TIME OPTION REBALANCING

by

Richard O. Michaud, Ph.D.
ACADIAN ASSET MANAGEMENT
Boston, MA 02110

September 6, 1990

Portfolio rebalancing is one of the most important investment management activities. Rebalancing is used to control portfolio risk, reflect new information and maintain the asset allocation structure to meet the fund's long-term investment objectives.

The objective is to illustrate rules that optimize asset allocation rebalancing over time A fund's Investment objectives are generally defined in terms of the normal percent of assets in various asset classes such as (domestic or foreign) stocks, bonds, cash, real estate, venture capital. Such guidelines are often the result of an intensive study that may include careful definition of assets, liabilities and objectives.

INVESTMENT POLICY AND PORTFOLIO REBALANCING

"Investment policy" or "strategic asset allocation," the expected long term average asset allocation of a fund, is widely considered to be the single most important investment decision (Brinson and Diermeier, 1985). Investment policy defines the "normal" asset mix that serves as a benchmark for rebalancing the portfolio over time.

Investment policy studies to define appropriate asset allocation weights can differ widely in scope. Common features include the use of historical monthly or quarterly total return data to simulate investment performance under a variety of assumptions. While the studies are designed to be "long term," their effective planning horizon is generally three to five years. This is because assumptions concerning regulation, risk and return structure of financial assets, plan liabilities and business risks, are unlikely to be valid over longer time periods. Therefore, implementation of the recommendations of an investment policy study ("normal policy rebalancing") generally requires periodic rebalancing to the (same) normal asset mix over a three to five year time horizon.

"INFORMATION BASED" REBALANCING

Many investment managers will revise a fund's asset allocation on the basis of short term market "forecasts." The process is usually described as "investment strategy" or "market timing" and attempts to use near term information to enhance fund performance.

An alternative approach, called "tactical asset allocation" (TAA), does not use a "traditional" forecasting process. Decisions are based on models of "equilibrium" relationships and asset class "yields" that depend solely on "current" information. Because of the absence of traditional "forecasting," proponents have claimed that TAA is a more reliable rebalancing strategy.

While it is arguable that the two "information based" rebalancing strategies differ with respect to assumed risk, both are, by definition, excess risk and return strategies when compared to the normal policy rebalancing benchmark. In addition, both imply the existence of time sensitive information that may lead to frequent portfolio rebalancing and excess transactions costs.

A THIRD ALTERNATIVE: "TIME OPTION" REBALANCING (TOR)

A third alternative for defining an optimal rebalancing strategy, called "time option" rebalancing (TOR), occupies an intermediate position between normal policy and "information" based rebalancing strategies. TOR maximizes expected return while maintaining the same level of risk as normal policy without using time sensitive information. This is accomplished by using "time option" information -- number of periods remaining until the end of the planning horizon and past experience -- that is ignored in normal policy rebalancing. It can be implemented as an overlay strategy relative to normal policy. The optimization technique for defining an optimal rebalancing process based on the information contained in past experience and time remaining is "dynamic programming."

REBALANCING STRATEGY SPECTRUM

It may be useful to develop a framework for describing a spectrum of portfolio rebalancing strategies. Normal policy rebalancing is the benchmark that defines a zero excess risk and return strategy consistent with the fund's investment policy. Time option rebalancing is defined to have an optimized level of expected return at the same level of risk as that implied by the fund's investment policy. Tactical asset allocation based only on current information may have less risk and return than rebalancing based on a pure forecasting process. These hypothesized relationships are illustrated in Figure 1 below.

MEASURING INVESTMENT PERFORMANCE OVER TIME

In order to compute an optimal rebalancing policy, we must first define the appropriate measure of return over time. A simple example should convince the reader that average return is not appropriate. If an investor experiences a 100% return in one period and a -50% return in the next, the average return over the two periods is 25%. However the actual return is 0%. This is because a dollar invested has grown to two at the end of the first period and returned to the original level of investment at the end of the second.

The appropriate way to measure return over time is the geometric mean (compound return, growth rate). The definition of geometric mean return G_N over N periods is given by:

$$G_N(R) = \{(1 + r_1)^* \ldots ^*(1 + r_N)\}^{1/N} - 1$$

where $r_i > -1$ are the returns in periods 1, ..., N.

FIGURE 1

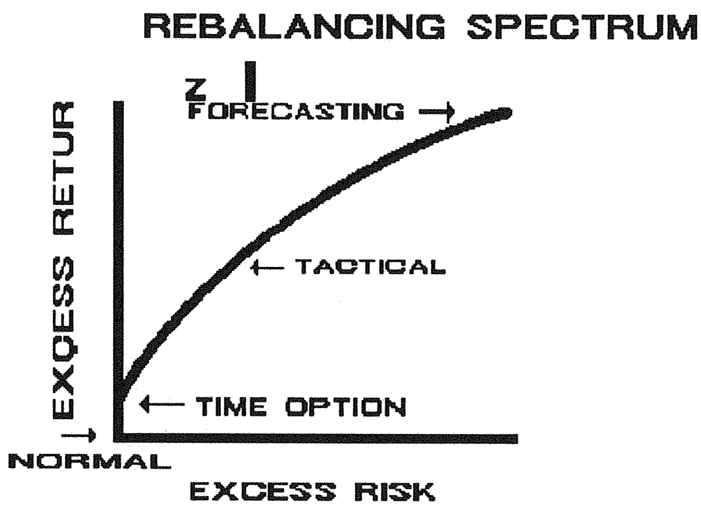

TOR AND THE "HAKANSSON" EFFICIENT FRONTIER

TOR is a rebalancing strategy that maximizes the expected (geometric mean) return at normal policy risk (variance of geometric mean) over the planning horizon. TOR leads naturally to consideration of the "Hakansson" (1971a) efficient frontier.

Hakansson efficiency is a generalization of Markowitz (1959) mean- variance efficiency for multiple time periods. It assumes that investors prefer maximum expected returns or minimum risk (variance) where return is defined as the geometric mean.[1]

TOR vs. normal policy rebalancing are illustrated in the Hakansson efficient frontier framework in Figure 2 below. As the Figure indicates, normal policy rebalancing implies an expected geometric mean return and variance. The point on the Hakansson efficient frontier at the same level of risk but above the expected return implied by normal rebalancing (along the line parallel to the Y axis) is the Hakansson time option rebalancing strategy.

FIGURE 2
HAKANSSON EFFICIENT FRONTIER
TIME OPTION VS. NORMAL REBALANCING

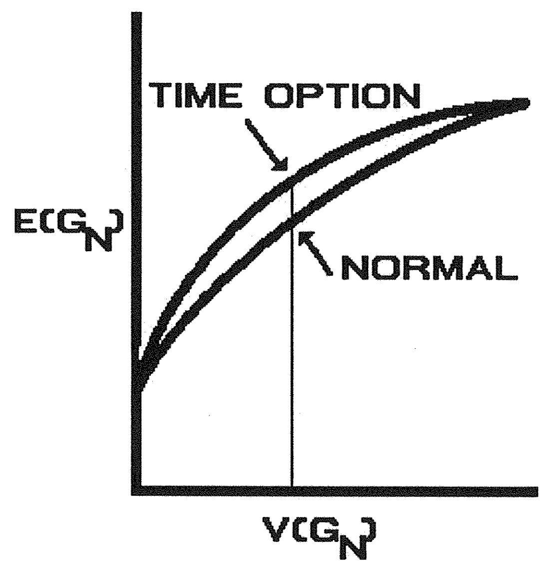

HAKANSSON EFFICIENT FRONTIER: THEORETICAL ISSUES

Time option rebalancing is the same as normal rebalancing strategy in two cases (Hakansson, 1971b): 1) At zero risk, assuming existence of a riskless asset over the horizon (lower point in Figure 2); 2) Maximum expected geometric mean (upper point in Figure 2). In most cases of practical interest, normal and time option rebalancing will lead to different strategies.

SIMPLIFYING ASSUMPTIONS

Technical details and computational difficulties limit the size of the problems that can be conveniently discussed. In this report, the scope is constrained to simple cases to illustrate the essential characteristics of the time option rebalancing strategy.

CASE 1: A SIMPLE ASSET ALLOCATION PROBLEM

Assume:

1) Two assets -- one risky (the market) and one riskless
2) Two rebalancing periods (initial and subsequent)
3) Risky asset has only two outcomes in each period:

R_w ("winning" outcome) probability p
R_l ("losing" outcome) probability 1-p

Let X be the variable that defines rebalancing strategy. For normal policy rebalancing,

$$X_f$$

is the normal or fixed proportion of assets in the risky security in each time period.

In general, an optimal rebalancing strategy depends on the period and on prior investment experience. Optimal policy in the initial period is a special case since there is no prior investment experience. However, optimal strategy in the second period depends on whether R_w or R_l occurred in the first period. Formally, we want to compute the optimal rebalancing strategy variables:

X_o First period
X_w Second period, conditional on R_w
X_l Second period, conditional on R_l

DECISION TREES: AN ILLUSTRATION

Decision trees are sometimes helpful in visualizing a multiperiod decision process. In Figure 3, a decision tree describing the two period returns for normal policy rebalancing is given. Each node indicates the possibility of either a winning (R_w) or losing (R_l) outcome in each period for the risky asset. There are two sets of nodes representing the two rebalancing periods. The ends of the nodes describe the return associated with the outcomes for the given investment policy decision X_f.

FIGURE 3

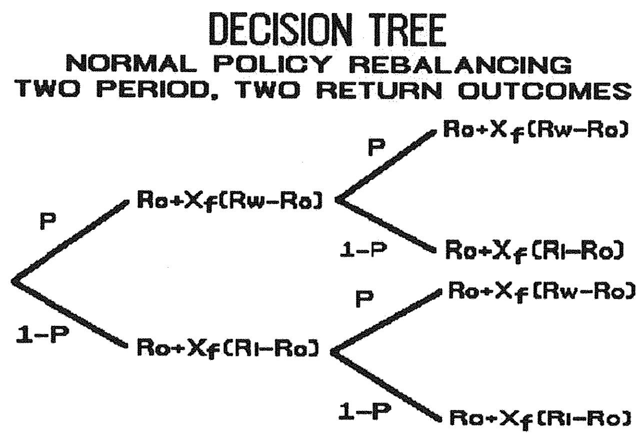

Figure 4 illustrates the same process for time option rebalancing. In this case, the investment policy decision variable changes from one time period to the next. In the second period, the investment policy decision variable also depends on the investment return in the previous period.

FIGURE 4

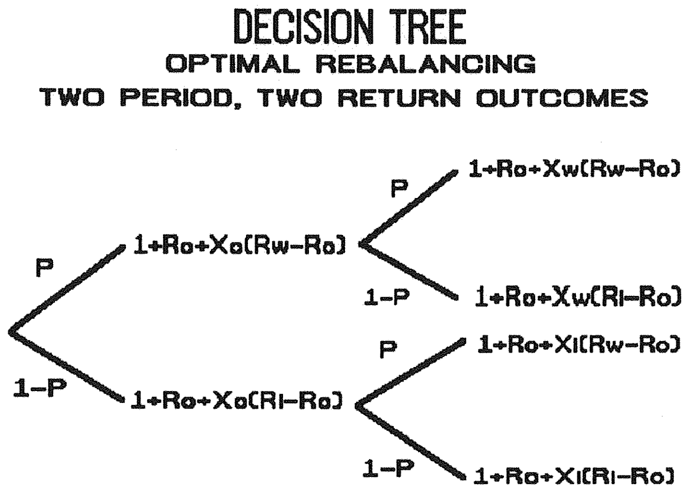

RESULTS: CASE 1

As Figure 2 shows, an optimized rebalancing strategy is defined with respect to the fund's normal investment policy. Since normal policy is assumed to be a given, the asset manager is primarily concerned with how much an optimized asset reallocation deviates from normal policy. Consequently, in the tables that follow, it will be convenient to report optimal time option rebalancing policies in terms relative to the value of the normal policy weight: $X_{(\ )}/X_f$.

Table 1 describes the optimal time option rebalancing policies with respect to a 50/50 normal investment policy in each period. The three lines in the table present the (relative) weights for an optimized rebalancing strategy for three different sets of market returns. For example, the second line assumes that the market either increases by 30% or declines by 10% in each period. The heading indicates that the probability of each outcome is one- half. Consequently, the second line reflects a market with a 10% expected return and a 20% standard deviation. The heading also shows that the riskless asset return is assumed to be 5%.

The data in the second line of Table 1 shows that an optimized rebalancing strategy should have 4% more than normal in the market at the beginning of the first period (column 1) (52% in the risky asset). The results (columns 2 and 3) show that an optimized rebalancing strategy will have 21% less than normal given a winning outcome in the first period (40% in the market) or 20% more than normal given a losing outcome in the first period (60% in the market).

THREE PRINCIPLES OF TIME OPTION REBALANCING

Examination of the results in Table 1 show that an optimized asset allocation rebalancing process will have the following characteristics:

1. Investment experience:
   Optimal rebalancing is a contrarian strategy; i.e., reduce the level of risk given favorable returns or increase the level of risk given unfavorable returns (columns 2 and 3).

2. Initial period:
   Optimal rebalancing requires taking "higher than normal" risk in the initial period (column 1).

3. The Effect of the Investment Environment:
   The more attractive the investment environment (higher expected return or lower volatility) the more optimal rebalancing deviates from normal policy, (rows 3 to 1).

The first result shows that TOR is fundamentally a contrarian strategy; i.e., when the market advances, take your gains and reduce risk, and conversely. The second result shows that TOR uses time efficiently; i.e., take more risk in the initial period to increase the likelihood of higher return while using subsequent periods to adjust risk depending on investment experience. The third result shows that an optimal rebalancing strategy increasingly deviates from normal policy as the opportunity for enhancing performance increases.

MORE REALISTIC ASSUMPTIONS

Loosening the restrictions in Case 1 can provide further insight into the characteristics of an optimized rebalancing strategy.

Case 2: Three market return outcomes -- win, expected, lose.
Includes the additional possibility that market return may be equal to expected return. The resulting distribution can be "more" normal. Consequently, we have an additional market return outcome in each period and another second period decision variable:

R_e expected outcome

X_e second period, conditional on R_e

Case 3: Three Periods.
Case 1 with an additional rebalancing period. Consequently, we have four additional third period decision variables:

x_{ww} R_w in 1st and 2nd periods
x_{wl} R_w in 1st, R_l in 2nd period
x_{lw} R_l in 1st, R_w in 2nd period
x_{ll} R_l in 1st and 2nd periods

REVISED PRINCIPLES OF OPTIMAL REBALANCING

The results in columns 2-4 in Table 2 significantly enhance our understanding of the rebalancing process. As before, columns 2 and 4 show that optimized rebalancing is contrarian when investment experience is counter to expectations. However, column 3 shows that optimal rebalancing is similar to normal policy if returns are consistent with expectations. Consequently:

1. Investment Experience
The investor should reduce the level of risk given favorable returns, increase the level of risk given unfavorable returns and maintain "near" normal levels of risk given expected returns.

Column 1 in Table 3 can be directly compared to Table 1 to show that increasing the number of periods increases the optimality of taking excess risk in the initial periods. Also columns 2 and 3 in Table 3 can be compared to the same columns in Table 1 to provide additional evidence of a tendency to increase risk in early periods as the number of periods remaining increases. Consequently:

2. Initial and intermediate periods:
The investor should increase exposure to risk in the early periods. In general, risk exposure should be increased as the number of adjustment periods available increases.

All three Tables confirm that optimal rebalancing will increasingly differ from normal policy in investment environments that provide significant opportunity to enhance performance, and conversely. Consequently:

3. Investment Environment:
The investor should increasingly deviate from normal policy weights as expected market return increases and as uncertainty decreases.

TAA VS. TOR

There are important similarities between TOR and TAA asset allocations. For TAA, favorable investment experience generally leads to a reduction in allocation to risky assets. This is because the computation of current asset class yields generally decline when assets have risen in value. Consequently, both strategies are contrarian asset allocation decision rules.

On the other hand, there are some interesting differences between the two strategies. TAA is time sensitive, requires the assumption of excess risk above normal policy, may result in substantial excess returns and uses models of capital market equilibrium to determine relative attractiveness. In contrast, TOR is not time sensitive (end of period rebalancing only), requires no excess normal policy risk, is likely to result in small though significant excess returns and uses no information not already available within the normal policy framework.

It is interesting to speculate that the popularity of TAA may be more strongly related to its contrarian character than to its forecasting ability. Many asset managers are attracted to a rebalancing process that reduces risk when performance has been favorable and conversely. TOR provides a rationale and procedure for contrarian rebalancing without forecasting that may be more consistent with investment intuition.

TOR AND MULTIPERIOD "DIVERSIFICATION"

It is of interest to note that the multiperiod performance of a normal policy rebalancing strategy resembles that of an undiversified portfolio, even if the underlying portfolio was well diversified in the traditional sense. In contrast, a TOR strategy leads to multiperiod returns that resembles that of a diversified portfolio. It is convenient to introduce the discussion by reviewing some basic empirical facts concerning traditional (single-period) diversification.

Annual equity total returns are highly right skewed or lognormal in shape. The typical or median return is much less than the average (Fisher and Lorie, 1970). This implies that an investor who invests in a single stock portfolio will have a relatively high probability of doing very well, a substantial probability of doing very poorly and a large probability of doing below market performance.

By aggregating stock returns into portfolios, Fisher and Lorie examined the performance of using increasingly diversified portfolios. They found that the returns of diversified portfolios resembled those of a normal distribution. This implies that a diversified, relative to an undiversified, investor will have a reduced probability of doing very well, a substantially increased probability of doing as well as the market and a reduced probability of doing very poorly. These concepts are illustrated in Figure 4 below.

Normal policy rebalancing over time leads to highly right skewed (log normal) portfolio returns.[2] This is true even if the original portfolio is well diversified in the traditional sense. This theoretical result is observed empirically in Fisher and Lorie (1970). Consequently, implementation of normal policy will lead to performance that is similar to that of an undiversified portfolio.

FIGURE 4

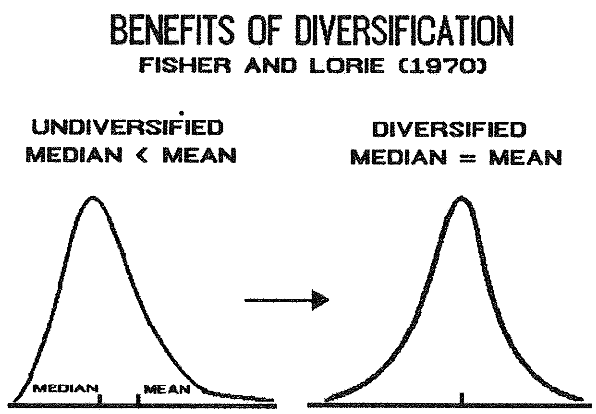

Alternatively, the multiperiod returns of a TOR strategy are far more "normal" than normal rebalancing. Using the data in the Tables and compounding shows that TOR multiperiod returns have lower probability of doing very well and substantially increased probability of performing as well as market averages. There is, however, an important exception to the traditional diversification analogy; TOR multiperiod returns are not really symmetric and time option rebalancing leads to an increased probability of doing very poorly.

These issues are further illustrated in Figure 5 below. Normal policy rebalancing implies a lognormal, right skewed distribution of portfolio returns over time. Time option rebalancing shifts the right and left tail to the left, and the median is shifted to the right and is approximately equal to the average.

SUMMARY: TIME OPTION REBALANCING BENEFITS

Time option rebalancing increases expected long-term return at the same level of risk as normal policy. The strategy is strongly consistent with contrarian investing -- when ahead, trim the sails; when behind, put sail to the wind. The results provide a framework for rationalizing the contrarian practices of many investment managers. The solutions help to quantify and expand intuition on the optimal application of contrarian asset allocation. Time option rebalancing avoids the "undiversification" effects of normal policy rebalancing. It requires only current information, plus a market environment assumption similar to that commonly used in investment policy studies. Rebalancing alternatives are either "no information" based and are therefore suboptimal or are "information" based and require assuming increased risk.

FIGURE 5

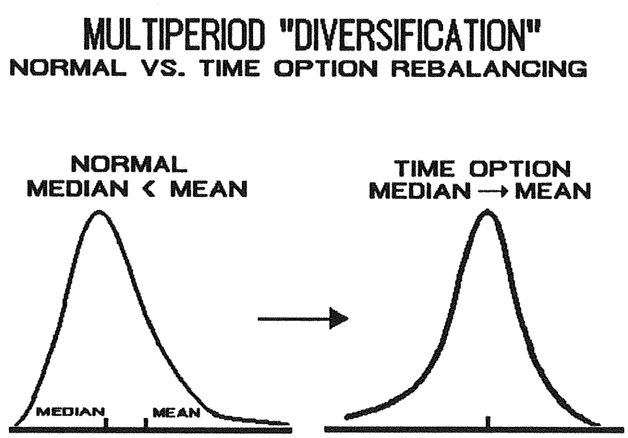

SUMMARY: TIME OPTION REBALANCING LIMITATIONS

Time option rebalancing increases the probability of poor performance relative to normal policy, but generally less than information based alternatives. Indeed, increasing the probability of inferior performance is probably inherent in any procedure that enhances median return over time. TOR will lead to more transactions costs than normal rebalancing, but generally less than information based alternatives. TOR requires careful specification of objectives and the planning horizon, discipline in implementing the strategy, careful performance monitoring and computational and analytic sophistication.

OPEN ISSUES

An in-depth analysis of the return enhancements likely under TOR are beyond the scope of the report. For the cases examined, return enhancements are small; the maximum observed was 40 basis points per period. Considerations should include the benefits of small increases in geometric mean return on terminal wealth over long time periods and the effect of more realistic assumptions.

Computed TOR results reflect the simplistic character of the assumptions. More realistic assumptions are likely to lead to less frequent large deviations from normal weights.

Practical implementation may require large scale problem solving. The nature of the process is beyond the scope of the report.

Hakansson efficiency is of interest in its own rights and as a intuitive guide for defining optimality. However, the underlying dynamic programming solution process can be used with other decision criteria providing a framework for understanding the characteristics of many definitions of optimality over time. The consequences of different criteria on time option rebalancing are beyond the scope of the report.

CONCLUSIONS

Normal policy rebalancing, implied by many investment policy asset allocation studies, leads to "undiversified" returns over the planning horizon. TOR does not ignore time remaining or investment experience in attempting to optimally meet fund objectives. The results are consistent with much intuition and investment practice. The additional demands it makes on the institution is analytic sophistication and increased attention to properly defining objectives. The benefits are that, if used with care, time option rebalancing will, on average, significantly enhance performance with little increase in investment risk.

FOOTNOTES

[1] Hakansson efficiency, and its antecedents (Markowitz, 1959, Chs. 6, 13; Latane, 1959) has been the subject of critiques by Samuelson and Merton (1974) and Merton and Samuelson

(1974), particularly with respect to using the criterion as a universal surrogate for expected utility maximization. The pros and cons of the controversy are beyond the scope of this report. Extended references, including discussion of practical implications, are given in Michaud (1981).

[2] Asymptotic property of functions of the product, assuming independent, identically distributed returns.

TECHNICAL NOTE

The dynamic programming algorithm is based on solving for the roots of a non linear system of M equations in M unknowns. Each equation is the partial derivative of the objective function with respect to the decision variables and a Lagrange variable.

The algorithm solver used in this report is the Minerr program from MathSoft's MathCAD. An earlier and much more limited version of the results were reported in Michaud and Monahan (1981) and solved with a generalized Newton-Raphson technique.

All reported computations were performed on a Northgate Elegance 1000 (25-MHz Intel 386 Microprocessor) and subsequently on a Compaq 386/25 both using 25-MHz Intel 387 Math Coprocessors.

TABLE 1
TIME OPTION ALLOCATION RELATIVE TO NORMAL*#
TWO ASSET, TWO PERIOD, TWO POINT RETURN DISTRIBUTION
$p = .5, R_o = 5\%$

R_m+	$x_o/x_f\%$	$x_w/x_f\%$	$x_l/x_f\%$
(20,0)	18	-51	43
(30,-10)	4	-21	21
(40,-20)	1	- 9	11

* Results reported for $x_f = .5$; results similar for policies x_f of interest ranging within the end points of the frontier

+ Two point risky asset return distribution (%)

See technical note for details of these and other computed results

Note: (30,-10) case is most consistent with historical annual market returns (mean = 10%, std. = 20%)

TABLE 2
TIME OPTION ALLOCATION RELATIVE TO NORMAL*
TWO ASSET, TWO PERIOD, THREE POINT RETURN DISTRIBUTION
Outcomes Equally Probable, R_o = 5%

R_m+	x_o/x_f%	x_w/x_f%	x_e/x_f%	x_l/x_f%
(20,10,0)+	18	-58	- 5	58
(30,10,-10)+	3	-18	1	22
(40,10,-20)+	0	- 2	0	3

* Results reported for x_f = .5; results similar for policies x_f of interest ranging within the end points of the frontier

+ Three point risky asset return distribution (%).

Note: End points of distribution increased so that mean and standard deviation of distribution comparable to Table 1.

TABLE 3
TIME OPTION ALLOCATION RELATIVE TO NORMAL*
TWO ASSET, THREE PERIOD, TWO POINT DISTRIBUTION
p = .5, R_o = 5%

R_m+	x_o/x_f%	x_w/x_f%	x_l/x_f%	x_{ww}/x_f%	x_{wl}/x_f%#	x_{ll}/x_f%
(20,0)	36	-39	64	-75	-27	99
(30,-10)	6	-14	20	-32	- 1	36
(40,-20)	1	- 4	6	-10	1	11

* Results reported for x_f = .5; results similar for policies x_f of interest ranging with the end points of the frontier

+ Two point risky asset return distribution (%)

$x_{wl} = x_{lw}$

REFERENCES

Brinson, G. and J. Deirmeier, 1985, "The Multiple Asset Investment Setting." Presented to: The Institute for Quantitative Investment Research in Finance, Spring.

Fisher, L. and J. Lorie, 1970, "Some Studies of Variability of Returns on Investments in Common Stocks," Journal of Business, April.

Hakansson, N., 1971a, "Capital Growth and the Mean-Variance Approach to Portfolio Selection," Journal of Financial and Quantitative Analysis, January.

-----, 1971b, "Multi-period Mean Variance Analysis: Toward a General Theory of Portfolio Choice", Journal of Finance, September.

Markowitz, H., 1959, Portfolio Selection: Efficient Diversification of Investments, Wiley, NY.

Merton, R. and P. Samuelson, 1974,

Michaud, R., 1981, "Risk and Long-Term Investment", Journal of Financial and Quantitative Analysis, June.

----- and J. Monahan, 1981, "Comparisons of Optimal vs. Stationary Investment Policies Over Time." Presented to: The Institute for Quantitative Research in Finance, Spring.

Samuelson, P. and R. Merton, 1974.