

MANAGED MONEY

Sortino ratio: A better measure of risk

BY TOM ROLLINGER & SCOTT HOFFMAN

The purpose of risk management and subsequent risk measures is to avoid large downside risk, so a performance metric that punishes excessive positive returns while missing excessive downside risk is flawed. The Sortino ratio offers a better measure of risk.

Many traders and investment managers want to measure and compare commodity trading advisors (CTAs) or trading systems. While there are many ways to measure an investment's performance, risk-adjusted returns are one of the most important measures to consider because, given the inherent free leverage of the futures markets, more return can be earned by taking more risk. The most popular measure of risk-adjusted performance is the Sharpe ratio. While the Sharpe ratio definitely is the most widely

Positively skewed vs. negatively skewed strategies

Desired target return

used, it is not without its issues and limitations. Because of the way the Sharpe ratio is calculated, it tends to punish upside volatility in a trading program. We believe the Sortino ratio improves on the Sharpe ratio in a few areas. The purpose of this article is to review the Sortino ratio's definition and present how to calculate it properly, because we have seen its calculation often done incorrectly.

Sharpe ratio

The Sharpe ratio is a metric that aims to measure the desirability of an investment by dividing the average period return in excess of the risk-free rate by the standard deviation of the return generating process. Devised in 1966 by Stanford Finance Professor William F. Sharpe as a measure of performance for mutual funds, it undoubtedly has some value as a measure of investment "quality," but it has a few limitations also.

The most glaring flaw is that it does not distinguish between upside and downside volatility (see "Good news, bad news," right). In fact, high outlier returns can have the effect of increasing the value of the denominator (standard deviation) more than the value of the numerator, thereby lowering the value of the ratio. For some positively skewed return distributions such as that of a typical trend-following CTA strategy, the Sharpe ratio can be *increased* by removing the largest positive returns. This is nonsensical because investors generally welcome large positive returns.

Additionally, to the extent that the distribution of returns is non-normal, the Sharpe ratio falls short. It is a par-



For more on risk measures, go to futuresmag.com/Risk

ticularly poor performance metric when comparing positively skewed strategies like trend-following to negatively skewed strategies like option selling (see "Bigger winners vs. more winners," page 42). In fact, for positively skewed return distributions, performance actually is achieved with *less risk* than the Sharpe ratio suggests. Conversely, standard deviation understates risk for negatively skewed return distributions, i.e., the strategy actually is *more risky* than the Sharpe ratio suggests. Typical long-term, trend-following CTAs, especially those with longer track records, generally have Sharpe ratios in the 0.50 – 0.90 range. However, negatively skewed programs (convergent strategies) like option writing will produce high Sharpe ratios, 3.0 and above, up until a devastating drawdown. The Sharpe ratio often misses the inherent risk of convergent strategies.

Sortino ratio

In many ways, the Sortino ratio is a better choice, especially when measuring and comparing the performance of managers

whose programs exhibit positive skew in their return distributions. The Sortino ratio is a modification of the Sharpe ratio, using downside deviation rather than standard deviation as the measure of risk — i.e., only those returns falling below a user-specified target (“Desired Target Return”) or required rate of return are considered risky (see “Good news, bad news”).

It is interesting to note that even Nobel laureate Harry Markowitz, when he developed Modern Portfolio Theory (MPT) in 1959, recognized that because only downside deviation is relevant to investors, using it to measure risk would be more appropriate than using standard deviation. However, he used variance (the square of standard deviation) in his MPT work because optimizations using downside deviation were computationally impractical at the time.

In the early 1980s, Dr. Frank Sortino had undertaken research to come up with an improved measure for risk-adjusted returns. According to Sortino, it was Brian Rom’s idea at Investment Technologies to call the new measure the Sortino ratio. The first reference to the ratio was in *Financial Executive Magazine* (August 1980) and the first calculation was published in a series of articles in the *Journal of Risk Management* (September 1981).

The Sortino ratio, S, is defined as:

$$S = \frac{R - T}{TDD}$$

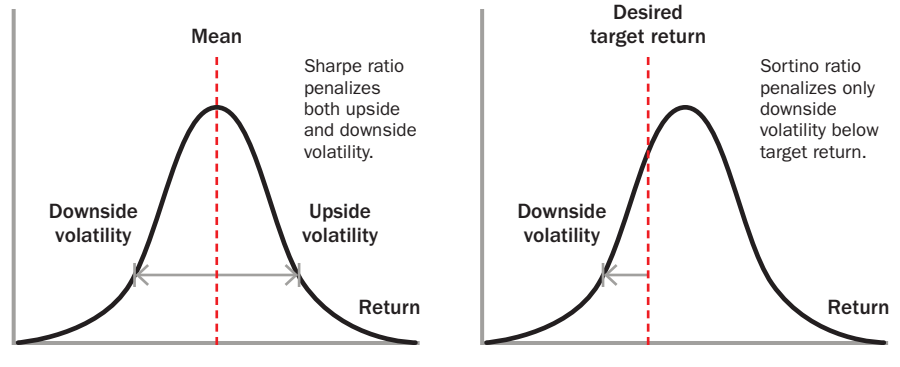
where

- R is the average period return;
- T is the target or required rate of return for the investment strategy under consideration (originally T was known as the minimum acceptable return, or MAR. In his more recent work, MAR is now referred to as the Desired Target Return).
- TDD is the target downside deviation.

The target downside deviation is defined as the root-mean-square, or RMS, of the deviations of the realized return’s underperformance from the target return where all returns above the target return are treated as underperformance of 0. Mathematically:

GOOD NEWS, BAD NEWS

Large return swings are a sign of volatility and risk, but if a program consistently produces strong up moves with lower down moves, it should not be punished for those strong moves in its favor.



Target Downside Deviation =

$$\sqrt{\frac{1}{N} \sum_{i=1}^N (\text{Min}(0, X_i - T))^2}$$

where

$X_i = i^{\text{th}}$ return

N = total number of returns

T = target return

The equation for TDD is very similar to the definition of standard deviation:

Standard Deviation =

$$\sqrt{\frac{1}{N} \sum_{i=1}^N (X_i - u)^2}$$

where

$X_i = i^{\text{th}}$ return

N = total number of returns

u = average of all X_i returns.

The differences are:

- 1) In the target downside deviation calculation, the deviations of X_i from the user selectable target return are measured, whereas in the Standard Deviation calculation, the deviations of X_i from the average of all X_i is measured.
- 2) In the target downside deviation calculation, all X_i above the target return are set to zero, but these zeros still are included in the summation. The cal-

culatation for Standard Deviation has no Min() function.

Standard deviation is a measure of dispersion of data around its mean, both above and below. Target downside deviation is a measure of dispersion of data below some user-selectable target return with all above target returns treated as underperformance of zero. Big difference.

Sortino ratio calculation

In this example, we will calculate the annual Sortino ratio for a hypothetical CTA with positive skew with the following set of annual returns:

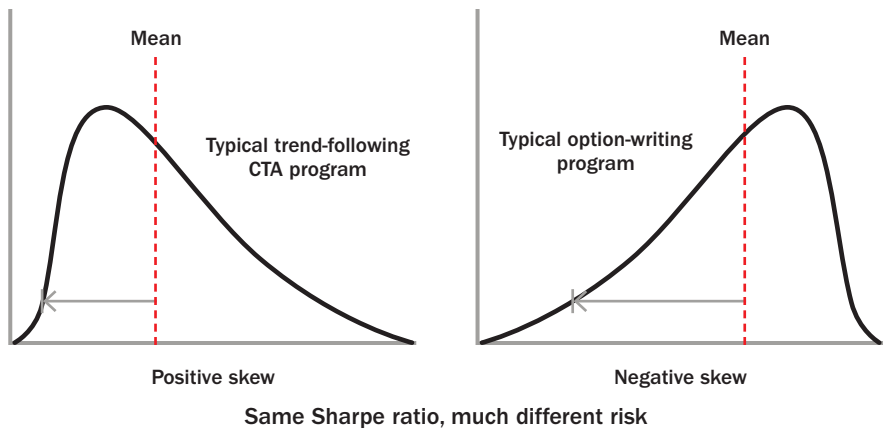
Annual Returns: 2%, 1%, -1%,
18%, 8%, -2%, 1%, -1%
Target Return: 0%

Although in this example we use a target return of 0%, any value may be selected, depending on the application, i.e., a futures trading system developer comparing different trading systems vs. a pension fund manager with a mandate to achieve 8% annual returns. Of course using a different target return will result in a different value for the target downside deviation. If you are using the Sortino ratio to compare managers or trading systems, you should be consistent in using the same target return value.

First, we will calculate the numerator of the Sortino ratio, the average period return minus the target return:

BIGGER WINNERS VS. MORE WINNERS

The Sharpe ratio assumes a normal distribution and tends to give a false sense of security to negatively skewed strategies that tend to produce consistent small positive returns but can be punctuated by rare, painful large negative returns.



Average annual return – Target return
 = 3.25% – 0% = 3.25%

Next, we will calculate the target downside deviation:

- 1) For each data point, calculate the difference between that data point and the target level. For data points above the target level, set the difference to 0. The result of this step is the underperformance data set.

- min(0, 2% – 0%) = 0%
- min(0, 1% – 0%) = 0%
- min(0, -1% – 0%) = -1%
- min(0, 18% – 0%) = 0%
- min(0, 8% – 0%) = 0%
- min(0, -2% – 0%) = -2%
- min(0, 1% – 0%) = 0%
- min(0, -1% – 0%) = -1%

- 2) Next, calculate the square of each value in the underperformance data set determined in Step 1.

- 0% ^ 2 = 0%
- 0% ^ 2 = 0%
- 1% ^ 2 = 0.01%
- 0% ^ 2 = 0%
- 0% ^ 2 = 0%
- 2% ^ 2 = 0.04%
- 0% ^ 2 = 0%
- 1% ^ 2 = 0.01%

- 3) Then, calculate the average of all squared differences determined in Step 2. Notice that we do not “throw away” the 0% values.

Average = (0% + 0% + 0.01% + 0% + 0% + 0.04% + 0% + 0.01%) / 8 = 0.0075%

- 4) Then, take the square root of the average determined in Step 3. This is the target downside deviation used in the denominator of the Sortino ratio.

Target Downside Deviation =
 Square root of 0.0075% = 0.866%

Finally, we calculate the Sortino ratio:

Sortino Ratio =
 3.25% / 0.866% = 3.75

This is a strong score and indicative of the return stream from which we calculated it. Calculating the Sharpe ratio on the same set of returns would have produced a Sharpe ratio (0% RFR) of 0.52, a mediocre one that indicates more volatility by penalizing the outsized positive returns.

Sortino vs. Sortino

Often in trading literature and trading software packages we have seen

the Sortino ratio, and in particular the target downside deviation, calculated incorrectly. Most often, we see the target downside deviation calculated by “throwing away all the positive returns and taking the standard deviation of negative returns.” We hope that by reading this article, you can see how this is incorrect. Specifically:

In Step 1, the difference with respect to the target level is calculated, unlike the standard deviation calculation where the difference is calculated with respect to the mean of all data points. If every data point equals the mean, then the standard deviation is zero, no matter what the mean is. Consider the following return stream: [-10, -10, -10, -10]. The standard deviation is 0; while the target downside deviation is 10 (assuming target return is 0).

In Step 3, all above target returns are included in the averaging calculation. The above target returns set to 0% in Step 1 are not thrown away.

The Sortino ratio takes into account both the frequency of below-target returns as well as the magnitude of them. Throwing away the zero underperformance data points removes the ratio’s sensitivity to frequency of underperformance. Consider the following underperformance return streams: [0, 0, 0, -10] and [-10, -10, -10, -10]. Throwing away the zero underperformance data points results in the same target downside deviation for both return streams, but clearly the first return stream has much less downside risk than the second.

In this article we presented the definition of the Sortino ratio and the correct way to calculate it. While the Sortino ratio addresses and corrects some of the weaknesses of the Sharpe ratio, we feel there is one measure that is even better yet: The Omega Ratio. We look forward to tackling the Omega Ratio in our next article. **F**

Tom Rollinger is director of new strategies development for Sunrise Capital Partners. Previously he was a portfolio manager for quantitative hedge fund legend Edward O. Thorp. Scott Hoffman is the founder of CTA Red Rock Capital Management.