working copy

STATISTICAL MEAN-VARIANCE ASSET ALLOCATION

by

Richard O. Michaud, Ph.D., S.V.P.

Acadian Asset Management
260 Franklin Street
Boston, MA 02110
RMICHAUD@MCIMAIL.COM

Draft: September 29, 1994

This paper was prepared for the Institute for Quantitative Research in Finance (The Q Group), Fall Seminar, Lake Tahoe, Nevada, October 2-4, 1994.

ABSTRACT

Probably the single most important limitation of mean-variance optimization as a practical tool for investment management is the statistical instability of mean-variance optimized portfolios. Mean-variance optimizers often function as a chaotic investment decision system. The algorithm overuses statistically estimated information and magnifies the impact of estimation errors. It is not a simple matter of garbage in, garbage out, but, rather, a small amount of garbage in, a large amount of garbage out. The result is that optimized portfolios often have little, if any, reliable investment value. Indeed, research has shown that an equally weighted portfolio may often be substantially closer to true mean-variance optimality than an optimized portfolio. The failure of optimized portfolios to meet investment objectives in many cases has led a number of sophisticated institutional investors to abandon the technology for simpler alternatives and to an increased reliance on intuition and priors. The problems of mean-variance optimization are not easily solved with alternative risk measures or objective functions. They are endemic to many optimization procedures.

This report will focus on some key statistical characteristics of mean-variance optimization. The scope is limited to traditional mean-variance asset allocation. This area of application includes many investment management situations of practical interest. Methods for reducing the instability of the optimization process and enhancing its investment value are introduced and illustrated with historical data. The two major categories of enhancements are: Statistical estimation techniques and investment related priors. Such methods can reduce the impact of estimation errors, enhance the investment meaning of the results, provide an understanding of the degree of precision, and stabilize the optimization. In isolation, each procedure can be helpful; together they may have a substantial impact on enhancing the investment value of mean-variance asset allocation.

The problem of mean-variance optimization instability is ultimately one of over-fitting data. Assuming mean-variance efficiency is the appropriate investment objective, the procedure emerges in a fundamentally altered role: Mean-variance optimization as the basis for tests of efficiency rather than as a prescription for optimal portfolio construction. "Statistically efficient" asset allocations are unlikely to reflect error maximization characteristics. The stability of the new procedure is reflected in the additional investment benefit of a significantly reduced need to trade.

Markowitz (1959) mean-variance efficiency is the classic financial paradigm for efficiently allocating capital among risky assets. Given estimates of the mean, standard deviation and correlation for a set of assets, it provides precise prescriptions for optimal asset allocation. The Markowitz efficient frontier represents the collection of all mean-variance efficient portfolios in the sense that all other portfolios have less expected return for a given level of risk or, equivalently, more risk for a given level of expected return. Risk is defined as the variance of returns.

Mean-variance efficiency has been applied to many investment situations. An active international equity manager may want to find optimal asset allocations among international equity markets based on market index historical returns. A plan sponsor may want to find an optimal long-term investment policy for allocating among domestic and foreign bonds, equities and other asset classes. An active domestic equity manager may want to find the optimal equity portfolio based on forecasts of return and estimated risk models. In these cases, and many others, mean-variance optimization serves as a framework for optimally allocating capital. The paradigm is sufficiently flexible to consider various trading costs, constraints and levels of risk.

The focus of the report is traditional asset allocation for major asset classes. The number of risky assets in a study rarely exceed fifty, and typically is in the range of ten to twenty. Asset classes often considered include U. S. equities, corporate and government bonds, international equity and bond indices, real estate and venture capital. Historical data is often used to estimate the input parameters. The asset allocation problem can be contrasted to the typical equity portfolio optimization situation. Domestic equity portfolio optimizations typically consider a hundred to five hundred stocks; international equity portfolio optimizations may include a thousand or more stocks. Equity portfolios normally include the influence of a commercial factor risk model and proprietary return forecasts, substantially increasing the complexity of the analysis of the optimization. In contrast, the asset allocation problem reflects a purer framework for the analysis of mean-variance optimization.

We begin the report with a typical mean-variance asset allocation problem and use it as the basis for comparing alternatives and enhancements. Traditional objections to mean-variance optimization are described and reviewed. The central theme of the paper is then proposed: The most serious limitation of mean-variance optimization as a practical tool of investment management is the statistical instability of the optimization process, leading to optimized portfolios that often have little, if any, investment value. Various proposals for enhancing the investment value of mean-variance optimized portfolios are given. Such procedures stabilize the process and minimize the influence of statistical errors in the estimates by introducing investment relevant priors, more powerful and appropriate statistical estimation procedures, and methods for dealing directly with the statistical variability of the prescribed allocations. Each proposal is illustrated using the data in the original example. An alternative view of the practical investment value of mean-variance efficiency is presented: Mean-variance optimization is more reliably a statistical test for mean-variance efficiency than a prescription for finding optimal allocations. A procedure based on the notion of "statistical efficiency" is proposed for defining more reliably efficient asset allocations. The procedure also has the additional benefit of requiring less trading. An equivalence relationship between mean-variance optimization and a suitably defined constrained linear regression is noted that sheds light on both procedures. Finally, a summary of the results is provided.

CLASSICAL MEAN-VARIANCE ASSET ALLOCATION

1. An Example

Consider an active global asset manager allocating capital to the following eight major asset classes: U. S. stocks and government/corporate bonds, Euro bonds, and the Canadian, French, German, Japanese and United Kingdom equity markets. The optimization inputs are computed from 198 months of index total returns in U. S. dollars for all eight asset classes and for U. S. 30 day t-bills, for the historical period: January 1978 to June 1994. The efficient frontier solutions assume no short selling (all allocations are non-negative). A quadratic programming algorithm is used to find the optimal mean-variance efficient frontier asset allocations under the assumptions. The results of the analysis is given in Figure 1. The Markowitz mean-variance efficient frontier is plotted with "+" signs. The (monthly) means and standard deviations of the assets are given in Table 1 and plotted in the figure and labeled as indicated.

Since the Japanese market index had the highest average monthly return for the period, it is on the efficient frontier at the most northeast point of the curve. The minimum risk portfolio is nearly one hundred percent Euro bonds, with 0.84% average return and 1.41% standard deviation. Other points on the efficient frontier will lie between these two extremes. For example: The efficient frontier asset allocation with average return 1.56% and standard deviation 4.58% (about half way between the largest and smallest return efficient portfolio) is approximately 20% U. S. equities, 20% Euro bonds, 20% French stocks, 25% Japanese stocks, and 15% U. K. stocks.

Some results of interest include: U. S. bonds significantly underperformed all other assets and an efficient portfolio for its level of risk. Except for Japan, all equity markets significantly underperformed the efficient portfolios at their level of risk. In this period, it paid to be efficiently diversified.

2. Reference Portfolios and Portfolio Analysis

Mean-variance analysis can have more practical value if it includes reference portfolios. We define the three reference portfolios as given in Table 1:

TABLE 1
HISTORICAL DATA AND REFERENCE PORTFOLIOS (%)
JANUARY 1978-JUNE 1994

	U. S. STOCKS	U. S. BONDS	EURO BONDS	CANADA STOCKS	FRANCE STOCKS	GERMAN STOCKS	JAPAN STOCKS	U. K. STOCKS
MEAN	1.19	0.81	0.82	0.99	1.58	1.17	2.21	1.63
ST.DEV.	4.37	2.04	1.41	5.64	7.16	6.36	9.91	7.98
INDEX	30	0	0	5	10	10	35	10
CURRENT	25	20	5	5	10	0	20	15
EQUAL	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8

U. S. 30 Day T-Bill rate Mean = 0.60 St.Dev. = 0.24

Figure 2 provides the results of including the reference portfolios in the efficient frontier analysis. As shown in the figure, all three reference portfolios plot close to the efficient frontier. The near efficiency of the current portfolio may lead us to conclude that it is sufficiently close that it does not need to be rebalanced. However, compare the current portfolio's composition to that of the efficient portfolio nearest it: 15% U. S. stocks, 35% Euro bonds, 15% French stocks, 25% Japanese stocks, 10% U. K. stocks. The portfolios are not particularly similar. Also note that an equal weighted portfolio is nearly mean-variance efficient as well. Does the fact that very different portfolios may have nearly similar risk and return indicate that the apparent optimality of the current portfolio is more apparent than real? Such issues will be further investigated below.

3. Return Premium Efficient Frontiers

It is often convenient to use return premiums, instead of total returns, as the basis of the mean-variance analysis. The return premium is simply the return minus the (nominal) risk free rate. For monthly returns, the short term rate is usually defined as the U. S. t-bill 30 day rate of return. The mean and standard deviation over this period is given in Table 1. Return premiums provide a more direct and investment meaningful measure of return for investment in risky assets. Return premiums are also similar to real rates of return, which are defined as excess return with respect to the inflation rate. Real rates, and consequently return premiums, may be relatively more stable than total returns across time. Return premiums are also convenient for comparing the results of one time period by eliminating the effect of changing riskless rates. Figure 3 provides an asset allocation analysis for the data in Figure 2 when return premiums are used. The data in Figure 3 will be used as the basis of many of the subsequent examples.

TRADITIONAL CRITICISMS OF MEAN-VARIANCE ASSET ALLOCATION

A number of criticisms have been directed at mean-variance efficient frontier analysis. Many can be classified in the following four categories: 1) Alternative measures of risk;

2) Alternative utility functions; 3) Long- vs. short-term investment horizons; 4) Alternative procedures. Each are discussed separately below.

1. Alternative Measures of Risk

The variance or standard deviation of return measures variability above and below the mean. From an investment point of view, the variability of returns above the mean can hardly be considered "risk". One obvious risk measure alternative, discussed in Markowitz (1959), is the semi-variance or semi-standard deviation of return. In this case, only returns below the mean are included in the estimation of risk. Such an approach is similar to an optimization with an objective function based on "downside" risk. There are two problems with the semi-variance alternative. If the distribution of asset returns are approximately symmetric, both measures of risk lead to the same conclusion. Also, efficient frontier algorithms based on the semi-variance are not as well developed as they are for the variance.

Many other risk measures are also available. Some of the more important include the mean absolute deviation and range measures. The pros and cons of various alternatives depend critically on the assumed nature of the asset return distribution. Because historical equity index returns are often nearly symmetric, and because of the convenience and relatively high state of development of mean-variance based algorithms, many analysts prefer to use a mean-variance efficient frontier approach.

2. Alternative Utility Functions

Unless asset returns are normally distributed, a Markowitz mean-variance analysis is rigorously consistent only with quadratic utility. The many limitations of a quadratic utility function as a representative of investor behavior are well known. Numerous alternative functions have been proposed as the basis of more rational investor decision making. Such approaches optimize expected utility directly leading to asset allocations that may bear little resemblance to efficient frontier portfolios.

The problem with the utility function approach is that an investor's utility is seldom, if ever, known with any specificity. The lack of specificity can be a more daunting problem than it may appear. This is because a class of utility functions can have similar functional forms, perhaps differing in the value of one two parameters, yet reflecting a very wide spectrum of attitudes towards bearing risk and implying sharply different investment behavior (Rubinstein, 1973). As a practical matter, the problem of specifying with sufficient accuracy the appropriate utility function for a given investor appears to be no less a limitation than those posed by traditional mean-variance analysis.

On the other hand, a strong case for quadratic utility, not as an exact representation of investor behavior, but as a useful approximation to expected utility maximization, can be made. At a given point, reasonable (non-pathological) utility functions are often well approximated by a (quadratic) function of the mean and the variance. While the coefficients of the best approximating mean-variance function may vary at different points of the expected utility function, in many cases increased expected utility is consistent with more expected return and less variance. Consequently, mean-variance efficient portfolios may be useful as approximations for maximizing expected utility (see Levy and Markowitz, 1979).

The use or non-use of utility functions in investment analysis is often a dividing line between practitioners and academics. From a rigorous point of view, only the specification of a utility function will do for solving optimality. However, given the difficulty of estimating utility functions with sufficient precision, mean-variance efficiency is often the practical tool of choice.

3. Long- Versus Short-Term Investment Horizon

One significant limitation of Markowitz mean-variance efficiency is that it is formally a single-period investment tool. In practice, institutional investors, such as endowment funds and pension plans, are often interested in return over long investment horizons. Portfolios that are short-term optimal need not be optimal in the long run.

The geometric mean or compound return is the statistic of choice for summarizing the investment implications of long-term return over time. In his classic study, Markowitz (1959, ch. 6) examined the long-term or geometric mean return associated with mean-variance efficient frontier portfolios. Using an example he showed that geometric mean return did not always increase as the risk of efficient frontier portfolios increased. He found that, beyond a certain point, increasingly risky efficient portfolios led to a reduction in long term return. Taking Markowitz' result to a logical conclusion, Hakansson (1971a) astonished many financial economists by giving an example of a mean-variance efficient frontier with negative geometric mean returns at all points. This implied, in this case, that all mean-variance efficient frontier portfolios led to ruin with probability one over sufficiently long time periods. Clearly, optimal single-period decisions could lead to serious negative long-term investment consequences.

These negative results associated with mean-variance analysis can easily be remedied. As a practical matter, the Hakansson efficient frontier is far from typical. In addition, Markowitz analysis can easily be supplemented and the Hakansson objections largely avoided with the use of geometric mean return analysis. Numerous mean-variance approximations of the geometric mean are available (e.g., Young and Trent, 1969). From such analyses one can determine the subset of efficient frontier portfolios that are long-term (geometric mean) efficient. Consequently, if long-term return considerations are important, limit consideration of Markowitz efficiency to the subset of efficient frontier portfolios that are long-term geometric mean return efficient as well.

Objections to the geometric mean criteria have been raised on other grounds. In particular, one major controversy surrounded the proposal (Hakansson, 1971b) of using the (long-term) geometric mean as a surrogate for expected utility. The controversy cannot be repeated here.[1] The position taken here (as in Markowitz, 1977 and Michaud, 1981) is that the geometric mean can be useful as practical investment information and as a supplement to efficient frontier analysis.

4. Alternative Methods: Monte Carlo Financial Planning Studies

Monte Carlo simulation studies have often be used to evaluate the funding implications of investment return assumptions and portfolio decisions on the operation of the fund and the funding of various cash flow liabilities. For example, in the case of an endowment fund, a monte carlo simulation study can be useful for evaluating the consequences of a particular asset allocation to endowment spending and fund value over various time horizons based on assumptions of asset expected return and risk. By varying the portfolio composition, the relative benefits of various asset allocations can be evaluated and an optimal allocation chosen. The evaluation of the consequences of investment decisions in terms of their impact on plan funding objectives, cash flows and period-by-period operation of the fund over time is arguably a more financially meaningful benchmark than mean-variance analysis.

However, the benefits of monte carlo simulation studies for defining an optimal asset allocation are closely related to the issues of long- vs. short-term return in the previous section. The monte carlo simulated relative benefits of alternative asset allocations will generally simply reflect the mean-variance efficiency of the distribution of the N-period geometric mean (Michaud 1976). Michaud (1981) provides a number of analytical tools that can be used to forecast the mean and variance of the N-period geometric mean as a function of the single-period mean and variance. These results can be helpful for the design and analysis of monte carlo simulations. While monte carlo simulations are often useful as a financial planning tool, their value for asset allocation is largely a function of the implications of asset risk and return assumptions on the N-period geometric mean return, which can be anticipated with analytic methods.

THE FUNDAMENTAL LIMITATIONS OF MEAN-VARIANCE EFFICIENCY

For practical asset management, the most serious limitation of mean-variance analysis is probably its statistical instability. Mean-variance optimizers function as "error maximizers" or chaotic investment management systems by over using statistical information in the estimated parameters (Michaud 1989). It is not simply a matter of garbage-in, garbage-out, but, rather, a small amount of garbage-in, a lot of garbage-out.

Demonstrations of the magnitude of the statistical instability of mean-variance optimization, biases produced by mean-variance optimizers, and the likely investment irrelevance of mean-variance optimized portfolios has been given by Jobson and Korkie (1980, 1981). They showed, for example, that an equally weighted portfolio may often be significantly more optimal than unconstrained mean-variance optimized portfolios. Such results serve to rationalize the behavior of many institutional investors who have experienced the limitations of optimized portfolios and have voted with their feet by abandoning mean-variance optimization for simpler, less error prone, alternatives.

THE STATISTICAL EQUIVALENCE ~~EFFICIENT FRONTIER~~ Region

An analysis of the statistical characteristics of mean-variance optimization may usefully begin with a description of the variability in the procedure. Because a mean-variance efficient frontier is a computation based on statistically estimated parameters, it has a variance. While the true values of the input parameters are unknown, the variance of the optimization process can be measured ~~indirectly~~ using monte carlo simulation. Using the estimated input parameters as a basis, we can monte carlo simulate new optimization inputs and compute new efficient frontiers. If the simulated optimization inputs are the same as the originals, the resulting simulated efficient frontier is the same as the original and no variability is observed. Because the monte carlo simulated optimization estimates are likely to be different, the resulting efficient frontier with the original efficient frontier inputs will track below the original efficient frontier. By computing many monte carlo simulated efficient frontiers, the amount of variability implicit in the data and the efficient frontier procedure can be observed.

The statistical variability of the efficient frontier estimation process can be illustrated for the historical data and assumptions in Figure 3. An exact replication of the data sampling process is used in Figure 4 for each monte carlo simulated efficient frontier: optimization inputs are computed for 198 monte carlo simulated monthly return premiums for the eight asset classes from a distribution based on the original (198 months) estimated means, variances and correlations. Figure 4 displays the results of 100 monte carlo simulated efficient frontiers. Each simulated frontier is as likely to be efficient as the original mean-variance efficient frontier. Consequently, all the simulated portfolios are mean-variance equivalent and the area they occupy below the original efficient frontier can be described as the "statistical equivalence" region.[2] The size of the statistical equivalence region indicates the level of variability inherent in the data and the optimization process.

One issue of practical interest is the level of variability indicated by the size of the statistical equivalence region as a function of the number of sample periods. In many cases, smaller sample periods are of interest, since shorter time periods may represent more homogeneous epochs for the return generating process. Often some compromise must be made between using long time periods with variable economic and market conditions and shorter time periods with too few data points for reliable estimation. Figure 5 provides the results of estimating efficient frontier variability with 60 months of sample data. These results are directly comparable to Figure 4 except for the number of sample periods. Noteworthy in Figure 5 is the substantial expansion in the size of the statistical equivalence region and that a number of asset classes are within or on the edge when compared to Figure 4.

EFFICIENT-FRONTIER ANALYSIS AS A STATISTICAL PROCEDURE

Mean-variance optimization is normally perceived as a tool for constructing optimal efficient portfolios. However, due to the substantial statistical instability of the process, the analyst must come to terms with the likely investment limitations of mean-variance optimized portfolios.

An alternative approach, one with reduced possibility of investment irrelevance, is to use the efficient frontier statistical equivalence region as the basis of a statistical test of candidate portfolios. The approach is to select the appropriate optimization framework and estimation procedure and estimate the statistical equivalence region. In most cases of practical interest, a number of candidate portfolios are available in an investment organization. The test proceeds by determining the (after cost) position of the candidate portfolios relative to the statistical equivalence region. If the choice is between two portfolios, one within the statistical equivalence region and one outside, the portfolio within is preferred. Alternatively, for a given portfolio, the statistical equivalence region can be used to determine if it should be rebalanced. Heuristics guided by investment priors and intuition are unlikely to be affected by statistical estimation errors and investment irrelevant error maximized portfolios can be avoided. Some portfolio heuristic procedures are discussed in the final section of this report.

The test procedure described in the previous paragraph is not statistically rigorous. A formal test requires an understanding of the distribution of the statistical efficient frontiers and the optimal partitioning of the statistical equivalence region based on type 1 error assumptions. More formal methods for inference can be based on Jobson and Korkie (1985) and are discussed in the final section. The informal procedure is useful for pedagogical and illustrative purposes.

Jobson (1991) provides an analytically derived estimate of the "sample acceptance region" when portfolios are not sign constrained. As a consequence of the portfolio construction differences, Jobson's sample acceptance region is significantly different from the shape of the statistical equivalence region in Figures 4 and 5: Jobson's region is open ended at the high risk return area of the efficient frontier; the Figure 4 and 5 equivalence regions are small at both end points.

The results in Figures 4 and 5 show that all three reference portfolios are well within the statistical equivalence region. On the other hand, in the case of Figure 4, except for Japan, the remaining assets can not be considered statistically efficient. Note however that, in the case of Figure 5, many of the assets are within or near the statistical equivalence region. This is consistent with the notion that the test region for Figure 5 has less power to distinguish non-efficient assets and portfolios than Figure 4 because it is based on a smaller sample size.

In both Figure 4 and 5, the current portfolio is well within the region of statistical equivalence. Accordingly, we may conclude that it is "statistically efficient" and does not need to be rebalanced. One important investment consequence of statistical equivalence region analysis is the likelihood of a significant reduction in the need for trading and portfolio turnover. Consequently, the procedure can have an important practical impact on enhancing performance.

ADMISSIBLE ESTIMATORS AND EMPIRICAL BAYES-STEIN ESTIMATION.

Admissibility is a minimal condition for an estimator. An estimator is admissible if no other estimator is uniformly better for a given loss function. Stein (1955) has shown that the classical method for estimating the means as inputs for optimization is not admissible. Consequently, there are uniformly better methods for estimating optimization inputs than those in common use by many practitioners and financial economists. Admissible estimators of the mean "shrink" the estimated returns back to a global prior. From an investment point of view, admissible estimators reduce the amount of instability in the optimization process by reducing the amount of statistical noise in the inputs.

A number of admissible estimation procedures for mean-variance optimization have been proposed and tested (e.g., Brown, 1976, Jorion, 1986). Many of these methods address shrinkage of the mean vector, leaving the covariance matrix unchanged. One consequence of this approach is that, while the shape of the efficient frontier curve may change, and the investor's optimal efficient portfolio from an expected utility point of view may change, the efficient frontier portfolios may remain the same.[3] Frost and Savarino (1986) have proposed an empirical bayes-stein procedure which shrinks both the mean vector and the covariance matrix. They use an informative prior that asserts that the more the sample estimates differ from the average, the more likely there is estimation error. Their method shrinks the estimated parameters to the common sample value as a function of the extent those characteristics differ from the average for all stocks.

Figure 6 illustrates the results of using a Frost-Savarino Bayes-Stein Estimator for the data in Figures 3 and 4. Comparing Figure 6 to 3 we find that average return for the efficient frontier and the assets lie in a narrower range than in the original data. This is because the Bayes-Stein process shrinks the sample parameter values to the sample average prior. When compared to Figure 4, the statistical equivalence region is slightly more compact reflecting the shrinkage in the parameter estimates and the reduced amount of variability in the inputs. One consequence of Frost-Savarino Bayes-Stein estimation is that assets and portfolios are more likely to lie in the statistical equivalence region, reducing the need for portfolio turnover and trading.

BENCHMARK OPTIMIZATION AND PRIORS

Benchmark optimization methods are among the most powerful for reducing the instability of mean-variance optimization and for enhancing the investment value of the solutions (see e.g., Michaud, 1989c). The increase in reliability is analogous to that of classical statistical procedures when a reliable bayesian prior is introduced. In the following, three benchmark methods are discussed: 1) Implied Return; 2) Index-Relative Return; 3) Economic Liabilities. Each method restructures the optimization process by the introduction of a prior in the form of a portfolio that is defined to be efficient or no-information optimal. Imposing the prior may lead to a more investment relevant definition of return. The prior may or may not imply redefining risk. The definition of the prior and the formulation of the optimization will also depend on investment objectives.

IMPLIED RETURN EFFICIENT FRONTIER

The implied return procedure assumes that a given portfolio or index is efficient. The prior is imposed on the optimization by solving for the "implied" returns given the efficient index and the sample covariance matrix. The procedure can be described as "backward" optimization and is given in Fisher (1975) and Sharpe (1974). The implied return efficient frontier can be computed based on the implied returns and the given covariance matrix. The implied returns are interpretable as return premiums. Note that the implied return procedure leaves risk unchanged. By imposing the prior, the optimization has more investment relevant information to define optimal asset allocations. As with all benchmark methods, an optimization with additional reliable information is less likely to be unstable and reflect statistical estimate errors, and more likely to provide investment meaningful results.

The procedure requires the identification of an efficient portfolio. One natural approach is to assume that the manager's benchmark index is mean-variance efficient. For practical purposes, a broad based market index may be sufficiently representative of investor interest to be assumed approximately mean-variance efficient. Alternatively, for active managers, an index may be efficient by definition since it is the benchmark used for performance evaluation and for defining (active) risk and return.

In Black and Litterman (1992), the primary motivation for the procedure is that the efficiency of the prior portfolio or index is a necessary condition for the optimization to be well defined. It is also of interest to note that the implied return procedure can be motivated from a more practical point of view. Historical average returns are highly period dependent and unstable when compared to covariance estimates. Mismatches between the relative reliability of input estimates can be a significant source of the instability of the optimization process (Michaud 1989a). The backward optimization procedure formulates return estimates that are consistent with the information level of the sample covariance matrix.

Figure 7 illustrates the result of this procedure for the data in Figure 3. Note that the index portfolio is now on the efficient frontier. It is, by construction, the maximum Sharpe ratio portfolio, as indicated by the tangent line drawn from the origin. Comparing the results with Figure 3 shows that the implied return procedure fundamentally alters the relative relationship of many assets and portfolios with respect to each other and to the efficient portfolio set.

The implied return efficient frontier in Figure 7 has unintuitive investment characteristics. The estimated implied returns range from 0% to 50%, which are inconsistent with the scale of the standard deviation of monthly returns on the horizontal axis. However, implied returns are unique up to a linear transformation. Consequently, they can be linearly rescaled to be formally consistent with monthly data without changing the composition of the efficient frontier or relative relationships. In Figure 8, the implied returns are linearly scaled so that the maximum and minimum values are the same as the sample average monthly return premiums. Rescaling is a convenience that allows useful comparisons with risk estimates and trading costs.

Figure 9 presents an estimate of the statistical equivalence region based on the scaled implied returns. The results can be directly compared to Figure 4. As can be seen from inspection, the impact of the implied return procedure is to make the statistical equivalence region more compact and more inclusive. If we assume that assets are, by themselves, unlikely to be efficient, the implied return procedure appears to have little power to reject mean-variance efficiency. Consequently, the implied return efficient frontier may not often be the framework of choice for analyzing the efficiency of the total return and risk of portfolios.

MIXED ESTIMATION AND ACTIVE RETURN FORECASTS

As with all benchmark procedures, the implied return efficient frontier provides a useful framework for including views of active return. This is because, in the absence of an active return forecast, the optimal portfolio is the index. What remains is to specify a procedure for rigorously integrating active views into an asset allocation.

In many cases, managers simply add their active return forecasts historical returns. The limitations of this naive procedure are substantial. In many cases, active forecasts are not equally reliable. In addition, the weighting given to historical data relative to the forecasts should be carefully considered. Given our discussions on the instability of mean-variance optimization, it would be surprising that the naive procedure is stable and reliable and provides consistent superior investment performance.

A rigorous procedure for mixing forecast views and uncertainty with historical data in the context of linear regression is given in Theil (1971). Black and Litterman (1992) provide a slightly revised formula for mean-variance optimization. The data in Table 2 illustrates the mixed estimation procedure for the historical data in Figure 3. The first line in Table 2 presents an analyst's forecasts of monthly index excess return or alpha: 0.2% for U. S. equities, 0.1% for France and -0.2% for Japan. The scaled implied return premiums from

Figure 8 are given in line two. The analyst's forecast alpha standard error is given in the third line. The resulting Theil-Black-Litterman mixed estimate excess returns are given in line four.

Table 2
Mixed Estimation Excess Returns
Based on Prior Monthly Alphas, Scaled Implied Returns, Std Errors (%)

	U. S. Stocks	U. S. Bonds	Euro Bonds	Canada Stocks	France Stocks	Germany Stocks	Japan Stocks	U. K. Stocks
Alpha	0.2	0	0	0	0.1	0	-0.2	0
Implied	0.46	0.08	0.06	0.57	0.71	0.62	1.47	0.72
Std. Er.	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Mixed	0.06	0	0	0.03	0.07	0	-0.14	0

Figure 10 illustrates how the mean-variance efficient frontier changes based on the mixed estimation data in Table 2 and the implied return efficient frontier in Figure 8. The "o" indicates the Figure 8 efficient frontier and asset risk and return estimates; the "+" denotes the mixed estimate frontier and asset estimates based on Table 2. The results show that the right arm of the mixed estimate efficient frontier rotates in a southeast direction about the index to accommodate the changes in expected return premiums. Since the index-weighted sum of active returns must be zero, the index is a fixed efficient point.

It should be noted that the mixed estimation procedure described in Table 2 and Figure 8 is not a particularly robust. Small changes in assumptions can lead to significantly different optimal solutions. The instability may be due in part to the fact that the implied return procedure does not change the definition of risk. In many cases the impact of the value of parameters required in Theil-Black-Litterman mixed estimation are difficult to anticipate and require considerable care if useful solutions are to result. The instability of the mixed estimation frontier, and its statistical significance, can be determined by computing the statistical equivalence region.

INDEX-RELATIVE EFFICIENT FRONTIER

For active asset management, index-relative optimization may be a valuable alternative to the implied return framework. In this case, the no active risk portfolio is defined as the index or benchmark. The efficiency of the index is imbedded in the optimization process by defining the index as a negative asset and estimating optimization parameters from the residual or index-relative returns for each asset in each time period.

Unlike implied return efficiency, index-relative optimization redefines the notion of both return and risk. Index-relative also requires estimating the risk of an asset in a different manner. If an asset has a 30% weight in the benchmark, an asset allocation consisting of solely the asset implies a 70% active weight with respect to the index. Accordingly, some convention must be used to allocate the 70% underweight in the remaining assets when computing individual asset returns and risk. The natural approach is to use the negative of the index weights, which sum precisely to -70%, which is the convention in this report. The optimization procedure must also be defined to reflect such considerations.

Figure 11 illustrates the index-relative efficient frontier and asset returns and risk for the data in Figures 3 or 1 and the definition of the index in Table 1. Comparing Figure 11 to Figures 1 or 3 and 8 illustrates the significant differences between the three optimization frameworks. In particular, Figure 11 indicates that many assets and reference portfolios may be relatively less active risk-return efficient than return premium or implied return efficient.

The statistical equivalence region in Figure 12 when compared to Figure 4 and 9 further illustrates the significant differences between the three approaches to asset allocation. In the case of index-relative return, the statistical equivalence region is more compact than in Figure 4 and less inclusive than in Figure 4 or 9. In particular, note that, except for the index, which is at the zero risk and return point, and Japan, the highest return point, no other portfolio or asset class is within the estimated statistical equivalence region. The results suggest that the index-relative efficient frontier equivalence region may have more power to reject the lack of efficiency in portfolios and assets than the implied return or return premium efficient frontiers. Figure 13 further illustrates the variability inherent in index-relative optimization. The increased volatility of the efficient frontiers when compared to Figure 12 indicates the significant reduction in the power of the procedure when smaller samples sizes are employed.

Finally, it should be noted that the mixed estimation procedure can be applied to the Figure 11 index-relative efficient frontier in the same way that it was with implied returns. This may be a very natural framework for making active asset allocation decisions. One key mixed estimation decision is the appropriate balance of historical with forecast active risk and return information.

ECONOMIC LIABILITY EFFICIENT FRONTIER

An alternative framework to the implied return efficient frontier for defining an optimal investment policy is benchmark optimization based on economic liabilities. Institutional investors, such as corporate pension plan sponsors and endowment fund trustees, are often interested in defining an optimal long-term investment policy. This is because investment policy is often considered the single most important investment decision. The investment objective in this case is not to beat an index but to best allocate assets in the light of long-term funding of liabilities.

Economic liability optimization is, in concept, quite simple (Michaud 1989c). It consists of defining the benchmark in terms of a model of the economic (non-actuarial) risk and return characteristics of fund liabilities. The benchmark liability may be a function of the returns of various financial assets and other economic and financial factors in each time period. As in the case of index-relative returns, the economic liability is introduced into the optimization as a negative asset and liability-relative returns are used to compute optimization input parameters.

The success of the economic liability optimization depends critically on defining an economically relevant model of fund liabilities. The primary concern is to capture how assets and liabilities interact and change dynamically in time. The key to the model building process is a sufficiently comprehensive understanding of the economic risk characteristics of the liabilities.

Endowment fund investment policy can provide a simple example of the economic liability procedure. An important endowment funding objective may be to maintain a stable level of purchasing power over time. One simple benchmark approach is to model the liability with the inflation rate as a convenient surrogate for changes in purchasing power. An economic liability efficient frontier provides optimal asset allocations for meeting the funding objective.

For defined benefit pension plans, the liability modeling process is likely to be more complex. One popular approach is to treat pension liabilities as interest rate sensitive. This is because the accrued or plan termination liabilities have essentially the same financial characteristics as a suitably defined portfolio of bonds. However, for an ongoing firm, plan termination liabilities do not include the valuation of all promised benefits. The present value of all current and future promised benefits, the "total actuarial liability", includes the additional component of future promised but unvested benefits. These additional "variable" liabilities can be substantial: For a typical final average pay plan, the size of promised benefits has been estimated at roughly 70% of termination liabilities (see Michaud, 1989b).

If plan termination is not a significant consideration, total pension liability may be the primary funding objective. Variable liabilities may have sharply different economic characteristics from vested liabilities. In particular, variable liabilities are not always interest rate sensitive. Variable liabilities consist of unvested pension benefits that depend on withdrawal rates, final average pay and other variables. Their risk characteristics depend on the business risks of the firm, corporate objectives, and other considerations, many of which may not be related to interest rate risks. Michaud (1989b) gives an example where variable pension liabilities may be highly correlated with equity market returns. Modeling pension liabilities with equities can have a fundamental impact on the asset allocation process. In particular, a well diversified equity portfolio may be the *low risk* asset allocation of choice for funding variable liabilities.

As a very simple example of the benchmark liability procedure for defined benefit pension plans, assume that a firm's total liability consists of 60% plan termination liabilities and 40% variable liabilities and that the plan is fully funded.[4] In addition, assume that plan termination liabilities can be modeled with a portfolio of government and corporate bonds and variables liabilities are highly correlated with U. S. equities.

Figure 14 illustrates the result of using this simple model of pension economic liability.[5] These results can be compared to those in Figures 3 and 7. While it is not surprising that U. S. stocks and bonds are closer to the efficient frontier, it is less anticipatable that asset classes like French and U. K. equities are also relatively closer to the efficient frontier.

Figure 15 estimates the variability of the economic liability efficient frontier. Comparing the statistical equivalence region to Figure 4 suggests that the power of the procedure is no less and may be more relevant to the objective of defining investment policy in the context of the assumed risk characteristics of the liabilities.

As a practical matter, the mixed estimation procedure is likely to be important in developing the economic liability efficient frontier. This is because, few investment policy studies will want to consider only historical data in developing a recommended policy. The "active" views may more accurately reflect adjustments to long-term historical data that may be indicative of relatively permanent changes in the underlying economy and market structure. In this case, mixed estimation is an appropriate procedure for modifying the long-term historical data.

STATISTICAL INFERENCE AND CONSTRAINED LINEAR REGRESSION

It is of interest to observe that mean-variance optimization is equivalent to a suitably defined constrained linear regression.[6] Intuitively, both procedures have the same constraints and variance. Except for the definition of the dependent variable, the mean-variance optimization equivalent regression is a stochastic independent variable constrained regression. The variability of the estimation process comes solely from the stochastic character of the independent variables.

The equivalence relationship can provide a fresh perspective on linear regression and mean-variance optimization. For example, the instability associated with mean-variance optimization is inherited by the stochastic constrained linear regression procedure. However, in general, the estimation of a constrained linear regression with a non-stochastic dependent variable has less variability than a corresponding unconstrained regression with a stochastic dependent variable. Consequently, the instability of mean-variance optimization is of a lower order than stochastic independent variable regression estimation, all other things the same.

The constrained linear regression can also shed light on the mean-variance optimization. Many linear regression statistics can be computed and may be useful for inference. In particular, under reasonable assumptions,[7] the usual test for the value of coefficients has a t-distribution with N-K+q degrees of freedom, where N is the number of sample periods, K is the number of variables and q=2, the number of constraints. The results may be used for a better understanding of an optimized portfolio and may serve as the basis for heuristics in guiding the construction of candidate optimal portfolios.

The following data in Table 3 illustrates these concepts with an analysis of the 9th and 12th least risky efficient portfolios in Figure 3 based on the t-statistics of the corresponding mean-variance equivalent constrained linear regression. The efficient allocations for the indicated efficient portfolios are given in lines one and five in the table. For these portfolios, the sign constrained efficient frontier solutions did not include the three assets: U. S. bonds, Canadian and German stocks. From Figure 3 note that the current portfolio and equal weighted portfolio are close to the 9th efficient portfolio, while the index portfolio is close to the 12th efficient portfolio. Line two displays the t-statistics for the 9th efficient portfolio. Line three and four display the t-statistics of the coefficients of the 9th portfolio relative to either the current or equal weighted portfolio weights. Line six provides the t-statistics for the 12th portfolio relative to the index.

Table 3
Mean-Variance Efficient Portfolio T-Statistics
Return Premium Efficient Frontier Portfolios (9th and 12th)

	U. S. Stocks	U. S. Bonds	Euro Bonds	Canada Stocks	France Stocks	Germany Stocks	Japan Stocks	U. K. Stocks
Eff-9th %	15	0	35	0	15	0	25	10
t-stat	2.0	0	7.0	0	3.2	0	9.7	3.4
t-current	-1.4	0	6.0	0	1.0	0	1.4	-0.6
t-equal	0.3	0	4.5	0	0.4	0	4.5	0.1
Eff-12th %	20	0	10	0	20	0	30	20
t-index	-0.9	0	1.5	0	1.8	0	-1.0	1.4

Figure 3 indicates that the structure of the current and equal weighted portfolios should be similar to the 9th efficient portfolio, the index similar to the 12th portfolio. These hypotheses are largely borne out in the t-statistics of the coefficients in Table 3. In particular, only one asset weight, Euro bonds, is significantly different from the current portfolio; the Euro bond and Japanese stock weight are the only ones significantly different from the equal weighted portfolio; no asset weight in the 12th portfolio is significantly different from the index portfolio. The t-statistics are useful as portfolio construction heuristics, guiding the analyst to the asset weights with the most impact on mean-variance efficiency. It should be noted that the proper interpretation of a t-statistic in the context of a mean-variance optimization is properly interpretable only with respect to a multivariate context. For example, a statistically significant t-value implies the significance of the portfolio weight given that the other (N-1) assets are present. While t-statistic analysis may be useful for understanding the micro structure of portfolio efficiency, F-tests may be useful for estimating the overall efficiency of a given portfolio.

SUMMARY AND CONCLUSION

The most serious limitation of mean-variance optimization as a practical tool of investment management is probably the statistical instability of the process. Because a mean-variance optimizer will often maximize the use of statistical errors in parameter estimates, a mean-variance optimized asset allocation may often be investment irrelevant. A consideration of the statistical characteristics of the mean-variance optimization process leads to a number of procedures for dealing with optimization stability.

Fundamentally, mean-variance optimization is a forecasting process and optimization errors the result of over-fitting data. The in-sample instability and out-of-sample inefficiency of mean--variance optimizations are analogous to the better known behavior of stepwise regression and other in-sample optimized linear regression forecasts. The forecast errors created are not indicative of any fundamental flaw in mean-variance optimization but the manner that the procedure has been applied.

A "statistical equivalence" region describes the variability inherent in mean-variance optimization and leads to an alternative approach to optimal asset allocation: as a statistical test of the efficiency of candidate portfolios. In many practical investment situations, this is all that is required from mean-variance efficiency analysis. This "statistical" asset allocation approach, that includes techniques for enhancing input estimation and investment relevance, provides a practical, more reliable and effective framework for defining optimal asset allocations.

FOOTNOTES

[1] The literature on the geometric mean portfolio selection criteria is extensive. It includes Merton and Samuelson (1974) and Samuelson and Merton (1974) on the one hand and Markowitz (1959, ch. 6, 1977), Latane (1959), Hakansson (1971a, 1971b, 1977, 1979), Hakansson and Miller (1975) on the other. Michaud (1981) attempts to summarize the issues from the point of view of a practitioner. The extensive references in the cited papers can be used to further explore the issues.

[2] The concept of the "statistical equivalent" region, discussed in Michaud (1989), has important antecedents in the work of Jobson and Korkie (1981). Also, see Jobson (1991) and Jorion (1992).

[3] Jobson (1994) proposes an alternative bayesian estimator for the optimization parameters that has many attractive properties.

[4] See Michaud 1989c for a benchmark optimization procedure when plans are not fully funded.

[5] It should be emphasized that the model presented is only a very simple approximation to economic (as opposed to actuarial) pension liability modeling and is not relevant for many firms. To be successful, the process will often require extensive analysis of the business risks of the firm, corporate objectives and other financial, economic, as well as actuarial, considerations.

1 This is a simple result of Taylor's Theorem discussed in Freshman calculus.

[6] Assume a sum to one coefficient constrained mean-variance optimized portfolio with a given mean return without positivity constraints (as noted below, this condition can be removed) on the coefficients. Define the following constrained linear regression using the same sample data as the independent variables: a) sum to one coefficient constraint; b) mean return constraint; c) dependent variable equal to zero; d) sample data with the mean of the independent variables removed. The constrained linear regression leads to the same values for the coefficients as the mean-variance optimization and conversely. Note that condition of non-positivity constraints in the mean-variance optimization can be removed. When positivity constraints are included in the optimization, the net effect is that some variables ~~in the original set~~ may be excluded from the mean-variance efficient set for a given mean. By defining the constrained linear regression so that only the variables in the efficient set at a point are used, all previous results apply. The regression data in Table 3 were computed using this result.

[7] The properties for inference include the standard linear regression assumptions generalized for stochastic multivariate normal variables and constrained linear regression (Jobson and Korkie 1985; see also Theil 1971, ch. 6, Green 1990, ch. 10).

REFERENCES

Black, F. and R. Litterman, 1992, "Global portfolio optimization." Financial Analysts Journal, September/October.

Brown, S., 1976, "Optimal portfolio choice under uncertainty." Unpublished Ph.D. diss., Univ. of Chicago.

Chopra, V., C. Hensel and A. Turner, 1993, "Managing mean-variance inputs: Returns from alternative global investment strategies in the 1980s." Management Science, pp. 835-845.

Fisher, L., 1975, "Using modern portfolio theory to maintain an efficiently diversified portfolio." Financial Analysts Journal, May/June.

Frost, P. and J. Savarino, 1986, "An empirical bayes approach to efficient portfolio selection." Journal of Financial and Quantitative Analysis, pp. 293-305.

Green, W., 1990, Econometric Analysis. New York: Macmillan.

Jobson, J. D., 1991, "Confidence regions for the mean-variance efficient set: An alternative approach to estimation risk." Review of Quantitative Finance and Accounting, pp. 235-257.

....., 1994, "Estimating the mean-variance efficient frontier: The Markowitz criterion is not enough." Presented to The Q Group, Fall Seminar.

. ....., and Bob Korkie, 1980, "Estimation for Markowitz efficient portfolios." Journal of the American Statistical Association, September.

...., and ....., 1981, "Putting Markowitz theory to work." Journal of Portfolio Management, Summer.

...., and ....., 1985, "Some tests of linear asset pricing with multivariate normality." Canadian Journal of Administrative Sciences.

Jorion, P. 1986, "Bayes-Stein estimation for portfolio analysis." Journal of Financial and Quantitative Analysis, pp. 279-292.

...., 1992, "Portfolio optimization in practice." Financial Analysts Journal, January/February.

....., 1994, "Estimation risk in portfolio selection." Presented to The Q Group, Fall Seminar.

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.

....., 1977, "Comment on Merton and Samuelson." In Risk and Return in Finance, I. Friend and J. Bicksler, eds. Ballinger: Cambridge, Mass.

....., 1979, "A characterization of optimal multi-period portfolio policies." In Portfolio Theory: 25 Years After, N. Elton and M. Gruber, eds., North Holland: New York.

....., and B. Miller, 1975, "Compound return mean-variance efficient portfolios never risk ruin." Management Science, December.

Latane, H., 1959, "Criteria for choice among risky ventures." Journal of Political Economy, pp. 144-155.

Levy, H. and H. M. Markowitz, 1979, "Approximating expected utility by a function of the mean and variance." American Economic Review, pp. 308-317.

Markowitz, H., 1959, Portfolio selection: Efficient diversification of investments. New York: John Wiley.

....., 1977, "Investment for the long run." In Risk and Return in Finance, I. Friend and J. Bicksler, eds. Ballinger: Cambridge, Mass.

Merton, R. and P. Samuelson, 1974, "Fallacy of the log-normal approximation to optimal portfolio decision making over many periods." Journal of Financial Economics, March.

Michaud, R. O., 1976, "Pension fund financial planning." Presented to the Q group Fall Seminar.

....., 1981, "Risk policy and long-term investment." Journal of Financial and Quantitative Analysis, June.

....., 1989a, "The Markowitz optimization enigma: Is optimized optimal?" Financial Analysts Journal, January/February.

......, 1989b, "Economic surplus and pension asset management." Pension Executive Review: Merrill Lynch, January.

....., 1989c, "Pension policy and benchmark optimization." Investment Management Review, March/April.

Rubinstein, M., 1973, "A comparative statics analysis of risk premiums." Journal of Business, pp. 602-615.

Samuelson, P. and R. Merton, 1974, "Generalized mean variance tradeoffs for best perturbation corrections to approximate portfolio decisions." Journal of Finance, March.

Sharpe, W., 1974, "Imputing expected security returns from portfolio composition." Journal of Financial and Quantitative Analysis, pp. 463-472.

Stein, C., 1955, "Inadmissibility of the usual estimator of the mean of a multivariate normal distribution." Proceedings of the 3rd Berkeley Symposium on Probability and Statistics. Berkeley: Univ. of California Press.

Theil, H., 1971, Principles of Econometrics. New York: John Wiley.

White, H., 1984, Asymptotic Theory for Econometricians. New York: Academic Press.

Young, W., and R. Trent, 1969, "Geometric mean approximations of individual security and portfolio performance." Journal of Financial and Quantitative Analysis, pp. 179-199.

FIGURE 1
CLASSICAL MEAN-VARIANCE EFFICIENT FRONTIER
January 1978-June 1994 Monthly Data

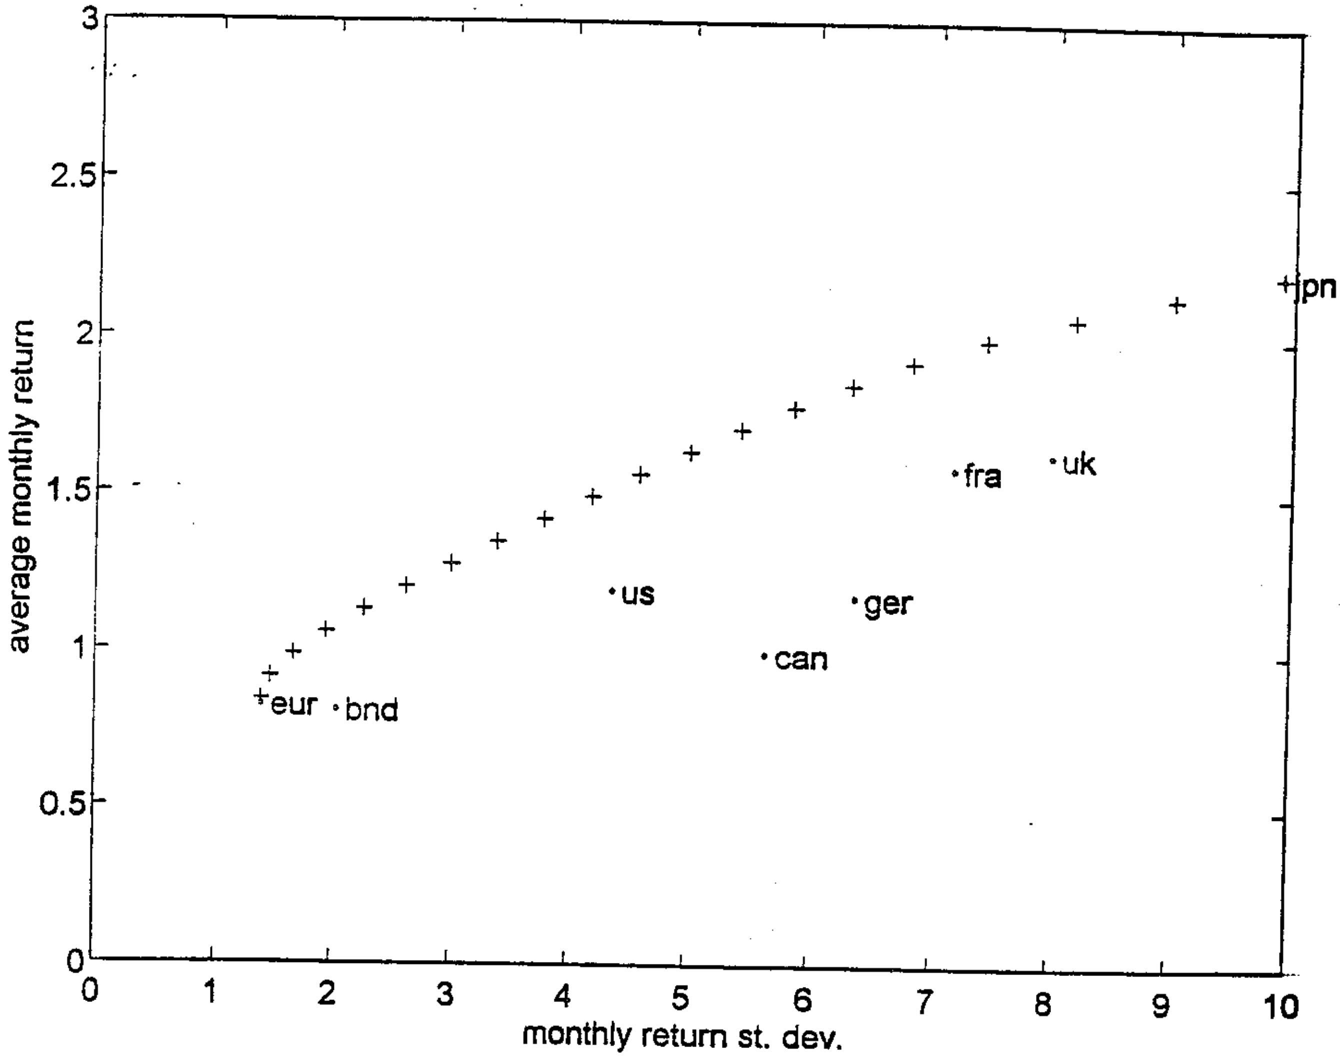

FIGURE 2
MEAN-VARIANCE EFFICIENT FRONTIER AND PORTFOLIO ANALYSIS
January 1978-June 1994 Monthly Data

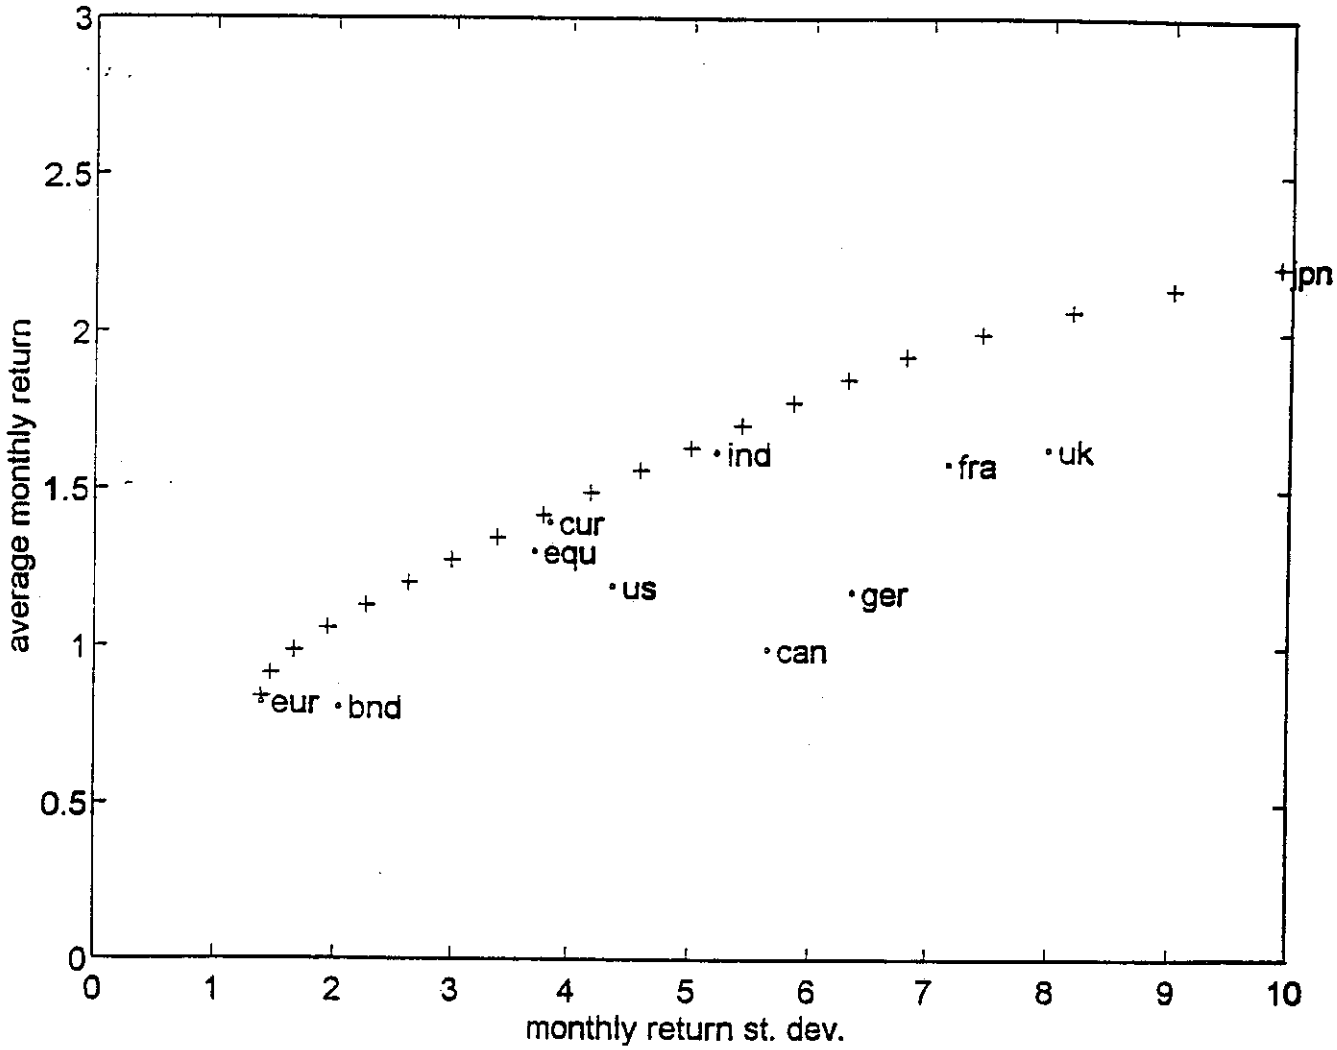

FIGURE 3
RETURN PREMIUM EFFICIENT FRONTIER
January 1978-June 1994 Monthly Data

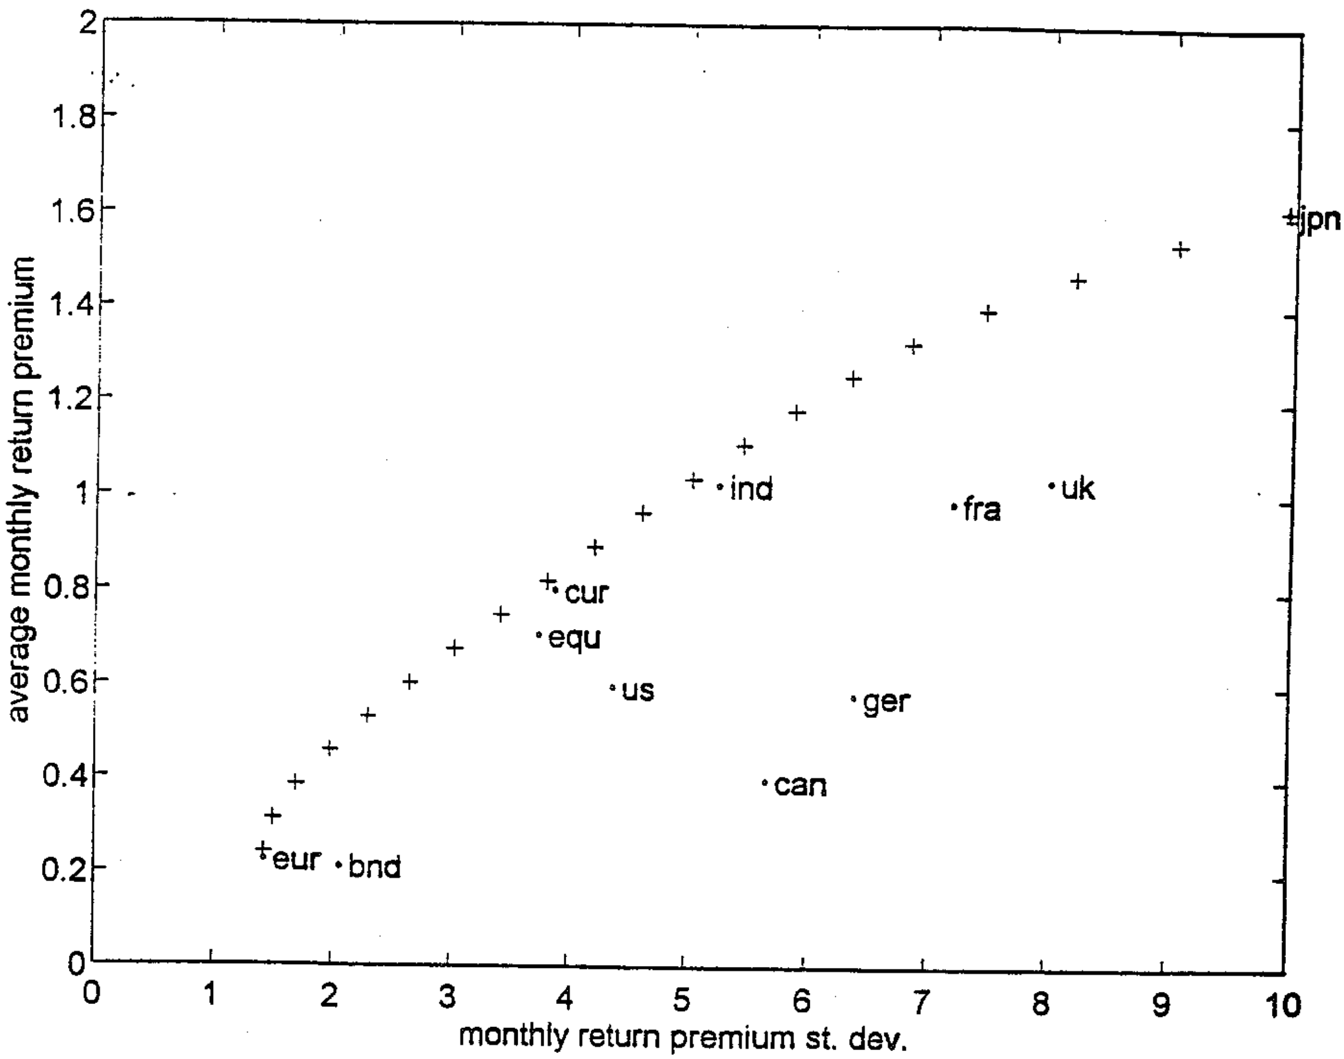

FIGURE 4
STATISTICAL EQUIVALENCE EFFICIENT FRONTIERS
100 SIMULATIONS, 198 TIME PERIODS
January 1978-June 1994 Monthly Data

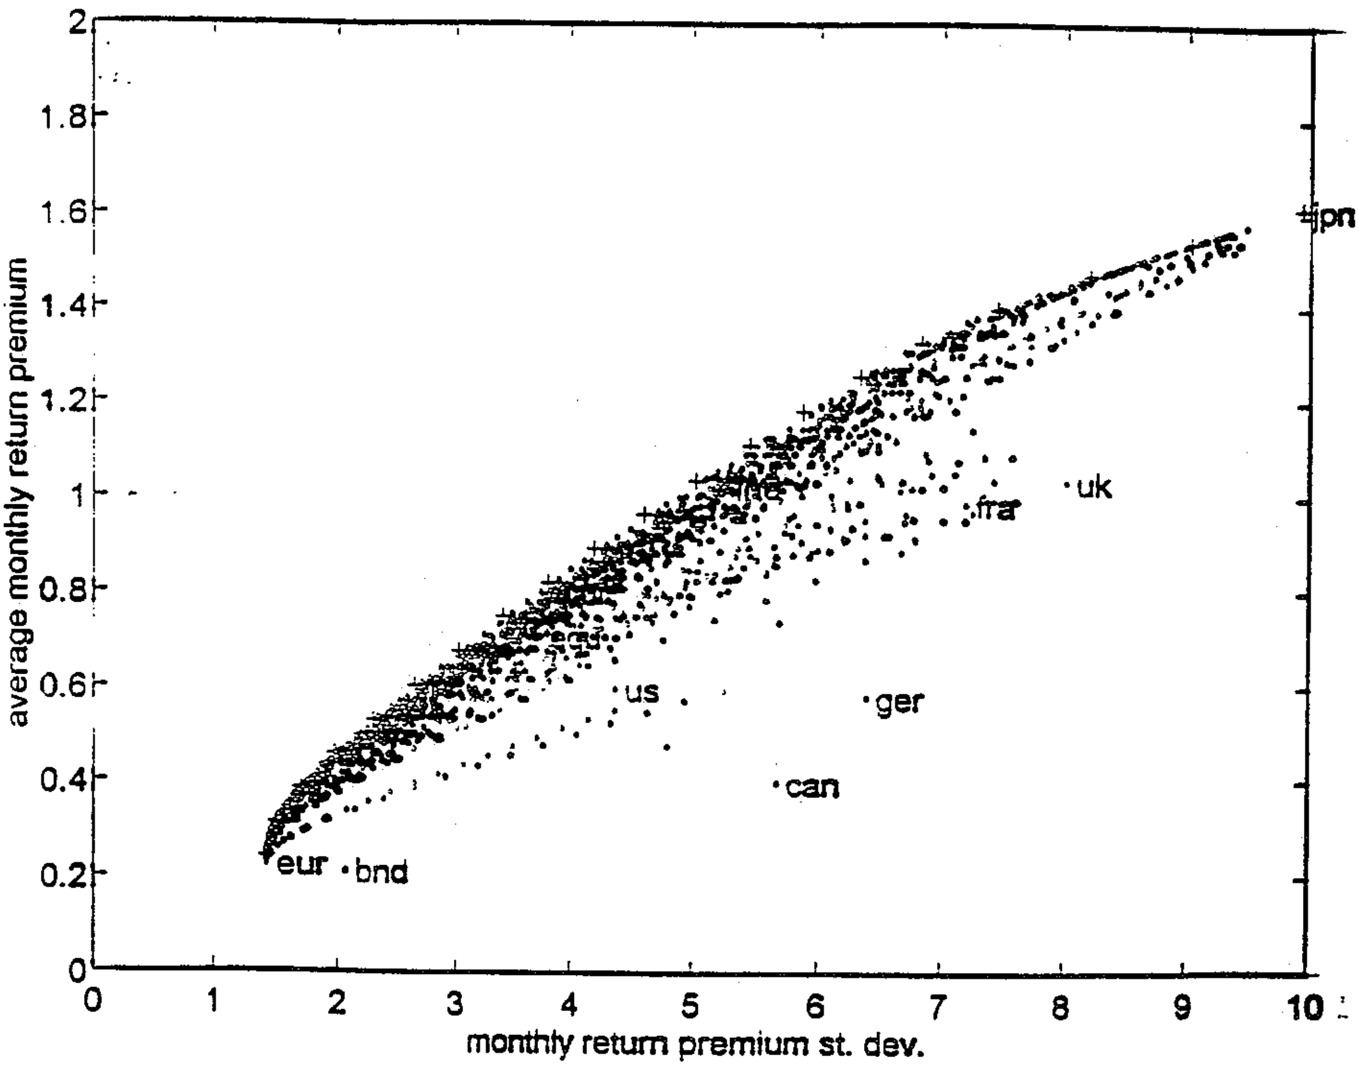

FIGURE 5
STATISTICAL EQUIVALENCE EFFICIENT FRONTIERS
100 SIMULATIONS, 60 TIME PERIODS
January 1978-June 1994 Monthly Data

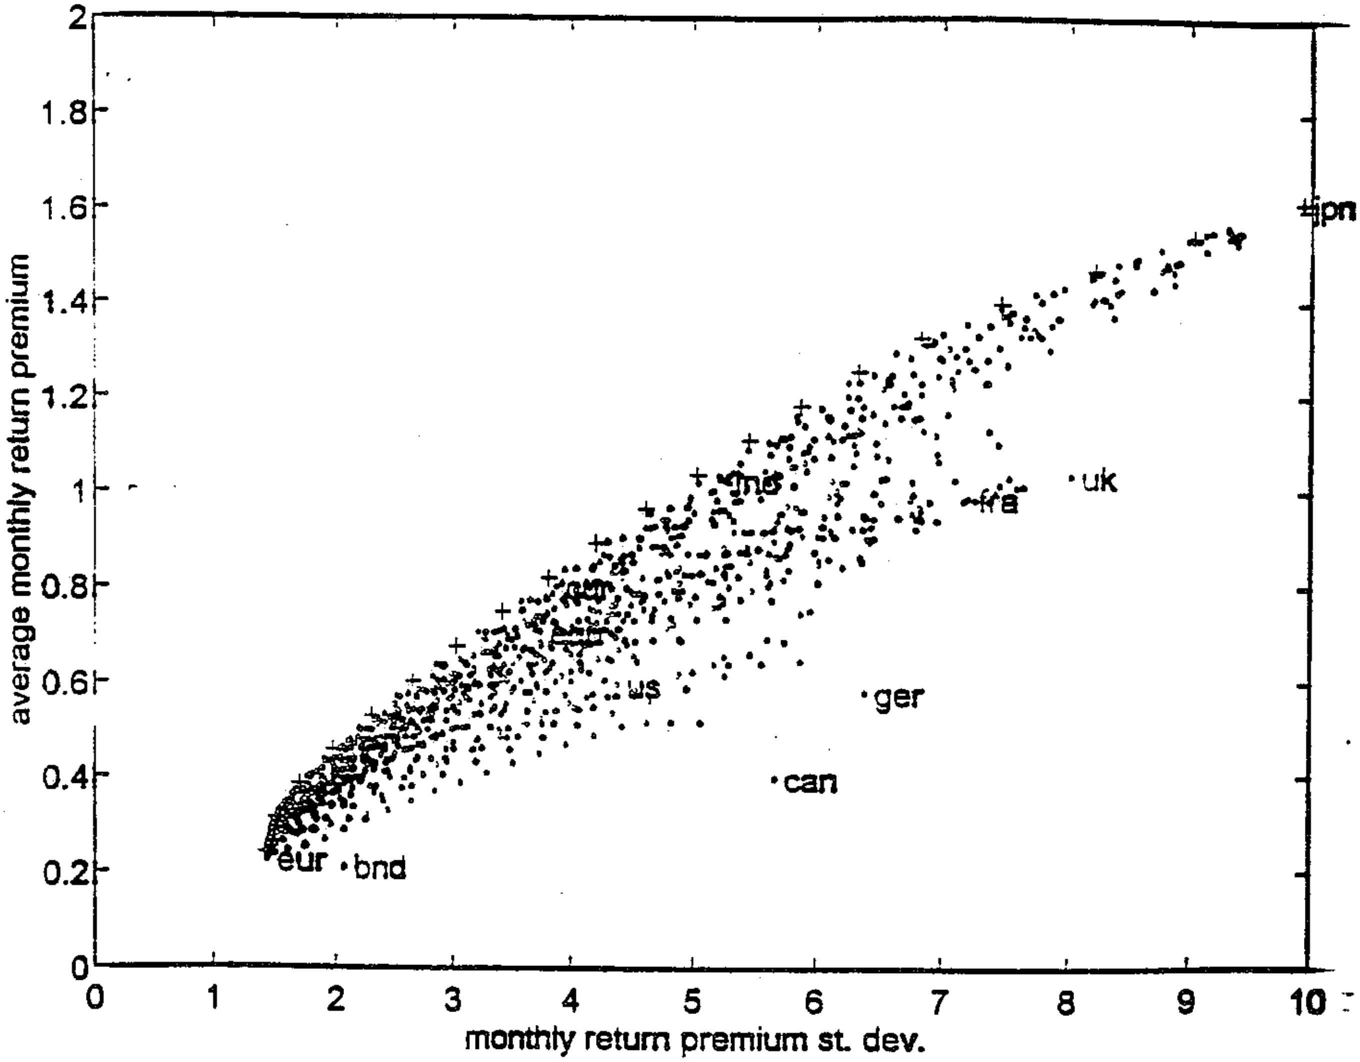

FIGURE 6
BAYES-STEIN ADJUSTED EFFICIENT STATISTICAL FRONTIERS
100 SIMULATIONS, 198 TIME PERIODS
January 1978-June 1994 Monthly Data

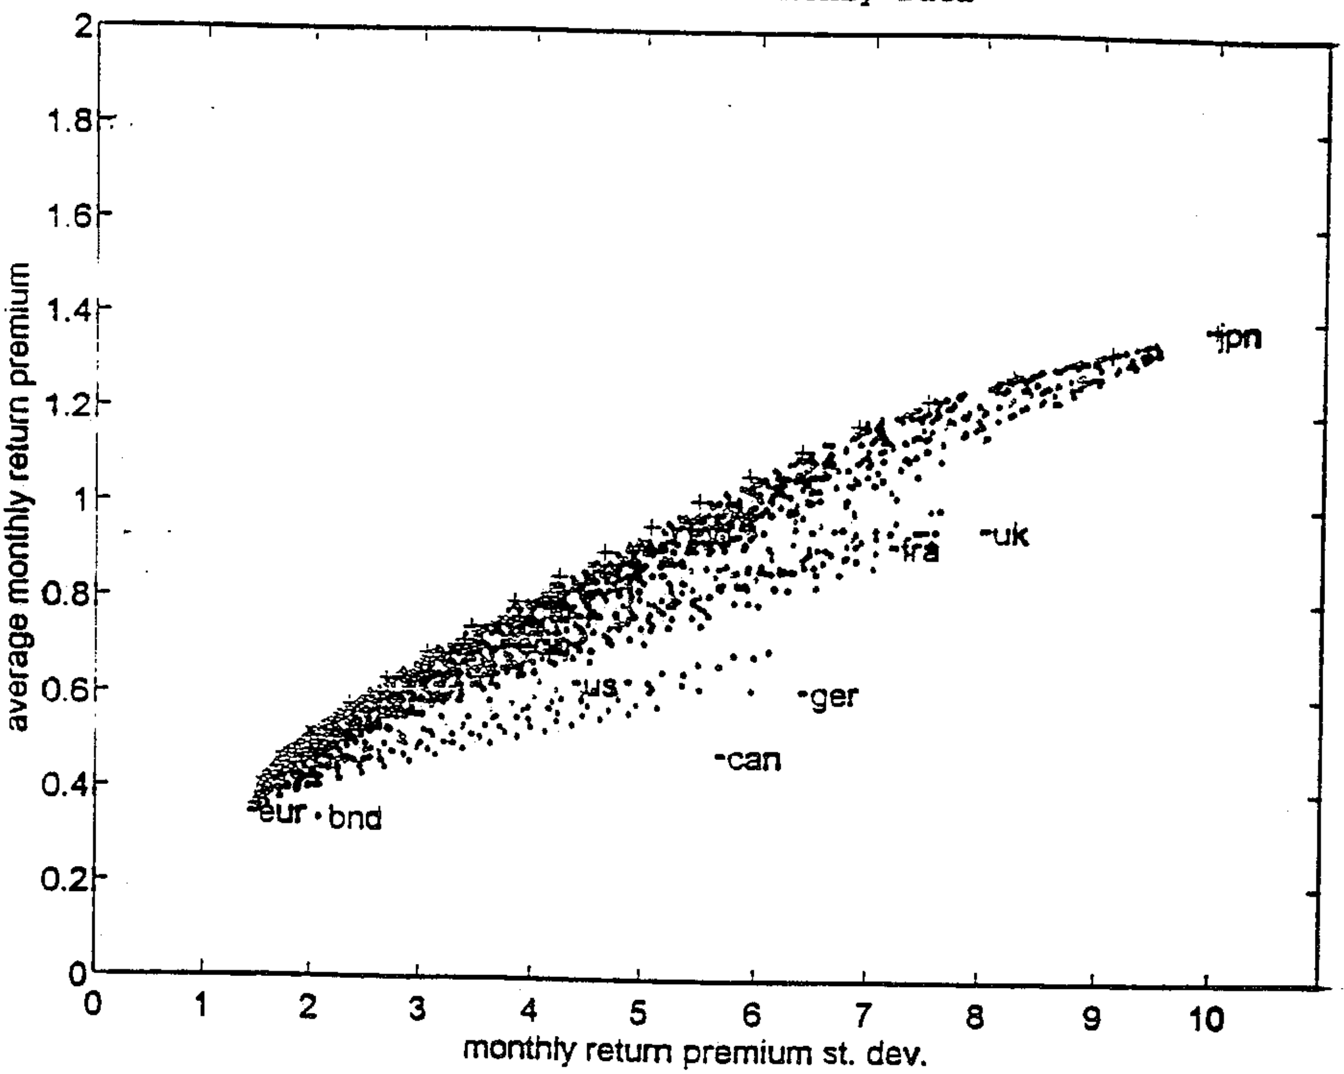

FIGURE 7
IMPLIED RETURN PREMIUM EFFICIENT FRONTIER
January 1978-June 1994 Monthly Data

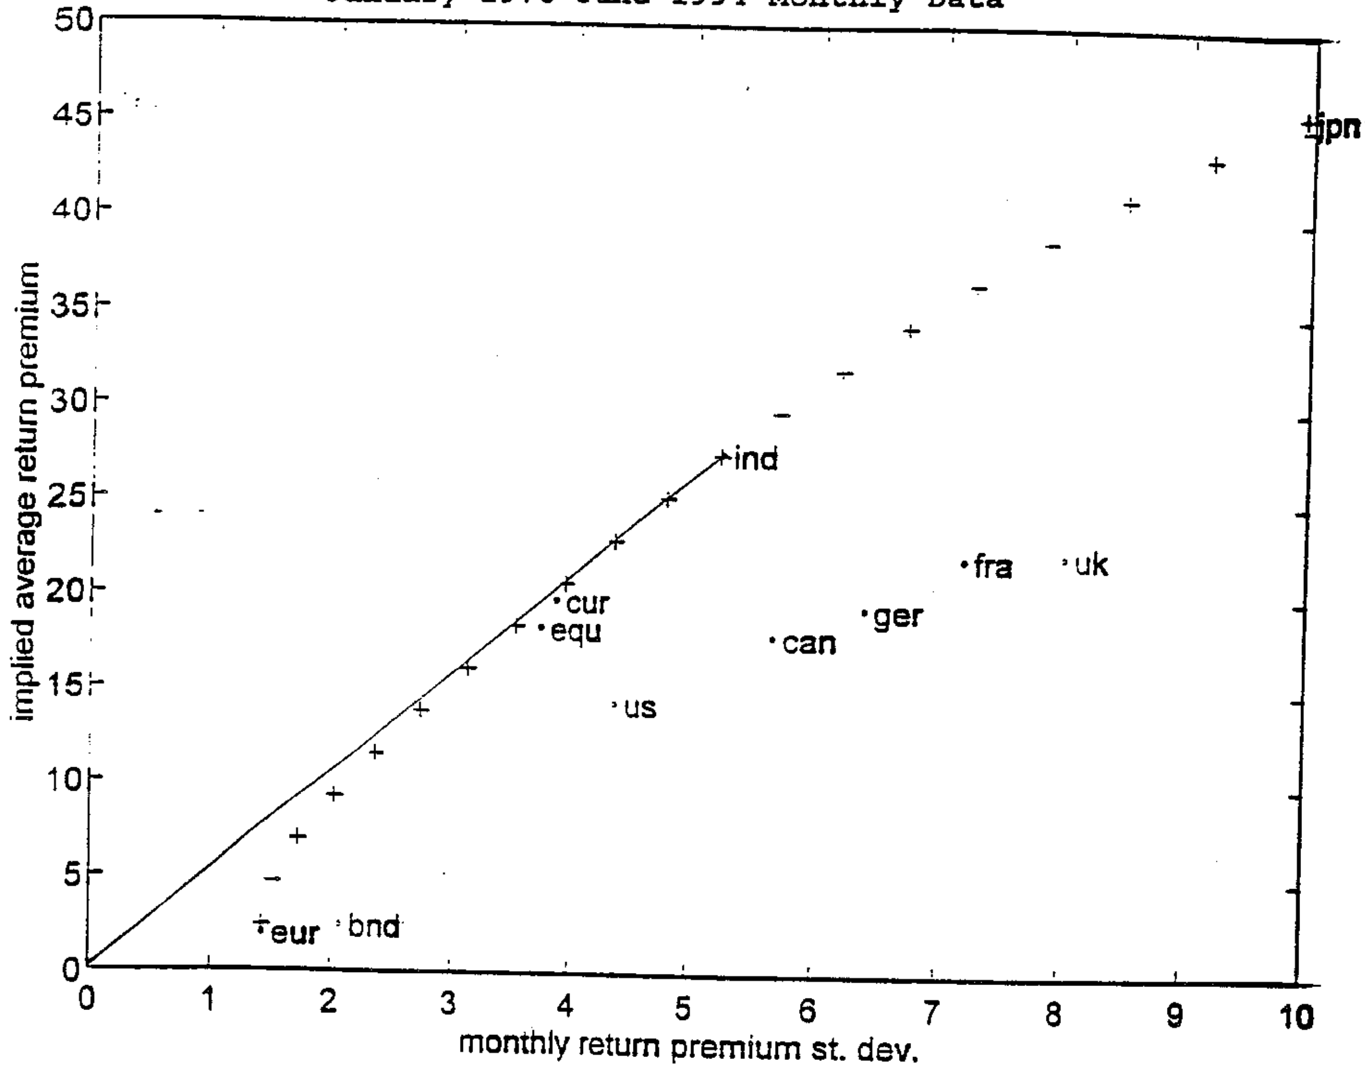

FIGURE 8
SCALED IMPLIED RETURN PREMIUM EFFICIENT FRONTIER
January 1978-June 1994 Monthly Data

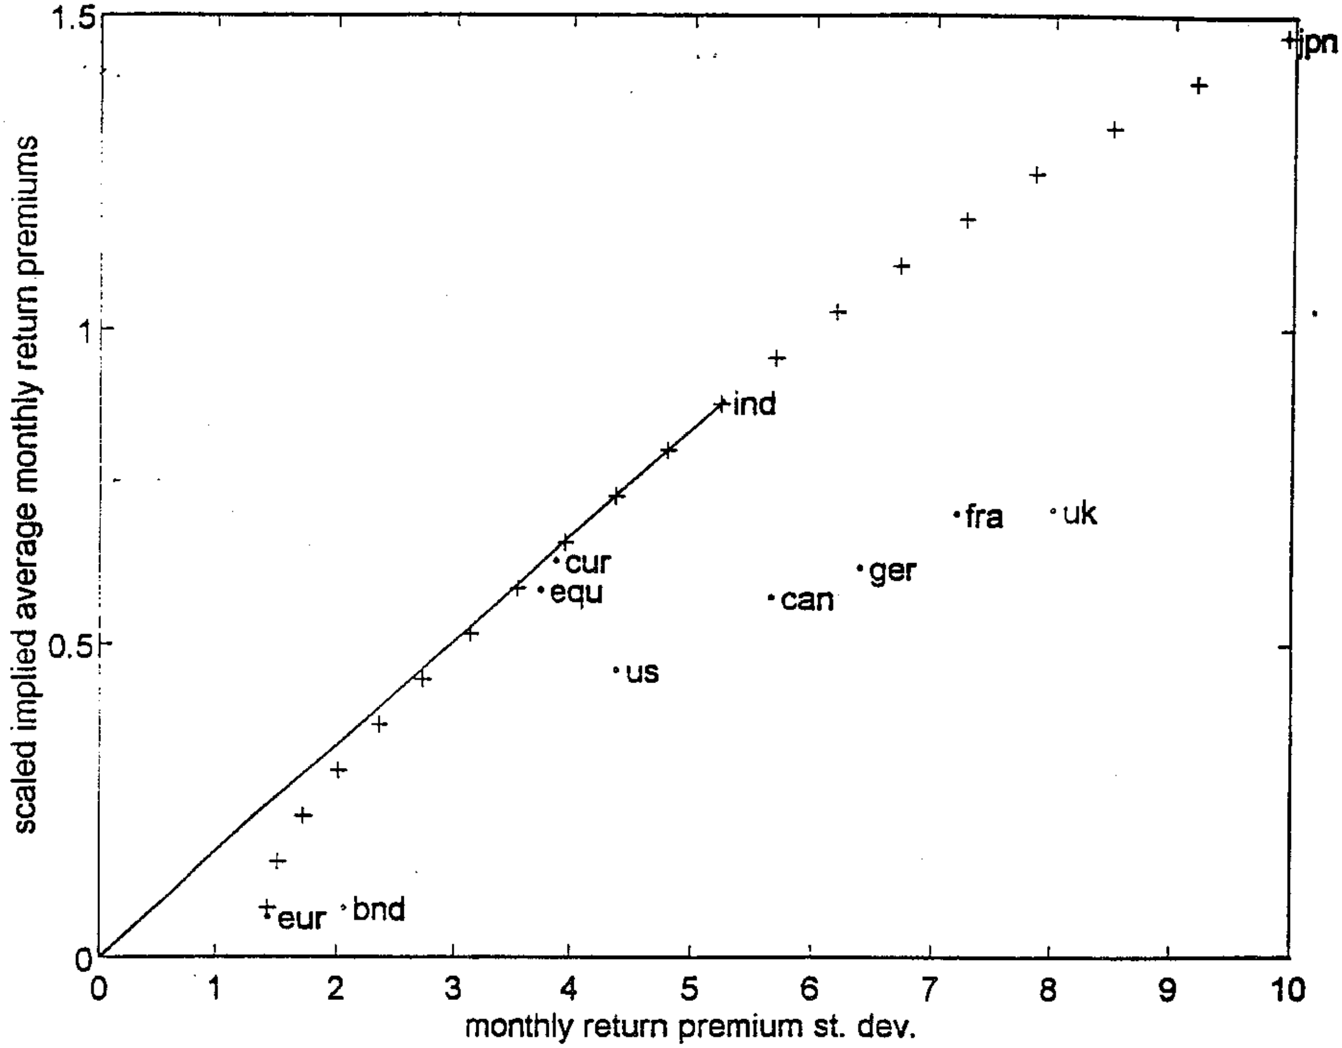

FIGURE 9
SCALED IMPLIED RETURNS EFFICIENT STATISTICAL FRONTIERS
100 SIMULATIONS, 198 TIME PERIODS
January 1978-June 1994 Monthly Data

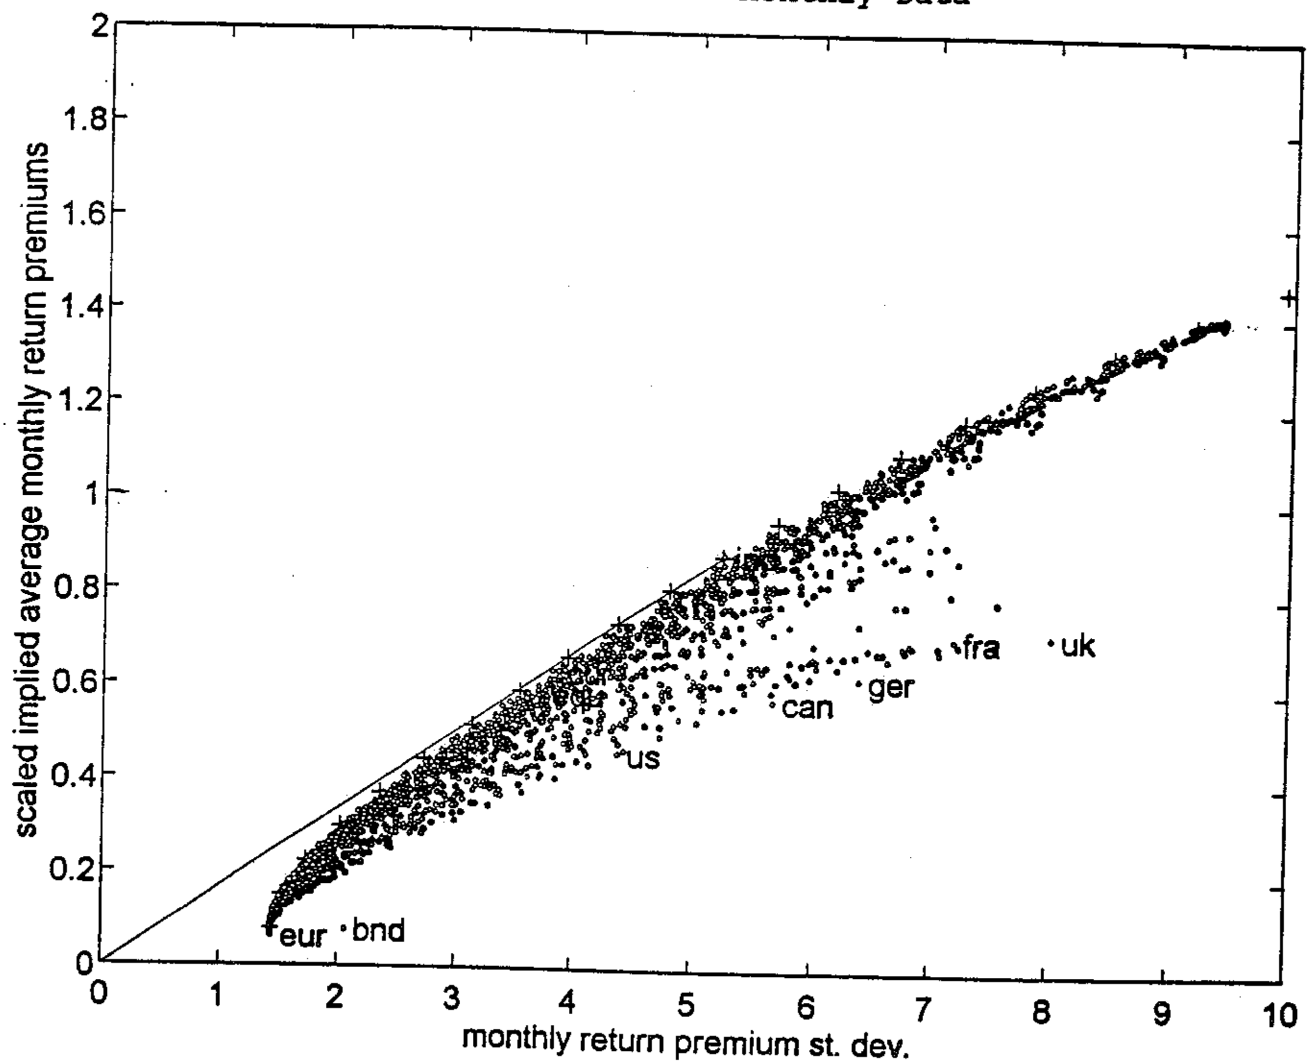

FIGURE 10
MIXED (+) AND SCALED (O) IMPLIED RETURN EFFICIENT FRONTIER
January 1978-June 1994 Monthly Data

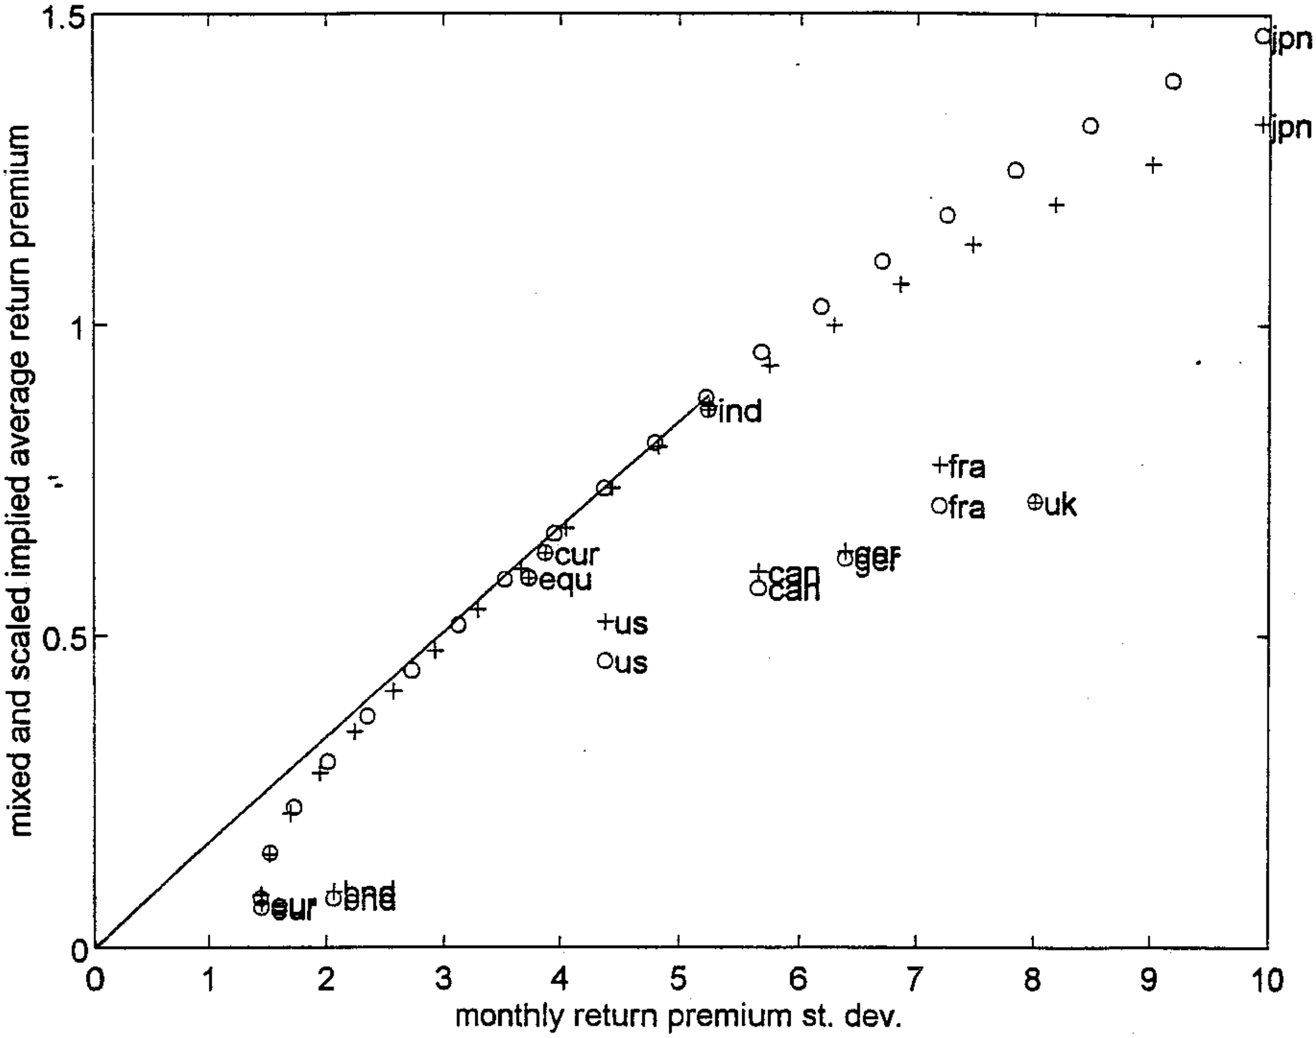

STATISTICAL ASSET ALLOCATION
Richard O. Michaud

FIGURE 11
INDEX-RELATIVE EFFICIENT FRONTIER
January 1978-June 1994 Monthly Data

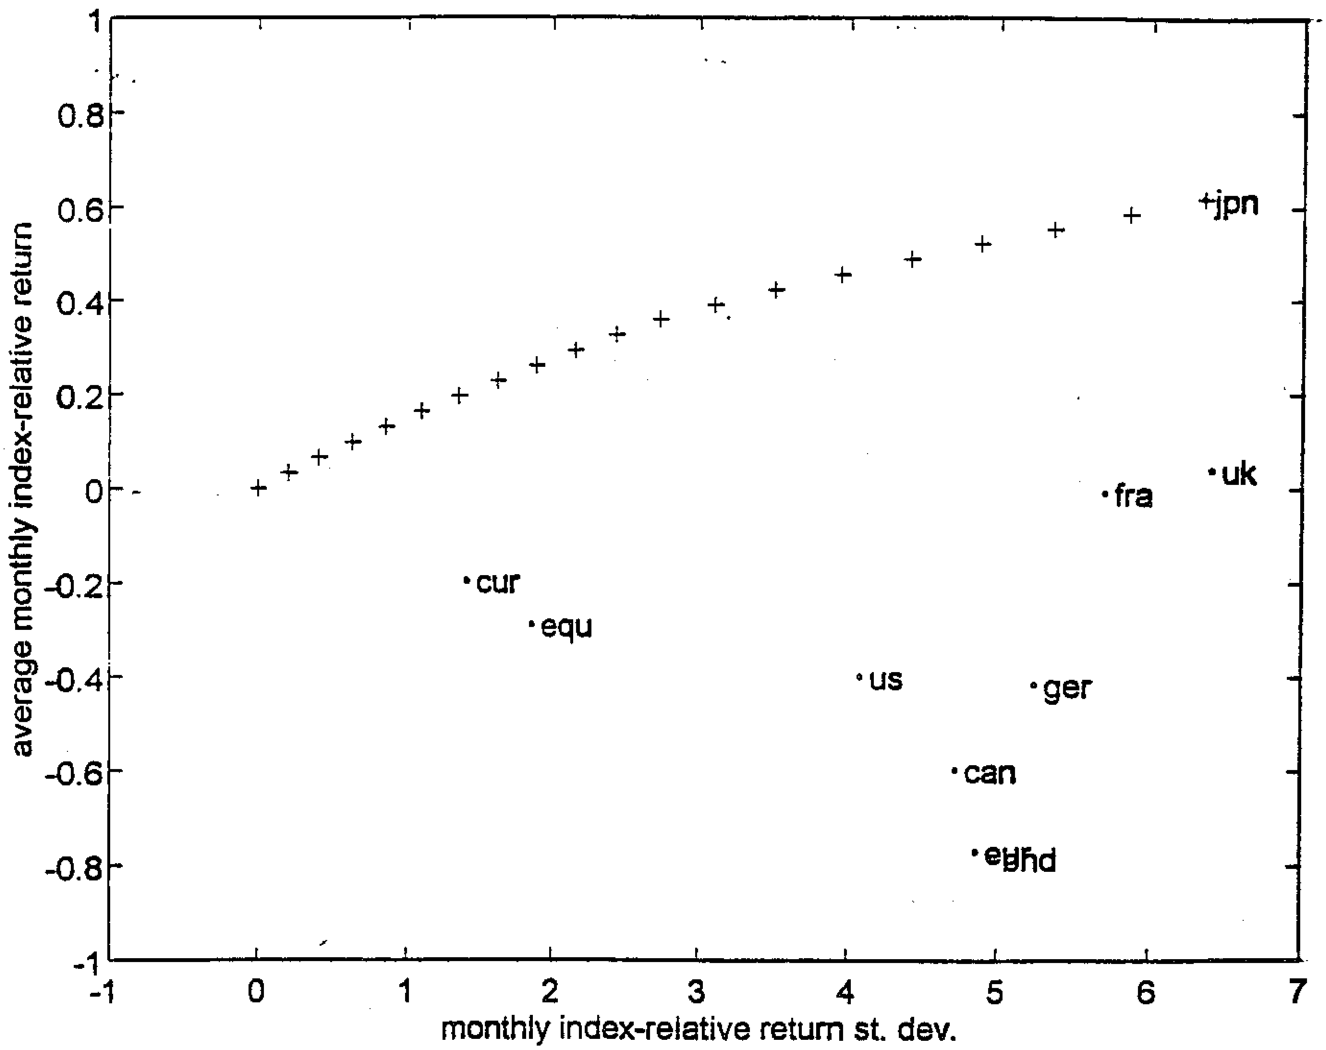

FIGURE 12
INDEX-RELATIVE EFFICIENT STATISTICAL FRONTIERS
100 SIMULATIONS, 198 TIME PERIODS
January 1978-June 1994 Monthly Data

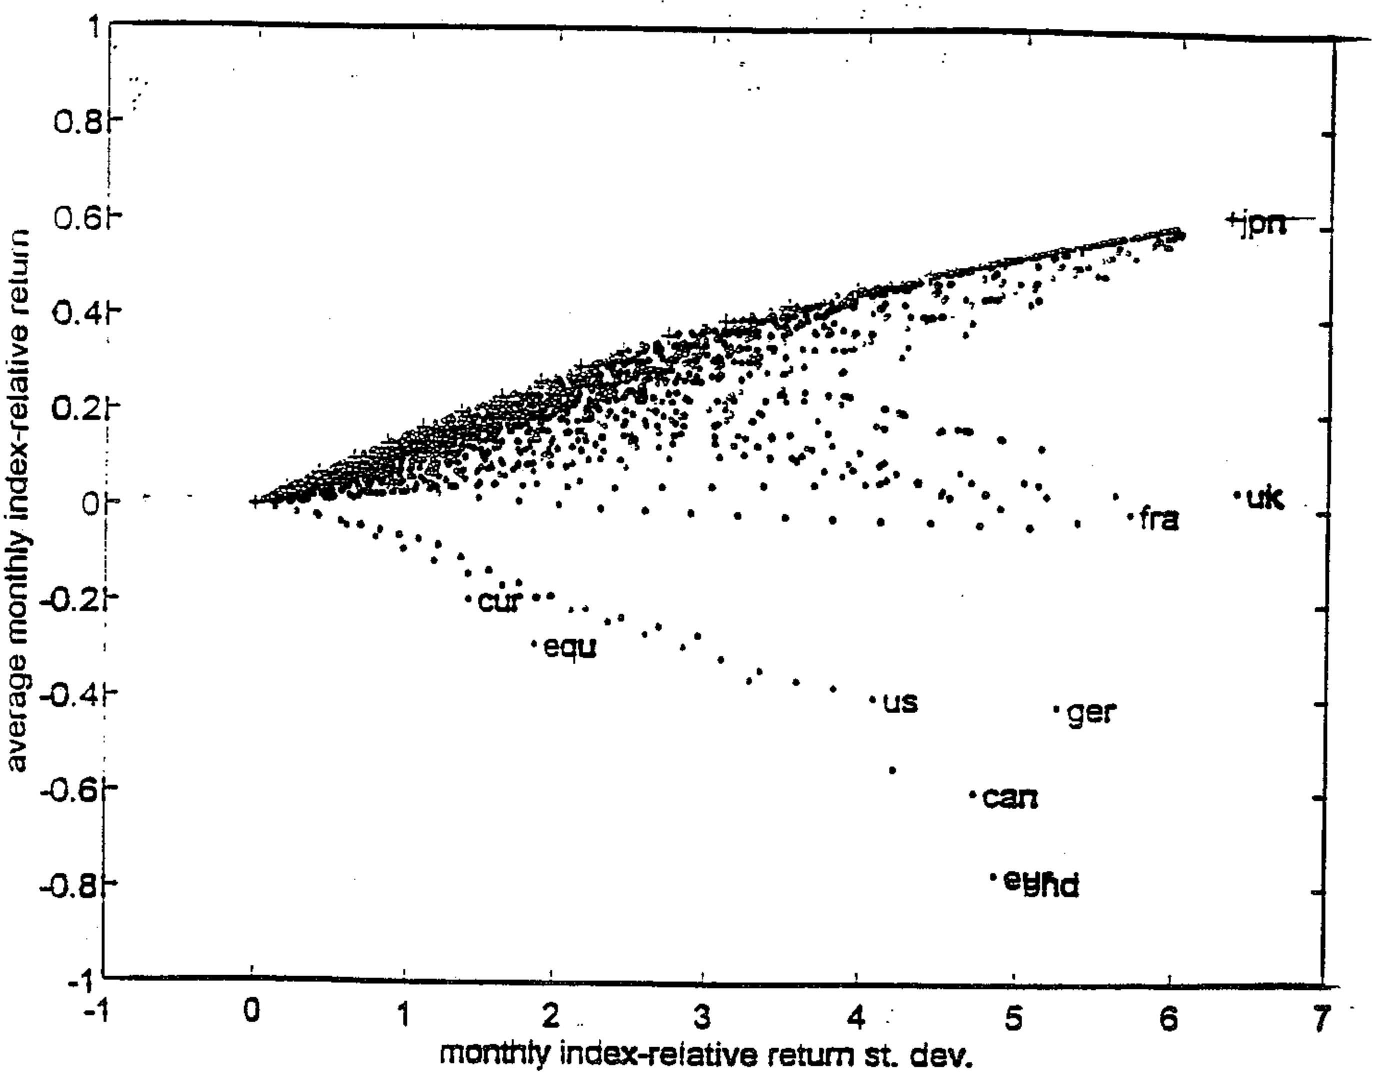

FIGURE 13
INDEX-RELATIVE EFFICIENT STATISTICAL FRONTIERS
100 SIMULATIONS, 60 TIME PERIODS
January 1978-June 1994 Monthly Data

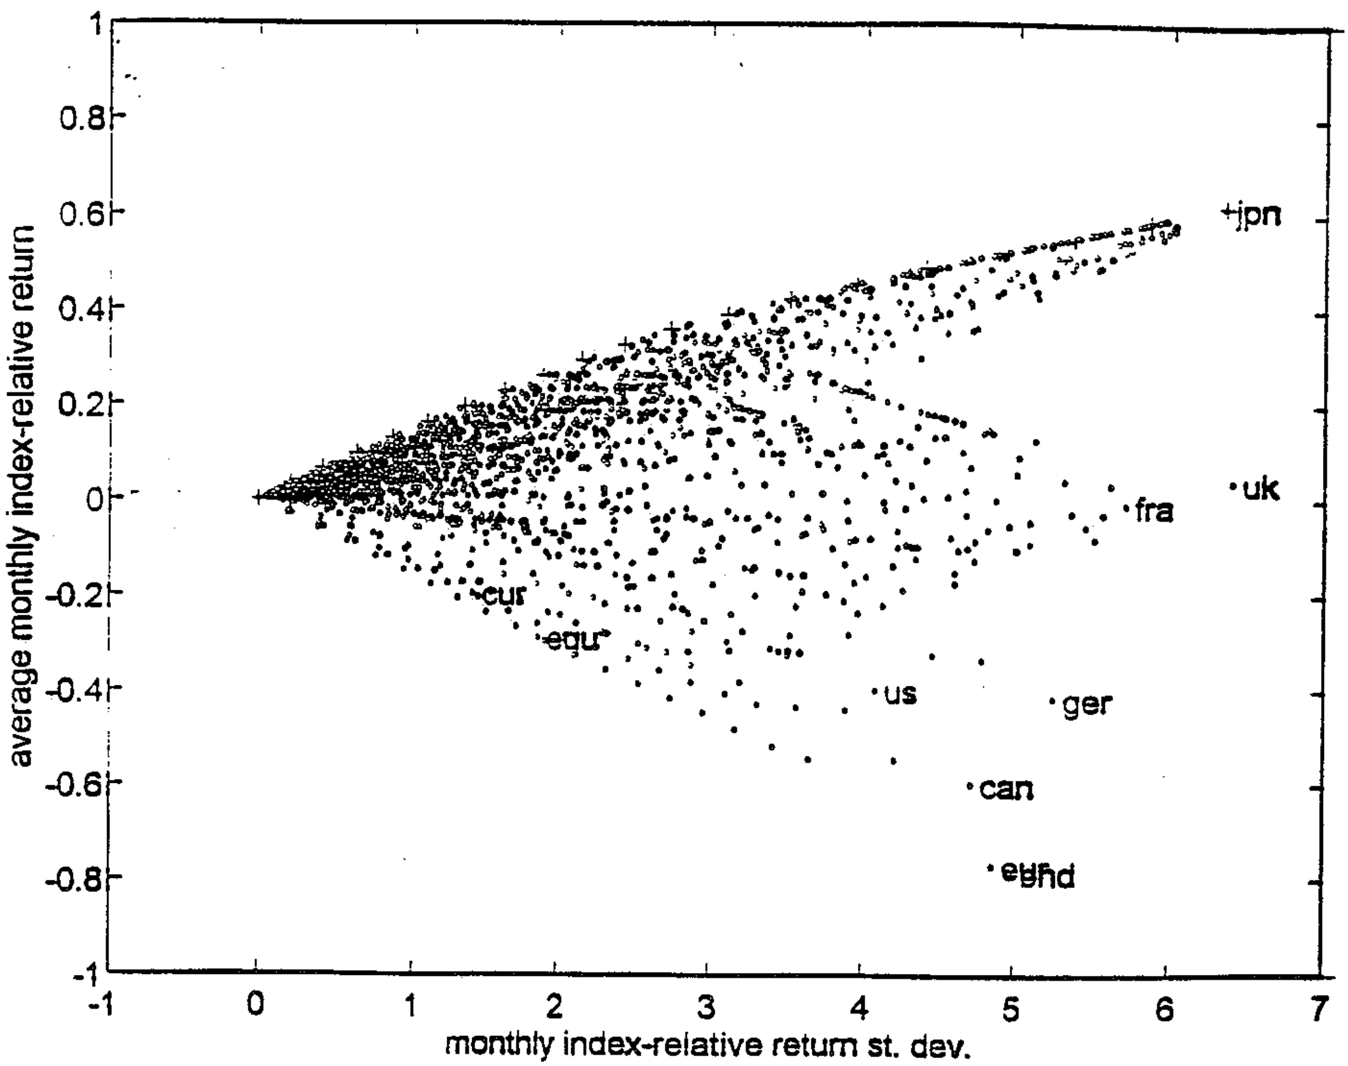

FIGURE 14
BENCHMARK LIABILITY EFFICIENT FRONTIER
January 1978-June 1994 Monthly Data

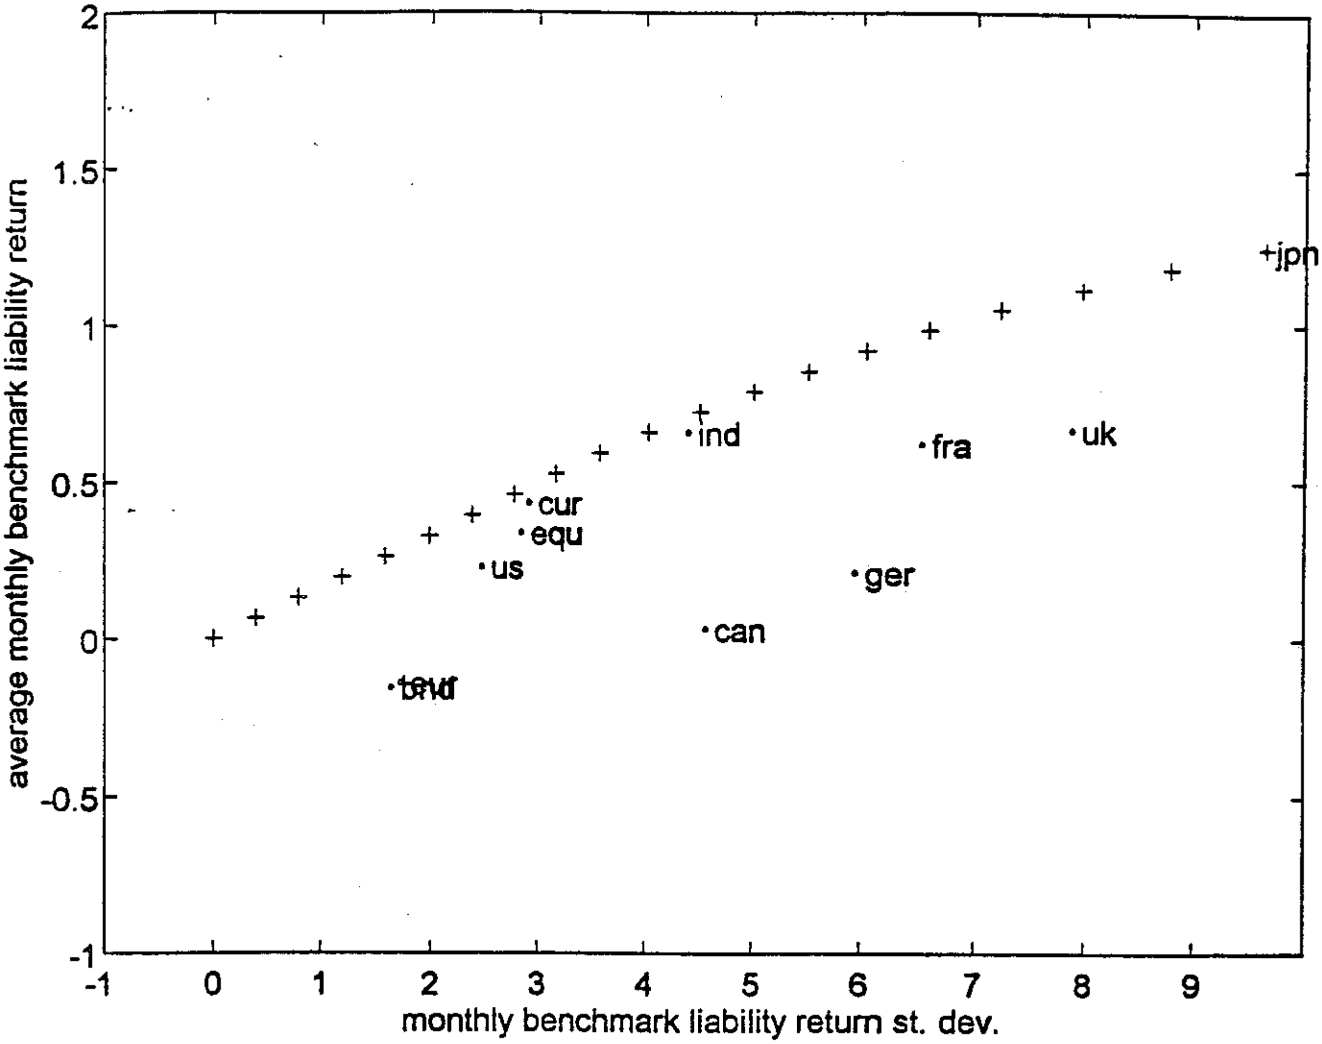

FIGURE 15
BENCHMARK LIABILITY EFFICIENT STATISTICAL FRONTIERS
100 SIMULATIONS, 198 TIME PERIODS
January 1978-June 1994 Monthly Data