Asset Allocation in a Downside-Risk Framework

W. V. Harlow

Asset Allocation in a Downside-Risk Framework

H. V. Harlow is a Vice President and portfolio manager with Fidelity Management Trust Company in Boston. He joined Fidelity from Salomon Brothers, where he was responsible for the design of quantitative investment and trading strategies. He serves on the CFA Candidate Curriculum Committee of the Association for Investment Management and Research and was a member of the AIMR's Body of Knowledge Task Force.

On his article: "A downside-risk framework focuses on return deviations below some target return level. This view accords with most investment managers' perception of risk. Downside risk offers an attractive approach to asset allocation decisions. Theoretically more general than the traditional mean-variance technique, it also promises significant improvement in the risk-reward tradeoffs afforded the investor."

by W. V. Harlow

Asset Allocation in a Downside-Risk Framework

A downside-risk approach to investment decisions uses intuitive measures of risk that focus on return dispersions below a specified target or benchmark return. Downside-risk measures are attractive not only because they are consistent with investors' perception of risk, but also because the theoretical assumptions required to justify their use are very simple. Equally important, a number of well known risk measures, including the traditional variance (standard deviation) measure, are special cases of the downside-risk approach. Asset allocation in a downside-risk framework therefore determines an investment opportunity set for downside-averse investors that is at least as efficient as that derived using conventional techniques.

A set of international asset allocation examples demonstrates the benefits of the downside-risk framework. Specifically, optimizations based on downside measures produce portfolio strategies with realized returns that have less downside risk exposure than those determined using variance. Thus investors averse to below-target return dispersions achieve a more attractive risk-return tradeoff within this framework. Moreover, in the asset allocation examples considered, the downside-risk approach produces a significantly higher average bond allocation relative to stocks. This difference in asset composition increases downside protection while offering the same or a greater level of expected return.

CENTRAL TO MODERN portfolio theory is the premise that investment decisions are made to achieve an optimal risk/return tradeoff from the available opportunities. In order to meet this objective, the portfolio manager must first evaluate capital market information and quantify *ex ante* measures of both risk and expected return for the appropriate set of assets. The next task is to isolate those combinations of assets that are the most "efficient," in the sense of providing the lowest level of risk for a desired level of expected return, and then to select one combination that is consistent with the risk tolerance of the investor.

While the principle of identifying portfolios with the required risk and return characteristics is certainly clear, the appropriate definition of risk is more ambiguous. One manager might view risk as the probability of shortfall below some benchmark level of return, for example, while another may be more sensitive to the overall magnitude of a loss, should one occur. These seemingly disparate notions of risk, as well as other possible definitions, serve as a reminder that simple return variance (or standard deviation)—the traditional measure of risk—is sometimes deficient for dealing with the rich set of portfolio objectives and constraints that investment managers often formulate.

This article discusses and demonstrates a general approach to asset allocation based on definitions of risk that are attractive alternatives to variance.[1] These alternatives all capture the appealing notion of "downside risk" and provide a more robust approach to portfolio optimization. Using an **asymmetric measure of risk** that focuses on the returns below a specified target or benchmark return level, this framework includes as special cases such well known measures as the probability of loss, expected loss

1. Footnotes appear at end of article.

Glossary

Asymmetric Measure of Risk: Any measure of risk that focuses on a portion of the return distribution rather than the spread or dispersion of the overall distribution. Downside-risk measures are asymmetric and isolate return deviations in the left tail of the distribution that fall below a specified target rate.

Semivariance: An asymmetric measure of risk that focuses on squared return deviations below the mean of the distribution. Target semivariance is a similar but more general measure in which return dispersions are considered below any arbitrary target or benchmark level of return.

Target Shortfall: A concept of downside risk that captures the severity of not achieving some minimum target or benchmark return. The target shortfall represents the expected deviation of returns falling below the target rate.

Mean-Variance Formulation: The traditional formulation of the investment decision problem, stated in terms of the expected mean return and variance of a portfolio of assets. In order to achieve the most efficient portfolio, assets are combined so as to minimize variance for a given level of return.

Downside-Risk Optimization: An alternative formulation of the investor's decision problem using a downside measure of risk as opposed to the variance or standard deviation of asset returns. Thus, instead of minimizing variance for a given level of return, a downside measure such as target semivariance is minimized.

and semivariance.[2] It also includes the traditional variance measure as a special case, thus ensuring that the downside-risk approach determines an efficient frontier for a downside-averse investor that is at least as optimal as that derived using existing techniques.

Although the concept of downside risk is older than modern portfolio theory itself, the theoretical development and extensions of this risk framework are relatively recent. Unfortunately, empirical investigations of this approach within the context of asset allocation remain somewhat limited. To obviate this deficiency, the present study examines the downside-risk framework by considering an international asset allocation example across 11 countries, using returns that span the last decade. In addition, it compares the downside-risk results with those derived using the mean-variance approach.

We find that empirical optimizations based on downside measures are more efficient than mean-variance measures in the sense that the selected portfolios have less downside exposure than those determined using variance. Thus investors averse to below-target return dispersions achieve a more attractive risk/return tradeoff with a downside-risk framework. Moreover, in the asset allocation examples considered, the downside-risk approach produces a significantly higher average bond allocation than a traditional approach. This difference in asset composition increases downside protection while offering the same or a greater level of expected return.

In view of its appealing theoretical and intuitive features, the downside-risk framework should provide a useful set of tools for portfolio managers considering a broad set of problems. Perhaps equally important, this approach can be implemented in a manner similar to the standard asset allocation procedures currently in place, with little added complexity.

Investment Risk

In his seminal contribution to modern portfolio theory, Harry M. Markowitz pointed out that investors could reduce the overall risk of their investments by forming well diversified portfolios.[3] As a part of his analysis, he considered various alternative risk measures in addition to variance and concluded that the most theoretically robust measure was **semivariance** (i.e., the expected value of the squared negative deviations about a specified "target" rate of return).

Semivariance captures the notion of downside risk and is an appropriate characterization of investment risk because investors are often concerned about losses relative to a threshold return level. It is important to note that, unlike the variance measure, semivariance does not increase with greater "upside potential." Upside potential is, rather, captured by the mean of the return distribution.

Because of purported computational problems associated with calculating the semivariance statistic, Markowitz adopted variance as the risk measure in his analyses. As a result, much of the initial research in finance focused on the issues surrounding a simpler mean-variance framework. Unfortunately, the theoretical assumptions necessary to support variance as a risk measure (discussed below) are somewhat restrictive. Moreover, variance is not consistent with investors' actual perception of risk.

Figure A Downside Risk for a Target Rate of Return of 6 Per Cent

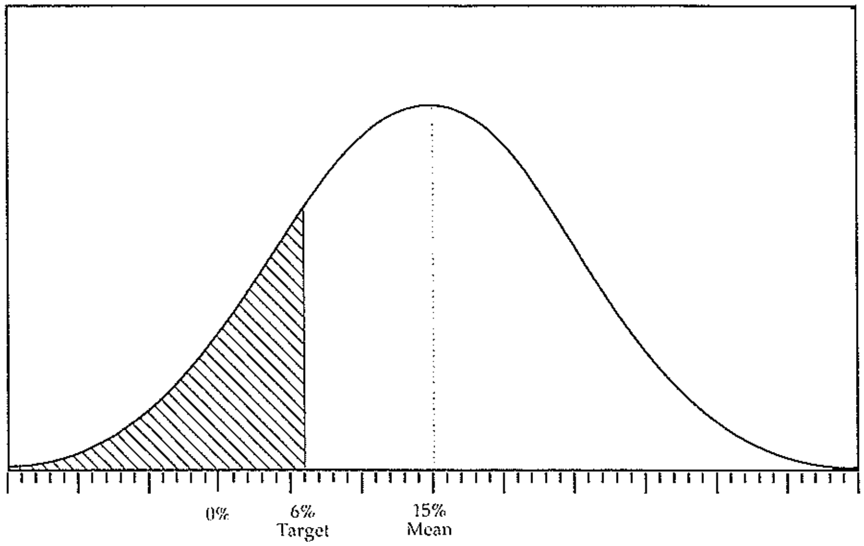

As researchers in finance, economics and psychology have noted over the past three decades, individuals view return dispersion in an asymmetric manner; that is, losses weigh more heavily than gains.[4] The examples of risk mentioned earlier—the probability of shortfall below some benchmark level of return and the expected magnitude of a loss—typify the concerns of portfolio managers. While these measures differ significantly, they both capture some element of downside risk. We focus our attention on downside-risk measures in order to alleviate some of the shortcomings of variance and establish a more general approach to asset allocation.

Downside-Risk Measures

Several classes of downside-risk measures are of particular interest in finance. All involve the tail of the relevant distribution of returns below some specific threshold level or target rate. These risk measures are referred to as "lower partial moments" (LPMs), because only the left-hand tail of the return distribution is used in the calculation.[5]

Computationally, the LPM for an empirical (discrete) distribution of portfolio returns, R_p, with a target rate, τ, is described by:

$$LPM_n = \sum_{R_p = -\infty}^{\tau} p_p(\tau - R_p)^n, \qquad (1)$$

where p_p is the probability that return, R_p, occurs.[6] The type of "moment," n, specified in Equation (1) captures an investor's preferences by determining the manner in which the return dispersion below the target is characterized.[7]

Figure A depicts a normal distribution with a mean of 15 per cent and a standard deviation of 10 per cent. For a target rate of 6 per cent, the applicable downside portion of the distribution is represented by the shaded area. For $n = 0$, the risk measure becomes a 0th-order moment (denoted LPM_0), with the term in brackets being raised to the 0th power (i.e., equal to 1). Hence the measure is simply the probability of falling below the target rate. For this particular case, LPM_0 equals 0.184. In other words, there is an 18.4 per cent chance that the return performance will fall short of the desired minimum level signified by the target rate.

For higher-order moments, the shaded area in Figure A remains pertinent for the risk calculation. However, for $n = 1$, LPM_1 becomes the expected deviation of returns below the target, or the **target shortfall**.[8] LPM_1 for this example equals 1.01 per cent. For $n = 2$, LPM_2 is analogous to variance, in that it is a probability weighting of squared deviations. Rather than computed around the mean of the distribution, however, the deviations are determined with respect to the target rate. LPM_2 can thus be referred to as a target semivariance. In our example, this measure equals 9.38 (or a 3.06 per cent target semideviation, using the square root of the LPM_2 measure).

Many popular notions of risk are special cases of the generalized LPM_n measure. For example, with n = 0 and a target rate equal to 0 per cent, LPM_0 is simply the probability of a loss. For n = 2 and a target rate equal to the mean of the distribution, LPM_2 becomes the traditional semivariance measure. Furthermore, for normal, or symmetric, distributions, LPM_2 is exactly proportional to variance [LPM_2 ($\tau = \bar{R}$) is equal to one-half of variance]. Using it as a risk measure would be equivalent to using variance and result in the same ordering of risky assets.

While the LPM_n measure of risk has obvious intuitive appeal, it is important to consider the economic justification for its use and the general conditions under which it is appropriate. The distinction between this framework and the traditional mean-variance approach lies in the assumptions regarding the distributional properties of returns and investor preferences. As mentioned earlier, the use of variance as a measure of risk requires a somewhat restrictive set of assumptions. Specifically, either returns must be normally distributed, *or* investors have to exhibit behavior describable by a quadratic utility function.[9] Within the LPM_n framework, distributions can be any one of a class characterized by a location and scale parameter (e.g., mean and standard deviation).[10] In addition, the downside-risk framework makes only general assumptions regarding investor utility functions (e.g., risk aversion and skewness preference).

The power and flexibility of the downside-risk framework stem from the *joint* set of its assumptions regarding investor preferences and asset return distributions. Unlike other frameworks, which place restrictions on preferences or on distributions, the LPM_n approach uses a combined set of reasonable and less restrictive assumptions. As a result, downside-risk analysis is not only more attractive in terms of its consistency with the way investors actually perceive risk, but it is also valid under a broader range of conditions.

Portfolio Optimization in a Downside-Risk Framework

Having established a formal definition of downside risk, we now consider its use in the context of portfolio optimization. The objective within this framework is essentially the same as in any approach—i.e., select a portfolio of assets in some combination so as to minimize risk subject to a specified level of expected return. In this case, however, LPM_n is used as the appropriate characterization of risk.

Figure B illustrates this decision problem graphically. If R_{p^*} is the investor's desired level of return, then P^* represents the portfolio that provides the lowest risk. For alternative values of R_{p^*}, the resulting set of solutions trace out a convex mean-LPM_n efficient frontier (denoted $MLPM_n$), reflecting the optimal risk/return tradeoffs in exactly the same manner as the traditional mean-variance approach.

Stated more formally, an investor who is averse to downside risk and who has a target rate of return, τ, must determine the allocation weight, X_j, for each relevant asset, j, to achieve an efficient point within the investment opportunity set. The nonlinear $MLPM_n$ optimization problem is represented by:

Select X to minimize: $LPM_n(\tau;X)$

$$= \sum_{R_p < \tau}^{\tau} p_p(\tau - R_p)^n \qquad (2)$$

Subject to: $n = 1$ or 2

$\{\Sigma_j X_j E(R_j) = R_{p^*}\}$ and

$\{\Sigma_j X_j = 1,\ X_j > 0\}$ without short sales or

$\{\Sigma_j X_j = 1\}$ with short sales,

where p_p is the probability of portfolio return, R_{p^*}.

Once again, we note that the optimization process uses the entire distribution of returns. Information contained in the right tail of the distribution (the returns above the target) does not contribute to risk, but it is captured in the mean of the distribution. Thus, all things being equal, two distributions with the same LPM_n but different means are not the same. In such a case, the distribution with the higher mean has a greater degree of positive skewness.

One necessary restriction on the optimization in Equation (2) is that LPM_n must be of first order or higher (i.e., n must be greater than zero) if any form of risk aversion is to be considered relevant in the decision-making process. The probability of target shortfall, therefore, cannot be used for the general purpose of

Figure B Downside-Risk ($MLPM_n$) Efficient Frontier

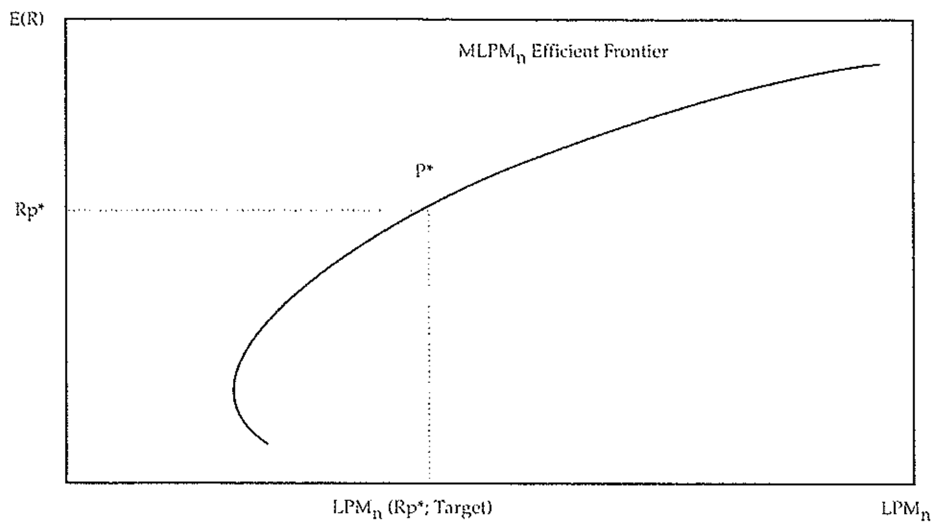

portfolio construction. The primary reason for this is that the LPM_0 measure does not differentially weight the returns in the lower tail of the distribution based on their "distance" from the target. In other words, the shortfall probability is an incomplete measure of risk, because it fails to provide any indication of how severe the shortfall will be, should it occur. As such, it is inappropriate to use this measure directly to define an efficient frontier.[11]

Note that, while the probability of target shortfall is not a complete measure of risk, it can be incorporated into the asset allocation decision. Specifically, LPM_0 provides a means of isolating the segment of the investment opportunity set most relevant to the investor. As Leibowitz, Henriksson and Kogelman have shown, the shortfall probability can be specified as a constraint in a portfolio optimization problem that utilizes some other measure of risk—for example, variance—to delineate the efficient frontier.[12] This shortfall restriction essentially provides information regarding the risk tolerance of the investor.

For the **mean-variance formulation** of Equation (2), the shortfall constraint is equivalent to an upward-sloping line in Figure B, with an intercept on the return axis representing the investor's target rate and a slope tied to the probability of falling below the target. The investor specifies both the probability and the target, thus establishing a level of aversion to below-target exposure. This constraint intersects or is tangent to the efficient frontier, indicating which portfolios are attractive to the investor. The probability of target shortfall, therefore, can be thought of as a risk-tolerance assessment tool when coupled with another measure of risk. This is true whether the measure is variance or a downside-risk criterion such as LPM_1 or LPM_2.

The appeal of the downside-risk framework is certainly not diminished because target shortfall probabilities cannot be used directly as a risk measure in a portfolio optimization problem. Indeed, beyond LPM_0, there exists a very large class of theoretically attractive and flexible risk measures. LPM_1 (target shortfall) and LPM_2 (target semivariance) provide an intuitive set of risk definitions that are more useful than traditional approaches and are valid under a broader set of economic conditions.

Downside-Risk Optimization

To appreciate the downside-risk framework more fully, it is useful to examine its implementation under alternative formulations. As with any asset allocation decision, the analysis can proceed in a variety of ways, based on the manner in which informational inputs are developed. In particular, the optimization problem stated in Equation (2) can be solved using explicit return forecasts with probabilities of realization derived through some independent modeling process. Alternatively, the analysis can be accomplished using historical data as proxies for *ex ante* asset behavior.

Equation (2) can be modified to utilize historical returns, as follows:

Select X to minimize: $LPM_n(\tau;X)$

$$= \sum_{R_p < \tau}^{\tau} \frac{1}{T-1} (\tau - R_p)^n \tag{3}$$

Subject to: $n=1$ or 2

$\{\Sigma_j X_j E(R_j) = R_{p^*}\}$ and

$\{\Sigma_j X_j = 1, X_j > 0\}$,

where T is the number of return observations. In other words, within the optimization process, LPM_n can be computed by constructing an empirical portfolio distribution given a set of asset weights, X, and T previously observed returns. This value is then minimized subject to an expected return constraint, R_{p^*}. The expected return for each asset, $E(R_j)$, is determined using its historical mean risk premium in conjunction with the current risk-free rate of interest, R_{f0}. That is, the asset's excess return at time t, $R_{jt} - R_{ft}$, is first averaged over T observations. This value is then added to R_{f0} to obtain the expected return, so that $E(R_j) = [\Sigma_T (R_{jt} - R_{ft})/T] + R_{f0}$.

Keep in mind that, for purposes of this discussion, we are interested primarily in evaluating the usefulness of an alternative portfolio construction technique, as opposed to a method of producing return forecasts. As most asset allocation procedures rely to some extent on the information contained in historical returns, the problem formulated in Equation (3) is representative of typical implementations and is thus the focus of our attention.

We consider a global asset allocation problem relevant to many portfolio managers. We include fully currency-hedged equity and fixed income markets in 11 countries—the United States, the United Kingdom, Japan, West Germany, Switzerland, France, the Netherlands, Sweden, Australia, Canada and Hong Kong. The returns used in this analysis span the 11-year period from January 1980 to December 1990.

Figure C presents actual $MLPM_n$ efficient frontiers for an investor who defines risk as any loss of initial capital and thus has a target rate of return, τ, equal to 0 per cent. These frontiers are computed using Equation (3) and 60 months of recent historical return data (i.e., $T = 60$) to derive estimates of the LPM_n risk measure and expected asset returns. Panel A depicts the $MLPM_1$ (target shortfall) frontier, while Panel B illustrates the $MLPM_2$ frontier using the same target rate of 0 per cent. In the latter case, the square root of the LPM_2 risk value is plotted to provide a measure comparable to percentage return (i.e., target semideviation). Both panels also indicate the risk and return of a global capitalization-weighted benchmark composed of 60 per cent equity and 40 per cent fixed income securities, based on the Salomon-Russell Global Equity Index and the Salomon Brothers World Government Bond Index.[13]

In Panel A, the 60/40 benchmark has an expected return of 13.5 per cent and a target shortfall of 9.07 per cent. In other words, because the investor's target rate has been specified as 0 per cent, the magnitude of the expected probability-weighted loss is 9.07 per cent. By comparison, an $MLPM_1$ efficient portfolio delivering the same benchmark expected return of 13.5 per cent has only a 4.53 per cent expected target shortfall. In Panel B, which uses target semideviation to define the frontier, the risk of the benchmark is estimated at 23.20 per cent, while the risk of the comparable $MLPM_2$ efficient portfolio is projected at 12.40 per cent.

Figure D indicates the effect that altering the target rate of return, τ, has on the magnitude of the estimated risk. For ease of comparison, we have included the previous frontiers, with target rates equal to 0 per cent as well as those constructed using target rates of 8 and 16 per cent. For the latter two cases, the specification of risk is based on measures of dispersion that capture underperformance relative to higher benchmark target levels of return. While the investor with a target return equal to 0 per cent is concerned with the loss of any principal, these higher targets represent the risk of returning less than the risk-free asset ($\tau = 8\%$) or less than the market portfolio ($\tau = 16\%$).

Inspection of the figures indicates that, as the target rate increases, both the $MLPM_1$ and $MLPM_2$ frontiers shift to the right. This occurs because more of the distribution of returns falls below the specified target rate. The downside component of returns therefore becomes larger, increasing the numerical value of the risk mea-

Figure D Downside – Risk Efficient Frontiers for Alternative Target Rates of Return

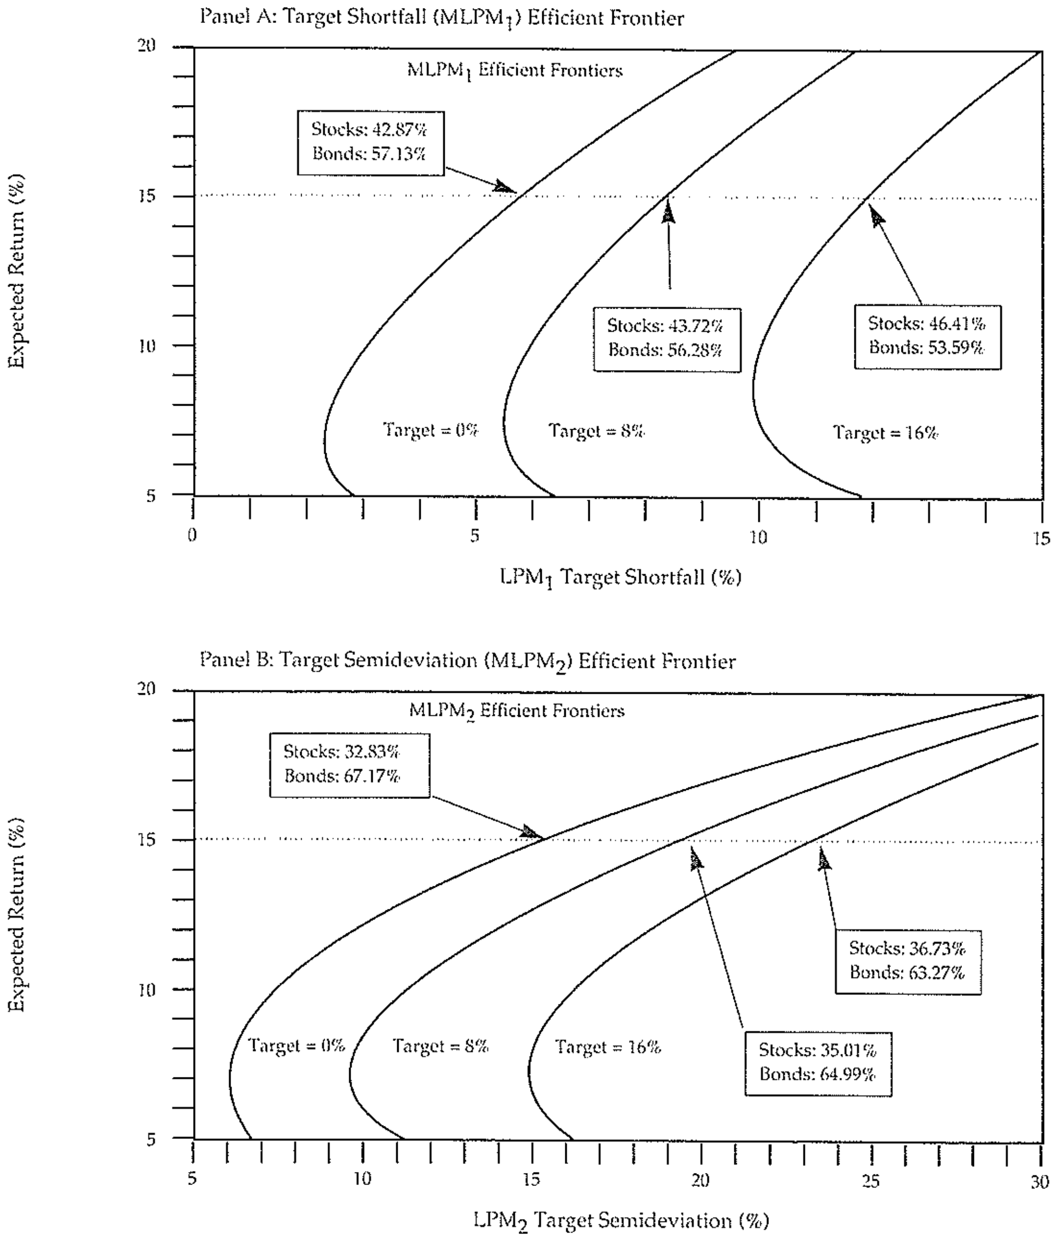

Simply changing the focus of attention from target shortfall to target semideviation dramatically alters the composition of the investor's portfolio. Indeed, this choice has a much larger effect on the allocation decision than the choice of a target rate for a specific type of risk measure.

The actual benefits of the downside-risk framework relative to the traditional mean-variance setting must be determined empirically. If return distributions are approximately normally distributed, then variance is a sufficient measure of risk, and differences between the two approaches will be small. But if returns are not symmetrically distributed about the mean, but instead possess some degree of skewness, then the asset allocation decision reached within the downside-risk framework can be