

All Hail the Sharpe Ratio

The Sharpe Ratio is the metric of choice amongst hedge funds. This article examines the uses and abuses of the Sharpe Ratio in hedge fund analysis.

In a 1966 paper, Dr. William F. Sharpe proposed the “reward-to-variability” ratio as a tool for comparing mutual funds. The idea was to measure the rate of return achieved by a fund relative to the amount of risk taken en route. The ratio measures return generated per unit of risk taken. Conceptually simple, rich in meaning and, originally, unambiguously defined, it can provide investors with an objective, quantitative measure of performance. In all but name, the ratio has been widely adopted by investors as a universal measure of the quality of a fund.

Sharpe Ratio Defined

An investor can always achieve some rate of return by taking absolutely no risk. Obviously a fund can too. Therefore, the taking of risk should result in returns which exceed the risk free rate. The more risk taken, the greater the excess should be. Thus a fund’s excess return (over the risk free rate) is the data that should interest an investor.¹

As defined by its creator, the Sharpe Ratio is simply the average of excess returns divided by the volatility of excess returns. Thus the Sharpe ratio increases as excess return increases or as volatility decreases. In theory, the higher the Sharpe ratio, the better the manager.

To the extent that the past volatility captures the risk taken by a fund, which will be

explored further, the measure answers the question, “Which fund provided the most return relative to risk taken.” And this is indeed the fundamental question any investor should ask.

The Problem

While the fundamental question is the right one, a Sharpe Ratio does not answer it as well as often presumed. The problem is that volatility is not risk. Volatility is a mere proxy for risk, and a poor one at that.

When the Sharpe Ratio is used to compare mutual funds who are buying stocks from the same universe, the extent to which volatility misstates the risks that were taken is mitigated by the consistency of the misstatement across the various funds being examined. If risk was higher than implied by volatility for one fund, then it was probably also higher for other funds too. The absolute correctness of the measure may be compromised, but relatively, the ratio still does well to rank the funds.

But even in the world of simple, long only mutual funds, there could be a problem. Consider a fund which takes a position in a company which is to be acquired. Assuming that the transaction ultimately succeeds, the fund will be rewarded for bearing the risk of its failure. But the risk of that failure will not be fully reflected in the volatility of the stock. The risk was latent: There was always a chance that the stock could drop suddenly and severely if the acquisition failed. This was a real risk, but it does not show up in volatility unless it actually happens. As a result, the fund’s enhanced

¹ The ratio of raw returns to raw volatility, often called the “information ratio”, is utterly meaningless. See Sharpe [1996].

return shows up in the numerator of the Sharpe ratio, but the increased risk does not appear in the denominator. The fund appears to have taken less risk than it actually did. The fund's Sharpe ratio has been given a boost by taking a risk that does not necessarily show up in its volatility. Compared to its peers via Sharpe Ratios, the fund looks better than it actually is.

The above scenario is an example of a very important concept: There are risks that can be taken which may never show up as volatility. But they are risks nonetheless and the investor must weigh them against returns lest she deceive herself. In the domain of mutual funds, and most certainly in the arenas where hedge funds play, there are all kinds of risks that are not necessarily reflected in return volatility. This is why impressive Sharpe Ratios are often available.

Event Risk

Event Risk is the threat of a rare but cataclysmic event occurring. Sovereign defaults, (Russia, 1998), Terrorist attacks (2001), corporate defaults (Enron), and stock market crashes (1987) are famous examples. There are endless examples of event risk manifestations that are contained, but no less catastrophic to a small segment of investors. On any given day there are single stocks which jump way up or way down to the glee or despair of certain investors.

An insurer is the canonical example of someone who takes event risk. Insurers collect premiums to assume the risk of having to make a large payout if some rare event occurs. Month to month, the insurer's net cash flow is positive, but occasionally, when a rare event occurs, they make large payouts, possibly bigger than what they have collected cumulatively over many months. Insurers are paid to bear event risk.

Investors too are compensated for bearing event risk in much the same way. They receive small, regular payments or large payments upfront to take the financial risk

of an improbable but expensive scenario occurring. Theoretically, in the long run, these payments should equal the large loss the investor could take should an "event" occur.² But in the short term, the downside may not manifest itself. As such, the investor can generate a seemingly nice income stream indefinitely, though not perpetually.

An investor holding a risky corporate bond is paid a premium for bearing the risk that the issuer may default. The seller of a put option is paid upfront to bear the risk of the underlying security dropping in price. The seller of a call option on a stock takes the risk that the stock jumps up. These types of positions are "short" event risk. They generate an income which, presumably, is fair compensation for bearing the catastrophe risk. However, sooner or later, the catastrophe might actually occur. Hence the investor is "short".

Any investment that has some component that is short event risk will necessarily get a Sharpe Ratio boost—at least until the event occurs. A fund that is chronically short event risk can thus produce an incredibly high Sharpe Ratio. Unfortunately the investor, and even the manager himself, often don't realize that the Sharpe Ratio is misstating the risk/return trade off.

It should be said that there is nothing inherently evil about being short event risk. Indeed, investors get paid well to do it. Good managers are the ones who find the investments that overpay relative to the level of event risk taken. It can indeed be a profitable strategy. The sin is to reap the rewards of this strategy while understating the true risk. And one of the most common ways to do this is by using Sharpe Ratios to describe the risk/return profile of a short event risk strategy.

Given enough time, a seller of event risk will eventually face the "event" and take


² In a (theoretical) expected value sense, the income received by the investor should sum to the expected loss.

losses. If he is lucky or good, it might take time. But when it does happen, his returns will finally contain all the “information” about the risks he is taking. Furthermore, his Sharpe Ratio will then truly reflect his risk/return trade off.


Any hedge fund which uses derivatives with nonlinear profiles, (like options) or engages in any sort of credit plays is necessarily buying or selling event risk. As such, the volatility of their returns is not representative of the risks they are taking. This is why comparing hedge funds on the basis of Sharpe Ratios is often pointless. Furthermore, hedge funds which are chronically short event risk—and there are many—are being deceptive when they let investors think of their returns volatility as an indication of the risks they take.

An Example

Consider a hedge fund which holds a collection of lower quality corporate bonds. The true nature of the returns distribution of this fund would look something like this:



Most of the time the fund achieves a small positive return. But occasionally, a bond defaults resulting in a big loss. It is very possible, (in this example, 86% probable), that a default event does not occur in the first two years of employing the strategy. Hence the fund’s observed return distribution will look like this:



Which means the two year old (hypothetical) hedge fund will have the following vital stats:

Average Annual Excess Return	10.1%
Volatility	4.2%
Sharpe Ratio	2.4

This product, for sale at 2/20, has seemingly great numbers. This is because the strategy, when summarised by an abridged data set, looks better than it really is. The true stats are destined to resemble that which comes from the fund’s true distribution:

Average Annual Excess Return	4.3%
Volatility	11.0%
Sharpe Ratio	0.4

Summary

When evaluating a hedge fund’s past performance, an investor must weigh the track record against the risk that was taken to achieve it. In theory, the Sharpe Ratio measures exactly this; return per unit risk. But often there were other risks that never showed up in volatility. The investor must understand these latent risks to really know the return per unit of risk that the fund has achieved. And to this cause, the Sharpe Ratio is at best meaningless and at worst deceiving. This is why there is no substitute for seeking a comprehensive understanding of what the hedge fund manager is doing.

Epilogue: The “Sortino Ratio” and Downside Deviation

In recent years, investors have started to talk about good volatility (stemming from

returns above some threshold) and bad or “downside” volatility (stemming from returns below some threshold). Furthermore, some funds are calling the ratio of the mean return to “bad volatility” their “Sortino Ratio”. (This is not, by the way, how Dr. Frank A. Sortino defines the ratio, hence the quotes.)

The argument in favour of “Sortino Ratios” over Sharpe Ratios is that high volatility can be created because the fund’s returns jumped up—something that nobody minds! Thus the Sharpe ratio is punished by “good” volatility. The “Sortino Ratio” uses only downside volatility as its denominator. Thus, if a fund’s volatility is higher because of positive jumps, the “Sortino Ratio” will be better than the Sharpe ratio.

The problem with the ratio stems from an age old human flaw: We understate the possibility of past risks that did not harm us. Figure 1, from Dr. Sortino’s web site³, says it all.

The little information that can be found in historical volatility is diminished further by removing data so as to “improve” the

number. By differentiating between good volatility and bad volatility, we assume that something other than what actually happened could not possibly have occurred. And the reality that it didn’t occur, re-enforces this delusion. In most cases, any fund that returned a stunning +12% in one month was taking some sort of risk that could have resulted in at least some loss. But when calculating “downside volatility” we ignore this +12% because it “unfairly” inflates volatility. But ignoring the +12% is tantamount to ignoring the risks that were taken to generate it. A responsible investor does not forgive excessive risk taking just because it generated profit. But that is exactly what the “Sortino Ratio” does.



Figure 1: Just because you got away with it doesn’t mean you didn’t take any risk.