Optimal Rebalancing Strategy for Institutional Portfolios

Walter Sun[1,†], Ayres Fan[1], Li-Wei Chen[1], Tom Schouwenaars[1], Marius A. Albota[2], Ed Freyfogle[3], Josh Grover[3]

Abstract

Institutional fund managers generally rebalance using *ad hoc* methods such as calendar basis or tolerance band triggers. We propose a different framework that quantifies the cost of a rebalancing strategy in terms of risk-adjusted returns net of transaction costs. We then develop an optimal rebalancing strategy that actively seeks to minimize that cost. We use certainty equivalents and the transaction costs associated with a policy to define a cost-to-go function, and we minimize this expected cost-to-go using dynamic programming. We apply Monte Carlo simulations to demonstrate that our method outperforms traditional rebalancing strategies like monthly, quarterly, annual, and 5% tolerance rebalancing. We also show the robustness of our method to model error by performing sensitivity analyses.

[1]Laboratory for Information and Decision Systems, Massachusetts Institute of Technology

[2]Research Laboratory of Electronics, Massachusetts Institute of Technology

[3]Sloan School of Management, Massachusetts Institute of Technology

[†]Corresponding author: waltsun@mit.edu

CONTENTS

I. INTRODUCTION

Institutional money managers develop risk models and optimal portfolios to match a desired risk/reward profile. Utility functions provide one method for modeling these risk preferences. Implicitly determined by the views of trustees or directors, utility functions yield a target portfolio which is a set of optimal weights for different asset classes. Given the fact that different asset classes can exhibit different rates of return, a manager cannot maintain this target of weights over time without active rebalancing. Furthermore, managers also must rebalance if the weights in the target portfolio are altered. This occurs when the model for expected returns of asset classes change or the risk profile is altered (as indicated by a change in the utility function).

Most academic theory ignores frictional costs and assumes that a portfolio manager can simply readjust their holdings dynamically without any problems. In practice, trading costs are non-zero and affect the decision to rebalance. The transactions costs involve the actual cost paid for the trades as well as the cost of manpower and technological resources. If the transactions costs exceed the expected benefit from rebalancing, then no adjustment should be made. However, without any quantitative measure for this benefit, we cannot accurately determine whether or not to trade.

Conventional approaches to portfolio rebalancing include periodic and tolerance band rebalancing [1], [2]. With periodic rebalancing, the portfolio manager adjusts the current weights back to the target weights at a consistent time interval (*e.g.*, monthly or quarterly). The drawback with this method is that trading decisions are independent of market behavior. So rebalancing may occur even if the portfolio is nearly optimal. Tolerance band rebalancing requires managers to rebalance whenever any asset class deviates beyond some pre-determined tolerance band (*e.g.*, $\pm 5\%$). When this occurs, the manager fully rebalances to the target portfolio. While this method reacts to market movements, the threshold for rebalancing is fixed, and the process of rebalancing involves trading all the way back to the optimal portfolio.

Previous research on dynamic strategies for asset allocation [3] has established the existence of a no-trade region around the optimal target portfolio weights [4]. If the proportions allocated to each asset at any given time lie within this region, trading is not necessary. However, if current asset ratios lie outside the no-trade region, Leland has shown that it is optimal to trade but only to bring the weights back to the nearest edge of the no-trade region and not all the way to the target ratios. The optimal strategy has been shown to reduce transaction costs by approximately 50%. However, the full analytical solution involves a complicated system of partial differential equations in multiple dimensions.

Mulvey and Simsek [5] have modeled the problem of rebalancing in the face of transaction costs as a generalized network with side conditions and developed an algorithm for solving the resulting problem. Meanwhile, Mitchell and Braun [6] have described a method for finding an optimal portfolio when proportional transactions costs have to be paid. More recently, Donohue and Yip [1] have confirmed the results of Leland [4] and have characterized the shape and size of the no-trade region and compared the performance of different rebalancing strategies.

In this paper, we present an approach that explicitly weighs transaction costs and tracking error costs. We define a cost for a suboptimal portfolio using certainty equivalents [7], and we use dynamic programming to develop a policy that trades only when the cost of trading is less than the cost of doing nothing. In addition, given a decision to rebalance, we do not enforce a constraint of rebalancing to the optimal portfolio. The importance of relaxing this constraint is that in many situations, the cost of fully rebalancing is more than the benefit obtained. We show that our method performs better than traditional methods of rebalancing and is robust to model error. Note that we assume that the portfolios are either tax-free or tax-deferred, which is the case for endowments, charities, pension funds, and most individual retirement funds.

In Section II, we discuss utility functions and the method of dynamic programming. We introduce certainty equivalents and discuss how we use them to track costs and determine our optimal rebalancing strategy in Section III. We then demonstrate the rebalancing problem on a simple two-asset example in Section IV to illustrate our algorithm and provide some simple sensitivity analyses. Section V examines the more general case of mean-variance optimization on multiple assets over long periods of time. We conclude the paper in Section VI.

II. Background

A. Utility Functions

Evaluating individual preferences to risk and return and making corresponding portfolio allocation decisions is a difficult task. Investment professionals need to analyze various factors to develop portfolios that allow clients to reach their investment goals while taking into account risks associated with bear markets and singular events such as crashes.

No single portfolio can meet the needs of every investor. As mentioned in the introduction, one way to specify an investor's risk preference is through the use of utility functions [8]. A utility function indicates how much satisfaction (utils) we get for a given level of return x. Clearly, most people prefer higher levels of return to lower levels of return, so generally utility functions monotonically increase with x. If the marginal utility decreases with x (*i.e.,* utility grows sublinearly), then an individual is said to be *risk*

	Utility Function	Expected Utility
Quadratic	$f_{\mathrm{q}}(x) = x - \frac{\alpha}{2}(x - x_0)^2$	$U_{\mathrm{q}}(\mu, \sigma) = \mu - \frac{\alpha}{2}\sigma^2$
Log wealth	$f_{\mathrm{l}}(x) = \log(1 + x)$	$U_{\mathrm{l}}(\mu, \sigma) = \log(1 + \mu) - \frac{\sigma^2}{2(1+\mu)^2}$
Power	$f_{\mathrm{p}}(x) = 1 - 1/(1 + x)$	$U_{\mathrm{p}}(\mu, \sigma) = 1 - \frac{1}{(1+\mu)} - \frac{\sigma^2}{(1+\mu)^3}$

TABLE I

UTILITY FUNCTIONS AND THEIR CORRESPONDING APPROXIMATE EXPECTED UTILITIES USED. THE UTILITY FUNCTIONS f_i ARE EXPRESSED IN TERMS OF THE RETURN x. THE EXPECTED UTILITY FUNCTIONS U_i ARE SPECIFIED IN TERMS OF THE MEAN RETURN μ AND THE STANDARD DEVIATION σ. FOR QUADRATIC UTILITY, α IS THE RISK AVERSION PARAMETER.

averse. Numerous other risk characteristics can be imparted through the shape of the utility function. A manager then chooses portfolio weights w so that expected utility is maximized.

Because future returns are unknown, we need to use expected utility to create an optimal portfolio or to decide on a rebalancing policy. It has been shown by Levy and Markowitz [9] that for most relevant utility functions, this expected utility U can be approximated using truncated Taylor series expansions to be a function of mean and standard deviation, $\mathrm{U}(\mu, \sigma)$.

In Table I, we list three utility functions and the corresponding expected utilities that we use [10]. For each utility, $f_i(x)$ for $i = \{\mathrm{q}, \mathrm{l}, \mathrm{p}\}$ (where q indicates quadratic, l indicates logarithmic, and p indicates power) represents the utility in utils given a return x, which we also refer to as the empirical utility. $\mathrm{U}_i(\mu, \sigma)$ for $i = \{\mathrm{q}, \mathrm{l}, \mathrm{p}\}$ is the expected utility. Figure 1 plots the three empirical utility functions as a function of return. The relative difference, not the absolute value, in utility for different x is what is important (we could scale the utility functions without affecting the corresponding optimal portfolio).

Quadratic utility is a commonly used function, and using it in portfolio construction is akin to doing standard mean-variance optimization. Regardless of whether or not the assets are Gaussian-distributed, the expected utility only involves the first two moments, so any higher order moments are ignored. The α parameter can be adjusted to indicate risk tolerance. A larger number indicates that an investor is more risk averse. One main difficulty with quadratic utility is that it has the odd behavior that for a large enough return, it is *too* risk averse and the utility function actually prefers a smaller return because $\lim_{x \to \infty} f_q(x) = -\infty$. This behavior begins at $x = x_0 + \frac{1}{\alpha}$, the maximum of the quadratic function. The log wealth and power utility functions do not exhibit this behavior.

Even though it is true that the expected utility can be written just in terms of the mean and variance,

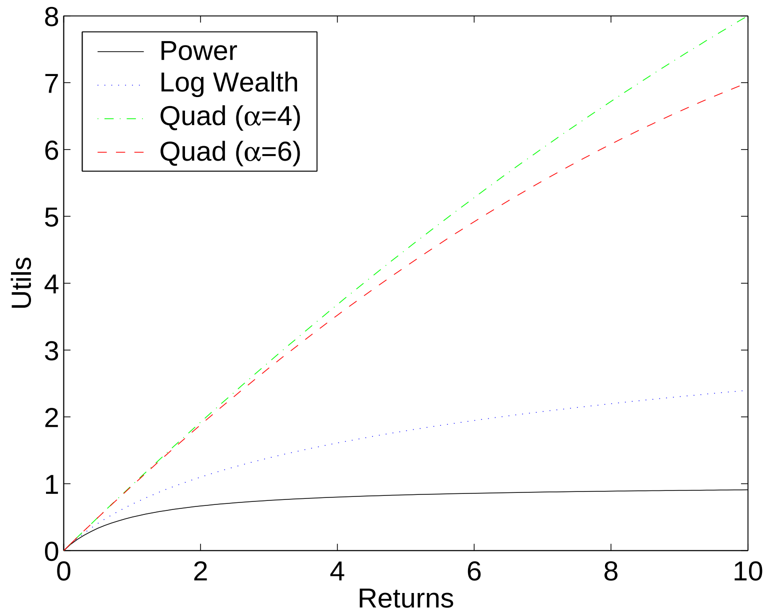

Fig. 1. Plots of the quadratic utility (for two different α's), log wealth utility, and power utility as a function of returns.

the expression shown for $U_q(\mu, \sigma)$ is only an approximation. The true expression is:

$$\mathrm{U}_q(\mu, \sigma) = \mu - \frac{\alpha}{2}(\sigma^2 + (\mu - x_0)^2). \tag{1}$$

Note that if we had *a priori* knowledge of the portfolio return μ, we would just choose $x_0 = \mu$. Unfortunately, μ is a function of the portfolio weights w, so we cannot fix it ahead of time. In the typical operating regime, $\mu(w) \approx \mu(w^*)$ because we would typically rebalance before the portfolios become too unbalanced. So if we choose $x_0 = \mu(w^*)$, then the $(\mu - x_0)^2$ term in (1) is small, and U_q as a reasonable approximation to the true expected utility. This leaves a much simpler expected utility function (especially in terms of μ).

The derivation of the expected utility functions for log wealth and power is non-obvious. Let's examine log wealth utility. We can expand the utility function around the point $x = \mu$ using a Taylor series:

$$\begin{aligned} f_l(x) = \log(1 + x) &= \log(1 + \mu) + \frac{1}{1!} f_l'(1 + \mu)(x - \mu) + \frac{1}{2!} f_l''(1 + \mu)(x - \mu)^2 + \cdots \\ &\approx \log(1 + \mu) + \frac{x - \mu}{1 + \mu} - \frac{(x - \mu)^2}{2(1 + \mu)^2}. \end{aligned} \tag{2}$$

Thus we see that

$$\begin{aligned} U_l(\mu,\sigma) &= E[\log(1+x)] \\ &\approx E\left[\log(1+\mu) + \frac{x-\mu}{1+\mu} - \frac{(x-\mu)^2}{2(1+\mu)^2}\right] \\ &= \log(1+\mu) - \frac{\sigma^2}{2(1+\mu)^2}. \end{aligned}$$

Additional terms of the Taylor expansion may be used to improve the approximation. These will then involve the skewness and the kurtosis and other higher-order moments. A similar method is applied to derive the approximation for power utility.

B. Dynamic Programming

Dynamic programming [11], [12], [13] is an optimization technique that finds the policy that minimizes expected cost given a cost functional and a dynamic model of state behavior. At time t, w_t is our state, u_t is our policy, and n_t is the state uncertainty. The state transition is defined by an arbitrary function h:

$$w_{t+1} = \mathrm{h}(w_t, u_t, n_t), \tag{3}$$

where w_{t+1} represents the new state which is influenced by the prior state w_t, the action taken u_t, and the uncertainty in the system dynamics n_t. We write the cost functional recursively as:

$$\mathrm{J}_t(w_t) = \mathrm{E}\left[\mathrm{G}(w_t, u_t, n_t) + \mathrm{J}_{t+1}(w_{t+1})\right], \tag{4}$$

where G is the cost for the current period and J_t is the so-called *cost-to-go* function. J_t is the expected future cost from t onwards given all future decisions. So, the cost at any given period is the expected cost from t to $t+1$ along with the expected cost from $t+1$ onwards. At each time t, the optimal strategy is to choose u_t such that the cost is minimized:

$$\mathrm{J}^*_t(w_t) = \min_{u_t} \quad \mathrm{E}\left[\mathrm{G}(w_t, u_t, n_t) + \mathrm{J}_{t+1}(w_{t+1})\right]. \tag{5}$$

Equation (5) is the discrete-time Bellman Equation. Assuming convergence, this recursion approaches a fixed point such that $\mathrm{J}^*_t(w) = \mathrm{J}^*_{t+1}(w) = \mathrm{J}^*(w)$. The challenge is therefore to determine the cost-to-go values $\mathrm{J}^*(w)$. Once these values are known, the optimal rebalancing decision is to choose the policy u^*_t that minimizes (5).

We can determine the cost-to-go values using a technique called *value iteration*. The idea behind value iteration is to choose an arbitrary set of cost-to-go values $\mathrm{J}_t(w)$ for some time t that we imagine to be very far in the future. We then repeatedly apply (5) to obtain cost-to-go values successively closer to the present. After a sufficient number of iterations, we will approach a steady-state, and the cost-to-go values should converge to the optimal values $\mathrm{J}^*(w)$.

III. Optimal Rebalancing Using Dynamic Programming

In this section, we investigate optimal rebalancing strategies for portfolios with transaction costs. In general, we consider a multi-asset problem where we are given an optimal portfolio consisting of a set of target portfolio weights $w^* = \{w_1^*, \ldots, w_N^*\}$, where N is the total number of assets. The optimal strategy should be to maintain a portfolio that tracks the optimal portfolio as closely as possible while minimizing the transaction costs.

We consider a model where we observe the contents of the portfolio w_t at the end of each month. At that point, we have the option of rebalancing the portfolio (i.e. apply our policy, or control, u_t). Thus, the portfolio at the beginning of the next month is $w_t + u_t$. Then we assume normal returns in the process noise n_t. We use a simple multiplicative dynamic model so that $w_{t+1} = (1 + n_t)(w_t + u_t)$, although in general, w_{t+1} can be an arbitrary function of w_t, u_t, and n_t.

In general, the decision to rebalance should be based on a consideration of three costs: the tracking error associated with any deviation in our portfolio from the optimal portfolio, the trading costs associated with buying or selling any assets during rebalancing, and the expected future cost from next month onwards given our actions in the current month. The optimal strategy dynamically minimizes the total cost, which is the sum of these three costs.

To apply dynamic programming, we must specify the cost function in the Bellman Equation. In our case, we write:

$$\mathrm{E}\left[\mathrm{G}(w_t, u_t, n_t)\right] = \tau(u_t) + \epsilon(w_t + u_t), \tag{6}$$

where $\tau(u_t)$ is the trading cost associated with applying our rebalancing decision u_t. This can include tangible costs such as commissions and market impact, but can also model indirect costs such as employee labor. $\epsilon(\cdot)$ represents the suboptimality cost, the cost of not having an optimal portfolio. $\epsilon(w_t + u_t) = 0$ whenever $w_t + u_t = w^*$ (i.e. choose u_t so that we rebalance to the target portfolio); otherwise, $\epsilon(\cdot) > 0$.

A. Modeling Tracking Error using Certainty Equivalents

Note that the cost-to-go values, and hence the optimal strategy, will depend on the cost functions $\tau(\cdot)$ and $\epsilon(\cdot)$ chosen. In this section, we discuss strategies for modeling tracking error.

In the certainty equivalence approach, we model the investor's preferences using a utility function (see Section II-A). For any portfolio weights w, we can express the expected utility as $\mathrm{U}(\mu^T w, w^T \Lambda w)$. We observe that there exists a risk free rate (which we will denote as $r_{CE}(w)$) that produces an identical expected utility. We therefore call $r_{CE}(w)$ the *certainty equivalent return* for the weights w. The condition for this is $\mathrm{U}(r_{CE}, 0) = \mathrm{U}(\mu^T w, w^T \Lambda w)$. The certainty equivalents for the three expected utility functions that we are using are:

1) Quadratic: $r_{CE}(w) = \mathrm{U_q}(\mu^T w, w^T \Lambda w)$
2) Log wealth: $r_{CE}(w) = \exp(\mathrm{U_l}(\mu^T w, w^T \Lambda w)) - 1$
3) Power: $r_{CE}(w) = 1/(1 - \mathrm{U_p}(\mu^T w, w^T \Lambda w)) - 1$.

One interpretation of the certainty equivalent then is as a risk-adjusted rate of return given the risk preferences embedded in the utility function.

If we hold a suboptimal portfolio w, the utility of that portfolio $\mathrm{U}(w)$ will be lower than $\mathrm{U}(w^*)$, with a correspondingly lower certainty equivalent return. We can interpret this as losing a riskless return (equal to the difference between the two certainty equivalents) over one period, corresponding to the penalty paid for tracking error. Therefore, under the certainty equivalence approach, the tracking error has the cost function

$$\epsilon(w) = r_{CE}(w^*) - r_{CE}(w). \tag{7}$$

The reason why we use a certainty equivalent is because in our cost function, $\tau(\cdot)$ and $\epsilon(\cdot)$ must have commensurate values. We know that the cost will be in terms of dollars or basis points or some other absolute measure. It is more straightforward to then convert portfolio tracking error into a similar absolute measure using certainty equivalents rather than trying to express the trading costs in terms of diminished expected utility.

B. Modeling Transaction Costs

Assume that we have a portfolio w and we want to go to another portfolio w'. The simplest model for transaction costs is simply to assume a linear cost. Under this model, we assume that for asset i we pay a transaction cost of c_i per dollar to buy or sell the asset. Under this model,

$$\tau(w', w) = c^T |w' - w|, \tag{8}$$

where $c^T = [c_1, \dots, c_N]$ is the vector of transaction cost coefficients. A variant of the linear cost model,

$$C(w', w) = c_+^T \max\{w' - w, 0\} + c_-^T \max\{w - w', 0\}, \tag{9}$$

allows for different costs to buy (c_+) and sell (c_-) assets.

Building on the linear cost model, we can also allow for fixed costs in rebalancing as well. This model discourages frequent rebalancing and can be used to model the significant dislocation costs associated with transferring assets from one manager to another. This model can be written as

$$C(w', w) = c_{+}^{T} \max\{w' - w, 0\} + c_{-}^{T} \max\{w - w', 0\} + c_{f}^{T} I(w - w'), \tag{10}$$

where c_f is the vector of fixed costs associated with trading each asset, and the indicator function $I(x)$ is 1 if $x > 0$ and 0 otherwise.

IV. Two-Asset Model

To introduce the problem of portfolio rebalancing, we first consider an example involving two risky asset classes. The benefits of the two risky asset model are that the optimal portfolio can be computed in closed form (see [14] for a derivation), we can visually examine the changes in portfolio weights (since a single asset's weight represents the full description of our portfolio), and the parameters are few enough that we can easily perform sensitivity analyses. We follow this example with extension simulations of a multi-asset model.

We assume that we can invest in (1) US Equity or (2) Private Equity. We assume that returns are normal[1], so the mean and covariance statistics sufficiently characterize the assets. We obtain 6.84% and 12.76% expected annual returns for US Equity and Private Equity, respectively, based on historical observations, and acquire the covariances from Terhaar *et al.* [16]. Specifically, they cite that US Equity has standard deviation of 12.80% annually, while Private Equity has standard deviation of 21.00%. The correlation coefficient between the two assets is -0.46. Given this information, we create the efficient frontier (Figure 2).

A. Simulation

For brevity, we only consider quadratic utility in this section. Suppose the risk aversion parameter α is 0.33. Using this assumption, the optimal portfolio balance is 41.00% in US Equities and 59.00% in Private Equity. To provide an example of our rebalancing method, we simulate the returns of the two equities over a ten year period, assuming normal distribution of returns with means and covariance as described earlier. Computationally, we obtain a random sample from the normal distribution for each month of our simulation.

[1]Because we consider monthly data, a normal assumption is reasonable. For longer time periods, lognormality would be a better assumption.

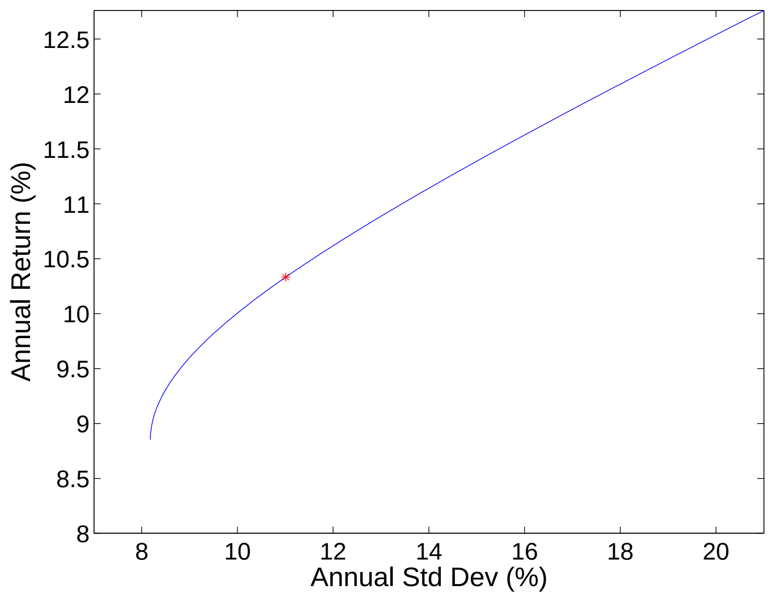

Fig. 2. Efficient frontier of two risky asset case with the optimal portfolio for quadratic utility ($\alpha = .33$) labeled.

Figure 3 shows how the portfolio weight of US Equities moves over one 120 month sample path. With no rebalancing (Figure 3(a)), the weight drifts from the optimal amount of 41.00% down to under 20%, resulting in large suboptimality costs (the exact costs are described below). Our optimal rebalancing strategy (Figure 3(b)) rebalances only when necessary. During months 40 to 45 and 90 to 110, the portfolio rebalances nearly every month to handle sharp changes in the portfolio, while for months 45 to 80, the lack of strong market movements in either direction allow us to avoid any transaction costs. The market movement during the times cited can be seen by examining the change in portfolio weights in Figure 3(a) where there is no rebalancing.

Table II shows the costs of trading using different strategies. Trading costs are assumed to be 20 bps for buying or selling public equity, and 40 bps for buying or selling private equity. The suboptimality cost was determined using the idea of certainty equivalents. For each portfolio, a certainty equivalent can be computed (in terms of monthly returns). The difference between the certainty equivalent of a non-optimal portfolio and that of the optimal portfolio is defined as the cost of not being optimal.

From the table, we observe that the aggregate monthly cost is minimized by our method. Over a ten year period, the cost of our algorithm, assuming $100 million invested, is $281,300. The next best method for this example, that of yearly rebalancing, costs $337,100. The results for each rebalancing method make intuitive sense. Monthly rebalancing leads to no deviation from optimality, but at the cost

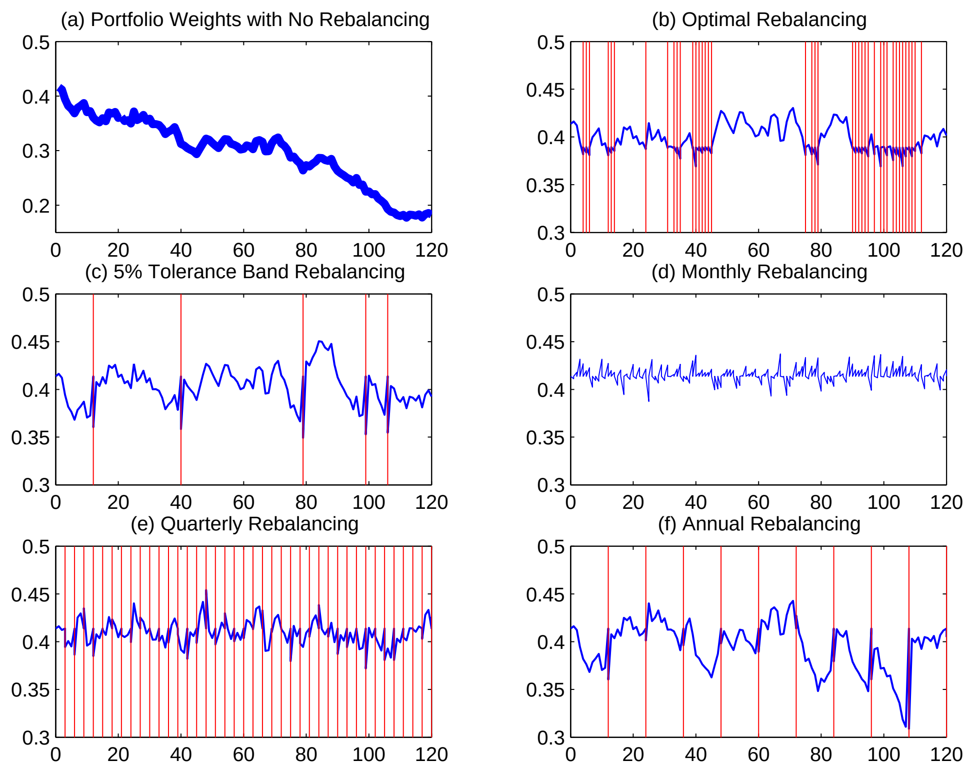

Fig. 3. Plots of US Equities weighting in the two asset example using different rebalancing models. The vertical lines indicate months where rebalancing was done (for monthly rebalancing, this is omitted since trading occurs in every month).

	Trading Cost (bps)	Suboptimality Utility Cost (bps)	Aggregate Cost (bps)
Optimal DP	18.14	9.99	28.13
No Trading	0.00	1509.74	1509.74
5% Tolerance	17.59	16.83	34.43
Monthly	62.61	0.00	62.61
Quarterly	37.78	1.74	39.52
Annual	20.29	13.42	33.71

TABLE II

TRADING COST, SUBOPTIMALITY COST, AND AGGREGATE COST USING SIX DIFFERENT REBALANCING STRATEGIES ON TWO RISKY ASSETS OVER A TEN YEAR PERIOD.

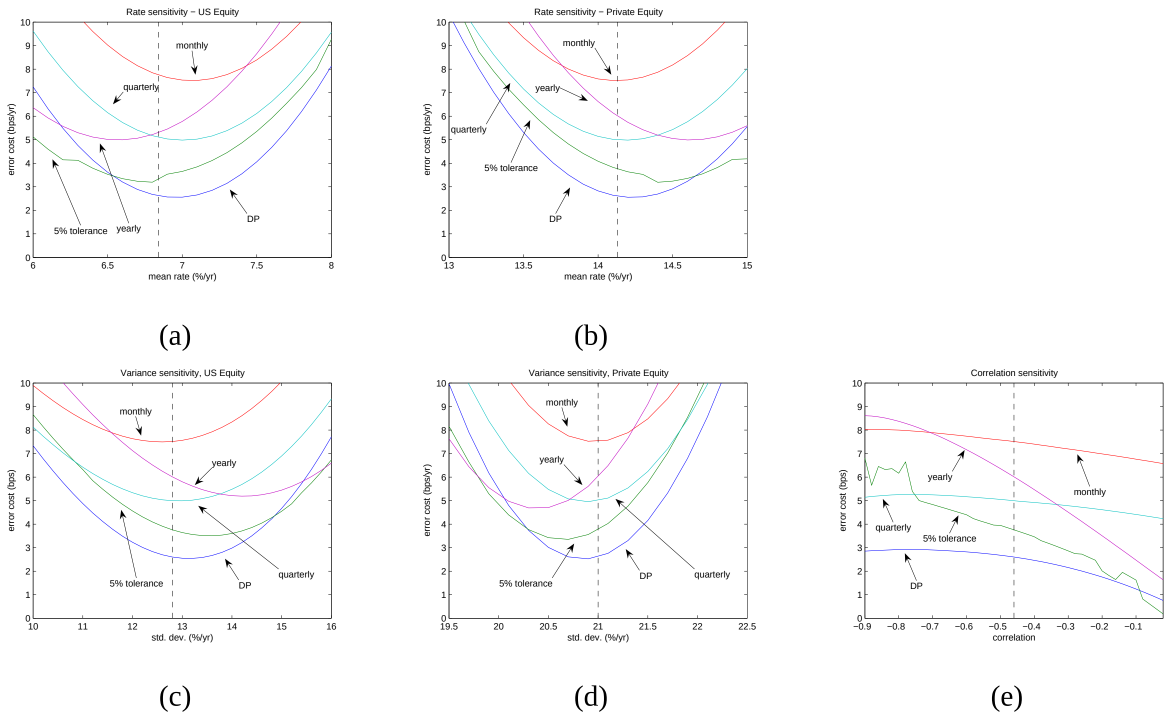

Fig. 4. Sensitivity analysis on mean, standard deviation, and correlation for two asset example. The vertical line indicates the rate used by the rebalancing strategies. (a) Sensitivity to mean rate for US Equity. (b) Sensitivity to mean rate for Private Equity. (c) Sensitivity to variance for US Equity. (d) Sensitivity to variance for Private Equity. (e) Sensitivity to correlation between US and Private Equity.

of high trading fees. Infrequent trading yields smaller trading costs, but higher non-optimality certainty equivalent costs. Our method of rebalancing whenever the cost of non-optimality exceeds the trading costs allows us to adequately trade-off the cost of non-optimality with that of trading.

B. *Sensitivity Analysis*

So far, we have assumed that the model for each asset is accurate. In practice, this is usually not the case – mean and variance of the returns of each asset as well as the correlation between assets must be estimated (e.g. using historical observations), and there is usually some error associated with each estimate. Errors in the parameter estimate will cause inaccuracies in the cost-to-go values obtained from the dynamic program, leading to suboptimal rebalancing. In this section, we investigate the impact of errors in each of these parameters on the rebalancing strategy.

We investigate a total of three parameters – mean, variance, and correlation. In each simulation, two parameters are held constant while the third is allowed to vary around the estimated value. The costs-to-go are used to calculate the estimated values. We characterize the performance of the strategy when the estimated parameters differ from the actual parameters. For calendar and tolerance band strategies, we assume that they would also rebalance to an optimal portfolio calculated from the estimated parameters. We expect the performance of all strategies to degrade when the parameter estimate is inaccurate. However, the main issue is whether some strategies are relatively more robust to inaccuracies than others. We expect our approach to be most sensitive to model error because the model parameters are used to both generate the target portfolio and the rebalancing strategy. The other methods only rely on the model for the target portfolio.

First, we assume that the variance and correlation are correctly estimated and investigate estimation errors in the mean. For each point, 100 sequences of ten year monthly returns were generated, and the performance of the dynamic rebalancing strategy was averaged over each sequence. The dynamic rebalancing strategy was reasonably insensitive to errors in estimating the mean: from Figure 4(a) and (b), we can observe that the DP approach outperforms all other strategies over a range of several percentage points of inaccuracies in the estimation. We can conclude that as long as the mean can be accurately estimated to within a few percentage points, the dynamic programming-based approach is still the best choice. From Figure 4(c) and (d), we see that the dynamic programming approach again outperforms the other approaches even if there are large errors in estimating the standard deviation – it remained the best performer even given inaccuracies in the standard deviation of several percentage points per year. Finally, in Figure 4(e), we observe that the dynamic programming approach is insensitive to errors in estimation of the correlations between assets – the approach outperforms all others in a wide range of correlations. This suggests that correlations do not need to be accurately estimated for the purposes of the DP.

V. Multi-Asset Model

Now that we have described and analyzed the simple two-asset model, we proceed to examine the general case of N risky assets. Unlike the two-asset scenario, the optimal portfolio cannot be computed in closed-form for any $N > 2$. In this section, we consider the case of five risky assets and assert that another choice of $N > 2$ would proceed similarly with the main difference being computation time. For our analysis, we concentrate on generating optimal portfolios with five asset classes: US Equity, Developed Market Equity, Emerging Market Equity, Private Equity, and Hedge Funds. In Table III, we show the mean and standard deviation of each asset class along with higher order statistics such as skewness and kurtosis for historical monthly returns from January 1990 to March 2004. Normally distributed data have zero skewness and a kurtosis of 3. Most of the assets exhibit approximately normal returns with the exception of Hedge Funds, which has high kurtosis indicating a heavy-tailed distribution.

	Index as Proxy (Source)	Mean Return (%)	Std. Dev. (%)	Skewness (normal = 0)	Kurtosis (normal = 3)
US Equity	Russell 3000 (Datastream)	6.84	14.99	-0.57	3.67
Developed Mkt Equity	MSCI EAFE+Canada (Datastream)	6.65	16.76	-0.20	3.27
Emerging Mkt Equity	MSCI EM (Datastream)	7.88	23.30	-0.73	4.71
Private Equity	Wilshire LBO (Bloomberg)	12.76	44.39	-0.40	3.82
Hedge Funds	HFR Mkt Neutral (Bloomberg)	5.28	10.16	-0.83	7.04

TABLE III

ANNUAL MEAN RETURNS, ANNUAL STANDARD DEVIATIONS, SKEWNESS, AND KURTOSIS FOR THE ASSET CLASSES.

	US Equity	Developed Markets	Emerging Markets	Private Equity	Hedge Fund
US Equity	1.00	0.46	0.45	0.64	0.29
Developed Markets	0.46	1.00	0.42	0.38	0.09
Emerging Markets	0.45	0.42	1.00	0.40	0.21
Private Equity	0.64	0.38	0.40	1.00	0.36
Hedge Fund	0.29	0.09	0.21	0.36	1.00

TABLE IV

CORRELATION COEFFICIENT MATRIX.

The correlation matrix used is shown in Table IV. Of the different assets, Private Equity provides the most expected return, but has the greatest amount of risk. On the other extreme, Hedge Funds have both the least expected return and the least amount of variability. The mean returns were provided by State Street Associates and the variances and correlations were computed empirically from data acquired from Datastream and Bloomberg.

It is known [10] that standard mean-variance portfolio optimization produces optimal portfolios only if returns are normally distributed or if quadratic utility is assumed. Otherwise, full-scale optimization must be performed to compute optimal portfolios when using more advanced utility functions such as log wealth or power utility. However, recent work by Cremers *et al.* [17] indicates that except when returns are highly non-normal, it is sufficient to perform mean-variance optimization on a Markowitz-style approximate expected utility function (see Section II-A) in terms of just the mean and standard deviation. They show

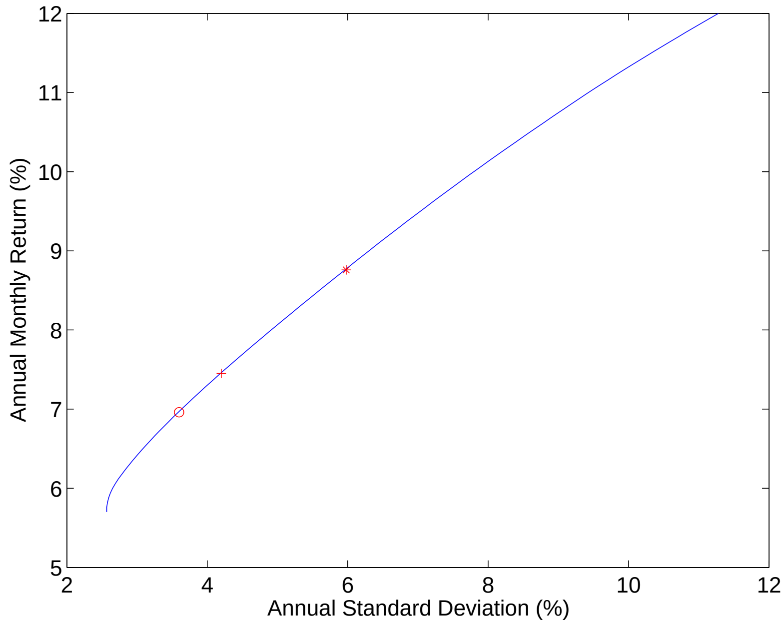

Fig. 5. Efficient frontier and optimal portfolios for the different utility functions discussed. (o) indicates the optimal portfolio for power utility. (+) indicates that for quadratic utility with $\alpha = 1.5$. (*) indicates the optimal portfolio for log wealth utility.

that the performance of the resulting portfolios and the performance of those generated from full-scale optimization do not differ significantly. When performing this approximate mean-variance optimization, the optimal portfolio lies on the efficient frontier[2] [15]. Therefore, to construct optimal portfolios for different utility functions, we first compute the efficient frontier by solving a quadratic programming problem and then search over those portfolios to find the one with the highest expected utility.

Figure 5 displays the efficient frontier for the five asset classes when short sales are not allowed. Searching over this frontier for each of the utility functions results in the optimal portfolios as indicated in the figure with the weights shown in Table V. These weights are the optimal weights we use throughout our analysis.

A. Results

Table VI show the results of our dynamic programming algorithm and some existing rebalancing methods on Monte Carlo simulations. We generated 10,000 sample paths, each for ten years of monthly

[2]Risk-averse expected utility functions are monotonically increasing in terms of return and monotonically decreasing in terms of risk. Hence if a portfolio is not on the efficient frontier, there exists a portfolio with equivalent return and less risk or more return and the same risk. Therefore this portfolio cannot be optimal.

	Quadratic (α=1.5)	Logarithmic	Power
US Equity	0.194	0.160	0.210
Developed Market Equity	0.222	0.240	0.213
Emerging Market Equity	0.185	0.275	0.143
Private Equity	0.156	0.292	0.093
Hedge Funds	0.243	0.033	0.341

TABLE V

OPTIMAL PORTFOLIO WEIGHTS FOR DIFFERENT UTILITY FUNCTIONS $U(\mu, \sigma)$.

return data. Each month is sampled independently from the others, so we do not model effects such as trends, momentum, or mean reversion. For each sample path, we simulate the various rebalancing methods by generating a return value for each month that is net of transaction costs.

We measure performance as a shortfall relative to an idealized rebalancing strategy which rebalances to the optimal portfolio every month for free. All real world strategies will suffer shortfalls due to trading cost, suboptimal portfolios, or both. We measure this shortfall using the actual trading cost that we incur (column (a)) and the decrease in certainty equivalent from the optimal portfolio (column (b)). These combine to give the aggregate shortfall in column (c).

Note that this aggregate shortfall is exactly what our dynamic programming algorithm is trying to minimize. All of the real-world algorithms are really trying (implicitly or explicitly) to minimize this aggregate cost by trying to balance trading cost and suboptimality cost. At one extreme is monthly rebalancing which has zero suboptimality cost but requires a lot of trading to implement. At the other extreme is no rebalancing which does not cost anything in trading cost to implement, but a severe penalty is paid in terms of risk adjusted return. All of the other methods fall somewhere in between.

The aggregate cost is perhaps a bit unsatisfying as a metric because the risk-adjusted return component is based on *expected* month-to-month returns (certainty equivalents) rather than the *actual* returns observed in the Monte Carlo simulation. Another approach to evaluate the results is to then use the actual sample returns. We can use the sample returns to compute the empirical utility at each time period (using the empirical utility functions in Table I), and then take the sample average of this to obtain an average utility shortfall (column (d)). Of course, since the data are ergodic[3], we expect the sample average shortfall and

[3]Ergodicity certainly holds in our Monte Carlo simulations because we sample independently.

		(a)	(b)	(c)	(d)	(e)	(f)
		Trading Cost (bps)	Suboptimality Cost (bps)	Aggregate Cost (bps)	Utility Shortfall (utils x 10^4)	Net Returns (%)	Standard Deviation (%)
Quadratic $\alpha = 1.5$	Ideal	0.00	0.00	0.00	0.00	7.45	14.84
	Optimal DP	4.04	1.72	5.75	5.55	7.40	14.86
	No Trading	0.00	71.72	71.72	71.36	6.77	14.96
	5% Tolerance	7.39	0.70	8.09	8.03	7.37	14.83
	Monthly	23.66	0.00	23.66	23.72	7.22	14.84
	Quarterly	13.68	0.28	13.96	14.28	7.32	14.85
	Annual	6.84	1.55	8.39	8.24	7.40	14.94
Power	Ideal	0.00	0.00	0.00	0.00	6.89	12.38
	Optimal DP	3.47	1.21	4.67	4.43	6.87	12.48
	No Trading	0.00	81.70	81.70	82.31	6.77	14.95
	5% Tolerance	5.30	0.83	6.13	5.75	6.83	12.36
	Monthly	20.05	0.00	20.05	19.96	6.69	12.38
	Quarterly	11.59	0.18	11.78	11.90	6.77	12.39
	Annual	5.82	1.02	6.84	6.64	6.84	12.46
Log Wealth	Ideal	0.00	0.00	0.00	0.00	8.65	20.57
	Optimal DP	4.87	2.26	7.13	7.09	8.57	20.49
	No Trading	0.00	91.51	91.51	87.82	6.77	14.98
	5% Tolerance	11.99	0.44	12.43	12.74	8.53	20.60
	Monthly	28.14	0.00	28.14	28.18	8.37	20.58
	Quarterly	16.25	0.40	16.65	17.13	8.49	20.59
	Annual	8.06	2.17	10.22	10.18	8.57	20.67

TABLE VI

QUADRATIC ($\alpha = 1.5$), POWER, AND LOG WEALTH UTILITY: ANNUALIZED TRADING COST, NON-OPTIMAL UTILITY COST, AND AGGREGATE COST USING SIX DIFFERENT TRADING STRATEGIES ON FIVE RISKY ASSETS SIMULATED OVER A TEN YEAR PERIOD 10,000 TIMES. THE UNITS ON UTILS MULTIPLIED BY 10^4 (IN COLUMN (D)) ARE SIMILAR TO BASIS POINTS IN COLUMNS (A)-(C). THIS IS OBVIOUS FOR THE QUADRATIC CASE WHERE THE CERTAINTY EQUIVALENT IS EQUAL TO THE UTILITY. FOR THE OTHER TWO CASES, TAKING A LINEAR APPROXIMATION AROUND $x = 0$ SHOWS THAT THE UTILITIES ARE PROPORTIONAL TO x. SO, UTILS TIMES 10^4 PROVIDES A REASONABLE APPROXIMATION TO BASIS POINTS.

the expected shortfall to produce similar results.

In columns (e) and (f) we list the sample means and standard deviations of the strategies. These are primarily illustrative to show the tradeoffs that different rebalancing strategies make. Most of the strategies fare a bit worse than the idealized rebalancing strategy on both net returns and risk.

When comparing the results from the different utility functions, recall that each utility function has a different risk level which in turn induces a different optimal portfolio. This is why the net returns and standard deviations (columns (e) and (f)) can vary so much from utility function to utility function.

B. Quadratic Utility

For the quadratic utility case shown in Table VI, we see that our optimal DP method performs 29% better in terms of expected cost and 31% better in terms of average utility over the next-best method, 5% tolerance bands. If we examine the costs, we see, as expected, that monthly trading incurs no suboptimality at the expense of high trading costs. The other extreme of no trading incurs an extremely large suboptimality cost because over a ten year period, assets can become quite unbalanced if unadjusted. It also should be noted that our method can be thought of as a dynamic tolerance band approach. Thus, since the 5% tolerance method is a subset of our algorithm, it can never do better.

C. Power Utility

For power utility, the results are shown in Table VI. As with quadratic utility, our expected loss is 24% less than the runner-up, 5% tolerance band rebalancing. The sample-based empirical utility shortfall is reduced by 22%. The benefits for this method are reduced from the quadratic utility case primarily because less rebalancing is needed overall because the power utility portfolio has the lowest variance.

Note that even though tolerance bands do better than annual rebalancing in this example (and also for quadratic utility, but not for log wealth), this should not necessarily be taken as an indicator that tolerance bands are a superior method to periodic rebalancing. Better performance can be obtained by tweaking the threshold parameter or the periodicity of rebalancing. For instance, setting the rebalancing time to two years for the power utility case results in an expected loss of 6.32 bps per annum, a savings of 1.74 bps over the annual strategy. This is achieved by accruing more than twice as much expected suboptimal risk-adjusted return (2.21 bps versus 1.03 bps), but also reducing trading costs by 29% (4.11 bps versus 5.81 bps). A more exhaustive search of possible fixed-interval rebalancing strategies could presumably yield an even better result.

D. Log Wealth Utility

For the log wealth utility case, the results are shown in Table VI. Again, our expected loss is 30% less than the best alternative (annual rebalancing). And the average simulated utility deficit is also 30% less than annual rebalancing. This is a clear win as we tie for the highest net return while we have the lowest standard deviation (except for the no rebalance case where in many cases, the high-variance/high-return assets become small quickly, and without rebalancing, we are stuck in low-variance/low-return assets). You can see the effect of the higher-variance portfolio in the trading cost numbers for the 5% Tolerance method. In the quadratic case, the trading costs are only marginally higher than the annual rebalance method. But in the log wealth case, they are 49% higher because the tolerance bands are breached more often. It's possible that better performance could be achieved by loosening the tolerance band as there is currently very little loss to portfolio suboptimality.

E. Computational Complexity

To provide some information regarding the computational complexity of our approach, we first state that we allow on the order of 15 possible weights for each asset. For five assets, we have an observation space of approximately 750,000 points (we must develop the optimal policy for each point). Our current implementation processes around 600,000 points per hour on a single PC (this problem can be easily parallelized; so, the total processing time also depends on the number of machines available). Thus, the run-time estimate for five assets is 75 minutes. If we assume the possibility of M different weights for an additional asset, the addition of this asset into our N-asset model would increase computation by a factor of M. Note that this is detailing the computation for learning the optimal policy. Once that is done, actually applying the policy is very fast.

F. Alternate Cost Functions

Before we complete this section, we address the possibility of a different trading cost function. In particular, while the numbers used are consistent with trading costs cited in other research papers [4], some may wonder if the results would be different for alternate trading costs. Table VII shows the results when we reduce the proportional trading costs in half and apply it to the quadratic utility strategy. We do only 20% better in expected cost, and 21% better in average utility, down from a 30% advantage with the original costs. Transaction costs for the other methods are cut in half, while suboptimality remains the same. Because in the original version transaction costs ranged from 82% of the aggregate cost for annual rebalancing to 100% of the cost for monthly rebalancing while they were only 70% for our method. If

	(a)	(b)	(c)	(d)	(e)	(f)
	Trading Cost (bps)	Suboptimality Cost (bps)	Aggregate Cost (bps)	Net Returns (%)	Standard Deviation (%)	Utility Shortfall (utils x 10^4)
Ideal	0.00	0.00	0.00	7.45	14.84	0.00
Optimal DP	2.64	0.87	3.51	7.42	14.85	3.42
No Trading	0.00	71.72	71.72	6.77	14.96	71.36
5% Tolerance	3.69	0.70	4.39	7.41	14.84	4.35
Monthly	11.83	0.00	11.83	7.34	14.84	11.86
Quarterly	6.84	0.28	7.12	7.38	14.85	7.44
Annual	3.42	1.55	4.97	7.43	14.94	4.83

TABLE VII

QUADRATIC UTILITY ($\alpha = 1.5$): TRADING COSTS, NON-OPTIMAL UTILITY COSTS, AND AGGREGATE COST USING SIX DIFFERENT TRADING STRATEGIES ON FIVE RISKY ASSETS SIMULATED OVER A 10 YEAR PERIOD $10{,}000$ TIMES. TRANSACTION COSTS ARE HALVED FROM THE PREVIOUS EXPERIMENTS.

our transaction costs were simply cut in half and we did not alter our trading strategy, we would expect the aggregate cost to decline by 35%. It actually declines by 39% because we adjust our strategy to trade more frequently and incur smaller suboptimality penalties.

VI. CONCLUSION

The *ad hoc* methods of periodic and tolerance band rebalancing provide simple but suboptimal ways to rebalance portfolios. Calendar-based approaches rely on the fact that, on average, we expect the portfolio to become less and less optimal as time goes on, but they do not use any knowledge about the actual state of the portfolio. The tolerance band approach does use the current portfolio to make a decision, but there is no sense of what the proper tolerance band setting is, or even how to choose it. In this work, we have shown that by formulating the rebalancing problem as an optimization problem and solving it using dynamic programming, we reduce the overall costs of portfolio rebalancing. We have demonstrated that the reduced costs hold for different investor risk preferences. Namely, we have compared the performance of our technique with others for three different utility functions: quadratic, log wealth, and power utility.

The costs of transactions are much more tangible than those for being suboptimal. However, through the use of certainty equivalents, we have provided a method that reasonably quantifies the cost of being suboptimal. Our simulations have confirmed that this optimal method provides gains over the best of the traditional techniques of rebalancing.

It is worth noting that in our analysis we assume returns over different intervals are independent. It has been discussed in the literature that mean reversion may exist. Under such circumstances, we expect our method to perform even better in comparison with periodic rebalancing because our algorithm would likely rebalance even less frequency.

Several extensions exist from our work. First, we may want to consider affine transaction costs. This model is appropriate if we believe that there is a fixed cost to making each and every transaction. Such an adjustment would likely favor dynamic trading methods over periodic rebalancing. Next, we may want to examine rebalancing over taxable portfolios. Asset managers of such funds have the additional consideration of tax consequences when a decision to transact needs to be made. The relaxation of the short sales constraint is another possible extension to the work. Although many tax-deferred funds do not allow short sales, several either explicitly do allow short selling or implicitly participate in short sales through investments into hedge funds.

In our work, we assume an instantaneous rebalancing at the end of each month. We may want to incorporate more general trading models which consider the effects of price impact. Finally, for the multi-asset case, we search a one-dimensional policy space representing portfolios which are a linear combination of the current portfolio and the target portfolio. We ideally want to search over the entire space of possible portfolios around the optimal portfolio. This would be particularly useful when trading costs have a fixed component. In these situations, it may be better to trade on only a subset of assets rather than a portion of all asset classes.

Acknowledgments

The authors would like to thank Sebastien Page (State Street Associates) and Mark Kritzman (Windham Capital Management) for introducing us to this problem and for their valuable guidance and advice. In addition, we note that this paper originated from a project performed for a course at the Massachusetts Institute of Technology's Sloan School of Management entitled *Proseminar in Financial Engineering*.

Appendix I
Efficient Frontier using Mean-Variance Optimization

Computing mean-variance efficient frontiers is a relatively straightforward process. This is an essential part of computing optimal portfolios for the non-normal returns or non-quadratic utility cases in order to avoid using full-scale optimization. We solve a series of quadratic programs [18], each minimizing the variance for a given expected portfolio return μ_p. Because we did not allow short sales, the optimization problem has the following form:

$$\begin{aligned} \min_w \quad & w'\Lambda w \\ \text{s.t.} \quad & w'\mu = \mu_p, \quad \textstyle\sum_i w_i = 1, \quad w \geq 0, \end{aligned} \tag{11}$$

where w are the unknown portfolio weights, Λ is the covariance matrix of the available assets and μ is the vector of expected asset returns. This optimization can be efficiently performed using Matlab's `Quadprog.m` function. For the quadratic utility function, it is not necessary to compute the entire efficient frontier. The optimal weights can directly be determined by solving a different quadratic program:

$$\begin{aligned} \max_w \quad & w'\mu - \tfrac{\alpha}{2} w'\Lambda w \\ \text{s.t.} \quad & \textstyle\sum_i w_i = 1, w \geq 0. \end{aligned} \tag{12}$$

REFERENCES

[1] C. Donohue and K. Yip, "Optimal Portfolio Rebalancing with Transaction Costs," *Journal of Portfolio Management*, (Summer 2003) 49-63.

[2] S. J. Masters, "Rebalancing," *Journal of Portfolio Management*, (Spring 2003) 52-57.

[3] A. F. Perold and W. F. Sharpe, "Dynamic Strategies for Asset Allocation," *Financial Analysts Journal*, (Jan-Feb 1995) 149-160.

[4] H. E. Leland, "Optimal Portfolio Management with Transaction Costs and Capital Gains Taxes," *Haas School of Business Technical Report,*, UC Berkeley, Dec 1999.

[5] J. M. Mulvey and K. D. Simsek, "Rebalancing Strategies for Long-term Investors," Kontoghiorghes, E.J., Rustem, B. and Siokos, S., eds., *Computational Methods in Decision-Making, Economics and Finance: Optimization Models*, Kluwer Academic Publishers, New York, (2002) 15-33.

[6] J. E. Mitchell and S. Braun, "Rebalancing an Investment Portfolio in the Presence of Transaction Costs," *Working Paper, Department of Rensselaer Polytechnic Institute*, Nov. 2002.

[7] D. Bernoulli, "Exposition of a New Theory on the Measurement of Risk," Translated from original Latin text written in 1738. *Econometrica*, 22:1 (Jan 1954) 23-36.

[8] R. Duncan Luce, *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches* Lawrence Erlbaum Associates, Inc., 2000.

[9] H. Levy and H. M. Markowitz, "Aproximating Expected Utility by a Function of Mean and Variance," *American Economic Review*, 69:3 (June 1979) 308-317.

[10] J. Cremers, M. Kritzman, S. Page, "Portfolio Formation with Higher Moments and Plausible Utility," *Revere Street Working Paper Series, Financial Economics 272-12*, November 22, 2003.

[11] R. Bellman, *Dynamic Programming,* Princeton University Press, Princeton, NJ, 1957.

[12] R. Bellman and S. Dreyfus, *Applied Dynamic Programming,* Princeton University Press, Princeton, NJ, 1962.

[13] D. P. Bertsekas, *Dynamic Programming and Optimal Control,* Athena Scientific, Belmont, MA 2000.

[14] R. Brealey and S. Myers, *Principles of Corporate Finance,* 5th Ed., Ch. 8, McGraw-Hill Companies, New York, 1996.

[15] H. M. Markowitz, "Portfolio Selection," *Journal of Finance*, March 19, 1952.

[16] K. Terhaar, R. Staub, and B. Singer, "Appropriate Policy Allocation for Alternative Investments," *The Journal of Portfolio Management,* (Spring 2003) 101-110.

[17] J. Cremers, M. Kritzman, S. Page, "Optimal Hedge Fund Allocations: Do Higher Moments Matter?" *Revere Street Working Paper Series, Financial Economics 272-13*, September 3, 2004.

[18] D. P. Bertsekas with A. Nedic and A. E. Ozdaglar, *Convex Analysis and Optimization*, Athena Scientific, Belmont, MA, 2003.