INVESTMENT PERFORMANCE MEASUREMENT

BY DEBORAH KIDD, CFA

The Sortino Ratio:
Is Downside Risk the Only Risk that Matters?

This article, the third in a series to examine risk-adjusted performance metrics, focuses on a commonly used measure of downside risk—the Sortino ratio. Useful as a complement to other risk measures, the Sortino ratio is not a complete measure of risk in itself. Investors should be aware that variations of the ratio can significantly affect the information contained in the measure.

Defining risk by measuring only the "bad" volatility of a distribution is intuitively appealing. Behavioral finance tells us that large negative surprises do not produce the same sorts of emotions as large positive surprises. Harry Markowitz recognized this distinction back in 1959 when he proposed a measure of downside variability he called semivariance. At the time, the calculation was too complex without the aid of computers, so Markowitz focused his work on mean variance instead. Standard deviation, the square root of variance, has been the predominant measure of risk ever since. Despite its widespread popularity, standard deviation is subject to two limitations: the assumption of a symmetrical return distribution, which penalizes upside deviations and downside deviations equally, and the use of the mean as a target return.

As computational ability and the use of computers evolved over the past few decades, so too did the development of downside risk measures. Semideviation and downside deviation are two risk measures that quantify the downside portion of the return distribution. The calculation for semideviation is similar to that of standard deviation but uses returns below the mean. The difference between the mean and each observation below the mean is squared, and the sum of the squared differences is divided by the total number of returns. This operation gives the semivariance, of which semideviation is the square root. An investment with normally distributed returns will have the same level of riskiness under both standard deviation and semideviation because the upside and downside volatility will be equivalent.

Downside deviation addresses not only the first limitation of standard deviation but also the second—using the mean of a distribution as the target return. Downside deviation measures risk in a manner similar to semideviation except that it substitutes an investor-defined target return for the mean return. Called the minimum accepted return (MAR), the target rate may be an absolute return, an index return, the risk-free rate, or zero (as in zero tolerance for principal loss). A frequently used performance metric designed to measure downside deviation and the risk of failing to achieve an investor's MAR, the Sortino ratio, is the focus of this article.

THE SORTINO RATIO

The Sortino ratio is a variation of the Sharpe ratio, the most universal measure of return to risk. The Sortino ratio was created in recognition of the realizations that large positive performance deviations should not be penalized in the same way as large negative performance deviations and that failing to earn the mean return is not how most investors define risk. Named after Dr. Frank Sortino of the Pension Research Institute, the ratio measures excess return to the risk of not meeting an investor's MAR. The formula for the Sortino ratio is simple; it is calculated in the same way as the Sharpe ratio except that the MAR replaces the risk-free rate in the numerator and downside

deviation replaces standard deviation in the denominator. The calculation for the Sortino ratio is as follows:

$$S = (\text{Mean portfolio return} - \text{MAR})/\text{Downside deviation}.$$

Because the Sharpe ratio defines risk as standard deviation, it falls prey to the same shortcomings as standard deviation. The Sortino ratio appears to resolve several of the issues inherent in the Sharpe ratio: It incorporates a relevant return target, in both the numerator and the denominator; it quantifies downside volatility without penalizing upside volatility; and because of its focus on downside risk, it is more applicable to distributions that are negatively skewed than measures based on standard deviation. Yet downside deviation has its own set of limitations. Furthermore, several variations of the downside deviation calculation exist, and which one an investor uses matters a great deal. Amelia Hopkins, Senior Vice President at Granville Capital, uses the Sortino ratio when evaluating hedge funds. She explains, "While the statistic, downside deviation, is easy to understand it is also easy to miscalculate."

DOWNSIDE DEVIATION—SOME CAVEATS

Although volatility statistics measure past returns, investors calculate them to help forecast future returns. What an investor really wants to know is, what is the asset's *expected* volatility? The variations in the calculation for downside deviation can have a considerable impact on the answer to this question. Despite its similarity to standard deviation, Sortino and Forsey (1996) caution that the proper calculation of downside deviation is more complex and that the widespread method of simply using the historical returns that fall below the MAR can significantly underestimate downside risk. "Calculation error," they say, "is due primarily to measuring only what did happen (discrete), instead of what could have happened (continuous)" (p. 37).

There are several problems with the discrete method. Investors know that historical returns do not predict future returns. Additionally, performance measurement results are highly dependent on the time period under consideration; excluding upside deviations further limits the data sample. Moreover, if the majority of the returns are positive, downside deviation can be significantly understated. Sortino and Forsey (1996) illustrate the shortcomings of the discrete method with an example based on the performance of the Japanese equity market during the 1980s. They compare this method with two methods based on a continuous distribution.

From 1980 through 1989, the Japanese equity market had 10 years of only positive annual returns, although it had 36 months of negative returns. In 1990, the year following the sample period, the market plunged 39 percent. Such an extreme loss would not have been predicted by the small sample of negative monthly returns, the largest of which was −12 percent. In fact, the monthly downside deviation of the sample period was 2.74 percent calculated under the discrete method, using an assumed monthly MAR of 0.08 percent and 45 return observations. The monthly downside deviation rises to 3.20 percent when calculated according to a method that involves fitting a continuous probability curve over the distribution of returns and using integral calculus for the computation.

The downside risk becomes even more evident when a procedure called bootstrapping is applied to the monthly returns. The bootstrap method produced a continuous return distribution, indicating a remote possibility for the Japanese equity market to plunge 42.7 percent in one year, despite the market's previous 10 years of annual gains. By randomly selecting and combining historical monthly returns, the bootstrap procedure can generate thousands of simulated annual returns. A drawback of this method is that it assumes the observed returns are the only returns possible. If there are no negative historical returns, bootstrapping does not create them.

Sortino and Forsey also noted other mistakes involving discrete data, which Hopkins sees often. She describes a fairly common mistake as dividing the sum of squared deviations (from MAR) by the number of observations below the MAR—rather than the total number of observations. "Another problem is calculating downside deviation on different time intervals (monthly versus annually)," she says. She provides an example in which an investor's goal is to earn 5 percent annually and the downside deviation is being calculated for equities:

> Using the monthly returns of the S&P 500 Index for 10 years (2001-2010) and a monthly MAR of 0.4167% (5 percent annually) gives a downside deviation of 3.71%. Annualizing the number by multiplying by the square root of 12 (which is another problem) gives 12.84%. If, however, 10 discrete periods of annual returns are used then there are only four observations below 5%, which produces a dramatically different number. I have read that you cannot annualize downside deviation in the same manner as standard deviation. . . .but analytical packages I've seen do it anyway.

A final caveat: Because it is incorrect to use only historical returns in the downside deviation calculation, it is also incorrect to annualize the number based on these returns. Annualizing discrete data will overstate risk. For the most precise measurement, Sortino and Forsey (1996) recommend fitting a continuous curve to a bootstrapped distribution and using integral calculus to make the calculation. The continuous methods for calculating downside deviation incorporate a forward-looking element into the measure, as opposed to providing an estimate of risk based on a limited set of historical returns.

USING THE SORTINO RATIO

Like the Sharpe ratio, the Sortino ratio is intended to be used in a relative context to compare a portfolio or fund with another fund, a benchmark index, or a manager universe. And like the Sharpe ratio, a higher ratio indicates better risk-adjusted performance. To compare funds, the Sortino ratio of each fund must use the same MAR. In addition, because there are variations in the downside deviation calculation, investors should understand how a Sortino ratio is calculated when using it for evaluation purposes. Comparing the Sharpe ratio and the Sortino ratio for a fund can give an indication of what portion of a fund's volatility is related to outperformance versus underperformance.

Because the Sortino ratio uses an investor-defined target return for the benchmark, the ratio is not as widely reported as other *ex post* risk-adjusted performance measures. A variation of the Sortino ratio used to facilitate the comparison of various funds incorporates the risk-free rate in the numerator and semideviation in the denominator. If the MAR is equivalent to the risk-free rate and the investment's returns are normally distributed, the Sortino ratio will assign the same ranking to a portfolio as the Sharpe ratio. Sortino and Price (1994) noted that using the risk-free rate detracts from the ratio's usefulness as a goal-oriented performance measure; they suggest using the average annual return for a market index instead to allow for broad comparisons.

Although downside deviation is better able to capture the risk–reward trade-off of a non-normal distribution than standard deviation, caution must be taken when applying the Sortino ratio to strategies with known asymmetric return distributions, such as hedge funds. Lo (2002) discovered grossly exaggerated Sharpe ratios among hedge funds resulting from serial correlation in the monthly returns. He noted that for hedge funds investing in illiquid and restricted securities, serial correlation—the correlation of a variable with itself over successive time intervals—can arise from the use of stale prices because of a lack of pricing information from publically traded markets. By smoothing returns and reducing variability, serial correlation can be a problem for any measure based on return variability, including downside deviation. Rogers and Van Dyke (2006) caution that investors should also be wary of strategies that exhibit exceptionally positive asymmetric return characteristics; such performance may indicate pricing issues or may be the result of a single performance event, signifying either a one-off opportunity or luck rather than skill.

Is downside risk the only risk that matters? Investors generally do care more about losses than gains. Yet downside risk is difficult to quantify, and upside deviations can provide valuable information about future performance. The Sortino ratio is a useful tool in assessing the riskiness of an investment, but it is not a complete measure of risk. Outperformance should not be ignored because these gains were generated by risk taking; the same strategy might produce corresponding losses at some point in the future. Said differently by Sortino, "Just because nothing bad happened doesn't mean you didn't take any risk."[1]◆

NOTES

1. Quote from www.sortino.com.

BIBLIOGRAPHY

Chance, Don M., Grant, Kenneth, and Marsland, John. "Risk Management." 2007. In *Managing Investment Portfolios: A Dynamic Process.* 3rd ed. Edited by John L. Maginn, Donald L. Tuttle, Jerald E. Pinto, and Dennis W. McLeavey (Hoboken, NJ: John Wiley & Sons):633.

Feibel, Bruce J. 2003. *Investment Performance Measurement.* Hoboken, NJ: John Wiley & Sons, 155–164, 199–200.

Lo, Andrew. 2002. "The Statistics of Sharpe Ratios." *Financial Analysts Journal*, vol. 58, no. July/August:36–50.

Nawrocki, David. 1999. "A Brief History of Downside Risk Measures." *Journal of Investing*, vol. 8, no. 3 (Fall):9–25.

Rogers, Douglas S., and Van Dyke, Christopher J. 2006. "Measuring the Volatility of Hedge Fund Returns." *Journal of Wealth Management*, vol. 9, no. 1 (Summer):45–53.

Sortino, Frank A., and Forsey, Hal J. 1996. "On the Use and Misuse of Downside Risk." *Journal of Portfolio Management*, vol. 22, no. 2 (Winter):35–42.

Sortino, Frank A., and Price, Lee N. 1994. "Performance Measurement in a Downside Risk Framework." *Journal of Investing*, vol. 3, no. 3 (Fall):59–64.

Deborah Kidd, CFA, is senior vice president at Boyd Watterson Asset Management, LLC.