The 4% Strategy Revisited: A Pre-Commitment Mean-Variance Optimal Approach to Wealth Management*

Duy-Minh Dang[†] P.A. Forsyth[‡] K.R. Vetzal[§]

April 2, 2016

Abstract

In contrast to single-period mean-variance (MV) portfolio allocation, multi-period MV optimal portfolio allocation can be modified slightly to be effectively a down-side risk measure. With this in mind, we consider multi-period MV optimal portfolio allocation in the presence of periodic withdrawals. The investment portfolio can be allocated between a risk-free investment and a risky asset, the price of which is assumed to follow a jump diffusion process. We consider two wealth management applications: optimal de-accumulation rates for a defined contribution pension plan and sustainable withdrawal rates for an endowment. Several numerical illustrations are provided, with some interesting implications. In the pension de-accumulation context, Bengen (1994)'s historical analysis indicated that a retiree could safely withdraw 4% of her initial retirement savings annually (in real terms), provided that her portfolio maintained an even balance between diversified equities and U.S. Treasury bonds. Our analysis does support 4% as a sustainable withdrawal rate in the pension de-accumulation context (and a somewhat lower rate for an endowment), but only if the investor follows an MV optimal portfolio allocation, not a fixed proportion strategy. Compared with a constant proportion strategy, the MV optimal policy achieves the same expected wealth at the end of the investment horizon, while significantly reducing the standard deviation of wealth and the probability of shortfall. We also explore the effects of suppressing jumps so as to have a pure diffusion process, but assuming a correspondingly larger volatility for the latter process. Surprisingly, it turns out that the MV optimal strategy is more effective when there are large downward jumps compared to having a high volatility diffusion process. Finally, tests based on historical data demonstrate that the MV optimal policy is quite robust to uncertainty about parameter estimates.

Keywords: multi-period mean-variance optimal, asset allocation strategy, sustainable withdrawal rate, endowment, pension de-accumulation

AMS Classification: 65N06, 93C20
JEL Classification: G11

*This work was supported by the Bank of Nova Scotia, the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Social Sciences and Humanities Research Council of Canada (SSHRC).

[†]School of Mathematics and Physics, The University of Queensland, Brisbane, QLD 4072, Australia, duyminh.dang@uq.edu.au

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

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

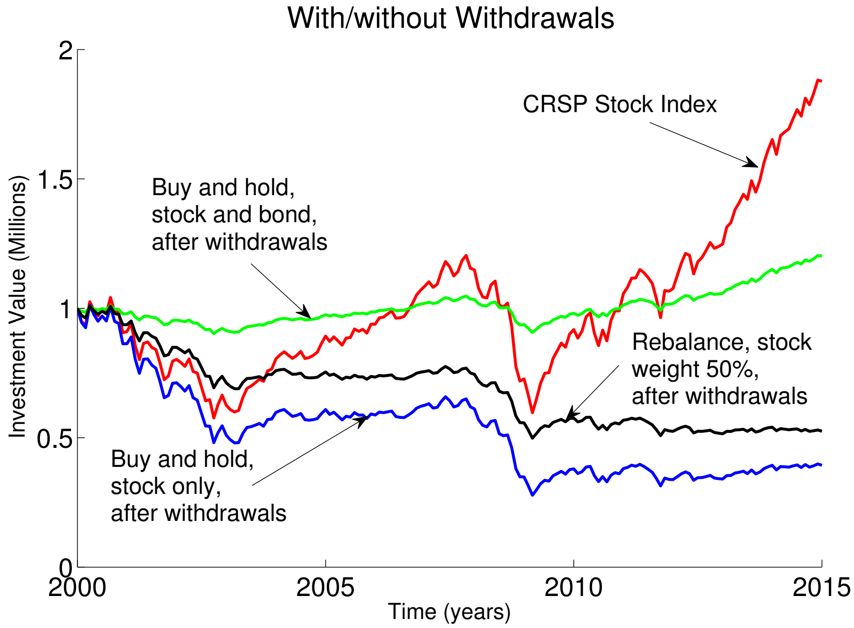

FIGURE 1.1: *Investor withdraws $50,000 per year, initial investment $1 million. CRSP Stock Index: value of investment with entire capital invested in equity index and without withdrawals. Buy and hold, stock only, after withdrawals: entire capital invested in equity index, with regular withdrawals. Buy and hold, stock and bond, after withdrawals: capital invested in U.S. Treasuries and equity index, withdrawals financed from Treasuries with regular withdrawals. Rebalance, stock weight 50%, after withdrawals: 50% stock and 50% short-term U.S. Treasuries, rebalanced monthly, with regular withdrawals.*

1 Introduction

As a motivational illustration, consider a hypothetical investor who retired on January 1, 2000 with retirement savings of U.S. $1,000,000. Knowing that the long-term returns of the stock market have generally exceeded those of government bonds, suppose that this investor invested all of these savings in a broadly diversified U.S. stock market index (a buy and hold strategy). Also suppose that this retiree made monthly withdrawals totalling $50,000 per year for living expenses. In order to examine the performance of this investor's portfolio, we use historical data for a total return stock index from the Center for Research in Security Prices (CRSP, see `www.crsp.com`). The CRSP VWD index is a capitalization-weighted index of all domestic equities trading on major U.S. exchanges, and includes dividends and other distributions. Figure 1.1 shows the performance of the index itself (i.e. without withdrawals) and the buy and hold (stock only) strategy (after withdrawals) during the period 2000-2015. It turned out that the regular withdrawals and two major market shocks (dot-com and financial crisis) hit investors following this strategy very hard. By 2015, this retiree was left with about $400,000.

On the other hand, suppose the investor was very cautious. In 2000, long-term U.S. Treasuries were yielding about 6.5%. Assume that the investor bought $770,000 of U.S. Treasuries maturing in 2015, which would generate $50,000 per year. In addition, suppose that the remaining $230,000 was invested in the stock index. The investor's portfolio in this case is also shown in Figure 1.1 (buy and hold, stock and bond, after withdrawals). The investor would have fared much better by

Strategy	Initial stock investment (2000)	Initial bond investment (2000)	Total withdrawals (2000-2015)	Final portfolio value (2015)
Buy and hold	$1,000,000	$0	$0	$1,876,844
Buy and hold	$1,000,000	$0	$750,000	$393,370
Buy and hold	$230,770	$769,230	$750,000	$1,202,349
Rebalance	$500,000	$500,000	$750,000	$525,159

Table 1.1: *Comparison of strategies which generate $50,000 per year, based on historical data for 2000-2015. In the rebalance case, the rebalancing is done monthly, and the bond component of the portfolio is invested in short term T-bills.*

following this approach, having around $1.2 million by 2015.[1]

A more classic strategy involves investing equal amounts in the stock index and short term bonds, with periodic rebalancing (Graham, 2014). Figure 1.1 shows the performance of this strategy, rebalancing monthly (rebalance, stock weight 50%, after withdrawals). For this illustration, historical short term T-bill rates were used for the bond investment. By following this strategy, the investor would have ended up with about $525,000 after withdrawing $50,000 annually, quite an improvement over the stock only case. The results for these cases are summarized in Table 1.1.

We can see from this example that the choice of an asset allocation strategy combined with regular withdrawals can have a very large effect on the terminal wealth. However, in view of the fact that defined benefit pension plans are rapidly disappearing, many individuals who are planning for retirement are faced with this predicament. The basic decision variables are (i) an asset allocation strategy; and (ii) a sustainable withdrawal rate.

It is not surprising that this wealth management issue has received much attention. A classic analysis is described in Bengen (1994). Essentially, Bengen examined historical data to determine the maximum inflation-adjusted withdrawal rate that a retiree can safely use without exhausting her assets over a 35 year period. It was assumed that the initial endowment was invested in a constant mix of 50% stocks and 50% intermediate-term U.S. Treasuries. The main conclusion reached was that a 4% withdrawal rate (escalated by the rate of inflation) could be considered to be quite safe. This 4% rule is frequently cited by financial planners, and optimal withdrawal rules under various assumptions have been the subject of many other studies. For a small selection of this large literature, see sources such as Milevsky and Young (2007), Scott et al. (2009), Horneff et al. (2010), and Milevsky and Huang (2011).

In this paper, we consider an asset allocation problem in which an investment manager can dynamically switch total wealth between risk-free assets (e.g. short-term government bonds) and a risky asset (e.g. an index ETF). The investment fund is also subject to periodic withdrawals that are inflation-adjusted. We consider two concrete applications of optimal asset allocation. In the first case, we consider the management of an endowment fund. The objective is to manage the fund so that the expected real value of the endowment is maintained at the end of a relatively long time horizon, which might typically be 20-30 years. More precisely, we seek to find the asset allocation strategy which permits specified periodic withdrawals from the endowment while preserving the

[1] This ballpark estimate is conservative in the following sense. The Treasury does not actually issue 15-year bonds. If the investor had bought newly issued 30-year bonds for their par value at the start of 2000, the same $50,000 of coupon income would have been collected each year but there would also have been a significant capital gain due to declining interest rates between 2000 and 2015. In other words, as of 2015 the investor would have owned bonds with a remaining maturity of 15 years that were worth substantially more than par.

real value of the endowment at the specified time horizon with the smallest possible risk.

We also consider the retirement spending problem addressed by Bengen (1994). There are several possible ways to formulate this problem. One way would be to specify an estimate of the longevity of the retiree, plus a longevity buffer. This might give us a target of 30-40 years for withdrawals. Bengen (1994) essentially determined the withdrawal rate (using historical data) which gave a worst case withdrawal longevity of about 35 years. This would be a conservative target for a 65-year old retiree. Another possibility would be to examine withdrawal rates which result in an expected value of zero wealth at the 35-year mark. Given this target wealth, we would then determine the withdrawal rate which hits this expected value of terminal wealth with an acceptable level of risk.

However, this strategy seems somewhat unsatisfactory. Due to the risk of exhausting assets before death, most individuals would probably use an overly conservative longevity estimate. Consequently, in this paper we pose the problem somewhat differently. We consider an initial investment horizon of 20 years. Our target expected wealth value at the end of this time is one-half of the original wealth (in real terms). Many 65-year olds can expect to live for at least 20 years. At the end of this time, if all goes according to plan, then the retiree will have half of her initial real wealth. At that point, the retiree can re-evaluate her personal situation, in terms of health, spending habits and bequest motivation. It seems to us that this target is a reasonable compromise allowing a conservative buffer after 20 years, without being unnecessarily cautious. The strategy can then be re-evaluated in light of changing circumstances. The 20-year initial time horizon is long enough to allow the optimal strategy to recover from downward market shocks, if any. To summarize, the retirement withdrawal rate (pension de-accumulation) problem is formulated as follows. Given a specified real withdrawal rate, we determine the optimal asset allocation strategy which results in an expected value of one-half the real initial wealth with the smallest possible risk after 20 years.

In this study, for either the endowment problem or the pension de-accumulation problem, we determine the optimal asset allocation which minimizes risk in terms of a multi-period MV strategy. Variance has been criticized as a risk measure since it penalizes the upside as well as the downside. However, the analysis of Zhou and Li (2000) and Li and Ng (2000) shows that continuous time MV asset allocation is equivalent to specifying a wealth target with a quadratic shortfall penalty. Vigna (2014) notes that the quadratic wealth target is never exceeded in the case where continuous rebalancing is allowed and the price of the risky asset is assumed to follow geometric Brownian motion (GBM). In this sense, continuous time MV asset allocation seeks to hit an expected value target while simultaneously minimizing two risk measures: variance and quadratic loss with respect to the quadratic wealth target, which is slightly above the expected value. See Vigna (2014) for a discussion of the practical implications of this result.

The strategy used in this paper is a pre-commitment policy. As noted in Basak and Chabakauri (2010), this is not time consistent. However, as pointed out in Wang and Forsyth (2011), a time consistent policy can be generated by adding a constraint to the pre-commitment algorithm. Hence, the time consistent strategy will generally be sub-optimal compared to the pre-commitment policy (Wang and Forsyth, 2011, 2012). We take the point of view that forcing time consistency is expensive and thus undesirable for the long-term investor.

If we permit rebalancing of the assets only at discrete intervals (e.g. yearly) and we use a jump diffusion model for the underlying risky asset in order to model the possibility of market crashes, then it is possible to exceed the quadratic wealth target. However, based on an observation of Cui et al. (2012), if we allow the possibility of optimally withdrawing cash from the investment portfolio, we can achieve an investment strategy which is never inferior and usually is superior in the MV sense to a policy which does not permit cash withdrawals (Dang and Forsyth, 2016; Forsyth and Vetzal, 2016). In this way, we can ensure that the quadratic wealth target is never exceeded at the end of the investment horizon, i.e. we do not penalize the upside.

The remainder of the paper is structured as follows. Section 2 describes the formulation of the MV wealth management problem, which requires solving a partial integro-differential equation (PIDE). Section 3 discusses various relevant details about the specification of withdrawal rates. Section 4 presents extensive numerical results for both the endowment and the pension de-accumulation problems. Several interesting properties of the MV optimal strategy are shown. We also demonstrate its superiority over constant proportion strategies and its robustness to parameter and model uncertainty. Section 5 provides a concluding summary.

2 Preliminaries

2.1 Assets

For simplicity, we assume that just two assets are available in the financial market, namely a risky asset and a risk-free asset. Let $S_t \equiv S(t)$ and $B_t \equiv B(t)$ respectively be the *amounts* (i.e. total dollars) invested in the risky asset and the risk-free asset at time $t \in [0, T]$, where T is the time horizon of the investment. In the following, we are interested in the terminal value of the total wealth $W_T = S_T + B_T$.[2]

First consider the risky asset. Define $t^- = t - \epsilon$, i.e. t^- is the instant of time before the (forward) time t, and let ξ be a random number representing a jump multiplier. When a jump occurs, $S_t = \xi S_{t^-}$. Allowing discontinuous jumps permits us to explore the effects of severe market crashes on the risky asset holding. As a specific example, as in Merton (1976) we assume that ξ follows a log-normal distribution $p(\xi)$ given by

$$p(\xi) = \frac{1}{\sqrt{2\pi}\zeta\xi} \exp\left(-\frac{(\log(\xi) - \nu)^2}{2\zeta^2}\right), \tag{2.1}$$

with mean ν and standard deviation ζ, with $\mathrm{E}[\xi] = \exp(\nu + \zeta^2/2)$, where $\mathrm{E}[\cdot]$ denotes the expectation operator. In the absence of control (i.e. if we do not adjust the amount invested according to our control strategy), the amount invested in the risky asset S follows the process

$$\frac{dS_t}{S_{t^-}} = (\mu - \lambda\kappa)dt + \sigma dZ + d\left(\sum_{i=1}^{\pi_t}(\xi_i - 1)\right). \tag{2.2}$$

where $\kappa = \mathrm{E}[\xi] - 1$, dZ is the increment of a Wiener process, μ is the real world drift rate, σ is the volatility, π_t is a Poisson process with positive intensity parameter λ, and ξ_i are i.i.d. positive random variables having distribution (2.1). Moreover, ξ_i, π_t, and Z are assumed to all be mutually independent.[3]

Also, it is assumed that in the absence of control the dynamics of the amount invested in the risk-free asset B are given by

$$dB_t = rB_t dt, \tag{2.3}$$

[2] Unlike previous work (e.g. Björk, 2009; Vigna, 2014), we do not assume that the portfolio is continuously rebalanced. As a result, we cannot specify the total wealth process in terms of a single stochastic differential equation. Consequently, it is simpler to define $S(t)$ and $B(t)$ in terms of dollar amounts invested, rather than prices of a unit investment in each assets, as is typically done with continuous rebalancing.

[3] One may argue that it would be preferable to include stochastic volatility effects in the S process. However, recent tests indicate that stochastic volatility has little effect on long-term dynamic MV optimal strategies (Ma and Forsyth, 2016).

where r is the (constant) risk-free rate. We make the standard assumption that the real world drift rate of S is strictly greater than r. Since there is only one risky asset, it is never optimal in an MV setting to short stock, i.e. $S_t \geq 0$, $t \in [0, T]$. However, we do allow short positions in the risk-free asset, i.e. it is possible that $B_t < 0$, $t \in [0, T]$.

In some of the examples considered in this paper, we assume that the dynamics for S_t (absent control) follows GBM. This is implemented by suppressing any possible jumps in (2.2), i.e. by setting the intensity parameter λ to zero.

2.2 A discrete rebalancing/withdrawal model

To avoid the unrealistic assumption of continuous rebalancing or withdrawals, we consider a set of pre-determined *intervention times* denoted by $\mathcal{T}$,

$$\mathcal{T} \equiv \{t_0 < \cdots < t_M = T\}. \tag{2.4}$$

At these intervention times, cash withdrawals can be made and the investor's portfolio may be rebalanced. To keep transaction costs to a minimum, these times would be typically at an annual or quarterly frequency. As discussed in Forsyth and Vetzal (2016), with long-term investment horizons, the result for an optimal strategy with yearly rebalancing is quite close to the result obtained using continuous rebalancing.

Let t_0 (i.e. $t = 0$) be the inception time of the investment. For simplicity, we specify the set of intervention times (2.4) to be equidistant with $t_m - t_{m-1} = \Delta t = T/M$, $m = 1, \ldots, M$.

At an intervention time, the investor withdraws an amount of cash, denoted by a_m, from the risk-free asset. The amount withdrawn at time t_m is denoted by a_m and is determined by

$$a_m = \begin{cases} a(t_m - t_{m-1})\mathrm{e}^{I\, t_m} = a\Delta t e^{I\, t_m} = a\Delta t e^{m\, I\, \Delta t}, & m = 1, \ldots, M \\ 0 & m = 0 \end{cases} \tag{2.5}$$

where a is the (continuous and constant) withdrawal rate and I is the (continuous and constant) inflation rate. Note that we assume there is no withdrawal at t_0. These periodic withdrawals are used to fund living expenses (in the pension de-accumulation case) or endowment cash flows. The presence of the inflation factor I preserves the real value of the withdrawals over time.

In addition, at intervention times $t_0, \ldots, t_{M-1}$, the investor adjusts the amounts in the stock and bond (i.e. rebalances the portfolio). At intervention times $t_1, \ldots, t_{M-1}$, where both the specified cash withdrawal and rebalancing occur, we assume that the cash is withdrawn first and then the portfolio is rebalanced.

2.3 Controls at each rebalancing date

We denote by $X(t) = (S_t^c, B_t^c)$, $t \in [0, T]$, the multi-dimensional (controlled) underlying process. The control generates a new allocation of the stock and bond. Let $c(\cdot) \equiv (\hat{b}(\cdot), \hat{f}(\cdot))$ denote the control as a function of the current state at $t \in [0, T]$, i.e.

$$c(\cdot) : (X(t^-), t^-) \mapsto c = c(X(t^-), t^-) \equiv (\hat{b}(X(t^-), t^-), \hat{f}(X(t^-), t^-)) \equiv (\hat{b}(t), \hat{f}(t)). \tag{2.6}$$

At each rebalancing time $t_m \in \mathcal{T}$, let the control be denoted by c_m, $m = 0, \ldots, M$, where

$$c_m = \begin{cases} (\hat{b}_m, \hat{f}_m) & m = 0, \ldots, M-1 \\ (0, 0) & m = M \end{cases}, \tag{2.7}$$

where we assume no rebalancing at the terminal time T. In (2.7), $\hat{b}_m$ is the amount of the risk-free asset after rebalancing and $\hat{f}_m$ is the free cash flow generated. The optimal withdrawal of free cash is discussed in Cui et al. (2012) and Dang and Forsyth (2016). We will describe $\hat{f}_m$ in detail below in Section 2.8.

Let $x \equiv (s, b) = (S^c_{t^-}, B^c_{t^-})$ denote the state of the portfolio at time t^-, $t \in [0, T]$. We denote by $(S^+, B^+) \equiv S^+(s, b, c, t),\ B^+(s, b, c, t)$ the state of the system immediately after application of the control $c \equiv (\hat{b}, \hat{f})$. After a scheduled withdrawal a_m and application of the controls $(\hat{b}_m, \hat{f}_m)$ at time $t_m \in \mathcal{T}$, we have

$$\begin{aligned} S^+(s, b, c, t_m) &= S^+_{t_m} = s + b - a_m - \hat{b}_m - \hat{f}_m \\ B^+(s, b, c, t_m) &= B^+_{t_m} = \hat{b}_m\ . \end{aligned} \tag{2.8}$$

2.4 Allowable controls

Let the (controlled) wealth of the portfolio at time $t \in [0, T]$ be given by

$$W^c_t \equiv W(S^c_t, B^c_t) = S^c_t + B^c_t, \quad t \in [0, T].$$

We strictly enforce the solvency condition, i.e. the investor can continue trading only if $((s, b) = (S^c_{t^-}, B^c_{t^-}))$

$$W(s, b) = s + b > 0. \tag{2.9}$$

In the event of insolvency, we require that the investor immediately liquidate all investments in the risky asset and stop trading, i.e.

$$S^+ = 0; \quad B^+ = W(s, b); \quad \text{if } W(s, b) \le 0. \tag{2.10}$$

Equation (2.10) holds for all $t \in [0, T]$. We also constrain the leverage ratio, i.e. the investor must select an allocation satisfying $(t_m \in \mathcal{T})$

$$\frac{S^+_{t_m}}{S^+_{t_m} + B^+_{t_m}} \le q_{\max}, \tag{2.11}$$

where $q_{\max}$ is a specified positive constant. In particular, for the endowment scenario we use $q_{\max} = 1$, whereas in the pension de-accumulation scenario we set $q_{\max} = 1.5$.

More precisely, define the solvency region $\mathcal{N}$ as

$$\mathcal{N} = \{(s, b) \in [0, \infty) \times (-\infty, +\infty) : s + b > 0\}. \tag{2.12}$$

The insolvency (or bankruptcy) region $\mathcal{B}$ is defined as

$$\mathcal{B} = \{(s, b) \in [0, \infty) \times (-\infty, +\infty) : s + b \le 0\}. \tag{2.13}$$

Let

$$\mathcal{Z}_{\mathcal{N}} = \Big\{ c \equiv (B, \hat{f}) \in [-\infty, +\infty) \times (0, +\infty) : S = (s + b) - a_m - \hat{f} - B, \\ \text{where } t \in \mathcal{T}, S \ge 0, \text{ and } 0 \le \frac{S}{S + B} \le q_{\max} \Big\},$$

and

$$\mathcal{Z}_{\mathcal{B}} = \begin{cases} \Big\{ c \equiv (\hat{b}, \hat{f}) = (s + b, 0) \Big\} & ; t \in [0, T] \setminus \mathcal{T} \\ \Big\{ c \equiv (\hat{b}, \hat{f}) = (s + b - a_m, 0) \Big\} & ; t \in \mathcal{T} \end{cases}. \tag{2.14}$$

The set of admissible controls $\mathcal{Z}$ is then

$$\mathcal{Z} = \begin{cases} Z_{\mathcal{N}} & \text{if } (s+b) \geq 0 \\ \mathcal{Z}_{\mathcal{B}} & \text{if } (s+b) < 0 \end{cases} . \tag{2.15}$$

2.5 Efficient frontiers and embedding methods

We now discuss how MV efficient frontiers can be determined in our setting. Let $\mathrm{E}^{t,x}[W_T^c]$ and $\mathrm{Var}^{t,x}[W_T^c]$ respectively denote the expectation and the variance of the controlled terminal wealth W_T^c conditional on the state (t,x) and on the control $c(\cdot)$. We denote the initial state by $(t_0, x_0) = (t = 0, X_0)$. Then the *achievable MV objective set* $\mathcal{Y}$ is

$$\mathcal{Y} = \{(\mathrm{Var}^{t_0,x_0}[W_T^c], \mathrm{E}^{t_0,x_0}[W_T^c]) : c \in \mathcal{Z}\}, \tag{2.16}$$

where $\mathcal{Z}$ is the set of admissible controls (2.15). For each point $(\mathcal{V}, \mathcal{E}) \in \mathcal{Y}$, and for an arbitrary scalar $\rho > 0$, define the set of points $\mathcal{Y}_{P(\rho)}$ as

$$\mathcal{Y}_{P(\rho)} = \big\{(\mathcal{V}_*, \mathcal{E}_*) \in \bar{\mathcal{Y}} : \rho\mathcal{V}_* - \mathcal{E}_* = \inf_{(\mathcal{V},\mathcal{E})\in\mathcal{Y}} \rho\mathcal{V} - \mathcal{E}\big\}. \tag{2.17}$$

Here, $\bar{\mathcal{Y}}$ denotes the closure of $\mathcal{Y}$, and ρ can be viewed as a risk-aversion parameter which governs how the investor trades off expected return (reward) and variance (risk). For a given ρ, $\mathcal{Y}_{P(\rho)}$ represents Pareto efficient points in that, given the variance of any point in $\mathcal{Y}_{P(\rho)}$, the corresponding expectation is the largest expectation that can be obtained for that variance.

Note that we have made no assumptions about the convexity (or lack thereof) of the achievable set. If the upper boundary of $\bar{\mathcal{Y}}$ is not strictly convex, use of the scalarization method (2.17) may not generate all possible Pareto points, but any points in $\mathcal{Y}_{P(\rho)}$ are sure to be Pareto optimal. In addition, $\mathcal{Y}_{P(\rho)}$ may not be a singleton.

The set of points on the efficient frontier, denoted by $\mathcal{Y}_P$, is just the collection of efficient points for all values of $\rho > 0$, i.e.

$$\mathcal{Y}_P = \bigcup_{\rho>0} \mathcal{Y}_{P(\rho)}.$$

In the context of MV optimal asset allocation, one of the primary objectives is to determine the efficient frontier $\mathcal{Y}_P$. However, as noted in the literature (see, e.g. Zhou and Li, 2000; Li and Ng, 2000; Basak and Chabakauri, 2010), the presence of the variance term in (2.17) causes difficulty if we try to determine $\mathcal{Y}_{P(\rho)}$ by solving the associated value function problem using dynamic programming. This problem can be circumvented by using the embedding result in Zhou and Li (2000) and Li and Ng (2000). More specifically, consider the set

$$\mathcal{Y}_{Q(\gamma)} = \inf_{c(\cdot)\in\mathcal{Z}} \left\{ \mathrm{E}^{t_0,x_0}[(W_T^c - \gamma/2)^2] \right\}, \tag{2.18}$$

where the parameter $\gamma \in (-\infty, +\infty)$, and the set

$$\mathcal{Y}_Q = \bigcup_{-\infty<\gamma<+\infty} \mathcal{Y}_{Q(\gamma)}.$$

The embedding result implies that there exists a $\gamma \equiv \gamma(t, x, \rho)$, such that for a given positive ρ, an optimal control c^* of (2.17) is also an optimal control of (2.18). Furthermore, we have the relation (Zhou and Li, 2000):

$$\frac{\gamma}{2} = \frac{1}{2\rho} + \mathrm{E}^{t_0,x_0}[W_T^{c^*}],$$

which implies that $\mathcal{Y}_P \subseteq \mathcal{Y}_Q$. In the following, we refer to $\gamma/2$ as the *quadratic wealth target* and

$$E^{t_0,x_0}[W_T^{c^*}]$$

as the *expected wealth target*.

2.6 Value function and its solution

We define the value function

$$V(t,x) = V(t,s,b) = \inf_{c(\cdot)\in\mathcal{Z}} \left\{ E^{t,x}[(W_T^c - \gamma/2)^2] \right\} . \tag{2.19}$$

We solve for the value function using dynamic programming, backwards from the terminal time $t = T$ to the initial time $t = 0$. Define the following operator

$$\mathcal{L}V \equiv \frac{\sigma^2 s^2}{2}\frac{\partial^2 V}{\partial s^2} + (\mu - \lambda\kappa)s\frac{\partial V}{\partial s} + rb\frac{\partial V}{\partial b} - \lambda V + \int_0^\infty p(\xi)V(\xi s, b, t)\, d\xi. \tag{2.20}$$

For a given value of γ, we compute the corresponding point on the efficient frontier by first solving problem (2.19). We use the following algorithm to solve for the value function (2.19) using dynamic programming. In reverse time order, at times $t_m, m = M, \ldots, 0$, we enforce the following conditions:

1. If $(s,b) \in \mathcal{B}$, we enforce the liquidation condition

$$V(t_m^-, s, b) = V(t_m, 0, s + b - a_m). \tag{2.21}$$

2. If $(s,b) \in \mathcal{N}$, we determine the optimal control c_m^*

$$\begin{aligned} V(t_m^-, s, b) &= \inf_{c_m \in \mathcal{Z}} V(t_m, S_m^+, B_m^+) \\ c_m^* &= (\hat{b}_m, \hat{f}_m) \\ S_{t_m}^+ &= s + b - a_m - \hat{b}_m - \hat{f}_m \ ; \ B_{t_m}^+ = \hat{b}_m. \end{aligned} \tag{2.22}$$

   As noted above, $\mathcal{Z}$ is the set of *admissible* controls, defined in equation (2.15). Note that for the special case of $t_m = T$, we have $V(T,s,b) = (W(s,b) - \gamma/2)^2$.

Within each time period $[t_{m-1}, t_m)$, $m = M, \ldots, 1$, we have

1. If $(s,b) \in \mathcal{B}$, we enforce the liquidation condition

$$V(t,s,b) = V(t,0,s+b). \tag{2.23}$$

2. If $(s,b) \in \mathcal{N}$, $V(t,s,b)$ satisfies the PIDE

$$\frac{\partial V}{\partial t} + \mathcal{L}V = 0, \tag{2.24}$$

   subject to the initial condition (2.22). We solve PIDE (2.24) from $t_m^- \to t_{m-1}$.

Equation (2.24) follows from equations (2.2) and (2.3) using standard dynamic programming arguments (Øksendal and Sulem, 2009). See Dang and Forsyth (2014) for relevant details regarding a derivation of the localized problem. We numerically solve this localized problem using finite differences with a semi-Lagrangian timestepping method as described in Dang and Forsyth (2014).

2.7 Constructing the efficient frontier

We denote by $c^*_\gamma(\cdot)$ the optimal control of problem (2.19). Once we have determined $c^*_\gamma(\cdot)$ from the solution process described above, we use this control to determine

$$U(t,x) = \mathrm{E}^{t,x}[W_T^{c^*_\gamma}], \tag{2.25}$$

since this information is needed to determine the corresponding MV point on the efficient frontiers. This essentially involves solving an associated linear partial differential equation (PDE), details of which are similar to those described in Dang and Forsyth (2014) and hence are omitted here. Using numerical solutions for (2.19) and (2.25) evaluated at (t_0, x_0), we compute the variance and expectation point $(\mathrm{Var}^{t_0,x_0}[W_T^{c^*_\gamma}], \mathrm{E}^{t_0,x_0}[W_t^{c^*_\gamma}])$. Repeating this procedure for different values of γ traces out the efficient frontier.

This procedure for constructing the efficient frontier generates points that are MV optimal with respect to the embedding problem. While all the points in the original MV efficient frontier $\mathcal{Y}_P$ are MV optimal with respect to the embedding problem, note that the converse does not necessarily hold. This is an important issue in the context of a numerical algorithm. An algorithm for removing spurious points and relevant discussions are presented in Tse et al. (2014) and Dang et al. (2016).

2.8 Semi-self financing: optimal free cash withdrawal

In the solution process for the value function (2.18), we employ the semi-self-financing strategy discussed in Dang and Forsyth (2016). More specifically, given a rebalancing time t_k, $k = 1, \ldots, M$, the time-t_k value of all specified cash withdrawals a_m made on or after time t_k, denoted by w_k, is computed by

$$w_k = \sum_{m=k}^{M} a_m \mathrm{e}^{-r\,(t_m - t_k)} = \sum_{m=k}^{M} (a\Delta t)\mathrm{e}^{I\,t_m}\mathrm{e}^{-r\,(t_m-t_k)} = (a\Delta t)\sum_{m=k}^{M} \mathrm{e}^{\Delta t\,(I\,m - r(m-k))}. \tag{2.26}$$

At time t_k, if $W_{t_k} > \frac{\gamma}{2}\mathrm{e}^{-r\,(T-t_k)} + w_k$, where w_k is defined in (2.26), we (i) withdraw $W_{t_k} - (\frac{\gamma}{2}\mathrm{e}^{-r(T-t_k)} + w_k)$ from the portfolio, and (ii) invest the remaining wealth $(\frac{\gamma}{2}\mathrm{e}^{-r(T-t_k)} + w_k)$ in the risk-free asset for the balance of the investment horizon. We refer to the amount of (i) as "free cash flow" to clearly distinguish it from the specified cash withdrawal a_m.

As shown in Dang and Forsyth (2016), this strategy is MV optimal. This is easy to see: suppose that at time t_k, after withdrawal of free cash, we have precisely $W^*_{t_k} = \frac{\gamma}{2} + w_k$, with w_k given from equation (2.26). If $W^*_{t_k}$ is invested in risk-free bonds, then after withdrawals we will have $W_T = \frac{\gamma}{2}$ with certainty. From (2.19), we have $V(x,t) \equiv 0$. Since $V(x,t) \geq 0$, this is an optimal strategy. Assuming that this value of γ generates a valid point on the original MV efficient frontier (Tse et al., 2014), then this strategy must also be MV efficient.

Since the free cash flow falls outside the scope of the MV framework, we do *not* include the expected value of the free cash flow and accumulated interest in the terminal portfolio wealth. We also do not use the free cash flow to fund any of the withdrawals. More precisely, the control variable $\hat{f}_m$ in (2.22) is given by

$$\hat{f}_m = \max\bigg((s+b) - \frac{\gamma}{2}e^{-r\,(T-t_m)} - w_m, 0\bigg) \tag{2.27}$$

If the free cash flow is available at a rebalancing time, it might make sense to substitute the free cash flow for part of the actual cash withdrawal at that time. However, our numerical experiments indicate that with annual rebalancing and downward (on average) jumps, the expected free cash flow is very small compared to the actual cash withdrawal. Consequently, using the free cash flow to reduce the withdrawals has very little impact on the results. For simplicity, we ignore these possible benefits.

We conclude this section by noting that for a given withdrawal rate a and provided W_0 is large enough, there is an obvious strategy which generates zero variance: at each rebalancing time t_m, $m = 1, \ldots, M$, invest in the risk-free asset all wealth after the withdrawal amount a_m is made. The certain value of the terminal portfolio wealth corresponding to this risk-free strategy, denoted by $\mathcal{E}_{\text{rf}}$, is

$$\mathcal{E}_{\text{rf}} = W_0 \mathrm{e}^{rT} - a\Delta t \sum_{m=1}^{M} \mathrm{e}^{I\, t_m} \mathrm{e}^{r\,(T-t_m)} = W_0 \mathrm{e}^{rT} - a\Delta t \sum_{m=1}^{M} \mathrm{e}^{\Delta t\,(I\,m+(M-m)\,r)}. \tag{2.28}$$

For future reference, let a_{one} be the second term on the right side of (2.28), i.e.

$$a_{\text{one}} = a\Delta t \sum_{m=1}^{M} \mathrm{e}^{\Delta t\,(I\,m+(M-m)\,r)}. \tag{2.29}$$

The quantity a_{one} is the time T value of all cash withdrawals a_m, $m = 1, \ldots, M$.

2.9 Target driven investing

Recall that the optimal control for multi-period MV objective functions can be found by determining the optimal control for (2.19). Since we use the semi-self-financing policy in Dang and Forsyth (2016), then $W_T \leq \gamma/2$. Hence multi-period MV optimal strategies can be interpreted as minimizing the quadratic loss with respect to $\gamma/2$.

As a result, we can view our strategies as a form of target driven investing, i.e. find the strategy $c(\cdot)$ which solves

$$\begin{aligned} &\inf_{c(\cdot)\in\mathcal{Z}} \left\{ \mathrm{E}^{t_0,x_0}[(W_T^c - \mathcal{W})^2] \right\} \\ &\mathcal{W} \text{ determined from the constraint } \mathrm{E}^{t_0,x_0}[W_T^c] = d \\ &d = \text{ specified by the investor }. \end{aligned} \tag{2.30}$$

We can identify $\mathcal{W} = \gamma/2$, and the quantity $(\mathcal{W} - E^{t_0,x_0}[W_T^c]) = 1/(2\rho)$ can be regarded as safety factor, which ensures that we achieve the specified expected value with the smallest possible quadratic loss with respect to the target $\mathcal{W}$. Of course, the control which solves problem (2.30) is also MV optimal.

3 Withdrawal rates

Our objective is to investigate the effects of withdrawal rates on the risk associated with the expectation of terminal portfolio wealth. In particular, we study this issue from the perspective of an investor who wants to maintain a specified level of the expected terminal wealth $\mathrm{E}[W_T]$. We denote this expected wealth target by W_{spec}.[4] In our endowment context,

$$W_{\text{spec}} = W_{\text{spec}}^{(1)} = W_0\, \mathrm{e}^{I\,T}, \tag{3.1}$$

[4]We are careful to distinguish the expected wealth target W_{spec} from the quadratic wealth target $\gamma/2$ in equation (2.18).

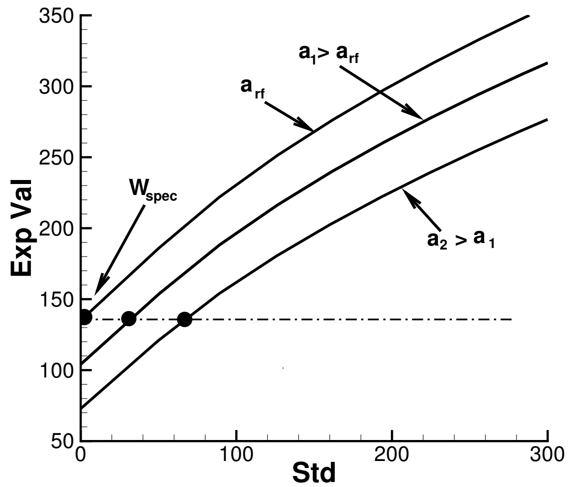

FIGURE 3.1: *Supporting higher withdrawals while keeping the same level of W_{spec} requires taking on more risk.*

i.e. the investor wants to maintain the real value of the endowment. In the pension de-accumulation case, we use

$$W_{\text{spec}} = W_{\text{spec}}^{(2)} = \frac{W_0}{2}\,\mathrm{e}^{I\,T}. \tag{3.2}$$

That is, the individual wants to maintain one half of the real value of the original wealth at the end of the time horizon T. As discussed above in the Introduction, if T is of the order of half of the remaining maximum life expectancy, then this objective allows the investor to re-evaluate her strategy in light of health and bequest motives, while still allowing enough time for the optimal asset allocation strategy to recover from possible downward jumps in the risky asset.

It follows from (2.28) that there exists a withdrawal rate at which W_{spec} can be achieved with zero variance if $W_0\,\mathrm{e}^{r\,T} - W_{\text{spec}} > 0$. This withdrawal rate, denoted by a_{rf}, can be computed as

$$a_{\text{rf}} = \frac{W_0\,\mathrm{e}^{r\,T} - W_{\text{spec}}}{\sum_{m=1}^{M} \mathrm{e}^{\Delta t\,(I\,m + (M-m)\,r)}}. \tag{3.3}$$

We denote by $a_{\text{rf}}^{(1)}$ and $a_{\text{rf}}^{(2)}$ the values of a_{rf} in the endowment and de-accumulation scenarios respectively.

In Figure 3.1 we plot three efficient frontiers corresponding to three different values of the withdrawal rate, namely $a = a_{\text{rf}}$, $a = a_1 > a_{\text{rf}}$ and $a = a_2 > a_1$. The left most point on the efficient frontier when $a = a_{\text{rf}}$ corresponds to a portfolio that follows a zero variance strategy. The certain value of the terminal portfolio wealth in this case is W_{spec}. However, when $a = a_1$, the efficient frontier shifts downward. Hence, in this case a zero variance strategy results in $\mathrm{E}[W_T] < W_{\text{spec}}$. Thus, to maintain the same level of W_{spec}, some risk must be taken. For the higher withdrawal rate a_2, the amount of risk taken must be larger.

We are interested in a numerical study of the following issue. Given a withdrawal rate $a > a_{\text{rf}}$, what is the minimum amount of risk that must be taken to maintain the same level of W_{spec}? To answer this, given W_{spec} and a fixed withdrawal rate $a > a_{\text{rf}}$, we can find an optimal control $c^*(\cdot)$ which guarantees $\mathrm{E}^{t_0,x_0}[W_T^{c^*}] = W_{\text{spec}}$ with the smallest variance. This task can be embedded into

Parameters	Base case	Jump diffusion	Diffusion with effective volatility
μ (drift)	0.10	0.10	0.10
σ (volatility)	0.15	0.15	$0.23 = \sigma_{\text{eff}}$
λ (jump intensity)	N/A	0.10	N/A
ν (jump multiplier mean)	N/A	-0.50	N/A
ζ (jump multiplier std. dev.)	N/A	0.20	N/A
r (risk-free interest rate)	0.03	0.03	0.03
I (inflation rate)	0.02	0.02	0.02
W_0 (initial wealth)	100	100	100
T (investment horizon - years)	20	20	20
$t_{i+1} - t_i$ (rebalance interval - years)	1	1	1

TABLE 4.1: *Input parameters for the various cases. See definitions of jump diffusion parameters in* (2.1). *With these parameters, the expected jump multiplier is* $E[\xi] \simeq .62$.

the problem of finding the parameter γ from (2.19) for which $\text{E}^{t_0,x_0}[W_T^c] = W_{\text{spec}}$ can be achieved with the smallest variance. This can be solved using Newton's method (see Algorithm A.1 in Appendix A). To compare the effects of the withdrawal rate on the minimum amount of risk for these different cases, it is convenient to plot the standard deviations $\sqrt{\text{Var}^{t_0,x_0}[W_T^{c*}]}$ obtained with different withdrawal rates versus the withdrawal rates.

4 Numerical results

Default parameters for our experiments are given in Table 4.1. The *base case* is GBM with drift $\mu = .10$, volatility $\sigma = .15$, and a risk-free rate of $r = .03$. These parameter values are similar to those estimated by Forsyth and Vetzal (2016) using U.S. market data from the past 60 years. In addition, to examine the effects of market crashes we consider a *jump diffusion* case, with parameters selected so that jumps occur on average about once per decade and jump sizes which are on average strongly negative (on the order of about -40%), but with a fairly large standard deviation. We also include a *diffusion* case with an *effective* volatility which approximates the behavior of the jump diffusion model by a pure diffusion process (Navas, 2000). It is interesting to include this case as conventional wisdom asserts that over long times, jump diffusions can be approximated by diffusions with enhanced volatility. In our experiments, the effective (enhanced) volatility is computed as in Navas (2000), i.e.

$$\begin{aligned}\sigma_{\text{eff}} &= \sqrt{\sigma^2 + \lambda(\nu^2 + \zeta^2)} \\ &= \sqrt{0.15^2 + 0.10((-0.5)^2 + (0.2)^2)} \approx 0.23 \; .\end{aligned} \tag{4.1}$$

Our two scenarios of endowment and pension de-accumulation use different W_{spec} and $q_{\max}$. Table 4.2 lists these values, along with the quantity a_{rf} defined in (2.29). Our PIDE solutions below use 120 timesteps and 245 and 117 nodes in the b and s directions, respectively. Numerical tests show that this level of grid refinement gives about three digits of accuracy in the mean and standard deviation.

Scenario	W_{spec}	$q_{\max}$	a_{rf}
Endowment	$W_{\text{spec}} = W_{\text{spec}}^{(1)} = 149.2$	$q_{\max} = q_{\max}^{(1)} = 1.0$	$a_{\text{rf}} = a_{\text{rf}}^{(1)} = 1.0$
Pension de-accumulation	$W_{\text{spec}} = W_{\text{spec}}^{(2)} = 149.2/2$	$q_{\max} = q_{\max}^{(2)} = 1.5$	$a_{\text{rf}} = a_{\text{rf}}^{(2)} = 3.3$

TABLE 4.2: *Levels of W_{spec} and maximum leverage constraints in the endowment and wealth management scenarios.*

4.1 Effects of withdrawal rates

Figure 4.1 presents plots of the standard deviations vs. withdrawal rates to show the amount of risk required to maintain the same level of W_{spec} for different withdrawal rates. The withdrawal rates can be interpreted as the real withdrawal rate expressed as a per cent of initial wealth. Panel (a) shows the endowment case for a wide range of withdrawal rates, while panel (b) zooms in to provide a clearer comparison for relatively low withdrawal rates. Panels (c) and (d) are similar, but for the pension de-accumulation scenario. We observe that the effect of the withdrawal rate is quite substantial in both cases. For the same withdrawal rate under either scenario, the base case is the least risky. This is followed by the jump diffusion case and then the diffusion with effective volatility case. It seems clear that approximating the jump diffusion by a pure diffusion with enhanced volatility overstates the risk. This should be borne in mind in any practical application of these results since an empirical estimate of historical market volatility (based, e.g., on GBM) will produce an effective volatility that includes both diffusive and jump effects.

For the pension de-accumulation scenario with the input data considered, it appears that withdrawing at any rate above 5% probably involves unacceptably high risk, especially in the jump diffusion setting. A 4% withdrawal rate does seem to be a reasonable compromise here between risk and reward, but note that this is only under the assumption of an optimal asset allocation policy. For the endowment case, a 4% withdrawal rate seems to generate quite a bit of risk, while a 3% withdrawal rate would be considerably safer.

4.2 Order of random returns

As noted by Milevsky and Salisbury (2006), the order of random returns is irrelevant for long-term investors with no need to generate income each year. However, under both of our scenarios the investor does need to generate income each year, and so is exposed to risk embedded in the order of random returns. Losses in the early years of investment can be devastating, resulting in a rapid depletion of the fund.

To investigate the effects of the order of random returns, we carry out the following experiment. Instead of withdrawing the amount a_m every year, $m = 1, \ldots, M$, we withdraw only once at time $t_M = T$ the amount a_{one} defined in (2.29). In this one-off withdrawal, we take into account the time value of money of the annual withdrawals. We focus here exclusively on the endowment scenario.[5] Note that under this one-off withdrawal $a_{\text{rf}}^{(1)}$ remains unchanged at one, as indicated in Table 4.2, and W_{spec} is set to be the same as for the yearly withdrawals, i.e. 149.2.

Figure 4.2 shows plots of standard deviation vs. withdrawal rates for this experiment. Panel (a) shows the tradeoff between risk and withdrawal rate for the base, jump diffusion, and diffusion with effective volatility cases. Panels (b)-(d) respectively consider these three modeling cases, comparing the risk required to maintain the same level of W_{spec} for the one-off withdrawal and with that for the

[5] While similar numerical results are obtained for the pension de-accumulation case, the one-off withdrawal is unlikely to be practically feasible in that setting since pensioners rely on periodic withdrawals.

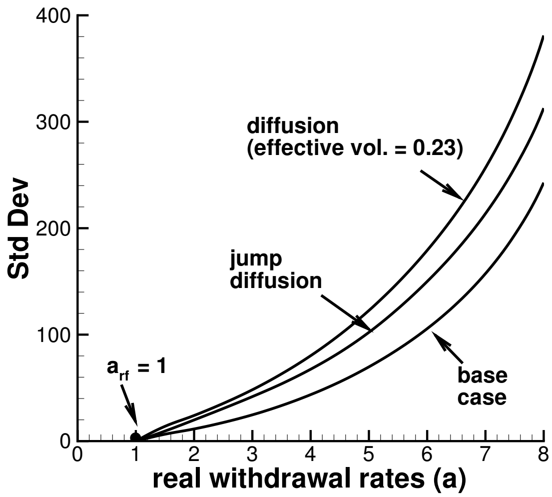

(a) *Endowment case.*

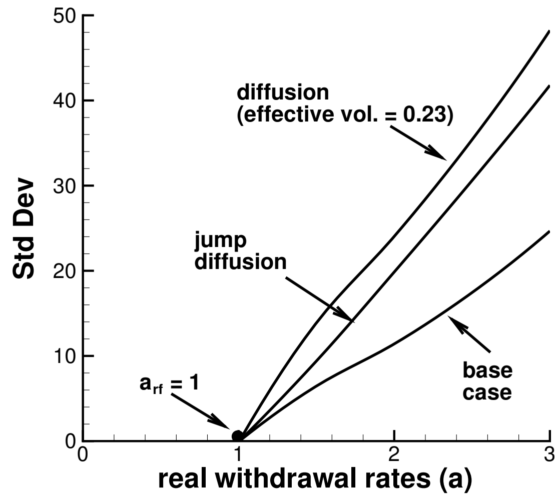

(b) *Endowment case: zoom in.*

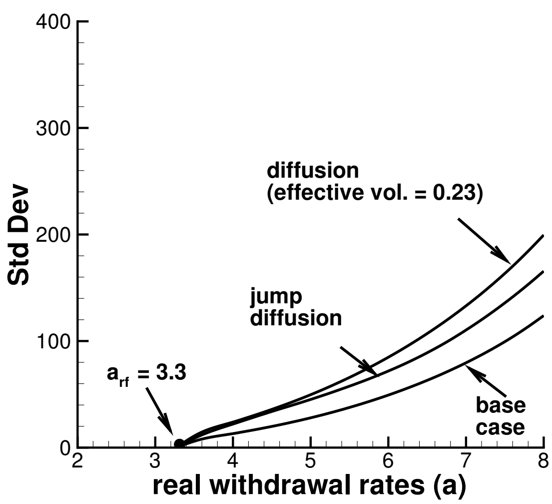

(c) *Pension de-accumulation case.*

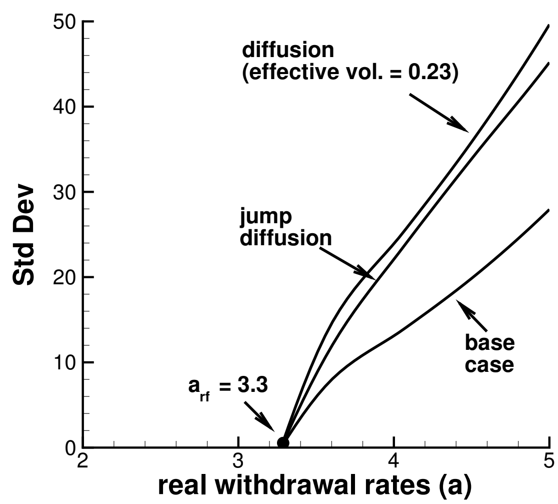

(d) *Pension de-accumulation case: zoom in.*

FIGURE 4.1: *Standard deviations vs. withdrawal rates. The withdrawal rates are expressed in real terms as a per cent of initial capital. Input data are given in Tables 4.1 and 4.2.*

annual withdrawals. These plots illustrate that order of return risk is highly significant, especially for large withdrawal rates.

4.2.1 An endowment strategy

The above results suggest a potentially interesting strategy for a charitable endowment which is concerned about risk due to the order of random returns. Suppose that the endowment has real assets (e.g. office buildings). The order of return risk could be eliminated by doing the following:

- Take out a bullet loan in the amount of $a_{\text{one}}\, e^{-rT}$ using the real assets as collateral, where a_{one} is given in equation (2.29). This loan is to be repaid entirely at $t = T$.

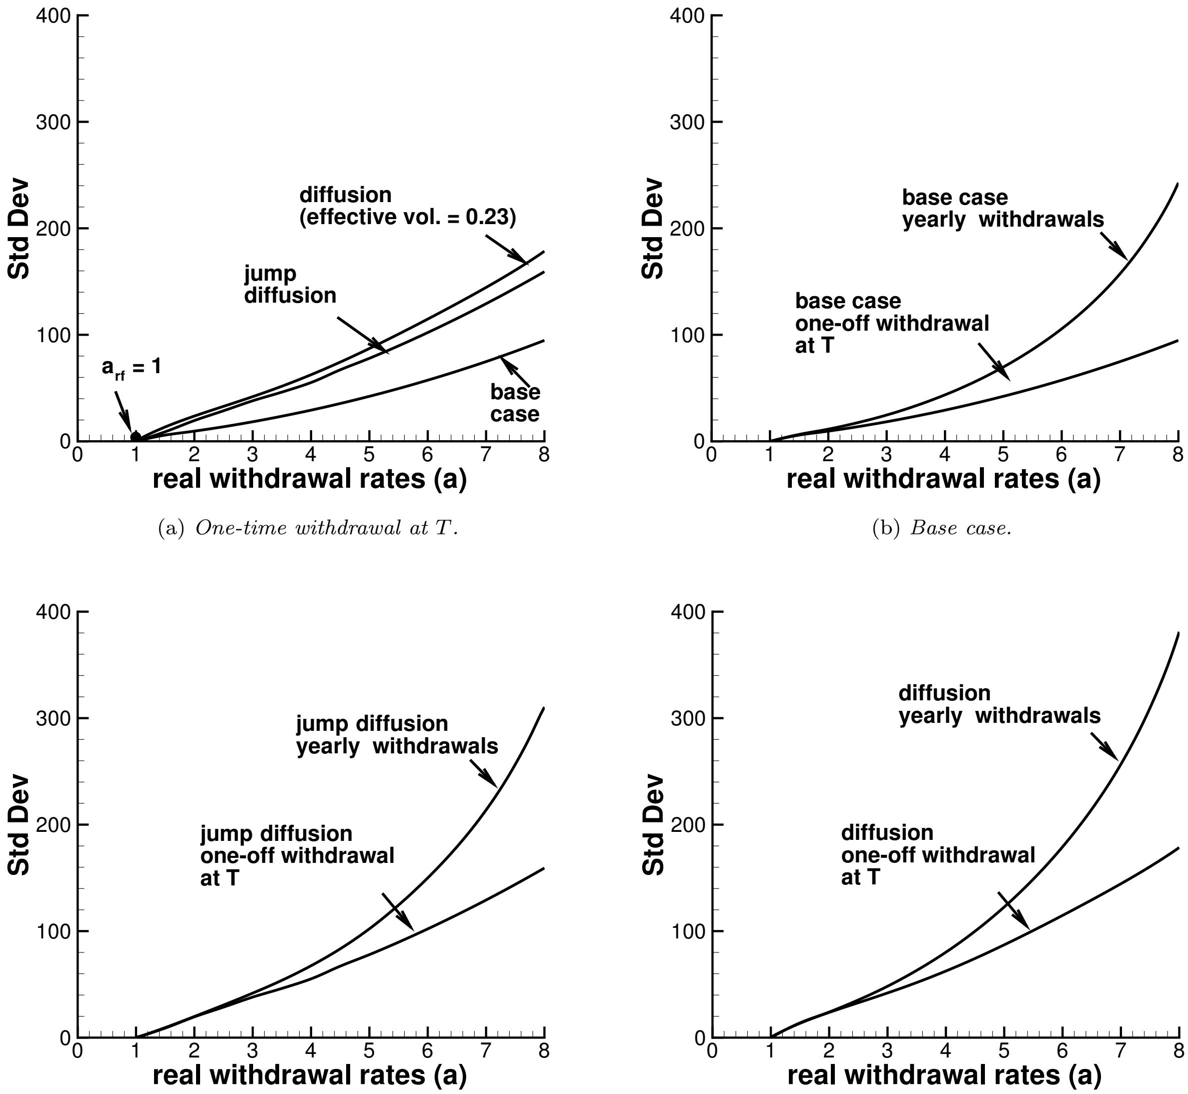

(a) *One-time withdrawal at T.* (b) *Base case.*

(c) *Jump diffusion case.* (d) *Diffusion with effective volatility case.*

FIGURE 4.2: *Standard deviation vs. withdrawal rate, comparing the risks of annual withdrawals with a one-off withdrawal at the investment horizon T for the endowment case. Input data are from Tables 4.1 and 4.2.*

- Invest the loan proceeds in risk-free assets to fund the endowment cash flows in $[0, T]$.
- Manage the investment portfolio using the optimal dynamic strategy over $[0, T]$. Note there are no withdrawals from the portfolio.
- At time $t = T$, make a balloon payment a_{one} to repay the loan.

Based on our results reported above, this strategy would reduce risk compared to funding the fixed cash flows from the investment portfolio.

Of course, since this strategy essentially involves borrowing, this can be reproduced by simply allowing more leverage in the set of admissible strategies. However, many endowments specifically

No. of MC simulations	No. of timesteps	Base case		Jump diffusion		Diffusion with effective volatility	
		Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
Endowment scenario							
4×10^4	80	148.7	103.1	148.6	156.8	146.5	178.8
16×10^4	160	149.1	104.9	149.0	158.9	148.7	180.8
64×10^4	320	149.2	105.7	149.2	159.9	149.2	182.1
PDE-computed		149.2	105.6	149.2	159.9	149.2	182.7
Pension de-accumulation scenario							
4×10^4	80	72.7	45.1	74.0	78.8	69.6	80.0
16×10^4	160	73.8	44.9	76.5	78.0	72.9	80.5
64×10^4	320	74.5	44.8	77.4	77.4	74.5	80.8
PDE-computed		74.6	44.7	74.6	77.1	74.6	80.7

TABLE 4.3: *Convergence of MC-computed and PDE-computed means and standard deviations. The withdrawal rate is $a = 6$. Other input data are from Tables 4.1 and 4.2.*

prohibit use of leverage. Hence the strategy of borrowing against real assets may be more acceptable to endowment trustees, even though this is clearly economically equivalent to use of leverage.[6]

4.3 Monte Carlo results

In this section, we carry out Monte Carlo (MC) simulations using the optimal controls generated by our PIDE solver. This provides further insight into the controlled investment process.

4.3.1 MC validation

As a model validation check, we proceed here to compute the mean and standard deviation of terminal wealth using both an MC method and the PDE approach (see Section 2.7). In particular, we proceed as follows. For each fixed value of $W^{(1)}_{\text{spec}}$ and $W^{(2)}_{\text{spec}}$ (see Table 4.2), we use the PIDE method described above in Section 2.6 to find optimal strategies which achieve this value with the smallest possible variance. These controls are stored for each discrete state value and timestep. We then carry out MC simulations from $t = 0$ to $t = T$ following these stored PIDE-computed optimal strategies. If necessary, we use interpolation to determine the controls for a given state value. For the MC computations, we use different timestep sizes and numbers of simulations. See Appendix B for details. We then compare MC-computed means and variances with the corresponding values calculated using the PDE approach of Section 2.7. As an illustrative example, Table 4.3 provides means and standard deviations for both the endowment and pension de-accumulation scenarios for a withdrawal rate of $a = 6$. In all cases it is clear that the MC-computed means and standard deviations converge consistently to the respective PDE-computed values.

[6] An alternative perspective on this strategy is to consider it as a form of debt restructuring. The commitment to withdraw cash on an annual basis for spending purposes has similar effects to those generated by incurring annual interest payments on a loan. The strategy outlined here effectively substitutes long-term zero-coupon debt for the leverage inherent in having yearly withdrawals.

4.3.2 Probability density functions

To avoid notational clutter, we will drop the superscript from W_T^c, with the understanding that all references to W_T in the following refer to the controlled wealth. Figure 4.3 presents plots of the probability density functions of the terminal wealth W_T for several different withdrawal rates a for both the endowment and the pension de-accumulation scenarios. These are obtained using MC simulations with 320 timesteps and 64×10^4 replications, as described in the previous subsection. For brevity, we present just the jump diffusion (panels (a) and (c)) and diffusion with effective volatility cases (panels (b) and (d)). The corresponding γ values used for this experiment are given in Table 4.4.

	Endowment			Pension de-accumulation	
a	Jump diffusion	Diffusion with effective volatility	a	Jump diffusion	Diffusion with effective volatility
4	384.4	419.9	4	158.4	164.5
6	509.8	699.9	7	353.9	435.0
8	1011.1	1332.9	8	510.0	648.5

TABLE 4.4: *Values of γ used to produce Figure 4.3.*

We make the following observations regarding Figure 4.3:

- The shape of the density functions is typically highly skewed. This is due to the optimal control, which attempts to minimize the quadratic loss with respect to the wealth target of $\gamma/2$, as in (2.18). Note that the quadratic wealth target $\gamma/2$ is an increasing function of the withdrawal rate a.

- The shape of the probability density function depends on the withdrawal rate. Note the change of the shape of the density function from single-peaked to double-peaked as a increases, with the second peak centered at a small negative value. This behavior is observed for both the jump diffusion and diffusion with effective volatility cases.[7] When a is small enough (e.g. $a = 4$), the chance of bankruptcy is quite low and so the density has a single peak near W_{spec}. As a increases (e.g. to 6 and 7), the chance of bankruptcy rises. This happens for two reasons: (i) the amounts withdrawn are larger; and (ii) the optimal strategy is to invest more in the risky asset over longer periods of time. Moreover, once bankruptcy occurs, the insolvency condition (2.14) leaves no scope for action: the investor has to liquidate all investments in the risky asset, and is not allowed to make further trades. Subsequent withdrawals are financed by borrowing, but the portfolio remains insolvent. This results in a clustering of values of terminal wealth in a narrow range below zero, resulting in a second peak in the left tail of the W_T distribution.

- Comparing the jump diffusion and diffusion with effective volatility cases, it appears that the jump diffusion setting is less risky as the density function has thinner tails and a higher peak. Note that the compensated drift for the jump diffusion specification is higher than for the pure diffusion case, as indicated by equation (2.2) (recall $\lambda > 0$ and $\kappa < 0$). Consequently, since jumps are comparatively rare, the investor has a higher probability of de-risking as the target is approached, compared with the effective volatility case. Hence, in the jump diffusion

[7] The same effect also occurs in the base case.

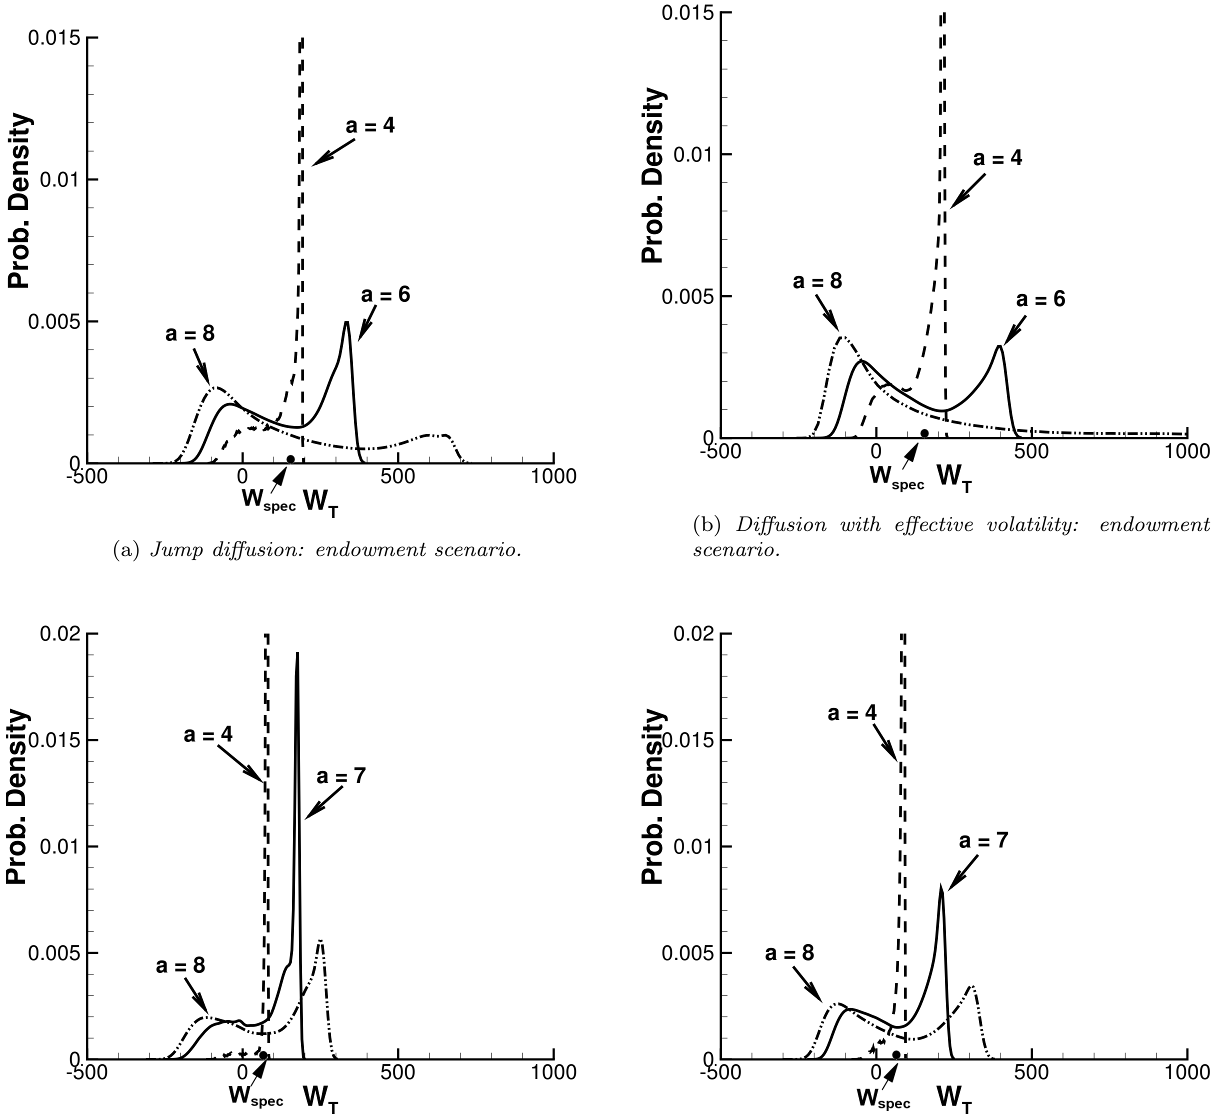

(a) *Jump diffusion: endowment scenario.*

(b) *Diffusion with effective volatility: endowment scenario.*

(c) *Jump diffusion: pension de-accumulation scenario.*

(d) *Diffusion with effective volatility: pension de-accumulation scenario.*

FIGURE 4.3: *Probability density function of W_T for several different withdrawal rates a. For $a = 4$, the density plots are clipped. A total of* 320 *timesteps and* 64×10^4 *replications are used. Other input data are from Tables 4.1 and 4.2.*

scenario, if a downward jump occurs, there is little negative effect on a de-risked portfolio. In contrast, if there is a high effective volatility, it is more difficult (and less likely) for the investor to de-risk. This is a bit counterintuitive, as it suggests that the optimal strategy can overcome sudden market drops more easily than continuous large volatility.

- Comparing the endowment and the pension de-accumulation scenarios for the same a, we observe that the densities for the endowment case have heavier tails and more pronounced left peaks. This follows since for the same a, the expected terminal wealth for the endowment case is larger than for the de-accumulation case. Hence, we expect more risk for the endowment case (i.e. a more spread out density function).

4.3.3 Means and standard deviations of the MV optimal control

To gain further insight into the optimal control strategy, we perform additional MC simulations using the same steps outlined in Section 4.3.1. For each rebalancing time t_m, $m = 0, \dots, M$, we compute the mean and standard deviation of the portion of the portfolio wealth that is invested in the risky asset after an optimal allocation has been applied. That is, we compute the mean and standard deviation of $S_{t_m}/(S_{t_m} + B_{t_m})$. We then plot these two quantities vs. rebalancing times $t_0, \dots, t_M$. Figures 4.4 and 4.5 respectively show illustrative results for the endowment and pension de-accumulation scenarios with $a = \{4, 8\}$.

We make the following observations based on these figures:

- The diffusion with effective volatility cases have the highest average allocation to the risky asset in both the endowment and the pension de-accumulation scenarios (see panels (a) and (c) of Figures 4.4 and 4.5). In addition, in all cases the mean fraction of wealth invested in the risky asset decreases over time. This is because we expect that on average

$$\frac{\gamma}{2} e^{-r(T-t)} + w_k - W_t^c \tag{4.2}$$

  will decrease with increasing time (Vigna, 2014). The intuition here is that on average controlled wealth will get closer to the sum of the discounted final quadratic wealth target and the discounted value of specified cash withdrawals. As indicated above (Section 2.8), if controlled wealth reaches this level, the optimal strategy is to de-risk completely.[8]

- Figure 4.5(a) shows that for the pension de-accumulation scenario with $a = 4$ the fraction invested in the risky asset is, on average, never larger than 0.5 and declines to a relatively low level over time.

- Figure 4.4(c) indicates that for the endowment scenario with $a = 8$, the investor will need to maintain the maximum allowable leverage ratio for quite a long time before switching to investing more in the risk-free asset. In this figure, the mean fraction of wealth invested in the risky asset for all modeling cases starts at $q_{\max}^{(1)} = 1$ and remains there for several years.

4.4 Fixed proportion rules

In this section, we compare the performance of the MV optimal asset allocation strategy to a simple fixed proportion rebalancing rule of Graham (2014), with the constant fraction being p. For this experiment, we proceed as follows:

1. Step 1: carry out MC simulations for the portfolio under this fixed proportion strategy. At each rebalancing time, we adjust the asset allocation so that the constant fraction p of the wealth is invested in the risky asset.

2. Step 2: given the value of expected terminal wealth calculated for the constant proportion strategy in Step 1, we then use PIDE methods to determine the MV optimal strategy which generates the same value for $E^{t_0,x_0}[W_T^c] = W_{\text{spec}}$ with smallest amount of risk, as described earlier.

[8] Of course, if controlled wealth is ever greater than the amount needed to be invested in the risk-free asset which ensures that all remaining withdrawals can be made and the final wealth target can be reached for certain, the excess is a free cash flow.

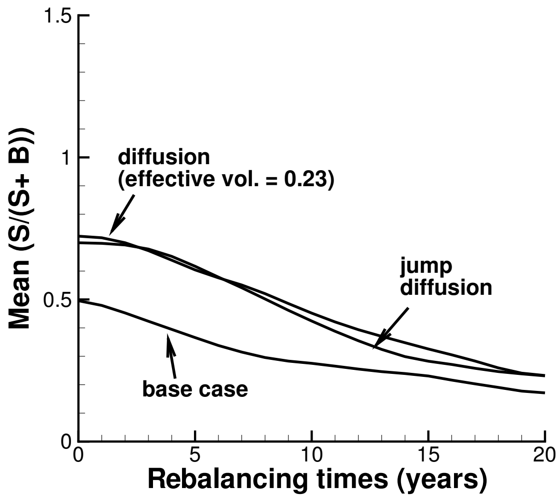

(a) $a = 4$: *endowment scenario; mean.*

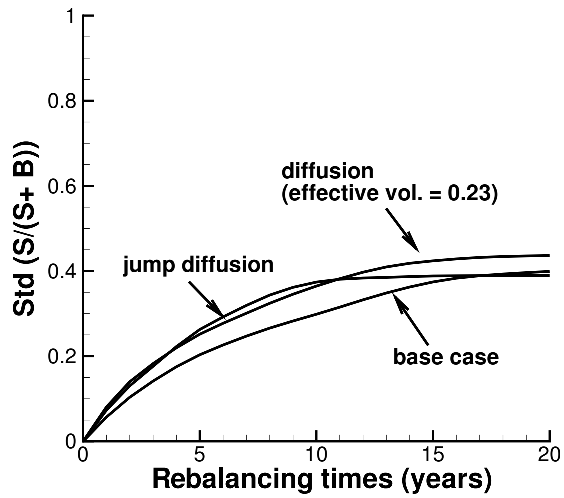

(b) $a = 4$: *endowment scenario; std. dev.*

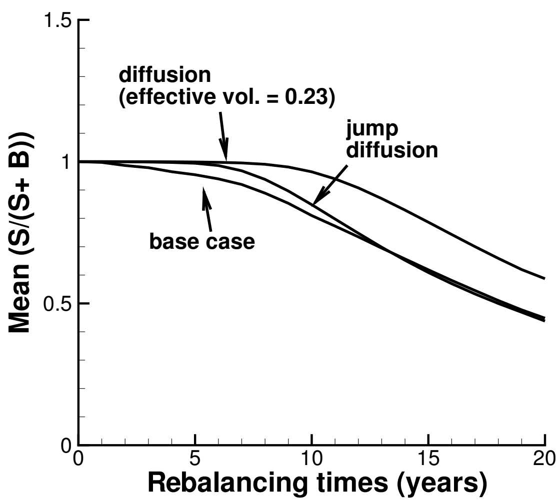

(c) $a = 8$: *endowment scenario; mean.*

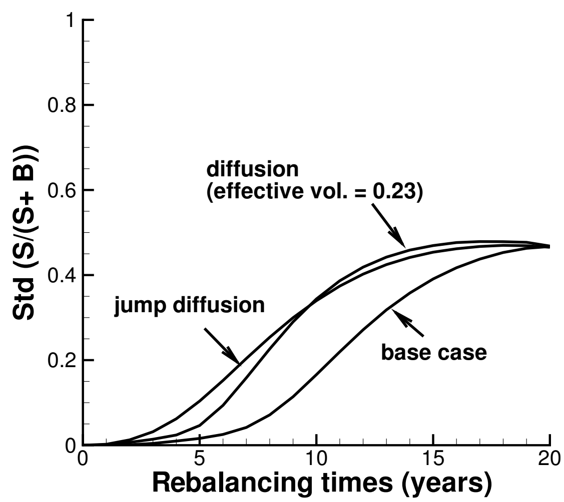

(d) $a = 8$: *endowment scenario; std. dev.*

FIGURE 4.4: *Means and standard deviations of the fraction of wealth invested in the risky asset at each rebalancing time for the endowment scenario. Withdrawal rates are $a = \{4, 8\}$. Other input data are from Tables 4.1 and 4.2.*

We emphasize that we do not specify W_{spec} exogenously for the MV optimal strategy, as in the previous numerical tests. Rather, W_{spec} is determined from the expected terminal wealth of the constant proportion strategy. We compare the MV optimal and the fixed proportion rule strategies in terms of standard deviation and probability of shortfall. For brevity, we only present the results for the pension de-accumulation scenario.[9]

As an illustrative example, consider a fixed proportion rule with $p = 0.5$, $T = \{20, 30\}$ years, $a = 4$ and remaining parameters from Table 4.1. Results for this case are presented in Table 4.5. We note that the value of expected terminal wealth from the constant proportion strategy is approximately

[9]We obtain qualitatively similar results for the endowment scenario.

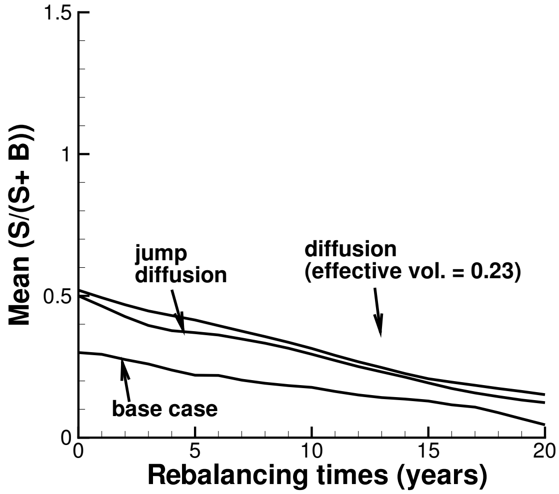

(a) $a = 4$: *pension de-accumulation scenario; mean.*

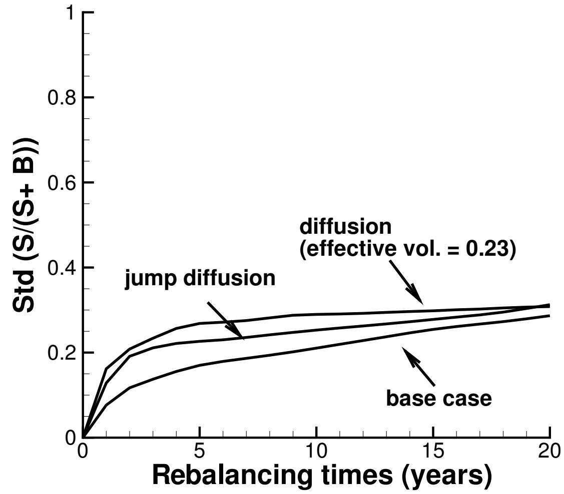

(b) $a = 4$: *pension de-accumulation scenario; std. dev.*

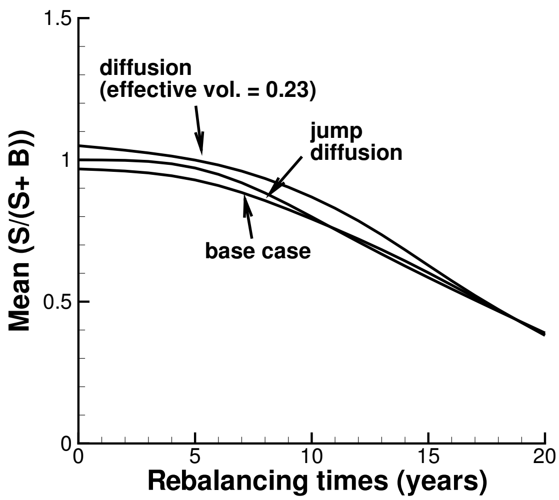

(c) $a = 8$: *pension de-accumulation scenario; mean.*

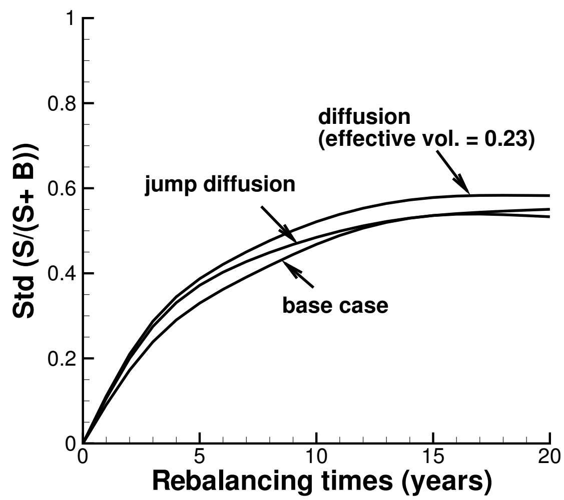

(d) $a = 8$: *pension de-accumulation scenario; std. dev.*

FIGURE 4.5: *Means and standard deviations of the fraction of wealth invested in the risky asset at each rebalancing time for the pension de-accumulation scenario. Withdrawal rates are $a = \{4, 8\}$. Other input data are from Tables 4.1 and 4.2.*

180 and 259 for $T = \{20, 30\}$, respectively.

Table 4.5 shows that for the case with an investment horizon of $T = 20$ years and given the same mean, the MV optimal strategy always results in smaller standard deviations than those obtained under the fixed proportion rule. While this is expected, we emphasize that the differences are quite substantial (e.g. $50.1/98.7 \approx 50\%$ for the base case; $104/158 \approx 65\%$ for the diffusion with effective volatility case). In addition, the shortfall probabilities under the MV optimal strategy are also much reduced compared to the constant proportion policy. When $T = 30$ years, the shortfall probabilities under the MV optimal strategy are less than one-half those for the constant proportion

	Base case		Jump diffusion		Diffusion with effective volatility	
	Fixed prop.	MV optimal	Fixed prop.	MV optimal	Fixed prop.	MV optimal
Investment horizon: $T = 20$ years						
Mean	180.1	180.1	181.2	181.2	180.9	180.9
Std. dev.	98.7	50.1	133.3	91.8	157.8	104.0
Prob. shortfall: $P(W_T < 180)$	0.57	0.20	0.58	0.28	0.62	0.33
Investment horizon: $T = 30$ years						
Mean	258.4	258.4	259.6	259.6	259.1	259.1
Std. dev.	210.4	92.1	287.3	170.1	346.5	195.7
Prob. shortfall: $P(W_T < 250)$	0.58	0.15	0.60	0.26	0.65	0.31

TABLE 4.5: *Comparison of the fixed proportion rule with $p = 0.5$ and the MV optimal strategy for the pension de-accumulation scenario with investment horizons of 20 and 30 years and a withdrawal rate of $a = 4$. Other input data are from Tables 4.1 and 4.2, except that W_{spec} for the MV optimal strategy is determined from the expected terminal wealth under the constant proportion strategy.*

strategy (e.g. $0.15/0.58 \approx 25\%$ for the base case; $0.31/0.65 \approx 45\%$ for the diffusion with effective volatility case). These results demonstrate the superiority of MV optimal strategies, especially for longer investment horizons.

4.5 Historical data tests: robustness to mis-specified parameters

Our previous numerical examples used base case parameters which were approximately equivalent to those estimated by Forsyth and Vetzal (2016) using the last six decades of U.S. market experience. In this subsection, we use parameters determined from longer term historical data to explore the robustness of the investment strategy to parameter estimation.

The parameters for the base case and the jump diffusion case are calibrated to historical data as discussed in Section 7 of Dang and Forsyth (2016). More specifically, we use daily and monthly total return data for the CRSP VWD index. This is the same index from Figure 1.1 above, but extended to cover the period from 1926 through 2014.

The parameters μ and σ for the base case (GBM) can be determined by maximum likelihood estimation. In order to determine the set of parameters for the jump diffusion model, the use of maximum likelihood methods is well-known to be problematic, due to multiple local maxima and the ill-posedness of trying to distinguish high frequency small jumps from diffusion (Honore, 1998). From the perspective of a long-term investor, the most important feature of a jump diffusion model is that it incorporates the effects of infrequent large jumps in asset prices. Small, frequent jumps look like enhanced volatility. When examined on a large scale, these effects are probably not important when building a long-term investment strategy. In calibrating jump diffusion models, we use the thresholding technique described in Mancini (2009) and Cont and Mancini (2011). This technique is considered to be more efficient for low frequency data. The reader is referred to Dang and Forsyth (2016) for details of the calibration techniques.

Since we consider annual rebalancing, we use the average one year T-bill rate over the period from 1934 to 2014, obtained from the U.S. Federal Reserve.[10] For inflation, we use the average CPI

[10]See `www.federalreserve.gov/releases/h15/data.htm`. This data series is only available starting in 1934.

inflation rate from the U.S. Bureau of Labor Statistics over the 1926-2014 period.[11]

The calibrated parameters are given in Table 4.6.[12] The base case parameters are estimated using the daily data. For the jump diffusion case, both the daily and monthly data are used. Note that the drift rates and volatilities are similar for all cases, but the jump parameters are quite different when estimated using daily vs. monthly data. Since the jump parameters are difficult to estimate, we can examine the effect of differing estimates on the investment results.

Parameters	Base case	Jump diffusion	
	Daily	Daily	Monthly
μ (drift)	0.1119	0.1120	0.1122
σ (volatility)	0.1862	0.1631	0.1715
λ (jump intensity)	N/A	1.528	0.0899
ν (jump multiplier mean)	N/A	-0.00759	-0.2631
ζ (jump multiplier std. dev.)	N/A	0.0733	0.0476
r (risk-free interest rate)	0.0499	0.0499	0.0499
I (inflation rate)	0.029	0.029	0.029
W_0 (initial wealth)	100	100	100
T (investment horizon)	30 years	30 years	30 years
$t_{i+1} - t_i$ (rebalance interval)	1 year	1 year	1 year

TABLE 4.6: *Parameters for the empirical data tests for three cases: GBM with daily data and jump diffusion using both daily and monthly returns.*

A comparison of the jump diffusion results for daily and monthly data in Table 4.6 shows some apparent substantial differences. The daily data implies a much higher frequency of jumps, but the jumps are on average of a much smaller magnitude. However, applying formula (4.1) gives effective volatilities of 0.1868 and 0.1893 for the daily and monthly data respectively.

In this test of the method's robustness to mis-specified parameters, we consider the pension de-accumulation case with a withdrawal rate $a = 4\%$. The value for W_{spec} in this case is $100e^{0.029\times 30}/2 \approx 119.2$. For the probabilities of shortfall, we compute $P(W_T < 0.9 \times W_{\text{spec}}) \approx 107$. For the models, we consider only the base case and the jump diffusion case. We proceed as follows:

- Step 1: Assume a model for the risky asset under which MV optimal strategies are computed (e.g. GBM). Under this *strategy computing model*, compute and store the MV optimal strategies from $t = 0$ to $t = T$ for which $\mathrm{E}^{t_0,x_0}[W_T^c] = W_{\text{spec}} = 119.2$ can be achieved with the smallest variance.

- Step 2: Carry out MC simulations for the portfolio from $t = 0$ to $t = T$ following the stored optimal strategies from Step 1, but assuming that the real world's dynamics of the risky asset follow a *different* model (e.g. jump diffusion). This different model is referred to as the *real world model*.

- Step 3: Compare the MC-computed mean, variance, and probability of shortfall for each pair of strategy computing model and real world model.

We begin by noting that across all different strategy computing models in Step 1, the parameter γ for which $\mathrm{E}^{t_0,x_0}[W_T^c] = W_{\text{spec}} = 119.2$ can be achieved with the smallest variance is approximately

[11] In particular, we use the annual average of the all urban consumers index (CPI-U), see `http://www.bls.gov/cpi`.
[12] This table is reproduced from Table 8.1 of Dang and Forsyth (2016).

258. With this value of γ, the smallest variance is about 34.1 across all different strategy computing models.

Table 4.7 shows the results for all combinations of representative test cases. These results clearly demonstrate that the MV optimal strategy results in very similar means, standard deviations, and shortfall probabilities for terminal wealth in all cases. This implies that the MV optimal strategy is quite robust to parameter mis-specification (e.g. if we compute the strategy assuming jump diffusion based on parameter estimates using daily data but run the simulations using a real world model which has parameter estimates based on monthly data).[13]

	Real world model					
Strategy computing model	Mean	Std. dev.	Prob. shortfall $P(W_T < 107)$	Mean	Std. dev.	Prob. shortfall $P(W_T < 107)$
GBM	Jump diffusion (daily)			Jump diffusion (monthly)		
	119.1	34.1	0.26	118.9	33.9	0.26
Jump diffusion (daily)	GBM			Jump diffusion (monthly)		
	118.9	33.9	0.26	119.0	33.8	0.27
Jump diffusion (monthly)	GBM			Jump diffusion (daily)		
	119.2	33.8	0.27	119.1	34.1	0.26

TABLE 4.7: *MC-computed mean and variance for each pair of different strategy computing and real world models. Same level of refinement as in Table 4.3 is used. Input data are provided in Table 4.6. The withdrawal rate is $a = 4\%$, and $\gamma = 258$ for all strategy computing models.*

5 Conclusions

It can be argued that a reasonable long-term model for a stock index uses a jump diffusion process. In this study, we use parameters which generate a jump about once per decade and such that when a jump occurs, the average result is a decline of about 40% in the value of the risky asset. With this jump diffusion model, we investigate the impact of periodic withdrawals on target final real wealth for two scenarios: de-accumulation of a defined contribution pension plan and operation of an endowment. In both scenarios, we use the optimal dynamic asset allocation determined using a multi-period MV objective function. Using either standard deviation or probability of shortfall to measure risk, the optimal strategy considerably outperforms a standard constant proportion strategy.

Under our assumed market parameters, we observe that withdrawal rates of 4% are probably reasonable for the pension de-accumulation scenario. This is, perhaps somewhat surprisingly, consistent with the results in Bengen (1994). However, we emphasize that the results in this paper are based on an optimal asset allocation strategy, not a simple constant mix rule. The optimal asset allocation for this scenario has a fairly low average allocation to risky assets. On the other hand, we find that 4% withdrawals for an endowment are probably not sustainable. This finding might raise some concerns for managers of endowments who are using a 4% withdrawal rate.

We also note that the same general conclusions are obtained if we use either a low risk diffusion model (our base case), or a diffusion model with an effective volatility that matches the total volatility of the jump diffusion model. Although the results are qualitatively similar for these

[13] A different type of robustness test for long-term MV optimal strategies has recently been reported by Ma and Forsyth (2016). They compare a stochastic volatility model with GBM, and find that the two models produce very similar results.

other specifications, we reiterate that the use of effective volatility diffusion model to approximate the jump diffusion model seems to lead to a significant overestimate of risk, at least for the cases considered here. Overall, the MV optimal policy under jump diffusion seems to be quite resilient to relatively rare jumps which on average represent significant market downturns.

Finally, we calibrate our jump diffusion and GBM models to long-term historical data. The jump diffusion parameter estimates are sensitive to the sampling frequency. Nevertheless, MC tests show that the resulting distribution of terminal wealth for the MV optimal strategy is robust to parameter uncertainty. This suggests that our conclusions with respect to withdrawal rates are robust as well.

References

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

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

Björk, T. (2009). *Arbitrage Theory in Continuous Time* (3rd ed.). Oxford University Press.

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

Cui, X., D. Li, S. Wang, and S. Zhu (2012). Better than dynamic mean-variance: time-inconsistency and free cash flow stream. *Mathematical Finance* 22, 346–378.

Dang, D. M. and P. A. Forsyth (2014). Continuous time mean-variance optimal portfolio allocation under jump diffusion: An 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 Y. Li (2016). Convergence of the embedded mean-variance optimal points with discrete sampling. *Numerische Mathematik* 132, 271–302.

Forsyth, P. A. and K. R. Vetzal (2016). Robust asset allocation for long-term target-based investing. Working paper, Cheriton School of Computer Science, University of Waterloo.

Graham, B. (2014). *The Intelligent Investor.* HarperBusiness.

Honore, P. (1998). Pitfalls in estimating jump diffusion models. Working paper, Center for Analytical Finance, University of Aarhus.

Horneff, W., R. Maurer, and R. Rogalia (2010). Dynamic portfolio choice with deferred annuities. *Journal of Banking and Finance* 34, 2652–2664.

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

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

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

Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.

Milevsky, M. and H. Huang (2011). Spending retirement on the planet Vulcan: the impact of longevity risk aversion on optimal withdrawal rates. *Financial Analysts Journal* 67, 45–58.

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

Milevsky, M. A. and T. S. Salisbury (2006). Financial valuation of guaranteed minimum withdrawal benefits. *Insurance: Mathematics and Economics* 38, 21–38.

Navas, J. (2000). On jump diffusion processes for asset returns. Working paper, Instituto de Empresa.

Øksendal, B. and A. Sulem (2009). *Applied Control of Jump Diffusions*. Springer.

Scott, J., W. F. Sharpe, and J. Watson (2009). The 4% rule–at what price? *Journal of Investment Management* 7, 31–48.

Tse, S., P. Forsyth, and Y. Li (2014). Preservation of scalarization optimal points in the embedding technique for continuous time mean variance optimization. *SIAM Journal on Control and Optimization* 52, 1527–1546.

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

Wang, J. and P. Forsyth (2012). Comparison of mean variance like strategies for optimal asset allocation problems. *International Journal of Theoretical and Applied Finance* 15:2. DOI: 10.1142/S0219024912500148.

Wang, J. and P. A. Forsyth (2011). Continuous time mean variance asset allocation: A time-consistent strategy. *European Journal of Operational Research* 209, 184–201.

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

Appendices

A A numerical algorithm to find γ to achieve W_{spec} with smallest variance

Algorithm A.1 describes a Newton's method to find γ for which given a withdrawal rate $\hat{a} > a_{\text{rf}}$, W_{spec} is achieved with the smallest variance:

Algorithm A.1 A Newton algorithm to find a γ value for which, given a withdrawal rate $\hat{a} > a_{\text{rf}}$, W_{spec} is achieved with the smallest variance.

1: $\gamma^{(0)} = \gamma_0$ with γ_0 being an initial guess;
2: **for** $k = 0, 1, \ldots,$ until convergence **do**
3: solve value function problem (2.18) with $\gamma = \gamma^{(k)}$ and $a = \hat{a}$ to obtain the optimal control $c^*_{\gamma^k}(\cdot)$
4: use the control from Line 3 to compute $E(\gamma^{(k)}) = \mathrm{E}^{t_0,x_0}_{c^*_{\gamma^k}(\cdot)}[W_T]$;
5: repeat Lines 3-4 with $\gamma = \gamma^{(k)} + \epsilon$, where $0 < \epsilon \ll 1$, to obtain $E(\gamma^{(k)} + \epsilon)$;
6: compute $E'(\gamma^{(k)}) \approx \dfrac{E(\gamma^{(k)} + \epsilon) - E(\gamma^{(k)})}{\epsilon}$;
7: compute

$$\gamma^{(k+1)} = \gamma^{(k)} - \frac{E(\gamma^{(k)}) - W_{\text{spec}}}{E'(\gamma^{(k)})};$$

8: **if** (converged) **then**
9: break from the iteration;
10: **end if**
11: **end for**
12: return $\gamma = \gamma^{(k+1)}$;

B Monte Carlo algorithm

This appendix provides a Monte Carlo algorithm for simulating the portfolio allocation problem under the jump diffusion model (2.2)-(2.3), assuming an allocation rule. We denote the rule by $\mathfrak{R} \equiv \{\mathfrak{R}^m\}_{m=0}^{M}$. Here, $\mathfrak{R}^m$ is the allocation rule for the rebalancing time t_m. Each of $\mathfrak{R}^m, m = 0, \ldots, M$, can be expressed in the form

$$\mathfrak{R}^m \equiv \left\{ (s_k, b_l, p^m_{k,l}) \right\}, \quad k = 1, \ldots, k_{\max}, \quad l = 1, \ldots, l_{\max}, \tag{B.1}$$

where s_k and b_l respectively denote the PDE grid point values in the s and b directions, and $p^m_{k,l}$ denotes the time-t_m optimal proportion of the portfolio wealth invested in the risky asset if the stock and bond amounts are s_k and b_l, respectively. Note that the $p^m_{k,l}$ are computed during the solution process of the PDE method.

We denote by $\mathrm{S}^{m,-}$ and $\mathrm{B}^{m,-}$ simulated values of S and B at the rebalancing time t^-_m, $m = 0, \ldots, M$, after the withdrawal a_m has been made. At t^-_m, if the portfolio is still solvent it is then rebalanced according to the rule $\mathfrak{R}^m$. Since $\mathrm{S}^{m,-}$ and $\mathrm{B}^{m,-}$ may not be exactly s_k, b_l, for some k and l, we employ $\mathfrak{R}^m$ and linear interpolation along the s and b directions to compute an optimal allocation of the portfolio. On the other hand, for the fixed proportion rule with parameter p, we simply have

$$\mathrm{S}^m = (\mathrm{S}^{m,-} + \mathrm{B}^{m,-})p; \quad \mathrm{B}^m = (\mathrm{S}^{m,-} + \mathrm{B}^{m,-})(1-p). \tag{B.2}$$

An MC simulation of the portfolio allocation problem under the jump diffusion model (2.2)-(2.3) is given in Algorithm B.1. In the algorithm, N is the number of timesteps, and I is the number of replications. We denote by $\mathrm{S}_i^{\cdot,-}$, $\mathrm{B}_i^{\cdot,-}$, $i = 1, \ldots, I$, the i-th replication of $\mathrm{S}^{\cdot,-}$ and $\mathrm{B}^{\cdot,-}$, respectively. Also, $\mathbb{I}_A$ denotes an indicator function.

Algorithm B.1 A Monte-Carlo algorithm for simulating the portfolio allocation problem under the jump diffusion model (2.2)-(2.3).

compute S_0 and B_0 using linear interpolation or (B.2);
set $\mathrm{B}_i^0 = B_0$, $\mathrm{S}_i^0 = S_0$, $\log\mathrm{S}_i^0 = \log(S_0)$, $\mathrm{AlreadyLiquidated}_i = 0$; $i = 1, 2, \ldots, I$;
set $dt = T/N$;
for $n = 1, 2, \ldots, N$ **do** {Timestep loop}
  set $t_n = ndt$;
  for $i = 1, 2, \ldots, I$ **do** {Simulation loop}
    set $\mathrm{B}_i^{n,-} = \mathrm{B}_i^{n-1} \mathrm{e}^{rdt} - \sum_{m=1}^{M} a_m \mathbb{I}_{t_n = t_m, t_m \in \mathcal{T}}$; {interest and withdrawal}
    if $\mathrm{AlreadyLiquidated}_i = 1$ **then**
      $\mathrm{S}_i^n = 0$; $\log\mathrm{S}_i^n = -\infty$; $\mathrm{B}_i^n = \mathrm{B}_i^{n,-}$;
    else
      generate $K \sim \mathrm{Poisson}(\lambda dt)$;
      set $\log\mathrm{S}_i^{n,-} = \log\mathrm{S}_i^{n-1} + (\mu - \lambda\kappa + \sigma^2/2)dt + \sigma\sqrt{dt}\mathrm{Normal}(0,1) + \nu K + \xi\sqrt{K}\mathrm{Normal}(0,1)$;
      if $\mathrm{S}_i^{n,-}, \mathrm{B}_i^{n,-} \in \mathcal{B}$ **then**
        set $\mathrm{B}_i^n = \mathrm{S}_i^{n,-} + \mathrm{B}_i^{n,-}$, $\mathrm{S}_i^n = 0$, $\log\mathrm{S}_i^n = -\infty$, $\mathrm{AlreadyLiquidated}_i = 1$;
        {liquidate the portfolio}
      end if
      set $\mathrm{S}_i^{n,-} = \mathrm{e}^{\log\mathrm{S}_i^{n,-}}$;
      if ($(t_n \in \mathcal{T})$ and $(\mathrm{AlreadyLiquidated}_i \neq 1)$) **then**
        compute S_i^n and B_i^n using interpolation or (B.2);
        {rebalance the portfolio}
        if $\dfrac{\mathrm{S}_i^n}{\mathrm{S}_i^n + \mathrm{B}_i^n} > q$ **then**
          let $W_i^n = \mathrm{S}_i^n + \mathrm{B}_i^n$; {enforce leverage cond.}
          set $\mathrm{S}_i^n = qW_i^n$ and $\mathrm{B}_i^n = (1-q)W_i^n$;
        end if
      else
        set $\mathrm{S}_i^n = \mathrm{S}_i^{n,-}$, and $\mathrm{B}_i^n = \mathrm{B}_i^{n,-}$; {not a rebalancing time or liquidated}
      end if
      set $\log\mathrm{S}_i^n = \log(\mathrm{S}_i^n)$;
    end if
  end for{End Simulation loop}
end for{End Timestep loop}
set $\mathrm{Portfolio}_i = \mathrm{S}_i^N + \mathrm{B}_i^N$, $i = 1, \ldots, I$;
return E(Portfolio) and Var(Portfolio);