Optimal Rebalancing for Institutional Portfolios

Minimizing costs using dynamic programming.

Walter Sun, Ayres Fan, Li-Wei Chen, Tom Schouwenaars, and Marius A. Albota

Walter Sun is a research assistant at the MIT Laboratory for Information and Decision Systems in Cambridge, MA.
waltsun@alum.mit.edu

Ayres Fan is a research assistant at the MIT Laboratory for Information and Decision Systems in Cambridge

Li-Wei Chen is a research assistant at the MIT Laboratory for Information and Decision Systems in Cambridge

Tom Schouwenaars is a research assistant at the MIT Laboratory for Information and Decision Systems in Cambridge

Marius A. Albota is a research assistant at the MIT Research Laboratory of Electronics in Cambridge

Institutional money managers develop risk models and optimal portfolios to match a desired risk/reward profile. Utility functions express risk preferences and implicitly reflect the views of fund trustees or directors.

Once a manager determines a target portfolio, maintaining this balance of assets is non-trivial. A manager must rebalance actively because different asset classes can exhibit different rates of return. Managers also must rebalance if weights in the target portfolio are altered. This occurs when the model for expected returns of asset classes changes or the risk profile changes.

Most academic theory ignores frictional costs, and assumes that a portfolio manager can simply readjust holdings dynamically without any problems. In practice, trading costs are non-zero and affect the decision to rebalance. Transaction costs involve commissions and market impact as well as cost of personnel and technological resources. If the transaction costs exceed the expected benefit from rebalancing, no adjustment should be made, but 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 (see Donahue and Yip [2003] and Masters [2003]). With periodic rebalancing, the portfolio manager adjusts 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.

Thus, rebalancing may occur even if the portfolio is nearly optimal.

Tolerance band rebalancing requires managers to rebalance whenever any asset class moves beyond some predetermined tolerance band (e.g., ±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.

Research on dynamic strategies for asset allocation has established a so-called no-trade region around the optimal target portfolio weights (see Perold and Sharpe [1995] and Leland [1999]). If the proportions allocated to each asset at any given time lie within this region, trading is not necessary. If current asset ratios lie outside the no-trade region, though, Leland [1999] has shown that it is optimal to trade but only to bring the weights back to the nearest edge of the no-trade region rather than to the target ratios.

The optimal strategy has been shown to reduce transaction costs by approximately 50%, but the full analytical solution involves a complicated system of partial differential equations in multiple dimensions.

Mulvey and Simsek [2002] model the problem of rebalancing in the face of transaction costs as a generalized network with side conditions and develop an algorithm for solving the resulting problem. Mitchell and Braun [2002] describe a method for finding an optimal portfolio when proportional transactions costs have to be paid.

Since then, Donohue and Yip [2003] have confirmed the results of Leland [1999]. They characterize the shape and size of the no-trade region, and compare the performance of different rebalancing strategies.

We present an approach that explicitly weighs transaction costs and portfolio tracking error. We assume we are living in a CAPM world, which means that asset returns are stationary, and mean and variance are the primary portfolio statistics of interest (see Markowitz [1952]). Utility functions coupled with asset return models then yield a target portfolio that is a set of optimal weights for different asset classes. We also assume the portfolios are either tax-free or tax-deferred, which is the case for endowments, charities, pension funds, and most individual retirement funds.

The main difficulty with reconciling transaction costs and tracking error is that they are expressed in different units. *Transaction costs* are measured tangibly by dollars. *Tracking error* is a more abstract concept. Because the optimal portfolio is the portfolio that maximizes our given utility function, we can express tracking error as the shortfall in utility from our current portfolio to the optimal portfolio.

Our first contribution is applying the concept of *certainty-equivalents* to create risk-adjusted returns that allow us to convert tracking error into a dollar-denominated cost (see Bernoulli [1954]). Note that we are not restricted to quadratic utility but can use arbitrary utility functions. Once we have a dollar cost, we can directly compare the transaction costs for rebalancing with the suboptimality costs for not rebalancing.

Yet this is leaving out an essential piece of the puzzle: Our actions this period also affect outcomes and decisions in future periods. Our second contribution is to then apply the method of *dynamic programming* to minimize a cost function that explicitly models this point. Thus our optimal policy trades only when the expected cost of trading is less than the expected cost of doing nothing, evaluating costs over the next period and all future periods.

In addition, we search over the rebalancing space from 0% (no rebalancing) to 100% rebalancing (full rebalancing) and the points in between. In most cases, partial rebalancing can provide nearly the same utility as full rebalancing while saving on transaction costs.

Our third contribution is a framework to evaluate rebalancing strategies quantitatively. We show that our method performs better than traditional methods of rebalancing and is robust to model error.

OPTIMAL REBALANCING USING DYNAMIC PROGRAMMING

We consider a multiasset problem where we are given an optimal portfolio consisting of a set of target portfolio weights $w^* = \{w_1, w_2, \ldots, w_N\}$, 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.

The model allows us to observe the contents of the portfolio w_t at the end of each month. At this 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 $w_t + u_t$. Assuming normal returns in the process noise are subject to 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.

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 onward, given our actions in the current month. The optimal strategy, determined through dynamic programming, minimizes the sum of these three costs.

Dynamic programming is an optimization technique that finds the policy that minimizes expected cost, given a cost function and a dynamic model of state behavior (see Bellman [1957], Bellman and Dreyfus [1962], and Bertsekas [2000]). The cost at any given period is the expected cost from t to $t + 1$ along with the expected cost from $t + 1$ onward.

Assuming convergence, this recursion approaches a fixed point known as the *cost-to-go* value. Once these values are known, the optimal rebalancing decision is to choose the policy that achieves this minimum.

To apply dynamic programming, we specify the cost function as a sum of trading and suboptimality costs where the trading costs include tangible fees such as commissions and market impact, but also indirect costs such as employee labor, while the suboptimality cost represents the cost of not having an optimal portfolio.

Note that the cost-to-go values, and hence the optimal strategy, will depend on the cost functions chosen. In the certainty-equivalence approach, we model the investor's preferences using a utility function (see Luce [2000]).

Because future returns are unknown, we need to use expected utility to create an optimal portfolio or to decide on a rebalancing policy. Levy and Markowitz [1979] have shown 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, $U(\mu, \sigma)$.

Exhibit 1 lists the three utility functions and the corresponding expected utilities that we use (see Cremers, Kritzman, and Page [2004]). For each utility, $f_i(x)$ for $i = \{q, l, p\}$ (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. $U(\mu, \sigma)$ for $i = \{q, l, p\}$ is the expected utility. Details and derivations may be found in Sun et al. [2004].

EXHIBIT 1
Utility Functions with Corresponding Approximate Expected Utilities

	Utility Function	Expected Utility
Quadratic	$f_q(x) = x - \alpha/2\,(x - x_0)^2$	$U_q(\mu,\sigma) = \mu - \alpha/2\,\sigma^2$
Log wealth	$f_l(x) = \log(1 + x)$	$U_l(\mu,\sigma) = \log(1 + \mu) - \sigma^2/[2(1 + \mu)^2]$
Power	$f_p(x) = 1 - 1/(1 + x)$	$U_p(\mu,\sigma) = 1 - 1/(1 + \mu) - \sigma^2/(1 + \mu)^3$

For any portfolio with weights 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 call this the *certainty-equivalent return* for the weights w.[1]

One interpretation of the certainty-equivalent then is 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 $U(w)$ will be lower than $U(w^*)$, with a correspondingly lower certainty-equivalent return. We can interpret this as losing a risk-free return (equal to the difference between the two certainty-equivalents) over one period, corresponding to the penalty paid for tracking error. The difference between the certainty-equivalents of a non-optimal portfolio and the optimal portfolio is defined as the cost of not being optimal.

We use a certainty-equivalent because the trading and suboptimality costs 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 then to convert the portfolio tracking error into a similar absolute measure using certainty-equivalents rather than try to express the trading costs in terms of diminished expected utility (although this can be done as well).

Assume we have a portfolio w, and we want to go to another portfolio w_0. The most basic model for transaction costs is simply to assume a linear cost. This is just one specific choice of transaction cost. Alternative transaction cost functions can be chosen as appropriate. For example, price impact models can be used, or a proportional plus fixed cost model can be used to discourage frequent trading.

EXPERIMENTAL RESULTS

We use five asset classes: hedge funds, developed markets equity, emerging markets equity, U.S. equity, and private equity. Exhibit 2 shows the mean and standard deviation of each asset class for historical monthly

EXHIBIT 2
Annual Mean Returns and Standard Deviations for Asset Classes

	Index as Proxy (Source)	Mean Return (%)	Std Dev (%)
Hedge Funds	HFR Mkt Neutral (Bloomberg)	5.28	10.16
Developed Markets	MSCI EAFE + Canada (Datastream)	6.65	16.76
Emerging Markets	MSCI EM (Datastream)	7.88	23.30
US Equity	Russell 3000 (Datastream)	6.84	14.99
Private Equity	Wilshire LBO (Bloomberg)	12.76	44.39

EXHIBIT 3
Correlation Coefficients

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

returns over the last decade. The correlation matrix is shown in Exhibit 3.

Of the different assets, private equity provides the highest expected return, but it also has the greatest amount of risk. At the other extreme, hedge funds have both the lowest expected return and the least amount of variability. (The mean returns were provided by State Street Associates, and the variances and correlations are computed empirically from data acquired from Datastream and Bloomberg.)

To introduce the problem of portfolio rebalancing, we first consider an example involving only two asset classes: developed markets and hedge funds. The benefits of the two-risky asset model are that: the optimal portfolio can be computed in closed form (see Brealey and Myers [1996] 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 there are few enough parameters that we can easily perform sensitivity analyses. We later carry out extensive simulations with a multiasset model.

At first we consider only quadratic utility with risk aversion parameter $\alpha = 1.5$. Using this assumption, the optimal portfolio balance is 51% in developed markets and 49% in hedge funds. To illustrate the behavior of our rebalancing method, we simulate the returns of the two asset classes over a single 20-year realization.

Exhibit 4 shows how the portfolio weight of developed markets moves over one 240-month sample path. With no rebalancing (Panel A), the weight drifts from the optimal amount of 49% to under 25%, resulting in high suboptimality costs.

The optimal rebalancing strategy rebalances only when necessary (Panel B). During months 110–120 and 170–215, the portfolio partially rebalances nearly every month to handle sharp changes in the portfolio, while for months 120–160, the lack of strong market movements in either direction allows us to avoid any transaction costs. The market movement during these times can be seen by examining the change in portfolio weights in Panel A, where there is no rebalancing.

We evaluate different rebalancing algorithms using a Monte Carlo simulation process. 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.

Exhibit 5 shows the annualized costs of different rebalancing strategies. Trading costs are 40 basis points for buying or selling developed markets and 60 bp for buying or selling hedge funds.

We show two metrics to evaluate performance. In both cases, the numbers are in terms of shortfall from an idealized rebalancing strategy that is allowed to rebalance to the optimal portfolio for free every month.

The first metric is to aggregate expected cost shown in the third column. This is simply the sum of the trading costs in the first column and the suboptimality costs in the second column. The suboptimality cost is computed at the end of each month as the difference between the expected risk-adjusted returns for the optimal portfolio and the current portfolio after the rebalancing policy has been applied. Note that this is precisely the metric that our dynamic programming approach is designed to minimize.

The second metric is what we call *empirical utility shortfall*. For each month, we compute the return net of transaction costs and then compute the *utils* associated with that return using our empirical utility function f_i (see Exhibit 1). This metric is similar to aggregate cost, but works with actual utility rather than the expected utility used to compute certainty-equivalents. We expect the sample average and the expected value to converge, which they do in our case.[2]

EXHIBIT 4
Developed Market Weights in Two-Asset Example—Different Rebalancing Models

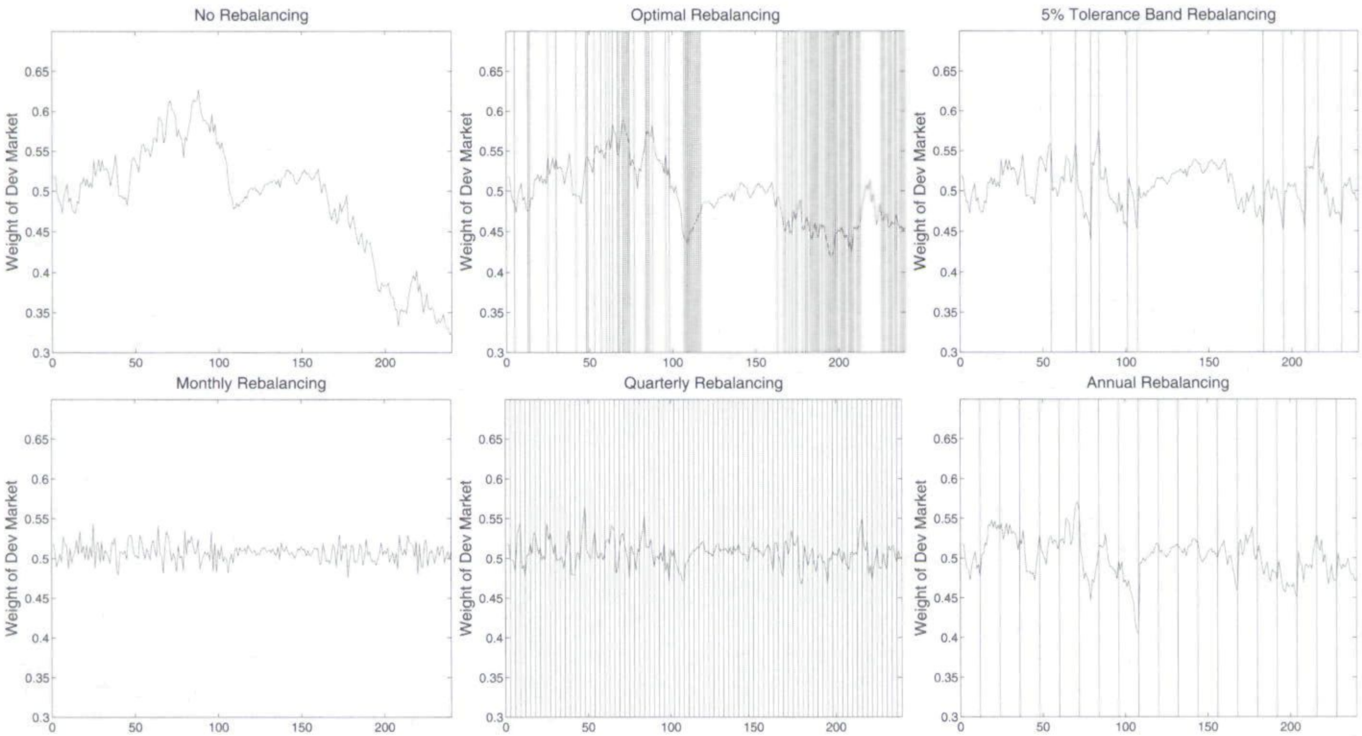

Vertical lines indicate rebalancing events (omitted in monthly rebalancing since it occurs at every time point).

EXHIBIT 5
Annualized Costs and Utility Shortfall—Six Different Rebalancing Strategies—Two Risky Assets

	Trading Cost (bps)	Suboptimality Cost (bps)	Aggregate Cost (bps)	Utility Shortfall (utils × 10^4)
Ideal	0.00	0.00	0.00	0.00
Optimal DP	1.57	0.46	2.03	2.15
No Trading	0.00	4.74	4.74	4.56
5% Tolerance	3.68	0.13	3.81	3.65
Monthly	12.92	0.00	12.92	12.94
Quarterly	7.46	0.05	7.51	7.46
Annual	3.71	0.26	3.97	4.05

Results are an average of 10,000 realizations.

From Exhibit 5 we observe that our method minimizes the aggregate cost. Assuming a portfolio of \$1 billion, the aggregate annual cost by our algorithm is \$203,000. The cost for rebalancing using 5% tolerance, the next-least expensive method, is \$381,000 annually.

The results for each rebalancing method make intuitive sense. Monthly rebalancing produces no deviation from optimality, but at the cost of high trading fees. Infrequent trading yields lower trading costs, but higher suboptimality 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.

So far we have assumed that the model for each asset is known. In practice, the mean and variance of each asset's returns as well as the correlation between the assets must be estimated (using historical observations). Errors in the parameter estimate, such as deviation from the true unknown value, will cause inaccuracies in the costs-to-go obtained from the dynamic program, resulting in suboptimal rebalancing.

To investigate the impact of errors in each of these parameters on the rebalancing strategy, we consider the means, standard deviations, and correlation parameters. For each simulation, only one of the three parameters is varied to isolate the effects of changes to each of these

EXHIBIT 6
Sensitivity Analysis for Two-Asset Example

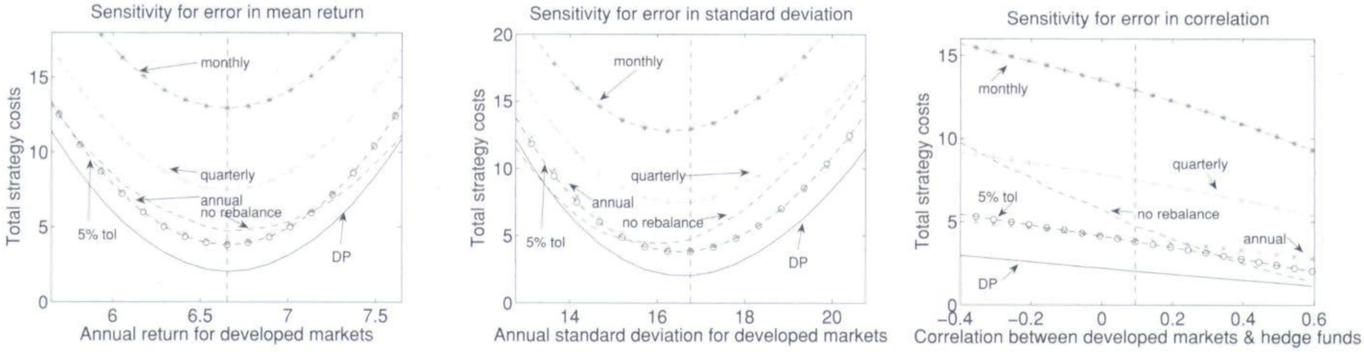

A. Mean return. B. Standard deviation. C. Correlation coefficient.
Vertical line indicates actual rate used by the rebalancing strategies.

factors. We then compute annualized aggregate cost, while the true value of the third parameter varies around the assumed value. This cost is computed relative to an idealized portfolio that knows the correct value and rebalances for free to the true optimal portfolio every month.

Most rebalancing strategies are inherently model-based because the target portfolio is determined by the model. When the target portfolio is determined with incorrect parameters, the target portfolio is itself suboptimal. Therefore we would expect all the rebalancing methods to be sensitive to model error. The exception to this among the rebalancing methods we examine is the strategy of no-rebalancing.

For our testing methodology, the assumed model plays no role except in determining the initial portfolio. Therefore we would expect no rebalancing to be least sensitive to model error. Our algorithm should be even more strongly influenced by the model because incorrect assumptions can also result in suboptimal rebalancing decisions from the dynamic program. Other methods such as calendar and tolerance band strategies are heuristic, and do not directly rely on the model when forming the rebalancing policy.

The results for mean sensitivity are shown in Exhibit 6, Panel A. For each point, 10,000 sequences of 20 years of monthly returns are generated, and the performance of each rebalancing strategy is averaged over each sequence.

Our dynamic rebalancing strategy is as robust as tolerance band and calendar rebalancing. We perform better than no-rebalancing until the annual return exceeds 7.6%.

This is because the annual return for developed markets is higher than the return for hedge funds. So, without rebalancing, as the annual return for developed markets increases, the average weight for developed markets tends to increase as well. At the same time, it becomes more advantageous to hold developed markets due to the higher return, so the true optimal portfolio will include a higher percentage of developed markets. Thus, in this case, a passive strategy will typically let the portfolio move inadvertently toward the optimal portfolio.

In Exhibit 6, Panel B, we see that the dynamic programming approach again outperforms the other approaches despite large errors in estimating the standard deviation—it still outperforms the other methods even for inaccuracies in the standard deviation of several percentage points per year. In this instance, the strategy of no-rebalancing performs better for lower standard deviations.

As in the mean sensitivity case, because developed markets have a higher average return than hedge funds, their weight tends to increase with time. So as the standard deviation for developed markets falls, the true optimal portfolio will again have a higher weight in developed markets, and the portfolio will again tend to drift toward the optimal portfolio.

Finally, in Exhibit 6, Panel C, we observe that the dynamic programming approach is relatively insensitive to errors in estimation of the correlations between the assets. As the cost does not change much for different true correlations, we conclude that correlations need not be accurately estimated for the purposes of our approach.

EXHIBIT 7
Efficient Frontier and Optimal Portfolios for Different Utility Functions Discussed

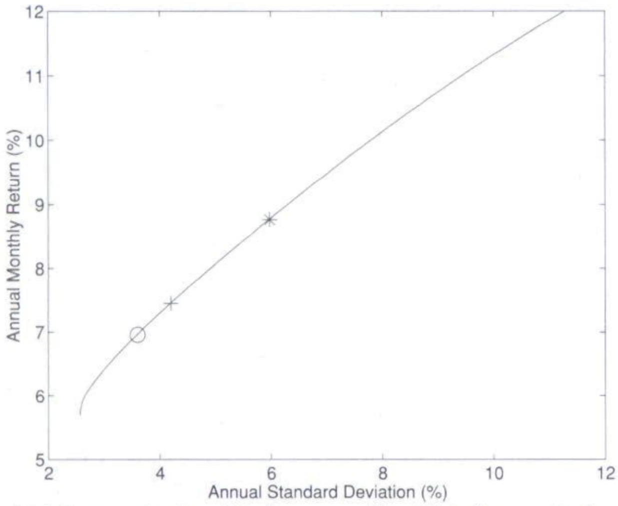

(o) indicates optimal portfolio for power utility. (+) indicates optimal portfolio for quadratic utility with $\alpha = 1.5$. (*) indicates optimal portfolio for log wealth utility.

In the general case of N risky assets, we examine five risky assets, and assert that another choice of $N > 2$ would proceed similarly; the main difference is computation time. According to Cremers, Kritzman, and Page [2003], 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 one is using more advanced utility functions such as log wealth or power utility.

Later work by Cremers, Kritzman, and Page [2004] 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 in terms of just the mean and standard deviation. They show that the resulting portfolios and portfolios generated from full-scale optimization do not perform significantly differently.

Under this approximate mean-variance optimization, the optimal portfolio lies on the efficient frontier (see Markowitz [1952]).[3] Therefore, to construct optimal portfolios for different utility functions, we first compute the efficient frontier by solving a quadratic programming problem and then search those portfolios to find the one with the highest expected utility.

Exhibit 7 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 with the weights shown in Exhibit 8. These weights are the optimal weights we use throughout the analysis. Note that power utility is the most risk-averse utility, and log wealth is the risk-seeking utility.

Monte Carlo Simulations

Exhibit 9 shows the results of various rebalancing algorithms for the five-asset case. The results are generated in a manner analogous to the two-asset case with Monte Carlo simulations over 20 years for 10,000 sample paths. Trading costs are 60 bp for hedge funds, emerging markets, and private equity; 40 bp for developed markets; and 30 bp for U.S. equity.

For the quadratic utility case, we see that our optimal dynamic programming (DP) method performs approximately 30% better than the next-best method, which is 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 high suboptimality cost because over a 20-year period assets can become quite unbalanced if unadjusted.

For power utility, our expected loss is 24% less than the runner-up, 5% tolerance band rebalancing. The

EXHIBIT 8
Optimal Portfolio Weights for Different Utility Functions

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

Exhibit 9

Quadratic, Power, and Log Wealth Utility Costs and Shortfall—Six Different Trading Strategies—Five Risky Assets

		Trading Cost (bps)	Suboptimality Cost (bps)	Aggregate Cost (bps)	Utility Shortfall (utils × 10^4)
Quadratic $\alpha = 1.5$	Ideal	0.00	0.00	0.00	0.00
	Optimal DP	3.97	1.49	5.47	5.23
	No Trading	0.00	30.18	30.18	32.89
	5% Tolerance	7.29	0.70	7.99	7.88
	Monthly	23.67	0.00	23.67	23.73
	Quarterly	13.69	0.28	13.96	14.21
	Annual	6.84	1.55	8.39	8.43
Power	Ideal	0.00	0.00	0.00	0.00
	Optimal DP	3.39	1.15	4.54	4.32
	No Trading	0.00	25.99	25.99	26.04
	5% Tolerance	5.19	0.82	6.01	5.95
	Monthly	20.05	0.00	20.05	19.96
	Quarterly	11.59	0.18	11.78	11.87
	Annual	5.81	1.02	6.84	6.80
Log wealth	Ideal	0.00	0.00	0.00	0.00
	Optimal DP	4.80	1.93	6.72	6.57
	No Trading	0.00	30.52	30.52	32.01
	5% Tolerance	11.95	0.43	12.38	12.72
	Monthly	28.15	0.00	28.15	28.19
	Quarterly	16.27	0.40	16.67	16.80
	Annual	8.05	2.17	10.22	10.65

Simulated over a 20-year period 10,000 times.

benefits for this method are reduced from the quadratic utility case primarily because power utility is more risk-averse, so the optimal portfolio has a lower variance. Therefore less rebalancing is needed overall, because the portfolio simply does not move as much. This can be seen in the quarterly and annual rebalancing methods, which trade less and thus suffer lower suboptimality costs.

For the log wealth utility case, our algorithm results in an expected loss that is 35% less than the best alternative, annual rebalancing in this case. The 5% tolerance method falls short in this case simply because the log wealth portfolio is a high-variance portfolio. In the quadratic case, the trading costs are only marginally higher than the annual rebalancing method. But in the log wealth case, they are 48% higher because the tolerance bands are breached more often.

As a general rule, we see it is more important to get the rebalancing right when dealing with higher-variance portfolios simply because many more rebalancing opportunities arise.

Note that even though tolerance bands do better than annual rebalancing for quadratic utility and power utility and worse for log wealth utility, it is not a valid conclusion that tolerance band methods perform better for low-variance portfolios and calendar rebalancing methods perform better for high-variance portfolios. There is a degree of freedom in each algorithm.

For tolerance bands, the bands can be loosened or tightened, depending on the variance of the optimal portfolio. For example, in the log wealth case, it is clear that the tolerance bands are too tight because of the sheer imbalance between trading costs and suboptimality costs.

For calendar algorithms, the period between rebalancing can be adjusted. For instance, setting the rebalancing time to two years for the power utility case results in an expected loss of 6.32 bp per year, a savings of 0.52 bp over the annual strategy. This is achieved by accruing more than twice as much expected tracking error (2.21 bp versus 1.02 bp) but also reducing trading costs by 29% (4.11 bp versus 5.81 bp). A more exhaustive search of possible fixed-interval rebalancing strategies could presumably yield an even better result.

Optimal Heuristic Algorithms

Our 5% tolerance and periodic rebalancing periods of monthly, quarterly, and annually are a subset of general tolerance band and calendar-based rebalancing algorithms. We make no claims that the parameters chosen are optimal but rather that the settings seem to be common in the literature.

Our work can be thought of in two ways. First, we show that our method is superior to any tolerance band or calendar-based rebalancing method. It should be noted that our method may be thought of as a dynamic tolerance band approach. Fixed-tolerance methods are a subset of the controls available to our algorithm, so they can never do better.

Second, the full dynamic programming algorithm is computationally intensive, especially with many assets. If we instead confine the search to a smaller subset of acceptable policies, we have a strategy that is not necessarily optimal, but perhaps good enough.

EXHIBIT 10

Annualized Strategy Costs for Different Calendar-Based and Tolerance Band Settings—Quadratic Utility

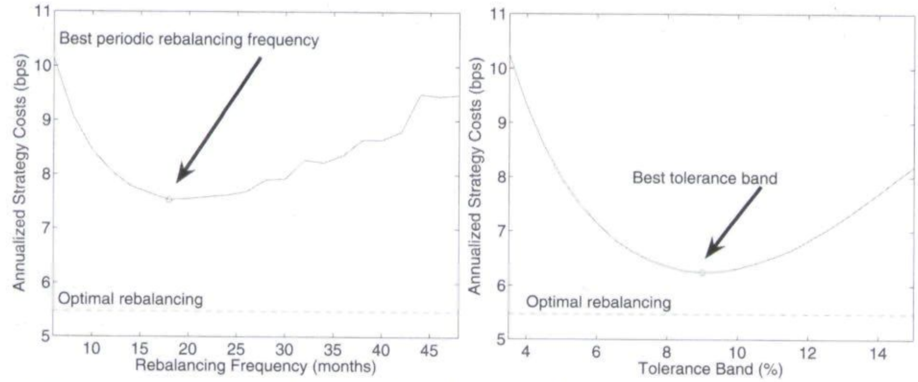

A. Plot of costs as function of different periodic rebalancing periods.

B. Plot of costs as function of different tolerance band windows.

Exhibit 10 graphs the results for the quadratic utility case. Each point is obtained by performing Monte Carlo simulations over 20 years for 10,000 realizations. In Panel A, we plot the total costs as a function of the number of months between rebalancing. The lack of smoothness in the curve, particularly between less frequent rebalancing, arises because the amount of rebalancing in a 240-month period must necessarily change in discrete steps.

We see that the best performance occurs for 18-month periods at a cost of 7.53 bp. Note this is still over 2 bp worse than our algorithm's performance (shown as the horizontal dashed line) but nearly 1 basis point better than annual rebalancing.

In Panel B, we vary the tolerance setting that triggers rebalancing. We find that the best setting occurs for a tolerance band of 9% and results in a cost of 6.25 bp, again worse than the 5.47 bp obtained by our algorithm.

Computational Complexity

To provide some information regarding the computational complexity of our approach, note 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 in which we must develop the optimal policy. Our current implementation processes around 600,000 points per hour on a single PC, although value iteration can be easily parallelized, so the total processing time also depends on the number of machines available.

For a non-parallelized implementation, the runtime estimate for five assets is 75 minutes. If we assume the possibility of M different weights for an additional asset, addition of this asset to our N-asset model would increase computation time by a factor of M. Memory requirements increase by a similar amount.

Note this includes detailing the computation for learning the optimal policy. Once that is done, actually applying the policy is very fast.

Alternative Cost Functions

Some may wonder whether the results would be different for alternative trading costs. Exhibit 11 shows the results when we cut the trading costs in half and apply this to the quadratic utility strategy. 5% tolerance bands remain the next-best strategy, but our advantage has narrowed from 30% to 20%. The reason is that the other algorithms trade too much, and now they are penalized less for it.

Transaction costs for the other methods are cut in half, while suboptimality remains the same. This produces reductions in aggregate cost ranging from 41% for annual rebalancing to 50% for monthly rebalancing (ignoring no-rebalancing, which obviously does not benefit).

Our algorithm actually reduces trading costs by less than half. This means that we are sensibly trading more, now that transaction costs have been lowered—so our algorithm automatically trades off a little extra trading to reduce the suboptimality costs by a greater amount.

CONCLUSION

The ad hoc methods of periodic and tolerance band rebalancing provide simple but suboptimal ways to

EXHIBIT 11
Quadratic Utility Costs and Shortfall—Transaction Costs Halved

	Trading Cost (bps)	Suboptimality Cost (bps)	Aggregate Cost (bps)	Utility Shortfall (utils × 10^4)
Ideal	0.00	0.00	0.00	0.00
Optimal DP	2.60	0.85	3.45	3.75
No Trading	0.00	30.18	30.18	32.89
5% Tolerance	3.64	0.70	4.34	4.53
Monthly	11.83	0.00	11.83	11.86
Quarterly	6.84	0.28	7.12	7.18
Annual	3.42	1.55	4.97	5.15

rebalance portfolios. Calendar-based approaches rely on the fact that, on average, we expect a 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 uses the current portfolio to make a decision, but this method has no sense of the proper tolerance band setting, or even how wide this band should be.

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. The reduced costs hold for different investor risk preferences, whether quadratic, log wealth, or power utility functions.

The costs of transactions are much more tangible than the cost of being suboptimal. Through the use of certainty-equivalents, however, we have provided a method that quantifies the cost of being suboptimal. Our simulations confirm that this optimal method provides gains over the best of the traditional techniques of rebalancing.

Note that the analysis assumes returns are independent across different intervals. The literature observes that mean reversion may exist. Under such circumstances, we expect our method to perform even better than periodic rebalancing, because our algorithm would likely rebalance even less often.

There are several possible extensions of our work. First, one may want to consider proportional plus fixed transaction costs. This model is appropriate if one believes there is a fixed cost in making each and every transaction. Such an adjustment would likely favor dynamic trading methods over periodic rebalancing.

Next, we might want to examine rebalancing over taxable portfolios. Asset managers of such funds have the additional consideration of tax consequences when they decide to transact.

The relaxation of the short sales constraint is another possible extension. Although many tax-deferred funds do not allow short sales, several either explicitly allow short-selling or implicitly participate in short sales through investments in hedge funds.

We also assume an instantaneous rebalancing at the end of each month. One could incorporate more general trading models that consider the effects of price impact.

Finally, for the multiasset case, we search a one-dimensional policy space representing portfolios that 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 this case, it may be better to trade on only a subset of assets rather than a portion of all asset classes.

ENDNOTES

The authors thank Sebastien Page and Mark Kritzman for their guidance and advice and Edward Freyfogle and Joshua Grover for their input. This article originated in a course project at the Massachusetts Institute of Technology's Sloan School of Management.

[1]The condition for this is $U(r_{CE}, 0) = U(\mu^T w, w^T \Lambda w)$. The certainty-equivalents for the three expected utility functions that we are using are:

Quadratic: $r_{CE}(w) = U_q(\mu^T w, w^T \Lambda w)$

Log wealth: $r_{CE}(w) = \exp(U_l(\mu^T w, w^T \Lambda w)) - 1$

Power: $r_{CE}(w) = 1/(1 - U_p(\mu^T w, w^T \Lambda w)) - 1$

[2]The units on *utils* multiplied by 10^4 in the last column of Exhibit 5 are similar to the basis points in the first three columns. This is clear for the quadratic case where the certainty-equivalent is equal to the empirical 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 is

reasonably commensurate with basis points and explains why the numbers in the last two columns are similar.

[3]Risk-averse expected utility functions increase monotonically in terms of return and decline monotonically in terms of risk. Hence, if a portfolio is not on the efficient frontier, there is a portfolio with equivalent return and less risk or more return and the same risk; therefore, this portfolio cannot be optimal.

REFERENCES

Bellman, R. *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957.

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

Bernoulli, D. "Exposition of a New Theory on the Measurement of Risk." *Econometrica*, 22 (1954), pp. 23-36.

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

Brealey, R., and S. Myers. *Principles of Corporate Finance*, 5th ed. New York: McGraw-Hill, 1996.

Cremers, J., M. Kritzman, and S. Page. "Optimal Hedge Fund Allocations: Do Higher Moments Matter?" Revere Street Working Paper Series, Financial Economics, 272 (2004).

——. "Portfolio Formation with Higher Moments and Plausible Utility." Revere Street Working Paper Series, Financial Economics, 272 (2003).

Donahue, C., and K. Yip. "Optimal Portfolio Rebalancing with Transaction Costs." *The Journal of Portfolio Management*, Summer 2003, pp. 49-63.

Leland, H. E. "Optimal Portfolio Management with Transaction Costs and Capital Gains Taxes." Haas School of Business Technology Report, University of California at Berkeley, 1999.

Levy, H., and H. M. Markowitz. "Approximating Expected Utility by a Function of Mean and Variance." *American Economic Review*, 69 (1979), pp. 308-317.

Luce, R. D. *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc., 2000.

Markowitz, H. M. "Portfolio Selection." *Journal of Finance*, Volume 7, Number 1, 1952, pp. 77-91.

Masters, S. J. "Rebalancing." *The Journal of Portfolio Management*, Spring 2003, pp. 52-57.

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

Mulvey, J. M., and K. D. Simsek. *Rebalancing Strategies for Long-Term Investors*. New York: Kluwer Academic Publishers, 2002.

Perold, A. F., and W. F. Sharpe. "Dynamic Strategies for Asset Allocation." *Financial Analysts Journal*, January/February 1995, pp. 149-160.

Sun, W., A. Fan, L. Chen, T. Schouwenaars, and M. Albota. "Optimal Rebalancing Strategy Using Dynamic Programming for Institutional Portfolios." Social Science Research Network Electronic Library, MIT, 2004. Available at http:// ssrn.com/ abstract=639284.