Dynamic Mean Variance Asset Allocation: Tests for Robustness*

Peter A. Forsyth† Kenneth R. Vetzal‡

Draft version: May 30, 2017

Abstract

We consider a portfolio consisting of a risk-free bond and an equity index which follows a jump diffusion process. Parameters for the inflation-adjusted return of the stock index and the risk free bond are determined by examining 89 years of data. The optimal dynamic asset allocation strategy for a long-term pre-commitment mean variance (MV) investor is determined by numerically solving a Hamilton-Jacobi-Bellman partial integro differential equation. The MV strategy is mathematically equivalent to minimizing the quadratic shortfall of the target terminal wealth. We incorporate realistic constraints on the strategy: discrete rebalancing (yearly), maximum leverage, and no trading if insolvent. Extensive synthetic market tests and resampled backtests of historical data indicate that the multi-period MV strategy achieves approximately the same expected terminal wealth as a constant weight strategy, but with much smaller variance and probability of shortfall.

Keywords: mean-variance, dynamic asset allocation, resampled backtests

1 Introduction

Due to the disappearance of traditional defined benefit pension plans, individual investors must now manage their own defined contribution (DC) retirement savings plans. Accumulation periods of thirty or more years are not uncommon, followed by a decumulation period of another twenty to thirty years. Retail investors are thus faced with a fifty year investment cycle, making them truly long-term investors. It is therefore of interest to determine optimal dynamic asset allocation strategies for such long-term investors.

A strategy which is implementable by individual investors should have the following characteristics:

- The portfolio should consist of a small number of widely available components, such as a low-risk bond portfolio and an equity index.
- The strategy should be based on a small number of parameters, which are relatively stable over the long term.
- The strategy should incorporate realistic solvency and leverage constraints.

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

†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.

- The strategy should require infrequent (e.g. yearly) calendar-based rebalancing.

Based on the embedding approach devised in Li and Ng (2000) and Zhou and Li (2000), it is known that pre-commitment multi-period MV optimization is equivalent to a quadratic target-based strategy (Vigna, 2014; Dang and Forsyth, 2016). Using the numerical methods in Dang and Forsyth (2014), a recent study by Forsyth and Vetzal (2017a) showed that the pre-commitment MV strategy was able to produce the same expected terminal wealth as a constant weight strategy, but with greatly reduced variance and probability of shortfall. This paper made the assumption that the risky asset followed geometric Brownian motion (GBM).

In (Graf, 2017) (GBM process) and (Forsyth and Vetzal, 2017b) (jump diffusion process), in the case of a lump sum investment with continuous rebalancing, it is shown that *any* deterministic glide path strategy is inferior to a constant weight strategy by the MV criterion. In this case, the optimal constant weight is simply the average over time of the weight of the glide path strategy.

For the more realistic case of periodic contribution and discrete rebalancing, a constant weight strategy is only slightly inferior to the optimal deterministic glide path, with the same expected terminal wealth. For example, Forsyth and Vetzal (2017b) show that for realistic market parameters and long term (greater than 20 years) investments, the ratio of the standard deviation of the optimal deterministic strategy to the standard deviation of the constant proportion strategy (same expected terminal wealth) is in the range $0.97 - 0.99$. In practice, most glide path strategies are derived based on heuristic reasoning and are unlikely to be optimal. This result is confirmed by empirical studies such as Esch and Michaud (2014). Deterministic glide path strategies are commonplace in target date funds. Consequently, in this study, we focus on comparing optimal MV strategies with constant weight strategies.

The contributions of this study are as follows:

- To investigate the effects of alternative assumptions to GBM for the equity index, we model the stock index as a jump diffusion, where the jump size follows a double exponential distribution. We determine the parameters for the model by examining a long history (89 years) of U.S. stock market data.[1]

- Our objective function is based on attempting to achieve a specified real (i.e. inflation-adjusted) target wealth with a small probability of shortfall.

- For a given strategy, computed using the methods in Dang and Forsyth (2014), we carry out extensive Monte Carlo tests in a synthetic market. To test robustness, we assume that the synthetic market parameters used in the Monte Carlo tests are different from the parameters used to generate the strategy.

- Based on strategies generated using parameters estimated (using various methods) from long term data, we test these strategies on historical data. Representative historical paths are examined. Bootstrap resampling (Annaert et al., 2009; Bertrand and Prigent, 2011; Cogneau and Zakalmouline, 2013) tests of the historical data are also carried out.

- In addition to the lump sum investment case discussed in Forsyth and Vetzal (2017a), we also consider the accumulation and decumulation phases, which would perhaps be more typical of a DC investor.

[1] An alternative potential extension to GBM would be to incorporate stochastic volatility. However, the tests in Ma and Forsyth (2016) indicate that stochastic volatility effects are unimportant for long-term (10 years or more) investors.

Our results confirm that the pre-commitment multi-period MV strategy is surprisingly robust. It is quite insensitive to parameter estimates and it significantly outperforms a constant weight strategy, in both synthetic markets and bootstrap resampled historical markets. We remind the reader that constant weight strategies outperform deterministic glide path strategies (lump sum investment), or are similar in performance to an optimal glide path strategy (periodic contributions). Consequently, our results should be of particular interest to those investors currently holding target date funds using a glide path strategy[2].

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. The investment horizon is T. As described in more detail below, we assume that the portfolio is rebalanced at discrete points in time in accordance with a control strategy. We denote by S_t and B_t the *amounts* invested in the risky and risk-free assets respectively at time t, $t \in [0, T]$. In the absence of control (i.e. between rebalancing times), S follows the jump diffusion process

$$\frac{dS_t}{S_{t^-}} = (\mu - \lambda\kappa)\ dt + \sigma\, dZ + d\left(\sum_{i=1}^{\pi_t} (\xi_i - 1)\right) \tag{2.1}$$

where μ is the (uncompensated) drift rate, σ is the volatility, dZ is the increment of a Wiener process, π_t is a Poisson process with intensity λ, and ξ denotes the random number representing the jump multiplier. When a jump occurs, we have $S_t = \xi S_{t^-}$, and $\kappa = E[\xi - 1]$ where $E[\cdot]$ is the expectation operator. As a specific example in this paper, we consider ξ_i as being i.i.d. positive random variables characterized by a double exponential distribution (Kou and Wang, 2004). The density function for $y = \log(\xi)$ is then

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

Given that a jump occurs, p is the probability of an upward jump and $(1-p)$ is the probability of a downward jump. Note that

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

It is assumed that the dynamics of the amount invested in the risk-free asset B (in the absence of control) are

$$dB_t = rB_t\, dt,$$

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

2.1 Rebalancing Times: Scheduled Injections/Withdrawals

We consider a set of pre-determined *rebalancing times*, denoted by $\mathcal{T}$,

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

[2]According to Morningstar, at the end of 2015, there was over USD 750 billion invested in target date funds in the U.S.

Let $t_0 = 0$ be the inception time of the investment. For simplicity, we specify the set of rebalancing times (2.4) to be equidistant with $t_m - t_{m-1} = \Delta t = T/M$, $m = 1, \ldots, M$. At each rebalancing time t_m, $m = 0, \ldots, M$, the investor (i) withdraws an amount of cash, denoted by a_m, from the risk-free asset, and then (ii) rebalances the portfolio. Note that we have defined a withdrawal as $a_m > 0$. A cash injection has $a_m < 0$. In the case of a lump-sum investment with no further cash injections/withdrawals, we have $a_m = 0$. For a full lifecycle case, we would have both injections and withdrawals.

2.2 Allowable Portfolios

Let $X_t = (S_t, B_t)$, $t \in [0, T]$, be the multi-dimensional (controlled) underlying process, and let $x = (s, b)$ denote the state of the system. At each rebalancing date, the investor chooses a control $c(s,b)$ which is a function of the current portfolio state (s,b). The control generates a new allocation of the stock and bond (this is defined more precisely in Appendix A).

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

$$W_t \equiv W(S_t, B_t) = S_t + B_t.$$

We strictly enforce a solvency condition, i.e. the investor can continue trading only if

$$W(s, b) = s + b > 0. \tag{2.5}$$

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

$$S = 0; \quad B = W(s,b); \quad \text{if } W(s, b) \leq 0.$$

We also constrain the leverage ratio, i.e. the investor must select an allocation satisfying

$$\frac{S}{S + B} \leq q_{\max}, \tag{2.6}$$

where $q_{\max}$ is a specified positive constant. Note that setting $q_{\max} = 1$ implies that $B \geq 0$, thus ruling out the use of any leverage.

2.3 Informal Derivation of the Objective Function

In order to give some intuition for the embedded MV optimal objective function, we will consider a simple case in this section. Suppose the risky asset follows pure GBM (i.e. $\lambda = 0$ in equation (2.1)), and that the portfolio is continuously rebalanced. Let $E^t_{c(\cdot)}[W_T]$ and $\text{Var}^t_{c(\cdot)}[W_T]$ denote the mean and variance of terminal wealth W_T as observed at time t. Now, consider an objective function of the form

$$\min_{c(\cdot)} \left\{ E^{t_0}_{c(\cdot)} \left[(W^* - W_T)^2 \right] \right\}. \tag{2.7}$$

In other words, we want to find the rebalancing strategy $c(\cdot)$ that minimizes the quadratic loss with respect to the quadratic wealth target W^*. To avoid trivial cases, we assume that the initial wealth $W_{t_0} < W^* e^{-r(T-t)}$.

It is clear that equation (2.7) penalizes shortfall, i.e. the amount by which the terminal wealth is less than W^*. It also seems at first glance to penalize terminal wealth values larger than W^*. However, if there are no jumps and we use continuous rebalancing, then Vigna (2014) proves that the optimal strategy $c(\cdot)$ which minimizes (2.7) is such that $W_t \leq W^*$ (assuming of course that

$W_{t_0} < W^* \mathrm{e}^{-r(T-t_0)}$). In other words, the wealth target is always approached from below. Hence, the optimal policy only penalizes losses with respect to the target.

The issue which then immediately arises is how to select W^*. A natural choice might seem to be $W^* = \mathrm{E}^{t_0}_{c(\cdot)}[W_T]$, so that we are minimizing the semi-variance. However, since $W_T \leq W^*$, then clearly $\mathrm{E}^{t_0}_{c(\cdot)}[W_T] < W^*$. An alternative method is to determine W^* such that $\mathrm{E}^{t_0}_{c(\cdot)}[W_T] = d$, with d specified by the investor. In other words, the investor is specifying a desired expected terminal wealth, but minimizing the quadratic loss w.r.t. $W^* > \mathrm{E}^{t_0}_{c(\cdot)}[W_T]$. In this case, our strategy aims at W^*, which has a built-in *safety factor* $(W^* - \mathrm{E}^{t_0}_{c(\cdot)}[W_T])$ to ensure that we obtain the specified expected terminal wealth on average. We can then formulate the optimization problem for $c(\cdot)$ as

$$\min_{c(\cdot)} \left\{ \mathrm{E}^{t_0}_{c(\cdot)} \left[(W^* - W_T)^2 \right] \right\} \tag{2.8}$$

where W^* is determined from the constraint $\mathrm{E}^{t_0}[W_T] = d$, with d specified by the investor.

It is now fascinating to observe from Li and Ng (2000) and Zhou and Li (2000) that the control which solves problem (2.8) is also the control which solves

$$\begin{aligned} &\min_{c(\cdot)} \left\{ \rho \mathrm{Var}^{t_0}_{c(\cdot)}[W_T] - \mathrm{E}^{t_0}_{c(\cdot)}[W_T] \right\} \\ &\text{where } \rho = \frac{1}{2(W^* - \mathrm{E}^{t_0}_{c(\cdot)}[W_T])} = \frac{1}{2(W^* - d)}. \end{aligned} \tag{2.9}$$

This means that the optimal control $c^*(\cdot)$ which minimizes the quadratic loss with respect to the quadratic target wealth W^* is also the control $c^*(\cdot)$ which solves the dynamic MV problem (2.9). For a discussion of the intuition underlying the identity of problems (2.8) and (2.9), see Forsyth and Vetzal (2017a). In essence, the optimal control $c^*(\cdot)$ is simultaneously MV optimal and quadratic loss optimal. This makes this choice of objective function potentially very desirable.

2.4 Efficient Frontiers and Embedding Methods

Having supplied some basic intuition in the prior section which assumed GBM, we return here to the jump diffusion context and proceed somewhat more formally. We start by describing how MV efficient frontiers can be determined. Let $\mathrm{E}^{x,t}_{c(\cdot)}[W_T]$ and $\mathrm{Var}^{x,t}_{c(\cdot)}[W_T]$ respectively denote the expectation and the variance of the terminal wealth W_T conditional on the state (x,t) and on the control $c(\cdot)$, $t \in \mathcal{T}$. We denote the initial state by $(x_0, t_0) = (X_0, t = 0)$. Then the *achievable MV objective set* $\mathcal{Y}$ is

$$\mathcal{Y} = \left\{ \left(\mathrm{Var}^{x_0,t_0}_{c(\cdot)}[W_T],\ \mathrm{E}^{x_0,t_0}_{c(\cdot)}[W_T] \right) : c \in \mathcal{Z} \right\}, \tag{2.10}$$

where $\mathcal{Z}$ is the set of admissible controls, reflecting the constraints discussed above in Section 2.2. 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

$$\mathcal{Y}_{P(\rho)} = \left\{ (\mathcal{V}_*, \mathcal{E}_*) \in \bar{\mathcal{Y}} : \rho \mathcal{V}_* - \mathcal{E}_* = \inf_{(\mathcal{V},\mathcal{E}) \in \mathcal{Y}} \rho \mathcal{V} - \mathcal{E} \right\}. \tag{2.11}$$

Here, $\bar{\mathcal{Y}}$ is the closure of $\mathcal{Y}$ and ρ can be interpreted as a risk-aversion parameter that reflects how the investor trades off expected return (reward) and variance (risk). For a given ρ, $\mathcal{Y}_{P(\rho)}$ represents Pareto efficient points: given the variance of any point in $\mathcal{Y}_{P(\rho)}$, the corresponding expectation is the highest expectation obtainable with that variance. The set of points on the efficient frontier, $\mathcal{Y}_P$, is the collection of efficient points for all values of $\rho > 0$, i.e.

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

Of course, a primary objective in MV optimal asset allocation 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 (2.11) 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, define the value function as

$$V(x,t) = \min_{c(\cdot)\,\in\,\mathcal{Z}} \left\{ \mathrm{E}^{x,t}_{c(\cdot)} \left[(W^* - W_T)^2 \right] \right\}, \tag{2.12}$$

where the parameter $W^* \in (-\infty, +\infty)$. The embedding result implies that there exists a $W^* \equiv W^*(x,t,\rho)$, such that for a given positive ρ, a control c^* which minimizes the objective function of (2.11) also minimizes that of (2.12). Based on the intuitive description in the simpler context of Section 2.3, we can identify W^* as the *quadratic wealth target*, and

$$\mathrm{E}^{x,t}_{c(\cdot)}[W_T]$$

as the *expected wealth target*. Following Zhou and Li (2000) we can connect these two quantities and the risk-aversion parameter ρ using

$$W^* = \frac{1}{2\rho} + \mathrm{E}^{x,t}_{c(\cdot)}[W_T],$$

which is just equation (2.9) expressed in an alternative form.

Up to this point, there may not seem to be much at all different between the discussion in Section 2.3 and that provided here. There is, however, a key difference. In Section 2.3 it was assumed that the risky asset evolved according to GBM and that the portfolio was continuously rebalanced. Effectively, this was a complete market, and this drives the result that the target will always be approached from below. If the market is incomplete, then it may be optimal to withdraw additional cash from the portfolio, over and above beyond the scheduled withdrawals described in Section 2.1 (Cui et al., 2012; Baeuerle and Grether, 2015; Dang and Forsyth, 2016). Suppose that we are at a rebalancing date t_k and that the quadratic wealth target W^* can be hit exactly by investing an amount $\hat{W}_k$ in the risk-free asset.[3] If the wealth of the portfolio at time t_k, W_{t_k}, exceeds $\hat{W}_k$, then the optimal strategy will be to

(i) withdraw $(W_{t_k} - \hat{W}_k)$ from the portfolio, and

(ii) invest the remaining $\hat{W}_k$ in the risk-free asset.

We refer to the amount withdrawn (and accumulated interest earned on it over the remainder of the investment horizon) as the *free cash flow*. This cash withdrawal is clearly MV optimal because the quadratic wealth target will be hit for certain, and the free cash flow is a bonus.

Appendix A describes the formulation of problem (2.12) as the solution of a Hamilton-Jacobi-Bellman (HJB) equation. We solve the HJB equation numerically, following Dang and Forsyth (2014). At each rebalancing time, we determine the optimal control for the value function (2.12) using the strategy discussed in Dang and Forsyth (2016). This strategy explicitly takes into account the possibility of the bonus free cash flow.

[3] A precise definition of $\hat{W}_k$ is given in equation (A.5) in Appendix A.

3 Parameters

To estimate the parameters of equations (2.1)-(2.2), we use the following data:

- CPI data from the U.S. Bureau of Labor Statistics.[4]
- The total monthly returns (including dividends and any other distributions) of an index of all domestic stocks traded on major U.S. exchanges for the *1926:1-2014:12* period. This is the VWD index from the Center for Research in Security Prices (CRSP).[5].
- Short-term (3-month) U.S. Treasury bill rates. For the period 1934-2014, this is from the U.S. Federal Reserve.[6] This series is not available prior to 1934. In order to include the effects of the 1929 stock market crash in our sample period, we use data on short-term government bond yields from the National Bureau of Economic Research for the 1926-1933 time frame.[7]

From the inputs listed above, real (inflation-adjusted) total equity and T-bill return indexes were constructed on a monthly basis for *1926:1 - 2014:12*. As our model assumes that the risk-free rate is constant, we just calculate the sample average monthly return from our real T-bill index. On an annualized basis, this turns out to be 0.6299%.

We next discuss the estimation of the parameters of the jump diffusion process (2.1)-(2.2). Consider a discrete series of index prices $S(t_i) = S_i, i = 1, \dots, N+1$ that are observed at equally spaced time intervals $\Delta t = t_{i+1} - t_i, \forall i$, with $T = N\Delta t$.[8] Let log returns be

$$\Delta X_i = \log\left(\frac{S_{i+1}}{S_i}\right), \tag{3.1}$$

and define the detrended log returns $\Delta \hat{X}_i$ as

$$\begin{aligned} \Delta \hat{X}_i &= \Delta X_i - \hat{m}\, \Delta t \\ \hat{m} &= \frac{\log(S_{N+1}) - \log(S_1)}{T}. \end{aligned} \tag{3.2}$$

Figure 3.1 shows a histogram of the monthly log returns from the real total return equity index, scaled to unit standard deviation and zero mean. We superimpose a standard normal density onto this histogram. The plot shows that the empirical data has a higher peak and fatter tails than a normal distribution, consistent with previous empirical findings for virtually all financial time series. This figure strongly suggests that it is worth investigating jump diffusion models, rather than relying solely on GBM.

Maximum likelihood (ML) estimation is easily applied to the case of GBM, i.e. if $\lambda = 0$ in (2.1).[9] However, estimating the parameters for the full jump diffusion model is more challenging. Use of ML estimation 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). Various methods

[4] In particular, we use the annual average of the all urban consumers (CPI-U) index. See `http://www.bls.gov/cpi`.

[5] See `http://www.crsp.com`. More specifically, results presented here were calculated based on data from Historical Indexes, ©2015 Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business. Wharton Research Data Services was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

[6] See `http://research.stlouisfed.org/fred2/series/TB3MS`.

[7] See `http://www.nber.org/databases/macrohistory/contents/chapter13.html`.

[8] We assume equal spacing for ease of exposition.

[9] Using the real stock index return over the full sample period from 1926-2014 gives ML estimates of $\hat{\mu} = 0.0816$ and $\hat{\sigma} = 0.1863$ when we impose $\lambda = 0$.

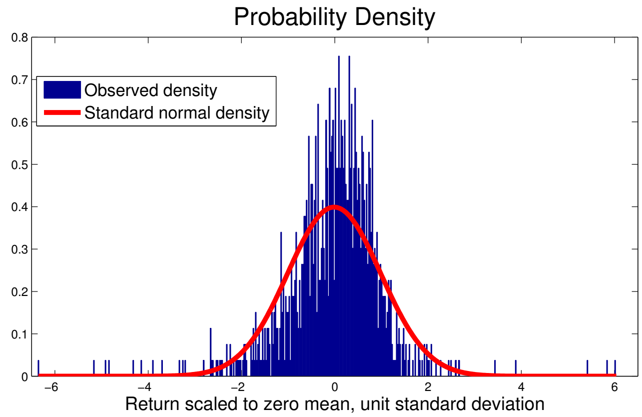

Figure 3.1: *Probability density of* log *returns for real CRSP VWD index. Monthly data, 1926-2014, scaled to unit standard deviation and zero mean. Standard normal density also shown.*

have been suggested to alleviate these problems (see, e.g. Hanson and Zhu, 2004; Synowiec, 2008). However, these methods typically involve fixing some of the parameters, which are estimated using other techniques. Since we are interested in the robustness of our investment strategy, we include a pure ML estimate of the jump diffusion parameters, knowing full well that it is unreliable. Further details about ML estimation are given in Appendix B.1.

Econometric techniques have been developed for detecting the presence of jumps in high frequency data, i.e. for a time scale of seconds (Aït-Sahalia and Jacod, 2012). However, from the perspective of a long-term investor, the most important feature of a jump diffusion model is that it allows modelling of infrequent large jumps in asset prices. Small and frequent jumps look like enhanced volatility when examined on a large scale, hence these effects are probably insignificant when constructing a long-term investment strategy. Consequently, as an alternative to ML estimation, we use the thresholding technique described in Mancini (2009) and Cont and Mancini (2011). This procedure is considered to be more efficient for low frequency data.

Suppose we have an estimate for the diffusive volatility component $\hat{\sigma}$. Then we detect a jump in period i if

$$\left|\Delta \hat{X}_i\right| > \mathcal{A}\,\hat{\sigma}\,\frac{\sqrt{\Delta t}}{(\Delta t)^{\beta}} \tag{3.3}$$

where $\beta, \mathcal{A} > 0$ are tuning parameters (Shimizu, 2013). The intuition behind equation (3.3) is simple. If we choose $\mathcal{A} = 4$, say, and $\beta \ll 1$, then equation (3.3) identifies an observation as a jump if the observed log return exceeds a 4 standard deviation Brownian motion change, which would be extremely improbable. If we choose a smaller time interval, i.e. reduce Δt, keeping the total time T fixed, then equation (3.3) indicates that we should raise the threshold which filters out the Brownian motion increments. Intuitively, this is because we have increased the number of samples and we expect to observe some large deviation events purely by chance. Typically, β in equation (3.3) is quite small, $\beta \simeq .01 - .02$. We provide a detailed discussion of how we calculate thresholding estimates in Appendix B.2. As described there, we replace $\mathcal{A}/(\Delta t)^{\beta}$ by the single parameter α.

Estimated parameters using various values of α and also ML are provided in Table 3.1. As might be expected, some of these estimates vary considerably across α, but the ML estimates are notably different in most cases. Note, however, that all of the values of $\hat{p} < 0.5$, indicating that downward jumps are more likely than upward jumps.

Method	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\lambda}$	$\hat{p}$	$\hat{\eta}_1$	$\hat{\eta}_2$
ML	0.0823	0.1111	5.1162	0.0946	13.8325	21.9910
Threshold ($\alpha = 2$)	0.0896	0.0970	2.3483	0.4258	11.2321	10.1256
Threshold ($\alpha = 3$)	0.0874	0.1452	0.3483	0.2903	4.7941	5.4349
Threshold ($\alpha = 4$)	0.0866	0.1584	0.1461	0.3846	3.7721	3.9943

Table 3.1: *Estimated parameters for double exponential jump diffusion model. Monthly real log total returns, CRSP VWD index. Sample period 1926-2014.*

Investment horizon	30 years
Market parameters	Threshold($\alpha = 3$, see Table 3.1)
Risk-free rate r	0.00623
Initial investment W_0	100
Rebalancing interval	1 year
Maximum leverage ratio $q_{\max}$	1.5
If insolvent	Trading stops

Table 4.1: *Base case data.*

Figure 3.2 shows the normalized probability densities for the estimated parameters given in Table 3.1. Panel (a) displays the results based on the ML and threshold ($\alpha = 2$) estimates, while panel (b) is derived from the threshold estimates when $\alpha = 3, 4$. In each case, a standard normal density is also shown, and the high peaks of the densities for the jump diffusion model relative to the standard normal case are readily apparent. Of course, the jump diffusion model also has fat tails relative to the normal distribution. These are not as visible in panels (a) and (b), but panel (c) provides a zoomed-in view of panel (a) which displays the fat tails much more clearly.

4 Base Case: Lump Sum Investment in a Synthetic Market

As a first example, we consider the investment problem where we assume the market follows the stochastic process (2.1), with the parameters given by the threshold ($\alpha = 3$) case from Table 3.1. Other base case investment parameters are given in Table 4.1.

We start by taking the parameters from Table 4.1 and simulating a constant proportion strategy, whereby the fraction of wealth invested in the stock index p is reset to 0.5 every year. We then determine the value of W^* for the MV optimal strategy which generates the same $E[W_T]$ as the constant proportion strategy. Some summary statistics for these strategies are shown in Table 4.2. The MV optimal strategy standard deviation is about 2.6 times smaller than the constant proportion strategy. Probabilities of shortfall for the MV strategy are also several times smaller than for the constant proportion strategy. Note that the last row of Table 4.2 shows results when $q_{\max} = 1$, which implies that leverage is prohibited. Ruling out leverage results in slightly worse performance compared to the $q_{\max} = 1.5$ case, but still markedly better performance compared to the constant proportion case. The magnitude of the expected free cash bonus is relatively small for the two MV optimal cases (and obviously zero for the constant proportion strategy).

Figure 4.1 plots cumulative distribution functions for the internal rate of return (IRR) for the MV optimal strategy compared to the constant proportion benchmark. Note that the free cash

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 300)$	$Pr(W_T < 400)$	Expected Free Cash
Constant proportion ($p = 0.5$)	417	299	0.41	0.60	0.0
MV optimal ($q_{\max} = 1.5$)	417	117	0.13	0.22	10.5
MV optimal ($q_{\max} = 1.0$)	417	128	0.18	0.26	9.4

TABLE 4.2: *Investment results for a constant proportion strategy and two MV optimal strategies with different leverage constraints. Based on input data from Table 4.1.*

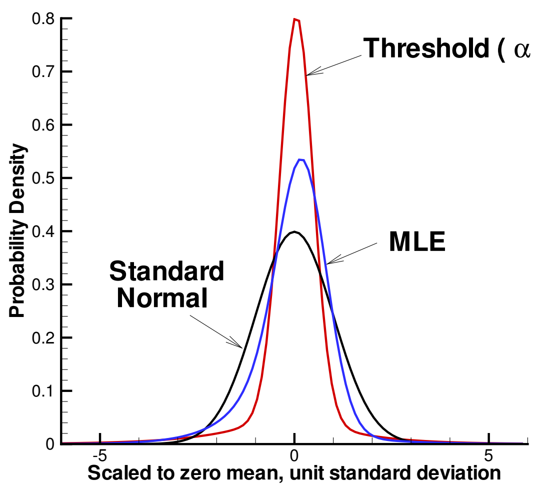

(a) *ML and threshold ($\alpha = 2$).*

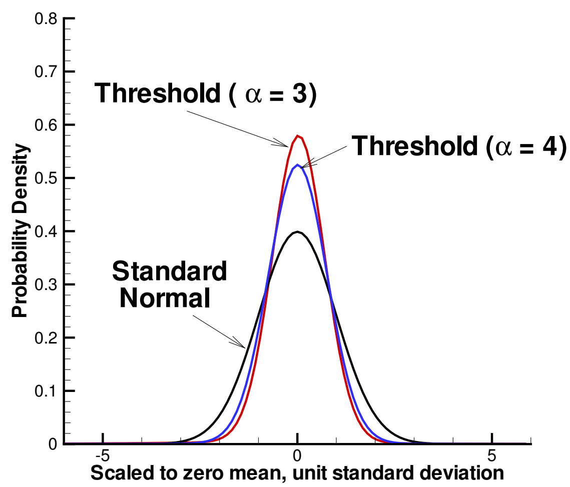

(b) *Threshold ($\alpha = 3, 4$).*

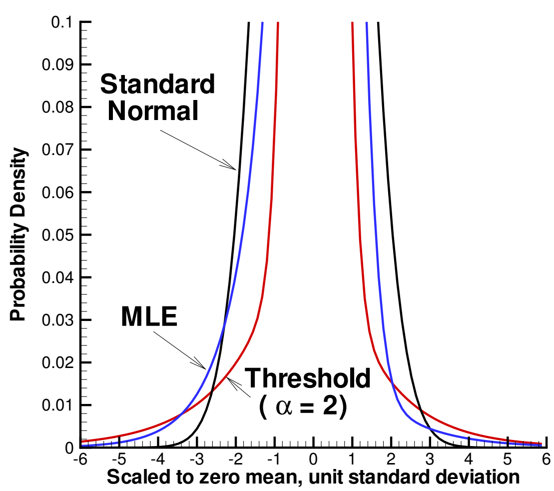

(c) *Zoom-in of panel (a).*

FIGURE 3.2: *Probability density plots for double exponential jump diffusion model. Parameters given in Table 3.1. Empirical histogram plotted in Figure 3.1.*

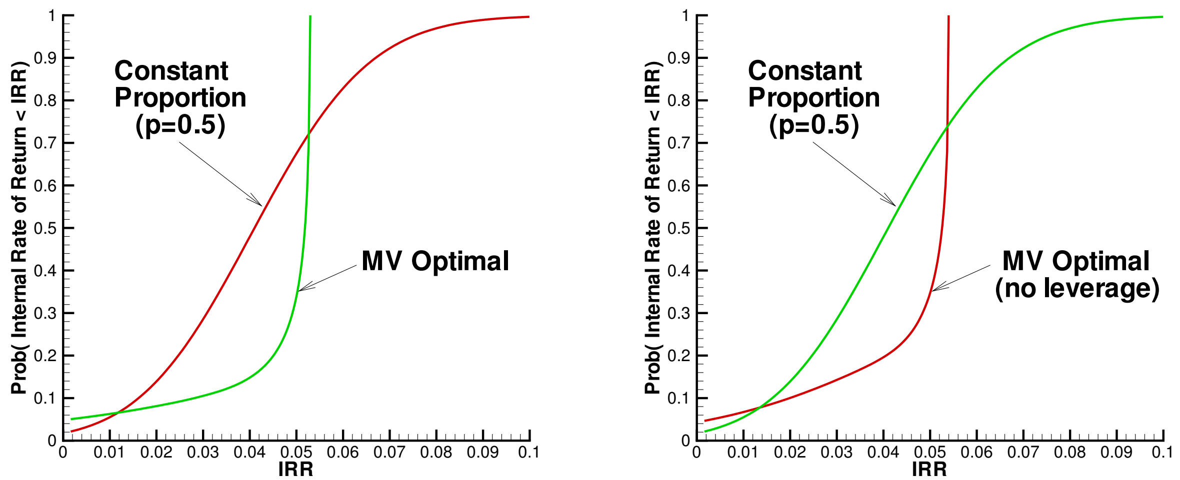

(a) *Constant proportion ($p = 0.5$) and MV optimal ($q_{\max} = 1.5$) strategies.*

(b) *Constant proportion ($p = 0.5$) and MV optimal ($q_{\max} = 1.0$) strategies.*

FIGURE 4.1: *Cumulative distribution functions of IRR. Input data given in Table 4.1.*

is not included in the IRR for the MV optimal case. Panel (a) is for the base case parameters (Table 4.1). We can see that the MV optimal strategy clusters the returns near the target level. For example, there is only a 15% probability that the IRR for the MV strategy will be less than .04, compared to a 40% probability that the constant proportion strategy will generate an IRR of less than .04.

However, there is no free lunch here. From Figure 4.1(a), we can also see that there is zero probability that the MV optimal control will give an IRR $>$ .06, while the constant proportion policy has about a 25% probability of achieving an IRR $>$ .06. In other words, the lower risk of underachieving the target return produced by the MV optimal control is obtained (in part) by sacrificing gains larger than the target.

Note that Figure 4.1(a) indicates that the constant proportion strategy has a smaller probability of very low IRR compared to the MV optimal policy. If we enforce a no-leverage constraint ($q_{\max} \leq 1.0$), then this probability is reduced. This case is shown in Figure 4.1(b). However, it is still larger than the constant proportion probability at low IRR. At low values of wealth, the MV strategy is fully invested in the risky asset, and hence a downward jump has a large adverse impact on wealth. However, as we shall see below, this effect seems to disappear when using resampled market data.

Figure 4.2 plots the mean and standard deviation of the fraction of wealth invested in the risky asset p over the 30-year investment horizon. Clearly, the MV optimal strategy has a higher proportion invested in the stock index at early times, and a lower fraction as we near the investment horizon. However, this is only true on average. The MV strategy is truly stochastic, i.e. it responds to the current attained wealth, in contrast to a deterministic glide path strategy (Forsyth and Vetzal, 2017a). In general, the MV optimal policy is contrarian: if the equity index drops, then more wealth is invested in the risky asset.[10] Conversely, if the risky asset increases in value, wealth is transferred to the risk-free asset (Dang and Forsyth, 2014, 2016; Forsyth and Vetzal, 2017a).

[10]This is not the case, however, if the attained level of wealth is sufficiently high that the quadratic wealth target can be achieved by investing solely in the risk-free asset. Obviously, if the portfolio has already been de-risked, it has no remaining exposure to the equity index.

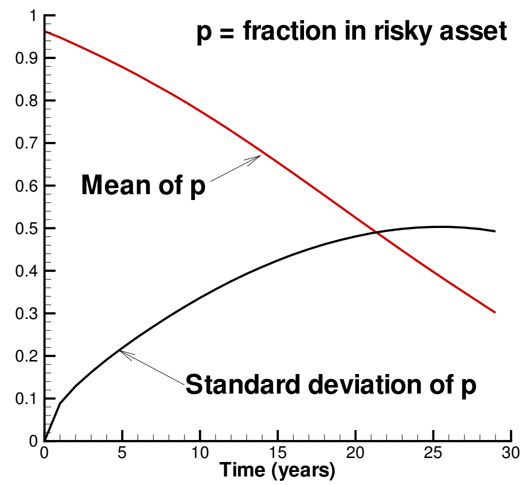

FIGURE 4.2: *Mean and standard deviation of the fraction of real wealth invested in the risky asset under the MV optimal policy. Input data given in Table 4.1.*

Figure 4.3 shows a heat map of the optimal control (in terms of the fraction of real wealth invested in the risky asset) as a function of time and real wealth. This figure shows that the optimal policy will require that a high fraction be put at risk as the remaining horizon shortens if past returns have been unfavorable. On the other hand, high early returns lead to lower optimal risk exposure, even zero exposure if real wealth is sufficiently high.[11]

4.1 Robustness in Synthetic Markets

In order to test robustness to parameter estimation ambiguity, we compute and store the MV optimal strategy for all the cases in Table 3.1 (strategy generating parameters), with the constraint that $E[W_T] = 417$. Then we carry out Monte Carlo (MC) simulations, assuming that the synthetic market parameters are different from the parameters used to compute the MV optimal strategies. To illustrate, consider the first set of results where it is indicated that the synthetic market parameters correspond to the threshold estimates with $\alpha = 2$. Three MV optimal strategies are computed, two based on the threshold estimates with $\alpha = 3$ and $\alpha = 4$, and one based on the ML parameter estimates. As noted above, in each case $E[W_T] = 417$. We then simulate the performance of these three strategies, and the constant proportion strategy with $p = 0.5$, assuming that the synthetic market evolves according to the threshold estimates with $\alpha = 2$.

This type of procedure is repeated for other possible synthetic market parameters. In each case, the basic idea is to examine what happens when the MV optimal strategy is determined using one set of parameters but then the market is simulated with a different set of parameters. The results in Table 4.3 show that the MV optimal strategy is very robust to parameter ambiguity. Even though the strategy generating parameters are different from the true synthetic market parameters, the results for the MV optimal control are very close in all cases.

[11]Note that the bright area along the bottom edge of Figure 4.3 is due to the insolvency constraint. This constraint requires that all investment in the risky asset be immediately liquidated if real wealth is non-positive.

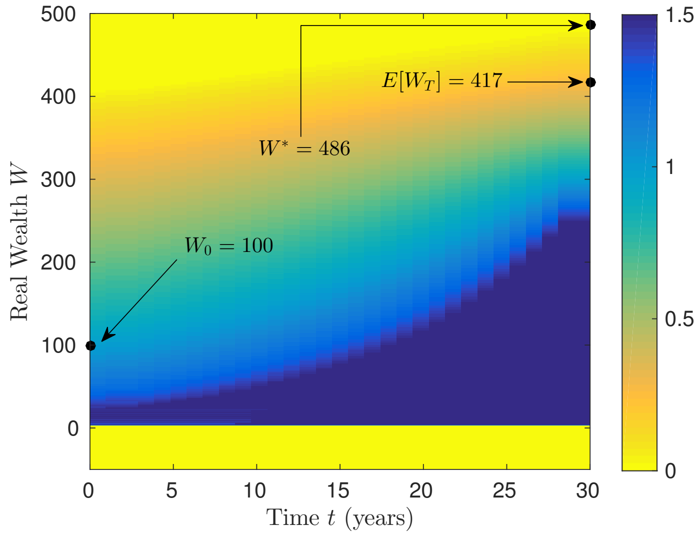

FIGURE 4.3: *Heat map of the fraction of real wealth invested in the risky asset. Input data given in Table 4.1.*

5 Lump Sum Investment: Historical Backtests

5.1 Single Paths

It is interesting to see how the MV optimal strategy would have fared along historical asset price paths. We calculate the optimal controls for the base case data (Table 4.1) with threshold ($\alpha = 3$) estimated parameters (Table 3.1). As above, the expected wealth target is $E[W_T] = 417$ and the quadratic wealth target is $W^* = 486$. We then apply these optimal controls to two 30-year series of historical monthly equity returns and T-bill rates: *1930:1 - 1959:12* and *1985:1 - 2014:12*. The results are plotted in Figure 5.1. The set of returns starting in 1930 are noteworthy, since this was probably one of the worst times in the last 100 years to make equity investments. We can see that the MV strategy was more (locally) volatile than the the constant proportion strategy in the early times, due to the fact that the MV strategy would have had a higher weight in equities in those years (see Figure 4.2).

For both historical periods, we can see that the MV optimal strategy would have produced a final wealth near the quadratic wealth target of $W^* = 486$. Note that for the period *1985:1 - 2014:12*, the MV strategy had a very small weight in the risky asset over the last 15 years. As a result, this strategy exhibited very small drawdowns during the bursting of the dot-com bubble in 2000 and the financial crisis of 2008.

5.2 Resampled Historical Data

Although it is fascinating to examine how the MV strategy would have performed for a given historical path, this is perhaps not so useful for risk management purposes. In order to develop a more quantitative approach to historical backtesting, we will use a resampling approach. To that end, we will use the moving block bootstrap resampling method to analyze the behaviour of the M-V optimal strategy.[12]

[12]See sources such as Sanfilippo (2003), Annaert et al. (2009), Bertrand and Prigent (2011), and Cogneau and Zakalmouline (2013) for discussions of bootstrap resampling applied to financial data.

Strategy Generating Parameters	Expected Value	Standard Deviation	$Pr(W_T < 300)$	$Pr(W_T < 400)$	Expected Free Cash
Synthetic Market Parameters: Threshold ($\alpha = 2$)					
Constant proportion ($p = .5$)	432	311	.40	.58	0.0
MV optimal: threshold ($\alpha = 3$)	419	123	.14	.21	5.1
MV optimal: threshold ($\alpha = 4$)	418	123	.13	.20	5.8
MV optimal: ML	420	123	.13	.19	6.9
Synthetic Market Parameters: Threshold ($\alpha = 3$)					
Constant proportion ($p = .5$)	417	299	.41	.60	0.0
MV optimal: threshold ($\alpha = 2$)	418	117	.13	.24	8.9
MV optimal: threshold ($\alpha = 4$)	418	117	.13	.21	11.0
MV optimal: ML	420	119	.13	.20	12.5
Synthetic Market Parameters: Threshold ($\alpha = 4$)					
Constant proportion ($p = .5$)	413	308	.42	.62	0.0
MV optimal: threshold ($\alpha = 2$)	418	114	.13	.24	12.0
MV optimal: threshold ($\alpha = 3$)	418	114	.13	.22	13.0
MV optimal: ML	421	116	.12	.20	15.0
Synthetic Market Parameters: ML					
Constant proportion ($p = .5$)	385	223	.42	.63	0.0
MV optimal: threshold ($\alpha = 2$)	420	107	.12	.24	1.0
MV optimal: threshold ($\alpha = 3$)	421	107	.12	.22	1.0
MV optimal: threshold ($\alpha = 4$)	421	106	.12	.21	2.0

TABLE 4.3: *Robustness check for synthetic market. Input data from Table 4.1. MV optimal strategies are computed using estimated strategy generating parameters from Table 3.1, constrained so that $E[W_T] = 417$. Investment performance is then assessed using MC simulations based on a different set of synthetic market parameters (also from Table 3.1. All strategies have annual rebalancing. The maximum leverage ratio for the MV optimal strategies is $q_{\max} = 1.5$.*

A single resampled path is constructed as follows. Suppose the investment horizon is T years. We divide this total time into k blocks of size b years, so that $T = kb$. We then select k blocks at random (with replacement) from the historical data. Each block starts at a random quarter. We then form a single path by concatenating these blocks. Note that the since we sample with replacement, the blocks may overlap. To avoid end effects, the historical data is wrapped around.[13] We then repeat this procedure for many paths. The sampling is done in blocks in order to account for (possible) serial dependence effects in the historical time series.

We compute and store the MV optimal parameters for the threshold ($\alpha = 3$) estimated parameters in Table 3.1 and the input data in Table 4.1. The strategy is determined so that $E[W_t] = 417$ in this synthetic market. We then apply this strategy using the bootstrap resampling method, based on the historical data for the period *1926:1 - 2014:12*. A total of 10,000 bootstrap samples are used, with results for various blocksizes given in Table 5.1.[14]

The main result from Table 5.1 is that for all choices of blocksize, the MV optimal strategy

[13] Since the great depression data of 1929-1931 appears near the start of our dataset, wrapping around will produce more blocks of bad returns compared to truncating the blocks.

[14] As discussed in Cogneau and Zakalmouline (2013), the choice of the blocksize is crucial. If the blocksize selected is too small, then serial dependence in the data is destroyed. If the blocksize is too large, then the variance estimates are unreliable. Unfortunately, the correct blocksize depends on the unknown stochastic process properties of the historical data. Consequently, we report results across a range of blocksizes.

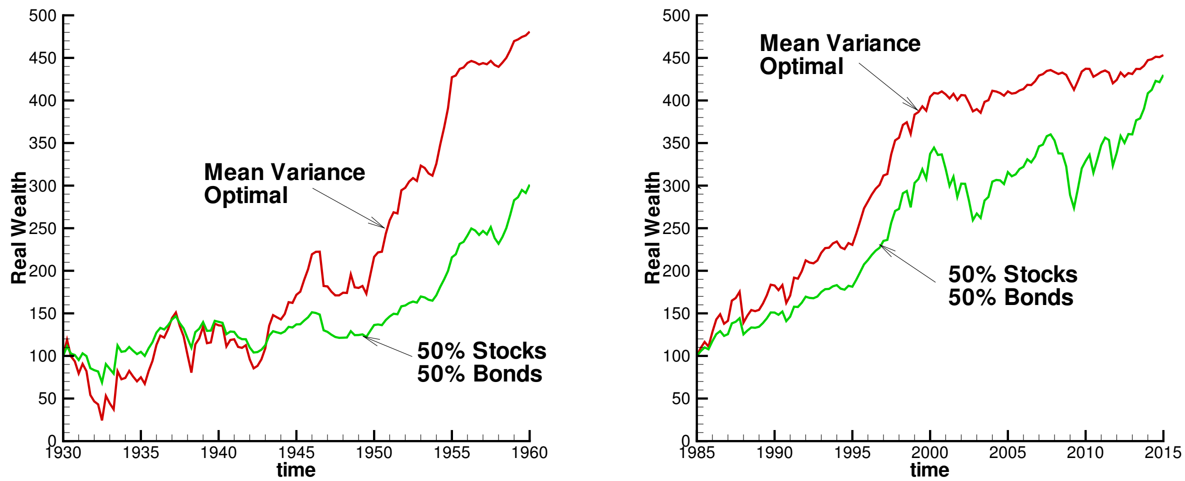

(a) *1930:1 - 1959:12.*

(b) *1985:1 - 2014:12.*

FIGURE 5.1: *Control generated using threshold ($\alpha = 3$) parameters from Table 3.1, applied to historical real returns. Input data given in Table 4.1. Quadratic wealth target $W^* = 486$.*

outperforms the constant proportion strategy by a wide margin. The results for a blocksize of one year seem to produce an extra level of randomness, perhaps due to the fact that serial correlation effects are reduced, while the blocksize of 30 years is clearly too large, since the randomness (as measured by standard deviation) seems artificially too small, as would be expected. In the following, we will use a blocksize of ten years, as this is the midpoint of the *reasonable* values of blocksize {5,10, 15} years.[15]

5.3 Robustness in Resampled Markets

In order to test robustness in the resampled historical data, we compute and store the MV optimal strategy for all sets of parameter estimates in Table 3.1. Each time, we impose the constraint that $E[W_T] = 417$. Then we carry out MC simulations, using the bootstrap resampling method described in Section 5.2, using the observed market data from 1926:1 to 2014:12. Similar to the results obtained above for the synthetic market (Table 4.3), the resampled results in Table 5.2 appear to be quite robust to parameter estimation ambiguity. A noteworthy aspect of this table is that the expected free cash for the MV optimal policy is somewhat higher than that seen earlier in the synthetic market experiments (Table 4.3).

Figure 5.2 presents some of the properties of the MV optimal strategy when applied to resampled historical data. Panel (a) shows the cumulative distribution function for the IRR. Comparing this plot with the corresponding plot for the synthetic market (Figure 4.1(a)), we can see that the low IRR cases where the constant proportion strategy outperforms the MV optimal policy occur with extremely low probability here. For example, $Pr(\text{IRR} < 0 = .0016)$ for the MV case compared with $Pr(\text{IRR} < 0 = .0007)$ for the constant proportion case. Given that the MV optimal policy produces higher expected free cash under resampling as compared to the synthetic market, we also show the cumulative distribution function for the IRR of the MV optimal strategy if we include

[15]We have tested all our results in subsequent sections for all blocksizes {1,5,10,15,30} and in all cases, the MV strategy outperforms the constant proportion policy.

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 300)$	$Pr(W_T < 400)$	Expected Free Cash
		Blocksize = 1 year			
Constant proportion ($p = .5$)	402	250	.41	.61	0
MV optimal	411	112	.13	.23	21
		Blocksize = 5 years			
Constant proportion ($p = .5$)	396	212	.40	.62	0
MV optimal	423	95	.09	.23	34
		Blocksize = 10 years			
Constant proportion ($p = .5$)	385	183	.38	.61	0
MV optimal	431	84	.07	.21	40
		Blocksize = 15 years			
Constant proportion ($p = .5$)	379	171	.37	.63	0
MV optimal	436	74	.06	.19	36
		Blocksize = 30 years			
Constant proportion ($p = .5$)	351	77	.31	.73	0
MV optimal	453	33	.00	.07	24

TABLE 5.1: *Based on data in Tables 3.1 (threshold, $\alpha = 3$) and 4.1, the MV optimal control is computed for $T = 30$ years. Moving block bootstrap resampling results, based on historical data for the period 1926:1 to 2014:12. There are a total of 10,000 bootstrap resamples. Annual rebalancing for both strategies. Maximum leverage ratio of $q_{\max} = 1.5$ for the MV strategy.*

the free cash. This is clearly superior to the constant proportion strategy across a wide range of IRRs, though it does slightly underperform for very high IRRs. Figure 5.2(b) plots the mean and standard deviation of the fraction of wealth optimally put at risk over time in the resampling experiment. It bears a strong resemblance to Figure 4.2, which was determined from a synthetic market. The similarity of these two figures is further evidence of the robustness of the MV optimal trading strategy: while the synthetic market does allow for diffusive randomness and large jumps, it assumes that parameters such as volatility and the risk-free rate are constant when determining the MV optimal strategy. However, when we apply such a strategy to real historical data which includes periods of high and low volatility and interest rates, we obtain very similar results as we do in an idealized synthetic market.

5.4 Alternative to Risk-Free Asset: 10 Year U.S. Treasury Bonds

The analysis discussed above assumed that the risk-free asset was a short-term (3-month) Treasury bill. Arguably, it may be more realistic to consider investments in long-term Treasury instruments instead. We investigate this possibility here by replacing the short-term risk-free asset with an index constructed from 10-year U.S. Treasury notes.[16] Figure 5.3(a) plots the evolution of this

[16]We obtain total returns for 10-year U.S. Treasury notes from CRSP back to May of 1941. To go back further to the start of our equity index data in 1926, we use long-term government bond yields reported in Table 48 of Homer (1977). These values are annual averages of monthly yields. We assume that yield was the prevailing yield halfway through each year, and determine yields for remaining months by linear interpolation. Each month, we calculate a total return index by assuming that the yield is applicable to a newly issued 10-year note selling at par paying semi-annual coupons. We recalculate the price of this bond using the new yield for the next month and sell it, reinvesting proceeds in another newly issued note. We merge this total return index with the CRSP data (which starts in 1941:5), and deflate using the CPI to convert it to real terms.

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 300)$	$Pr(W_T < 400)$	Expected Free Cash
Constant proportion ($p = .5$)	385	183	.38	.61	0.0
MV optimal: threshold ($\alpha = 2$)	431	85	.07	.23	34
MV optimal: threshold ($\alpha = 3$)	431	84	.07	.21	40
MV optimal: threshold ($\alpha = 4$)	430	83	.07	.21	42
MV optimal: ML	431	86	.07	.19	48

TABLE 5.2: *From input data in Tables 3.1 and 4.1, the MV optimal control is determined such that* $E[W_T] = 417$. *Reported values are moving block bootstrap results (10,000 resamples, blocksize of 10 years) based on historical data from 1926:1 to 2014:12. All strategies are rebalanced annually. MV optimal strategies all have a maximum leverage ratio* $q_{\max} = 1.5$.

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 480)$	$Pr(W_T < 540)$	Expected Free Cash
		Synthetic Market			
Constant proportion ($p = .5$)	536	381	.56	.63	0
MV optimal: threshold ($\alpha = 3$)	536	185	.25	.32	12
		Resampled Backtest			
Constant proportion ($p = .5$)	500	270	.55	.63	0
MV optimal: threshold ($\alpha = 3$)	534	146	.26	.39	59

TABLE 5.3: *The MV optimal strategy is computed based on data in Tables 3.1 and 4.1, except that* $r = 0.02362$ *because the real 10-year U.S. Treasury index is used as the proxy for the risk-free asset. The moving block bootstrap resampling results are based on historical data for the period 1926:1 to 2014:12. There are 10,000 resamples, with a blocksize of 10 years.*

index over time, clearly showing the spectacular 30-year bull market associated with the lengthy period of declining yields that began in the early 1980s. On an annualized basis, the real mean return for this index is 2.362%.

Assuming a synthetic market with parameters given in Table 4.1, with the exception that $r = 0.02362$ (i.e. the mean annualized return of the real 10-year U.S. Treasury index), we compute the expected return of an equal weight strategy. We then determine the parameter W^* for the MV optimal strategy which gives this same expected value (in this synthetic market). The optimal MV strategy is then computed which gives this same expected value. This strategy is then tested using observed historical data with, as usual, 10,000 resamples and a blocksize of 10 years. Figure 5.3(b) shows the cumulative distribution function for the IRR, and it is clearly quite similar to earlier cases. As with Figure 5.2(a), the MV optimal strategy outperforms the constant proportion strategy except for extremely low or high IRRs (especially when free cash is included). Such cases occur with quite low probability, however. Table 5.3 shows the results for both the synthetic market and the bootstrap resamples, using the observed market returns and the observed 10-year Treasury returns. Once again, it is clear that the MV optimal strategy performs fairly similarly in the synthetic market as in the resampled simulations, though it again generates much higher free cash across the resamples.

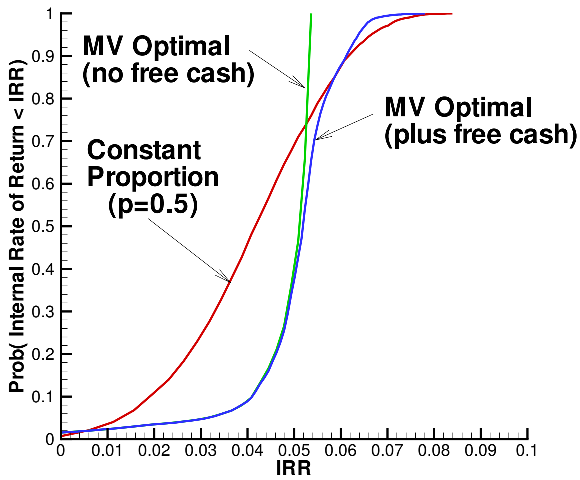

(a) *Cumulative distribution function for IRR.*

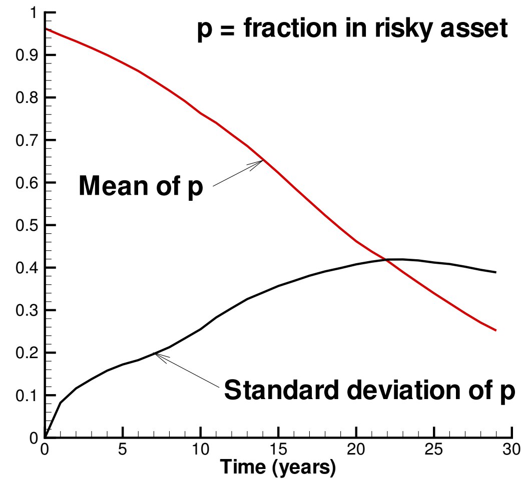

(b) *Mean and standard deviation of MV optimal control.*

FIGURE 5.2: *Properties of MV optimal strategy using input parameters from Tables 3.1 (threshold with $\alpha = 3$) and 4.1. Bootstrap resampled backtest (10,000 resamples, blocksize of 10 years) with historical data from 1926:1 to 2014:12.*

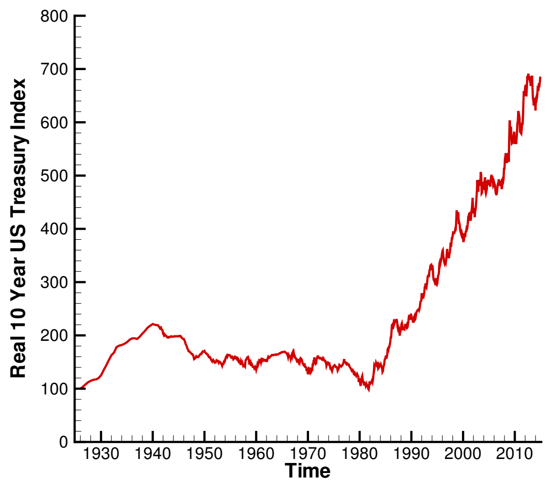

(a) *Real 10-year U.S. Treasury index, 1926-2014.*

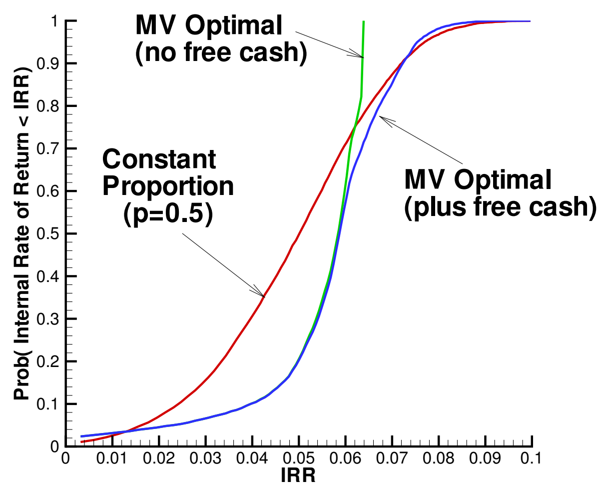

(b) *Cumulative distribution function for IRR.*

FIGURE 5.3: *MV optimal control calculated using data in Tables 3.1 and 4.1, except that $r = .02362$ since the real 10-year U.S. Treasury index is used to proxy for the risk-free asset. Moving block bootstrap results based on historical data from 1926:1 to 2014:12, with 10,000 resamples and blocksize of 10 years.*

Investment horizon	30 years
Market parameters	Threshold ($\alpha = 3$) (see Table 3.1)
Risk-free rate r	0.00623
Initial investment W_0	0.0
Investment at each year	10
Rebalancing interval	1 year
Maximum leverage $q_{\max}$	1.5
If insolvent	Trading stops

TABLE 6.1: *Accumulation case data. Cash is injected at $t = 0,1,\ldots,29$ years.*

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 580)$	$Pr(W_T < 680)$	Expected Free Cash
Synthetic Market					
Constant proportion ($p = .5$)	682	324	.45	.61	0
MV optimal: threshold ($\alpha = 3$)	682	124	.12	.24	16
Resampled Backtest					
Constant proportion ($p = .5$)	648	211	.42	.60	0
MV optimal: threshold ($\alpha = 3$)	683	92	.12	.32	61

TABLE 6.2: *Based on data in Tables 3.1 and 6.1, the MV optimal control is computed for $T = 30$ years with initial wealth of zero. Accumulation phase, cash injected at a rate of 10 per year. $E[W_T] = 682$ in the synthetic market. Moving block bootstrap resampling results are based on historical data for the period 1926:1 to 2014:12, with 10,000 resamples and a blocksize of 10 years.*

6 Wealth Accumulation

In this section, we consider a situation of wealth accumulation, rather than examining the strategy based on an initial lump sum investment. We suppose that the investor has zero initial wealth, and invests a fixed real amount over a 30-year period. The problem parameters are described in Table 6.1, and some representative results are shown in Table 6.2. Again, we can see that for both the synthetic market and the resampled historical market data, the MV optimal strategy considerably outperforms the constant proportion strategy. For the synthetic market case, both the MV optimal and constant proportion strategies have $E[W_T] = 682$. The MV optimal control for the synthetic market is computed and stored, and used in the resampled historical backtest. Other tests (not shown) indicate that the MV optimal strategy is robust to parameter estimation ambiguity (in the synthetic market and the resampled historical market). The MV strategy is also superior to the constant proportions strategy for all choices of blocksize in the resampled historical market.

Figures 6.1 and 6.2 show the cumulative distribution function of the IRR and the mean controls for the synthetic market and the resampled historical backtest respectively. Figure 6.1(a) illustrates that in the synthetic market the constant proportion strategy produces a superior IRR compared to the MV optimal policy for low probability ($< .02$) low return cases. The same effect can be seen for the resampled data in Figure 6.2(a), but here the probability of this occurring is very small ($< .005$). It is interesting to see that the historical resampled backtest appears to show less risk than the the synthetic market. We can see from Figures 6.1(b) and 6.2(b) that the MV strategy

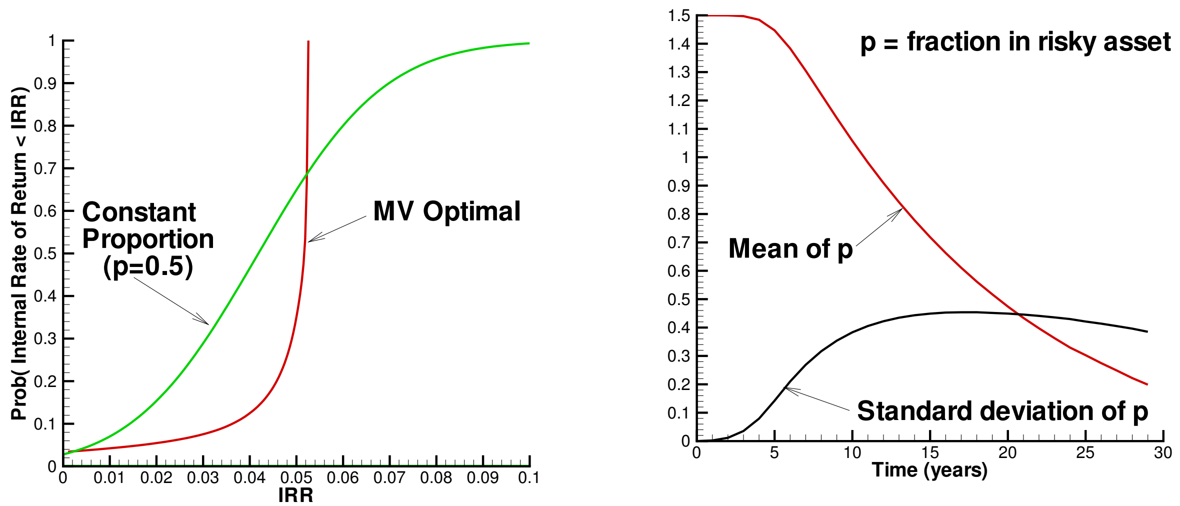

(a) *Cumulative distribution function of IRR.*

(b) *Mean and standard deviation of MV optimal control.*

FIGURE 6.1: *Properties of MV optimal strategy for the accumulation case using input parameters from Tables 3.1 and 6.1. Synthetic market.*

uses much more leverage at early times compared to the lump sum case (Figure 4.2). Intuitively, in the accumulation case, if insolvency occurs during the early stages of wealth accumulation, the debt can be paid off with the future cash injections.

6.1 Wealth Accumulation: Single Historical Paths

Figure 6.3 shows the results for both the constant proportion strategy and the MV optimal strategy (computed using the data in Table 6.1) for the historical returns from *1930:1 - 1959:12* (panel (a)) and *1985:1 - 2014:12* (panel (b)). Along these two historical paths, the MV optimal strategy clearly outperformed the constant proportion strategy in the long-run. Considerable patience may have been necessary during the earlier period, as very poor initial returns combined with high leverage for the MV optimal strategy put it in a trailing position during the first few years.

7 Wealth Decumulation

In this section, we consider the problem of wealth decumulation. We suppose that the investor has initial wealth 200, and withdraws at a real rate of 8 per year annually over 20 years (4% real of the original amount invested in line with the rule of thumb suggested in Bengen (1994)). The problem parameters are described in Table 7.1.

As a benchmark, we use a constant proportion strategy with 40% stocks and 60% bonds, rebalanced yearly. This gives a real expected value of $E[W_T] = 202$ at $T = 20$ years (i.e. close to the original wealth) in the synthetic market. We then determine the value of W^* which gives this same expected value for the MV optimal control, and store this strategy, which is used in the historical backtests.

Some representative results are shown in Table 7.2. Again, we can see that for both the synthetic market and the resampled historical market data, the MV optimal strategy considerably

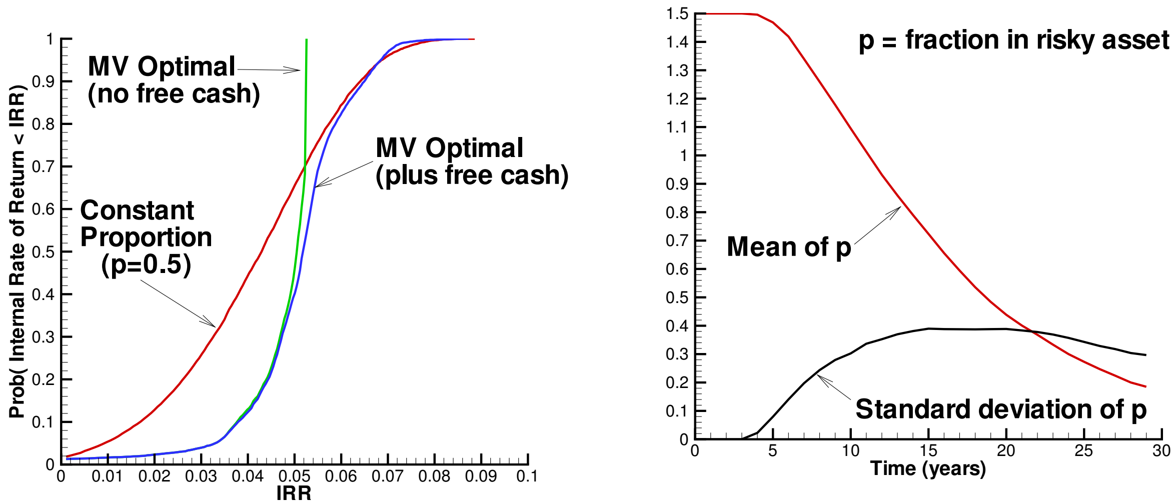

(a) *Cumulative distribution function for IRR.*

(b) *Mean and standard deviation of MV optimal control.*

FIGURE 6.2: *Properties of MV optimal strategy for the accumulation case using input parameters from Tables 3.1 and 6.1. Bootstrap resampled backtest, based on historical data from 1926:1 to 2014:12 with 10,000 resamples and blocksize of 10 years.*

Investment horizon	20 years
Market parameters	Threshold ($\alpha = 3$) (see Table 3.1)
Risk-free rate r	0.00623
Initial investment W_0	200
Withdrawal at each year	8
Rebalancing Interval	1 year
Maximum leverage $q_{\max}$	1.5
If insolvent	Trading stops

TABLE 7.1: *Decumulation case data. Cash is withdrawn at $t = 1, \ldots, 29, 30$ years.*

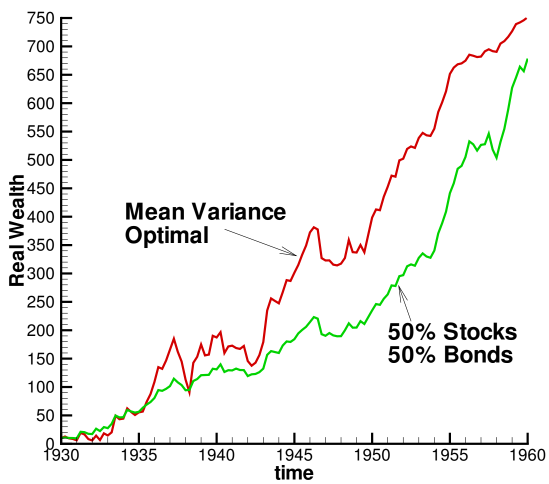

(a) *1930:1 - 1959:12.*

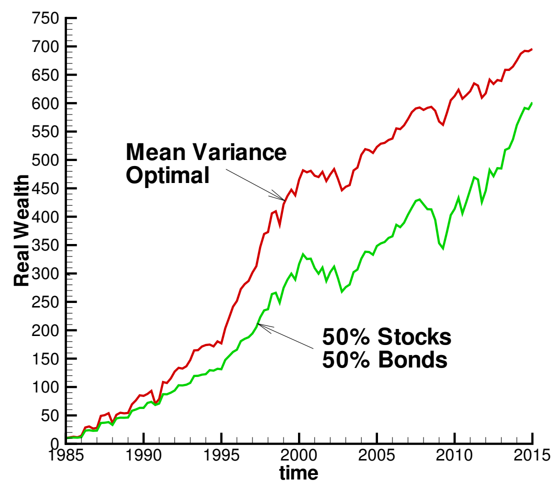

(b) *1985:1 - 2014:12.*

FIGURE 6.3: *Control generated using threshold ($\alpha = 3$) parameters from Table 3.1, applied to historical real returns. Input data given in Table 6.1. This wealth accumulation case has a quadratic wealth target* $W^* = 746$.

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 20)$	$Pr(W_T < 100)$	$Pr(W_T < 200)$	Expected Free Cash
			Synthetic Market			
Constant proportion ($p = .4$)	202	144	.02	.22	.60	0
MV optimal: threshold ($\alpha = 3$)	202	73	.05	.11	.27	6.4
			Resampled Backtest			
Constant proportion ($p = .4$)	194	109	.01	.22	.57	0
MV optimal: threshold ($\alpha = 3$)	212	51	.02	.05	.25	61

TABLE 7.2: *Based on input data in Tables 3.1 and 7.1, the MV optimal control is computed. Decumulation phase, cash withdrawn at a real rate of 8 per year. $E[W_T] = 202$ in the synthetic market. Moving block bootstrap resampling results are based on historical data for 1926:1 to 2014:12 using 10,000 resamples and blocksize of 10 years.*

outperforms the constant proportion counterpart (except for extremely low values of W_T which occur with very low probability). Other tests (not shown) demonstrate that the MV optimal strategy is robust to parameter estimation ambiguity in both the synthetic market and the resampled historical market). The MV strategy also turns out to be superior to the constant proportion strategy for all choices of blocksize in the resampled historical market.

Figure 7.1 shows the cumulative distribution function of the IRR for both the constant proportion strategy and the optimal M-V strategy, as well as the mean and standard deviation of the optimal control, for the synthetic market. Panel (a) shows that the MV optimal strategy again sacrifices upside potential in exchange for a large amount of downside protection, except for cases with extremely low terminal wealth. Panel (b) shows that the optimal policy here is much more cautious during the earlier years, as compared to the accumulation case depicted in Figure 6.1(b). The fraction of wealth optimally invested in the risky asset here is a bit over 50% initially, and it remains there for some time before gradually tapering off (on average). Similar results are shown for the resampled backtest in Figure 7.2.

7.1 Wealth Decumulation: Single Historical Paths

Figure 7.3 shows the results for both the constant proportions strategy and the MV optimal strategy (computed using the data in Table 7.1) for the historical returns from *1930:1 - 1949:12* and *1995:1 - 2014:12*. In both cases, following the MV optimal strategy would have produced a terminal wealth larger than that from the constant proportion strategy, with a bigger gap emerging over the earlier time period shown in panel (a).

8 Conclusions

Based on long-term historical data, we determine the average parameters for an inflation-adjusted stock market index, which is assumed to follow a jump diffusion process. Assuming a portfolio consisting of a risk-free bond and an equity index, we determine the optimal asset allocation strategy for an MV investor. We impose realistic constraints on the strategy (i.e. infrequent rebalancing, a maximum leverage ratio, and no trading if insolvent).

We reiterate that for a lump sum investment, assuming a continuously rebalanced portfolio containing a risk-free bond and an equity index following a jump diffusion process, any deterministic glide path strategy is inferior to a constant proportion strategy (Forsyth and Vetzal, 2017b). More

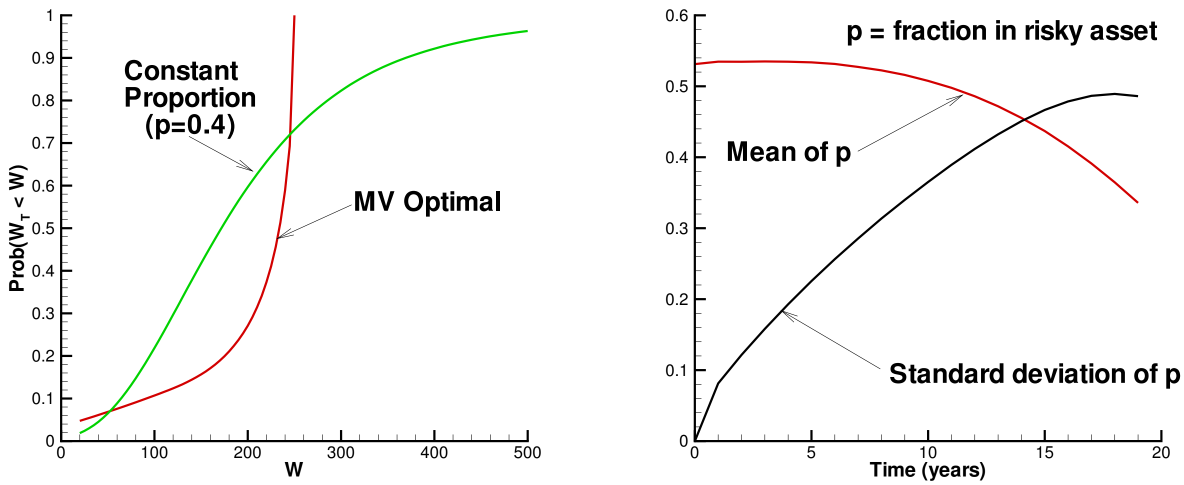

(a) *Cumulative distribution function of terminal wealth W_T.*

(b) *Mean and standard deviation of MV optimal control.*

FIGURE 7.1: *Properties of the MV optimal strategy for the decumulation case using input parameters from Tables 3.1 and 7.1. Synthetic market.*

specifically, when both strategies can have the same expected terminal wealth, the MV optimal strategy has a standard deviation of terminal wealth that is less than or equal to that for the constant weight strategy. The dominating constant weight strategy is simply the time average of the deterministic strategy.

If we consider discrete rebalancing, with periodic addition (accumulation) or withdrawal (decumulation), then the optimal deterministic glide path can produce a slightly smaller standard deviation (for a fixed expected terminal wealth) compared to a constant proportion strategy. However, this improvement is very small. Historical backtests and Monte Carlo simulations indicate that this result holds in more realistic cases as well (Graf, 2017; Esch and Michaud, 2014). Hence we have focused on comparing the optimal MV strategy with a constant weight strategy. The difference in performance metrics between the optimal MV and the constant weight strategy overwhelms the small differences between optimal glide path and constant weight, in all cases.

In a synthetic market with the parameters determined by the historical averages, we compare the MV optimal strategy with a constant weight strategy under the condition that both strategies have the same expected value. Extensive Monte Carlo tests show that:

- the MV optimal strategy has the same expected terminal wealth as the constant weight strategy, with much reduced variance and probability of shortfall;
- the MV strategy is robust to model parameter ambiguity. In other words, we can determine the MV strategy based on market parameters different from the synthetic market parameters used in the Monte Carlo tests, and the MV strategy continues to outperform the constant weight strategy.

Using the MV optimal strategy obtained using the historical average parameters, we have also carried out an extensive set of bootstrap resampled backtests based on historical data over the *1926:1 -2014:12* period. Again, the MV optimal strategy outperformed the constant weight strategy. Obviously, since we are using MV criteria to assess performance, the MV optimal strategy will

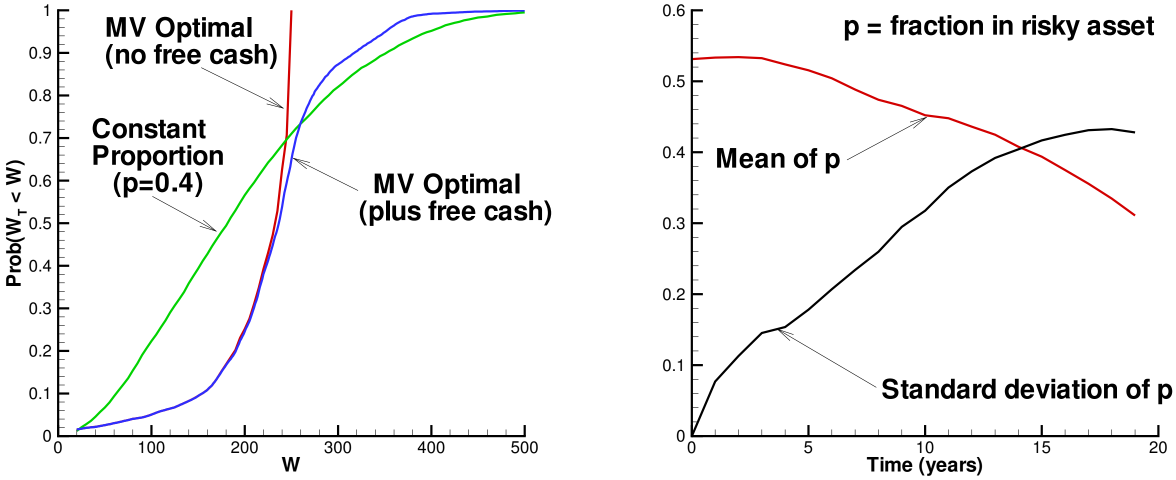

(a) *Cumulative distribution function of terminal wealth W_T.*

(b) *Mean and standard deviation of the MV optimal control.*

FIGURE 7.2: *Properties of the MV optimal strategy for the decumulation case using input parameters from Tables 3.1 and 7.1. Bootstrap resampled backtest, based on historical data from 1926:1 to 2014:12 with 10,000 resamples and blocksize of 10 years.*

outperform any other strategy, including the constant proportion strategy. However, we emphasize that it is the magnitude of the outperformance that is significant. The same average terminal wealth is attained by the MV optimal policy with much lower standard deviation and shortfall probability.

These results all hold for the case of lump-sum investment, accumulation and decumulation scenarios. This would suggest that investors should be using non-deterministic asset allocation strategies (i.e. strategies which respond to the path-dependent accumulated wealth state variable), in contrast to the omnipresent deterministic glide path strategies. Although the computations underlying the determination of the MV optimal strategy are complex, the strategy can ultimately be implemented simply by using a table such as that represented by the heat map in Figure 4.3.

The MV optimal strategy is highly contrarian: the risky asset fraction is increased when the total real portfolio value decreases, and the risky asset fraction is decreased as the total portfolio value nears the target. At all times, however, the strategy adapts to the difference between the target at the horizon and the current level of attained wealth in order to determine how much (if any) wealth to invest in the risky asset.

Although these results are very promising, the MV optimal strategy does have some drawbacks:

- the constant weight strategy outperforms the MV optimal strategy for very poor outcomes which occur with very low probabilities (e.g. 2-4%);
- in the case of a downward-trending market early during the investment period, the MV optimal policy buys increasing amounts of the risky asset, thereby causing the investor to *"catch a falling knife"*; and
- the MV optimal policy de-risks as the target is approached, thereby sacrificing the possibility of very large gains.

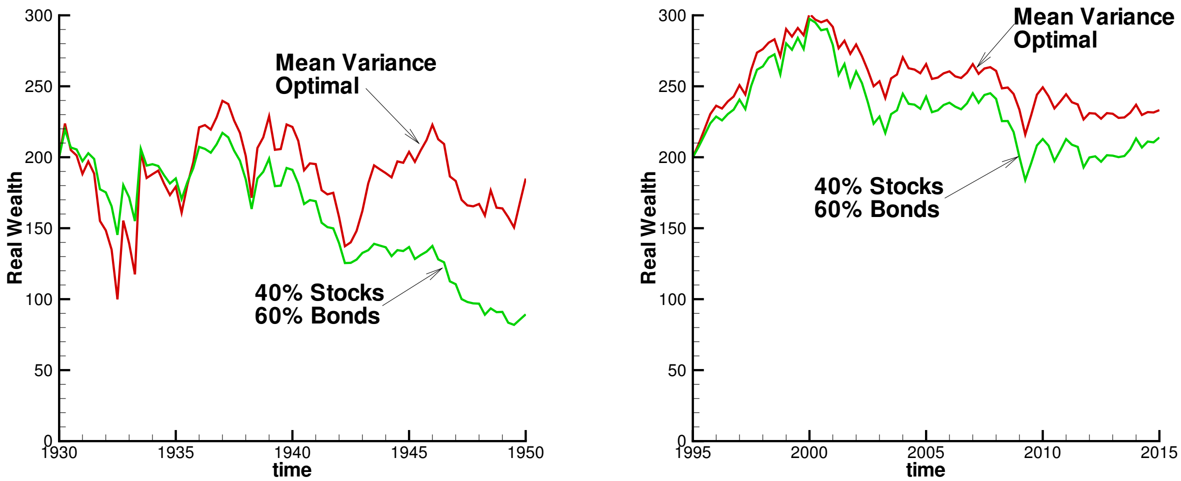

(a) *1930:1 - 1949:12.*

(b) *1995:1 - 2014:12.*

FIGURE 7.3: *Control generated using threshold ($\alpha = 3$) parameters from Table 3.1, applied to historical real returns. Input data given in Table 7.1. This wealth decumulation case has a quadratic wealth target $W^* = 248$.*

It is instructive to examine the application of the MV optimal strategy in the context of historical market returns in 1930-1950. In this case, the MV optimal investor *caught the falling knife* in the 1930s. By 1950, this investor was well ahead of a constant proportion investor, but at the cost of a rather rough ride. Conversely, if the market continued to trend upward, the MV optimal investor would de-risk and experience lower terminal wealth compared to the constant weight investor. Both of these situations will be psychologically difficult. Although following the MV optimal strategy puts the odds in our favour, this might be small comfort in some low-probability scenarios.

Finally, we observe an interesting phenomenon for all the resampled backtests. The amount of the free cash accumulated was always much larger in the backtests compared to the synthetic market tests. We conjecture that this is due to the high serial correlation of interest rates, which has been ignored in the strategy generating model (i.e. we assumed constant interest rates). However, it is clear that standard mean-reverting interest rate models will not be helpful here. Typical calibration of interest rate models results in mean-reversion times that are relatively fast compared to the long-term investment horizons considered here (see, e.g. Vetzal, 1997). The historical data implies that over the long term, interest rates are perhaps more appropriately modelled by a regime-switching process. Given such a process, it is straightforward to compute the MV optimal strategy. However, we are faced with two challenges in practical implementation of a regime-switching aware strategy. First, it will be difficult to estimate parameters. Second, in order to apply the strategy, we must determine which regime we are in at any given time. This requires determination of a hidden Markov chain state variable, which is a non-trivial task.

Appendices

A Formulation as an HJB Equation

Our problem is posed on the domain $0 \leq s \leq \infty$ and $-\infty \leq b \leq +\infty$. We define the solvency region $\mathcal{N}$ as

$$\mathcal{N} = \{(s,b) \in [0,\infty) \times (-\infty, +\infty) : W(s,b) > 0\}. \tag{A.1}$$

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

$$\mathcal{B} = \{(s,b) \in [0,\infty) \times (-\infty, +\infty) : W(s,b) \leq 0\}. \tag{A.2}$$

Let backwards time $\tau = T - t$, and $\tau_j = T - t_m$, $m = 0, \ldots, M$, $j = M - m$, be the time left to the investment horizon at the mth re-balancing time. Note that $\tau_0 = T$ and $\tau_M = 0$. In addition, for later reference let $\tau_j^+ = \tau_j + \epsilon$, where $\epsilon \to 0^+$. Also, denote by $\bar{a}_j$, $j = 0, \ldots, M-1$, the withdrawal amount in terms of backward time variable τ. Further, let the control variable at each re-balancing time be denoted by c_j, $j = 0, \ldots, M$, where

$$c_j = (\hat{b}_j, \hat{f}_j). \tag{A.3}$$

In (A.3), $\hat{b}_j$ is the amount of the risk-free asset after re-balancing and $\hat{f}_j$ is the free cash flow generated. The optimal withdrawal of free cash is discussed in Dang and Forsyth (2016). We will specify $\hat{f}_j$ precisely below in equation (A.6).

Given a re-balancing time t_k, $k = 0, \ldots, M$, the time-t_k value of all cash withdrawals made on or after time t_k, denoted by w_k, is computed by

$$w_k = \sum_{m=k}^{M} a_m \mathrm{e}^{-r\,(t_m - t_k)} \ . \tag{A.4}$$

Define

$$\hat{W}_k = W^* \mathrm{e}^{-r\,(T-t_k)} + w_k \ , \tag{A.5}$$

where w_k is defined in equation (A.4). At time t_k, if $W_{t_k} > \hat{W}_k$, we (i) withdraw $(W_{t_k} - \hat{W}_k)$ from the portfolio, and (ii) invest the remaining wealth $\hat{W}_k$ in the risk-free asset for the balance of the investment horizon Dang and Forsyth (2016). It is straightforward to see that $W_T = W^*$ under this strategy. We refer to the amount of (i) as "free cash flow" to clearly distinguish it from the cash withdrawal a_k. As shown in Dang and Forsyth (2016), this strategy is M-V optimal.

Since the free cash flow falls outside the scope of the MV framework, unless otherwise noted we do *not* include the expected value of the free cash flow and accumulated interest in the terminal portfolio wealth. More precisely, the control variable $\hat{f}_j$ in (A.9) is given by

$$\hat{f}_j = \max\Big((s+b) - W^* e^{-r\,(T-t_j)} - w_j, 0 \Big) \tag{A.6}$$

Although we have assumed that we have periodic withdrawals ($\bar{a}_j > 0$), all the above expressions are valid for the cash injection $\bar{a}_j < 0$ case as well.

In addition, define the following operators

$$\begin{aligned} \mathcal{L}V &\equiv \frac{\sigma^2 s^2}{2} V_{ss} + (\mu - \lambda\kappa) s V_s + r b V_b - \lambda V, \\ \mathcal{J}V &\equiv \int_0^\infty p(\xi) V(\xi s, b, \tau)\ d\xi. \end{aligned} \tag{A.7}$$

For a given value of W^*, we compute the corresponding point on the efficient frontier as follows. We first solve the value function problem (2.12) as follows. At backwards time $\tau = \tau_j$, $j = 0, \ldots, M$, we enforce the following conditions:

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

$$V(s,b,\tau_j^+) = V(0,s+b-\bar{a}_j,\tau_j). \tag{A.8}$$

2. If $(s,b) \in \mathcal{N}$, we enforce the re-balancing optimality condition

$$\begin{aligned} V(s,b,\tau_j^+) &= \min_{c_j \in \mathcal{Z}} V(S_j^+, B_j^+, \tau_j) \\ c_j &= (\hat{b}_j,\hat{f}_j) \\ S_j^+ &= s+b-\bar{a}_j-\hat{b}_j-\hat{f}_j \; ; \; B_j^+ = \hat{b}_j. \end{aligned} \tag{A.9}$$

   As noted above, $\mathcal{Z}$ is the set of *admissible* controls. In our case, the choice of control is subject to the leverage constraint

$$\frac{S_j^+}{S_j^+ + B_j^+} < q_{\max}. \tag{A.10}$$

   We also require that $\hat{f}_j \geq 0$, i.e. we allow only free cash withdrawals. Note that for the special case of τ_0, we have $V(s,b,\tau_0) = (W^* - W(s,b))^2$.

Within each time period $(\tau_j^+, \tau_{j+1}]$, $j = 0, \ldots, M-1$, we have

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

$$V(s,b,\tau) = V(0,s+b,\tau). \tag{A.11}$$

2. If $(s,b) \in \mathcal{N}$, $V(s,b,\tau)$ satisfies the partial integro-differential equation (PIDE)

$$V_\tau = \mathcal{L}V + \mathcal{J}V, \tag{A.12}$$

   subject to the initial condition (A.9).

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).

We denote by $c^*_{W^*}(\cdot)$ the optimal control of problem (2.12). Once we have determined $c^*_{W^*}(\cdot)$ from the solution process described above, we use this control to determine

$$E^{x,t}_{c^*_{W^*}}[W_T], \tag{A.13}$$

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) over each re-balancing time period $[\tau_j, \tau_{j+1}]$, $j = 0, \ldots, M-1$.

We store the optimal controls generated from solution of the HJB equation, and then compute other quantities of interest (e.g. cumulative probability distribution) using a Monte Carlo method.

B Parameter Estimation

This appendix provides details about the maximum likelihood and thresholding methods used for parameter estimation.

B.1 Maximum Likelihood

Let $\phi(x)$ be the standard normal density function with mean zero and unit variance, evaluated at point x. Let $\Phi(x)$ be the standard normal cumulative distribution function. Assuming that only one jump occurs in Δt, the density function for the log return ΔX is

$$\begin{aligned}\mathcal{P}(\Delta X) = &\frac{(1-\lambda\Delta t)}{\sigma\sqrt{\Delta t}}\phi\left(\frac{\Delta X - \mu'\Delta t}{\sigma\sqrt{\Delta t}}\right) \\ &+ \lambda\Delta t\left[p\eta_1 \exp\left(\eta_1^2\sigma^2\Delta t/2\right)\exp\left(-(\Delta X-\mu'\Delta t)\eta_1\right)\Phi\left(\frac{\Delta X-\mu'\Delta t-\eta_1\sigma^2\Delta t}{\sigma\sqrt{\Delta t}}\right)\right] \\ &+ \lambda\Delta t\left[(1-p)\eta_2 \exp\left(\eta_2^2\sigma^2\Delta t/2\right)\exp\left((\Delta X-\mu'\Delta t)\eta_2\right)\Phi\left(-\frac{\Delta X-\mu'\Delta t+\eta_2\sigma^2\Delta t}{\sigma\sqrt{\Delta t}}\right)\right],\end{aligned} \tag{B.1}$$

where $\mu' = \mu - \lambda\kappa - \sigma^2/2$. The ML estimate for $\Theta = \{\mu, \sigma, \lambda, p, \eta_1, \eta_2\}$ is determined from

$$\max_{\Theta}\left(\sum_i \log\mathcal{P}(\Delta X_i)\right). \tag{B.2}$$

B.2 Thresholding Estimates

Based on the monthly log returns in Figure 3.1, we set the jump detection indicator $\mathbf{1}_i$ as follows

$$\begin{aligned}\mathbf{1}_i^{up} &= \begin{cases}1 & \text{if } \Delta\hat{X}_i > \alpha\,\hat{\sigma}\sqrt{\Delta t} \\ 0 & \text{otherwise}\end{cases}, \\ \mathbf{1}_i^{dwn} &= \begin{cases}1 & \text{if } \Delta\hat{X}_i < -\alpha\,\hat{\sigma}\sqrt{\Delta t} \\ 0 & \text{otherwise}\end{cases}.\end{aligned} \tag{B.3}$$

Note that we have replaced the parameters $\mathcal{A}/(\Delta t)^{\beta}$ from equation (3.3) by the single parameter α in equation (B.3). Criteria (B.3) allows us to separate downward from upward jumps.

Define

$$\sum_{i=1}^{N}\mathbf{1}_i^{up} = N^{up}\ ; \quad \sum_{i=1}^{N}\mathbf{1}_i^{dwn} = N^{dwn}\ ; \quad N^{jmps} = N^{up} + N^{dwn}\ ; \quad \sum_{i=1}^{N}\left(1 - \mathbf{1}_i^{up} - \mathbf{1}_i^{dwn}\right) = N^{gbm},$$

where N^{jmps} is the total number of jumps detected and N^{gbm} is the number of GBM increments. Our estimate of the diffusive volatility is then

$$\hat{\sigma}^2 = \frac{1}{\Delta t}\,var\left(\left\{\Delta\hat{X}_i \mid \left(\mathbf{1}_i^{dwn} + \mathbf{1}_i^{up} = 0\right)\right\}\right). \tag{B.4}$$

Equations (B.3)-(B.4) constitute an implicit equation for $\hat{\sigma}$, which must be solved by an iterative method (Clewlow and Strickland, 2000). Having estimated the Brownian motion volatility $\hat{\sigma}$, we can estimate the other jump parameters under the assumption that only one jump occurs in $[t_i, t_{i+1}]$:

$$\lambda = \frac{N^{jmps}}{T}\ ; \qquad\qquad p = \frac{N^{up}}{N^{jmps}}\ ;$$
$$\eta_1 = \frac{1}{mean\left(\left\{\Delta\hat{X}_i \;\middle|\; (\mathbf{1}_i^{up} = 1)\right\}\right)}\ ; \quad \eta_2 = \frac{-1}{mean\left(\left\{\Delta\hat{X}_i \;\middle|\; (\mathbf{1}_i^{dwn} = 1)\right\}\right)}. \tag{B.5}$$

Once we fix the estimates for $\sigma, \lambda, p, \eta_1, \eta_2$, we can estimate the drift term μ. The simplest method is as follows. From equation (2.1), we have ($X = \log S$)

$$dX = \left(\mu - \lambda\kappa - \frac{\sigma^2}{2}\right)\, dt + \sigma\, dZ + d\left(\sum_{i=1}^{\pi_t} \log \xi_i\right). \tag{B.6}$$

Taking expectations of both sides of equation (B.6) (and continuing to assume that only one jump takes place in $[t, t + dt]$) gives

$$E[dX] = \left(\mu - \lambda\kappa - \frac{\sigma^2}{2}\right)\, dt + \lambda E[\log \xi]\, dt. \tag{B.7}$$

Recall that $\kappa = E[\xi] - 1$. Writing equation (B.7) in discrete time gives

$$\frac{mean\left(\{\Delta X_i\}\right)}{\Delta t} = \left(\mu - \lambda\kappa - \frac{\sigma^2}{2}\right) + \lambda\left(\frac{p}{\eta_1} - \frac{(1-p)}{\eta_2}\right). \tag{B.8}$$

Equation (B.8) can be solved for the drift parameter μ.

References

Aït-Sahalia, Y. and J. Jacod (2012). Analysing the spectrum of asset returns: jump and volatility of components of high frequency data. *Journal of Economic Literature 50*, 1007–1050.

Annaert, J., S. V. Osslaer, and B. Verstraete (2009). Performance evaluation of portfolio insurance strategies using stochastic dominance criteria. *Journal of Banking and Finance 33*, 272–280.

Baeuerle, N. and S. Grether (2015). Complete markets do not allow free cash flow streams. *Mathematical Methods of Operations Research 81*, 137–146.

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.

Bertrand, P. and J. Prigent (2011). Omega performance measure and portfolio insurance. *Journal of Banking and Finance 35*, 1811–1823.

Clewlow, L. and C. Strickland (2000). *Energy Derivatives - Pricing and Risk Management*. Lacima, London.

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

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

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. 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. 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.

Esch, D. N. and R. O. Michaud (2014). The false promise of target date funds. Working paper, New Frontier Advisors, LLC.

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

Forsyth, P. A. and K. R. Vetzal (2017b). A comparison of adaptive and non-adaptive asset allocation strategies. Working paper, Univesity of Waterloo.

Graf, S. (2017). Life-cycle funds: Much ado about nothing? European Journal of Finance, forthcoming.

Hanson, F. and Z. Zhu (2004). Comparison of market parameters for jump-diffusion distributions using multinomial maximum likelihood estimation. pp. 3919–3924. Proceedings of the 43rd IEEE Conference on Decision and Control, Bahamas.

Homer, S. (1977). *A History of Interest Rates* (2nd ed.). New Brunswick, NJ: Rutgers University Press.

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

Kou, S. 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.

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 20:1*, 1–37.

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

Sanfilippo, G. (2003). Stocks, bonds and the investment horizon: a test of time diversification on the french market. *Quantitative Finance 3*, 345–351.

Shimizu, Y. (2013). Threshold estimation for stochastic differential equations with jumps. pp. 747–752. Proceedings of the 59th ISI World Statistics Conference, Hong Kong.

Synowiec, D. (2008). Jump diffusion models with constant parameters for financial log-return processes. *Computers and Mathematics with Applications 56*, 2120–2127.

Vetzal, K. (1997). Stochastic volatility, movements in short term interest rates, and bond option values. *Journal of Banking and Finance 21*, 169–196.

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

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