by Martin L. Leibowitz and William S. Krasker

The Persistence of Risk: Stocks versus Bonds over the Long Term

One of the most common tenets of investing holds that the asset with the highest expected return over the long run is virtually certain to provide superior performance. Thus investors are willing to bear the greater short-term risks associated with equities as compared with bonds, say, because they know the equities will eventually outperform the bonds.

But risk doesn't disappear over time. According to standard analytical models, an equity portfolio is virtually certain to outperform a fixed income portfolio over a long enough period. But 10, or even 30, years may not be "long enough."

A simple model using standard assumptions about asset volatilities and risk premiums, for example, shows that a stock portfolio has a 32 per cent chance of underperforming a bond portfolio over a 10-year horizon. Even after 30 years, there remains a substantial 21 per cent probability that stocks will fall short of bonds.

INVESTMENT DECISIONS are often framed in terms of the appropriate tradeoff between risk and expected return. A further dimension arises when the investor must balance long-term goals with short-term concerns. Greater short-term risk is viewed as the price of better "long-term" returns. Long-term returns seem to be naturally associated with better "expected returns" and with the hope that short-term risks will fade when compared with the return growth over the long run.

Such long-term return/short-term risk rationales underlie many investment situations, including stock/bond allocations, yield curve maturity selections, yield pickup bond swaps, and high-yield/low-yield currency preferences. In all these situations, a key question is whether the greater expected returns will eventually overwhelm the impact of the associated risks. For most investors, intuition implies that they will. However, given the standard analytical models, we can show that there are many cases in which this intuition has theoretical justification only for the very long "long run." Over the more relevant intermediate term, the standard model gives rise to a more complex pattern that includes some strikingly counterintuitive results.

If we extrapolate the standard input assumptions, an equity portfolio is virtually certain to outperform any fixed income investment over a long enough period. The question is what constitutes "long enough." Most investors would be surprised to learn that these assumptions imply that, after 10 years, there is a 32 per cent probability that a stock portfolio will underperform a bond benchmark. After 20 years, the shortfall probability is still slightly above 25 per cent. Even more striking, after 30 years the probability that bonds will outrun stocks remains a hefty 21 per cent. These results follow directly from the evolution over time of the theoretical probability distribution for relative returns.

In general, the standard models focus on the variance of the short-term return as a measure of risk. Because our analysis is based on such models, the only long-term risk discussed below is the accumulated impact of this variability. It should be pointed out, however, that these standard models are patently simplistic. There are more facets to return than a constant, nominal-dollar expected return, and there are surely more dimensions to long-term risk than the

Martin Leibowitz is Managing Director, Bond Portfolio Analysis Group, and William Krasker is a Vice President at Salomon Brothers Inc.

accumulation of short-term variability (for example, credit risks, disaster scenarios, major secular reversals). Nevertheless, despite their limitations, these standard models have a widespread impact in practice, where they are both applied explicitly and form an implicit basis for intuitive decisions. It is therefore important to trace more fully their implications for the persistence of shortfall risks over the two to 30-year intermediate term that forms the relevant period for most of us mortals and our institutions.

Distribution of Returns Over Time

We base our analysis on a theoretical model of asset returns, in which successive returns are independent and have lognormal probability distributions. Specifically, the logarithmic return—the logarithm of one plus the ordinary return—has the familiar bell-shaped normal distribution. Figure A depicts the resulting lognormal distribution for the value of an equity portfolio in one year, assuming that its initial value is 100, the expected return is 13 per cent and the logarithmic volatility—the standard deviation of the logarithmic return—is 18 per cent.[1]

At the end of the one-year period, the expected value (or mean) of the portfolio is 113. As shown, there is a 0.25 probability that the portfolio value will be above 125.5 and a 0.25 probability that the portfolio value will be below 98.5. In addition, there is a 0.05 probability that the portfolio value will exceed 149.5, but also a 0.05 probability that it will fall below 82.7. Moreover, the 50th percentile (the median) occurs at 111.2. This value differs slightly from the mean of 113, reflecting a modest skewness in the distribution.

Figure B introduces an alternative graphic representation of the percentiles to illustrate how the probability distribution for the value of the equity portfolio changes as the number of years in the fund horizon increases. The boxes represent the distributions at one, two and three-year horizons and show explicitly the expected value as "handlebars," as well as the fifth, 25th, 50th, 75th and 95th percentiles of the distribution. (The percentiles in the first-year box correspond with those in Figure A.)

Although the expected portfolio value increases steadily as the horizon is lengthened, the variability also rises; the net effect on the downside risk thus depends on the particular percentile selected. For example, if we measure the risk by the 25th percentile—the level below which the portfolio has only a 0.25 probability of falling—this return threshold rises from one year to the next. If we focus instead on the fifth percentile—the bottom of the boxes—the return threshold value remains nearly the same.

The Stock/Bond Ratio

Figures A and B depict the absolute returns from all-equity portfolios. Figure C compares the probability distributions for the value of the equity portfolio at different horizons with the

1. Footnotes appear at end of article.

Figure A Distribution for Stock Portfolio Value in One Year

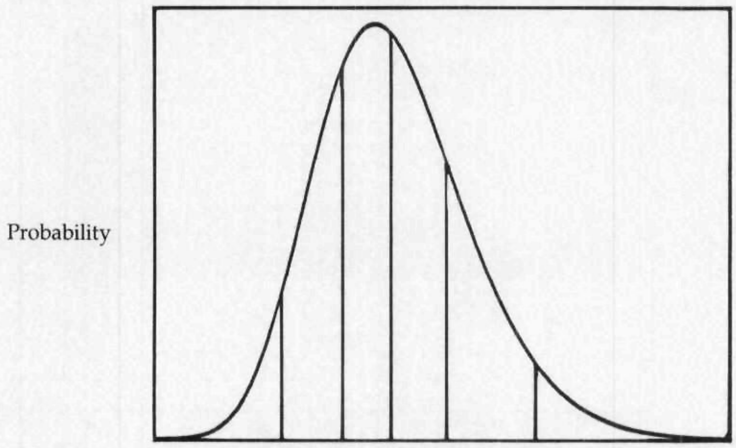

Figure B Distribution for Stock Portfolio Value at Different Horizons

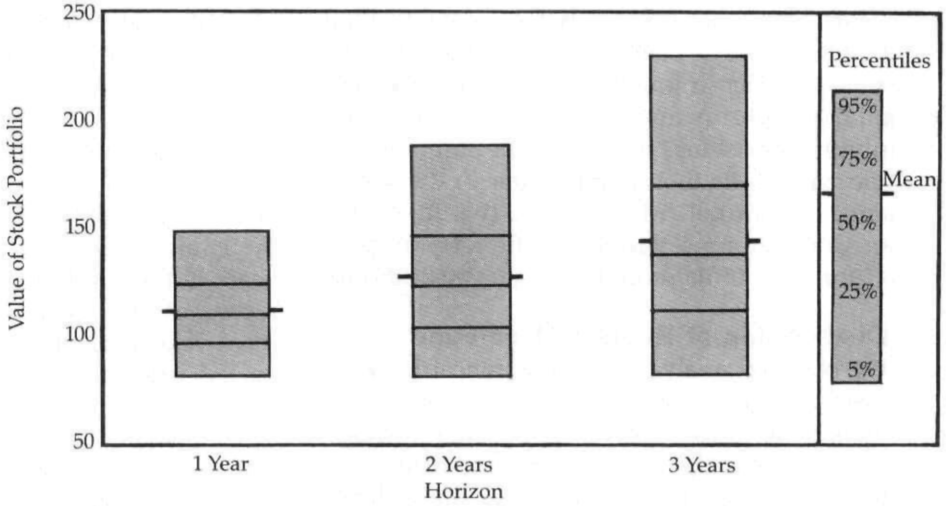

value of a bond portfolio having the same starting value of 100. The bond portfolio has an expected return of 9 per cent and a logarithmic volatility of 10 per cent, roughly corresponding to a portfolio with a six-year effective duration (or a par bond maturity of approximately 10 years at the outset).

Figure C shows that although the expected portfolio value is smaller for bonds, the risk is also substantially lower. At a three-year horizon, the bond value exceeds 95.9 with a 0.95 probability (that is, the bottom of the box). For this same probability, the wider stock distribution can provide a threshold of only 82.3.

With constant absolute levels of expected returns, these results cannot be extended beyond short horizons. Indeed, the very existence of uncertainty about bond returns implies changes in yields, which presumably imply changes in expected returns. Therefore, for the remainder of this article we shall assume that it is the *relative* risk premium that remains constant. More precisely, for technical reasons, we will keep the difference between the expected stock logarithmic return and the expected bond logarithmic return at a constant level that closely approximates a 4 per cent ordinary one-year risk premium. In addition, we will assume the

Figure C Stock and Bond Distributions

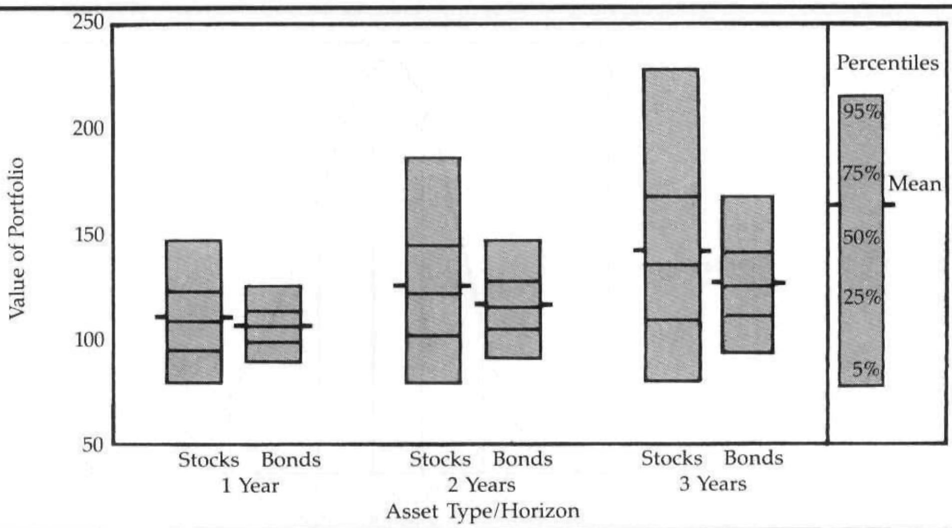

Figure D Ratio of Stock Portfolio Value to Bond Portfolio Value

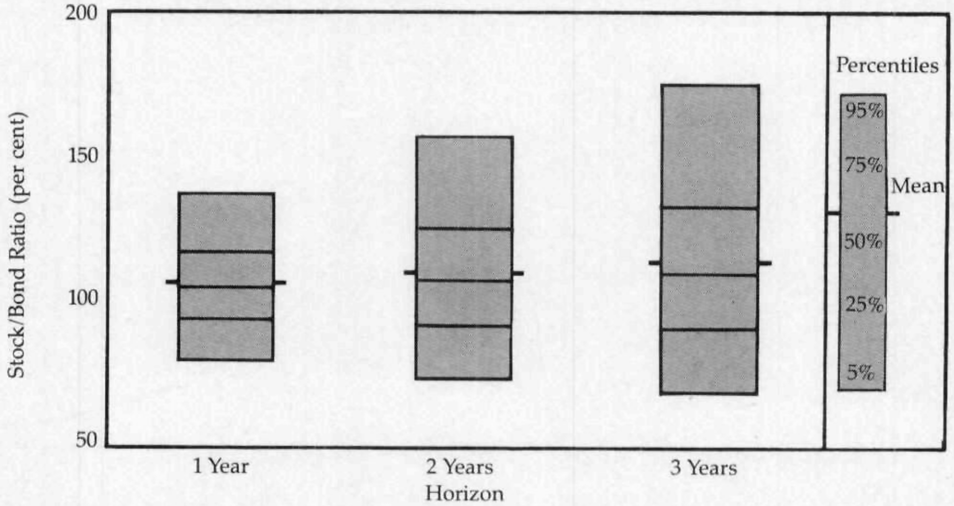

same asset volatilities used earlier and make a more or less standard correlation assumption of 0.4 between the stock and bond logarithmic returns.

Now consider the ratio of the cumulative value of a 100 per cent stock portfolio to the cumulative value of a 100 per cent bond portfolio. Figure D shows the probability boxes for the stock/bond ratio for horizons of up to three years. A ratio of 100 per cent means that the stock and bond portfolios are equal, while ratios greater than 100 per cent imply that the stock value exceeds the bond value. Clearly, on average, a pension fund would achieve a better performance from stocks over a three-year period: The expected value of the ratio of stocks to bonds grows to 112.4 per cent by the third year. However, the figure also shows that there is substantial risk at the three-year horizon. At that date, there is a 25 per cent probability that the ratio of stock value to bond value will be less than 88.6 per cent.

The probability that the stock value will fall below the bond value remains high even for what would normally be regarded as the long run. Figure E shows this shortfall probability as a function of the number of years in the horizon. At a 20-year horizon, there is still a 25 per

Figure E Probability that Bond Portfolio Value Exceeds Stock Portfolio Value

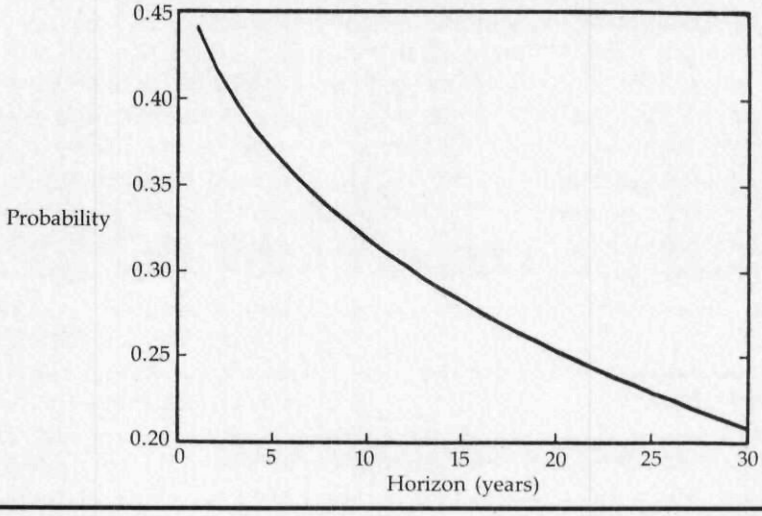

Figure F Distribution for Stock/Bond Ratio in 20 Years

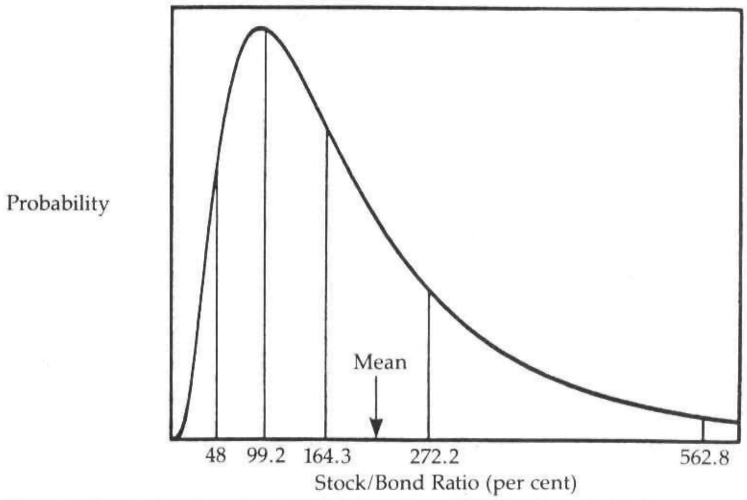

cent probability that the bond value will exceed the stock value. At the same time, the expected value of the stock/bond ratio grows to an attractive 217 per cent over 20 years. Clearly, for these long horizons, the prospect of exciting (and even stellar) relative performance appears to coexist with the significant probability of shortfalls. Although these precise numerical results depend on the selected volatilities, correlation and expected change in the stock/bond ratio, the qualitative conclusions about the persistence of risk hold for any reasonable range of these parameters.

The presence of a high expected value for the stock/bond ratio, despite a substantial probability that the ratio will be less than 100 per cent, is a manifestation of the pronounced skewness that develops as time passes. As Figure F shows, the distribution at the 20-year horizon simply does not have the symmetry that most people tend to visualize. The stock/bond ratio is considerably more likely to be below its expected value of 217 per cent than above, and, in fact, the median of the distribution is only 164 per cent. Offsetting this, however, is the fact that the distribution has a very "long tail" on the upside. For example, there is a 5 per cent chance that the ratio will exceed 563 per cent.

Shortfall Probabilities for Specified Stock/Bond Ratios

One way to evaluate the riskiness of stocks relative to bonds—or to pension liabilities, which have roughly similar risk characteristics in a nominal-dollar framework—is to choose a particular shortfall level and determine the probability that the ratio of stock value to bond value will fall below that level. Figure G shows this shortfall probability as a function of the years in the horizon for several specified stock/bond ratios. (The case in which the ratio equals 1.0 is the same as that in Figure F.)

The curves slope downward when the number of years to the horizon becomes very large. This is a consequence of the intuitive notion that, in the long run, the higher expected return of equities must dominate its larger standard deviation. However, Figure G contains two surprises. First, for certain outcomes, such as a ratio of 90 or 80 per cent, the probability actually increases at the beginning. (This can be interpreted in terms of the minimum time span required for bad shortfalls to develop.) This growth in risk continues for several years: The risk of falling below a 60 per cent shortfall level is substantially higher at a 10-year, or even a 15-year, horizon than at a one-year horizon. The second surprising finding, which emerges most clearly in the graph corresponding to the 90 per cent shortfall level, is that even when the curve becomes downward sloping, it takes a very long time to reach a low probability level.

Suppose, for example, that a pension fund manager views a funding ratio of 80 per cent of its current level as an unacceptable shortfall and would like to know the probability that this level will be penetrated. According to Figure G, the probability of such a shortfall is only 7 per

Figure G Shortfall Probabilities for Specified Stock/Bond Ratios

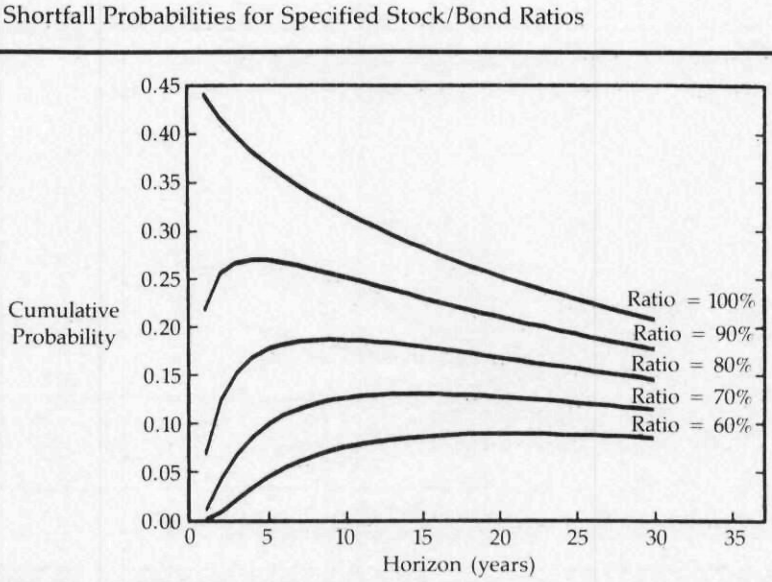

cent if the horizon is one year, but it rises to 19 per cent if the horizon is 10 years. The probability remains at a high 17 per cent even with a 20-year horizon.

Stock/Bond Percentiles Over Time

Another way to measure the shortfall risk is by percentiles of the probability distribution for the stock/bond ratio. Specifically, we choose a particular probability of a shortfall and then compute the stock/bond ratio to which it corresponds.

Figure H shows these percentile ratios as a function of the number of years in the horizon for several different shortfall probabilities. Once again, intuition would at first suggest that these shortfall percentiles should improve with longer time horizons. Specifically, for a given probability of shortfall, we might expect the shortfall level to be an increasing function of the time to the horizon.

This turns out to be consistently true only after a transition period, which can be surprisingly long. Over relevant horizons, such as 10 or 20 years, the relationship can be just the opposite. For example, the 10th percentile of the stock/bond ratio does not reach its minimum—63 per cent—until the horizon is 19 years.

Figure H Percentiles of Stock/Bond Ratio

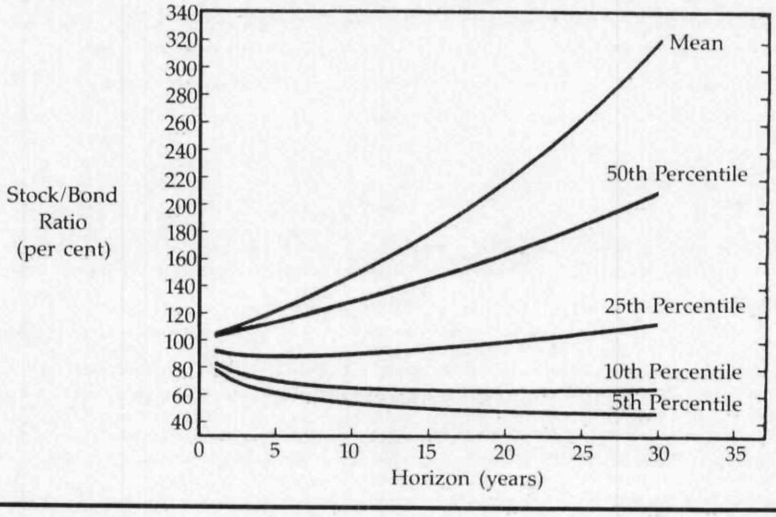

Figure I Shortfall Levels for Different Risk Premiums

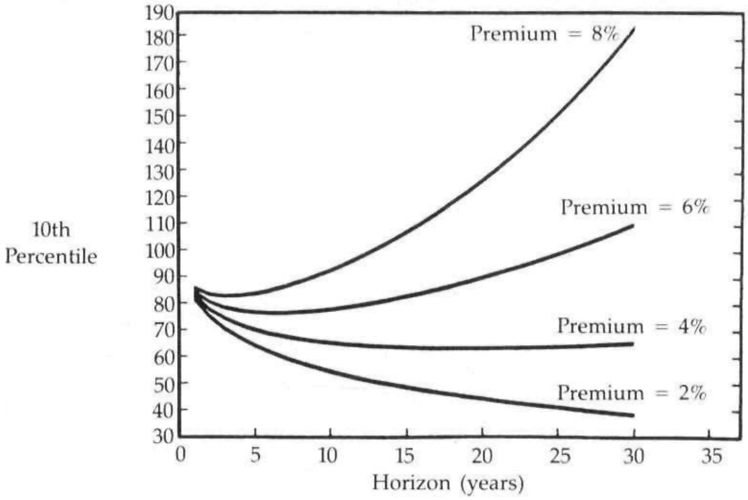

Effects of Different Risk Premiums and Correlations

The preceding analysis has been based on our assumption that the logarithmic risk premium is constant at a level that implies an ordinary risk premium of approximately 4 per cent. Not surprisingly, the results are sensitive to the value that is assumed for this risk premium.

Figure I graphs the 10th-percentile shortfall level as a function of the number of years in the horizon for several different logarithmic risk premiums. The curves are labeled by the ordinary risk premiums to which they correspond when the bond expected return is 9 per cent. (The curve for the risk premium of 4 per cent matches the 10th-percentile curve in Figure H.)

In effect, the larger the risk premium, the shorter the time period necessary for the "long-run" effects to take hold. For example, with the original risk premium of 4 per cent, the 10th percentile of the stock/bond ratio does not reach its minimum until the horizon reaches 19 years. However, with an 8 per cent risk premium, the 10th percentile achieves its smallest value at a three-year horizon.

In the standard asset allocation problem, lower correlations represent a desirable opportunity for diversification. However, when one is dealing with relative returns—either relative to a benchmark index or to a liability framework—it turns out that higher correlations are desirable.[2] These results are illustrated in Figure J, which shows the acceleration of the "long-run" benefits for the hypothetical case of perfect correlation (as well as the deceleration for the case of zero correlation). These correlation effects become significant in various bond-to-bond comparisons, or when a bond portfolio is compared with an interest-rate-related benchmark.

Conclusion

These results are based on a highly simplistic model, in which the probability distribution for percentage changes in the stock/bond ratio is constant over time. There is no explicit treatment for inflation, macroeconomic cycles or profit growth. We assume a rigid long-term adherence to the asset class, with all incremental returns implicitly reinvested in the original asset structure. There is no provision for active management or portfolio-rebalancing strategies, which could reduce the likelihood of the most adverse outcomes.

In other words, these results could be nothing more than a mathematical curiosity spawned by pushing short-term models too far into the future. Consequently, considerable caution should be exercised when taking the extrapolations of such models as a literal characterization of long-term return prospects for any asset class.

Nevertheless, our core finding sheds important light on certain perceptions that have become embedded in our investment mythology. One such article of faith is that a steadfast adherence to the "asset of choice" with the highest expected return will, over time, virtual-

Figure J Shortfall Levels for Different Correlations

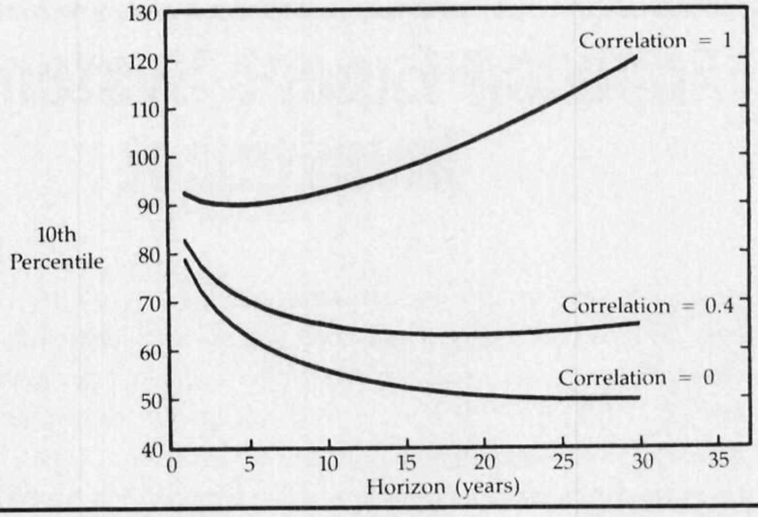

ly ensure a superior performance. However, over long time periods, even while return expectations grow ever more favorable, risks refuse to fade. Risk persists—and at surprisingly high levels of significance. Even the most determinedly long-run investor—one who has steeled himself to persevere relentlessly through repeated bouts of short-term turmoil—will find that he must still come to grips with long-term risk and its continuing uncertainty. ■

Footnotes

1. The ordinary expected return, R, the expected logarithmic return, μ, and the volatility, σ, are related by the following formula:

   $$\log (1 + R) = \mu + \tfrac{1}{2}\sigma^2.$$

2. See "Portfolio Optimization Under Shortfall Constraints: A Confidence Limit Approach to Managing Downside Risk" (Salomon Brothers Inc, New York, August 1987).

Ennis footnotes concluded from page 27.

grams in that year. The regression coefficient is actually a small negative value—bond ratings falling as funded ratios rise—but not significantly different from zero. Johnson & Higgins is the source of pension funding data. Fiduciary Management Trust Co., Boston, provided general obligation rates.

7. Between 1983 and 1986, the median state contribution to the PERSs in the 1987 Greenwich Associates survey declined from 18.3 to 16.2 per cent of payroll. This decline coincided with a period of sharply rising asset values.
8. The slope of the regression is −0.096 of the natural log of liability (in billions). The t-statistic of the slope coefficient is 3.6. Data for 42 states were taken from Johnson & Higgins' *1986 Pension Commission Clearinghouse Report on State Pension Commissions.* This survey contained data for 44 of the 50 states; two observations were deleted because they appeared to be spurious.
9. Although most PERS members make regular deposits with their PERS, they are not investors. They are using pretax dollars to purchase annuities that are subsidized by the state (taxpayers). Their deposits are credited with interest to compensate them in the event they withdraw from covered service before they vest.
10. The median and average values of pension assets of the 50 states were approximately \$3.5 billion and \$8.6 billion, respectively, in 1986.