Robust Asset Allocation for Long-Term Target-Based Investing

Peter A. Forsyth[a] Kenneth R. Vetzal[b]

November 8, 2016

Abstract

This paper explores dynamic mean-variance asset allocation over long horizons. This is equivalent to target-based investing with a quadratic loss penalty for deviations from the target level of terminal wealth. We provide a number of illustrative examples in a setting with a risky stock index and a risk-free asset. Our underlying model is very simple: the value of the risky index is assumed to follow a geometric Brownian motion diffusion process and the risk-free interest rate is specified to be constant. We impose realistic constraints on the leverage ratio and trading frequency. In many of our examples, the mean-variance optimal strategy has a standard deviation of terminal wealth less than half that of a constant proportion strategy which has the same expected value of terminal wealth, while the probability of shortfall is reduced by a factor of two to three. We investigate the robustness of the model through resampling experiments using historical data dating back to 1926. These experiments also show much lower standard deviation and shortfall probability for the mean-variance optimal strategy relative to a constant proportion strategy with approximately the same expected terminal wealth.

JEL Classification: C63, G11

Keywords: long-term investment, mean-variance optimal asset allocation, dynamic asset allocation, target-based investing, quadratic loss

Acknowledgements: This work was supported by the Natural Sciences and Engineering Research Council of Canada and the Social Sciences and Humanities Research Council of Canada. Results presented in this paper were calculated (or derived) based in part on data from Historical Indexes © 2015 Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business. In addition, Wharton Research Data Services (WRDS) was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

[a]David R. Cheriton School of Computer Science, University of Waterloo, 200 University Avenue West, Waterloo ON, Canada N2L 3G1, paforsyt@uwaterloo.ca, +1 519 888 4567 ext. 34415

[b]Corresponding author: School of Accounting and Finance, University of Waterloo, 200 University Avenue West, Waterloo ON, Canada N2L 3G1, kvetzal@uwaterloo.ca, +1 519 888 4567 ext. 36518

1 Introduction

Savings for long-term objectives such as retirement or higher education are of paramount importance to investors. The value of total assets held in U.S. retirement accounts at the end of 2014 was almost \$25 trillion.[1] The size of this sum implies that even modest improvements in the management of these assets can offer significant economic benefits.

Target date funds implement an asset allocation strategy intended to help investors who want to have money available at a specified future date. As the target date is approached, "glide path" formulas are applied to automatically reduce allocations to riskier investments such as equities. Assets under management in target date funds (in the U.S.) have grown from \$44 billion in 2004 to over \$700 billion by the end of 2014.[2] However, the efficacy of target date funds is questionable: from a mean-variance (MV) perspective, deterministic glide path strategies are dominated by a constant proportion strategy which continuously maintains fixed asset weights over time, at least in our modelling framework (Graf, 2016).

We explore optimal long-term asset allocation for an investor who seeks to achieve a targeted level of terminal wealth where a quadratic loss function penalizes deviations from the target. This is equivalent to using a multi-period MV framework (Vigna, 2014); hence, in the following, we will refer to optimal long-term target-based strategies as MV optimal strategies. We consider the case where the investor makes a single lump sum initial investment, leaving accumulations and withdrawals for future work.

We will demonstrate that our strategy outperforms the constant proportion benchmark. We emphasize that this benchmark dominates any deterministic approach such as a glide path strategy, for a lump sum investment (Graf, 2016). Tests conducted under the assumed modelling conditions indicate that MV optimal strategies that produce the same expected terminal wealth as a constant proportion benchmark have a standard deviation of terminal wealth that is less than half that of the corresponding constant proportion benchmark, and shortfall probability is reduced by a factor of between two and three. In addition, resampling experiments based on long-run historical data lead to very similar relative comparisons, so our results are not an artifact of the assumed modelling conditions.

An alternative to the MV or target-based approach is to assume a particular utility function over terminal wealth, as in the classic approach pioneered by Merton (1969, 1971). Of course, since the MV approach can be justified on the basis of a quadratic utility function, we are implicitly assuming here that the alternative being considered is non-quadratic. Using such an alternative offers some advantages, but also has various disadvantages:

- Alternative utility functions avoid perceived undesirable characteristics of quadratic utility. For example, quadratic utility implies that as wealth rises beyond a satiation point, investors prefer less of it. Under the target-based approach the target level of terminal wealth corresponds to the satiation point for quadratic utility. However, in a model with continuous trading where the value of the risky asset follows geometric Brownian motion (GBM), Vigna (2014) shows that under an optimal policy the target will never be reached. More generally, Cui et al. (2012) argue that the target cannot be reached if markets are complete. In the incomplete market case, the investor can simply withdraw cash if the target level is exceeded, leaving enough in the risk-free asset to reach the target.
- Utility-based models suffer in practice from difficulty in determining the form and parameters of the utility function. Of course, specific investment decisions in the MV framework require

[1] Investment Company Institute (2015), Figure 7.5.
[2] Investment Company Institute (2015), Figure 7.26.

some estimate of risk-aversion to balance risk and anticipated rewards. However, the target-based interpretation of the MV framework may offer some scope for estimating risk-aversion, as discussed in Vigna (2014).

- In the context of asset allocation targeted towards a long-term objective, some popular utility specifications may imply unrealistic behavior. For example, with constant relative risk aversion (CRRA), the marginal utility of wealth at zero is infinite. Investors become extremely cautious if there is a chance that wealth could be near zero. In contrast, the quadratic loss function implicit in the MV criterion gives strong incentives to avoid insolvency, but not to the extreme extent that CRRA utility does.

- Standard dynamic programming techniques are not applicable in the multi-period MV context. However, Li and Ng (2000) and Zhou and Li (2000) have shown that this problem can be circumvented using an embedding result. As noted by Basak and Chabakauri (2010), the dynamic MV approach is not *time-consistent*. In other words, dynamic MV analysis relies on *pre-commitment*, under which investors pick an initial policy and stick with it no matter what happens subsequently. Numerous experimental studies have found evidence of time-inconsistent preferences (see Frederick et al., 2002, for a review). If preferences are time-inconsistent there may be an expanded role for government intervention to help individuals to attain pre-commitment policies (Kocherlakota, 2001). In addition, a time-consistent strategy can be constructed from a pre-commitment policy by imposing a constraint (Wang and Forsyth, 2011). This implies that time-consistent strategies are *sub-optimal* relative to pre-commitment policies.

Thus, the problem of long-term savings oriented towards specific objectives makes the dynamic MV approach a sensible choice.

Recent research on multi-period MV strategies includes studies by Bielecki et al. (2005), Basak and Chabakauri (2010), Cui et al. (2012), Vigna (2014), and Shen et al. (2014). All of these papers are analytical in nature. However, analytic solutions generally do not exist when realistic restrictions are imposed. Such constraints can be significant. Bielecki et al. (2005) discuss the impact of a no-bankruptcy condition, noting that the risk-return tradeoff is clearly less favorable when this condition is imposed. However, the Bielecki et al. solution does not have any limits on leverage and assumes continuous trading.

More realistic constraints can be imposed if numerical solution methods are used. Reliable and efficient numerical partial differential equation (PDE) techniques which can be applied in the continuous time dynamic MV context have been recently developed (Wang and Forsyth, 2011; Dang and Forsyth, 2014). Alternatively, Monte Carlo methods have also appeared recently (Cong and Oosterlee, 2016). Such procedures can be expected to be less efficient than numerical PDE techniques in the low-dimensional case, but they do offer the ability to handle cases involving many risky assets. In this paper, we use numerical PDE methods to investigate a number of properties regarding MV optimal portfolios over long horizons in a setting with a single risky asset. In particular, we make the following contributions:

- We provide an extensive set of comparisons between MV optimal investment strategies and a constant proportion benchmark. The MV optimal strategies we consider incorporate realistic constraints on leverage and trading frequency. We show that these MV optimal strategies can offer dramatically reduced risk compared to constant proportion strategies with the same expected terminal wealth.

- We run additional tests which suggest that our MV optimal strategies are quite robust to parameter uncertainty and model mis-specification. This implies that extensions to cases such as stochastic volatility models are unlikely to matter very much over long horizons, at least if volatility is mean-reverting and the speed of reversion is not slow.

- We provide backtesting results based on historical data which confirm that the MV optimal strategy significantly outperforms the constant proportion strategy in terms of risk reduction, while maintaining approximately the same expected terminal wealth. Our backtesting results are actually quite similar to what would be expected under idealized GBM modelling conditions with a constant risk-free rate. In addition to demonstrating the robustness of our simple modelling strategy, this reinforces the implication that more complicated modelling approaches involving stochastic volatility or stochastic interest rates may not offer significant advantages over long-term investment horizons.

In an Appendix, we also present a simple and intuitive geometric description of the embedding result (Li and Ng, 2000; Zhou and Li, 2000) which facilitates the use of standard dynamic programming methods to solve multi-period MV portfolio optimization problems.

2 The Model

We consider a model in which the investor chooses a strategy to allocate funds between a risky asset and a risk-free asset, with a terminal horizon date T. Denote by S_t the amount invested in the risky asset at time t. S_t is assumed to follow the GBM process

$$dS_t = \mu S_t \, dt + \sigma S_t \, dZ, \tag{1}$$

where μ is the appreciation rate, σ is the volatility, and dZ is the increment of a Wiener process. Let B_t be the amount invested at time t in the risk-free asset. Assume that B_t follows the process

$$dB_t = r B_t \, dt, \tag{2}$$

where r is the risk-free interest rate. We make the standard assumption that $\mu > r$, so that it is never optimal to short the risky asset, i.e. $S_t \geq 0$, $t \in [0, T]$. On the other hand, we do allow short positions in the risk-free asset ($B_t < 0$ is admissible).

The wealth that the investor has invested in this portfolio is $W_t = S_t + B_t$. As elucidated by Benjamin Graham (Graham, 2003, p. 89), a classic investment strategy for the *defensive investor* is to choose a fraction of wealth to invest in the risky asset and to dynamically rebalance the portfolio to preserve this ratio.

It is well-known that it is optimal for an investor with CRRA utility in this modelling framework to maintain constant fractions of the portfolio in the two assets. Note that if constant fractions p and $1-p$ are respectively allocated to the risky and risk-free assets and the portfolio is continuously rebalanced, then the process followed by wealth W_t is

$$dW_t = [(1-p)r + p\mu] \, W_t \, dt + p\sigma W_t \, dZ \tag{3}$$

and a closed form expression for the probability density of W_t follows easily.

The dynamic MV criterion can be viewed as having a target-based objective with a quadratic loss function where any penalties for exceeding the target can be avoided by withdrawing cash and leaving enough money in the risk-free asset to reach the target over the remaining horizon (Cui et al., 2012). To allow for this possibility, we have to be more specific about how wealth is defined.

Allocated wealth is the wealth available for allocation into the portfolio, and is denoted at time t by $W_t = S_t + B_t$. *Non-allocated wealth*, denoted by W_t^n, contains any cash withdrawals from the portfolio and accumulated interest.

We only consider the case where the investor rebalances the portfolio at discrete points in time. Denote the set of rebalancing times by $\mathcal{T} = \{t_0 = 0 < t_1 < \cdots < t_M = T\}$. Let t^- and t^+ respectively represent instants in time just before and just after time t, and let $X(t) = (S(t), B(t))$ be the underlying process. The rebalancing decision is captured by a control containing two components. We use $c(\cdot) \equiv (d(\cdot), e(\cdot))$ to denote this control as a function of the current state at t^- for rebalancing time t, i.e. $c\left(X(t^-), t^-\right) \equiv (d(X(t^-), t^-), e(X(t^-), t^-)) \equiv (d(t), e(t))$ for $t \in \mathcal{T}$. The control component $d(\cdot)$ denotes the non-negative amount of cash withdrawn from the portfolio. Note that this cash withdrawal is considered to occur immediately before rebalancing occurs at time t. The control component $e(\cdot)$ is the amount of allocated wealth that is invested in the risk-free asset after rebalancing at time $t \in \mathcal{T}$. If desired, restrictions such as leverage constraints can be imposed by specifying the set of admissible controls $\mathcal{Z}$, and so we have $c \in \mathcal{Z}$.

Let $x = (s, b)$ denote the state of the system at time t for $t \in [0, T]$. Given the prevailing state $x \equiv (s, b) = (S(t^-), B(t^-))$ at time t^-, $t \in \mathcal{T}$, we denote by $(S^+, B^+) \equiv (S^+(s, b, c), B^+(s, b, c))$ the state of the system at time t immediately after the application of the control $c \equiv (d, e)$. We then have

$$\begin{aligned} S^+\left(s, b, c \equiv (d, e)\right) &= (s + b) - d - e, \\ B^+\left(s, b, c \equiv (d, e)\right) &= e. \end{aligned} \tag{4}$$

Following Dang and Forsyth (2016), the investment strategies under consideration can be characterized by $d(t)$. In particular, a *self-financing* strategy requires $d(t) = 0$ for all $t \in \mathcal{T}$. A *semi-self-financing* strategy allows cash withdrawals, i.e. a control with $d(t) > 0$, $t \in \mathcal{T}$, is admissible. In the case of a semi-self-financing strategy, let t_α be the times where $d(t_\alpha) > 0$.

We can now provide a precise definition of non-allocated wealth W_t^n. At time $t \in [0, T]$,

$$W_t^n = \sum_{t_\alpha \leq t} d(t_\alpha) e^{r(t - t_\alpha)}. \tag{5}$$

In the following, we will refer to W_T^n as the *free cash* from the investment strategy. We can now also precisely specify solvency and leverage constraints. In particular, we enforce the solvency condition that the investor can continue trading only if

$$W(s, b) = s + b > 0. \tag{6}$$

If insolvency occurs, the investor must immediately liquidate all investments in the risky asset and cease trading, i.e. if $W(s, b) \leq 0$ then

$$S^+ = 0; \qquad B^+ = W(s, b). \tag{7}$$

The leverage constraint is enforced as a ratio, i.e. the investor has to choose an allocation such that

$$\frac{S^+}{S^+ + B^+} \leq q_{\max}, \tag{8}$$

where $q_{\max}$ is a specified parameter.

We respectively denote by $E_{c(\cdot)}^{x,t}[W_T]$ and $Var_{c(\cdot)}^{x,t}[W_T]$ the expectation and the variance of the terminal allocated wealth conditional on the state (x, t) and the control $c(\cdot)$, $t \in \mathcal{T}$. Let $(x_0, 0) \equiv (X(t = 0), t = 0)$ denote the initial state. Then the achievable MV objective set is

$$\mathcal{Y} = \left\{ (Var_{c(\cdot)}^{x_0,0}[W_T],\, E_{c(\cdot)}^{x_0,0}[W_T]) : c \in \mathcal{Z} \right\}. \tag{9}$$

For simplicity, assume that $\mathcal{Y}$ is a closed set and let $Var_{c(\cdot)}^{x_0,0}[W_T] = \mathcal{V}$ and $E_{c(\cdot)}^{x_0,0}[W_T] = \mathcal{E}$. For each point $(\mathcal{V}, \mathcal{E}) \in \mathcal{Y}$ and for an arbitrary scalar $\rho > 0$, we define the set of points $\mathcal{Y}_{P(\rho)}$ to be

$$\mathcal{Y}_{P(\rho)} = \left\{ (\mathcal{V}_*, \mathcal{E}_*) \in \mathcal{Y} : \rho\mathcal{V}_* - \mathcal{E}_* = \min_{(\mathcal{V},\mathcal{E}) \in \mathcal{Y}} (\rho\mathcal{V} - \mathcal{E}) \right\}. \tag{10}$$

ρ can be viewed as a risk-aversion parameter governing how the investor trades off expected value and variance. For a given ρ, $\mathcal{Y}_{P(\rho)}$ represents an efficient point in that it offers the highest expected value given variance. The set of points on the efficient frontier $\mathcal{Y}_P$ is the collection of these efficient points for all values of ρ, i.e.

$$\mathcal{Y}_P = \bigcup_{\rho > 0} \mathcal{Y}_{P(\rho)}. \tag{11}$$

As discussed in the Appendix, the presence of the variance term in (10) precludes determining $\mathcal{Y}_{P(\rho)}$ by solving for the associated value function using dynamic programming, but this can be circumvented with the embedding result (Li and Ng, 2000; Zhou and Li, 2000). We define the value function $V(x,t)$ as

$$V(x,t) = \min_{c \in \mathcal{Z}} \left\{ E_c^{x,t} \left[(W_T - \gamma/2)^2 \right] \right\}, \tag{12}$$

where the parameter $\gamma \in (-\infty, +\infty)$. The embedding result implies that there exists a $\gamma \equiv \gamma(x,t,\rho)$ such that for a given positive ρ, a control $c^* \equiv (d^*, e^*)$ which minimizes (10) also minimizes (12). The value function $V(s,b,t)$ can be found by solving the associated Hamilton-Jacobi-Bellman (HJB) equation (described shortly below) backward in time with the terminal condition

$$V(s,b,T) = (W(s,b) - \gamma/2)^2. \tag{13}$$

During this solution process, the optimal control c^* can be determined. We then use this control to find the quantity $U(s,b,t) = E_{c^*}^{x,t}[W_T]$ since this information is needed to determine the corresponding variance and expected value point $\left(Var_{c^*}^{x_0,t=0}[W_T], E_{c^*}^{x_0,0}[W_T] \right)$ on the efficient frontier. This step involves solving an associated linear PDE (Dang and Forsyth, 2014). Given the optimal control c^*, it is straightforward to determine other quantities of interest by using Monte Carlo simulation.

We now describe the HJB PDE which is used to determine c^*. Define the *solvency* region as $\mathcal{S} = \{(s,b) \in [0,\infty) \times (-\infty,+\infty) : W(s,b) > 0\}$, and the *bankruptcy* region as $\mathcal{B} = \{(s,b) \in [0,\infty) \times (-\infty,+\infty) : W(s,b) \le 0\}$. Also define the diffusion operator $\mathcal{L}V$ as

$$\mathcal{L}V \equiv \frac{\sigma^2 s^2}{2} \frac{\partial^2 V}{\partial s^2} + \mu s \frac{\partial V}{\partial s} + rb \frac{\partial V}{\partial b}, \tag{14}$$

and the intervention operator $\mathcal{M}(c)V(s,b,t)$ as

$$\mathcal{M}(c)V(s,b,t) = V(S^+(s,b,c), B^+(s,b,c), t). \tag{15}$$

At each portfolio rebalancing time $t = t_i \in \mathcal{T}$, we apply the following conditions:

- If $(s,b) \in \mathcal{B}$, we enforce the insolvency condition (i.e. the investor must liquidate all investment in the risky asset and cease trading)

$$V(s,b,t_i^-) = V(0, W(s,b), t_i^+). \tag{16}$$

- If $(s, b) \in \mathcal{S}$, we impose the rebalancing optimality condition

$$V(s, b, t_i^-) = \min_{c \in \mathcal{Z}} \mathcal{M}(c) V(s, b, t_i^+). \tag{17}$$

For the special case where $t = t_M = T$, the terminal condition (13) holds. Within each time period $[t_{i-1}, t_i^-)$, $i = M, \ldots, 1$, the following considerations apply:

- If $(s, b) \in \mathcal{B}$, we enforce the insolvency condition

$$V(s, b, t) = V(0, W(s, b), t). \tag{18}$$

- If $(s, b) \in \mathcal{S}$, then $V(s, b, t)$ satisfies the PDE

$$V_t + \mathcal{L}V = 0, \tag{19}$$

  subject to the initial condition (17).

For computational purposes, we localize the above problem and apply suitable asymptotic boundary conditions (Dang and Forsyth, 2014)

Pre-commitment MV portfolio optimization is equivalent to maximizing the expectation of a quadratic utility function provided that $W_T \leq \gamma/2$. Hence, $\gamma/2$ can be viewed essentially as the terminal allocated wealth target for the portfolio. Define the *discounted terminal allocated wealth target* of the portfolio at time $t \in [0, T]$ as

$$F_t = \frac{\gamma}{2} e^{-r(T-t)}. \tag{20}$$

Dang and Forsyth (2016) note that for all $t \in [0, T]$, the state $(s, b) = (0, F_t)$ is a time t globally optimal solution to the value function $V(s, b, t)$, i.e. $V(0, F_t, t) = 0$ for all $t \in [0, T]$. The value function is nonnegative, so the point where it is zero is clearly a global minimum. Therefore any admissible policy which allows moving to this point is an optimal one. Once this point is attained, it is optimal to remain at it.

As previously mentioned, Vigna (2014) shows that under some conditions allocated wealth will never exceed this target, i.e. $W_t \leq F_t$ for all $t \in [0, T]$. More specifically, Vigna considers the case of continuously rebalanced self-financing strategies (i.e. $d(t) = 0$ for all $t \in [0, T]$) with no leverage constraints and trading continuing even under bankruptcy. With discrete rebalancing, we cannot be sure that $W_t \leq F_t$. However, Dang and Forsyth (2016) show that if W_t ever exceeds F_t at a discrete rebalancing time $t \in \mathcal{T}$, the MV optimal strategy is to withdraw a cash amount of $W_t - F_t$ from the portfolio and to invest the remainder F_t in the risk-free asset. We refer to this as the *optimal semi-self-financing strategy*. With this in mind, if we follow this strategy, then the terminal condition (13) is replaced by

$$V(s, b, T) = \left(\min(s + b, \gamma/2) - \gamma/2\right)^2. \tag{21}$$

3 Parameter Estimation

Our subsequent analysis is based on parameters estimated from U.S. market history. For the risky asset we use the NYSE/AMEX/NASDAQ/ARCA value-weighted index with distributions from the

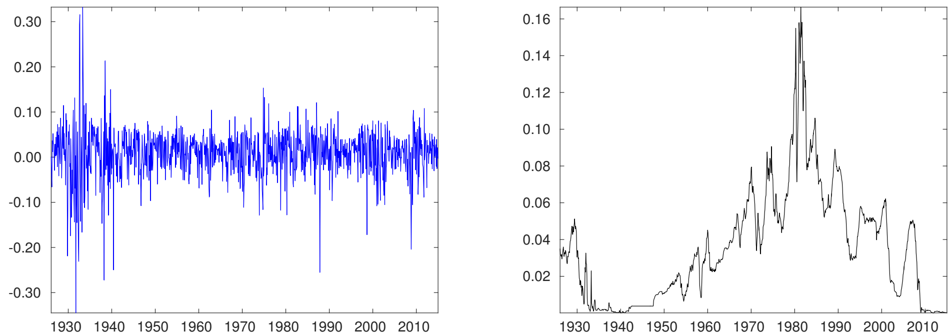

(a) *Stock index log returns R.* (b) *Annualized risk-free interest rate r.*

FIGURE 1: *Stock index returns and risk-free interest rate. Monthly data, 1926-2014.*

Center for Research in Security Prices (CRSP). The full sample period runs from January 1926 through December 2014.[3][4] Letting the level of this index as of month i be S_i, the log return R for month i is $R_i = \ln(S_i/S_{i-1})$, which is shown in Figure 1(a). We derive the risk-free interest rate from secondary market 3-month U.S. Treasury bill rates obtained from the Federal Reserve[5] (*1934:1-2014:12*) and the National Bureau of Economic Research (NBER) Macrohistory Database[6] (*1926:1-1933:12*). The parts of our analysis which cover the period prior to 1934 are then subject to the caveat that the risk-free rate measured over this period is somewhat different compared to other time periods. We convert the Treasury bill yields to continuously compounded annual rates and denote them by r. These are plotted in Figure 1(b).

Table 1 provides descriptive statistics for monthly stock index log returns R and the annualized risk-free rate r over the full 1926-2014 sample and three sub-periods. The first sub-period runs from 1926 to 1954. The second sub-period contains the first, but is extended through 1984.[7] The final sub-period (1955-2014) covers the most recent six decades and may perhaps be viewed as more representative of what to expect going forward. Average monthly stock index log returns are relatively stable at 0.7%-0.8% across the four periods considered in Table 1. In each case, median values are notably higher, reflecting negative skewness. Consistent with numerous prior studies, observed kurtosis levels far exceed that of the normal distribution. The risk-free interest rate averaged about 3.5% over the full sample. However, it varied significantly across time, ranging from just over 1% in the first sub-period to 4.7% in the last one. Median values of r lie somewhat below these means, consistent with observed positive skewness, and kurtosis is relatively high. Over all periods considered, R and r consistently display slightly negative correlations. Collectively,

[3]See Chapter 4 of CRSP (2012) for a detailed description.

[4]The number of securities included in this index changes significantly over time. Only NYSE-listed securities are included prior to July 1962, when AMEX-listed securities first become available. NASDAQ-listed securities are added as of December 1972, followed by stocks listed on the ARCA exchange in March 2006. However, our parameter estimates do not change appreciably if we restrict attention over the entire sample period to just NYSE listings, only NYSE and AMEX listings, or NYSE/AMEX/NASDAQ listings.

[5]See `www.federalreserve.gov/releases/h15/data.htm`.

[6]See `www.nber.org/databases/macrohistory/contents/chapter13.html` In particular, we use the `m13029a` series from January 1926 to December 1930 and the `m13029b` series from January 1931 to December 1933. Neither series is strictly comparable to the Federal Reserve Treasury bill yield data which starts in 1934. The `m13029a` series is based on U.S. Treasury 3-month and 6-month notes and certificates, except for the April-June 1928 period which uses 6-month and 9-month certificates. We switch to the `m13029b` series when it starts in January 1931 since it is based on 3-month maturities over the 1931-1933 period.

[7]The rationale for choosing these two sub-periods will become clear below in Section 5.

	1926:1 - 2014:12 ($N = 1068$)		1926:1 - 1954:12 ($N = 348$)		1926:1 - 1984:12 ($N = 708$)		1955:1 - 2014:12 ($N = 720$)	
	R	$\bar{r}$	R	$\bar{r}$	R	$\bar{r}$	R	$\bar{r}$
Mean	0.0078	0.0354	0.0070	0.0113	0.0073	0.0346	0.0082	0.0470
Median	0.0130	0.0311	0.0130	0.0043	0.0114	0.0274	0.0130	0.0464
Std. dev.	0.0539	0.0312	0.0707	0.0124	0.0577	0.0337	0.0436	0.0309
Skewness	-0.5651	1.0830	-0.3570	1.3806	-0.3718	1.3396	-0.8071	0.8388
Kurtosis	9.8982	4.3961	8.5445	3.9841	9.9557	4.6839	5.9129	4.2703
Correlation	-0.0164		-0.0090		-0.0351		-0.0381	
Drift $\hat{\mu}$	0.1108 (0.0206)		0.1134 (0.0476)		0.1077 (0.0271)		0.1095 (0.0200)	
Volatility $\hat{\sigma}$	0.1867 (0.0085)		0.2445 (0.0180)		0.1998 (0.0112)		0.1511 (0.0062)	

TABLE 1: *Descriptive statistics for stock index monthly log returns R and annualized risk-free interest rate $\bar{r}$. N is the number of monthly observations. $\hat{\mu}$ and $\hat{\sigma}$ are the annualized maximum likelihood estimates of the risky asset drift and volatility respectively, with robust standard errors in parentheses. $\bar{r}$ is the average continuously compounded risk-free interest rate over the indicated period.*

these statistics indicate that our model which assumes simple GBM for risky asset returns and a constant risk-free rate is seriously mis-specified from an econometric standpoint. However, we will argue below that this is not a major concern in our context of MV optimality with a long-term horizon.

Our model has three parameters: the risky asset drift μ, its volatility σ, and the risk-free interest rate r. Given the GBM specification, it is straightforward to calculate maximum likelihood estimates of μ and σ using the log returns R. For the risk-free rate, we simply use the average value of the continuously compounded annualized rate r. The estimated drift and volatility parameters given in Table 1 are expressed in annualized terms.

4 Illustrative Examples

We now present an extensive set of illustrative examples.[8] The main benchmark to which we compare the MV optimal strategy is a constant proportion strategy in which the investor continuously maintains a fixed fraction of wealth in the risky asset, which can be regarded as a common default strategy. We remind the reader that for a lump sum investment, the fixed fraction strategy MV-dominates any continuously rebalanced strategy where the fraction invested in the risky asset is a deterministic function of time (Graf, 2016).

We begin with a general comparison between a constant proportion strategy with an even split between the risky and risk-free assets and the MV optimal strategy derived by Bielecki et al. (2005). This latter strategy enforces the restriction that the investor's wealth cannot ever be negative, but assumes continuous rebalancing and allows for infinite leverage.

Given the parameters, we calculate the mean and standard deviation of terminal wealth for the constant proportion strategy. We next determine the standard deviation of terminal wealth for the MV optimal strategy subject to the restriction that the mean for this strategy be the same as that of the constant proportion strategy. We then calculate the ratio of the standard deviation of the

[8]From this point on, we will use "expected value" and $E[W_T]$ in place of the more cumbersome $E^{x_0,0}_{c(\cdot)}[W_T]$. Similarly, "standard deviation" will refer to the standard deviation of terminal wealth as of time 0 conditional on the initial state and the investment strategy.

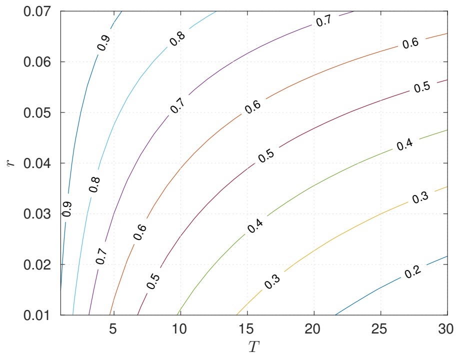

(a) *Risk-free rate r and horizon T.*

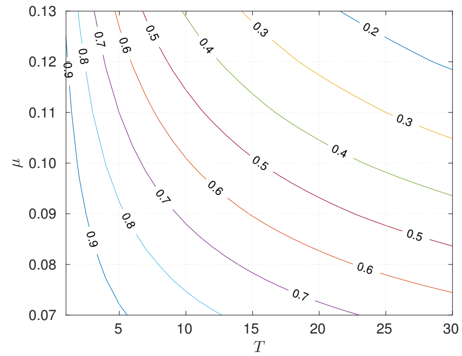

(b) *Drift rate μ and horizon T.*

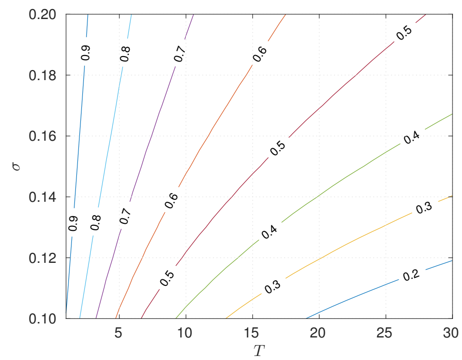

(c) *Volatility σ and horizon T.*

FIGURE 2: *Contour plots of the ratio of the standard deviation of terminal wealth for the MV optimal strategy from Bielecki et al. (2005) to the corresponding standard deviation for the constant proportion strategy ($p = .5$). Both strategies assume continuous rebalancing and have the same expected terminal wealth. The MV optimal strategy allows for infinite leverage.*

MV optimal strategy to that of the constant proportion strategy. Obviously, this ratio is at most unity.

Figure 2(a) plots contours of the ratio defined above for values of the risk-free rate r ranging from .01 to .07 and for horizons of $T = 1, 2, \ldots, 30$ years. The remaining parameters used are $\mu = .10$ and $\sigma = .15$, broadly consistent with our empirical estimates over the past 60 years from Table 1. As a guide to interpretation, along the curve marked "0.5" the standard deviation of terminal wealth for the MV optimal strategy is 50% of the corresponding standard deviation for the constant proportion strategy, even though both strategies have the same expected value of terminal wealth. If the horizon T is close to 30 years the ratio is around 20% when $r \simeq .02$. This means that the MV optimal strategy offers the same expected level of terminal wealth but has only about one-fifth of the risk, as measured by the standard deviation of terminal wealth. As another example, with $T = 10$ years the MV optimal strategy has less than half of the risk of the constant proportion strategy with equivalent expected terminal wealth if $r \simeq 2.5\%$ or lower.

Figure 2(b) provides a similar comparison for different levels of the risky asset drift rate μ, with $\sigma = .15$ and $r = .04$. This plot also shows some significant potential improvements over the constant proportion strategy, especially for long horizons and high values of μ. If $\mu \simeq .10$, the MV optimal strategy has at most 50% of the risk of the constant proportion strategy for $T \geq 15$ years. Figure 2(c) in turn shows contours for the ratio of standard deviations across a range of values for the volatility σ, given $\mu = .10$ and $r = .04$. These results again show marked improvement over

	Constant Proportion	MV Optimal (Analytic)	MV Optimal (Numerical)
Investment horizon T (years)	30	30	30
Initial investment W_0	100	100	100
Risky asset drift rate μ	0.10	0.10	0.10
Risky asset volatility σ	0.15	0.15	0.15
Risk-free rate r	0.04	0.04	0.04
Rebalancing frequency	Continuous	Continuous	Annual
Insolvency condition	n.a.	Yes	Yes
Maximum leverage ratio $q_{\max}$	n.a.	∞	1.5

TABLE 2: *Base case data. "MV Optimal (Analytic)" refers to the model of Bielecki et al. (2005) which has a closed form solution. In this model the insolvency condition ensures that wealth can never be negative. "MV Optimal (Numerical)" refers to the model in this paper which must be solved numerically. In this model the insolvency condition is that if wealth becomes negative, the investor must immediately liquidate the investment in the risky asset and stop trading. The maximum leverage ratio $q_{\max}$ is defined in equation* (8).

the constant proportion strategy, particularly for relatively low values of σ and over long horizons. With $\sigma \simeq .15$, the standard deviation ratio is at most .5 for $T \geq 15$ years. A clear pattern displayed in all three of these contour plots is that the ratio of the standard deviations drops markedly as maturity rises, indicating that the superiority of the MV optimal strategy increases significantly with the investment horizon.

Collectively, the plots in Figure 2 show that an MV optimal strategy can offer significant benefits over a constant proportion strategy, especially for longer horizons and for low values of r and σ. However, recall that this MV optimal strategy assumes continuous rebalancing (as does the constant proportion strategy) and it also allows for infinite leverage. This raises the possibility that imposing realistic constraints may significantly erode the advantages of the MV optimality criterion. Enforcing such restrictions requires the use of a numerical approach, as described above.

The examples to follow are thus based on three different types of strategies: (i) constant proportion with continuous rebalancing (analytic solution available); (ii) the MV optimal strategy of Bielecki et al. (2005) which assumes continuous rebalancing and allows for infinite leverage while enforcing an insolvency condition which ensures that the investor's wealth cannot become negative (analytic solution available); and (iii) our MV optimal strategy which imposes a different insolvency condition (i.e. if wealth becomes negative, the investor has to immediately liquidate the investment in the risky asset and stop trading) and in addition restricts the leverage ratio and the rebalancing frequency (numerical solution required). Base case input data for all three types are provided in Table 2. The market parameters μ, σ, and r are specified as .10, .15, and .04 respectively, reflecting the parameter estimates from Table 1 with particular focus on the most recent 60 years. We also designate an initial investment of 100 and an investment horizon of 30 years. The strategy based on the numerical solution has a base case maximum leverage ratio of $q_{\max} = 1.5$ and annual rebalancing.

As a benchmark base case, we show the results for the constant proportion strategy with continuous rebalancing in Table 3. Next, we consider MV optimal strategies under various assumptions: maximum leverage ratio of $q_{\max} = 1.5$ with annual rebalancing, unrestricted leverage with annual rebalancing, and unconstrained leverage with continuous rebalancing. When required, the numer-

Constant Proportion	Expected Value	Standard Deviation	Shortfall Probability
$p = 0.0$	332.01	n.a.	n.a.
$p = 0.5$	816.62	350.12	$Prob(W_T < 800) = 0.56$
$p = 1.0$	2008.55	1972.10	$Prob(W_T < 2000) = 0.66$

TABLE 3: *Results for the constant proportion strategy using base case data from Table 2.*

ical procedures used are described in Dang and Forsyth (2014). The semi-self-financing optimal withdrawal policy described in Section 2 is applied.

We specify the target expected value to be 816.62, as reported in Table 3 for the $p = 0.5$ case. Convergence test results are given in Table 4. For a given grid refinement level, we calculate the γ which generates this target expected value by Newton iteration. The fourth column shows the total expected value, including the free cash from the withdrawal that occurs if the allocated wealth exceeds the discounted target. Even with annual rebalancing, the magnitude of the expected free cash is not large, and so we exclude it from subsequent reported expected values in this section. In the case with unrestricted leverage and annual rebalancing, extrapolating the results given in Table 4 gives an ultimate standard deviation of about 124.51. Comparing this figure to the exact closed form solution (also with unlimited leverage), we see that the effect of annual rebalancing over the 30-year horizon compared to continuous rebalancing is quite small.[9] This is encouraging, as it suggests that transaction costs are unlikely to have a significant impact on our results since we do not need to trade frequently. All results provided subsequently use level 3 grid refinement.

As a point of comparison, consider the case with annual rebalancing and $q_{\max} = 1.5$. Table 4 reports an expected value of 816.62 and a standard deviation of 142.85 (refinement level 3). Although not shown in the table, the probability of having terminal wealth below 800 for this case is 19%. From Table 3, a continuously rebalanced strategy with fixed weight $p = 0.5$ has the same expected value, a standard deviation of 350.12, and a 56% chance of terminal wealth below 800. The MV optimal strategy considered here that produces the same expected wealth as the constant proportion strategy reduces the standard deviation by a factor of about 2.5 and the shortfall probability by a factor of almost 3. This is quite a dramatic improvement in terms of classical measures of portfolio efficiency.

As an additional comparison, Table 3 shows that a constant proportion strategy with $p = 1.0$ produces an expected value of 2008.55, a standard deviation of 1972.10, and a 66% probability of having final wealth lower than 2000. The MV optimal strategy for the same expected value of 2008.55 with annual rebalancing and $q_{\max} = 1.5$ turns out to have a standard deviation of 969.33 and a 40% probability of having terminal wealth below 2000. In other words, relative to the constant proportion case the standard deviation is halved and the shortfall probability reduced by a factor of about 1.65.

Figure 3 compares the cumulative distribution functions of W_T for the two cases considered above. Focusing on panel (a) where the constant proportion strategy has $p = 0.5$, we observe that the outcomes from the MV optimal strategy are much more clustered near the expected value of 816.62, with a very narrow range of possible outcomes above this amount. The MV optimal strategy produces a smaller probability that $W_T < W^*$ for $360 < W^* < 800$. For $W^* < 360$, the

[9]Dang and Forsyth (2014) have verified that the numerical method used here converges to the solution from Bielecki et al. (2005) under the same assumptions.

Refinement	Expected Value	Standard Deviation	Expected Value (Including Free Cash)	γ
		Maximum leverage $q_{\max} = 1.5$, annual rebalancing		
0	816.62	161.00	840.95	1804.03
1	816.62	151.09	832.32	1769.79
2	816.62	145.72	827.64	1756.83
3	816.62	142.85	824.92	1751.94
		Infinite leverage, annual rebalancing		
0	816.62	146.69	843.75	1746.73
1	816.62	136.47	833.38	1731.53
2	816.62	130.71	828.42	1723.94
3	816.62	127.61	825.37	1720.76
		Exact solution: infinite leverage, continuous rebalancing		
n.a.	816.62	118.84	n.a.	n.a.

TABLE 4: *Level 0 refinement used: s nodes: 65, b nodes: 129, timesteps: 60. Numbers of nodes, timesteps doubled on each refinement. The localized domain uses $s_{\max} = 20{,}000$, with $b_{\max} = |b_{\min}| = s_{\max}$. A nonuniform grid was used, with a fine mesh spacing near $s = W_0$ and $b = W_0$, and increasing exponentially as $s, |b| \to s_{\max}$. The base case input data used is provided in Table 2. The exact solution is from Bielecki et al. (2005). The equivalent wealth target is $\gamma/2$.*

fixed proportion strategy is better than the MV optimal one, but these are very low probability cases, i.e. $Prob(W_T < 360) \simeq .04$.

The results for the case with an expected value of 2008.55 (where the fixed weight strategy uses $p = 1.0$) shown in Figure 3(b) are qualitatively similar. At extremely low wealth levels, the constant proportion strategy outperforms the MV optimal allocation. Pre-committing to the MV optimal allocation reduces the chance of being significantly below the target over a wide range, but also sacrifices the upside with very high levels of wealth.

To gain insight into the properties of the controls for the MV optimal strategy, we return to the base case example (Table 2) with expected terminal wealth of 816.62. We numerically solve for the optimal controls and store them in a table. Figure 4(a) presents this table in the form of a heat map representing the fraction of wealth that is optimally invested in the risky asset. The darkest region near the top corresponds to wealth levels that are high enough to make investing fully in the risk-free asset optimal. In the brightest area near the bottom it is optimal to invest as much as possible in the risky asset by using the maximum permissible degree of leverage.[10] The dark strip along the bottom edge for non-positive wealth levels reflects the insolvency condition (7), which dictates that the investor must liquidate all investments in the risky asset if $W_t \leq 0$.

Figure 4(a) shows the optimal strategy across the range of wealth levels over time, but it is uninformative about probabilities. To address this, we run a Monte Carlo experiment. We begin with an initial investment of $W_0 = 100$, as in the lower left corner of Figure 4(a). We then simulate the performance of the risky stock index over the 30-year horizon across 1 million paths. Along each path, we calculate the value of wealth given our portfolio weights and the performance of the index. At the end of each year the portfolio is rebalanced in accordance with our optimal control strategy. We then calculate the mean and standard deviation of the fraction of wealth p in the

[10]The staircase pattern most visible near the edges of this area is due to the annual rebalancing frequency.

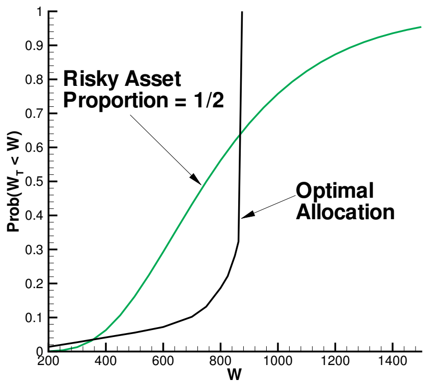

(a) $E[W_T] = 816.62$.

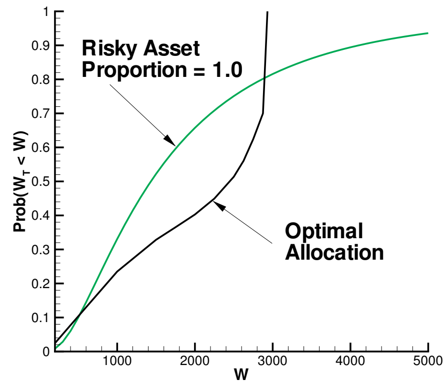

(b) $E[W_T] = 2008.55$.

FIGURE 3: *Comparison of cumulative distribution functions for the constant proportion policy and the MV optimal strategy. The base case data used are given in Table 2.*

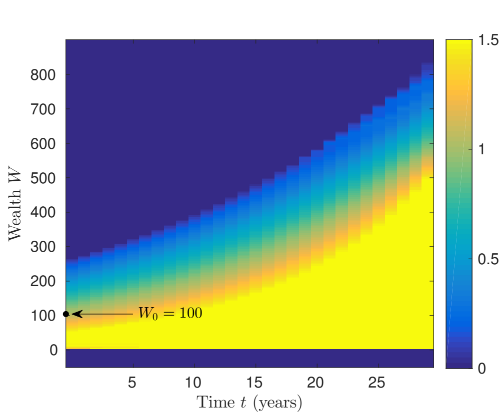

(a) *Heat map of the MV optimal control strategy for base case The shading indicates the fraction of wealth that is optimally invested in the risky asset.*

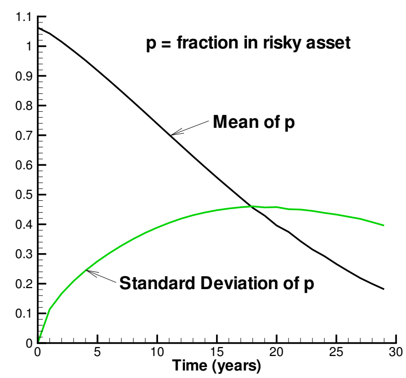

(b) *Evolution of the fraction of wealth invested in the risky asset (p) over time for the base case parameters wht an expected value of terminal wealth of 816.62.*

FIGURE 4: *Heat map and mean value of optimal MV strategy. Base case parameters (see Table 2).*

risky asset over time along each path. Figure 4(b) shows that the MV optimal strategy starts out with a mean of p slightly over one, indicating a modestly levered strategy. The standard deviation of p is quite low initially. As time passes, the mean of p drops considerably.

A deterministic linear glide path strategy can be defined as

$$p(t) \quad = \quad p_{\max} + (p_{\min} - p_{\max}) \frac{t}{T} \ , \tag{22}$$

which would behave in a roughly similar way to the mean of p in Figure 4(b), with suitable choices for $p_{\max}$ and $p_{\min}$. However, the MV optimal strategy improves on this by factoring in the target and prevailing levels of wealth, as well as time. A constant proportion strategy exploits long-run mean-reversion by selling assets following price increases and buying assets after their values decline. The MV optimal strategy also purchases assets following declines in their prices and selling assets after price increases, but in a more sophisticated way.

We next explore the effect of altering some of the constraints on the admissible controls. Table 5 shows the results obtained by perturbing some of the parameters for the base case MV optimal

MV Optimal Strategy	Expected Value	Standard Deviation	Shortfall Probability
Base case ($q_{\max} = 1.5$)	816.62	142.85	$Prob(W_T < 800) = 0.19$
No leverage ($q_{\max} = 1.0$)	816.62	162.54	$Prob(W_T < 800) = 0.21$
No withdrawal	816.62	144.49	$Prob(W_T < 800) = 0.20$

TABLE 5: *Perturbations of the MV optimal strategy. The base case data are given in Table 2. The base case results are the same as in Table 4, and are reproduced here for convenience. The no withdrawal case precludes using the optimal semi-self-financing strategy.*

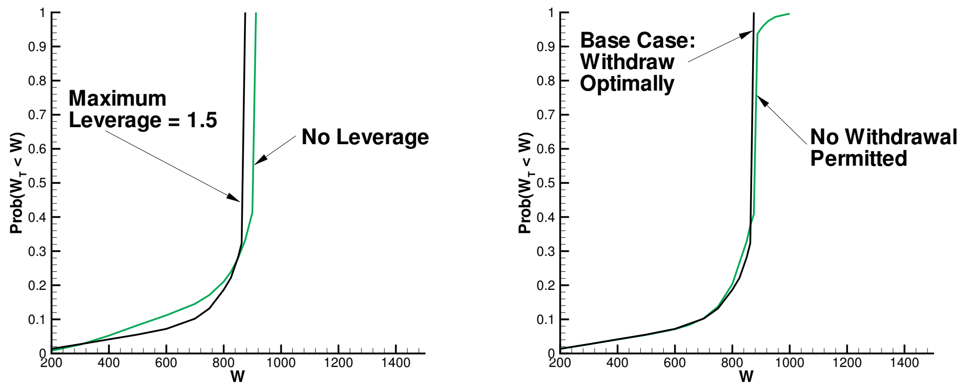

(a) *No leverage.* (b) *No withdrawal.*

FIGURE 5: *Comparison of cumulative distribution functions for the base case optimal strategy and the perturbations from Table 5. The base case data are given in Table 2.*

strategy. The corresponding cumulative distribution functions are compared with the base case in Figure 5. The general effect of reducing the maximum leverage permitted from $q_{\max} = 1.5$ to $q_{\max} = 1.0$ (i.e. no borrowing) is fairly small in terms of standard deviation. The effect of not using the globally optimal strategy of withdrawing cash when $W_t > F_t$ is negligible. An additional observation that is not apparent from Table 5 is that rebalancing more than once per year has little impact. In several numerical experiments we found that the difference between continuous and discrete rebalancing was very small if the number of rebalances exceeded 25, regardless of the time horizon.

Table 6 investigates the impact of shortening the investment horizon T. We use the data in Table 2, except that we reduce T to 15 years and we rebalance semi-annually rather than annually. As with the longer horizon cases considered earlier, the MV optimal policy which has the same expected value of terminal allocated wealth as the fixed proportion policy has significantly smaller standard deviation compared to that strategy. The probability of shortfall at the points given in Table 6 is also reduced substantially. The improvements are not as dramatic as those seen above for the longer time horizon, reaffirming the implication from Figure 2 that the superiority of MV optimal strategies increases significantly with the investment horizon.

We next consider the influence of the leverage ratio constraint. For long-term investments, Table 5 above indicates that the effect of allowing use of a maximum leverage ratio of $q_{\max} = 1.5$ as compared to no leverage ($q_{\max} = 1.0$) is not large. However, the impact of using leverage may be

Strategy	Expected Value	Standard Deviation	Shortfall Probability
Constant proportion $p = 0.5$	285.77	84.79	$Prob(W_T < 250) = 0.38$
MV optimal	285.77	48.96	$Prob(W_T < 250) = 0.13$
Constant proportion $p = 1.0$	448.17	283.95	$Prob(W_T < 400) = 0.54$
MV optimal	448.17	180.44	$Prob(W_T < 400) = 0.35$

Table 6: *Results for a time horizon of 15 years with other input data from Table 2. The MV optimal strategies use semi-annual rebalancing, while the constant proportion strategies are rebalanced continuously.*

Strategy	$q_{\max}$	Rebalancing	Expected Value	Standard Deviation
Constant proportion ($p = 1.0$)	n.a.	Continuous	271.83	88.15
MV optimal	1.0	Quarterly	271.83	88.09
MV optimal	1.5	Quarterly	271.83	34.90
MV optimal	10.0	Quarterly	271.83	12.67
MV optimal (exact)	∞	Continuous	271.83	11.59

Table 7: *Effect of leverage with extreme parameter values of $\sigma = .10$ and $r = 0$. The initial investment $W_0 = 100$, the investment horizon $T = 10$ years, and the risky asset drift rate $\mu = .10$. The MV optimal (exact) case is the analytic solution from Bielecki et al. (2005) which assumes continuous rebalancing and unrestricted leverage.*

bigger for shorter-term horizons. To illustrate, we consider a case with a 10-year horizon, quarterly rebalancing, and an initial investment of $W_0 = 100$. In addition, we keep μ at its base case value of .10 but we use extremely low values of $\sigma = .10$ and $r = 0$. The results shown in Table 7 indicate that allowing more leverage in this setting reduces risk dramatically.

Finally, we explore some effects of parameter uncertainty. We consider ranges for the market parameters of $r \in [.02, .06]$, $\mu \in [.06, .14]$, and $\sigma \in [.10, .20]$. We compute and store the optimal controls using the base case values for these parameters from Table 2, corresponding to the mean values of the ranges listed above if the parameters are uniformly distributed over the ranges. The remaining input data are as in Table 2, so these are actually the same controls used to produce the base case results reported earlier in Table 5. We then run 1 million Monte Carlo simulation paths. At every rebalancing time along each path, we randomly draw each of r, μ, and σ from its associated range, assuming uniform distributions, but we implement the control strategy based on the stored mean values. We find an average terminal wealth of 807, with a standard deviation of 145. There is a 19% probability that $W_T \leq 800$. These values are quite similar to the corresponding ones from Table 5, which were calculated assuming that the parameters were actually equal to the mean values. This indicates that parameter uncertainty does not have a large impact on the properties of the MV optimal strategy, at least under the conditions of this test. This suggests that modelling extensions such as mean-reverting stochastic interest rates and volatility may not be very important in this long-term setting, unless there is a significant chance of a large and persistent deviation from the long-run reversion level. Numerical tests which confirm that mean-reverting stochastic volatility effects are unimportant for a long term investor are given in Ma and Forsyth

Strategy	Expected Value	Standard Deviation	Shortfall Probability
	Estimation Period: 1926:1 to 1954:12		
Constant proportion ($p = .5$)	649	488	$Prob(W_T < 525) = .51$
MV optimal	649	166	$Prob(W_T < 525) = .15$
	Estimation Period: 1926:1 to 1984:12		
Constant proportion ($p = .5$)	845	499	$Prob(W_T < 725) = .50$
MV optimal	845	223	$Prob(W_T < 725) = .16$
	Estimation Period: 1926:1 to 2014:12		
Constant proportion ($p = .5$)	896	490	$Prob(W_T < 800) = .51$
MV optimal	896	197	$Prob(W_T < 800) = .14$

Table 8: *Estimated results for a continuously rebalanced constant proportion strategy with $p = .5$ and the MV optimal strategy with yearly rebalancing and maximum leverage ratio $q_{\max} = 1.5$. The investment horizon is $T = 30$ years and the initial investment is $W_0 = 100$. These results assume that the risk-free rate is constant and the value of the risky asset follows a GBM process. See Table 1 for each set of parameter estimates.*

(2016).

5 Historical Backtests

This section provides backtests of the MV optimal strategy using the same historical U.S. market data that was used above for parameter estimation. These tests will investigate the robustness of the strategy in the presence of an imperfectly known stochastic process for the risky asset and the risk-free rate.

While our earlier examples were based on parameter estimates that roughly reflected the last six decades (i.e. the sample period from 1955 through 2014) from Table 1, the tests reported in this section are based on the parameter estimates from the sub-periods 1926 to 1954 and 1926 to 1984, as well as the full 1926 to 2014 sample. Table 8 provides the results based on these estimation periods for a continuously rebalanced constant proportion ($p = .5$) strategy in terms of expected value, standard deviation, and shortfall probability. Also shown are the corresponding results for the MV optimal strategy that has the same expected value with annual rebalancing and a maximum leverage ratio $q_{\max} = 1.5$. In each case, Table 8 shows that the MV optimal strategy can be expected to significantly outperform the constant proportion benchmark. Of course, this assumes GBM and a constant risk-free rate. The tests to follow are based on observed historical values for risky returns and the risk-free rate.

5.1 Out of sample tests

Our ability to conduct out-of-sample tests is severely restricted by our focus on a long-term horizon, a relatively short period of available historical data, and the need for a reasonably long period for parameter estimation. Nevertheless, we can gain some useful insights about the MV optimal strategy by conducting two out-of-sample experiments.

First, consider an investor who uses the available data from 1926:1 to 1954:12 to estimate market parameters assuming that the market index follows GBM and estimating r as the average 3-month risk-free rate during that time. Table 1 shows that the resulting estimates are $\hat{\mu} = .1134$, $\hat{\sigma} = .2445$, and $\bar{r} = .0113$. From these values, the investor estimates the returns for a continuously rebalanced constant weight strategy ($p = .5$) over the next 30 years. Based on the estimated expected value of wealth for this strategy at the end of the horizon, the investor then determines the MV optimal strategy which generates this same expected wealth as well as the standard deviation and shortfall probability for this strategy. The results are shown in Table 8 under "Estimation Period: 1926:1 to 1954:12".

We store the MV optimal controls used to generate the results in Table 8 and then apply them using the actual historical returns and risk-free rate observed over the 30-year period from 1955:1 to 1984:12. This simulates the performance of an investor who (i) computes the MV optimal strategy assuming GBM for the risky asset and a constant risk-free interest rate estimated using known historical data as of 1954:12; and (ii) pre-commits to following this strategy for the period beginning in 1955:1 and ending in 1984:12. An initial investment of 100 at the start of 1955 is used for both the MV optimal strategy and the constant proportion strategy ($p = .5$). Since the risk-free rate is derived from 90-day Treasury bill yields, we assume that the rate prevailing at the start of each quarter holds throughout the quarter.[11] Each strategy is rebalanced annually. We do not apply the semi-self financing version of the MV optimal strategy (i.e. we do not withdraw funds in excess of the minimum amount that the model indicates is required to meet the target wealth level, instead we invest this free cash in the risk-free asset).[12] The results are plotted in Figure 6(a). Both strategies performed very well, exceeding the expected terminal wealth of 649 from Table 8 by large margins. However, the MV optimal strategy was clearly superior, following a path that was not only higher but also much smoother, particularly towards the end of the period. The primary reason for this was that the general level of the risk-free rate was much higher than its estimated value of 1.13%. In fact, the average level of the risk-free rate over the investment horizon was about 5.7%. We can also compare the observed properties of the stock index return series over the investment horizon with the implied values from Table 1 under the GBM model. The annualized standard deviation of monthly log stock index returns was 14.4%, considerably lower than the value of $\hat{\sigma} = .2445$. The annualized mean of this series was 9.17%, a little higher than the implied value of $\hat{\mu} - \hat{\sigma}^2/2 = .1134 - .2445^2/2 = .0835$.

We repeat this experiment with the parameters estimated over the period from 1926:1 to 1984:12. As indicated in Table 1, the resulting estimates are $\hat{\mu} = .1077$, $\hat{\sigma} = .1998$, and $\bar{r} = .0346$. We use these values to project the results for both the constant proportion strategy and the MV optimal strategy for the 30-year period from 1985:1 through 2014:12. These results are provided under "Estimation Period: 1926:1 to 1984:12" in Table 8. The MV optimal controls are then applied to historical market data over the period from 1985:1 through 2014:12, with results depicted in Figure 6(b). The constant proportion strategy ended up slightly outperforming the MV optimal one. The terminal wealth levels were 938 for the former strategy vs. 865 for the latter. Both exceeded the expected value of 845 (Table 8). Over most of the period, the MV optimal strategy had higher wealth, falling below that of the constant weight strategy in the last half of 2013.[13]

The MV optimal strategy was more heavily invested in the risky asset early on, reflected in the

[11] In other words, a long position in the bond component effectively amounts to buying a 3-month Treasury bill at the beginning of a quarter and holding it until maturity.

[12] This can be seen as being conservative. The investor could choose to invest some or all of the free cash in the risky asset instead. As long as the free cash investment is not levered, the investor would still be assured of reaching the target wealth level at the horizon date under the assumed modelling conditions.

[13] This is not a surprise as under the MV optimal strategy the investor is content with reaching the target. Upside potential above the target is sacrificed in exchange for downside protection below it.

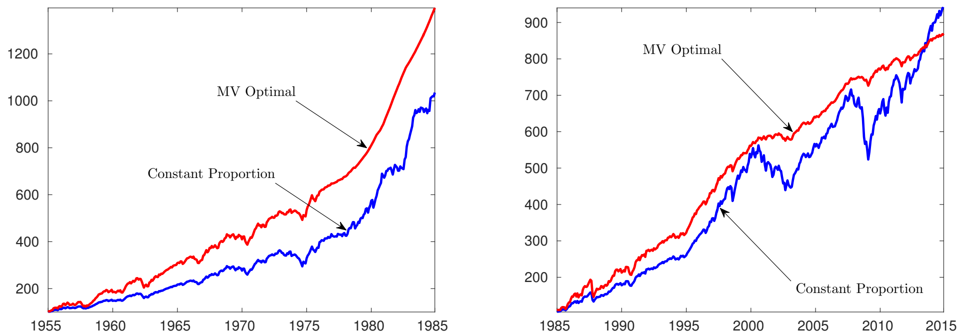

(a) *Out-of-sample experiment, 1955-1984. The MV optimal strategy relies on parameter estimates using data from 1926:1 to 1954:12, as given in Table 1.*

(b) *Out-of-sample experiment, 1985-2014. The MV optimal strategy relies on parameter estimates using data from 1926:1 to 1984:12, as given in Table 1.*

FIGURE 6: *Comparison of MV optimal and constant proportion ($p = .5$) strategies using observed data.*

larger drop (compared to that for the constant proportion strategy) seen in Figure 6(b) following the 1987 market crash. However, the de-risking approach of the MV optimal strategy pays off later on. The MV strategy moved smoothly through the the "tech bubble" in 2002, and the financial crisis in 2008, in marked contrast to the constant proportion strategy.

5.2 Bootstrap resampling

To provide more meaningful tests of the MV optimal strategy using historical data, we turn to a bootstrap resampling approach. This type of procedure has been applied in numerous prior studies to assess the performance of investment strategies; some relatively recent examples include Sanfilippo (2003), Ledoit and Wolf (2008), Annaert et al. (2009), Bertrand and Prigent (2011), and Cogneau and Zakamouline (2013).

It is well-known that an important issue when applying bootstrap resampling to time series data is that the standard bootstrap assumes independent observations, so using it does not preserve temporal dependence. Our stock index return data exhibits some dependence. The autocorrelations are small in magnitude (largest absolute value of about 0.1), but enough of them are statistically significant that the Ljung-Box test clearly rejects the null hypothesis of no autocorrelation for 20 monthly lags. More importantly, the level of the risk-free rate is quite persistent. The autocorrelation function decays slowly, remaining above 0.6 even after five years (60 lags).

To address this issue, we use a moving block approach. A single path is constructed as follows. The investment horizon of T years is divided into k blocks of size M years, so that $T = kM$. We then select k blocks at random (with replacement) from the historical data. Each block starts at a random quarter. We form a single path by concatenating these blocks. Since we sample with replacement, the blocks may overlap. To avoid end effects, the historical data is wrapped around. We then repeat this procedure for many paths.

In this case, parameter estimates are based on the full sample period from 1926 to 2014. From Table 1, $\hat{\mu} = .1108$, $\hat{\sigma} = .1867$, and $\bar{r} = .0354$. With an initial wealth of $W_0 = 100$ and an investment horizon of $T = 30$ years, we calculate the expected level of terminal wealth and its standard deviation for a continuously rebalanced constant proportion strategy with a 50/50 split between the risky stock index and the risk-free asset. We then determine the MV optimal strategy which produces the same expected terminal wealth, subject to our usual base case constraints of annual rebalancing and $q_{\max} = 1.5$, with results shown in Table 8 under "Estimation Period: 1926:1 to 2014:12".

We store the MV optimal controls used to calculate the results in Table 8 but then apply them to historical data, using the resampling approach with 10,000 bootstrap samples. As discussed in Cogneau and Zakamouline (2013), the choice of block size is important. If the block size used is too small, then the serial dependence in the data is not captured adequately. As an example in our context, a small block size can result in many unrealistically large changes in the risk-free rate across concatenated blocks. On the other hand, if the block size is too large then the variance estimates will be unreliable. Excessive homogeneity across the bootstrap samples will lead to a standard deviation that is too low.[14] Various techniques have been suggested to alleviate these effects, such as randomly selecting the block sizes within each sample (Politis and Romano, 1994) and the matched block size method (Sanfilippo, 2003). Cogneau and Zakamouline (2013) conclude that any improvements obtained using these methods could be achieved by selecting the *correct* block size. Unfortunately, this depends on the unknown stochastic process properties of the historical data.

Recognizing that the choice of block size can strongly affect results, we report results for a range of block sizes in Table 9. For all block sizes, the MV optimal strategy outperforms the constant proportion strategy in terms of shortfall probability and standard deviation. Note that the sum of expected terminal wealth and expected free cash is comparable for the MV optimal and constant proportion strategies in all cases. The overall results are highly favorable for the MV optimal strategy. Interestingly, the results for block sizes of 5 and 10 years are fairly close to the theoretical results from Table 8 for both strategies. This again demonstrates the robustness of the MV optimal strategy. Although it is clearly mis-specified in econometric terms since it assumes a constant risk-free rate and GBM, the overall long-term results under these assumptions are quite comparable to those observed in this resampling test for these two block sizes.

A noteworthy feature of Table 9 is the magnitude of the expected free cash for the MV optimal strategy. The values reported here are significantly higher than those reported above in Table 4. Table 9 is based on resampling with replacement, so there can be paths where the risk-free rate is significantly underestimated for a sizeable part of the 30-year horizon. This leads to some large outliers, and a high average value of free cash.[15]

We present detailed plots for the case with a block size of 10 years. Figure 7 depicts the cumulative probabilities from the 10,000 simulations for the constant proportion, MV optimal (no free cash) and MV optimal (plus free cash) cases. The results are quite similar to those seen in Figure 3. The constant proportion strategy outperforms the MV optimal strategy (no free cash) in the extreme left tail of the distribution where terminal wealth is very low, and also in the right tail once the target wealth level for the MV strategy is reached. As usual, the MV optimal strategy (no free cash) essentially gives up the possibility of extremely high wealth in return for achieving outperformance across a wide range of moderate wealth levels. However, if we add in the free cash, then the MV optimal strategy offers considerably more potential upside for high wealth levels. In fact, the cumulative distribution function is fairly close to that of the constant proportion strategy for high values of terminal wealth, while of course being equivalent to the optimal allocation at lower levels (where there is no free cash).[16]

---

[14] In the limit as block size and path size tend to the length of the time series, all samples are simply permutations of the entire time series.

[15] Table 9 reports an average free cash of 126 for a block size of 10 years. In this case, the maximum observed free cash across the 10,000 resamples was 1,761, but the median value was 21.

[16] Moreover, recall that we have simply invested any free cash in the risk-free asset. We could track the constant proportion strategy even more closely by allocating some of the free cash to the risky asset.

Strategy	Expected Value	Standard Deviation	$Prob(W_T < 800)$	Expected Free Cash
Block size: 1 year				
Constant proportion ($p = .5$)	966	600	.48	0
MV optimal	870	213	.15	72
Block size: 5 years				
Constant proportion ($p = .5$)	936	481	.46	0
MV optimal	888	191	.11	111
Block size: 10 years				
Constant proportion ($p = .5$)	923	470	.49	0
MV optimal	911	148	.08	126
Block size: 20 years				
Constant proportion ($p = .5$)	933	494	.48	0
MV optimal	929	104	.05	109
Block size: 30 years				
Constant proportion ($p = .5$)	923	420	.41	0
MV optimal	942	74	.03	122

TABLE 9: *Moving block bootstrap resampling results based on historical data for 1926:1 to 2014:12. The investment horizon is $T = 30$ years, and the initial investment is $W_0 = 100$. The MV optimal strategy has maximum leverage ratio $q_{\max} = 1.5$. Both strategies are rebalanced annually.*

Figure 8 provides several scatter plots of the terminal wealth across the 10,000 resamples for the constant proportion strategy (horizontal axis) vs. the MV optimal strategy, with and without the free cash component (vertical axis). In each case, the solid line marks where the two strategies being compared have the same terminal wealth. The region above and left of this line contains points where the MV optimal strategy has higher terminal wealth than the constant proportion strategy. The opposite holds for points below and right of the solid line. In each plot the vertical dashed line is the median value for the constant proportion strategy, while the horizontal dashed line is the median for the MV optimal strategy. Panel (a) plots all observations for the MV optimal strategy (no free cash) vs. those for the constant proportion strategy. Obviously, the spread of outcomes is far wider for the constant proportion strategy. Since we discard any free cash in excess of the target here, terminal wealth for the MV optimal strategy is capped at the target. Note how close the median line is to this upper bound: the maximum terminal wealth here is 986 but the median is 954. This is to be expected for a target-based approach, as there will be a large number of paths with final wealth just below the target. The median value for the constant proportion strategy is considerably lower at 814. In panel (a), there are 5,950 observations above and left of the solid line, indicating that the MV optimal strategy resulted in higher terminal wealth for almost 60% of the resamples. It is also worth observing that the MV optimal strategy resulted in very low terminal wealth for a number of resampled paths. The smallest values visible in Figure 8(a) for this strategy are noticeably lower than the smallest values for the constant proportion strategy. Since we resample with replacement, it is possible to have paths which repeat periods of very poor index returns. If the strategy happens to also be using leverage during such times, we can end up with very low terminal wealth. However, these are extreme outliers, as indicated by the lower bound

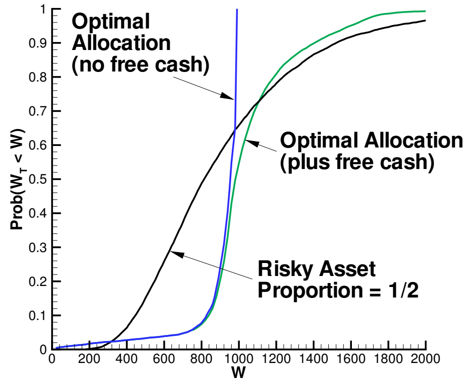

FIGURE 7: *Cumulative probability distributions from 10,000 resamples for constant proportion (p = .5), MV optimal (no free cash), and MV optimal (plus free cash) strategies with block size of 10 years.*

of the vertical range plotted in Figure 8(a), which excludes the worst 2.5% of observations for the MV optimal (no free cash) strategy.

Comparing panels (b) and (a), we can observe the significance of the free cash component. In Figure 8(b), many observations lie above and left of the solid line, indicating that the MV optimal strategy ended up with higher terminal wealth for a large number of the resampled paths. These plots clearly show a positive correlation between the performance of the constant proportion and MV optimal (plus free cash) strategies. Recall that we attribute a high free cash component to having a path with a high risk-free rate. As the constant proportion strategy also benefits from such an environment, this correlation is not surprising.

6 Summary and Conclusions

Compared to the constant proportion strategy, the MV optimal semi-self-financing strategy produces a smaller standard deviation for the same expected terminal value. A common criticism of the use of standard deviation as a risk measure is that it penalizes gains as well as losses, relative to the expected value. With continuous rebalancing and assuming that the value of the risky asset follows a diffusion process without jumps, the total wealth of the portfolio can never exceed the discounted target. As wealth approaches the discounted target, the optimal strategy is to move more wealth into the risk-free asset. This minimizes both the expected quadratic loss relative to the target and the variance. Under discrete rebalancing, cash is withdrawn from the portfolio if the total wealth exceeds the discounted target. This is MV optimal, as well as possibly providing the investor with a free cash bonus during the investment period.

Overall, the MV optimal strategy achieves excellent performance in our numerical simulations. The general intuition for this is as follows. Recall that we are dealing with pre-commitment strategies. Since MV optimality is equivalent to minimizing quadratic loss relative to a target, the investor effectively picks a target terminal wealth at the initial time. The investor pre-commits to being satisfied with this terminal wealth, in the sense that if the market has good returns and the target can be hit by switching to the risk-free asset, then that is the optimal policy. Again, this is optimal in terms of minimizing the probability of being below the target.

The key qualitative aspects of the MV optimal strategy are illustrated by Figures 4(a) and 4(b). The MV optimal strategy captures the advantages of both constant proportion and linear

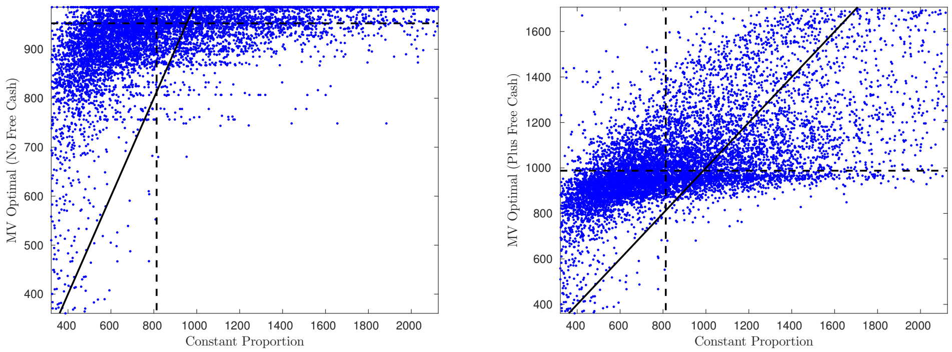

(a) *Constant proportion and MV optimal (no free cash), excluding outliers.*

(b) *Constant proportion and MV optimal (plus free cash), excluding outliers.*

FIGURE 8: *Scatter plots of terminal wealth for MV optimal and constant proportion strategies across 10,000 resamples with a block size of 10 years. The solid line in each plot indicates where the two cases considered have equal terminal wealth. The dashed vertical and horizontal lines in each plot are the median values of terminal wealth for the constant proportion and MV optimal strategies respectively. The plot range is restricted to the central 95% of observations for each strategy.*

glide path strategies. A constant proportion strategy shifts wealth from assets which have risen in value to assets which have declined in value. The MV optimal strategy also *buys low and sells high*, but this strategy also takes into account the accumulated wealth and time-to-go.

From the Appendix, we observe that the target $\gamma/2$ and the mean are related by $\gamma/2 = 1/(2\rho) + E[W_T]$. We can then immediately see that the target will be slightly larger than the expected value if the risk-aversion parameter ρ is large. In this case, we have a strategy which is MV efficient, quadratic loss efficient with respect to the target $\gamma/2$, and approximately mean-semi-variance efficient. This suggests that pre-commitment MV strategies are especially interesting in cases where the expected value is near the target.

To summarize, the MV optimal strategy will be useful under the following conditions:

- The investor commits to a long-term strategy. With typical market parameters, the MV optimal strategy can be expected to outperform a constant proportion strategy by a large margin for investment horizons of at least 10-15 years, though the degree of outperformance will depend on factors such as the maximum allowable leverage ratio.
- The investor pre-commits to a target wealth at the end of the investment horizon. The reduced risk associated with the MV optimal strategy comes at the cost of eliminating some investment paths with gains which would substantially exceed the target.
- The investor should be somewhat risk-averse, so that the target is relatively close to expected terminal wealth.
- The investor must accept that for some low probability cases where the risky investment returns are very poor, the constant proportion strategy will turn out to be superior to the MV optimal strategy.

These conditions could be applicable in scenarios such as individuals saving for retirement, pooled pension plans, or education savings plans.

A possible extension for future work is to incorporate randomness over time in parameters such as volatility or the risk-free rate, but it would probably be better to do this through a regime-switching model which allows for long-term persistence, rather than a mean-reverting diffusion model. However, the relatively simple specification used here which assumes GBM and a constant risk-free rate appears to be quite robust.

Another potential avenue for future research would be to augment the GBM process for the risky asset with jumps. We can expect the MV optimal strategy to display more cautious behavior if jumps are added to the model.

Observe that MV optimality can involve shifting wealth from assets which have risen in value to assets which have declined in value, and it can also entail sitting out bull markets entirely if the target wealth level can be achieved by investing fully in the risk-free asset. As such, it has distinctly contrarian characteristics. Of course, constant proportion strategies also require a high degree of commitment to "lean against the wind", since they involve buying more of risky assets following declines in their prices. Slifka (2013) points out the pitfalls of institutional investors chasing performance and herding, but also notes the practical difficulties in being contrarian.

In a quite different vein, Ko and Huang (2012) report results from an experiment in which subjects exhibited time-inconsistent risk preferences. In particular, participants took on more risk after losses than would be the case under time-consistency. Ko and Huang (2012) contend that their results "highlight the importance of pre-commitment to long-term financial planning" (p. 471). From our perspective the reasoning underlying this statement is suspect. Ko and Huang interpret risk-taking in response to poor initial returns as impulsive gambling behavior and advocate pre-commitment as a way to prevent it. The model described here leads to the same general conclusion regarding the importance of pre-commitment, but for entirely different reasons. MV optimality requires pre-committing to taking on risk in response to poor returns, as well as reducing risk and sacrificing upside potential after sufficiently high returns. As a final comment about pre-commitment, we note that it has been criticized in the literature since it is not time-consistent (e.g. Basak and Chabakauri, 2010). However, time-consistency can be viewed as just another constraint on the set of admissible investment strategies. As such, from the standpoint of long-term investment with a desired target, time-consistent policies cannot outperform pre-commitment strategies.

Appendix: General Intuition for the Embedding Result

Standard dynamic programming arguments cannot be applied to multi-period MV problems, but the work of Li and Ng (2000) and Zhou and Li (2000) shows that an embedding result can be applied to overcome this obstacle. This appendix provides a simple and intuitive geometric description of the embedding result. The notation used mimics that of Section 2, but we emphasize that we are considering a more general environment here. Readers interested in formal proofs are referred to Li and Ng (2000) and Zhou and Li (2000). In addition, a simple yet formal argument can be found in Appendix A of Dang and Forsyth (2016).

Let the wealth of an investor at time t be W_t and consider a general MV context in which this investor desires a high expected level of terminal wealth at some horizon date T. The investor also wants to avoid risk, measured by the variance of terminal wealth. Both the expected level of wealth and its variance will depend on the investment strategy of the investor, denoted by c. At time t, the conditional expected value of terminal wealth given c is $E_c^t[W_T]$ and the conditional variance of terminal wealth is $Var_c^t[W_T]$. For notational simplicity, let $E_c^{t=0}[W_T] = \mathcal{E}$ and $Var_c^{t=0}[W_T] = \mathcal{V}$. The MV achievable set is $\mathcal{Y} = \{(\mathcal{V}, \mathcal{E}) : c \in \mathcal{Z}\}$, where $\mathcal{Z}$ is the set of admissible strategies. $\mathcal{Y}$ represents the possible combinations of expected terminal wealth and its variance given any restrictions on the investment strategy that are captured by $\mathcal{Z}$. Since the investor seeks to both maximize expected wealth and to minimize its variance, we have a multi-objective optimization problem. As usual, we

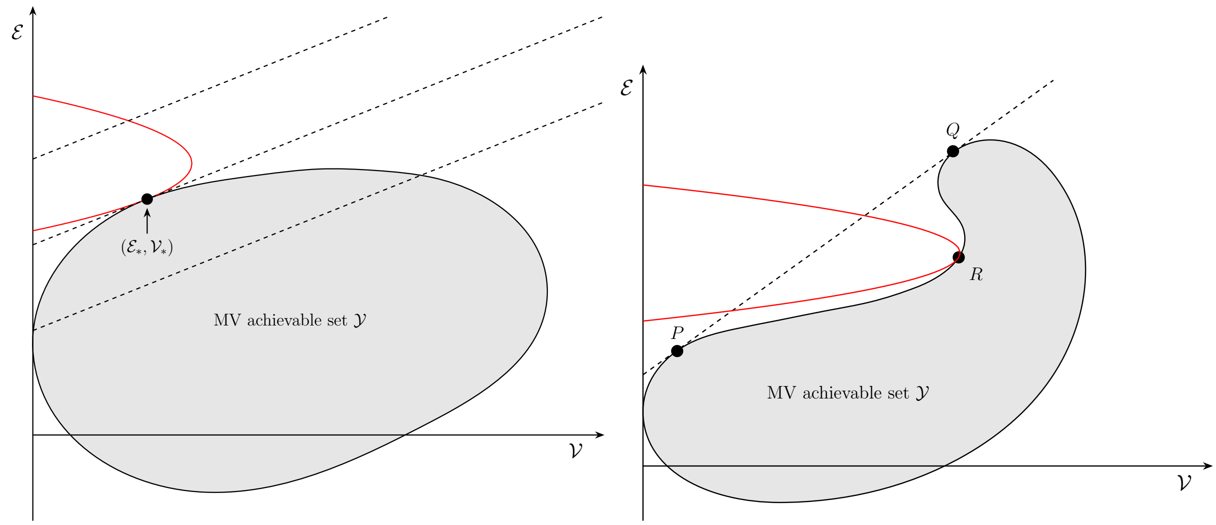

(a) *Problem (P1): dashed line intersects $\mathcal{Y}$. Problem (P2): solid line intersects $\mathcal{Y}$.*

(b) *A case where the solution of problem (P2) is not a solution of problem (P1).*

FIGURE 9: *Solutions of problems (P1) and (P2). The dashed lines are the investor's indifference curves. The shaded region is the MV achievable objective set $\mathcal{Y}$.*

convert this to the problem:

$$\min \rho\mathcal{V} - \mathcal{E}, \tag{P1}$$

where ρ is a risk-aversion parameter. In this setting, "indifference curves" are straight lines on a graph with $\mathcal{V}$ on the horizontal axis and $\mathcal{E}$ on the vertical axis. The slope of these indifference curves is determined by ρ. Given ρ, the optimal achievable combination $(\mathcal{V}_*, \mathcal{E}_*)$ is a tangency point between an indifference curve and the set of attainable combinations $\mathcal{Y}$.

Figure 9(a) illustrates this idea. The MV achievable set $\mathcal{Y}$ is the shaded region. In general, the shape and location of $\mathcal{Y}$ will be determined by parameters such as expected returns, variances of returns, and covariances of returns for investable assets, as well as constraints on strategies reflected in $\mathcal{Z}$. The investor wants a combination of risk and expected wealth that is as far to the upper left as possible, and so arrives at a point where an indifference curve is tangent to $\mathcal{Y}$. The entire set of MV (Pareto) optimal points can be traced out by solving (P1) repeatedly for different values of ρ.[17]

However, in a dynamic multi-period context, we cannot solve problem (P1) using standard dynamic programming techniques which rely on the Bellman principle of optimality (see, e.g. Basak and Chabakauri, 2010).

The basic idea of Li and Ng (2000) and Zhou and Li (2000) is to consider an alternative problem which *can* be solved via dynamic programming. Every solution of the original problem is also a solution of this alternative problem, so the original problem is embedded in the alternative one. This is the basis for referring to this idea as the "embedding result". Reconsidering problem (P1), imagine constructing the straight line in the $(\mathcal{V}, \mathcal{E})$ plane

$$\rho\mathcal{V} - \mathcal{E} = C_1, \tag{23}$$

where C_1 is a constant chosen so that the intersection of (23) and $\mathcal{Y}$ contains at least one point. Then reduce C_1 (i.e. move line (23) to the left) as much as possible, keeping at least one point in

[17] If $\mathcal{Y}$ is not a convex set, then this approach may not produce all of the Pareto optimal points. However, any points generated are certain to be valid Pareto points.

the intersection of $\mathcal{Y}$ and line (23). Any remaining points on line (23) that are in $\mathcal{Y}$ will be Pareto optimal.[18]

The difficulty with trying to minimize $\rho\mathcal{V} - \mathcal{E}$ directly arises because $\mathcal{V} = E_c^{t=0}\left[W_T^2\right] - \left(E_c^{t=0}\left[W_T\right]\right)^2$, and the $\left(E_c^{t=0}\left[W_T\right]\right)^2$ term cannot be handled with standard dynamic programming. However, consider an objective function of the form

$$\mathcal{V} + \mathcal{E}^2 - \gamma\mathcal{E} + constant \tag{24}$$

where γ is a parameter to be specified. The advantage of this expression is that $\mathcal{V}+\mathcal{E}^2 = E_c^{t=0}\left[W_T^2\right]$, which can be dealt with using dynamic programming. The idea of using objective function (24) is similar to the strategy involved in using (23). Consider the parabola (viewed in the $(\mathcal{E}, \mathcal{V})$ plane)

$$\mathcal{V} = -\left(\mathcal{E}^2 - \gamma\mathcal{E}\right) + C_2 = -\left(\mathcal{E} - \gamma/2\right)^2 + \gamma^2/4 + C_2 = -\left(\mathcal{E} - \gamma/2\right)^2 + C_3, \tag{25}$$

where C_2 and C_3 are constants. Informally, this is a *leftward looking* parabola in the $(\mathcal{V}, \mathcal{E})$ plane, as shown in Figure 9(a). Choose C_3 so that the intersection of the parabola (25) and $\mathcal{Y}$ contains at least one point. Then reduce C_3 as much as possible (i.e. shift the parabola to the left) while keeping at least one point in the intersection of the parabola and $\mathcal{Y}$. Suppose $\rho\mathcal{V}_* - \mathcal{E}_* = C_1$ where $(\mathcal{V}_*, \mathcal{E}_*)$ is a Pareto optimal point. We can pick γ so that the parabola (25) intersects $\mathcal{Y}$ at the single point $(\mathcal{V}_*, \mathcal{E}_*)$ by requiring that the parabola be tangent to the line (23) at $(\mathcal{V}_*, \mathcal{E}_*)$. The slope of the parabola (25) is $\partial\mathcal{E}/\partial\mathcal{V} = -1/(2\mathcal{E} - \gamma)$. Since the slope of line (23) is ρ, we have $-1/(2\mathcal{E}_* - \gamma) = \rho$ at $(\mathcal{V}_*, \mathcal{E}_*)$, implying that $\gamma = 1/\rho + 2\mathcal{E}_*$.

To recap, we would like to determine Pareto optimal points by solving problem (P1) directly, but this is not possible using standard dynamic programming. However, any strategy that solves (P1) leading to the point $(\mathcal{V}_*, \mathcal{E}_*)$ can be found by solving

$$\min \mathcal{V} + \mathcal{E}^2 - \gamma\mathcal{E} \tag{P2}$$

where $\gamma = 1/\rho + 2\mathcal{E}_*$. Unlike (P1), problem (P2) can be solved via traditional dynamic programming. Figure 9(a) shows the outcome of this procedure, where the parabola solving (P2) is tangent to the line from problem (P1) at the optimal point.

The entire efficient frontier can be traced out by solving (P2) repeatedly for different values of γ, but this raises another issue. Every strategy which solves (P1) also solves (P2), but the converse is not true: there can be solutions of (P2) which are not solutions of (P1). This can happen if the upper boundary of $\mathcal{Y}$ is not convex. Consider Figure 9(b). If we could solve (P1), we would have an efficient frontier with a gap in it, i.e. no solution to (P1) anywhere on the edge of $\mathcal{Y}$ between P and Q. But we could solve (P2) and reach a point such as R, which is clearly not Pareto optimal. However, a simple procedure can be used to eliminate any non-Pareto optimal points (i.e. points which cannot be found by solving (P1), if we could do so). We construct the upper left convex hull of the points produced by solving (P2) for all values of γ. The Pareto optimal points are the intersection of the points in the upper left convex hull and the points found by solving (P2).[19]

Note that there is no requirement that the original MV formulation reduce to a convex optimization problem.

References

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

[18]As drawn in Figure 9(a), there is a unique tangency point. In general, there could be more than one tangency point, depending on the shape of $\mathcal{Y}$.

[19]See Tse et al. (2014) for details.

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

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

Bielecki, T. R., H. Jin, S. R. Pliska, and X. Y. Zhou (2005). Continuous-time mean-variance portfolio selection with bankruptcy prohibition. *Mathematical Finance 15*, 213–244.

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

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

CRSP (2012). *Data Descriptions Guide: CRSP US Stock & US Index Databases*. Available at `www.crsp.com/files/data_descriptions_guide_0.pdf`.

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: A numerical impulse control approach. *Numerical Methods for Partial Differential Equations 30*, 664–698.

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

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature 40*, 351–401.

Graf, S. (2016). Life-cycle funds: Much ado about nothing? *European Journal of Finance*. forthcoming.

Graham, B. (2003). *The Intelligent Investor*. New York: HarperCollins. Revised edition, forward by J. Zweig.

Investment Company Institute (2015). *2015 Investment Company Fact Book*. Available at `www.icifactbook.org`.

Ko, K. J. and Z. Huang (2012). Time-inconsistent risk preferences in a laboratory experiment. *Review of Quantitative Finance and Accounting 39*, 471–484.

Kocherlakota, N. R. (2001, Summer). Looking for evidence of time-inconsistent preferences in asset market data. *Federal Reserve Bank of Minneapolis Quarterly Review 25*, 13–24.

Ledoit, O. and M. Wolf (2008). Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance 15*, 850–859.

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

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

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous time case. *Review of Economics and Statistics 51*, 247–257.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory 3*, 373–413.

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

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

Shen, Y., X. Zhang, and T. K. Siu (2014). Mean-variance portfolio selection under a constant elasticity of variance model. *Operations Research Letters 42*, 337–342.

Slifka, D. (2013). Improving investment behavior with pre-commitment. *Journal of Investing 22*(1), 83–87.

Tse, S. T., P. A. 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 efficiency of mean-variance based portfolio selection in defined contribution pension schemes. *Quantitative Finance 14*, 237–258.

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