A New View Of Mean Variance

September 1, 1998

Many authors have raised serious objections to mean-variance (MV) efficiency as a framework for defining portfolio optimality and have proposed a number of alternatives. Most of the alternatives can be classified in one of five categories: 1) non-variance risk measures, 2) utility function optimization, 3) multiperiod objectives, 4) Monte Carlo financial planning or 5) linear programming. Analysis shows that the alternatives often have their own serious limitations and that MV efficiency is far more robust than is appreciated. Although they are symptomatic of an underlying unease with MV efficiency, none of the proposals addresses the basic limitations of MV optimization.

Alternative Measures Of Risk

In MV efficiency, the variance, or standard deviation, of return is the measure of security and portfolio risk. The variance measures variability above and below the mean. From an investor's point of view, the variance of returns above the mean is often not "risk." One obvious and intuitively appealing non-variance measure of risk, discussed by Harry Markowitz as early as 1959, is the semi-variance or semi-standard deviation of return. In this risk measure, only returns below the mean are included in the estimate of variability.

The semi-variance is an example of a downside risk measure. In this case, downside risk is relative to the average or mean of return. There are many other ways to measure downside risk. A simple example is replacing average return with a specified level of return, such as zero or the risk-free rate.

Many other non-variance measures of variability are also available. Some of the more important include the mean absolute deviation and range measures. The pros and cons of various risk measures depend on the nature of the return distribution.

The return distribution of an asset or portfolio depends on several factors. Because the returns of diversified equity portfolios, equity indexes and other assets are often approximately symmetric over periods of institutional interest, efficiency based on non-variance risk measures may be nearly equivalent to MV efficiency.

An important issue is whether, in practice, non-variance risk measures lead to significantly different efficient portfolios. Figure 1 defines three reference portfolios used in subsequent analyses of MV portfolio efficiency: index, current and equal-weighted. The index portfolio is roughly consistent with a capitalization-weighted portfolio relative to a world equity benchmark for the six equity markets. The current portfolio represents a typical U.S.-based investor's global portfolio asset allocation. The most significant differences between the index and the current portfolios are the allocations to fixed-income assets. An equal-weighted portfolio is useful as a reference point.

Figure 2 displays the mean and standard deviation of the total monthly return premiums over the January 1978-December 1995 period for eight major asset classes: U.S. stocks and government/corporate bonds, euros, and the Canadian, French, German, Japanese and U.K. equity markets. The return premium is the return minus the risk-free rate. It is often convenient to use total return premiums, instead of total returns, as the basis of MV analysis in practice. Return premiums are similar to real rates of return. By removing the impact of varying risk-free rates, return premiums may be relatively more stable than total returns and more useful in a forecasting context.

The total return premium is the U.S. dollar total return minus the U.S. dollar short-term interest rate in each period. The monthly short-term interest rate for a U.S. dollar-based investor is usually defined as the U.S. T-bill 30-day return.

Figure 3 displays the MV efficient frontier associated with the historic premium return data, while Figure 4 provides an illustration, comparing the MV efficient frontier in Figure 3 with a mean-semi-variance efficient frontier based on the same historic data.

As Figure 4 shows, the two efficient frontiers are virtually identical, except at the low-risk end. The differences at low risk reflect the fact that bond returns are less symmetric than equities. Many currently fashionable risk alternatives have similar efficient frontier characteristics.

Some asset classes, such as options, do not have return distributions that are approximately symmetric. The return distributions of fixed-income indexes are not as symmetric as many equity asset classes. In addition, the return distribution of diversified equity portfolios becomes increasingly asymmetric over a long-enough period. Consequently, the variance measure for defining portfolio efficiency is not always useful or appropriate. For many applications of institutional interest, however, a variance-based efficient frontier is often little different (and even less often statistically significantly different) from frontiers that use other measures of risk. It is also generally a lot more convenient, especially when deriving statistical properties of efficient frontiers.

Utility Function Optimization

For many financial economists, maximizing expected utility of terminal wealth is the basis for all rational decision making under uncertainty. The issue of interest is whether Markowitz MV efficiency is consistent with expected utility maximization. If it is not, perhaps optimization based on specific utility functions should replace MV efficiency.

Markowitz MV efficiency is strictly consistent with expected utility maximization only under either of two conditions: normally distributed asset returns or quadratic utility functions. The normal distribution assumption is unacceptable to most analysts and investors. Although diversified equity portfolio and index returns are often reasonably symmetric, their distribution is not precisely normal. In addition, the limitations of quadratic utility as a representation of investor behavior are well known. Consequently, MV efficiency is not strictly consistent with expected utility maximization. (Note: Because a quadratic function is not monotone increasing as a function of wealth, from some point on, expected quadratic utility declines as a function of increasing wealth. Quadratic utility functions are primarily useful as approximations of expected utility maximization, usually in some particular region of the wealth spectrum.)

One alternative is to define portfolio optimality in terms of optimizing specific utility functions. Many analysts have suggested that utility functions are a more rational basis for investor decision making and portfolio structuring. Utility function optimizations need not resemble MV efficient portfolios.

There are, however, significant practical limitations to using utility functions as the basis of defining an optimization. One obvious limitation is the feasibility and viability of practical algorithms for computing optimal portfolios. Depending on functional form, nonlinear optimization methods may be required that may have significant limitations in many applications.

An equally important limitation of the utility function approach to portfolio optimization is utility function specificity. In practice, investor utility is unknown. The lack of specificity of the investor's utility function is a far more daunting practical problem than it may appear. This is because a class of utility functions can have similar functional forms, perhaps differing in the value of only one or two parameters, yet represent a very wide, even contradictory, spectrum of risk bearing and investment behavior. In these cases, even small errors in the estimation of utility function parameters can lead to very large changes in the investment characteristics of an optimal portfolio. As a practical matter, the problem of specifying with sufficient accuracy the appropriate utility function for a given investor appears to be a severe practical limitation of utility function-based portfolio optimization.

On the other hand, MV efficiency can be rationalized as a convenient approximation of expected utility maximization. For non-pathologic utility functions, quadratic utility functions are often useful approximations of maximum expected utility at a point. Note that the best-approximating quadratic function may not be the same at different points of the expected utility function. Consequently, MV efficient portfolios are often good approximations of maximum expected utility and a practical framework for portfolio optimization.

The use of utility functions in defining portfolio optimality often divides practitioners from academics. From a rigorous point of view, only the specification of an appropriate utility function will do for defining portfolio optimality. However, few practitioners use non-quadratic utility functions to find optimal portfolios. Given the difficulty of estimating utility functions with sufficient precision, the convenience of quadratic programming algorithms and the robustness of the approximating power of quadratic utility at a point, MV efficiency is often the practical tool of choice.

Multiperiod Investment Horizons

Markowitz MV efficiency is formally a single-period model for investment behavior. Many institutional investors, however, such as endowment and pension funds, have long-term investment horizons on the order of five, 10 or 20 years. How useful is MV efficiency for investors with long-term investment objectives?

One way to address long-term objectives is to base MV efficiency analysis on long-term units of time. MV efficiency, however, is probably most appropriate for relatively short-term periods. This is true because a quadratic approximation of maximum expected utility is most likely to be valid for monthly, quarterly, or yearly periods. In addition, lengthening the unit of time reduces the number of independent periods in a historic data set and the statistical significance of optimization parameter estimates. On the other hand, increasing the historic data period may diminish the relevance of the estimates for the forecast period.

An alternative approach is to consider the multiperiod distribution of the geometric mean of return. The geometric mean, or compound, return is the statistic of choice for summarizing portfolio return over multiple periods. For example, suppose an investor experiences a 100% return in one period and a -50% return in the next period. The two-period average return is 25%, but the two-period wealth is the same as at the beginning. Therefore, the true multiperiod return is 0%. The geometric mean provides the correct answer, whereas the average does not.

Assume that, in each period, MV efficiency defines optimal portfolio choice. Also assume that the distribution of single-period return does not vary (appreciably) over the multiperiod investment horizon. What are the long-term consequences of repeatedly investing in MV efficient portfolios?

Some essential results are due to Markowitz. He shows that 1) MV efficient portfolios need not be efficient in the long run and 2) long-term efficiency is not necessarily monotonic in portfolio risk. In particular, MV efficient portfolios on the upper segment of the efficient frontier may be less long-term efficient than portfolios with less risk.

Nils Hakansson, in a 1971 article in the Journal of Financial and Quantitative Analysis, gives an example of an MV efficient frontier in which repeated investing produces a negative long-term geometric mean at all points. This example shows that all MV efficient frontier portfolios may lead to ruin with probability equal to one over long enough investment horizons. However, the Hakansson example is neither typical nor likely.

Further analysis of the geometric mean criterion is useful. The mean and variance of N-period geometric mean return is a natural N-period generalization of Markowitz efficiency. (It may be fitting to call the objective Hakansson efficiency, after the researcher who has done much of the pioneering work in this area.) Various approximations show that portfolios on the (single-period) MV efficient frontier are often good approximations of N-period geometric mean efficient portfolios. Consequently, N-period geometric mean MV efficiency is roughly a special case of MV efficiency in many cases of practical interest.

Define the critical point as the MV efficient portfolio with the maximum N-period expected geometric mean return. The critical point is a useful construct for understanding and using N-period geometric mean efficiency. The N-period expected geometric mean is a positive function of the mean of (single-period) expected return and a negative function of the variance. Consequently, the critical point defines the boundary of portfolios on the lower segment of the MV efficient frontier that are N-period geometric mean MV efficient and those on the upper segment that are not. N-period horizon MV efficiency leads to the simple decision rule of considering only MV efficient portfolios on the lower segment of the efficient frontier up to the critical point efficient portfolio. Note that critical points that are not end points of the MV efficient frontier do not always exist.

A number of analysts have raised objections to the geometric mean as an investment criterion. In particular, a significant controversy emerged from the proposal of using the (long-term) expected geometric mean as a surrogate for expected utility. This controversy, although it is beyond the scope of this discussion, is essentially concerned with the limitations of using any investment rule, however attractive, as an alternative to expected utility maximization. The opposing view concerns the limitations of using utility functions in practice and the value of the MV geometric mean criterion as a convenient source of useful investment information. MV geometric mean investment objectives are often consistent with many institutional investment mandates.

One more issue may be of interest. The assumption has been that the investor repeatedly invests in the same efficient frontier portfolio over some investment horizon. However, optimal multiperiod investment with an MV geometric mean objective is a dynamic programming strategy that implies varying the choices of MV efficient portfolios in each period.

Multiperiod considerations are important issues for investors with long-term investment objectives. To avoid possible negative long-term consequences of MV efficiency, a simple solution is to limit consideration to efficient frontier portfolios at or below the critical point. As a useful approximation, it is convenient to consider long-term efficient portfolios as a subset and a special case of MV efficiency.

Asset-Liability Planning Studies

Many financial institutions invest substantial resources in defining an appropriate long-term average asset allocation or investment policy. They do this because the long-term average asset allocation is one of the most important investment decisions an institution or investor can make. The importance of defining an optimal investment policy has spurred alternative approaches to MV efficiency analysis. Probably the most important of these is an asset-liability financial planning study based on Monte Carlo simulation.

In a Monte Carlo financial planning study, a computer model simulates the random functioning of a fund and changes in its liabilities over time. Depending on the study and application, the liability model may be very detailed. For a defined benefit pension plan, it can include a comprehensive examination of corporate objectives, economic projections and future hiring policy as well as current workforce census. In some cases, liability modeling may affect asset allocation decisions in terms of feasibility, particularly for regulated firms, such as insurance companies.

Estimates of likely cash flows and funding status result from performing many simulations. By varying asset return and allocation assumptions, the simulation can evaluate the implications of various asset allocation decisions on the evolution of funding status and cash flows. Endowment fund simulations can provide useful information on likely levels of endowment spending and fund value over time. Similarly, defined benefit pension plan simulations can be useful for anticipating required contributions and plan funding status for various assumptions and investment periods.

The important issue is whether Monte Carlo asset-liability financial planning is a superior alternative to MV efficiency for defining an optimal asset allocation. Proponents argue that plan funding status and cash flow objectives are more meaningful than the MV efficiency of a feasible portfolio. The anticipation of likely cash flows and required contributions can provide valuable fund planning information. The problem is that such information may have relatively limited usefulness for defining an optimal long-term asset allocation.

Generally, only feasible MV efficient frontier asset allocations relative to fund liability are of interest. This is because feasible liability-relative allocations with more expected return for a given risk level are almost always preferable. Consequently, a valid Monte Carlo asset-liability simulation study still requires liability-relative MV efficiency analysis to determine candidate efficient allocations. Within the context of feasible liability-relative efficient allocations, consider the consequences of varying asset mixes. In general, the Monte Carlo results show that riskier efficient asset mixes lead to a greater likelihood of meeting or exceeding funding objectives and of increasing volatility. Evaluating the tradeoffs associated with funding status and cash flow volatility in various time periods is often of no less difficulty than evaluating the risk-return tradeoffs in an efficient frontier context. Monte Carlo simulation studies do little more than illustrate the simple principle that, for feasible efficient portfolios, more risk leads to more return on average, and more volatility.

There is an exception to these basic principles governing Monte Carlo asset-liability financial planning simulation. Analysts have seen that increasing efficient portfolio risk does not always lead to an increased likelihood of meeting fund objectives. Such results appear to rationalize the importance of the Monte Carlo procedure relative to efficiency analysis. However, our earlier discussion can help to explain this result.

Monte Carlo simulation studies generally assume repeated investment in candidate portfolios. If the liability-relative efficient frontier has an internal critical point, asset allocations on the long-term inefficient segment will exhibit the behavior that increasing risk leads to decreases in the ability of the fund to meet objectives. In many cases, such results can be anticipated analytically by computing the critical point of the efficient frontier (or of the candidate portfolios) and analyzing N-period geometric mean efficiency. However, the issue is more than simply a tool for rationalizing the results of a simulation study. The N-period geometric mean implications of input assumptions are the engine that drives the simulations and can lead to predefined conclusions.

Monte Carlo asset-liability simulation has many uses as a tool for financial planning. It is useful for understanding the likelihood of meeting funding objectives and likely cash flows associated with various fund investments and allocations. The procedure has limited value, however, as an alternative to MV efficiency for defining an optimal asset allocation. Many of its asset allocation benefits are analytically anticipatable in terms of the mean and variance of the multiperiod geometric mean distribution. On the other hand, the analytic tools for understanding the geometric mean distribution as a function of the MV efficient frontier portfolios over an N-period investment horizon can be useful for designing effective Monte Carlo simulation financial planning studies.

Linear Programming Optimization

The practical limitations of MV optimization as a tool of equity portfolio management have been familiar to many astute asset managers for many years. One common alternative is to optimize portfolios with linear programming.

Linear programming portfolio optimization is a special case of quadratic programming. The most significant difference is that linear programming does not include portfolio variance. In this procedure, the objective is to maximize expected equity portfolio return subject to a variety of linear equality and inequality constraints on portfolio structure. The procedure relies heavily on clever use of constraints on industries, sectors and stock weights to control portfolio risk and maximize expected return. The constraints also serve to design portfolios with various specific characteristics and objectives.

In the hands of a sophisticated analyst, linear programming is an optimization technique that may avoid many of the fundamental limitations of equity portfolio MV optimization. It has its own limitations, however. In practice, it is difficult to control the structure of a portfolio precisely. From a theoretical point of view, only an MV optimization framework can optimally use active forecast information. The issue remains whether any linear programming approach is to be preferred to a carefully defined, input-adjusted MV optimization. The problems that most limit the practical value of MV optimization may be more attributable to the return than to risk dimension, which suggests that the linear programming alternative may ultimately be of limited value.

A somewhat obvious final issue may be worthy of note. Several investment institutions use far less sophisticated optimization procedures than linear or quadratic programming. These "homemade" optimization alternatives are often not the product of a conscious effort to avoid MV optimizer limitations but reflect a lack of analytic sophistication in the organization. Technical limitations in an optimization algorithm are unlikely to enhance the investment value of a portfolio over standard procedures.

-Richard O. Michaud