The Hitchhiker's Guide to Risk-adjusted Returns

By Rex Macey, CIMA®, CFA®

In *The Hitchhiker's Guide to the Galaxy*, a computer—after working for 7.5 million years—finds that 42 is the answer to "life, the universe, and everything." By its very nature, a risk-adjusted return (RAR) promises to encapsulate return and risk into a single performance measure. But just as simple answers to questions about the meaning of life tend to be elusive, no single measure succinctly describes investment performance. With that important caveat, here is a brief guide to risk-adjusted returns, including their uses and misuses.

> In *The Hitchhiker's Guide to the Galaxy*, a computer—after working for 7.5 million years—finds that 42 is the answer to "life, the universe, and everything."

The Sharpe Ratio

Due to its simplicity or to the pedigree of its Nobel Laureate creator, the Sharpe ratio is the most commonly used RAR. William Sharpe introduced his measure more than 40 years ago (Sharpe 1966). As it is used today, the Sharpe ratio is defined as the return of a portfolio over and above a risk-free rate divided by the standard deviation of the portfolio.

$$\text{Sharpe ratio} = \frac{r_{portfolio} - r_{risk\text{-}free}}{s_{portfolio}}$$

The Sharpe ratio and other RARs demand a return differential, which also is called an excess return (Sharpe 1994). An example will illustrate why using return alone, rather than a return differential, might lead to the choice of an inferior investment.

Consider investments A and B and a risk-free security with the following characteristics (see table 1):

While A offers more return per unit of risk (8/5 v. 7/5), B is the better investment, a fact signaled by its superior Sharpe ratio. This is because an investor placing one-half of his assets in B and the other half in the risk-free security would earn a return of 9.5 percent. The portfolio would have a standard deviation of 5 percent (because the standard deviation of the risk-free investment is zero). The blend of the second investment and the risk-free security offers more return for the same level of risk as A, making the second investment preferable.

Modigliani and Modigliani Ratio

Because the Sharpe ratio is the excess return of a portfolio divided by its standard deviation, the units are difficult to interpret. The newest score on the block, the Modigliani and Modigliani (MM) ratio, proposes a twist on the Sharpe ratio to simplify interpretation (Modigliani and Modigliani 1997). The MM ratio is the Sharpe ratio multiplied by the standard deviation of the market portfolio. The rankings of a set of portfolios will be identical whether ordered by Sharpe or MM. MM provides a value that is the excess return that investors would have achieved relative to the market (i.e., the benchmark portfolio) if they had assumed a level of risk equal to that of the market. The MM risk-adjusted performance arrives at the score by leveraging or deleveraging the portfolio so that its risk is identical to that of the market (i.e., benchmark) portfolio. It assumes that investors may borrow or lend at a risk-free (standard deviation = 0) rate.

The Information Ratio

The information ratio (IR) is a more general version of the Sharpe ratio. To use the IR one chooses a benchmark, such as the S&P 500 Index. Each period (e.g., month) the difference between the return of the portfolio and that of the benchmark is computed. The IR is the average of the differences divided by the standard deviation of the differences. In other words, it is the excess return divided by the tracking error.

The Sharpe, MM, and IR measures all use standard deviation to represent risk. Standard deviation measures the variability (i.e., uncertainty) of portfolio returns around an average. A criticism of standard deviation stems from the

TABLE 1

Investment	Return	Std. Dev	Ret/Std. Dev	Sharpe
Risk-free security	5%	0%		
A	8%	5%	8/5	0.6=(8–5)/5
B	14%	10%	7/5	0.9=(14–5)/10
½ B + ½ risk-free security	9.5%	5%	9.5/5	0.9=(9.5–5)/5

fact that observations above the mean contribute as much to the statistic as observations an equal distance below the mean. Critics point out that investors generally don't mind unexpectedly high returns; they mind unexpectedly low returns.

Because standard deviation measures total variability, it is appropriate for measuring total portfolio return as long as the distribution of returns is symmetric. If returns are symmetric, above and below average observations occur with the same frequency and magnitude. If returns are asymmetric, then it makes sense to focus on downside risk. One downside risk statistic, called semi-deviation, considers only the observations that are negative. Similarly, one can choose to focus on observations that fall below a target or required rate of return.

Sortino Ratio

The Sortino ratio resembles the Sharpe ratio, except that it uses downside risk in place of standard deviation. Use of this statistic instead of Sharpe makes sense when the distribution of returns is asymmetric.

RARs that use standard deviation are appropriate for well-diversified portfolios and for entire portfolios. A different measure should be used for component styles or subasset classes (e.g., emerging markets). The fathers of modern portfolio theory, Harry Markowitz, Sharpe, and others, recognized that investors should expect to be compensated for accepting only systematic (i.e., market) risk. They distinguish between systematic risk—that of the market, which cannot be diversified away—and unsystematic risk, which can be reduced through diversification. Thus, investors who take the extra risk of holding a concentrated portfolio can expect greater variability in returns (i.e., risk) but not higher average returns.

According to modern portfolio theory, investors should expect to be compensated according to the capital asset pricing model (CAPM).

$$R_{portfolio} = R_{risk\text{-}free} + \beta\left(R_{market} - R_{risk\text{-}free}\right)$$

The expected return of a portfolio should be the risk-free rate plus the excess return of the market multiplied by beta, which is the systematic risk of a portfolio. Within the framework of this model, beta is the measure of risk.

Imagine two emerging market investments with identical returns and standard deviations. While they may have the same Sharpe ratio, an investor would prefer the one with the lower correlation to the rest of his portfolio. Beta incorporates the correlation to the rest of the portfolio.

Treynor Ratio

Like the Sharpe ratio, the Treynor ratio divides the excess return of the portfolio by its risk. The difference is that the Treynor ratio uses beta to represent risk. It also is useful for ranking investments.

$$\text{Treynor ratio} = \frac{r_{portfolio} - r_{risk\text{-}free}}{\beta_{portfolio}}$$

Jensen's Measure

Also called Jensen's alpha, this statistic measures a portfolio's return in excess of that predicted by the CAPM.

$$\text{Jensen's alpha} = R_{portfolio} - \left[R_{risk\text{-}free} + \beta\left(R_{market} - R_{risk\text{-}free}\right)\right]$$

(The term in the square brackets is the return predicted by the CAPM so that the expression is the portfolio return less the CAPM prediction.)

While CAPM was enough to win a Nobel Memorial Prize for Sharpe, the model, particularly beta, has been criticized on practical terms. Fama and French found that beta did not explain differences in the returns of stocks. In their study of stocks on the NYSE, AMEX, and NASDAQ between 1963 and 1990, high beta portfolios did not outperform low beta portfolios, on average (Fama and French 1992). Entire papers, and probably books, could be written about beta. I simply will state that beta must be used with extreme care.

Fama-French Factor Models

Many professionals use what is known as the Fama-French three-factor model to compensate for the flaws associated with the single beta used in the Jensen measure. Whereas Jensen assumes that risk comes from one source—market exposure as represented by beta—Fama-French adds two factors.[1] One is an exposure to size, which they define as the market capitalization, and the other is an exposure to value, defined as book value divided by market value. In addition, if fixed income is included in the portfolio, they add two more factors: One measures the time to maturity of the portfolio's holdings and the other assesses default risk. The general idea is to measure the excess return of a portfolio relative to what would have been earned by a benchmark portfolio taking the same systematic risks.

$$\text{Fama-French alpha} = R_{portfolio} - \left[R_{risk\text{-}free} + \beta_{market}\left(R_{market} - R_{risk\text{-}free}\right) + \beta_{size}\left(R_{large} - R_{small}\right) + \beta_{value}\left(R_{value} - R_{growth}\right)\right]$$

Shapes Matter

Users of risk-adjusted returns should consider probability distributions, which define the likelihood of each possible event. For example, the probability that any one face of a die will appear when it is rolled is 0.1667 (1/6). This is called a uniform distribution. Distributions can come in many shapes. Most people are familiar with the normal distribution, which is shaped like a bell. This is a symmetric distribution. Security returns are not normally distributed. A common assumption is that the logarithm of relative stock prices (a fancy way of saying 1 + r, where r is a return, such as 8 percent) is normally distributed. This is the assumption of the Black-Scholes options-pricing model. While this assumption is pretty good (but not exact) for stocks, it should not be made with the distributions of many other securities and strategies. For example, neither the logarithm of returns on options nor the logarithm of returns on bonds are normally distributed.

Risk-adjusted returns capture two properties of investments that investors care about: return and risk. There are other potential properties of an investment, such as the symmetry in its returns. Investors like positive skewness. With positive skewness they are more likely to enjoy large positive returns than large negative returns (e.g., lottery tickets). Buying call options has positive skewness; writing calls has negative skewness. Options often are embedded in securities. For example, the prepayment option held by homeowners means that mortgage-backed bonds have embedded short calls.

Another property one should consider is kurtosis. A distribution has excess kurtosis if the probability of extreme events is higher than it would be under a normal distribution. Compared to a normal distribution, one with excess kurtosis would have fatter tails. Investors don't like kurtosis. It turns out that certain strategies often employed by hedge funds have excess kurtosis.

Because of these issues, simple RAR measures fail to adequately depict the performance of a portfolio. In addition, past performance may not be indicative of future results. If returns are not predictive, then it follows that RARs won't be. Generally, RARs are computed using past data. The future may differ so much from the past that the predictive value of an RAR is insignificant.

An assumption behind RARs is that investors can borrow or lend easily. Imagine an investor who is comfortable taking a market level of risk. She could index and receive market levels of return and risk. Also assume a low-risk investment, which if leveraged to the risk of the market, offers higher return than the market. The low-risk investment will have a more attractive RAR. However, if the investor is prohibited from leveraging, then the RAR is misleading. The higher-risk, lower-RAR index fund may be more attractive for that investor.

Ex-post risk-adjusted returns measure *observed* risk, not *assumed* risk. Just because risk was not detected does not mean it did not exist. Investors may assume risk without knowing it. I once had a client who won $1 million in a Pepsi promotion. My client would have won $1 *billion* if he had selected the same number between 0 and 999,999 picked by a chimpanzee. Pepsi was insured against paying $1 billion by Warren Buffet's Berkshire Hathaway. While Berkshire assumed the entire risk of losing $1 billion,[2] this risk was not observed in the company's financial statements. The important point is that the risk measured may be substantially less than the risk taken.

> "One must take care to consider appropriate risk measures and time periods, realizing that simple answers, like "42," may not adequately reflect the complexities of investing."

Summary

RARs appear to be fairly simple. Divide return by risk to arrive at return per unit of risk. But such measures may be misleading. Theoretically correct RARs normalize the varying risks of different investments, accounting for borrowing, lending, and shorting, before comparing returns. One must take care to consider appropriate risk measures and time periods, realizing that simple answers, like "42," may not adequately reflect the complexities of investing. I&W

Rex P. Macey, CIMA®, CFA®, is director of equity management at Wilmington Trust in Atlanta, GA. He earned a B.A. in mathematics from Vanderbilt University and an M.B.A. from Kenan-Flagler Business School at the University of North Carolina at Chapel Hill. Contact him at rmacey@wilmingtontrust.com.

References

Adams, Douglas. 1979. *The Hitchhiker's Guide to the Galaxy.* London: Pan Books.

Fama, Eugene F., and Kenneth R. French. 1992. "The Cross-Section of Expected Stock Returns." *The Journal of Finance* 47, no. 2 (June): 427–465.

Modigliani, Franco, and Leah Modigliani. 1997. "Risk-Adjusted Performance: How to Measure It and Why." *Journal of Portfolio Management* 23, no. 2 (winter): 45–54.

Sharpe, William F. 1966. "Mutual Fund Performance." *Journal of Business* (January): 119–138.

Sharpe, William F. 1994. "The Sharpe Ratio." *Journal of Portfolio Management* (fall): 49–58.

Endnotes

1 For details of the additional factors, see Ken French's Web site, http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/f-f_bench_factor.html.

2 To be paid over 40 years.