

thinknotes

april 2004

A Closer Look at Risk-adjusted Performance Measures

- When analysing risk, we look at the factors that may cause retirement funds to fail in meeting their objectives. One of the more frequently used measures of risk is the volatility of returns.
- In this note we discuss volatility in detail, including what the volatility measure means, how it is calculated using standard deviation, how volatilities can be annualised as well common pitfalls one should be aware of in general when calculating them.
- We also elaborate on how the volatility measure is used to define various risk-adjusted performance ratios like the Sharpe, Information and Sortino ratios. These ratios are popular in investment practice as they aim to capture, in a single number, the consistency of an asset or a portfolio's performance in terms of its volatility.
- Although we do not necessarily agree with their more popular usage, asset managers frequently use these ratios to communicate the consistency of their performance to clients. We do, however, include accurate definitions and formulae in this document.

Research Analysts

Petri Greeff
Andrew van Biljon
Prasheen Singh
Laurett Jardim
+27 21 683 7111

Feedback?

research@riscura.com



good thinking

A Closer Look at Risk-adjusted Performance Ratios

Introduction

When analysing risk, we look at the factors that may cause investment funds to fail in meeting their objectives. Listed below are some of the risks that we consider when selecting and monitoring investment managers:

- Volatility of return;
- Lack of liquidity in assets;
- Credit risk;
- Sensitivity to the market i.e. beta and duration;
- Investment style;
- Sector and stock specific risk.


For the purposes of this document, we will only discuss one of the above risk measures i.e. the volatility of return and how it is used as a component in the calculation of some risk-adjusted performance ratios. These ratios are popular in investment practice as they attempt to capture the consistency of a portfolio's performance in terms of its risk (or volatility) in one single number.

Volatility

The volatility of an asset is measured by calculating the standard deviation of the asset's returns. This measure indicates how dispersed a set of returns is around their average or mean. To better understand what standard deviation represents, one can use the following example: given a large set of returns with a normal (symmetrical 'bell-shaped') distribution, one standard deviation is a numerical value that indicates the distance needed on either side of the average in order to include 68% of all observations. Similarly, two standard deviations is the distance needed in order to include 95% of all observations.

Figure 1 shows how the above example can be practically applied to a fictitious asset's set of monthly returns over a one-year period. The solid horizontal line at 3% percent represents an average of the monthly returns. Given a standard deviation of 5%, the two dashed horizontal lines at -2% and 8% are positioned one standard deviation away from the average. Therefore, assuming that the returns of this asset are distributed according to a normal distribution, 68% of all possible returns should lie between the margins of -2% and 8%.

Figure 1. A set of 12 monthly returns showing the average as well as the position of one standard deviation above and below the average. If the returns follow a normal distribution, 68% of all returns should be between -2% and 8%.



In the investment universe, the standard deviation calculation can be used to estimate various types of risk. For example, calculating the standard deviation of a set of returns yields what we refer to as the 'absolute volatility' of an asset or portfolio, as shown in Figure 1. This gives a measure of capital or return risk of the asset or portfolio. Furthermore, if we subtract a benchmark's returns from portfolio's returns and calculate the standard deviation thereof, we obtain a 'tracking error' which provides a measure of risk of the portfolio's performance relative to the benchmark performance. The tracking error is also sometimes referred to as the 'active risk'.

There are two types of standard deviation calculations: i.e. sample and population. Volatility of an asset is measured by using the sample deviation, as we cannot obtain the exact statistical distribution of the asset's returns. We therefore have to settle for only a relatively small sample of these returns to approximate the standard deviation of the entire population. Mathematically the sample standard deviation, indicated by the Greek letter σ , is defined as:

$$\sigma = \sqrt{\frac{\sum_{i=1}^n (x_i - \mu_x)^2}{n-1}}$$

where x_i is the return for observation i , μ_x is the average of all the observations and n is the total number of observations.

It is important to note that in order to combine volatilities, the individual volatilities should be squared, and then summed. The combined volatility is then obtained by taking the square root thereof¹.

Calculating the standard deviation of returns is of course not the only method to measure an investment's risk. When using standard deviation you are implicitly making the assumption that all the possible asset returns are normally distributed. Experience shows that this is perhaps a very ignorant view on the behaviour of asset returns, ignoring the possibility of performance 'skew' and 'fat-tails' which result in larger negative (or positive) returns than would be expected.

When dealing with asset returns, the observations can be made over any chosen period. The most frequently used period is monthly, but weekly or even daily periods can also be used. From a statistical viewpoint, the accuracy of the standard deviation calculation depends on the number of observations included in the calculation – the more, the better. It is difficult to state the exact number since it depends on the degree that the returns vary from each other, but a good rule of thumb is to use at least 36 returns.

It is widely accepted that any return or risk number is usually quoted or reported in an annualised format. In the case of standard deviation of monthly returns, it is converted into an annual number as follows:

$$\sigma_A = \sqrt{12}\sigma$$

where σ_A represents the annualised standard deviation². Standard deviations of returns over other periods can be converted into an annual number in a similar manner. For example, if quarterly returns are used the 12 will be replaced by 4 in the formula above.

Although annualising leads to uniformity, it also increases the danger of misinterpreting the numbers. For example, take two sets of returns from the same asset over a period – one calculated monthly, the other weekly. Annualising the volatilities of these returns might show that the annualised weekly value is higher than that of the annualised monthly one, even though both describe the same underlying asset. This is because weekly returns are in general exposed to more short-term volatility compared to monthly returns. One should therefore be aware of the period of the annualised numbers before comparing annualised tracking errors or absolute volatility numbers between different assets.

¹ This is assuming that there is no correlation between the assets.

² Note that if the returns were calculated using holding period returns, as is frequently the case, this is only an approximate method to annualise monthly volatilities.

Introducing Risk-adjusted Performance Ratios

The purpose of risk-adjusted performance ratios is to measure whether investors are sufficiently compensated for the volatility inherent in an investment i.e. calculating the return of the investment per unit of risk. There are several varieties of risk-adjusted performance ratios, the most popular being:

- Sharpe ratio;
- Sortino ratio;
- Information ratio

These risk-reward ratios are also generically referred to as efficiency ratios, each having the same basic definition of an average return over a chosen period divided by the risk over that period, as measured by the standard deviation of the returns. The different ratios are distinguished on the basis of exactly how the returns are defined. The exact definitions of these returns will be discussed below.

It is important to note that none of these ratios should ever be used in isolation. They should rather be used to compare different portfolios; or track a particular portfolio over time. When comparing portfolios, none of these ratios should be compared to similar ratios from other sources without the proper knowledge of exactly how these values were calculated, i.e. over what period the returns were calculated, how many returns were used in the volatility calculation, etc. Also when reporting any of these ratios, it is good practice to communicate the information, as mentioned above, regarding the calculation itself. Finally, all these measures provide an indication of historical performance if past data is used. Readers should be aware of the weakness of past performance (and any function thereof) as a guide to future performance.

Sharpe Ratio

This ratio measures the consistency of the performance of a portfolio in excess of the risk-free rate in risk-adjusted terms. Conceptualised nearly 40 years ago, the original idea was that the Sharpe Ratio would indicate whether or not the portfolio returns differed significantly from the given risk-free rate by using a definition very similar to the well-known Student's t-test statistic. Subsequently, the ratio in its original form has seen many transformations as assumptions and approximations were made in various parts of its calculation – so much so, that various popular definitions currently exist for the Sharpe Ratio.

Since the return of the portfolio above the risk-free rate is also referred to as the 'excess return', the Sharpe Ratio originally was defined as the average of the excess return divided by the volatility of the excess returns:

$$SR = \frac{Avg(R - r_f)}{\sigma}$$

where R is a vector containing the monthly portfolio returns over each period, r_f is a vector containing the risk-free rate of the corresponding periods and σ is the volatility of the portfolio returns minus the corresponding risk-free rate.

Although holding-period returns are nowadays frequently used to calculate the Sharpe Ratio, given the definition above, the original method actually used logarithmic returns³. The main advantage of using the latter type of returns was that the task of annualising the ratio was straightforward, due to the additive property of logarithmic returns when calculating total compounded returns. For example, a set of 36 monthly logarithmic returns can be annualised by either adding the monthly returns and dividing the answer by 3 or, equivalently, calculating the average of the 36 returns and multiplying the result by 12. A similar argument allows volatility, as calculated per standard deviation, to be annualised by multiplying the monthly standard deviation by the square root of 12. Dividing these two factors that appear in the numerator and denominator respectively, a Sharpe Ratio calculated using monthly returns can be annualised by multiplying by the square root of 12. Readers will notice that even when annualising, the Sharpe Ratio keeps its intended t-test statistic format of an average return divided by the standard deviation of returns.

As mentioned, various mutated definitions of the original Sharpe Ratio are regularly used in the investment industry today. For example, in the case of using holding-period returns as opposed to logarithmic returns with the original Sharpe Ratio formula, and the annualisation method described above, the approximation coincidentally yields more or less the same numerical result. This is because both types of return calculations are comparable with each other as long as the absolute return value remains relatively small⁴. In a more crude approximation however, the averaging of the returns in the numerator is ignored altogether when using holding-period returns. In this definition, the Sharpe Ratio is then popularly defined as the result of the annualised total return minus the annual risk-free rate divided by the annualised volatility of the return. This incorrect simplification in the numerator probably developed out of confusion between logarithmic returns, their additive compounding and implicit averaging property when annualising.

Information Ratio

It is difficult to trace the exact origin of this ratio, but it could be assumed that it developed out of the much older Sharpe Ratio definition. In the light of this, the Information Ratio should be similarly defined since it measures the performance of a portfolio relative to a benchmark instead of the risk-free rate in risk-adjusted terms:

³ As opposed to holding-period returns which are defined as:

$$R = P_{i+1} / P_i - 1,$$

logarithmic returns are defined as:

$$R = \ln(P_{i+1} / P_i)$$

⁴ This can be easily proved by writing out the Taylor expansion of the logarithmic return definition and comparing the first term with the holding-period return definition.

$$IR = \frac{Avg(R - r)}{\sigma}$$

where R is a vector containing the monthly portfolio returns over each period, r is a vector containing the benchmark returns over the corresponding periods and σ is the volatility of the difference between the portfolio returns and benchmark returns. The return of the portfolio in excess of the benchmark's return is also referred to as the 'active return', leading to the Information Ratio also being described as the average of the active return divided by the tracking error or active risk.

In the formula above it is assumed that the benchmark used in the calculation closely matches the style of the portfolio's management in terms of risk⁵. If this is not the case, the Information Ratio needs to be adjusted by an explicit beta parameter to accommodate for the fact that more risk than what is reflected in the benchmark is being taken on.

Similar approximations that exist for the Sharpe Ratio also exist for the Information Ratio. Holding-period returns are used most of the time in the calculations, but as mentioned, this is permissible when the absolute return values are small. The ratio can also be annualised by multiplying it with the square root of 12.

Sortino Ratio

The Sortino Ratio measures whether the portfolio's return in excess of a specified benchmark was sufficient to cover the downside risk inherent in the investment. Downside risk is measured by the volatility of negative active returns. The calculation is similar to calculating tracking error, except that the positive active returns are set equal to zero and still included in the standard deviation calculation.

In the light of the definitions for the Sharpe and Information Ratios, the Sortino Ratio is likewise defined as the average of the active return divided by the downside risk:

$$SorR = \frac{Avg(R - r)}{\sigma'}$$

where R is a vector containing the monthly portfolio returns over each period, r is a vector containing the benchmark returns over the corresponding periods and σ' is the downside risk of the portfolio.

In the calculation of the Sortino Ratio, the benchmark return can also be set to zero, which then indicates whether the portfolio's positive returns were sufficient to cover the risk of negative returns. It is therefore an indicator of capital preservation in nominal terms. The benchmark return can also be equated to the inflation benchmark so that the ratio then indicates whether real returns were sufficient to cover the risk of under-performing inflation. This is an important indicator of a fund's ability to match inflation-adjusted liabilities.

Once again, similar approximations and annualising, as mentioned for the Sharpe and Information ratio, can also be applied to the Sortino Ratio.

⁵ In the CAPM single-factor model framework, this amounts to the case where beta is equal to one.

Example Calculations


An example of calculating the risk-adjusted performance ratios for a fund is shown in Tables 1 and 2 below. The average excess return in Table 1 was calculated by subtracting the average risk-free return from the average fund return, while the average active return was similarly calculated by subtracting the average benchmark return from the average fund return. The risk of the fund, benchmark and the risk-free instrument was calculated using the standard deviation of their returns. The active risk and excess risk was calculated using the standard deviation of the excess and active returns. The downside risk is calculated similarly to the risk calculation using standard deviation of the returns, but with all positive returns set to zero.

Table 1. Average Return, Risk and Downside Risk calculated using the monthly returns for a portfolio, benchmark and risk-free instrument, as well as the excess and active returns of the portfolio.

	Portfolio	Benchmark	Risk-free	Excess	Active
Average Return	10.0%	8.0%	6.0%	4.0%	2.0%
Risk	7.0%	3.0%	1.0%	5.8%	6.5%
Downside Risk	6.5%	2.0%	0.0%	4.8%	3.5%

Table 2. Monthly and annualised risk-adjusted performance ratios calculated for a portfolio calculated using the return and risk values from Table 1. The annualised ratios were obtained by multiplying the monthly ratios by $\sqrt{12}$.

Ratios	Monthly	Annualised
Sharpe (SR)	0.69	2.39
Information (IR)	0.31	1.07
Sortino (SoR)	0.57	1.98



good **thinking**

Established in 1999 as an independent asset liability consultant and risk manager, RisCura provides a variety of solutions via consulting, auditing, risk management and software to pensions funds, asset managers, banks and other financial services businesses.

The RisCura research team is the engine room of our business, ensuring that we remain at the forefront of trends in the retirement fund and asset management industries.

The team focuses on producing high quality, value-added research. Most research topics are inspired by our pension fund clients and aimed at providing them with a sound platform from which to make decisions. In researching topics we evaluate current global ideas and thinking, and then tailor them to suit local market conditions.

This research report has been issued by RisCura Solutions (Pty) Ltd. It may not be reproduced or further distributed or published, in whole or in part, for any purposes. RisCura Solutions (Pty) Ltd has based this document on information obtained from sources they believe to be reliable but which have not been independently verified; RisCura Solutions (Pty) Ltd makes no guarantee, representation or warranty and accepts no responsibility or liability as to its accuracy or completeness. Expressions of opinion herein are those of the author only and are subject to change without notice. This document is not and should not be construed as an offer or the solicitation of an offer to purchase or subscribe or sell any investment. RisCura Solutions (Pty) Ltd and/or their officers, directors and employees may own or have positions in any investment mentioned herein or any investment related thereto and may from time to time add to or dispose of any such investment. RisCura Solutions (Pty) Ltd may have assumed an underwriting position in the securities of companies discussed herein (or investments related thereto) and may also perform or seek to perform investment banking or underwriting services for or relating to those companies.

RisCura Research Team: research@riscura.com

Tel: +27 21 683 7111, Fax: +27 21 683 8277

Colinton House, The Oval, 1 Oakdale Rd, Newlands, Cape Town

