Multi-Period Portfolio Optimization with Constraints and Transaction Costs

Joëlle Skaf and Stephen Boyd*

April 20, 2009

Abstract

We consider the problem of multi-period portfolio optimization over a finite horizon, with a self-financing budget constraint and arbitrary distribution of asset returns, with objective to minimize the mean-square deviation of final wealth from a given desired value. When there are no additional constraints, this problem can be solved by standard dynamic programming; the optimal trading policy is affine, *i.e.*, linear plus a constant. We describe a suboptimal policy that handles additional constraints on the portfolio or trading, such as linear transaction costs or a no-shorting constraint. The suboptimal policy involves solving an optimization problem, typically a convex quadratic program, at each step, using the Bellman (value) function for the associated unconstrained problem to approximately account for the value of future portfolios. Examples show that this suboptimal trading policy often obtains an objective value close to that for the associated problem without constraints, and is therefore nearly optimal. In particular we will see that even with transaction costs, our suboptimal trading policy performs almost as well as when there are no transaction costs.

1 Introduction

In this paper we formulate the multi-asset multi-period portfolio optimization problem as a stochastic control problem with linear dynamics and a convex quadratic objective, the mean-square error in achieving a desired final wealth. When there are no transactions costs, and the trading is self-financing, *i.e.*, the total revenue from sales equals the total cost of purchases, the optimal trading policy, which is affine (*i.e.*, linear plus a constant), can be found using dynamic programming (DP).

When transaction costs are present, or additional constraints are imposed, the optimal policy is very difficult to compute, even though it can be characterized easily using DP. (By gridding the value function, however, the case of one risk-free and one risky asset can be effectively solved.) In this paper we propose two suboptimal policies for the general

*Electrical Engineering Department, Stanford University.

constrained case. Both are based on the optimal policy for the associated unconstrained case. Simulations show that the more sophisticated suboptimal policy performs very well.

1.1 Previous and related work

There is a large body of work on dynamic portfolio optimization with constraints. Interest in the effects of transaction costs on portfolio optimization go back to Samuelson [28] and Constantinides [9, 23]. Literature varies with the choice of objective (usually, a utility to be maximized), continuous versus discrete time, finite versus infinite horizon, and so on. For a representative sample, see [1, 11, 12, 16, 21, 20] and [33]. Typical choices of utilities have been power (CRRA) utilities (*e.g.*, [21]) or log utilites (*e.g.*, [1]). Papers have most frequently dealt with the case of two assets (one risky and one risk-free) but there is also some published work for the case of multiple risky assets (*e.g.*, [20]). Our choice of objective, mean-square error in achieving a desired final wealth value, is not traditional in dynamic portfolio optimization, but is used in problems such as index tracking and portfolio replication. For more on these see, *e.g.*, [8, 13, 14].

2 Multi-period portfolio optimization problem

Portfolio evolution. We let $x_t \in \mathbf{R}^n$ be the vector (portfolio) of holdings (in dollars) in n assets, at the beginning of period t, for $t = 1, \ldots, T+1$, with negative entries denoting short positions. We assume that the initial portfolio x_1 is given. The wealth, or total portfolio value, at time period t is denoted $w_t = \mathbf{1}^T x_t$, where $\mathbf{1}$ is the vector with all components one. We let $u_t \in \mathbf{R}^n$, $t = 1, \ldots, T$, denote the vector of trades (in dollars) executed at the beginning of the period t, with positive entries denoting purchases and negative entries denoting sales. We let $x_t^+ = x_t + u_t \in \mathbf{R}^n$ denote the vector of holdings after the trades.

The holdings at the beginning of the next investment period are given by

$$x_{t+1} = A_t x_t^+ = A_t(x_t + u_t), \quad t = 1, \ldots, T,$$

where $A_t = \mathbf{diag}(r_t)$, with $r_t \in \mathbf{R}^n_+$ the vector of (random) asset returns. We will assume that r_t are independent, with known distributions. We denote the return mean and variances as

$$\bar{r}_t = \mathbf{E} r_t, \qquad \Sigma_t = \mathbf{E} r_t r_t^T - \bar{r}_t \bar{r}_t^T, \qquad t = 1, \ldots, T.$$

We assume that the second moment of r_t,

$$P_t = \Sigma_t + \bar{r}_t \bar{r}_t^T,$$

is positive definite, for $t = 1, \ldots, T$. (If this is not the case, there is a nonzero portfolio with certain return zero.) We can have, but do not require, an asset with a risk-free return. (This corresponds to a zero row and column in Σ_t.)

Trading policies. The goal is to find a trading policy, *i.e.*, functions $\varphi_1, \ldots, \varphi_T : \mathbf{R}^n \to \mathbf{R}^n$, with $u_t = \varphi_t(x_t)$. The trading policies must be such that the trades u_t satisfy

$$(u_t, x_t) \in \mathcal{C}_t, \qquad t = 1, \ldots, T, \tag{1}$$

where $\mathcal{C}_t \subseteq \mathbf{R}^{2n}$ is the constraint set for period t. In the simplest case, with no transaction costs and the requirement that the trading be self-financing, we have

$$\mathcal{C}_t = \mathcal{C}^{\text{basic}} = \{(x, u) \mid \mathbf{1}^T u = 0\}, \qquad t = 1, \ldots, T. \tag{2}$$

This constraint means that the total asset sales balances the total asset purchases in each period. We refer to this case as the *unconstrained case*.

We will always assume that

$$(x, u) \in \mathcal{C}_t \implies \mathbf{1}^T u \leq 0, \tag{3}$$

i.e., the total value of assets purchased is no more than the total value of the assets sold. We can interpret $-\mathbf{1}^T u_t$, which is the difference between the total value of assets sold and the total value of assets bought, as the transaction cost in period t. Our assumption is that the transaction costs are nonnegative; when $\mathcal{C}_t = \mathcal{C}^{\text{basic}}$, the transaction costs are zero.

For later reference, we describe some other possible constraint sets. We can model linear transaction costs by replacing the constraint $\mathbf{1}^T u = 0$ in (3) with

$$\mathbf{1}^T u + \kappa_{\text{buy}}^T u_+ + \kappa_{\text{sell}}^T u_- = 0, \tag{4}$$

where κ_{sell} is the (nonnegative) vector of selling transaction cost rates, κ_{buy} is the (nonnegative) vector of buying transaction cost rates, and $u_+ = \max(u, 0)$ and $u_- = \max(-u, 0)$ are the positive and negative parts of u, respectively. (We assume that $0 \leq \kappa_{\text{buy}} < 1$ and $0 \leq \kappa_{\text{sell}} < 1$.) The constraint (4) states that $-\mathbf{1}^T u$, which is the total gross proceeds from sales minus the total gross paid for purchases, equals $\kappa_{\text{buy}}^T u_+$, the total transaction cost for purchases, plus $\kappa_{\text{sell}}^T u_-$, the total transaction cost for sales.

We can impose a no-shorting constraint, as in

$$x_t^+ = x_t + u_t \geq 0. \tag{5}$$

This constraint states that after trading there are no short positions. Since the returns are nonnegative, this ensures that $x_{t+1} = A_t x_t^+ \geq 0$, and in particular, the wealth is always nonnegative, when the no-shorting constraint (5) is imposed.

In the general case, it can happen that there is no u_t for which $(u_t, x_t) \in \mathcal{C}_t$, which means that there is no feasible trade from the portfolio x_t. We refer to this event as *ruin*. For the basic unconstrained case (3) or with linear transactions costs (4), with or without the no-shorting constraint (5), ruin cannot occur, since $u_t = 0$ is always feasible, no matter what x_t is.

Objective. The final wealth is $w_{T+1} = \mathbf{1}^T x_{T+1}$. We take as objective the mean-square error,

$$J = \mathbf{E}(w_{T+1} - w^{\rm des})^2, \tag{6}$$

where $w^{\rm des} > 0$ is a desired final wealth. The square-root of J, which has units dollars, is the root-mean-square (RMS) error in achieving the desired wealth.

Our quadratic objective actually *penalizes* final wealth that exceeds the desired value, where of course we should be happy with such an outcome. This undesirable penalty on increasing final wealth, above the desired value, is shared with several other standard objectives, such as variance-adjusted mean return, which also penalizes large positive values of final wealth. For final wealth values less than the desired value, however, our objective provides the right incentive. A down-side mean-square error,

$$J^{\rm ds} = \mathbf{E}(w_{T+1} - w^{\rm des})_-^2,$$

better matches our true goals than does J. With this objective, however, we lose tractability of the unconstrained problem, which is the basis of our method for handling the constrained problem.

Our mean-square error objective is not a traditional one for dynamic portfolio optimization. A more typical objective is the expected value of a concave utility function of final wealth, which is to be maximized. Common examples include variance adjusted mean, log utility, and power (CRRA) utility. Our mean-square error objective is used in other contexts, such as index tracking and portfolio replication. Our goal in this paper is not to defend our choice of objective over others, but rather to point out that with this choice of objective, the optimal trading policy can be found when there are no transaction costs or other constraints, and what appear to be good suboptimal policies can be found when constraints are present.

Our mean-square error objective is easily related to the mean and variance of the final wealth, $\mathbf{E}w_{T+1}$ and $\mathbf{var} w_{T+1}$, which are traditional measures of portfolio performance, by expressing J as

$$J = (\mathbf{E}w_{T+1} - w^{\rm des})^2 + \mathbf{var} w_{T+1}.$$

Minimizing J is the same as maximizing

$$\mathbf{E}w_{T+1} - \gamma \mathbf{var} w_{T+1} - \gamma(\mathbf{E}w_{T+1})^2,$$

where $\gamma = 1/(2w^{\rm des})$. The first two terms here are a traditional variance adjusted mean utility.

Multi-period portfolio optimization problem. The multi-period portfolio problem is to determine trading policies $\varphi_1, \ldots, \varphi_T$, that satisfy the constraint (1), and minimize J. This is a stochastic control problem with linear dynamics (for more on stochastic control, see, *e.g.*, [3, 5, 6, 19, 25, 32]). We let $J^\star$ denote the optimal objective value, *i.e.*, the minimum possible value of J over all trading policies that satisfy the constraint.

3 Optimal policies for unconstrained case

For the unconstrained case, we can compute the optimal trading policies, which are affine, using DP (see, *e.g.*, [4, 5, 15, 26, 27]). Let $V_t(z)$ be the optimal objective value (*i.e.*, minimum possible value of J), of the truncated problem started in state $x_t = z$ at time period t. (Here we optimize over the policies $\varphi_t, \ldots, \varphi_T$.) Let $V_t^+(z)$ denote the optimal objective value (*i.e.*, minimum possible value of J) of the truncated problem started in post-trade state $x_t^+ = z$ at time period t. (Here we optimize over the policies $\varphi_{t+1}, \ldots, \varphi_T$.)

We will show that, for $t = 1, \ldots, T$, V_t and V_t^+ are convex quadratic functions with the specific forms

$$V_t(z) = a_t(\mathbf{1}^T z - w_t^{\mathrm{tar}})^2 + b_t, \tag{7}$$

and

$$V_t^+(z) = a_{t+1}\left((\bar{r}_t^T z - w_{t+1}^{\mathrm{tar}})^2 + z^T \Sigma_t z\right) + b_{t+1}, \tag{8}$$

where $a_1, \ldots, a_{T+1} > 0$, $w_1^{\mathrm{tar}}, \ldots, w_{T+1}^{\mathrm{tar}}$, and $b_1, \ldots, b_{T+1}$ will be defined below. Along the way we will show that the optimal policies are affine, *i.e.*, have the form

$$\varphi_t(z) = K_t(z - g_t),$$

where K_t and g_t will be defined below.

The form (7) states that V_t is a function of only the total portfolio value at time step t. This is easily explained. Going forward, no other attribute of the current portfolio matters: Since there are no transactions costs, we are free to select as the post-trade portfolio any one with the same total value. We can interpret w_t^{tar} as a time-varying target wealth. (As with the final desired wealth, it can also be interpreted as a wealth level above which our objective actually gives the wrong incentive.)

We will use induction, running backward from the last period $t = T$ to the first period $t = 1$, to establish (7) and (8). We first show that V_T^+ has the form (8). Then we show that, for $t = 1, \ldots, T$, if V_t^+ has the form (8), then V_t has the form (7). Finally, we show that, for $t = 2, \ldots, T$, if V_t has the form (7), then V_{t-1}^+ has the form (8).

Expression for V_T^+. To derive an expression for V_T^+, we assume that $x_T^+ = z$. Using $x_{T+1} = A_T x_T^+$, we have $w_{T+1} = \mathbf{1}^T A_T z = r_T^T z$, and so

$$V_T^+(z) = \mathbf{E}(w_{T+1} - w^{\mathrm{des}})^2 = (\bar{r}_T^T z - w^{\mathrm{des}})^2 + z^T \Sigma_T z,$$

which has the form (8), with

$$a_{T+1} = 1, \qquad w_{T+1}^{\mathrm{tar}} = w^{\mathrm{des}}, \qquad b_{T+1} = 0. \tag{9}$$

Expression for V_t from V_t^+. Now suppose that V_t^+ has the form (8). To find $V_t(z)$, we suppose that $x_t = z$ and $u_t = v$, which results in $x_t^+ = z + v$. From this state we follow the optimal policy, which yields objective $V_t^+(z+v)$ (by definition). So we must choose v to minimize $V_t^+(z+v)$ subject to $\mathbf{1}^T v = 0$:

$$\varphi_t(z) = \mathrm{argmin}_{\mathbf{1}^T v=0} V_t^+(z+v), \tag{10}$$

which results in optimal objective value

$$V_t(z) = \min_{\mathbf{1}^T v=0} V_t^+(z+v). \tag{11}$$

To minimize $V_t^+(z+v)$ we can just as well minimize

$$\begin{aligned}(V_t^+(z+v) - b_{t+1})/a_{t+1} &= (\bar{r}_t^T(z+v) - w_{t+1}^{\mathrm{tar}})^2 + (z+v)^T\Sigma_t(z+v)\\ &= (z+v)^T P_t(z+v) - 2w_{t+1}^{\mathrm{tar}}\bar{r}_t^T(z+v) + (w_{t+1}^{\mathrm{tar}})^2.\end{aligned}$$

A straightforward Lagrange multiplier argument tells us that the optimal post-trade state has the form

$$z + v = P_t^{-1}(\lambda\mathbf{1} + w_{t+1}^{\mathrm{tar}}\bar{r}_t),$$

where λ is chosen so that $\mathbf{1}^T v = 0$,

$$\lambda = \frac{\mathbf{1}^T z - w_{t+1}^{\mathrm{tar}}\mathbf{1}^T P_t^{-1}\bar{r}_t}{\mathbf{1}^T P_t^{-1}\mathbf{1}}.$$

Substituting this value of λ into our expression for $z+v$ we see that the optimal v is an affine function of z,

$$\varphi_t(z) = K_t(z - g_t), \tag{12}$$

where

$$K_t = -I + \frac{1}{\mathbf{1}^T P_t^{-1}\mathbf{1}} P_t^{-1}\mathbf{1}\mathbf{1}^T, \qquad g_t = w_{t+1}^{\mathrm{tar}} P_t^{-1}\bar{r}_t. \tag{13}$$

Using the optimal value of $z+v$ we obtain

$$\begin{aligned}V_t(z) &= a_{t+1}\left(\lambda^2(\mathbf{1}^T P_t^{-1}\mathbf{1}) + (w_{t+1}^{\mathrm{tar}})^2(1 - \bar{r}_t^T P_t^{-1}\bar{r}_t)\right) + b_{t+1}\\ &= a_t(\mathbf{1}^T z - w_t^{\mathrm{tar}})^2 + b_t,\end{aligned}$$

where

$$a_t = a_{t+1}/(\mathbf{1}^T P_t^{-1}\mathbf{1}), \tag{14}$$

$$w_t^{\mathrm{tar}} = w_{t+1}^{\mathrm{tar}}(\mathbf{1}^T P_t^{-1}\bar{r}_t), \tag{15}$$

$$b_t = b_{t+1} + a_{t+1}(w_{t+1}^{\mathrm{tar}})^2(1 - \bar{r}_t^T P_t^{-1}\bar{r}_t). \tag{16}$$

Since P_t is positive definite, $a_t > 0$, so $V_t(z)$ has the claimed form (7).

From (14), (15), and (9) we have the explicit expressions

$$a_t = \prod_{\tau=t}^{T}\frac{1}{\mathbf{1}^T P_\tau^{-1}\mathbf{1}}, \qquad w_t^{\mathrm{tar}} = w^{\mathrm{tar}}\prod_{\tau=t}^{T}\mathbf{1}^T P_\tau^{-1}\bar{r}_\tau. \tag{17}$$

From (16) and (9) we have

$$b_t = \sum_{\tau=t}^{T} a_{\tau+1}(w_{\tau+1}^{\mathrm{tar}})^2(1 - \bar{r}_\tau^T P_\tau^{-1}\bar{r}_\tau). \tag{18}$$

Expression for V_{t-1}^+ from V_t. Now suppose that V_t has the form (7). Assuming $x_{t-1}^+ = z$, we have $x_t = A_{t-1}z$ (which is random), from which point the optimal objective value is $V_t(x_t)$. It follows that the optimal objective value starting from $x_{t-1}^+ = z$ is

$$\begin{aligned} V_{t-1}^+(z) &= \mathbf{E}V_t(A_{t-1}z) \\ &= a_t\mathbf{E}(r_{t-1}^Tz - w_t^{\rm tar})^2 + b_t \\ &= a_t\left((\bar r_{t-1}^Tz - w_t^{\rm tar})^2 + z^T\Sigma_{t-1}z\right) + b_t. \end{aligned}$$

We have therefore shown that, if V_t has the form (7), V_{t-1}^+ has the form (8), for $t = 2, \ldots, T$.

Summary. We have shown that the pre- and post-trade optimal value functions V_t and V_t^+ have the forms given in (7) and (8), where a_t, $w_t^{\rm tar}$, and b_t are given in (17) and (18). The optimal policy is affine,

$$\varphi(z) = K_t(z - g_t),$$

where K_t and g_t are given in (13).

The optimal objective value for the stochastic control problem is given by

$$J^\star = V_1(x_1) = a_1(w_1 - w_1^{\rm tar})^2 + b_1. \tag{19}$$

Interpretations. We have seen that each post-trade state is a linear combination of $P_t^{-1}\bar r_t$ and $P_t^{-1}\mathbf{1}$. The portfolio $(1/\bar r_t^TP_t^{-1}\bar r_t)P_t^{-1}\bar r_t$ is the one that minimizes variance after one period of investment, $z^T\Sigma_tz$, subject to a mean return of one, $\bar r_t^Tz = 1$. (This constraint implies that $z^TP_tz = z^T\Sigma_tz + 1$.) When there is a risk-free asset, the minimum variance is zero and this portfolio is concentrated in the risk-free asset; see below. The portfolio $(1/\mathbf{1}^TP_t^{-1}\mathbf{1})P_t^{-1}\mathbf{1}$ is the portfolio that minimizes the return second moment z^TP_tz subject to unit investment, *i.e.*, $\mathbf{1}^Tz = 1$. It in turn can be expressed as a linear combination of the first portfolio, and one that minimizes one-step variance, subject to unit investment.

Each post-trade portfolio is on the mean-variance efficient frontier for the single period investment problem; $w_t^{\rm tar}$ determines which point on the frontier we choose in period t.

Risk-free asset. Suppose asset 1 is risk-free with return $(r_t)_1 = \mu_t > 0$. (We assume that there are no other risk-free assets.) This means that Σ_t has zero first row and column, with the remaining submatrix positive definite, and the first component of $\bar r_t$ is μ_t. Then $P_t^{-1}\bar r_t = (1/\mu_t)e_1$, where e_1 is the first standard unit vector. (This portfolio achieves zero one-step variance, with mean return one.)

Several other expressions appearing above simplify in this case. We have $\bar r_t^TP_t^{-1}\bar r_t = 1$, from which it follows that $b_t = 0$, and the optimal objective value is just $a_1(w_1 - w_1^{\rm tar})^2$. We have $\mathbf{1}^TP_t^{-1}\bar r_t = 1/\mu_t$, so $w_t^{\rm tar} = w^{\rm tar}\prod_{\tau=t}^T(1/\mu_\tau)$. In other words, the target just scales with the risk-free return at each step. The vector g_t has the simple form $g_t = \prod_{\tau=t}^T(1/\mu_\tau)e_1$.

Computation. The only significant computation is in computing $P_t^{-1}\bar{r}_t$ and $P_t^{-1}\mathbf{1}$. (If there is a risk-free asset, then only $P_t^{-1}\mathbf{1}$ involves computation.) If we exploit no structure in P_t, these can both be computed from one Cholesky factorization of P_t, which costs $n^3/3$ arithmetic operations. In the general case, then, the cost of determining the optimal policies is $O(Tn^3)$. The optimal gain matrix K_t is diagonal plus rank one, and can be stored in this form. Computing $u_t = K_t(x_t - g_t)$ can then be carried out with $O(n)$ cost. If P_t does not depend on time, we evidently need to compute $P_t^{-1}\bar{r}_t$ and $P_t^{-1}\mathbf{1}$ only once.

One common form for Σ_t, especially when n is large, is diagonal plus rank k, with $k \ll n$. (This corresponds to a factor model with k factors.) In this case P_t is diagonal plus rank $k+1$, so $P_t^{-1}\bar{r}_t$ and $P_t^{-1}\mathbf{1}$ can be computed efficiently using for example the Sherman Morrison Woodbury formula, at $O(nk^2)$ cost. For a problem with $n = 10000$ assets and $k = 30$ factors (say), $P_t^{-1}\bar{r}_t$ and $P_t^{-1}\mathbf{1}$ can be computed in a handful of milliseconds, on a typical 2GHz personal computer.

4 Suboptimal policies for the constrained case

In this section we describe two heuristics for finding suboptimal trading policies when the constraint sets $\mathcal{C}_t$ are not the basic one $\mathcal{C}^{\text{basic}}$, *e.g.*, if there are nonzero linear transaction costs or a no-shorting constraint is imposed.

The optimal objective value for the associated unconstrained case, obtained by replacing $\mathcal{C}_t$ with $\mathcal{C}^{\text{basic}}$, which is easily obtained using the methods described in §3, provides a lower bound on the optimal objective value for the constrained problem. To see this we note that an optimal policy for the constrained problem is also a feasible policy for the unconstrained problem, since $\mathcal{C}_t \subseteq \mathcal{C}^{\text{basic}}$, from which it follows that the optimal value obtained is no more than the optimal value for the unconstrained problem. For any suboptimal policy for the constrained case, we can compare its objective value obtained (evaluated using Monte Carlo, in general) to the optimal objective value for the associated unconstrained problem. If these numbers are close, we can be certain that the suboptimal policy is nearly optimal.

4.1 Projected affine policies

Our first suboptimal policy is simple: We simply project the optimal action (trade) for the associated unconstrained problem, given by $K_t(x_t - g_t)$, onto the constraint set:

$$u_t = \varphi_t^{\text{pa}}(x_t) = \text{argmin}_{(x_t,v)\in\mathcal{C}_t} \|v - K_t(x_t - g_t)\|_2. \tag{20}$$

The super script 'pa' stands for 'projected affine'. With the projected affine policies, we execute the feasible trade that is closest to the one for the associated unconstrained problem. In the terms of dynamic programming, φ_t^{pa} is a simple policy approximation method.

When $\mathcal{C}_t$ is convex, evaluating (20) requires solving a convex optimization problem, and so is tractable. When $\mathcal{C}_t$ is polyhedral, evaluating $\varphi_t^{\text{pa}}(x_t)$ involves solving a (convex) quadratic program (QP). In simple cases (*e.g.*, when the constraints include just linear transaction costs) we can work out an explicit form for, or simple algorithm for computing, $\varphi_t^{\text{pa}}(x_t)$.

When $\mathcal{C}_t$ is not convex, computing the projection can be a hard problem. In some cases, however, it can be efficiently computed, for example by solving a convex relaxation. For more on convex optimization, see, *e.g.*, [7].

Linear transaction costs. As an example, consider the case with linear transaction costs, but no other constraints,

$$\mathcal{C}^{\text{trans}} = \{(x,u) \mid \mathbf{1}^T u + \kappa_{\text{buy}}^T u_+ + \kappa_{\text{sell}}^T u_- = 0\}.$$

To evaluate u_t, we must solve the problem

$$\begin{array}{ll} \text{minimize} & \|v - u_t^{\text{opt}}\|_2 \\ \text{subject to} & \mathbf{1}^T v + \kappa_{\text{buy}}^T v_+ + \kappa_{\text{sell}}^T v_- = 0, \end{array} \qquad (21)$$

with variable v, where $u_t^{\text{uopt}} = K_t(x_t - g_t)$ is the optimal trade vector for the unconstrained problem. We first form the convex relaxation

$$\begin{array}{ll} \text{minimize} & \|v - u_t^{\text{opt}}\|_2 \\ \text{subject to} & \mathbf{1}^T v + \kappa_{\text{buy}}^T v_+ + \kappa_{\text{sell}}^T v_- \leq 0, \end{array} \qquad (22)$$

with variable v.

The relaxed problem (22) can be solved; using a simple Lagrange multiplier argument, it can be shown that

$$u_t = (u_t^{\text{uopt}} - \lambda(\mathbf{1} + \kappa_{\text{buy}}))_+ - (u_t^{\text{uopt}} - \lambda(\mathbf{1} - \kappa_{\text{sell}}))_-,$$

where $u_t^{\text{uopt}} = K_t(x_t - g_t)$ is the optimal trade vector for the unconstrained problem, and λ is the solution of the equation

$$\mathbf{1}^T u_t + \kappa_{\text{buy}}^T (u_t)_+ + \kappa_{\text{sell}}^T (u_t)_- = 0. \qquad (23)$$

When $u_t^{\text{uopt}} \neq 0$, the lefthand side is a decreasing piecewise linear function of λ, which is positive for $\lambda = 0$ and zero for $\lambda = \max_i u_i/(1 + (\kappa_{\text{buy}})_i)$, so the solution is readily found by bisection. Since u_t satisfies (23), it is evidently feasible for, and therefore optimal for, the nonconvex problem (20). (Note that we use $\mathbf{1}^T u_t^{\text{opt}} = 0$ to show that the relaxed solution solves the original problem; for general u_t^{opt}, the solutions of (20) and (22) need not be the same.)

No-shorting constraint. With no transaction costs and no-shorting constraints, *i.e.*,

$$\mathcal{C}_t = \{(x,u) \mid \mathbf{1}^T u = 0, \ x + u \geq 0\},$$

evaluating $\varphi_t^{\text{pa}}(x_t)$ entails solving a convex QP, which has the simple solution

$$u_t = \min(u_t^{\text{opt}}, -x_t).$$

4.2 Control-Lyapunov policies

Our second suboptimal trading policy is motivated by (12), which states that the optimal trade is the feasible one that maximizes the post-trade optimal value function. The policy is

$$u_t = \varphi_t^{\rm clf}(x_t) = \mathop{\rm argmin}_{(x_t,v)\in\mathcal{C}_t} V_t^+(x_t+v), \tag{24}$$

where V_t^+ is the optimal post-trade value function for the associated unconstrained problem, described in §3. When $\mathcal{C}_t$ is convex, evaluating $\varphi_t^{\rm clf}(x_t)$ is a convex optimization problem; when $\mathcal{C}_t$ is polyhderal, it is a QP. When $\mathcal{C}_t$ is not convex, evaluating $\varphi_t^{\rm clf}(x_t)$ can be hard, and we may need to resort to an approximation, for example by solving a convex relaxation.

The trading policy defined by (24) is called a *control-Lyapunov policy*; here V_t^+ is called the associated *control-Lyapunov function*. If V_t^+ were replaced by the true post-trade optimal value function for the constrained problem (which is not in general quadratic), then (24) would give the optimal policy. In the terms of dynamic programming, $\varphi_t^{\rm clf}$ is a simple value function approximation method. For more on control-Lyapunov policies, see [10, 17, 29, 30, 31].

We can give a simple interpretation of (24). By solving (24), we find the optimal trade at time period t, assuming that, from time period $t+1$ on, there are no further constraints or trading costs, *i.e.*, we replace $\mathcal{C}_{t+1},\ldots,\mathcal{C}_T$ with $\mathcal{C}^{\rm basic}$. Thus, we are underestimating the true optimal objective value, and the trading policy (24) can therefore, and roughly speaking, result in more trading than would take place under the optimal policy. But we will see that it still often results in a very good trading policy.

No trade zone. When $\mathcal{C}_t$ is a convex cone, the (necessary and sufficient) optimality condition for (24) is

$$(x_t,v)\in\mathcal{C}_t, \qquad \nabla V_t^+(x_t+v)\in\mathcal{C}_t^*, \qquad \nabla V_t^+(x_t+v)^Tv=0,$$

where $\mathcal{C}_t^*$ is the dual cone of $\mathcal{C}_t$. (See, *e.g.*, [7].) From this we can find the necessary and sufficient condition under which $v=0$ is a solution of (24):

$$(x_t,0)\in\mathcal{C}_t, \qquad \nabla V_t^+(x_t)\in\mathcal{C}_t^*. \tag{25}$$

If $\mathcal{C}_t^*$ has nonempty interior (which occurs when $\mathcal{C}_t$ is pointed) then (25) defines a cone of portfolios, with nonempty interior, for which no trading is done. In other words, (25) defines a no-trade zone.

Let us work out this condition more explicitly for the specific case of linear transaction costs, with

$$\mathcal{C}_t = \{u \mid \mathbf{1}^Tu + \kappa_{\rm buy}^T u_+ + \kappa_{\rm sell}^T u_- \leq 0\}.$$

The dual cone is

$$\mathcal{C}_t^* = \{c \mid -(\mathbf{1}+\kappa_{\rm buy}) \leq c/\nu \leq -(\mathbf{1}-\kappa_{\rm sell}) \text{ for some } \nu>0\}\cup\{0\}.$$

Therefore there is no trading in time period t if and only if

$$\max_{i=1,\ldots,n} \frac{d_i}{1+(\kappa_{\text{buy}})_i} \leq \min_{i=1,\ldots,n} \frac{d_i}{1-(\kappa_{\text{sell}})_i}.$$

where $d = -(P_t x_t - w_{t+1}^{\text{tar}} \bar{r}_t)$.

Self-validating property. The control-Lyapunov policy is self-validating, in the following sense. If we replace V_t^+ in (24) with the actual cost-to-go function for the constrained problem, then (24) defines the optimal policy for the constrained problem. If the actual cost-to-go function is close to the quadratic function V_t^+, we can guess that the control-Lyapunov policy will be close to optimal. From this we can guess that the performance will be nearly optimal, which in turn suggests that the actual cost-to-go function will be close to the quadratic function V_t^+. In summary, and very roughly: If the control-Lyapunov policy achieves performance that is close to the performance obtained for the unconstrained problem, the assumption on which it is based, *i.e.*, that V_t^+ is a good approximation of the actual cost-to-go function, will be valid.

5 Example

In this section we illustrate the trading algorithms described above with a numerical example with simulated returns. We first describe how the return model was generated. Our goal is to obtain a simple return model that captures (at least crudely) some of the typical features seen in real return models.

5.1 Return model

We will let p denote the number of trading days per year and q the number of years that will be considered, so the total number of trading periods is $T = pq$. We take the returns r_t to be independent identically distributed with a log-normal distribution, *i.e.*,

$$\log(r_t) \sim \mathcal{N}(\mu/p, S/p),$$

where μ and S are the mean and variance of the log annual returns. We choose S and μ as follows. Asset 1 will be a risk free asset with log annual return $\mu_1 = \mu_{\text{rf}}$, so $S_{ij} = 0$ for $i = 1$ or $j = 1$. For the remaining assets we set

$$S_{ii} = \left(\sigma_{\max}(i-1)/n\right)^2, \quad i = 1, \ldots, n,$$

so the assets have log annual return standard deviations linearly varying from 0 (the risk free asset) to $\sigma_{\max}^2$ (the riskiest asset). We set log annual return means as

$$\mu_i = \mu_{\text{rf}} + \rho S_{ii}, \quad i = 1, \ldots, n,$$

where ρ is the reward-to-risk ratio.

The off-diagonal elements of S are given by

$$S_{ij} = (S_{ii}S_{jj})^{1/2}\, C_{ij}, \quad i,j = 1,\ldots,n,$$

Here C is the matrix of log annual return correlation coefficients. We choose C so that the correlations between log annual returns range from around -0.1 to 0.9 or so. To do this we generate a matrix $Z \in \mathbf{R}^{n\times n}$ will all entries drawn from a standard Gaussian distribution, then form $Y = ZZ^T + \lambda \mathbf{1}\mathbf{1}^T$, where $\lambda > 0$, and finally we take

$$C = \mathbf{diag}(Y_{11}^{-1/2},\ldots,Y_{nn}^{-1/2})\, Y\, \mathbf{diag}(Y_{11}^{-1/2},\ldots,Y_{nn}^{-1/2}).$$

We choose λ so the minimum entry of C is around -0.1.

We now describe the particular numerical instance we use in the simulations. We have $n = 10$ assets, and trade monthly for 10 years, *i.e.*, $p = 12$, $q = 10$, so $T = 120$. The risk-free log annual return is $\mu_{\rm rf} = 2\%$, and the reward-to-risk ratio is $\rho = 0.4$. We take $\sigma_{\max} = 40\%$, so the asset log annual return standard deviations range from 0 to 25% in 4.44% increments, and the log annual mean returns range from 2% to 18% in 1.78% increments.

We emphasize that the details of our return statistics are given here only for completeness; we have tried our methods with several other return statistics, with similar results.

5.2 Simulations

Our initial portfolio has total value one entirely invested in the risk-free asset, *i.e.*, $x_1 = e_1$. We take the desired final wealth to be $w^{\rm des} = 2.5937$, which represents a log annual growth rate of 10% per year, which is in the middle of the range of asset returns.

We generate $N = 1000$ return realizations, and for each realization we simulate the portfolio evolution with several constraint sets and trading policies.

- *Unconstrained case.*
  - Optimal policies.
  - Optimal buy-and-hold strategy, in which we choose u_1 to optimize the objective with $u_2 = \cdots = u_T = 0$. (The optimal u_1 can be found by solving a QP.)
- *Linear transaction costs.* We impose a 0.25% transaction cost for buying and selling all assets except the risk-free one.
  - Projected affine policies.
  - Control-Lyapunov policies.
  - Optimal buy-and-hold strategy, which can be found by solving a QP.
- *Linear transaction costs with no-shorting constraint.*
  - Projected affine policies.
  - Control-Lyapunov policies.
  - Optimal buy-and-hold strategy, which can be found by solving a QP.

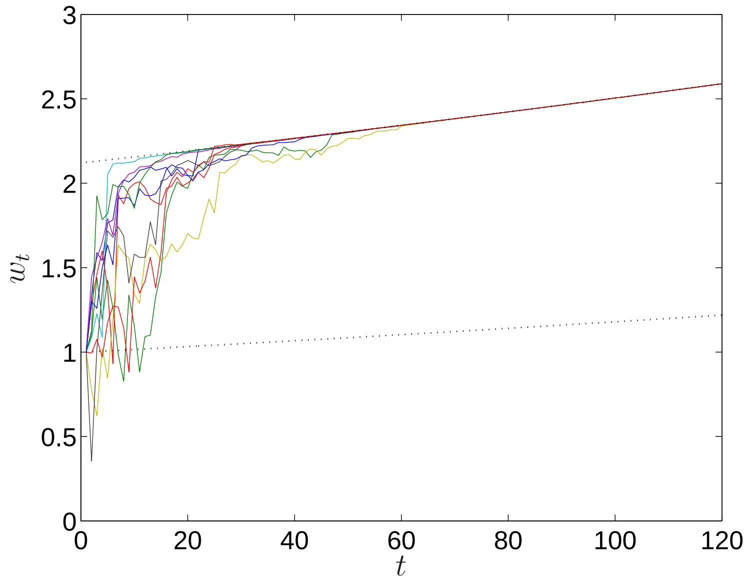

Figure 1: Wealth trajectories for unconstrained case with optimal policies.

5.3 Unconstrained case

From (19) we find that $J^\star = 2 \times 10^{-6}$, so the RMS final wealth error is $\sqrt{J^\star} = 0.0014$. This small value means that the final wealth obtained with the optimal policies is very near the desired final wealth. This is confirmed in our simulations of 1000 trajectories, for which the final wealth has average 2.5937 (which is the desired wealth to four significant figures) and (negligible) standard deviation 4×10^{-11}. Ten wealth trajectories are shown in Figure 1. The figure also shows two dotted curves: 1.02^t (the growth of one dollar invested in the risk-free asset), and $w_t^{\rm tar} = 2.5937(1.02)^{T-t}$ (the desired final wealth's present value at time t, using the risk-free rate).

The optimal trading policy starts with aggressive, highly leveraged trading. After the first trade (which is the same for all trajectories since the initial portfolio is always $x_1 = e_1$), the total short position, $\mathbf{1}^T(x_2)_-$, is around 60. Once the wealth gets near the target wealth value, however, most of the portfolio is shifted to, and maintained in, the risk-free asset.

For one of the ten trajectories plotted, the total wealth drops to smaller than half the original wealth before regaining value, and ultimately finishing with a wealth very close to the desired value (as all trajectories do). In fact, the wealth can drop to a *negative* value before recovering. This occurs in 7.7% of the trajectories. (The objective in our problem formulation does not depend on intermediate wealth values, so we cannot complain about this. Methods described in §6 can be used to reduce large fluctations in intermediate wealth.)

We also simulate the optimal buy-and-hold strategy, for comparison. This has optimal objective value (obtained from the QP) 0.4845, which is consistent with the empirical value from our simulations, 0.4820. The average final wealth is 2.2407, with standard deviation

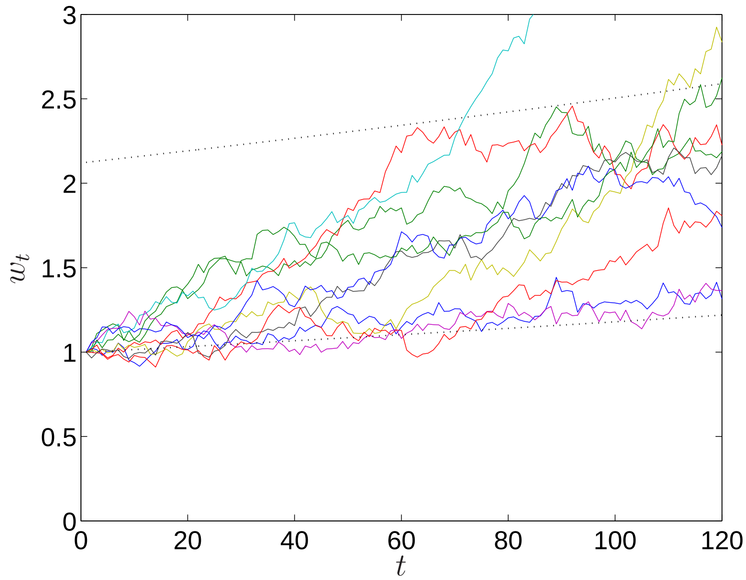

Figure 2: Wealth trajectories for unconstrained case, with optimal buy-and-hold policy.

0.5999. Figure 2 shows wealth trajectories for the optimal buy-and-hold strategy. One (favorable) trajectory goes off our plot, reaching a final wealth of $w_{T+1} \approx 4.5$.

We can see that the final wealth has a very large standard deviation, as we would expect. For example, the probability of the final wealth falling below a 9.5% annual return is 67.4% for the optimal buy-and-hold policy (whereas, in comparison, it is essentially zero for the the optimal trading policy).

5.4 Linear transaction costs

We now consider the case with 0.25% linear buying and selling transaction costs on all assets except the risk-free asset. The projected affine, control-Lyapunov, and optimal buy-and-hold strategies are simulated for each realization.

The projected affine policy is found using the simple method described in §4.1. For the control-Lyapunov policy, we solve the convex relaxation of the problem, using the (convex) constraint

$$\mathbf{1}^T u_t + \kappa_{\text{buy}}^T (u_t)_+ \kappa_{\text{sell}}^T (u_t)_- \leq 0, \qquad (26)$$

which yields a QP. This constraint allows for the possibility of discarding money, and while we cannot prove that the relaxation is always tight, this is the case is all of our simulations. Simulation of the control-Lyapunov policies required the solution of 120000 QPs. To do this in reasonable time, we implemented a basic primal-dual interior-point method for this specific QP, using the C code generation feature of CVXMOD [24]. This custom C code solves the QP in around 100μsec on a 2GHz PC. (The results were verified using CVX [18].)

The control-Lyapunov policies perform very well, with average cost $J = 2 \times 10^{-5}$, corresponding to RMS final wealth error $\sqrt{J} = 0.0045$, which is small enough to mean that, as with the optimal policy for the unconstrained case, the final wealth is always very close to the desired final wealth. Indeed, the final wealth has average (over the 1000 realizations) 2.5914 and standard deviation 0.0039.

The projected affine policies perform a bit worse, obtaining $J = 0.0053$, corresponding to RMS final wealth error $\sqrt{J} = 0.0728$. The final wealth has average 2.5506 and standard deviation 0.0583.

For the optimal buy-and-hold policy, the objective value is 0.4897, corresponding to RMS final wealth error $\sqrt{J} = 0.6998$. The final wealth has a mean 2.2369 and standard deviation 0.6020.

Figures 3, 4, and 5 show ten wealth trajectories when the projected affine, the control-Lyapunov, and the optimal buy-and-hold policies are used, respectively. The probability of the final wealth falling below 9.5% annual return is 67.3% for the optimal buy-and-hold policy. It is 5.2% for the projected affine policy, and 0% (none out of our 1000 trajectories) for the control-Lyapunov policy.

As in the optimal policies for the unconstrained case, there is agressive leveraging in the first step of the projected affine policy, with a total short position in x_2 around 60. With the control-Lyapunov policy, however, there is much less leveraging; the total short position in x_2 is around 3. The projected affine policies trade at each step, but the control-Lyapunov policies do not trade around 9% of the time.

The plots show that the total wealth can become negative at some intermediate time; for example, one trajectory in Figure 3 drops to a wealth of -0.5 dollars before recovering. This occurs in 23.3% of the trajectories for projected affine policies, 8.7% of the trajectories for control-Lyapunov policies, and never for the optimal buy-and-hold policy.

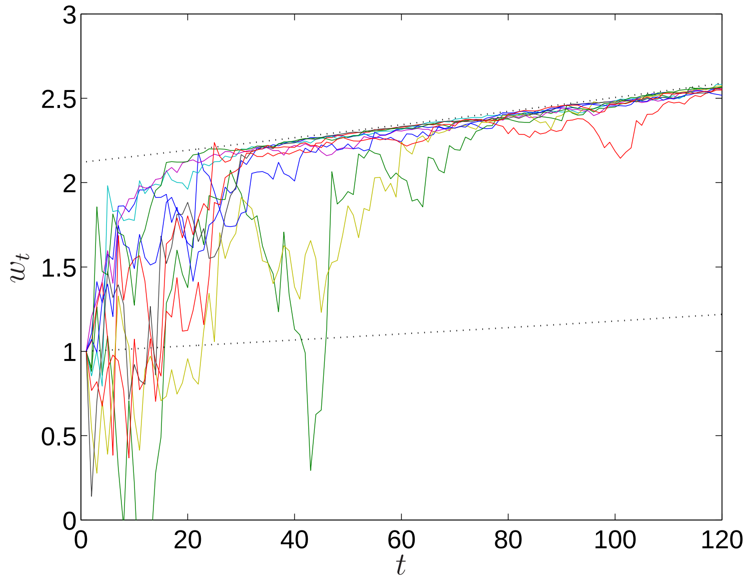

Figure 3: Some wealth trajectories for the projected affine policy with linear transaction costs.

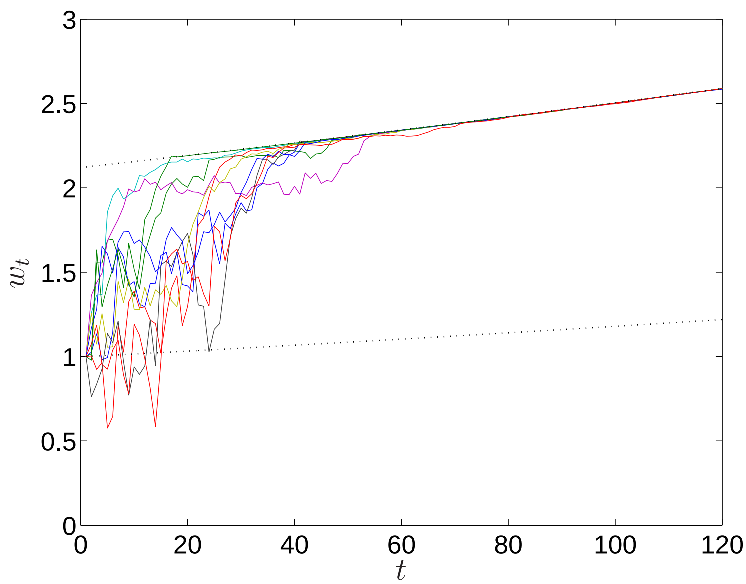

Figure 4: Some wealth trajectories for the control-Lyapunov policy with linear transaction costs.

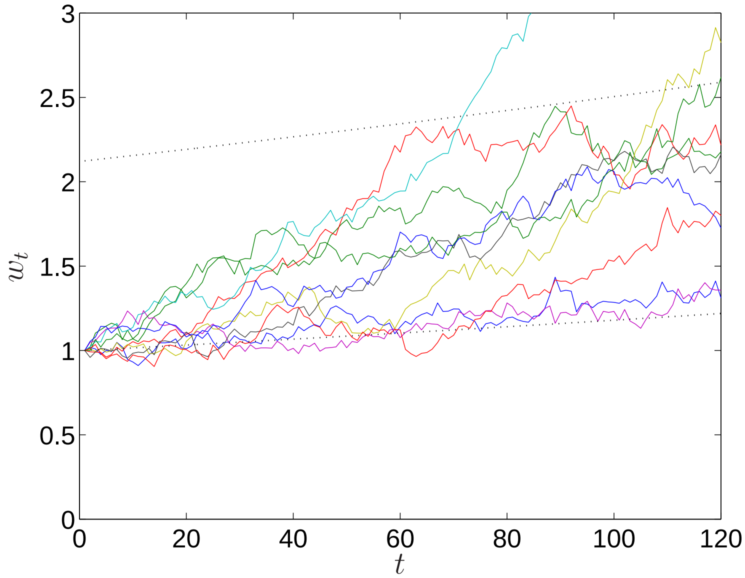

Figure 5: Some wealth trajectories for the optimal buy-and-hold policy with linear transaction costs.

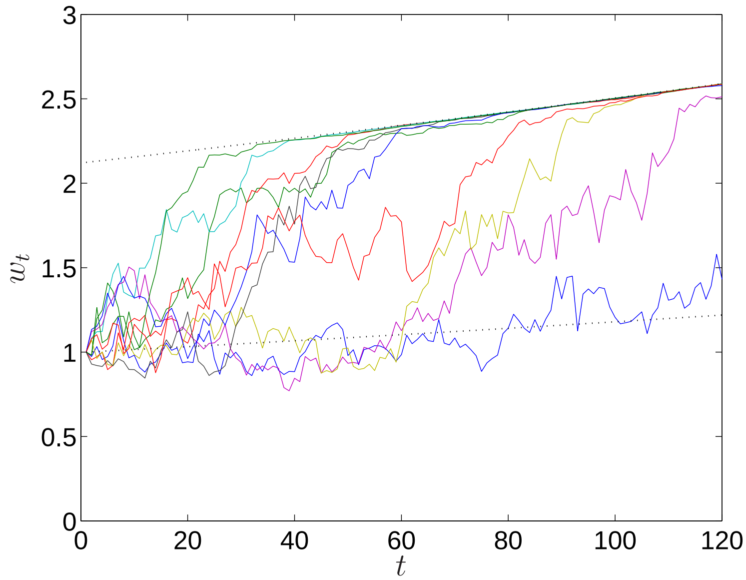

Figure 6: Some wealth trajectories for the control-Lyapunov policy with transaction costs and no-shorting constraints.

5.5 Linear transaction costs and no-shorting constraint

We now consider the case with 0.25% linear transaction costs (except on the risk-free asset) and a no-shorting constraint. The projected affine, control-Lyapunov, and buy-and-hold strategies are simulated for each realization, using the convex relaxation (26) instead of the linear transaction cost constraint. To solve the 240000 QPs for the simulations, we used a custom generated primal-dual solver as described above for the linear transactions case.

The control-Lyapunov policies obtain $J = 0.0773$, corresponding to RMS final wealth error 0.2780. This is significant, especially when compared to the unconstrained or linear transaction cost cases considered above. The final wealth has mean 2.5298 and standard deviation 0.2707.

The projected affine policies perform much worse, obtaining $J = 1.8842$. The final wealth has mean 1.2211 and standard deviation 0. In fact, the projected affine policies never trade; all final wealth values are an initial investment of 1 help in the risk-free asset the entire time. For the optimal buy-and-hold policy, the objective value is 0.5500; the final wealth has a mean 2.1485 and standard deviation 0.5931.

Figures 6 and 7 show ten wealth trajectories when the control-Lyapunov and the buy-and-hold policies are used, respectively.

The probability of the final wealth falling below 9.5% annual return is 73.7% for the optimal buy-and-hold policy, 100% for the projected affine policy, and 6.3% for the control-Lyapunov policy. The control-Lyapunov policy does not trade 19% of the time.

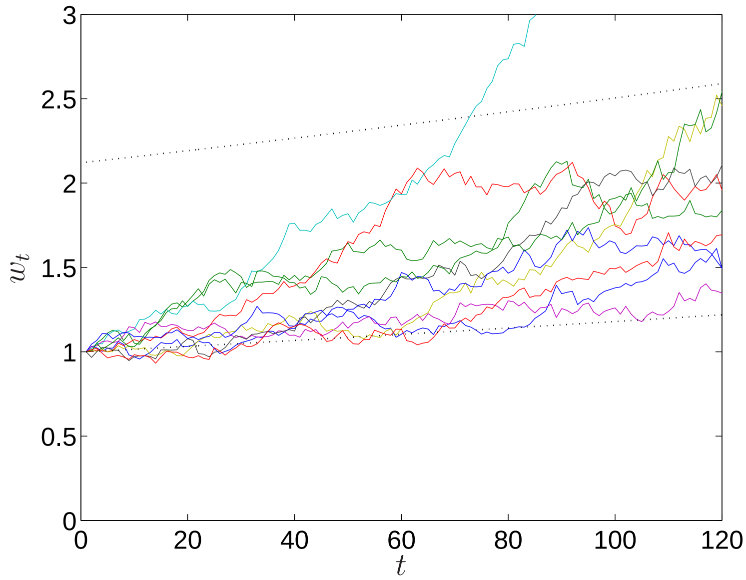

Figure 7: Some wealth trajectories for the optimal buy-and-hold policy with transaction costs and no-shorting constraints.

6 Variations and extensions

The methods described in this paper can be generalized and modified in many ways, some of which we describe here.

General linear dynamics. We have taken A_t to be diagonal, but everything works when A_t has off-diagonal elements (with suitable modification of the formulas). Off-diagonal elements can be used to model (random) dividend payments.

Cash in and out. So far we have required self financing. But we can allow cash in and cash out (as in retirement planning). In the unconstrained case we replace $\mathbf{1}^T u_t = 0$ with $\mathbf{1}^T u_t = d_t$, where d_t is the deposit into the account (if positive) or withdrawal (if negative).

Constraint set. A wide variety of constraints can be handled, either exactly (if convex) or approximately (if nonconvex). Examples include maximum and minimum allowed values for u_t and x_t^+, or a leverage limit such as

$$\mathbf{1}^T(x_t + u_t)_- \leq \eta \mathbf{1}^T(x_t + u_t)_+,$$

which limits the total short position to a factor η times the total long position. (This can be re-written as a convex cone constraint on (x_t, u_t).)

Other transaction costs. We have use a simple linear transaction cost model, but many others can be used, including quadratic or fixed-plus-linear transaction costs. (The latter can be handled using the methods described in [22].)

Objective function. We can add a cost for tracking a wealth trajectory along the way, such as

$$J = \mathbf{E} \sum_{t=2}^{T+1} \alpha_t (w_t - w_t^{\mathrm{des}})^2,$$

where $(w_2^{\mathrm{des}}, \ldots, w_{T+1}^{\mathrm{des}})$ is a desired trajectory that we want to track, and $\alpha_t \geq 0$ are weights. Taking $\alpha_{T+1} = 1$ and $\alpha_t = 0$ for $t = 2, \ldots, T$ recovers the problem presented in §2. One simple choice for the target wealth trajectory is $w_t^{\mathrm{des}} = w_1 (w^{\mathrm{des}}/w_1)^{t/T}$, which corresponds to a constant wealth growth rate.

We can also add a quadratic penalty to penalize large trades,

$$J = \mathbf{E} \sum_{t=2}^{T+1} \alpha_t (w_t - w_t^{\mathrm{des}})^2 + \sum_{t=1}^{T} \rho_t \|u_t\|^2, \tag{27}$$

where $\rho_t \geq 0$ are trading penalty weights.

In both of these cases, the methods described in §3 can be used to work out the optimal policy for the unconstrained case. Both V_t and V_t^+ are convex quadratic functions, which can be found using DP. (The particular formulas for these differ from the ones given in §3.)

7 Conclusions

We have shown that, for the unconstrained case with a mean-square final wealth error objective, the multi-period portfolio optimization problem can be solved exactly, using DP. The optimal trading strategy is affine, and the pre- and post-trade optimal value functions are convex quadratic.

We proposed two suboptimal policies for the case when there are transaction costs or constraints. Our examples show that the first one, which simply projects the optimal trade for the unconstrained case onto the constraint set, can work reasonably well in some cases, but can fail in others. The second policy, of the control-Lyapunov type, however, works well in more cases. In some cases (for example, with linear transaction costs), it can deliver performance that is very close to the unconstrained case. Even in cases where the basic idea behind the control-Lyapunov policy does not apply (*i.e.*, the actual value function for the constrained problem is not close to the value function for the unconstrained problem), the control-Lyapunov policies seem to do very well.

Our final comment is about how the methods of this paper might be used. Control policies designed by linear-quadratic stochastic control methods are widely used in traditional control engineering applications. In these applications, the quadratic objective functions are not considered to be the true engineering objectives; rather they are thought of as surrogates for the true engineering objectives. The weights in the objective function are then tuned to give good simulation results. For some description of this, see, *e.g.*, [2].

We suggest the same approach can be used for multi-period portfolio optimization, using control-Lyapunov policies to handle constraints. We use the objective (27) instead of the objective (6), and tune the parameters and weights (*i.e.*, w_t^{des}, α_t, ρ_t) to obtain good performance, as judged by Monte Carlo simulation.

Acknowledgment

We thank Jacob Mattingley for help in generating the custom QP solvers used for the simulations.

References

[1] M. Akian, A. Sulem, and M. Taksar. Dynamic optimization of long-term growth rate for a portfolio with transaction costs and logarithmic utility. *Mathematical Finance*, 11(2):152–188, 2001.

[2] B. Anderson and J. Moore. *Optimal Control - Linear Quadratic Methods*. Prentice-Hall, 1990.

[3] K. Åström. *Introduction to Stochastic Control Theory*. Dover Publications, 2006.

[4] R. Bellman. *Dynamic Programming*. Courier Dover Publications, 1957.

[5] D. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.

[6] J. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer, 1997.

[7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[8] I. Buckley and R. Korn. *Optimal Cash Management for Equity Index Tracking in the Presence of Fixed and Proportional Transaction Costs*. Johannes-Gutenberg-Univ., 1997.

[9] G. Constantinides. Multiperiod consumption and investment behavior with convex transaction costs. *Management Science*, 25:1127–1137, 1979.

[10] M. Corless and G. Leitmann. Controller design for uncertain system via Lyapunov functions. In *Proceedings of the American Control Conference*, volume 3, pages 2019–2025, 1988.

[11] J. Cvitanić and I. Karatzas. Hedging and portfolio optimization under transaction costs: A martingale approach. *Mathematical Finance*, 6:133–166, 1996.

[12] M. Davis and A. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713, 1990.

[13] R. Dembo. Optimal portfolio replication. *Algorithmics Technical paper series*, 9501, 1995.

[14] R. Dembo and D. Rosen. The practice of portfolio replication. *Algo Research Quarterly*, 3(2):11–22, 2000.

[15] E. Denardo. *Dynamic Programming: Models and Applications*. Prentice-Hall, 1982.

[16] B. Dumas and E. Luciano. An exact solution to a dynamic portfolio choice problem under transaction costs. *The Journal of Finance*, 46(2):577–595, 1991.

[17] R. Freeman and J. Primbs. Control Lyapunov functions: New ideas from an old source. In *Proceedings pf the 35th IEEE Conference on Decision and Control*, pages 3926–3931, Kobe, Japan, 1996.

[18] M. Grant, S. Boyd, and Y. Ye. CVX: Matlab software for disciplined convex programming, 2006. Available at http://www. stanford. edu/ boyd/cvx.

[19] P. Kumar and P. Varaiya. *Stochastic systems: Estimation, identification and adaptive control*. Prentice-Hall, 1986.

[20] H. Liu. Optimal consumption and investment with transaction costs and multiple risky assets. *The Journal of Finance*, 59(1):289–338, 2005.

[21] H. Liu and M. Loewenstein. Optimal portfolio selection with transaction costs and finite horizons. *Review of Financial Studies*, 15(3):805–835, 2002.

[22] M. Lobo, M. Fazel, and S. Boyd. Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research*, 152(1):376–394, 2007.

[23] M. Magill and G. Constantinides. Portfolio selection with transaction costs. *J. Econ. Theory*, 13:245–263, 1976.

[24] J. Mattingley and S. Boyd. CVXMOD - Convex optimization software in python, 2008. Available at http://cvxmod.net.

[25] A. Prekopa. *Stochastic Programming*. Kluwer Academic Publishers, 1995.

[26] M. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc. New York, NY, USA, 1994.

[27] S. Ross. *Introduction to Stochastic Dynamic Programming: Probability and Mathematical*. Academic Press, Inc. Orlando, FL, USA, 1983.

[28] P. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51:239–246, 1969.

[29] E. D. Sontag. A Lyapunov-like characterization of asymptotic controllability. *SIAM Journal on Control and Optimization*, 21(3):462–471, 1983.

[30] M. Sznaier, R. Suarez, and J. Cloutier. Suboptimal control of constrained nonlinear systems via receding horizon constrained control Lyapunov functions. *International Journal on Robust and Nonlinear Control*, 13(3-4):247–259, 2003.

[31] Y. Wang and S. Boyd. Performance bounds for linear stochastic control. *To appear, Systems and Control Letters*, 2008.

[32] P. Whittle. *Optimization Over Time: Dynamic Programming and Stochastic Control*. Wiley, 1982.

[33] W. Ziemba and R. Vickson. *Stochastic Optimization Models in Finance*. World Scientific, 2006.