Beating a constant weight benchmark: easier done than said

Peter A. Forsyth[a] Pieter M. van Staden[b] Yuying Li [c]

May 21, 2023

Abstract

We determine a simple dynamic benchmark for asset allocation by solving an optimal stochastic control problem for outperforming the traditional constant proportion benchmark. An objective function based on a time averaged quadratic deviation from an elevated benchmark is proposed. We argue that this objective function combines the best features of tracking error and tracking difference. Assuming parametric models of the stock and bond processes, a closed form solution for the optimal control is obtained. The closed form optimal control is then *clipped* to prevent use of excessive leverage, and to prevent trading if insolvent. Monte Carlo computations using this clipped control are presented which show that for modest levels of outperformance (i.e. 80-170 bps per year), this easily implementable strategy outperforms the traditional constant proportion benchmark with high probability. We advocate this clipped optimal strategy as a suitable benchmark for active asset allocation.

Keywords: optimal control, benchmark outperformance, asset allocation
JEL codes: G11, G22
AMS codes: 91G, 65N06, 65N12, 35Q93

1 Introduction

Many pension plans have a benchmark portfolio which is used to measure the efficiency of the realized investment strategy. These benchmark (or reference) portfolios are invariably based on publicly traded financial assets.

The Canadian Pension Plan (CPP) with CAD 540 billion assets under management has a base reference portfolio of 85% global equity and 15% Canadian government bonds (Canadian Pension Plan, 2022). The *non-base* CPP portfolio has a benchmark of 55% global equity and 45% government bonds.[1] Note that the CPP also allows use of leverage. According to the CPP Annual Report (2022), the CPP has outperformed its benchmark by an annualized 80 bps after fees over the last five years.

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

[b]David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, pieter.vanstaden@gmail.com

[c]David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, yuying@uwaterloo.ca

[1]The base portfolio of the CPP plan is much larger than the non-base (additional) portfolio. Clearly, the base portfolio benchmark is riskier than the non-base portfolio. This is rationalized by noting that the base CPP is only *"partially funded"*, (CPP Annual Report, 2022) This, of course, means that since the plan is not fully funded, greater risk must be taken to have a chance of meeting obligations. One of us (PAF) is currently receiving CPP benefits. PAF finds this comment somewhat disconcerting.

Another example is the Norwegian government pension plan, which has USD 1.35 trillion assets under management. The Norwegian plan has a benchmark of 70% equities and 30% bonds (Government Pension Fund Global, 2022).

Typically, these plans will report results relative to the constant proportion benchmark, in terms of performance and risk measures. These reports are used to justify active investment strategies and/or investment in alternative assets and use of leverage.

Investments in alternative assets are a major strategy in the *Endowment Model* for managing a portfolio. This model was based on the spectacular success of the Yale endowment over many years. However, a *post hoc* analysis of endowments and public pension plans which attempted to emulate the Yale model compared to a 70-30 (equity-bond) reference portfolio showed poorer returns compared to the benchmark (Ennis, 2021), post 2008.

Given the widely adopted industry practice in evaluating performance relative to the constant proportion benchmark strategy, the quest for computing a strategy which outperforms this benchmark becomes immediately relevant and important. Furthermore, one may ask whether it is time to revisit the ubiquitous constant proportion reference strategy.

A better benchmark strategy needs to satisfy at least two criteria: (i) it can be easily constructed (ii) it robustly performs better than the existing constant proportion benchmark. We aim to outperform the traditional benchmark and discover a more useful benchmark by solving a stochastic optimal control problem with suitable objective functions.

There is a large literature on techniques for constructing dynamic strategies for outperforming a benchmark. We refer the reader to (Browne, 2000; Oderda, 2015; Al-Aradi and Jaimungal, 2018; Ni et al., 2022) and the references cited therein.

In the context of measuring the efficiency of index exchange traded funds (ETFs), there are two common metrics: tracking *error* and tracking *difference*. The original use of these metrics was in the context of an index ETF, where the objective is to track the index closely, not outperform it.

In this context, tracking error of a portfolio relative to a benchmark is defined as

$$\text{Tracking Error} \quad = \quad \text{std}\left(\mathcal{R} - \hat{\mathcal{R}}\right) \ , \tag{1.1}$$

where std is the standard deviation, $\mathcal{R}$ is the return of the active portfolio, and $\hat{\mathcal{R}}$ is the return of the benchmark.

In fact, the Norwegian Pension plan (more properly referred to as Government Pension Plan Global) specifies a very tight tracking error of the plan portfolio relative to the 70-30 benchmark of publicly traded assets (Norges Bank, 2021).

The motivation for metric (1.1) is described in Wander (2000). Briefly, this metric might make sense if the investor wants to hire a portfolio manager who will outperform an index, without taking on too much risk. However, the tracking error (also known as the volatility of relative returns) has some odd properties. Suppose that total wealth in the active investment portfolio is denoted by W and the total wealth in the benchmark is $\hat{W}$. Assume that both portfolios follow geometric Brownian motion (GBM),

$$\frac{d\hat{W}}{\hat{W}} = \hat{\mu}\ dt + \hat{\sigma}\ d\hat{Z}\ ;\ \frac{dW}{W} = \mu\ dt + \sigma\ dZ\ ;\ d\hat{Z} \cdot dZ = \rho\ dt\ , \tag{1.2}$$

where $dZ, d\hat{Z}$ are increments of Wiener processes. From equation (1.1) we can see that the instantaneous tracking error per unit time, assuming processes (1.2) is

$$(\ \text{Tracking Error}\ )^2 \quad = \hat{\sigma}^2 + \sigma^2 - 2\rho\hat{\sigma}\sigma\ . \tag{1.3}$$

Note the unusual aspect of equation (1.3): tracking error decreases as correlation increases. This obviously rewards a manager whose active portfolio has a high positive correlation to the benchmark. Suppose $\sigma = \hat{\sigma}$, $\rho = 1$, $\mu \ll \hat{\mu}$. In this case the tracking error is identically zero, even though the managed portfolio severely underperforms the benchmark.

Tracking error (1.1) might be a valid criteria if we desire to track the benchmark as closely as possible, but this metric has been criticized (Johnson et al., 2013; Hougan, 2015; Charteris and McCullough, 2020; Boyde, 2021). In fact, these authors suggest that, in measuring the *post hoc* performance of ETFs, the simple tracking *difference* is a more appropriate metric. Tracking difference is simply the difference between the cumulative returns of the investment portfolio and the benchmark.[2]

Suppose that the amount invested in the benchmark at time t is $\hat{W}(t)$, and the amount in the active portfolio is $W(t)$, with the same amounts invested at time zero, i.e. $W(0) = \hat{W}(0)$. We measure the performance of the active portfolio, relative to the benchmark, over the time horizon $[0,T]$. Then, following the spirit of the tracking difference metric, van Staden et al. (2023) suggests the following control problem for outperformance relative to the benchmark

$$\min_{\mathbb{P}(\cdot)} E\left[\left(W(T) - e^{\beta T}\hat{W}(T)\right)^2\right] , \tag{1.4}$$

where $E[\cdot]$ is the expectation and $\mathbb{P}(\cdot)$ is the dynamic control strategy (i.e. the asset allocation), and $W(T)$ and $\hat{W}(T)$ are the terminal wealth associated with strategy $\mathbb{P}(\cdot)$ and the benchmark respectively.

The intuition behind objective function (1.4) is clear. We desire to outperform the benchmark cumulatively over the period $[0,T]$ by a factor of $e^{\beta T}$ (i.e. continuously compounded at a rate of β per year). We also desire to minimize the volatility relative to the elevated benchmark. The performance metric (1.4) thus directly targets an outperformance of β per year, and attempts to minimize the uncertainty (risk) associated with meeting this target. In a sense, this performance metric combines the desirable features of tracking error and tracking difference.

We can vary the amount of risk we are willing to take, relative to the benchmark, by adjusting β. As $\beta \to 0$, then the optimal solution to problem (1.4) is to simply invest in the benchmark portfolio. We implicitly assume that it is possible to invest in the benchmark directly, or an asset which closely replicates the benchmark. However, as β becomes large, we can expect to have to take on more risk than the benchmark, in order to increase outperformance.

A criticism of the objective function (1.4) is that it is symmetric with respect to the upside and the downside. This is, of course, a common problem with volatility-type performance criteria. In Ni et al. (2022), this objective function was modified to be

$$\min_{\mathbb{P}(\cdot)} E\left[\left(\max(0, e^{\beta T}\hat{W}(T) - W(T))\right)^2 + \max(0, W(T) - e^{\beta T}\hat{W}(T))\right] . \tag{1.5}$$

The objective function (1.5) has a quadratic penalty for underperformance, and a linear penalty for outperformance. However, use of objective function (1.5) does not permit closed form solutions, and requires use of machine learning techniques (Ni et al., 2022) in order to determine the optimal policy $\mathbb{P}(\cdot)$.

Another possible criticism of objective function (1.4) is that deviation from the elevated benchmark is only considered at the terminal time T. However, investment managers are usually required to report performance at regular intervals, perhaps quarterly or monthly. Therefore, deviations from the performance target throughout the investment horizon $[0,T]$ are also of concern.

[2] In practice, this is often reported in an annualized fashion.

To address this concern, the following objective function has been suggested in van Staden et al. (2022)

$$\min_{\mathbb{P}(\cdot)} E\left[\int_0^T \left(W(t) - e^{\beta t}\hat{W}(t)\right)^2 dt\right] , \tag{1.6}$$

which is the time averaged quadratic deviation from the elevated benchmark over the investment horizon. The main focus in van Staden et al. (2022) is on the use of machine learning methods, coupled with a data driven approach, to solve for optimal portfolios using objective function (1.6).

In contrast to van Staden et al. (2022), the objective of this note is to study properties of closed form optimal control solution to problem (1.6) for a simple two asset (stock index and bond) portfolio. We also provide a short, intuitive derivation of the optimal control. We use extensive Monte Carlo simulations to examine the properties of this closed form solution. Since the closed form solution permits infinite leverage, and trading can continue if bankrupt, we apply the clipping technique to the closed form optimal control, as in Vigna (2014), to approximate the solution to the constrained optimal control problem.

If we permit a modest amount of leverage, i.e., borrowing up to 30% of net wealth, then this level of leverage constraint appears to have a modest effect on the solution (compared to the unconstrained control case). This suggests that the clipped optimal control is a good approximation of the true constrained control. This simple closed form approximation can then be used to obtain an intuitive understanding for the control produced by the objective function (1.6) with realistic investment constraints.

One of our main conclusions is that the outperformance objective (1.6), for modest values of β, i.e. $\beta < 200$ bps per year, results in fairly conservative controls, which have a high probability of outperforming the benchmark, without requiring unreasonable amounts of leverage at any time in $[0,T]$.

We further demonstrate that the clipped optimal control of the optimal analytic strategy, using publicly traded stock and bond indexes, offers close to optimal performance. Since the clipped optimal strategy can be easily computed by an asset manager based on historical data, we advocate this strategy as an enhanced benchmark for an active asset manager, replacing the standard constant proportion strategy. We believe that this new dynamic benchmark would allow investors (and the taxpayers paying into public pension plans) to discern true investment skill (or the lack thereof) of the asset managers.

There is more room for success if we apply criteria (1.6) to cases where the investment portfolio has additional assets compared to the benchmark, which would normally be the case. In addition, it might be desirable to avoid postulating a parametric form for the portfolio constituents and generate market scenarios by bootstrapping historical returns (Ni et al., 2022). This requires use of numerical techniques (such as machine learning), which is the main topic of van Staden et al. (2022).

2 Investment Market

We assume that the investor has access to two funds: a broad market stock index fund and a constant maturity bond index fund, and the investment horizon is T. Let S_t and B_t respectively denote the real (inflation adjusted) amounts invested in the stock index and the bond index respectively. In general, these amounts will depend on the investor's strategy over time, as well as changes in the real unit prices of the assets. In the absence of an investor determined control (i.e. cash injections or rebalancing), all changes in S_t and B_t result from changes in asset prices.

We model the stock index as following a jump diffusion process. Let $S_{t^-} = S(t-\epsilon), \epsilon \to 0^+$, i.e. t^- denotes immediately before time t, and let ξ be a random jump multiplier. When a jump occurs, $S_t = \xi S_{t^-}$. We assume that $\log(\xi)$ follows a double exponential distribution (Kou, 2002; Kou and Wang, 2004) with parameters η_1 and η_2, respectively. The probability of an upward jump is $\mathcal{P}_u$, while $1 - \mathcal{P}_u$ is the probability of a downward jump. The density function for $y = \log(\xi)$ is

$$f(y) = \mathcal{P}_u \eta_1 e^{-\eta_1 y} \mathbf{1}_{y \geq 0} + (1 - \mathcal{P}_u) \eta_2 e^{\eta_2 y} \mathbf{1}_{y<0} \ . \tag{2.1}$$

Note that the density of ξ has the form $g(\xi) = f(\log \xi)/\xi$. Define $\kappa = E[\xi - 1]$, and assuming constant units of stock index holding,

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

where μ is the (uncompensated) drift rate, σ is the diffusive volatility, Z_t is a Brownian motion, π_t is a Poisson process with positive intensity parameter λ, and ξ_i are i.i.d. positive random variables having distribution (2.1). Moreover, ξ_i, π_t, and Z_t are assumed to all be mutually independent.

We assume the constant maturity bond index follows

$$dB_t \quad = \quad rB_t\ dt \quad . \tag{2.3}$$

Let p_t be the fraction of total wealth W_t invested in stock index at t. Assuming continuous rebalancing, the total wealth in the investment portfolio follows the process

$$dW_t \quad = \quad p_t W_t \left(\frac{dS_t}{S_t}\right) + (1 - p_t) W_t \left(\frac{dB_t}{B_t}\right) + q\ dt\ , \tag{2.4}$$

where q is continuous constant rate of cash injection into the portfolio. Similarly, let $\hat{W}_t$ be the total wealth invested in the benchmark portfolio, with $\hat{p}$ being the fraction of total benchmark wealth invested in the stock portfolio. We assume in the following that $\hat{p}$ is a constant, which is normally the case for large pension plan benchmarks. This makes our closed form solution final expressions quite simple. However, it is still possible to obtain closed form solutions if $\hat{p} = \hat{p}(t)$, but the final expressions for the optimal control become quite unwieldy.

Then, analogously to equation (2.4), the process followed by the benchmark wealth is

$$d\hat{W}_t \quad = \quad \hat{p}\hat{W}_t \left(\frac{dS_t}{S_t}\right) + (1 - \hat{p}) \hat{W}_t \left(\frac{dB_t}{B_t}\right) + q\ dt\ . \tag{2.5}$$

Remark 2.1 (Stochastic Bond Returns)**.** *Here we have assumed that bond process is non-stochastic, which is arguably a reasonable approximation for short term, low volatility bond indexes. However, it is possible to directly model real returns of a constant maturity bond index fund by a jump diffusion process (Lin et al., 2015; Forsyth, 2022).*

3 Cumulative Tracking Difference

We will now proceed to formally define the investment problem based on the objective function (1.6), which we will refer to as the cumulative tracking difference (CD) in the following.

Begin with equations (2.4 - 2.5). Define the value function $\tilde{V}(w, \hat{w}, t)$ as

$$\tilde{V}(w, \hat{w}, t) \quad = \quad \inf_p \left\{ E_p^{(w,\hat{w},t)} \left[\int_t^T (\hat{W}(s) e^{\beta s} - W(s))^2\ ds \middle| W(t) = w, \hat{W}(t) = \hat{w} \right] \right\}, \tag{3.1}$$

where $E_p^{(w,\hat{w},t)}[\cdot]$ denotes the expectation under the control $p(\cdot)$ as observed at $(w,\hat{w},t)$. For notational simplicity, we subsequently omit the dependence in $p(\cdot)$, when there is no confusion.

For $t \in [0,T-\Delta t]$ the tower property gives

$$
\begin{aligned}
\tilde{V}(w,\hat{w},t) &= \inf_p \Bigg\{ E_p^{(w,\hat{w},t)} \Bigg[\int_t^{t+\Delta t} (\hat{W}(s)e^{\beta s} - W(s))^2 \, ds \\
&\qquad + \int_{t+\Delta t}^{T} (\hat{W}(s)e^{\beta s} - W(s))^2 \, ds \Bigg| W(t)=w, \hat{W}(t)=\hat{w} \Bigg] \Bigg\} \\
&= \inf_p \Bigg\{ E_p^{(w,\hat{w},t)} \Bigg[\int_t^{t+\Delta t} (\hat{W}(s)e^{\beta s} - W(s))^2 \, ds \\
&\qquad + \tilde{V}(W(t+\Delta t), \hat{W}(t+\Delta t), t+\Delta t) \Bigg| W(t)=w, \hat{W}(t)=\hat{w} \Bigg] \Bigg\} . \qquad (3.2)
\end{aligned}
$$

It will be convenient to write the final equation in terms of *backward time* $\tau = T - t$. To this end, we define

$$
\begin{aligned}
V(w,\hat{w},\tau) &= \tilde{V}(w,\hat{w},T-\tau) \\
&\tau = T - t \, . \qquad (3.3)
\end{aligned}
$$

In Appendix A, we take the limit as $\Delta t \to 0$ in equation (3.2), use Ito's Lemma with jumps (Tankov and Cont, 2009), and write the final equations in terms of $V(w,\hat{w},\tau)$ (as in equation (3.3)) to obtain the Hamilton-Jacobi-Bellman (HJB) equation

$$
V_\tau = \inf_p \mathcal{L}_p V \; , \qquad (3.4)
$$

where $\mathcal{L}_p V$ is defined as

$$
\begin{aligned}
\mathcal{L}_p V \equiv\; & \left(w(r + (\mu - r - \lambda\kappa)p) + q\right) V_w + \left(\hat{w}(r + (\mu - r - \lambda\kappa)\hat{p}) + q\right) V_{\hat{w}} \\
& + \frac{p^2 w^2 \sigma^2}{2} V_{ww} + \frac{\hat{p}^2 \hat{w}^2 \sigma^2}{2} V_{\hat{w}\hat{w}} + \left(p\hat{p}w\hat{w}\sigma^2\right) V_{w\hat{w}} \\
& + \lambda \int_0^\infty V(w + pw(\xi - 1), \hat{w} + \hat{p}\hat{w}(\xi - 1), \tau) g(\xi) \, d\xi + (\hat{w}e^{\beta(T-\tau)} - w)^2 - \lambda V \; . \qquad (3.5)
\end{aligned}
$$

Here, $g(\xi)$ is the density of ξ and subscripts in V, e.g., V_τ, denote partial derivatives. Since there are no investment constraints, the domain of PDE (3.4) is $(w,\hat{w},\tau) \in (-\infty,+\infty) \times (-\infty,+\infty) \times [0,T]$.

In addition, note that from equation (3.1) we have

$$
\tilde{V}(w,\hat{w},T) = 0 \; , \qquad (3.6)
$$

hence

$$
V(w,\hat{w},0) = 0 \; . \qquad (3.7)
$$

3.1 Closed form solution

We give a brief overview of the method used to derive the closed form solution here. For a rigorous solution of problem (3.4), we refer the reader to van Staden et al. (2022).

It is convenient to define the following parameters. From equation (2.1), recalling that $y = \log(\xi)$, we have

$$\begin{aligned} \kappa &= E[\xi - 1] = \frac{\mathcal{P}_u \eta_1}{\eta_1 - 1} + \frac{(1 - \mathcal{P}_u)\eta_2}{\eta_2 + 1} - 1. \\ \kappa_2 &= E[(\xi - 1)^2] = \frac{\mathcal{P}_u \eta_1}{\eta_1 - 2} + \frac{(1 - \mathcal{P}_u)\eta_2}{\eta_2 + 2} - 2\kappa - 1 \ . \end{aligned} \tag{3.8}$$

Assume

$$\begin{aligned} V &= A(\tau)w^2 + B(\tau)w + C(\tau) + \hat{A}(\tau)\hat{w}^2 + \hat{B}(\tau)\hat{w} + D(\tau)w\hat{w} \\ V_\tau &= A_\tau w^2 + B_\tau w + C_\tau + \hat{A}_\tau \hat{w}^2 + \hat{B}_\tau \hat{w} + D_\tau w\hat{w} \end{aligned} \tag{3.9}$$

Recall that the subscript, e.g., V_τ, denotes partial derivative. Substitute equation (3.9) into equation (3.4). This is a quadratic function of the control p. It is easily verified (after the fact) that, for the objective function of the optimization problem in (3.4), the coefficient of p^2 is positive. Applying the first order condition determines the optimal control. This yields a system of ODEs for the unknown $A, B, \ldots$, with initial conditions for the ODEs determined by matching equation (3.9) with equation (3.7).

It turns out that the control $p(w, \hat{w}, \tau)$, $\tau = T - t$, depends only on A,B,D, which are given by

$$\begin{aligned} A_\tau &= (2r - \eta)A + 1 \ \ ; \ \ A(0) = 0 \\ D_\tau &= (2r - \eta)D - 2e^{\beta t} \ \ ; \ \ D(0) = 0 \\ B_\tau &= (r - \eta)B + 2qA + qD \ \ ; \ \ B(0) = 0 \\ &\eta = \frac{(\mu - r)^2}{\sigma_e^2} \ ; \ \sigma_e^2 = \sigma^2 + \lambda\kappa_2 \ ; \ \kappa_2 = E[(\xi - 1)^2] \end{aligned} \tag{3.10}$$

The solutions are

$$A = \frac{e^{(2r-\eta)\tau} - 1}{(2r - \eta)} \quad ; \quad D = 2e^{\beta T}\left(\frac{e^{-\beta\tau} - e^{(2r-\eta)\tau}}{2r - \eta + \beta}\right) \tag{3.11}$$

and

$$\begin{aligned} B &= \frac{2q}{2r - \eta}\left(\frac{e^{(2r-\eta)\tau} - e^{(r-\eta)\tau}}{r} - \frac{\left(e^{(r-\eta)\tau} - 1\right)}{r - \eta}\right) \\ &\quad + \frac{2qe^{\beta T}}{2r - \eta + \beta}\left(\frac{e^{(r-\eta)\tau} - e^{-\beta\tau}}{r - \eta + \beta} - \frac{\left(e^{(2r-\eta)\tau} - e^{(r-\eta)\tau}\right)}{r}\right) \ . \end{aligned} \tag{3.12}$$

The optimal control p^* (from the first order condition) is given by

$$\begin{aligned} p^* &= \frac{(\mu - r)}{w\sigma_e^2}\left(h(\tau) + (\hat{w}f(\tau) - w)\right) + \hat{p}\frac{\hat{w}}{w}f(\tau) \\ &h(\tau) = \frac{-B}{2A} \ ; \ f(\tau) = \frac{-D}{2A} \end{aligned} \tag{3.13}$$

Some algebra (van Staden et al., 2022) shows that

$$e^{\beta T} \geq f(\tau) \geq e^{\beta(T-\tau)} \quad ; \quad h(\tau) \geq 0 \ . \tag{3.14}$$

Remark 3.1 (Trading continues if bankrupt). *Since the closed form solution allows unbounded leverage, then there is nothing to prevent $W_t < 0$. This is similar to the closed form solution for multi-period mean variance optimization (Zhou and Li, 2000; Wang and Forsyth, 2010). In this case, equation (3.13) implies that $p^* < 0$, so that the amount in the stock index is $p^* w > 0$. In other words, the investor can continue to borrow and trade stocks, even if insolvent, which is unrealistic.*

3.1.1 Intuition from control (3.13)

Consider the simple case where there is no cash injection, i.e. $q = 0$, which implies that $B(\tau) \equiv 0$. For ease of exposition, make the assumptions that

$$
\begin{aligned}
(\mu - r) &\geq 0 \ , \\
\hat{p}\hat{w} &\geq 0 \ , \\
w &> 0 \ , \text{ but see Remark 3.1 } \ .
\end{aligned}
\tag{3.15}
$$

In this case, it then follows from equation (3.13) that

$$
p \quad \begin{cases} = \hat{p} & \text{if } w = \hat{w} f(\tau) \\ > \hat{p} & \text{if } w < \hat{w} f(\tau) \\ < \hat{p} & \text{if } w > \hat{w} f(\tau) \end{cases} \ . \tag{3.16}
$$

The strategy is fundamentally contrarian. If the active portfolio performs poorly relative to the benchmark, then the stock index weight is increased. On the other had, if we are fortunate, and the active portfolio does well relative to the benchmark, then the stock index weight is decreased.

Remark 3.2 (Robustness of control to misspecification). *In van Staden et al. (2021), it has been noted that optimal multi-period mean-variance strategies are robust to model misspecification errors, in contrast to the single period mean-variance case. This robustness can be traced to the nature of a contrarian control. We conjecture that, similarly to the multi-period mean-variance case, the optimal control (3.13) is also robust to model parameter misspecification. Some numerical tests verifying this conjecture are given in Section 4.1.5.*

3.2 Clipped control: Handling bankruptcy and bounded leverage

The closed form solution (3.13) is for the unconstrained optimal control problem, e.g. infinite leverage is allowed and trading can continue if bankrupt. This can produce unrealistically optimistic results for some closed form solutions of optimal control in financial applications, e.g., a closed form solution for multi-period Mean-Variance (MV) optimal strategies (Zhou and Li, 2000). In the MV case, if the MV control problem under a no-bankruptcy constraint (i.e. trading stops if bankrupt) is solved, then this constraint has a large effect on the solution, compared to the unconstrained solution (Wang and Forsyth, 2010).

In order to prevent unbounded leverage, we can require that the fraction of wealth invested in stocks satisfy the constraint

$$
p \in [0, p_{\max}] \ , \tag{3.17}
$$

where $p_{\max}$ is a bounded constant. This would then change equation (3.4) to

$$
\begin{aligned}
V_\tau &= \inf_{p \in \mathcal{Z}} \mathcal{L}_p V \\
&\qquad \mathcal{Z} = [0, p_{\max}] \ .
\end{aligned}
\tag{3.18}
$$

In general, it is not possible to obtain a closed form solution to equation (3.18). Smooth solutions to the HJB equation (3.18) may not exist. It is non-trivial to devise numerical techniques which ensure convergence to the viscosity solution (Wang and Forsyth, 2010; Ma and Forsyth, 2017) of equation (3.18).

Since we are also interested in discovering a better performing strategy satisfying the constraint (3.18), which can easily be computed by any asset manager, we consider a *clipping* procedure to the unconstrained optimal control (3.18). The optimal control and the clipped optimal control p_c^* are explicitly given below

$$\begin{aligned} p^* &= \frac{(\mu - r)}{\sigma_e^2 w}\Big(h(\tau) + (\hat{w} f(\tau) - w)\Big) + \frac{\hat{p}\hat{w} f(\tau)}{w} \\ p_c^* &= \min(\max(0, p^*), p_{\max}) \ , \end{aligned} \tag{3.19}$$

A similar idea was exploited in Vigna (2014) in the context of closed form solutions for multi-period mean-variance asset allocation.

Remark 3.3 (No trading in stocks if bankrupt)**.** *If there are no jumps (i.e. $\lambda = 0$ in equation (2.4)), then imposing a bounded leverage constraint ensures that $W_t \geq 0$, see Wang and Forsyth (2010). However, if $p_{\max} > 1$, this is no longer true if we permit jumps in the stock index. Note that equation (3.19) also imposes the condition $p_c^* \geq 0$. If $W_t \leq 0$, then this forces $p_c^* \equiv 0$, which means that the stock is liquidated, and debt accumulates with interest r until $t = T$.*

Note that if $w \gg \hat{w} f(t)$ in equation (3.19), then, it is possible that $p^* < 0$. In other words, the unconstrained control in this case shorts stocks. This problem can be attributed to the symmetric risk measure in equation (3.1), since extreme outperformance, i.e., $w > e^{\beta t}\hat{w}$, is also penalized. The clipped control prevents this sort of undesirable behaviour.

Some pension plans are required to undertake a policy of no-leverage, i.e. $p_{\max} = 1$, while other plans allow limited leverage. In our model set-up, we have only two assets: a stock index and a bond index for both the benchmark and the optimal portfolio. Usually the benchmark is a stock index and a bond index. However, many pension plans are using alternative assets, such as private equity and private credit. Although controversial, some authors have suggested that returns on private equity can be replicated using a leveraged small cap stock index (see Phalippou (2014); L'Her et al. (2016)). To this end, we set $p_{\max} = 1.3$ to approximate (very roughly) a portfolio with some exposure to alternative assets.

To summarize, we have clipped the unconstrained control to ensure that we have a feasible solution to the constrained problem (3.18). It is unlikely that a closed form solution exists for equation (3.18). The clipped control is almost certainly sub-optimal. In the following, we will carry out numerical simulations using both the unconstrained control, and the feasible, sub-optimal clipped control. In terms of the objective function (1.4), the unconstrained control solution will provide a lower bound for the true constrained control objective function. We can give a bound for the error in using the clipped control (3.19) by examining the difference between the clipped control objective function value and the unconstrained control objective function value.

However, it is of more practical interest to examine the performance, in terms of the usual investment metrics, of the clipped control strategy compared to the benchmark. We will see that this approximate control does surprisingly well.

4 Numerical Results

We use data from the Center for Research in Security Prices (CRSP) on a monthly basis over the 1926:1-2021:12 period. Our base case tests use the CRSP 30 day T-bill for the bond asset and the CRSP value-weighted total return index for the stock index. This latter index includes all distributions for all domestic stocks trading on major U.S. exchanges. All of these various indexes are in nominal terms, so we adjust them for inflation by using the U.S. CPI index, also supplied by CRSP.

We use the threshold technique (Cont and Mancini, 2011; Dang and Forsyth, 2016) to estimate the parameters for the parametric stochastic process models. Table 4.1 shows the results of calibrating the models to the historical data.

μ	σ	λ	$\mathcal{P}_u$	η_1	η_2	T-bill return r
0.0897	0.1464	0.3229	0.2258	4.3638	5.5316	0.0035

TABLE 4.1: *Estimated annualized parameters for double exponential jump diffusion model. Value-weighted CRSP index, 30 day US T-bill index deflated by the CPI. Sample period 1926:1 to 2021:12. The mean return of the 30-day T-bill index is* $r = 0.0035$.

4.1 Investment scenario

Table 4.2 shows our base case investment scenario. We consider $T = 10$ years, with an initial investment of 100. Cash injection occurs continuously at a rate of 10 per year. The target benchmark is $\hat{p} = 0.70$ in the stock index and 0.30 in bonds. Recall that this is the benchmark used by the Norwegian fund (Government Pension Fund Global, 2022). This is also the benchmark used in a study of the underperformance of endowments (Ennis, 2021).

Investment horizon T (years)	10.0
Equity market index	CRSP Cap-weighted index (real)
Bond index	30-day T-bill (US) (real)
Initial portfolio value W_0	100
Cash Injection per year q	10
Rebalancing times	Continuous
Outperformance target (per year) β	$\{.01, .02\}$
Benchmark fraction in stock index $\hat{p}$	.70
Market parameters	See Table 4.1

TABLE 4.2: *Input data for examples.*

4.1.1 Bounded leverage (clipped optimal control)

We carry out Monte Carlo simulations assuming the processes (2.2-2.3). We use 1000 timesteps and 6.4×10^5 simulations. We consider $\beta = \{.01, .02\}$. A desirable strategy should achieve high probability of $(W_T/\hat{W}_T) > 1$.

Rather than report the objective function value (1.6), We define a normalized dimensionless objective function as

$$\text{Normalized Objective Function} \;\; = \;\; \frac{1}{W_0}\sqrt{\left(\frac{1}{T}E\left[\int_0^T \left(W(t) - e^{\beta t}\hat{W}(t)\right)^2 dt\right]\right)} \qquad (4.1)$$

Table 4.3 shows the normalized objective function, comparing the results for the clipped and unconstrained controls. We remind the reader that the clipped control is only an approximation to the true control for equation (3.18). However, an upper bound for the error incurred by using the clipped control can be determined from the difference between the unconstrained objective function, and the objective function obtained using the clipped control. Table 4.3 shows that the worst case error from the approximate control (in terms of the normalized objective function) is of the order of one percent for $\beta = .01$ and five per cent for $\beta = .02$. We emphasize that this is very likely a gross overestimate of the error incurred using the clipped control to solve problem (3.18).

	$\beta = .01$	$\beta = .02$
unconstrained	0.07441	0.1556
Clipped $p \in [0,1.3]$	0.07540	0.1629

TABLE 4.3: *Normalized objective function (4.1). Scenario in Table 4.2. Clipped control refers to equation (3.19), unconstrained control equation(3.13). The target outperformance β is as shown.*

However, perhaps a more meaningful comparison is in terms of the usual investment statistics, which we show in Table 4.4. We can see from Table 4.4 that the statistics are very similar for the unconstrained control and the clipped control, for $\beta = .01$. As we might expect, the differences are somewhat larger for the more aggressive case of $\beta = .02$. In this case ($\beta = .02$), the largest difference occurs for the expected shortfall at the 5% level, which is about six per cent. Expected shortfall in this case is the mean of the worst 5% of the outcomes for W_T, which we denote by ES(5%).

However, we emphasize that the unconstrained control case is not an implementable trading strategy in practice.

Mean W_T	Median W_T	5^{th} percentile	95^{th} percentile	ES(5%)	Median IRR
Unconstrained optimal control $\beta = .01$					
352.38	325.45	165.65	623.29	131.99	0.062
Clipped optimal control: $p \in [0,1.3]$; $\beta = .01$					
352.17	325.43	164.43	623.26	129.27	0.062
Unconstrained optimal control $\beta = .02$					
377.47	349.07	162.04	681.18	117.15	0.071
Clipped optimal control $p \in [0,1.3]$; $\beta = .02$					
375.61	348.70	147.08	681.12	110.33	0.071

TABLE 4.4: *Statistics of W_T for the clipped optimal strategy, and the constant proportion benchmark. Scenario in Table 4.2. Clipped control (3.19) used for the stock index weight. ES(5%) is the mean of the worst 5% of the outcomes for W_T. IRR is the internal rate of return. We use* 1000 *timesteps and* 6.4×10^5 *simulations. The target outperformance β is as shown. Scenario in Table 4.2. Statistics for the benchmark portfolio given in Table 5.1*

Figure 4.1 shows the cumulative distribution function (CDF) of the ratio $(W_T/\hat{W}_T)$, for both the unconstrained control (3.13) and the clipped control (3.19), for $\beta = .01, .02$. A desirable outcome is that $(W_T/\hat{W}_T) > 1$ (the active portfolio has outperformed the benchmark).

Of course, the solution to the constrained control problem, analogous to (3.4) but with $p \in [0, p_{\max}]$, will differ from the clipped control solution. The clipped optimal control p_c^* in (3.19) only approximates the solution to the constrained optimal control problem. Consequently we would expect the true constrained solution CDF of $W/\hat{W}$ to differ from the clipped control solution. However, Figure 4.1(a) shows that, for $\beta = .01$, the CDFs from the clipped control and the unconstrained control overlap. This indicates that the clipped control (3.19) is almost exact in this case, since the constraints do not appear to be binding for this value of β. For the case of $\beta = .02$, Figure 4.1(b) shows that the CDFs for the clipped control approximation and the exact unconstrained control overlap, except for a small difference near $(W/\hat{W}) = 1$.

The implication is that the clipped optimal control is a reasonable approximation to the exact optimal control under constraints (3.18) at least for moderate levels of the outperformance target $\beta \leq 200$ bps per year.

From now on, we will show results using only the clipped approximate control p_c^* (3.19). We will refer to this as the clipped *optimal* control to distinguish this strategy from the benchmark. It will be understood that p_c^* is in fact only an approximation to the optimal control under constraints (3.18).

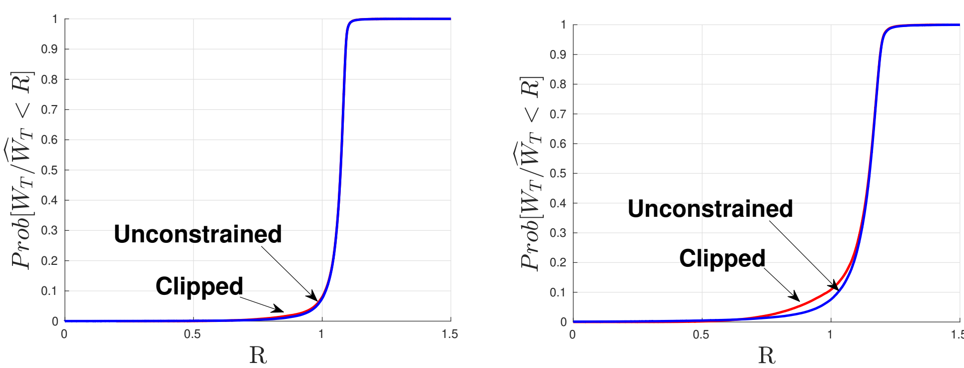

(a) $\beta = .01$ (b) $\beta = .02$

FIGURE 4.1: *CDF of the ratio* $R_T = W_T/\hat{W}_T$*, scenario in Table 4.2. Clipped optimal control* $\tilde{p}^*$ *with* $p_{\max} = 1.3$ *in (3.19). Unconstrained optimal control,* p^* *from equation (3.13). Outperformance is indicated if* $R > 1$*.* 1000 *timesteps and* 6.4×10^5 *Monte Carlo simulations.*

4.1.2 Wealth ratio

Figure 4.2 shows the time evolution of the wealth ratio $(W_t/\hat{W}_t)$, assuming the clipped control (3.19). Recall that outperformance at t is indicated when $(W_t/\hat{W}_t) > 1$. Observe that for $\beta = .01$, there is an 80% probability that the clipped control strategy generates wealth greater than 0.99 of the benchmark wealth, at all times during the ten year investment horizon. There is an 80% probability of outperforming the benchmark at all times greater than about 2.5 years, for both values of β. In addition, from Figure 4.1, we observe that the clipped control solution has a 90% probability of outperforming the benchmark at $t = T$, for both values of β. For $\beta = .01$, there is clearly a smaller spread of the wealth ratio around the median value (over time) compared with $\beta = .02$, in Figure 4.2. This corresponds to our intuition: as the outperformance target β is increased, it is necessary to take on more risk.

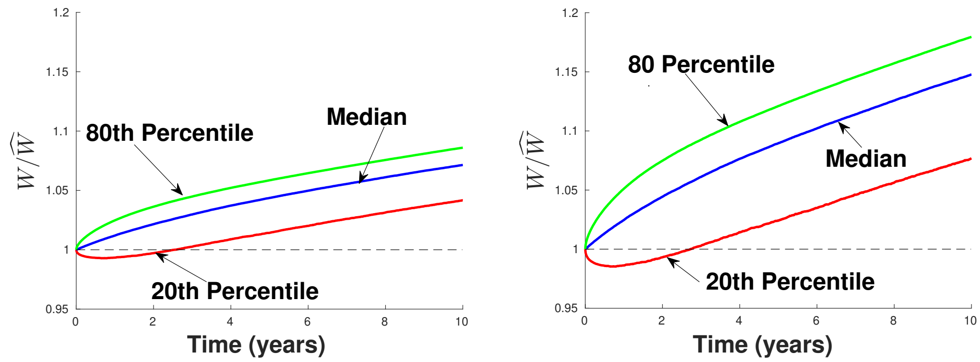

(a) $\beta = .01$ (b) $\beta = .02$

FIGURE 4.2: *Time evolution of the wealth ratio* $W_t/\hat{W}_t$*, clipped optimal strategy (3.19). Scenario in Table 4.2. Outperformance is indicated if* $R_t = (W_t/\hat{W}_t) > 1$*. 1000 timesteps and* 6.4×10^5 *Monte Carlo simulations.*

4.1.3 Fraction in stocks

Figure 4.3 shows the percentiles of the time evolution of the fraction in the stock index. In this case, there is a striking difference between Figure 4.3(a) ($\beta = .01$) and Figure 4.3(b) ($\beta = .02$). For $\beta = .01$, the median equity fraction starts off at about 0.83 and decreases as time goes on. The upper and lower percentiles are tightly clustered about the median. The 80th percentile fraction in equities never exceeds .90 (recall that the benchmark equity fraction is .70). In contrast, the $\beta = .02$ case shows a much wider variation about the median. At the 80th percentile level, the clipped optimal control in this case shows a modest amount of leverage ($p \leq 1.05$).

The reader should note that for any given stochastic path, the control does not stay at the percentile bounds, but responds to actual investment experience. For example, in the $\beta = .01$ case, Figure 4.3(a) can be interpreted as indicating that the fraction in equities never exceeds 0.88 at the 80th percentile and is never less than 0.75 at the 20th percentile, over the ten year horizon.

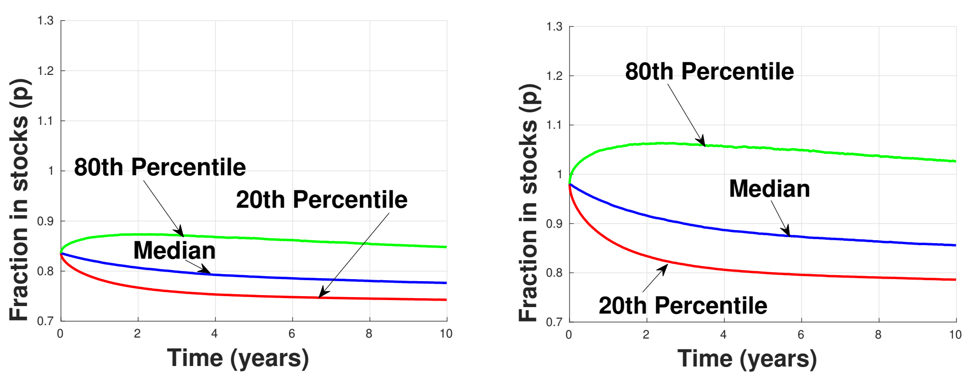

(a) $\beta = .01$ (b) $\beta = .02$

FIGURE 4.3: *Time evolution of the equity fraction, clipped optimal strategy equation (3.19). Scenario in Table 4.2. 1000 timesteps and* 6.4×10^5 *Monte Carlo simulations.*

4.1.4 Internal rate of return

Another way of examining the results is to compute the annualized pathwise internal rate of return (IRR), for both the clipped optimal strategy (3.19) and the benchmark, over the entire 10 year period.

Denote the IRR of the clipped optimal strategy by IRR_{co} and the IRR of the benchmark by IRR_{bench}. The pathwise difference IRR_{diff} is then determined by

$$IRR_{diff} \quad = \quad IRR_{co} - IRR_{bench} \ . \tag{4.2}$$

Figure 4.4 shows the CDF of IRR_{diff}. For the aggressive target outperformance $\beta = .02$ (Figure 4.4(b)), observe that there is an 80% chance that the IRR of the clipped optimal strategy beats the benchmark by more than 100 bps per annum. The median outperformance is about 170 bps per annum. As expected, the less aggressive case of $\beta = .01$ (Figure 4.4(a)), has about a 92% probability of beating the benchmark at ten years, with a median pathwise outperformance of about 85 bps per year.

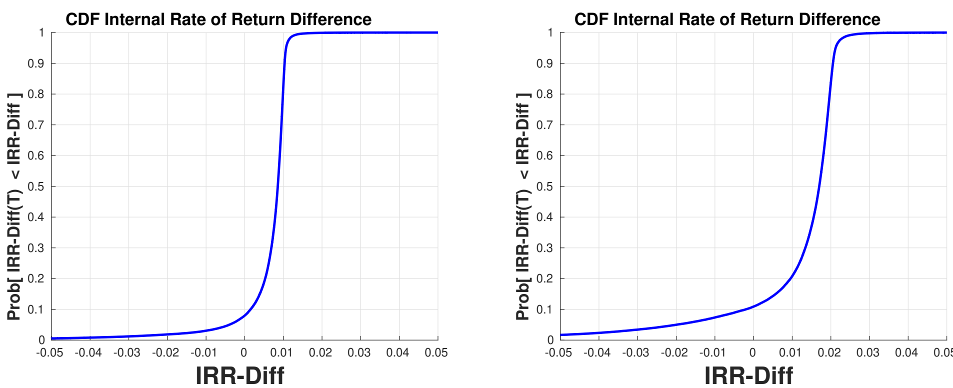

(a) $\beta = .01$ (b) $\beta = .02$

FIGURE 4.4: *CDF of the pathwise difference in terminal IRR (clipped optimal strategy compared to the benchmark), over* $[0.T]$*, see equation (4.2). Scenario in Table 4.2. Clipped control,* $p_{\max} = 1.3$ *in equation (3.19). Unconstrained control, equation (3.13). 1000 timesteps and* 6.4×10^5 *Monte Carlo simulations. Outperformance indicated by* $IRR_{diff} > 0$.

4.1.5 Parameter Misspecification

As an additional check on the robustness of this strategy, we will simulate a case where

- The active strategy is based on an assumed set of parameters
- The actual risky asset follows a different set of parameters

Numerical experiments reveal that, as might be expected, the most sensitive parameter is the stock drift μ in equation (2.2). We will use the base case parameters in Table 4.1, with the scenario in Table 4.2.

We will focus attention on the conservative outperformance target of $\beta = .01$ in equation (1.6). We will compute the optimal strategy (3.19) using the parameters in Table 4.1. In our simulations, we will reduce the simulated stock drift by 200 bps and 400 bps (annually). To be more precise, we replace μ in equation (2.2) by $\mu_a = \mu - .02$ and $\mu_a = \mu - .04$. Note that when we reduce the drift, we reduce the drift for both the controlled strategy and the benchmark. Table 4.5 shows that, even in this case where the parameters are misspecified, the active strategy continues to have a median outperformance of $70 - 80$ bps per year.

Mean W_T	Median W_T	5^{th} percentile	95^{th} percentile	ES(5%)	Median IRR
Reduce drift μ by 200 bps					
Benchmark $\hat{p} = .70$					
294.22	271.79	152.18	506.26	131.14	.040
Clipped optimal control					
311.87	289.21	140.11	552.66	107.31	.048
Reduce drift μ by 400 bps					
Benchmark $\hat{p} = .70$					
263.36	243.68	137.87	450.17	119.11	.026
Clipped optimal control					
275.87	256.84	115.79	490.98	88.31	.033

TABLE 4.5: *Statistics of W_T for the clipped optimal strategy, and the constant proportion benchmark. Scenario in Table 4.2. Clipped control (3.19) used for the stock index weight. ES(5%) is the mean of the worst 5% of the outcomes for W_T. We use 1000 timesteps and 6.4×10^5 simulations. The target outperformance $\beta = .01$. Control computed using the data in Table 4.1. Actual simulations used process (2.2) for the stock, except that the actual drift μ_a is $\mu_a = \mu - .02$ and $\mu_a = \mu - .04$.*

Figure 4.5 shows the CDFs of the wealth ratio $W_T/\hat{W}_T$. We compare the cases with and without stock drift reductions. The results are very close for a reduction of 200 bps. For the 400 bps case, the probability of underperformance has increased from ten per cent to twenty per cent. This is acceptable under this extreme stress test.

Figure 4.6 shows the time evolution of $W_t/\hat{W}_t$ for the cases with a drift reduction of 200 and 400 bps per year.

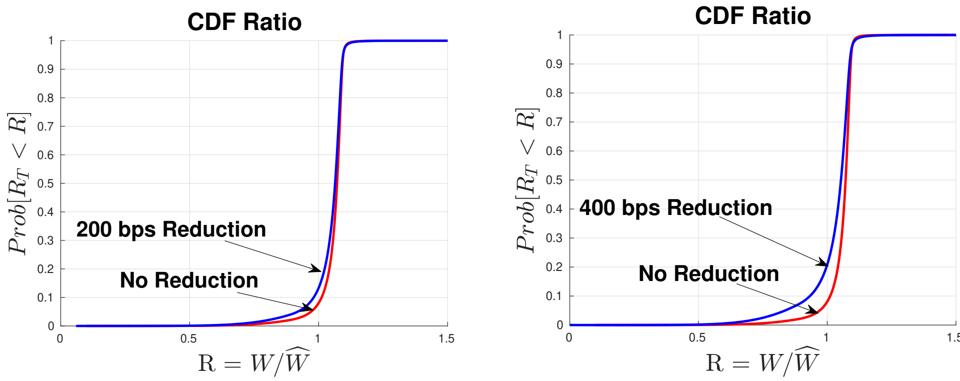

(a) $\mu_a = \mu - .02$ (b) $\mu_a = \mu - .04$

FIGURE 4.5: *CDF of the ratio $R_T = W_T/\hat{W}_T$ for the scenario in Table 4.2. No reduction: control computed using data in Table 4.1, stock price follows equation (2.2) for both controlled portfolio and benchmark. Reduction: control computed using data in Table 4.1, stock price follows equation (2.2) for both controlled portfolio and benchmark except that the stock drift is reduced by the amount shown. Outperformance is indicated if $R > 1$.*

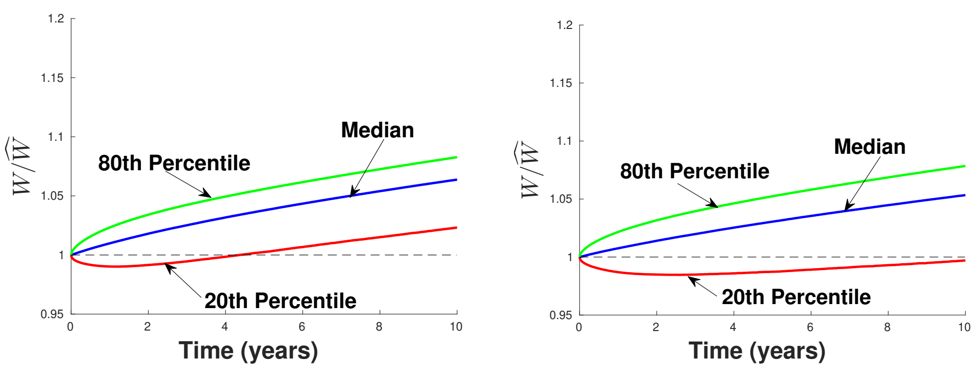

(a) $\mu_a = \mu - .02$ (b) $\mu_a = \mu - .04$

FIGURE 4.6: *Time evolution of the wealth ratio $W_t/\hat{W}_t$, strategy (3.19). Strategy computed using the data in Table 4.1. Simulated stock market follows (2.2), except that the stock drift μ_a is reduced by the amount shown. Scenario in Table 4.2. Outperformance is indicated if $W_t/\hat{W}_t > 1$. 1000 timesteps and 6.4×10^5 Monte Carlo simulations.*

5 Summary statistics

Table 5.1 shows summary statistics for the clipped optimal control and the constant proportion benchmark. We can see directly from this table that the median IRR for the clipped control for the aggressive case of $\beta = .02$ is about 170 bps higher than the benchmark. However, there is no free lunch here, the $5th$ percentile for the clipped control is 147 compared to the $5th$ percentile for the benchmark of 169. In Table 5.1, we include the expected shortfall at the 5% level, which is simply the mean of the worst five per cent of the terminal wealth values W_T. We denote this tail measure by $ES(5\%)$. For the $\beta = .02$ case, the $ES(5\%)$ for the benchmark is 145 compared to 110 for the clipped optimal control.

On the other hand, for the (relatively) conservative case of $\beta = .01$, (Table 5.1) the median for the clipped optimal strategy outperforms the benchmark by 80 bps per year, and has about the same result at the 5th percentile. The ES(5%) is, in this case, only slightly worse than the benchmark. In this case, the results using the clipped optimal strategy are quite impressive. If we target an outperformance of 100 bps per year, then the actual median outperformance is about 80 bps per year, with very little increase in the downside tail risk.[3] This is almost a free lunch.

We remind the reader that the total real amount invested over 10 years is 200, hence these tail outcomes (at the 5th percentile) are very poor, for both the constant proportion benchmark and clipped optimal strategy. While pension plan holders would be very disappointed in these results for either strategy (at the 5th percentile), in general, the clipped control strategy is preferable to the constant proportion strategy and can therefore serve as an enhanced benchmark for active asset managers.

Mean W_T	Median W_T	5^{th} percentile	95^{th} percentile	ES(5%)	Median IRR
Benchmark $\hat{p} = .70$					
329.38	303.66	168.6	570.35	144.97	0.054
Clipped optimal control $\beta = .01$					
352.17	325.43	164.43	623.26	129.27	0.062
Clipped optimal control $\beta = .02$					
375.61	348.70	147.08	681.12	110.33	0.071

TABLE 5.1: *Statistics of W_T for the clipped optimal strategy, and the constant proportion benchmark. Scenario in Table 4.2. Clipped control (3.19) used for the stock index weight. ES(5%) is the mean of the worst 5% of the outcomes for W_T. We use* 1000 *timesteps and* 6.4×10^5 *simulations. The target outperformance β is as shown. Scenario in Table 4.2. ES(5%) is the mean of the worst 5% of the outcomes. IRR is internal rate of return.*

6 Conclusions

In this paper, we have shown that the clipped form of the closed form control for the cumulative difference objective function can achieve a high probability (90%) of outperforming a benchmark, with a median outperformance of 80-170 bps per year. The clipped form of the control has the

[3] The annualized outperformance of the Canadian Pension Plan (CPP) relative to the benchmark (2017-2022), net of costs, is 80 bps. See page 46 in CPP Annual Report (2022).

desirable property that (i) leverage is bounded and (ii) no trading if bankrupt. Technically, the clipped control is suboptimal, but our Monte Carlo simulations indicate that the degree of sub-optimality is small. This property can be traced to the inherent conservative policy of the cumulative difference objective function.

Based on the assumption that the market dynamics are driven by equation (2.2), with known parameters, our simulations show that a dynamic trading strategy can beat a fixed weight benchmark by 80-170 bps per year with little risk. This is, of course, not surprising, since the admissible control set for a dynamic trading strategy is clearly larger than the singleton fixed weight control. Even in the case of misspecified parameters, the dynamic strategy still holds up well.

The optimal control solution reminds us of a very important fact. Any attempt to outperform a benchmark has some risk of underperforming the benchmark. To assert otherwise is to postulate an arbitrage opportunity. Hence, it is important to quantify this risk-reward tradeoff.

Consequently, we advocate the use of the clipped control from the cumulative difference objective function as a dynamic benchmark strategy. Since a closed form control is readily available, it would be straightforward to apply this clipped optimal control to historical return data of publicly traded assets.[4] This would then differentiate true investment skill from the easy gains due to dynamic trading.[5]

Of course, most of these pension plans employ a large universe of possible assets, including private equity and private credit. It is arguable that many of these alternative assets can be replicated using publicly traded factor portfolios (Ang, 2014). Hence, a better outperforming strategy would be an optimal dynamic strategy comprised of standard indexes and factor portfolios. We intend to report on this in our future work (van Staden et al., 2022).

7 Acknowledgements

Forsyth's work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2017-03760. Li's work was supported by he Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2020-04331.

The CRSP data were calculated based on data from Historical Indexes, ©2022 Center for Research in Security Prices (CRSP), the University of Chicago Booth School of Business. Wharton Research Data Services was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

8 Conflicts of interest

The authors have no conflicts of interest to report.

Appendices

[4] We could use historical data, known at the beginning of an investment period, to estimate market parameters. The parameter β in the objective function can then be adjusted to generate the desired IRR outperformance compared to the benchmark.

[5] It is interesting to note that the CPP 2021 annual report(Canadian Pension Plan, 2021) lists personnel costs as CAD 938 million, for 1,936 employees, giving an average cost of CAD 500,000 per employee-year.

A Informal derivation of equation (3.4)

We rewrite equation (2.2) informally as

$$\begin{aligned} \frac{dS_t}{S_{t^-}} &= (\mu - \lambda\kappa)\ dt + \sigma\, dZ_t + (\xi - 1) d\mathcal{Q} \\ d\mathcal{Q} &= \begin{cases} 0 & \text{Probability: } (1 - \lambda\ dt) \\ 1 & \text{Probability: } (\lambda\ dt) \end{cases} \end{aligned} \tag{A.1}$$

We can then write equations (2.4) and (2.5) as

$$\begin{aligned} dW_{t^-} &= W_{t^-}(r + p(\mu - r - \lambda\kappa))\ dt + q\ dt + pW_{t^-}\sigma\ dZ + pW_{t^-}(\xi - 1)d\mathcal{Q} \\ d\hat{W}_{t^-} &= \hat{W}_{t^-}(r + \hat{p}(\mu - r - \lambda\kappa))\ dt + q\ dt + \hat{p}\hat{W}_{t^-}\sigma\ dZ + \hat{p}\hat{W}_{t^-}(\xi - 1)d\mathcal{Q}\ . \end{aligned} \tag{A.2}$$

Now, given SDEs (A.2), value function $\tilde{V}(w, \hat{w}, t)$, with $W(t) = w, \hat{W}(t) = \hat{w}$, then Ito's Lemma (see (Tankov and Cont, 2009)) gives

$$\begin{aligned} d\tilde{V} &= \tilde{V}_t\ dt + (w(r + (\mu - r - \lambda\kappa)p) + q)\, \tilde{V}_w\ dt + (\hat{w}(r + (\mu - r - \lambda\kappa)\hat{p}) + q)\, \tilde{V}_{\hat{w}}\ dt \\ &\quad + \frac{p^2 w^2 \sigma^2}{2} \tilde{V}_{ww}\ dt + \frac{\hat{p}^2 \hat{w}^2 \sigma^2}{2} \tilde{V}_{\hat{w}\hat{w}}\ dt + \left(p\hat{p}w\hat{w}\sigma^2\right) \tilde{V}_{w\hat{w}}\ dt \\ &\quad + \left(\hat{p}\hat{w}\sigma\tilde{V}_{\hat{w}} + pw\sigma\tilde{V}_w\right)\ dZ_t \\ &\quad + \left(\tilde{V}(w + pw(\xi - 1), \hat{w} + \hat{p}\hat{w}(\xi - 1), t) - \tilde{V}(w, \hat{w}, t)\right)\ d\mathcal{Q}\ . \end{aligned} \tag{A.3}$$

Rewrite equation (3.2)

$$\begin{aligned} 0 &= \inf_p \Bigg\{ E_p^{(w,\hat{w},t)} \Bigg[\int_t^{t+\Delta t} (\hat{W}(s)e^{\beta s} - W(s))^2\ ds \\ &\qquad + \tilde{V}(W(t + \Delta t), \hat{W}(t + \Delta t), t + \Delta t) - \tilde{V}(w, \hat{w}, t) \Bigg| W(t) = w, \hat{W}(t) = \hat{w} \Bigg] \Bigg\} \\ &= \inf_p \left\{ E_p^{(w,\hat{w},t)} \left[(\hat{w}e^{\beta t} - w)^2\ dt + d\tilde{V} \right] \right\} \quad ; \quad \Delta t \to 0\ . \end{aligned} \tag{A.4}$$

Recall that $g(\xi)$ is the density of ξ. Substitute equation (A.3) into equation (A.4), noting that $E[dZ_t] = 0$ and $E[d\mathcal{Q}] = \lambda\ dt$ gives

$$\begin{aligned} 0 &= \inf_p \Bigg\{ \tilde{V}_t\ dt + (w(r + (\mu - r - \lambda\kappa)p) + q)\, \tilde{V}_w\ dt + (\hat{w}(r + (\mu - r - \lambda\kappa)\hat{p}) + q)\, \tilde{V}_{\hat{w}}\ dt \\ &\quad + \frac{p^2 w^2 \sigma^2}{2} \tilde{V}_{ww}\ dt + \frac{\hat{p}^2 \hat{w}^2 \sigma^2}{2} \tilde{V}_{\hat{w}\hat{w}}\ dt + \left(p\hat{p}w\hat{w}\sigma^2\right) \tilde{V}_{w\hat{w}}\ dt \\ &\quad + \lambda \left(\int_0^\infty \tilde{V}(w + pw(\xi - 1), \hat{w} + \hat{p}\hat{w}(\xi - 1), \tau) g(\xi)\ d\xi - \tilde{V} \right) \lambda\ dt + (\hat{w}e^{\beta t} - w)^2\ dt \Bigg\}\ . \end{aligned} \tag{A.5}$$

Now, define

$$\begin{aligned} \tau &= T - t \\ V(w, \hat{w}, \tau) &= \tilde{V}(w, \hat{w}, T - \tau)\ . \end{aligned} \tag{A.6}$$

Substitute equation (A.6) into (A.5) and divide by dt to obtain

$$V_\tau \;=\; \inf_p \mathcal{L}_p V \ , \tag{A.7}$$

where

$$\begin{aligned}\mathcal{L}_p V \;\equiv\; & \left(w(r + (\mu - r - \lambda\kappa)p) + q\right) V_w + \left((r + \hat{w}(\mu - r - \lambda\kappa)\hat{p}) + q\right) V_{\hat{w}} \\ & + \frac{p^2 w^2 \sigma^2}{2} V_{ww} + \frac{\hat{p}^2 \hat{w}^2 \sigma^2}{2} V_{\hat{w}\hat{w}} + \left(p\hat{p} w\hat{w}\sigma^2\right) V_{w\hat{w}} \\ & + \lambda \int_0^\infty V(w + pw(\xi - 1), \hat{w} + \hat{p}\hat{w}(\xi - 1), \tau) g(\xi) \ d\xi + (\hat{w}e^{\beta(T-\tau)} - w)^2 - \lambda V \ .\end{aligned} \tag{A.8}$$

References

Al-Aradi, A. and S. Jaimungal (2018). Outperformance and tracking: dynamic asset allocation for active and passive portfolio management. *Applied Mathematical Finance 25(3)*, 268–294.

Ang, A. (2014). *Asset Management: A Systematic Approach to Factor Investing*. New York: Oxford University Press.

Boyde, E. (2021). ETFs: why tracking difference usually matters more than tracking error. *Financial Times*. https://www.ft.com/content/44b08a6c-55d7-4841-8e8e-aad6451a4cc3.

Browne, S. (2000). Risk-constrained dynamic active portfolio management. *Management Science 46(9)*, 1188–1199.

Canadian Pension Plan (2021). CPP investments annual report. https://www.cppinvestments.com/the-fund/our-performance/financial-results.

Canadian Pension Plan (2022). Our investment strategy. https://www.cppinvestments.com/the-fund/how-we-invest/our-investment-strategy.

Charteris, A. and K. McCullough (2020). Tracking error vs. tracking difference: does it matter? *Investment Analysts Journal 49:3*, 269–287.

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

CPP Annual Report (2022). https://cdn4.cppinvestments.com/wp-content/uploads/2022/06/CPP-Investments_F2022-Annual-Report-EN.pdf.

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

Ennis, R. M. (2021). Failure of the endowment model. *The Journal of Portfolio Managment 47:5*, 128–143.

Forsyth, P. A. (2022). A stochastic control approach to defined contribution plan decumulation: the nastiest, hardest problem in finance. *North American Actuarial Journal 26:2*.

Government Pension Fund Global (2022). Investment strategy. https://www.nbim.no/en/the-fund/how-we-invest/investment-strategy/.

Hougan, M. (2015). Tracking difference: the perfect ETF metric. https://www.etf.com/sections/blog/tracking-difference-perfect-etf-metric.

Johnson, B., H. Bioy, A. Kellet, and L. Davidson (2013). On the right track: measuring tracking efficiency in ETFs. *The Journal of Index Investing 4:3*, 35–41.

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

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

L'Her, J. F., R. Stoyanova, K. Shaw, W. Scott, and C. Lai (2016). A bottom-up approach to the risk-adjusted performance of the buyout fund market. *Financial Analysts Journal 72:4*, 36–49.

Lin, Y., R. MacMinn, and R. Tian (2015). De-risking defined benefit plans. *Insurance: Mathematics and Economics 63*, 52–65.

Ma, K. and P. A. Forsyth (2017). An unconditionally monotone numerical scheme for the two factor uncertain volatility model. *IMA Journal on Numerical Analysis 37*, 905–944.

Ni, C., Y. Li, P. Forsyth, and R. Caroll (2022). Optimal asset allocation for outperforming a stochastic benchmark. *Quantitative Finance 22(9)*, 1595–1626.

Norges Bank (2021). Annual report. `https://www.nbim.no/en/publications/reports/2021/annual-report-2021/`.

Oderda, G. (2015). Stochastic portfolio theory optimization and the origin of rule based investing. *Quantitative Finance 15(8)*, 1259–1266.

Phalippou, L. (2014). Performace of buyout funds revisited. *Review of Finance 18*, 189–218.

Tankov, P. and R. Cont (2009). *Financial Modelling with Jump Processes.* New York: Chapman and Hall/CRC.

van Staden, P., D.-M. Dang, and P. A. Forsyth (2021). The surprising robustness of dynamic Mean-Variance portfolio optimization to model misspecification errors. *European Journal of Operational Research 289:1*, 774–792.

van Staden, P. M., P. A. Forsyth, and Y. Li (2022). Across-time risk-aware strategies for outperforming a benchmark. Working paper, Cheriton School of Computer Science, University of Waterloo.

van Staden, P. M., P. A. Forsyth, and Y. Li (2023). Beating a benchmark: dynamic programming may not be the right numerical approach. *SIAM Journal on Financial Mathematics 14:2*, 407–451.

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

Wander, B. H. (2000). The volatility of relative performance as a risk measure. *The Journal of Investing 9:2 (summer)*, 39–44.

Wang, J. and P. A. Forsyth (2010). Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation. *Journal of Economic Dynamics and Control 34*, 207–230.

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