What Practitioners Need To Know . . .

. . . About Return and Risk

Mark Kritzman
Windham Capital Management

At first glance, return and risk may seem to be straightforward concepts. Yet closer inspection reveals nuances that can have important consequences for determining the appropriate method for evaluating financial results. This column reviews various measures of return and risk with an emphasis on their suitability for alternative uses.

Return

Perhaps the most straightforward rate of return is the *holding-period return*. It equals the income generated by an investment plus the investment's change in price during the period the investment is held, all divided by the beginning price. For example, if we purchased a share of common stock for $50.00, received a $2.00 dividend, and sold the stock for $55.00, we would have achieved a holding-period return equal to 14.00%. In general, we can use Equation (1) to compute holding-period returns.

Eq. 1

$$HPR = (I + E - B)/B$$

where

HPR = holding-period return,
I = income,
E = ending price and
B = beginning price.

Holding-period returns are also referred to as *periodic returns*.

Dollar-Weighted versus Time-Weighted Rates of Return

Now let us consider rates of return over multiple holding periods. Suppose that a mutual fund generated the following annual holding-period returns from 1988 through 1992:

1988	−5.00%
1989	−15.20%
1990	3.10%
1991	30.75%
1992	17.65%

Suppose further that we had invested $75,000 in this fund by making contributions at the beginning of each year according to the following schedule:

1988	$5,000
1989	$10,000
1990	$15,000
1991	$20,000
1992	$25,000

By the end of 1992, our investment would have grown in value to $103,804.56. By discounting the ending value of our investment and the interim cash flows back to our initial contribution, we can determine the investment's *dollar-weighted rate of return*, which is also referred to as the *internal rate of return*:

$$5000 = -\frac{10000}{(1+r)} - \frac{15000}{(1+r)^2} - \frac{20000}{(1+r)^3} - \frac{25000}{(1+r)^4} + \frac{103805}{(1+r)^5},$$

DWR = 14.25%

We enter the interim contributions as negative values, because they are analogous to negative dividend payments. Although we cannot solve directly for the dollar-weighted rate of return, most financial calculators and spreadsheet software have iterative algorithms that quickly converge to a solution. In our example, the solution equals 14.25%.

The dollar-weighted rate of return measures the annual rate at which our cumulative contributions grow over the measurement period. However, it is not a reliable measure of the performance of the mutual fund in which we invested, because it depends on the timing of the cash flows. Suppose, for example, we reversed the order of the contributions. Given this sequence of contributions, our investment would have grown to a higher value—$103,893.76. The dollar-weighted rate of return, however, would have been only 9.12%:

$$25000 = -\frac{20000}{(1+r)} - \frac{15000}{(1+r)^2} - \frac{10000}{(1+r)^3} - \frac{5000}{(1+r)^4} + \frac{103894}{(1+r)^5},$$

DWR = 9.12%

In order to measure the underlying performance of the mutual fund, we can calculate its *time-weighted rate of return*. This measure does not depend on the timing of cash flows.

We compute the time-weighted rate of return by first adding one to each year's holding-period return to determine the return's *wealth relative*. Then we multiply the wealth relatives together, raise the product to the power 1 divided by the number of years in the measurement period, and subtract 1. Equation (2) shows this calculation:

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Eq. 2

$$TWR = \left[\prod_{i=1}^{n} (1 + HPR_i)\right]^{1/n} - 1$$

where

TWR = time-weighted rate of return,
HPR_i = holding-period return for year i and
n = number of years in measurement period.

If we substitute the mutual fund's holding-period returns into Equation (2), we discover that the fund's time-weighted rate of return equals 5.02%.

The time-weighted rate of return is also called the *geometric return* or the compound annual return. Although the geometric return and the compound annual return are often used interchangeably, technically the geometric return pertains to a population whereas the compound annual return pertains to a sample. I use the term geometric return to refer to both. It is the rate of return that, when compounded annually, determines the ending value of our initial investment assuming there are no interim cash flows. For example, suppose we invest $10,000 in a strategy that produces a holding-period rate of return of 50% in the first year and −50% in the second year. At the end of the second year, we will end up with $7500. The geometric return over the two-year measurement period equals −13.40%:

$$[(1 + 0.5)(1 - 0.5)]^{1/2} - 1 = -0.1340$$

If we multiply $10,000 times (1 − 0.1340) and then multiply this result again by (1 − 0.1340), we arrive at the ending value of this investment—$7500.00.

In order to manipulate geometric returns, we must first convert them to wealth relatives and raise the wealth relatives to a power equal to the number of years in the measurement period. Then we multiply or divide the cumulative wealth relatives, annualize the result, and subtract one in order to convert the result back into a geometric return.

Suppose, for example, that five years ago we invested $100,000 in a fund that we thought would earn a geometric return of 8.00% over a 20-year horizon so that, at the end of the horizon, we would receive $466,095.71. During the past five years, however, the fund's geometric return was only 6.50%. What must its geometric return be for the next 15 years if we are to reach our original goal of $466,095.71? We start by raising 1.08 to the 20th power, which equals 4.6609571. This value equals the cumulative wealth relative of the anticipated geometric return. We then raise the wealth relative of the geometric return realized thus far to the 5th power, which equals 1.3700867. We then divide 4.6609571 by 1.3700867, raise this value to the power 1 over 15, and subtract 1, to arrive at 8.50%.

Alternatively, we can convert wealth relatives to continuous returns by taking their natural logarithms and manipulating these logarithms. For example, the natural logarithm of 1.08 equals 0.076961, and the natural logarithm of 1.065 equals 0.062975. We multiply 0.062975 times 5 and subtract it from 0.076961 times 20, which equals 1.2243468. Then we divide this value by 15 and use the base of the natural logarithm 2.718281 to reconvert it to an annual wealth relative, which equals 1.085.[1]

Geometric Return versus Arithmetic Return

It is easy to see why the geometric return is a better description of past performance than the arithmetic average. In the example in which we invested $10,000 at a return of 50% followed by a return of −50%, the arithmetic average overstates the return on our investment. It did not grow at a constant rate of 0%, but declined by 13.40% compounded annually for two years. The arithmetic average will exceed the geometric average except when all the holding-period returns are the same; the two return measures will be the same in that case. Furthermore, the difference between the two averages will increase as the variability of the holding-period returns increases.

Table I Annual S&P 500 Returns

1983	22.63%	1988	16.58%
1984	5.86%	1989	31.63%
1985	31.61%	1990	−3.14%
1986	18.96%	1991	31.56%
1987	5.23%	1992	7.33%

If we accept the past as prologue, which average should we use to estimate a future year's expected return? The best estimate of a future year's return based on a random distribution of the prior years' returns is the arithmetic average. Statistically, it is our best guess for the holding-period return in a given year. If we wish to estimate the ending value of an investment over a multiyear horizon conditioned on past experience, however, we should use the geometric return.

Suppose we plan to invest $100,000 in an S&P 500 index fund, and we wish to estimate the most likely value of our investment five years from now. We assume there are no transaction costs or fees, and we base our estimates on yearly results ending December 31, 1992, which are shown in Table I. The arithmetic average equals 16.83%, while the geometric average equals 16.20%. Our best estimate for next year's return, or any single year's return for that matter, equals 16.83%, because there is a 1-in-10 chance of experiencing each of the observed returns. However, the best estimate for the terminal value of

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

our fund is based on the geometric average. It equals $211,866.94, which we derive by raising 1.1620 to the 5th power and multiplying this value by $100,000.

Here is what we mean by "best estimate." Consider an experiment in which we randomly choose five returns from a sample of 1000 returns and calculate the terminal value of $1 invested in these five returns. We repeat this experiment 200 times. The best estimate is the value that corresponds to the typical or most likely terminal wealth of these experiments. It equals 1 plus the geometric average of the 1000 returns raised to the 5th power.

Risk

Although the time-weighted rate of return measures the constant annual rate of growth that determines terminal wealth, it is nonetheless limited as a measure of performance because it fails to account for risk. We can adjust returns for risk in several ways. One approach is to compute a portfolio's return in excess of the riskless return and to divide this excess return by the portfolio's standard deviation. This risk-adjusted return, called the *Sharpe measure*, is given by Equation (3).[2]

Eq. 3

$$S = \frac{(R_P - R_F)}{\sigma_P}$$

where

S = the Sharpe measure,
R_P = portfolio return,
R_F = riskless return and
σ_P = the standard deviation of portfolio returns.

Because it adjusts return based on total portfolio risk, the implicit assumption of the Sharpe measure is that the portfolio will not be combined with other risky portfolios. Thus the Sharpe measure is relevant for performance evaluation when we wish to evaluate several mutually exclusive portfolios.

The Capital Asset Pricing Model (CAPM) assumes that risk consists of a systematic component and a specific component. Risk that is specific to individual securities can be diversified away, hence an investor should not expect compensation for bearing this type of risk. Therefore, when a portfolio is evaluated in combination with other portfolios, its excess return should be adjusted by its systematic risk rather than its total risk.[3]

The *Treynor measure* adjusts excess return for systematic risk.[4] It is computed by dividing a portfolio's excess return, not by its standard deviation, but by its beta, as shown in Equation (4).

Eq. 4

$$T = \frac{(R_P - R_F)}{\beta_P},$$

where

T = the Treynor measure,
R_P = portfolio return,
R_F = riskless return and
β_P = portfolio beta.

We can estimate beta by regressing a portfolio's excess returns on an appropriate benchmark's excess returns. Beta is the coefficient from such a regression. The Treynor measure is a valid performance criterion when we wish to evaluate a portfolio in combination with the benchmark portfolio and other actively managed portfolios.

The intercept from a regression of the portfolio's excess returns on the benchmark's excess returns is called alpha. Alpha measures the value-added of the portfolio, given its level of systematic risk. Alpha is referred to as the *Jensen measure*, and is given by Equation (5).[5]

Eq. 5

$$\alpha = (R_P - R_F) - \beta_P(R_B - R_F),$$

where

α = the Jensen measure (alpha),
R_P = portfolio return,
R_F = riskless return,
β_P = portfolio beta and
R_B = benchmark return.

The Jensen measure is also suitable for evaluating a portfolio's performance in combination with other portfolios, because it is based on systematic risk rather than total risk.

If we wish to determine whether or not an observed alpha is due to skill or chance, we can compute an *appraisal ratio* by dividing alpha by the standard error of the regression:

Eq. 6

$$A = \frac{\alpha}{\sigma_e},$$

where

A = the appraisal ratio,
α = alpha and
σ_e = the standard error of the regression (nonsystematic risk).

The appraisal ratio compares alpha, the average nonsystematic deviation from the benchmark, with the nonsystematic risk incurred to generate this performance. In order to estimate the likelihood that an observed alpha is not due to chance, we can test the null hypothesis that the mean alpha does not differ significantly from 0%. If we reject the null hypothesis, the alpha is not due to chance.

Suppose, for example, that a portfolio's alpha equals 3% and that its standard error equals 4%, so that the appraisal ratio equals 1.33. If we look up this number in

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

a t distribution table, we discover that, given the amount of nonsystematic risk, there is a 10% chance of observing an alpha of this magnitude by random process. Hence, we would fail to reject the null hypothesis that alpha does not differ significantly from 0%.

Downside Risk

In the previous example, we based the probability estimate on the assumption that alpha is normally distributed. This assumption is reasonable for evaluating returns over short measurement periods. Over multiple-year measurement periods, however, it is the logarithms of the wealth relatives that are normally distributed. The geometric returns themselves are lognormally distributed, which means that they are positively skewed. Therefore, in order to estimate the likelihood of experiencing particular outcomes over multiple-year horizons, we should calculate the normal deviate based on the mean and standard deviation of the logarithms of the wealth relatives.

Some investment strategies produce distributions that are skewed differently from a lognormal distribution. Dynamic trading strategies such as portfolio insurance or strategies that involve the use of options typically generate skewed distributions.

A distribution that is positively (right) skewed has a long tail above the mean. Although most of the outcomes are below the mean, they are of smaller magnitude than the fewer outcomes that are above the mean. A distribution that is negatively (left) skewed has a long tail below the mean. It has more outcomes above the mean, but they are smaller in magnitude than those below the mean.

Skewness is calculated as the sum of the probability-weighted cubed deviations around the mean. Risk-averse investors prefer positive skewness, because there is less chance of large negative deviations.

When we are evaluating strategies that have skewed distributions other than lognormal, we cannot rely on only return and standard deviation to estimate the likelihood of achieving a particular result. Nor can we compare portfolios or strategies that have different degrees of skewness using only these two characteristics. Instead, we have to specify a target return and base our evaluation on the dispersion of returns *below* this target, rather than the dispersion of returns around the mean.[6] An alternative method for dealing with skewness is to define utility as a function of mean, variance and skewness and then to evaluate portfolios and strategies based on their expected utilities.

I have attempted to shed some light on the subtleties that distinguish various measures of return and risk. When we rely on these summary statistics to evaluate past results or to predict future consequences, it is important that we understand their precise meaning.[7]

Footnotes

1. *For a review of logarithms and continuous returns, see M. Kritzman, "What Practitioners Need to Know About Lognormality,"* Financial Analysts Journal, *July/August 1992.*
2. *See W. Sharpe, "Mutual Fund Performance,"* Journal of Business, *January 1966.*
3. *For a discussion of the Capital Asset Pricing Model, see M. Kritzman, "What Practitioners Need to Know About the Nobel Prize,"* Financial Analysts Journal, *January/February 1991.*
4. *See J. Treynor, "How to Rate Management of Investment Funds,"* Harvard Business Review, *January-February 1965.*
5. *See M. Jensen, "The Performance of Mutual Funds in the Period 1945–1964,"* Journal of Finance, *May 1968.*
6. *See W. V. Harlow, "Asset Allocation in a Downside-Risk Framework,"* Financial Analysts Journal, *September/October 1991.*
7. *I thank Robert Ferguson for his helpful comments.*

Altman *footnotes concluded from page 60.*

of the S&P 500 and the Merrill Lynch High Yield Master Index.

6. *See Altman, Eberhart and Zekavat, "Priority Provisions,"* op. cit.
7. *The arithmetic unweighted index did considerably better in 1991, with a return of over 100%. This reflected some exceptionally high returns on a number of small issues.*
8. *The Merrill Lynch Master Index of High Yield Debt increased by 34.6% in 1991—also reversing relatively poor years in 1989 and 1990.*
9. *This rate does not include prior defaults in the population base of high-yield bonds; it only includes those bonds that could have defaulted during the year.*
10. *Ward and Griepentrog, "Risk and Return in Defaulted Bonds,"* op. cit.
11. *E. Altman,* Corporate Financial Distress and Bankruptcy, *2nd ed. (New York: John Wiley & Sons, 1993).*
12. *For an analysis of this market, see Altman,* The Market for Distressed Securities and Bank Loans, op. cit. *and Carlson and Fabozzi,* The Trading and Securitization of Senior Bank Loans, op. cit.
13. *I appreciate the support of the Foothill Group (Los Angeles) and Merrill Lynch & Co., as well as the assistance of Suzanne Crymes and Vellore Kishore of the NYU Salomon Center.*

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.