ c_i = c_i(W_i^- + q_i, t_i)\,;\ c_i \ge 0 \end{cases} . \tag{7.1}$$

We recognize objective function (7.1) as minimizing the probability that the terminal wealth W_T will be less than W^{min}. If $W^{min} = 0$, then this will minimize the probability of portfolio depletion.

In problem (7.1), we withdraw surplus cash $c_i(W_i^- + q_i, t_i)$ from the portfolio if investing in the risk-free asset ensures that $W_T \ge W^{min}$. More precisely, let

$$Q_\ell = \sum_{j=\ell+1}^{j=M-1} e^{-r(t_j - t_\ell)} q_j \tag{7.2}$$

be the discounted future contributions as of time t_ℓ. If

$$(W_i^- + q_i) > W^{min} e^{-r(T-t_i)} - Q_i, \tag{7.3}$$

then an optimal strategy is to (i) withdraw surplus cash $c_i = W_i^- + q_i - \left(W^{min} e^{-r(T-t_i)} - Q_i\right)$ from the portfolio; and (ii) invest the remainder $\left(W^{min} e^{-r(T-t_i)} - Q_i\right)$ in the risk-free asset. This is an optimal strategy in this case since $Pr[W_T < W^{min}] = 0$, which is the minimum of problem (7.1).

In the following, we will refer to $c_i > 0$ as *surplus cash*. We assume that any surplus cash is invested in the risk-free asset. Of course, it is also possible to invest it in the risky asset. Some experiments with this alternative approach showed a large effect on $E[W_T]$, but very little impact on $Median[W_T]$, $Pr[W_T < 0]$, and CVAR. Hence we assume that surplus cash is invested in the risk-free asset for simplicity.

If at any point surplus cash is generated (as defined in equation (7.3)), then the objective function is identically zero, and the surplus can be invested in any combination of the stock and bond. This is obviously not a unique strategy, since $E[W_T]$ will depend on how the surplus cash is invested. Hence we must precisely specify what we do with the surplus cash, in order to make Problem 7.1 well posed.

In our summary statistics, we will include surplus cash in measures such as $E[W_T]$, but we will exclude it from the standard deviation $std[W_T]$ since this is supposed to be a measure of risk. Along any path where surplus cash is generated, we have no probability of ruin. But including the surplus cash in $std[W_T]$ will generally increase $std[W_T]$, which seems counter-informative since there is no risk (in the sense of ruin) along this path. In any case, we do not believe that $std[W_T]$ is a very useful risk measure for these types of problems, due to the highly skewed distribution of terminal wealth.

We begin by computing and storing the optimal controls from solving problem (7.1) with $W^{min} = 0$. In other words, we try to minimize the probability of portfolio depletion before year 60. To assess this strategy, we use these controls as input to a Monte Carlo simulation in the synthetic market. Recall that in this case the simulated paths will have exactly the same statistical properties as those assumed when generating the optimal controls. The results are shown in the first row of Table 7.1. In this idealized setting, the final wealth distribution has a median that is almost zero, but also about a 2% chance of being less than zero. Figure 7.1 plots the cumulative distribution

$\hat{b}$	W^{min}	$Median[W_T]$	$Mean[W_T]$	$std[W_T]$	$Pr[W_T < 0]$	5% CVAR
			Synthetic market			
NA	0	3.67	20.7	88.3	.0195	−223
NA	200	204	217	134	.0284	−235
NA	400	406	410	185	.036	−310
NA	600	607	601	240	.042	−436
NA	1000	1009	973	360	.053	−562
			Historical market			
0.5	0	187	199	103	.07	−353
1.0	0	207	236	105	.068	−346
2.0	0	228	283	86	.053	−245
5.0	0	260	341	59	.034	−121
0.5	200	412	417	141	.043	−335
1.0	200	434	456	145	.040	−337
2.0	200	456	512	118	.030	−183
5.0	200	492	579	85	.017	−9.3

TABLE 7.1: *Optimal control determined by solving problem* (7.1), *i.e.* $\min Pr[W_T < W^{min}]$ *in the synthetic market, with* W^{min} *as indicated, assuming the scenario in Table 4.1.* W_T *denotes real terminal wealth after* 60 *years, measured in thousands of dollars. Statistics for the synthetic market case are based on* 6.4×10^5 *Monte Carlo simulation runs. Statistics for the historical market cases are based on* 10,000 *stationary block bootstrap resamples of the historical data from 1926:1 to 2016:12.* $\hat{b}$ *is the expected blocksize, measured in years. Surplus cash is included in the mean, median, CVAR and probability of ruin, but excluded from the standard deviation.*

function of W_T for this case. The sharp increase in the distribution function near $W_T = 0$ suggests that this strategy will be very sensitive to the asset market parameters. Table 7.1 also shows the results for increasing values of W^{min} in the synthetic market. As expected, increasing W^{min} increases the median of W_T, but this comes at the expense of increasing the probability of ruin. Figure 7.2 shows the percentiles of the total wealth (panel (a)) and the optimal fraction invested in equities (panel (b)) as a function of time. Figure 7.2(a) shows greater dispersion between the 5th and 95th percentiles during the accumulation phase ($t \leq 30$) than during the decumulation phase ($30 < t \leq 60$). From Figure 7.2(b), the median fraction invested in the risky stock index is surprisingly low, essentially de-risking completely by the end of the accumulation period. Figure 7.2(c) shows the heat map of the minimize ruin strategy.

We next test this strategy with $W^{min} = 0$ in the historical market. This implies using the same optimal controls as above, but instead simulating by bootstrap resampling of the historical data over the 1926:1 to 2016:12 period (see Appendix B). Results for several different expected blocksizes $\hat{b}$ ranging from 0.5 years to 5.0 years are provided in Table 7.1. These results differ substantially from the synthetic market case: $Median[W_T]$ and $Mean[W_T]$ are markedly higher in the historical market, but so are the risk measures $std[W_T]$, $Pr[W_T < 0]$, and 5% CVAR (except if $\hat{b} = 5$ years). Since we are directly trying to minimize $Pr[W_T < 0]$, it is worth emphasizing that this ruin probability is higher than in the synthetic market by a factor of more than 3 for the two shortest expected blocksizes. Even when $\hat{b} = 5$ years, the ruin probability is almost 75% higher in the historical market. These results are consistent with our earlier discussion about Figure 7.1:

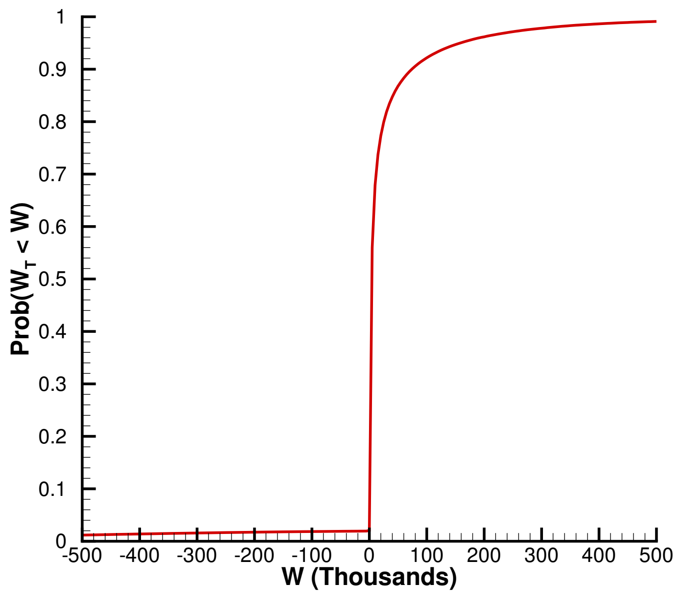

FIGURE 7.1: *Cumulative distribution function. Optimal control determined by solving problem* (7.1)*, i.e.* $\min Pr[W_T < W^{min}]$ *in the synthetic market, with* $W^{min} = 0$*, assuming the scenario in Table 4.1. Distribution computed from* 6.4×10^5 *Monte Carlo simulation runs in the synthetic market. Surplus cash is included in the distribution function.*

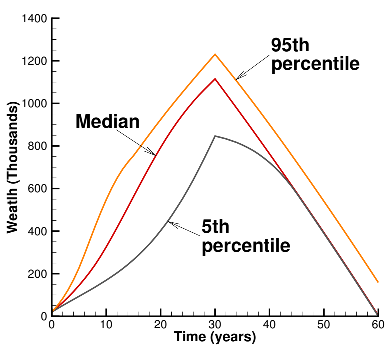

(a) *Percentiles of accumulated wealth.*

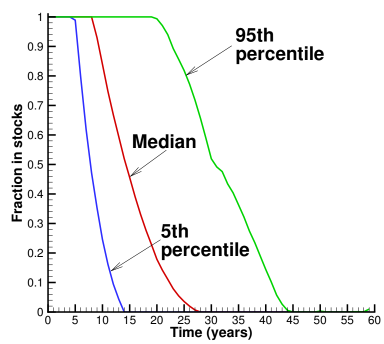

(b) *Percentiles of optimal fraction in equities.*

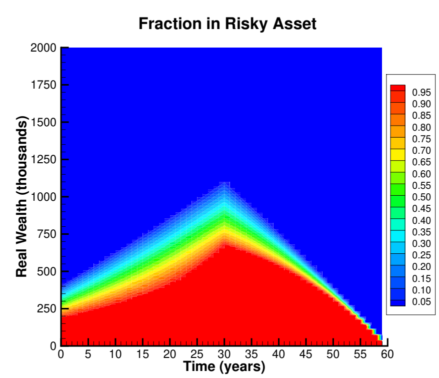

(c) *Optimal control heat map.*

FIGURE 7.2: *Percentiles of real wealth and the optimal fraction invested in equities. Optimal control computed by solving problem* (7.1)*, i.e.* $\min Pr[W_T < W^{min}]$ *in the synthetic market, with* $W^{min} = 0$*, assuming the scenario in Table 4.1. Statistics based on* 6.4×10^5 *Monte Carlo simulation runs in the synthetic market. Surplus cash is included in the real wealth percentiles.*

the very sharp increase in the cumulative distribution function at $W_T = 0$ for the synthetic market implies that performance is unlikely to be robust to departures from the statistical properties of the idealized synthetic market, which is exactly what happens in the historical market.

The instability here can be traced to the use of bootstrap historical real interest rates. For example, if the case with $\hat{b} = 2.0$ years is repeated using the fixed average historical real interest rate (i.e. $r = .004835$) for all time periods, but with the bootstrapped historical stock returns, then $Pr[W_T < 0] = .013$ compared to the value of .053 in Table 7.1. In this case, since $W^{min} = 0$ under the objective function (7.1), any errors in prediction of the real bond return become magnified, due to the very rapid de-risking. It could be argued that the use of bootstrapped real bond returns is very pessimistic with a blocksize of 2.0 years. Effectively, this simulates a market where the investor de-risks rapidly after the accumulation phase, but then the strategy fails due to real interest rate shocks.

In an effort to determine a more robust strategy, we experimented with setting $W^{min} > 0$, so as to provide a buffer of wealth as insurance against misspecification of real interest rates. The last four rows of Table 7.1 show the results obtained by computing and storing the optimal strategy from solving problem (7.1) with $W^{min} = 200$ in the synthetic market and then using this strategy in bootstrap resampling tests. As expected, this strategy is much more stable in terms of the probability of ruin compared to the $W^{min} = 0$ case. By any measure, the bootstrap results for $W^{min} = 200$ are superior to the those obtained with $W^{min} = 0$.[10]

We can summarize our attempts to minimize probability of ruin as follows. Although at first glance it would appear that minimizing the probability of negative terminal wealth (i.e. portfolio depletion) is a reasonable objective, our tests call this into question. Clearly, aiming for zero final wealth is too sensitive to modelling parameters to be useful. This sensitivity appears to be solely due to the use of bootstrapped bond return data and not due to the bootstrapped equity return data. Due to rapid de-risking, this strategy is sensitive to real interest rate shocks along any paths with early allocation to the bond index. The bootstrap resampling approach introduces random (and potentially large) real interest rate shocks into the market, which occur more often as the expected blocksize gets smaller. It could be argued that this is unduly pessimistic, but we contend that this is a useful stress test. This sensitivity to real interest rate shocks is ameliorated somewhat by setting the final wealth target to be a non-zero amount. However, comparing the historical market results in Tables 5.2 (constant weight allocations) and 7.1 (minimizing probability of ruin), it seems that the median terminal wealth is reduced significantly in order to reduce the probability of portfolio depletion.

8 Mean-CVAR optimization

As another possible objective, we consider minimizing the mean of the worst α fraction of outcomes (i.e. CVAR). Recall that we define CVAR in terms of terminal wealth, not losses, so we want to maximize CVAR.

Let $\mathcal{P} = \{p_0, p_1, \ldots, p_{M-1}\}$ be the set of controls at $t \in \mathcal{T}$. Let CVAR_α denote the CVAR at

[10] Experiments with larger values of W^{min} increased $Pr[W_T < 0]$ in the bootstrap tests.

level α. For a fixed value of α and a scalar κ, the mean-CVAR optimization problem is:

$$\max_{\mathcal{P}} \quad \text{CVAR}_\alpha^{\mathcal{P}} + \kappa E^{\mathcal{P}}[W_T]$$

$$\text{subject to} \begin{cases} (S_t, B_t) \text{ follow processes (2.1)-(2.4)}; \quad t \notin \mathcal{T} \\ W_i^+ = W_i^- + q_i\,;\ S_i^+ = p_i W_i^+\,;\ B_i^+ = W_i^+ - S_i^+\,; \quad t \in \mathcal{T} \\ p_i = p_i(W_i^+, t_i)\,;\ 0 \le p_i \le 1 \end{cases}, \tag{8.1}$$

where we use the notation $E^{\mathcal{P}}[\cdot]$ to emphasize that the expectation is computed using the control $\mathcal{P}$. We give a brief description of the algorithm used to solve problem (8.1) in Appendix C. Due to the leverage constraint imposed in equation (8.1), this optimization problem is well-posed without adding an additional funding level constraint on the terminal wealth (Gao et al., 2017).

Note that problem (8.1) is underspecified if $\kappa = 0$. By setting κ to a small positive number, e.g. $\kappa = 10^{-8}$, we can force the following strategy. Let W_α^* be the Value at Risk (VAR) at level α (see Appendix C). Along any path where we can achieve $W_T > W_\alpha^*$ with certainty by investing some amount in bonds, we then invest the remainder in stocks. More precisely, if

$$(W_i^- + q_i) > W_\alpha^* e^{-r(T-t_i)} - Q_i, \tag{8.2}$$

where Q_i is defined in equation (7.2), then the optimal strategy is to invest $W_\alpha^* e^{-r(T-t_i)} - Q_i$ in bonds and $(W_i^- + q_i) - W_\alpha^* e^{-r(T-t_i)} + Q_i$ in stocks. Effectively, we are maximizing CVAR_α (i.e. minimizing risk) with the tie-breaking strategy that if our wealth is large enough, then we invest the amount required to attain $W_T > W_\alpha^*$ in bonds and the excess in stocks. Conversely, if we set κ to a small negative number, then the optimal strategy along any path where equation (8.2) holds will be to switch all accumulated wealth to bonds.

It is well known that mean-CVAR optimization is not time consistent (Strub et al., 2017). In other words, if the optimization problem is restarted at some later time $t > 0$, then the strategy computed at this later time may not agree with the strategy computed at $t = 0$. The mean-CVAR strategy is termed a *pre-commitment* strategy, since the investor is committed to follow the strategy. However, we can view the pre-commitment mean-CVAR strategy, determined at $t = 0$, as the time consistent strategy for an alternative objective function for $t > 0$. This is discussed in Appendix C.

Table 8.1 shows the results. In the synthetic market, $Median[W_T]$, $Pr[W_T < 0]$, and 5% CVAR are the same for both $\kappa = \pm 10^{-8}$, but $Mean[W_T]$ and $std[W_T]$ are dramatically different. This indicates that the large mean of terminal wealth for $\kappa = +10^{-8}$ is due to small probability paths with extremely large values of W_T. The bootstrap (i.e. historical market) results are generally worse than the synthetic market results, except for an expected blocksize of 5 years. We also include a few representative results for non-trivial positive κ in Table 8.1. Note the very rapid increase in $E[W_T]$ as κ increases, and the enormous discrepancy between the mean and median values of W_T. The 5th, 50th, and 95th percentiles of wealth over time for the bootstrap tests are shown in Figure 8.1(a) for the case $\kappa = +10^{-8}$. Note the U-shape of the 95th percentile. This is due to the fact that on any path where the wealth satisfies equation (8.2), the optimal strategy is to invest the surplus in stocks since this will maximize expected terminal wealth. Contrast this with Figure 8.1(b), which shows the results when $\kappa = -10^{-8}$. Recall that this forces the strategy to invest in bonds along any path where the wealth satisfies equation (8.2). Figure 8.1(c) shows the heat map of the optimal mean-CVAR strategy.

9 Quadratic shortfall with expected value constraint

By now it seems clear that directly minimizing a measure of the risk of ruin is not a good strategy, since the results are not very stable under the bootstrap tests. Even in the synthetic market tests,

$\hat{b}$	κ	$Median[W_T]$	$Mean[W_T]$	$std[W_T]$	$Pr[W_T < 0]$	5% CVAR
			Synthetic market			
NA	10^{-8}	132	733	13844	.027	−185
NA	-10^{-8}	132	137	142	.027	−185
NA	.025	455	19935	104100	.029	−227
NA	.0375	1015	23593	111886	.031	−249
NA	0.05	1444	25297	115347	.033	−267
			Historical market			
0.5	10^{-8}	240	855	2957	.047	−283
0.5	-10^{-8}	165	193	182	.048	−285
1.0	10^{-8}	270	1053	2943	.046	−286
1.0	-10^{-8}	172	218	219	.048	−291
2.0	10^{-8}	320	1223	3343	.036	−184
2.0	-10^{-8}	186	259	253	.038	−189
5.0	10^{-8}	409	1434	3222	.024	−74
5.0	-10^{-8}	215	310	292	.025	−81

TABLE 8.1: *Optimal control determined by solving mean-CVAR problem* (8.1) *with* $\alpha = .05$ *in the synthetic market, assuming the scenario in Table 4.1.* W_T *denotes real terminal wealth after* 60 *years, measured in thousands of dollars. Statistics for the synthetic market cases are based on* 6.4×10^5 *Monte Carlo simulation runs. Statistics for the historical market cases are based on* 10,000 *stationary block bootstrap resamples of the historical data from 1926:1 to 2016:12.* $\hat{b}$ *is the expected blocksize, measured in years.* κ *specifies the asset allocation along paths where* $W_T > W_\alpha^*$ *with certainty; see equation* (8.2) *and accompanying discussion.*

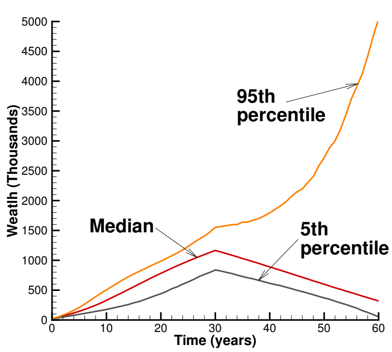

(a) $\kappa = +10^{-8}$.

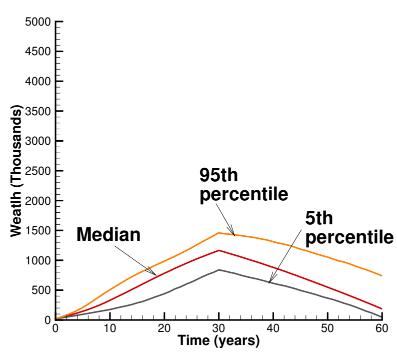

(b) $\kappa = -10^{-8}$.

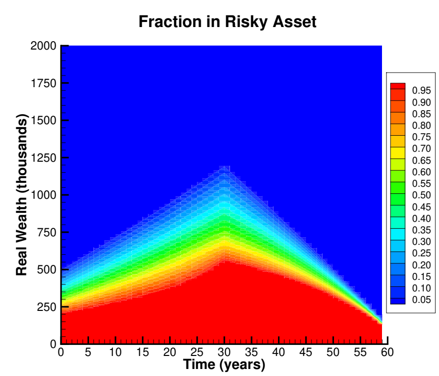

(c) *Optimal control heat map,* $\kappa = -10^{-8}$.

FIGURE 8.1: *Percentiles of real wealth in the historical market. Optimal control determined by solving mean-CVAR problem* (8.1) *with* $\alpha = .05$ *and* $\kappa = -10^{-8}$ *in the synthetic market, assuming the scenario in Table 4.1. Statistics based on* 10,000 *bootstrap resamples of the historical data from 1926:1 to 2016:12 with expected blocksize* $\hat{b} = 2$ *years.* κ *specifies the asset allocation along paths where* $W_T > W_\alpha^*$ *with certainty; see equation* (8.2) *and accompanying discussion.*

we can see that there is a very large cost incurred in terms of the median terminal wealth to reduce the probability of ruin by a small amount. It seems plausible to attempt to target a reasonable value of terminal wealth, and then to minimize the size of the shortfall. A natural candidate objective function in this case is minimizing the quadratic shortfall with respect to a target level of final wealth (W^*), as suggested by Menoncin and Vigna (2017) and others. Writing this problem more formally:

$$
\min_{\{(p_0,c_0),\ldots,(p_{M-1},c_{M-1})\}} E\left[(\min(W_T - W^*,0))^2\right]
$$

$$
\text{subject to} \begin{cases} (S_t, B_t) \text{ follow processes (2.1)-(2.4);} \quad t \notin \mathcal{T} \\ W_i^+ = W_i^- + q_i - c_i\,;\ S_i^+ = p_i W_i^+\,;\ B_i^+ = W_i^+ - S_i^+\,; \quad t \in \mathcal{T} \\ p_i = p_i(W_i^+, t_i)\,;\ 0 \le p_i \le 1 \\ c_i = c_i(W_i^- + q_i, t_i)\,;\ c_i \ge 0 \end{cases} . \tag{9.1}
$$

We can interpret problem (9.1) as minimizing the quadratic penalty for shortfall with respect to the target W^*. As in Section 7, we allow surplus cash withdrawals over and above the scheduled injections/withdrawals q_i. An optimal strategy is to withdraw

$$
c_i = \max\left[W_i^- + q_i - \left(W^* e^{-r(T-t_i)} - Q_i\right), 0\right] \tag{9.2}
$$

from the portfolio and invest the remainder in the bond index (Dang and Forsyth, 2016). Recall that Q_i is defined in equation (7.2). In addition, the following result due to Zhou and Li (2000) implies that problem (9.1) simultaneously minimizes two measures of risk: expected quadratic shortfall and variance.

Proposition 9.1 (Dynamic mean variance efficiency)**.** *The solution to problem* (9.1) *is multi-period mean variance optimal.*

Remark 9.1 (Time consistency)**.** *There is considerable confusion in the literature about pre-commitment mean-variance strategies. These strategies are commonly criticized for being time inconsistent (Basak and Chabakauri, 2010; Björk et al., 2014). However, the pre-commitment optimal policy can be found by solving problem* (9.1) *using dynamic programming with a fixed* W^**, which is clearly time consistent. Hence, when determining the time consistent optimal strategy for problem* (9.1)*, we obtain the optimal mean variance pre-commitment solution as a by-product. Vigna (2017) and Menoncin and Vigna (2017) provide further insight into this. As noted by Cong and Oosterlee (2016), the pre-commitment strategy can be seen as a strategy consistent with a fixed investment target, but not with a risk aversion attitude. Conversely, a time consistent strategy has a consistent risk aversion attitude, but it is not consistent with respect to an investment target. We contend that consistency with a target is more useful for life cycle investment strategies.*

We determine W^* in problem (9.1) by enforcing the constraint

$$
E[W_T] = W^{spec}, \tag{9.3}
$$

where W^{spec} is the desired expected value of W_T. Computationally, we do this by embedding problem (9.1) in a Newton iteration where we solve the equation $(E[W_T] - W^{spec}) = 0$ for W^*. Note that adjusting W^{spec} allows us to indirectly adjust $Median[W_T]$. We choose $W^{spec} = 1000$.[11] This gives a value of $W^* = 1123$. This choice gives an average allocation to the stock index of

[11] Recall that units are thousands of dollars, so this corresponds to real terminal wealth of $1,000,000.

$\hat{b}$	$Median[W_T]$	$Mean[W_T]$	$std[W_T]$	$Pr[W_T < 0]$	5% CVAR
Synthetic market					
NA	1123	1032	354	.042	−377
Historical market					
0.5	1144	1096	344	.041	−345
1.0	1155	1134	334	.038	−311
2.0	1169	1198	290	.026	−112
5.0	1200	1280	234	.015	+154

TABLE 9.1: *Optimal control determined by solving problem* (9.1) *(quadratic shortfall) with* $E[W_T] =$ 1000 *(excluding surplus cash) in the synthetic market, assuming the scenario in Table 4.1.* W_T *denotes real terminal wealth after* 60 *years, measured in thousands of dollars. Statistics for the synthetic market case are based on* 6.4×10^5 *Monte Carlo simulations. Statistics for the historical market cases are based on* 10,000 *stationary block bootstrap resamples of the historical data from 1926:1 to 2016:12.* $\hat{b}$ *is the expected blocksize, measured in years. Surplus cash is included in the mean, median, CVAR, and probability of ruin, but excluded from the standard deviation.*

about 0.42.[12] Note that in this case, the mean and median of the terminal wealth are comparable hence it is of interest to specify the mean value of the terminal wealth. It results in a median final wealth that is roughly comparable in the synthetic market to that seen earlier in Table 5.1 for the case with a constant equity weight of $p = 0.4$.[13]

Table 9.1 presents the results. Note that the constraint in equation (9.3) is the mean without surplus cash, while the means reported in this table include surplus cash. However, the average value of surplus cash is not very large ($1032 - 1000$ in the synthetic market). Unlike for the previous objective functions considered, in this quadratic shortfall case the results in the historical market are generally superior to those in the synthetic market.

As a point of reference, if we determine the optimal strategy with the constraint that $E[W_T] = 500$, then, in the synthetic market, we find that $Median[W_T] = 556$ and $Pr[W_T < 0] = .035$, compared with $Median[W_T] = 1123$ and $Pr[W_T < 0] = .042$ for the $E[W_T] = 1000$ case. This is a large reduction in median terminal wealth for a fairly small improvement in probability of ruin. This effect of a relatively modest reduction in probability of ruin at a cost of a steep reduction in median terminal wealth is representative of many test cases we have run, using a wide range of parameters.

Figure 9.1 shows the percentiles of the wealth (panel (a)) and the fraction invested in stocks (panel (b)) for the historical market with expected blocksize $\hat{b} = 2.0$ years. In Figure 9.1(a), the 5th percentile represents a very poor outcome. However, in this case there is still a reasonably large buffer of remaining wealth at the end of 60 years. Figure 9.1(b) shows that the optimal strategy for this quadratic shortfall objective starts out with 100% invested in the equity index over the first several years. If market returns are very favourable during that period, there will be a sharp fall in the equity fraction (e.g. the 5th percentile case), to the point of possibly being completely

[12]This is the time average of the median value of the equity weight p.

[13]We experimented with other ways of specifying W^*. For example, rather than using the value which resulted in $E[W_T] = 1000$, we determined the value which minimized $Pr[W_T < 0]$. Although this looked promising in the synthetic market, its performance in the historical market tests was worse compared to the strategy which set $E[W_T] = 1000$.

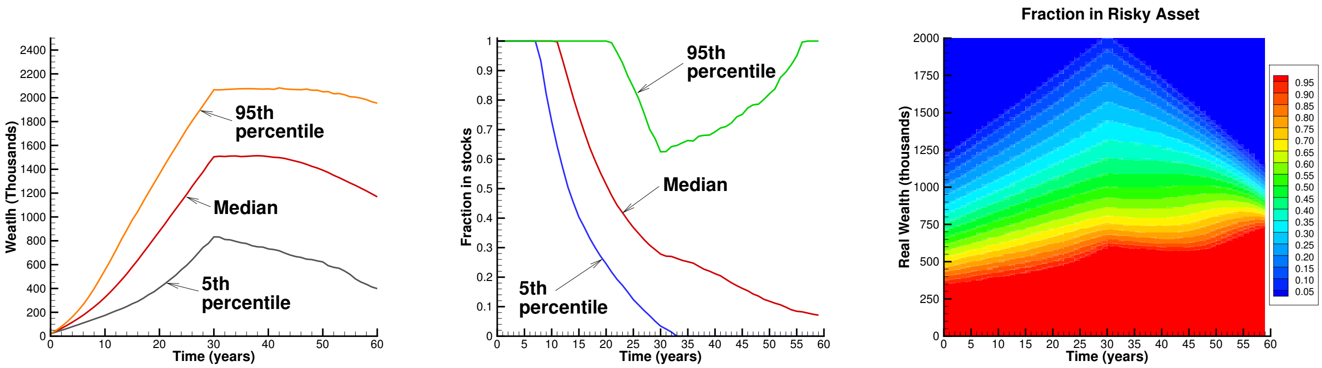

(a) *Percentiles of accumulated wealth.*

(b) *Percentiles of optimal fraction in equities.*

(c) *Optimal control heat map.*

FIGURE 9.1: *Percentiles of real wealth and the optimal fraction invested in equities. Optimal control computed by solving the quadratic shortfall problem* (9.1) *with the constraint that* $E[W_T] = 1000$ *in the synthetic market, assuming the scenario in Table 4.1. Statistics based on* 10,000 *stationary block bootstrap resamples of the historical data from 1926:1 to 2016:12. Expected blocksize* $\hat{b} = 2$ *years.*

de-risked for the last 25 years of the 60 year horizon. In this case, the $5th$ percentile represents a favourable investment outcome. The median case illustrates the same de-risking, but to a lesser extent (approximately 10% invested in the equity index over the last decade). On the other hand, the 95th percentile maintains the initial 100% allocation to equities for much longer, starts to de-risk, but then turns around with an increasing allocation to equities over approximately the last 25 years. It appears that withdrawals coupled with poor returns require higher equity exposures in order to reach the target. Figure 9.1(c) shows the heat map of the optimal quadratic shortfall strategy.

Unlike most other strategies, the quadratic shortfall strategy produces results in the historical market which are generally as good or better than in the synthetic market. The significant terminal wealth target adds robustness, ameliorating the effect of stochastic interest rates (which are introduced in the bootstrap resampling tests). In addition, as shown in Figure 9.1(c), the basic strategy is heavily contrarian: when wealth is low (e.g. due to equity market drops) invest more in equities, and then capture gains (by moving cash to bonds) when wealth is high (e.g. due to favourable equity market returns). Historically, this has been a winning strategy: over very long periods the market does move up, but retirees do not enjoy an unlimited time horizon. While there will always be a risk that equities do not achieve the growth target, the quadratic shortfall strategy greatly reduces risk by recognizing the opportunity of shifting the allocation to bonds when the target is within sight.

Overall, it seems that these strategies, which can be interpreted as minimizing the expected quadratic shortfall with respect to a target, with an expected value constraint, are fairly robust. The ruin probabilities in the historical market are $Pr[W_T < 0] \simeq .03$ ($\hat{b} = 2$), which may be acceptable in practice. Recall that in the synthetic market, the best possible strategy gives $Pr[W_T < 0] \simeq .02$. The quadratic shortfall strategies give up only a small amount in terms of probability of failure.[14] In return we have a good chance of a large bequest (or a safety buffer for longevity), i.e. $Median[W_t] > 1{,}000$.

[14] Recall that the optimal strategy for minimizing $Pr[W_T < 0]$ was not very robust in terms of the bootstrap stress tests.

10 Some alternative strategies

We now briefly discuss some other strategies which we have considered. First, we have tested strategies where we replace the objective function in the quadratic shortfall problem (9.1) by

$$\min_{\{(p_0,c_0),\ldots,(p_{M-1},c_{M-1})\}} E\left[|(\min(W_T - W^*,0)|^{\beta} \right], \tag{10.1}$$

for powers $\beta \in \{1,3,4\}$, in addition to the $\beta = 2$ case considered in detail in Section 9. Similar results were obtained for all choices of β, with $\beta = 2$ having a slight edge.

Another target-based objective function has been recently suggested in Zhang et al. (2017). This is the *sharp target* objective. It seeks to maximize expected terminal wealth over a specified target range, where the upper end of the range corresponds to a wealth goal and the lower end represents a desired conservative minimum. We give a brief overview of our results using this objective function in Appendix D. This objective function produced results similar to the quadratic shortfall criteria, but with noticeably worse CVAR. Hence, it appears that the quadratic shortfall (expected value constraint) objective function discussed in Section 9 gives somewhat better overall results.

11 Comparison of strategies

Table 11.1 compares representative results for several of the strategies discussed earlier.[15] This comparison is in the historical market, with an expected blocksize of $\hat{b} = 2$ years. The focus is on the two risk measures which we view as most important in this context: ruin probability and 5% CVAR. A secondary criterion is the median terminal wealth (since mean terminal wealth can be misleading due to a small number of simulated paths with extreme results). Table 11.1 shows that in terms of minimizing risk, the quadratic shortfall objective function with an expected value constraint from Section 9 seems to be superior to the other objective functions. It also offers a relatively high median terminal wealth. It is outperformed significantly on this dimension by the constant equity fraction strategies with $p = 0.60$ and $p = 0.80$, but these constant weight strategies also have much higher risk exposures.

Figure 11.1 plots kernel-smoothed probability densities of terminal wealth W_T in the historical market for the three constant weight strategies and the quadratic shortfall strategy from Table 11.1. This figure highlights some of the differences between the simpler constant weight approaches and the quadratic shortfall strategy. This latter strategy clearly sacrifices a lot of upside potential in exchange for downside protection, concentrating the wealth distribution in a narrow range, compared to the constant weight cases.

If we are concerned that too much upside is sacrificed for the quadratic shortfall method, we can try using a higher expected value constraint. Suppose, for example, that we target $E[W_T] = 2500$ in the synthetic market. Then in the historical bootstrap market ($\hat{b} = 2$ years), we obtain $Median[W_T] = 2961$, which is approximately the median obtained for the constant weight $p = .6$ case in Table 11.1. The quadratic shortfall risk measures in this case are $Pr[W_T < 0] = .04$, and 5% CVAR$= -331$. These results are still superior to the constant weight $p = .6$ case, but the quadratic shortfall strategy has to maintain a relatively high allocation to equities in order to hit the expected value target, so that there is less freedom to reduce risk.

As an additional stress test, we consider a case where the optimal strategy was computed with the historical parameters, but, going forward, the stock returns are reduced by 200 basis points per

[15] Table 11.1 excludes some strategies which performed relatively poorly, such as minimizing the probability of ruin with $W^{min} = 0$ and the mean-CVAR strategy with $\kappa = -10^{-8}$.

Strategy	$Median[W_T]$	$Pr[W_T < 0]$	5% CVAR
Const. equity fraction $p = .40$	961	.15	−461
Const. equity fraction $p = .60$	2931	.07	−389
Const. equity fraction $p = .80$	6151	.054	−411
Minimize probability of ruin (Section 7) $\min Pr\left[W_T < W^{min}\right]$; $W^{min} = 200$	456	.030	−183
Mean-CVAR (Section 8) $\max \text{CVAR}_\alpha + \kappa E[W_T]$; $\alpha = .05$; $\kappa = +10^{-8}$	320	.036	−184
Sharp target (Appendix D) $W_L = 100, W_U = 1178$	1138	.031	−204
Quadratic shortfall (Section 9) $\min E\left[(\min(W_T - W^*, 0))^2\right]$: W^* : $E[W_T] = 1000$	1169	.026	−112

TABLE 11.1: *Optimal controls determined by solving for strategies in the synthetic market, assuming the scenario in Table 4.1. Reported results use these controls in the historical market and are based on* 10,000 *stationary block bootstrap resamples of the historical data from 1926:1 to 2016:12 with expected blocksize $\hat{b}$ = 2 years. W_T denotes real terminal wealth after* 60 *years, measured in thousands of dollars. Surplus cash is included in the median terminal wealth, where applicable.*

year relative to the historical average. Obviously all strategies in this case are adversely affected, but the quadratic shortfall strategy computed using incorrect parameter estimates is still superior to the constant weight strategies.

12 Conclusion

DC pension plan holders generally have no choice but to invest in risky assets in order to achieve even minimal salary replacement levels. We make the conservative assumption that the DC plan holder requires fixed cash flows for 30 years after retirement, after an accumulation period of 30 years. We also assume that the holder does not choose to annuitize, which is consistent with observed behaviour.

Our main result is that an objective function which focuses purely on a risk measure such as minimizing the probability of ruin or maximizing CVAR[16] performs well in a synthetic market, but poorly in bootstrap backtests (the historical market). The main problem seems to be that these strategies are not robust due to real interest rate shocks introduced by the resampling process.

In addition, even in the synthetic market, we observe that the small decreases in the probability of ruin come at the cost of drastically reducing the median terminal wealth (i.e. a bequest or an additional longevity safety valve). Greater robustness is achieved by targeting a final wealth greater than zero, which acts as a buffer against uncertainties in market parameters. Another side effect of this is that a significant terminal wealth acts as additional buffer for longevity risk (e.g. the risk of living for more than 30 years of retirement).

Minimizing the quadratic shortfall with an expected terminal wealth constraint appears to be a good strategy in general, as long as the expected terminal wealth constraint is sufficiently large

[16]Recall that we define CVAR as the mean of the worst α fraction of terminal wealth, not the losses, so we want to maximize CVAR to minimize risk.

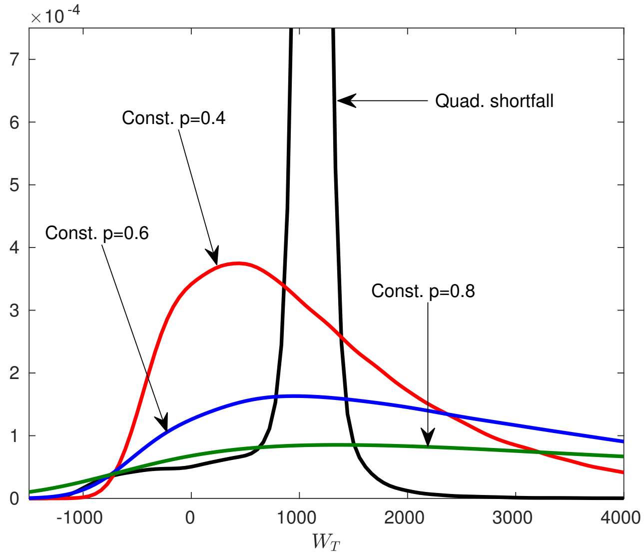

FIGURE 11.1: *Kernel smoothed probability densities for three constant weight strategies and the quadratic shortfall strategy, assuming the scenario in Table 4.1. Densities based on stationary block bootstrap resampling of the historical data from 1926:1 to 2016:12 with an expected blocksize of* $\hat{b} = 2.0$ *years.* W_T *denotes real terminal wealth after* 60 *years, measured in thousands of dollars. The quadratic shortfall method enforces the constraint that* $E[W_T] = 1000$ *in the synthetic market used to determine the optimal control for that strategy.*

to buffer the real interest rate shocks. This method results in an acceptable probability of ruin, and a significant median terminal wealth. This strategy is also robust to the misspecification of the drift of the risky asset, and is superior (by most measures) to standard constant weight strategies. However, this approach requires some experimentation in order to set the expected terminal wealth constraint appropriately.

It is interesting to observe that a robust strategy involves aiming for a significant size of terminal wealth (which may turn out to be a bequest) in order to have a small probability of ruin. In this instance, the investor and her heirs are likely to agree on the strategy.

We would like to emphasize that it is important to stress test any strategy, e.g. by bootstrapping the historical data. Some strategies which appear to work very well in the synthetic market fail in the bootstrap stress tests. However, we believe that our tests point the way to some promising choices of objective function for full life cycle DC plan asset allocation.

Any strategy which involves investing in risky assets to meet fixed cash flows has a non-zero probability of portfolio depletion before the horizon date. The best that can be done is to make this probability acceptably small. Nevertheless, failure can occur, which begs the question of what happens then. A possible backup in many cases would be the use of the retiree's other assets, such as real estate. For example, it may be possible to use a reverse mortgage to monetize the retiree's home. As long as the value of any real estate asset is larger than (the negative of) the 5% CVAR, then we can regard the real estate asset as at least a partial hedge against the tail risk.

Our basic question in this work was whether a suitably chosen investment strategy would offer a DC plan member the opportunity to have a similar retirement income stream as provided by a traditional DB plan. The quadratic shortfall strategy produces a 30-year real annuity with a low probability of ruin, not a guaranteed life annuity (assuming DB pension plan solvency). In this respect, it falls a bit short of providing a fully comparable retirement income stream. Offsetting this, however, is a reasonably good chance of a large buffer, which could be used to pay for higher than anticipated expenses, or as a significant bequest, or as a hedge against extreme longevity. The quadratic shortfall strategy can also be regarded as being superior to annuitization, since it preserves liquidity and is defined real terms, whereas in practice annuities are almost invariably defined in nominal terms, and often considerably overpriced compared to their actuarial value (MacDonald et al., 2013).

We have restricted attention in this paper to requiring a fixed (real) withdrawal during the decumulation phase. Another alternative is to allow the withdrawal to vary in response to the then current portfolio value, based on an estimate of remaining lifetime. This shifts the risk from portfolio depletion to volatile decumulation cash flows (Waring and Siegel, 2015; Westmacott and Daley, 2015). In this case, the control problem objective function would be to minimize the withdrawal volatility and maximize the cumulative withdrawals. We intend to study this approach in the future.

Acknowledgements

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

Conflicts of interest

The authors have no conflicts of interest to report.

Appendices

A Calibration of model parameters

To estimate the jump diffusion model parameters, we use the thresholding technique described in Mancini (2009) and Cont and Mancini (2011). This procedure is considered to be relatively efficient for fairly low frequency data, such as the monthly frequency used here. For details, see Dang and Forsyth (2016) and Forsyth and Vetzal (2017a). We use a threshold parameter $\alpha = 3$ in our estimates.[17]

Table A.1 provides the resulting annualized parameter estimates for the double exponential jump diffusion given in equation (2.2). The drift rate μ corresponds to an expected annual return of almost 9%. The diffusive volatility σ might seem slightly low at less than 15%, but recall that the overall effective volatility includes this amount plus the contribution to volatility from jumps. The jump intensity λ implies that jumps can be expected to occur approximately every 3 years. When a jump happens, it is about 3 times more likely to be a move down than a move up. Upward jumps are a little larger on average than downward jumps.

Figure A.1 shows the normalized histogram of real CRSP Deciles (1-10) index log returns for the period 1926:1-2016:12. The standard normal density and scaled jump diffusion density are also shown. The improved fit from the jump diffusion model is readily apparent.

[17]This parameter has the intuitive interpretation that if the absolute value of the log return in a period is larger an α standard deviation Brownian motion return, then it is identified as a jump.

μ	σ	λ	p_{up}	η_1	η_2
.08753	.14801	.34065	.25806	4.67877	5.60389

TABLE A.1: *Estimated annualized parameters for the double exponential jump diffusion model given in equation* (2.2) *applied to the value-weighted CRSP Deciles (1-10) index, deflated by the CPI. Sample period 1926:1 to 2016:12. CRSP and CPI data from Dimensional Returns 2.0 under licence from Dimensional Fund Advisors Canada.*

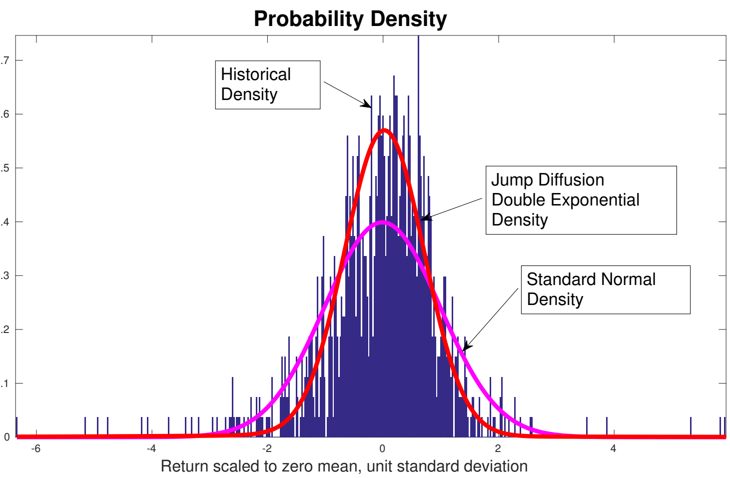

FIGURE A.1: *Histogram of real log returns of CRSP Deciles (1-10) index, scaled to zero mean and unit standard deviation. Standard normal density shown, as well as the fitted jump diffusion, double exponential distribution, also scaled to zero mean and unit standard deviation. Jump diffusion parameters: threshold ($\alpha = 3$) from Table A.1. Sample period 1926:1 to 2016:12. Source: CRSP and CPI data from Dimensional Returns 2.0 under licence from Dimensional Fund Advisors Canada.*

The historical average annualized real interest rate for one-month US T-bills from 1926:1 to 2016:12 was $r = 0.004835$. The volatility of the one-month T-bill return was about .018, which justifies ignoring the randomness of short term interest rates, at least as a first approximation. We test the effect of this assumption on optimal strategies by applying the computed strategies to the historical market, which is constructed using bootstrap resamples of the data series and so includes the effect of stochastic real interest rates.

B Bootstrap resampling

We use bootstrap resampling to study how the various strategies would have performed on actual historical data. A single bootstrap resampled path can be constructed as follows. Divide the total investment horizon of T years into k blocks of size b years, so that $T = kb$. We then select k blocks at random (with replacement) from the historical data (from both the deflated equity and T-bill indexes). Each block starts at a random month. A single path is formed by concatenating these blocks. The historical data is wrapped around to avoid end effects, as in the circular block bootstrap (Politis and White, 2004; Patton et al., 2009). This procedure is then repeated for many paths.

The sampling is done in blocks in order to account for possible serial dependence effects in the historical time series. The choice of blocksize is crucial and can have a large impact on the results (Cogneau and Zakalmouline, 2013). We simultaneously sample the real stock and bond returns from the historical data. This introduces random real interest rates in our samples, in contrast to the constant interest rates assumed in the synthetic market tests and in the determination of the optimal controls.

To reduce the impact of a fixed blocksize and to mitigate the edge effects at each block end, we use the stationary block bootstrap (Politis and White, 2004). The blocksize is randomly sampled from a geometric distribution with an expected blocksize $\hat{b}$. The optimal choice for $\hat{b}$ is determined using the algorithm described in Patton et al. (2009).[18] Calculated optimal values for $\hat{b}$ were 57 months for the T-bill index and 3.5 months for the real CRSP index. We adopt a paired sampling approach whereby we sample simultaneously from both stock and bond indexes, so we must use the same blocksize for both indexes. Since the recommended blocksizes are quite different for the two indexes, we sidestep this issue by presenting results for a range of blocksizes.

C Definition of CVAR

Let $p(W_T)$ be the probability density function of wealth at $t = T$. Let

$$\int_{-\infty}^{W_\alpha^*} p(W_T)\, dW_T = \alpha, \tag{C.1}$$

i.e. $Pr[W_T > W_\alpha^*] = 1 - \alpha$. We can interpret W_α^* as the Value at Risk (VAR) at level α. The Conditional Value at Risk (CVAR) at level α is then

$$\text{CVAR}_\alpha = \frac{\int_{-\infty}^{W_\alpha^*} W_T\, p(W_T)\, dW_T}{\alpha}, \tag{C.2}$$

which is the average of the worst α fraction of outcomes. Typically $\alpha = .01, .05$. Note that the definition of CVAR in equation (C.2) uses the probability density of the final wealth distribution, not the density of *loss*. Hence, in our case, a larger value of CVAR (i.e. a larger value of worst case terminal wealth) is desired. In our examples, we have both positive and negative values of CVAR.

Given an expectation under control $\mathcal{P}$, $E^{\mathcal{P}}[\cdot]$, as noted by Rockafellar and Uryasev (2000) and Miller and Yang (2017), the mean-CVAR optimization can be expressed as

$$\max_{\mathcal{P}} \sup_{W^*} E^{\mathcal{P}}\left(W^* + \frac{1}{\alpha}\left[(W_T - W^*)^- \right] + \kappa W_T \right). \tag{C.3}$$

Following Miller and Yang (2017), we interchange the max and sup operations in equation (C.3), which allows us to rewrite the objective function (C.3) as

$$\sup_{W^*} \left\{ \max_{\mathcal{P}} E^{\mathcal{P}}\left(W^* + \frac{1}{\alpha}\left[(W_T - W^*)^- \right] + \kappa W_T \right) \right\}, \tag{C.4}$$

and solve the inner optimization problem using an HJB equation (Dang and Forsyth, 2014; Forsyth and Labahn, 2018). Standard methods can then be used to solve the outer optimization problem.

Remark C.1 (Time consistency of mean-CVAR strategies)**.** *Suppose that we solve the mean-CVAR problem at $t = 0$, for a given confidence level α. This determines a value of W^* in equation (C.4).*

[18]This approach has also been used in other tests of portfolio allocation problems recently (e.g. Dichtl et al., 2016).

If we fix the value of W^*, then the pre-commitment mean-CVAR strategy (computed at $t = 0$), is the time consistent solution for the objective function

$$\left\{ \max_{\mathcal{P}} E^{\mathcal{P}} \left((W_T - W^*)^- + \kappa' W_T \right) \right\}$$
$$\kappa' = \alpha\kappa \; ; \; \alpha > 0 \; ; \; W^* = \textit{ fixed} \; , \tag{C.5}$$

for all $t > 0$. Hence the pre-commitment mean-CVAR solution is time consistent in terms of linear shortfall with respect to a fixed target W^*. Alternatively, the pre-commitment mean-CVAR policy can also be seen as a time consistent mean-CVAR strategy if we allow time dependent confidence level α and wealth dependent expected wealth target (Strub et al., 2017).

D Sharp target

Another possible objective function is the *sharp target* suggested in Zhang et al. (2017):

$$\max_{\{(p_0,c_0),\ldots,(p_{M-1},c_{M-1})\}} E\left[(W_T - W_L)\mathbf{1}_{W_L \le W_T < W_U}\right]$$
$$\text{subject to} \begin{cases} (S_t, B_t) \text{ follow processes (2.1)-(2.4)}; \quad t \notin \mathcal{T} \\ W_i^+ = W_i^- + q_i - c_i \, ; \; S_i^+ = p_i W_i^+ \, ; \; B_i^+ = W_i^+ - S_i^+ \, ; \quad t \in \mathcal{T} \\ p_i = p_i(W_i^+, t_i) \, ; \; 0 \le p_i \le 1 \\ c_i = c_i(W_i^- + q_i, t_i) \, ; \; c_i \ge 0 \end{cases} , \tag{D.1}$$

where W_L, W_U are parameters. We can think of W_L as a minimum required value of the final wealth and W_U as the desired value. We withdraw cash from the portfolio if investing the remaining amount in the risk-free asset (along any given path) ensures that $W_T > W_U$. The surplus (withdrawn amount) is also invested in the risk-free asset. Note that we have to specify what rule to use if a risk-free investment results in $W_T > W_U$, since otherwise the problem is not fully specified.

The idea of objective (D.1) is to reward outcomes between $W_L < W_T < W_U$, with higher reward for outcomes near W_U. There is no reward for outcomes $W_T > W_U$. A possible problem is that all outcomes $W_T < W_L$ are penalized equally. To be comparable with the results in Section 9 (quadratic shortfall with expected value constraint), we fix $W_L = 100$ and determine W_U so that $E[W_T] = 1000$ in the synthetic market. This gives $W_U = 1178$.

The results for the sharp target strategy are shown in Table D.1. Comparing the historical market results from this table with those for the quadratic shortfall strategy in Table 9.1, we see that the sharp target gives similar results, except that the 5% CVAR is notably worse. This can be traced the the fact that all shortfalls below W_L are weighted equally in the sharp target objective, while larger shortfalls are increasingly penalized with the quadratic shortfall objective function.

References

Arnott, R. D., K. F. Sherrerd, and L. Wu (2013). The glidepath illusion and potential solutions. *The Journal of Retirement 1*(2), 13–28.

Basak, S. and G. Chabakauri (2010). Dynamic mean-variance asset allocation. *Review of Financial Studies 23*, 2970–3016.

Björk, T., A. Murgoci, and X. Y. Zhou (2014). Mean variance portfolio optimization with state dependent risk aversion. *Mathematical Finance 24*, 1–24.

$\hat{b}$	$Median[W_T]$	$Mean[W_T]$	$std[W_T]$	$Pr[W_T < 0]$	5% CVAR
			Synthetic market		
NA	1102	1023	372	.047	−507
			Historical market		
0.5	1108	1061	365	.045	−449
1.0	1120	1100	357	.042	−413
2.0	1138	1171	315	.031	−204
5.0	1177	1265	251	.018	+156

TABLE D.1: *Optimal control determined by solving problem* (D.1) *with* $W_L = 100$ *and* $W_U = 1178$*, so that* $E[W_T] = 1000$ *(excluding surplus cash) in the synthetic market, assuming the scenario in Table 4.1.* W_T *denotes real terminal wealth after* 60 *years, measured in thousands of dollars. Statistics for the synthetic market case are based on* 6.4×10^5 *Monte Carlo simulation runs. Statistics for the historical market cases are based on* 10,000 *stationary block bootstrap resamples of the historical data from 1926:1 to 2016:12.* $\hat{b}$ *is the expected blocksize, measured in years. Surplus cash is included in the mean, median, CVAR, and probability of ruin, but excluded from the standard deviation.*

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.

Cairns, A. J. G., D. Blake, and K. Dowd (2006). Stochastic lifestyling: Optimal dynamic asset allocation for defined contribution pension plans. *Journal of Economic Dynamics and Control 30*, 843–977.

Campanele, C., C. Fugazza, and F. Gomes (2015). Life-cycle portfolio choice with liquid and illiquid financial assets. *Journal of Monetary Economics 71*, 67–83.

Chen, A. and L. Delong (2015). Optimal investment for a defined-contribution pension scheme under a regime switching model. *Astin Bulletin 45*, 397–419.

Christiansen, M. C. and M. Steffensen (2018). Around the life cycle: Deterministic consumption-investment strategies. *North American Actuarial Journal*, forthcoming.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*, 491–533.

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

Cong, F. and C. W. Oosterlee (2016). Multi-period mean variance portfolio optimization based on Monte-Carlo simulation. *Journal of Economic Dynamics and Control 64*, 23–38.

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

Dahlquist, M., O. Setty, and R. Vestman (2018). On the asset allocation of a default pension fund. *Journal of Finance 73*, 1893–1936.

Dammon, R. M., C. S. Spatt, and H. H. Zhang (2004). Optimal asset location and allocation with taxable and tax-deferred investing. *Journal of Finance 59*, 999–1037.

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.

Dang, D.-M., P. A. Forsyth, and K. R. Vetzal (2017). The 4% strategy revisited: a pre-commitment optimal mean-variance approach to wealth management. *Quantitative Finance 17*, 335–351.

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

Donnelly, C., R. Gerrard, M. Guillén, and J. P. Nielsen (2015). Less is more: Increasing retirement gains by using an upside terminal wealth constraint. *Insurance: Mathematics and Economics 64*, 259–267.

Donnelly, C., M. Guillén, J. P. Nielsen, and A. M. Pérez-Marín (2017). Implementing individual savings decisions for retirement with bounds on wealth. *Astin Bulletin 48*, 111–137.

Fagereng, A., C. Gottlieb, and L. Guiso (2017). Asset market participation and portfolio choice over the life-cycle. *Journal of Finance 72*, 705–750.

Forsyth, P. and G. Labahn (2018). ϵ-monotone Fourier methods for optimal stochastic control in finance. to appear, Journal of Computational Finance.

Forsyth, P. A. and K. R. Vetzal (2017a). Dynamic mean variance asset allocation: tests for robustness. *International Journal of Financial Engineering 4*(2). 1750021 (electronic), `https://doi.org/10.1142/S2424786317500219`.

Forsyth, P. A. and K. R. Vetzal (2017b). Optimal asset allocation for retirement savings: deterministic vs. adaptive strategies. Working paper, University of Waterloo.

Forsyth, P. A. and K. R. Vetzal (2017c). Robust asset allocation for long-term target-based investing. *International Journal of Theoretical and Applied Finance 20*. 1750017 (electronic), `https://doi.org/10.1142/S0219024917500170`.

Forsyth, P. A., K. R. Vetzal, and G. Westmacott (2017). Target wealth: the evolution of target date funds. White paper, PWL Capital, available from `www.pwlcapital.com/en/The-Firm/White-Papers`.

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

Gao, J., K. Zhou, D. Li, and X. Cao (2017). Dynamic mean-LPM and mean-CVaR portfolio optimization in continuous-time. *SIAM Journal on Control and Optimization 55*, 1377–1397.

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.

Graf, S. (2017). Life-cycle funds: Much ado about nothing? *European Journal of Finance 23*, 974–998.

Guan, G. and Z. Liang (2014). Optimal management of DC pension plan in a stochastic interest rate and stochastic volatility framework. *Insurance: Mathematics and Economics 57*, 58–66.

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.

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

Li, D. and W.-L. Ng (2000). Optimal dynamic portfolio selection: multiperiod mean-variance formulation. *Mathematical Finance 10*, 387–406.

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

Ma, K. and P. A. Forsyth (2016). Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation under stochastic volatility. *Journal of Computational Finance 20*(1), 1–37.

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.

Michaelides, A. and Y. Zhang (2017). Stock market mean reversion and portfolio choice over the life cycle. *Journal of Financial and Quantitative Analysis 52*, 1183–1209.

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

Miller, C. and I. Yang (2017). Optimal control of conditional value-at-risk in continuous time. *SIAM Journal on Control and Optimization 55*, 856–884.

O'Hara, M. and T. Daverman (2017). Reexamining "to versus through". White Paper, BlackRock, www.blackrock.com/investing/literature/whitepaper/to-versus-through-whitepaper.pdf.

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. M. 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.

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.

Rupert, P. and G. Zanella (2015). Revisiting wage, earnings, and hours profiles. *Journal of Monetary Economics 72*, 114–130.

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

Strub, M., D. Li, X. Cui, and J. Gao (2017). Discrete-time Mean-CVaR portfolio selection and time-consistency induced term structure of the CVaR. SSRN Abstract: 3040517.

Tretiakova, I. and M. S. Yamada (2011). What DC plan members really want. *Rotman International Journal of Pension Management 4*(2), 60–70.

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

Vigna, E. (2017). Tail optimality and preferences consistency for intertemporal optimization problems. Working paper, Collegio Carlo Alberto, Università Degli Studi di Torino.

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

Westmacott, G. and S. Daley (2015). The design and depletion of retirement portfolios. White paper, PWL Capital, available from `www.pwlcapital.com/en/The-Firm/White-Papers`.

Wu, H. and Y. Zeng (2015). Equilibrium investment strategy for defined-contribution pension schemes with generalized mean-variance criterion and mortality risk. *Insurance: Mathematics and Economics 64*, 396–408.

Yao, H., Y. Lai, Q. Ma, and M. Jian (2014). Asset allocation for a DC pension fund with stochastic income and mortality risk: A multi-period mean-variance framework. *Insurance: Mathematics and Economics 54*, 84–92.

Zhang, R., N. Langrené, Y. Tian, Z. Zhu, F. Klebaner, and K. Hamza (2017). Sharp target range strategy with dynamic portfolio selection: decensored least squares Monte Carlo. Working paper, Monash University.

Zhou, X. Y. and D. Li (2000). Continuous-time mean-variance portfolio selection: a stochastic LQ framework. *Applied Mathematics and Optimization 42*, 19–33.