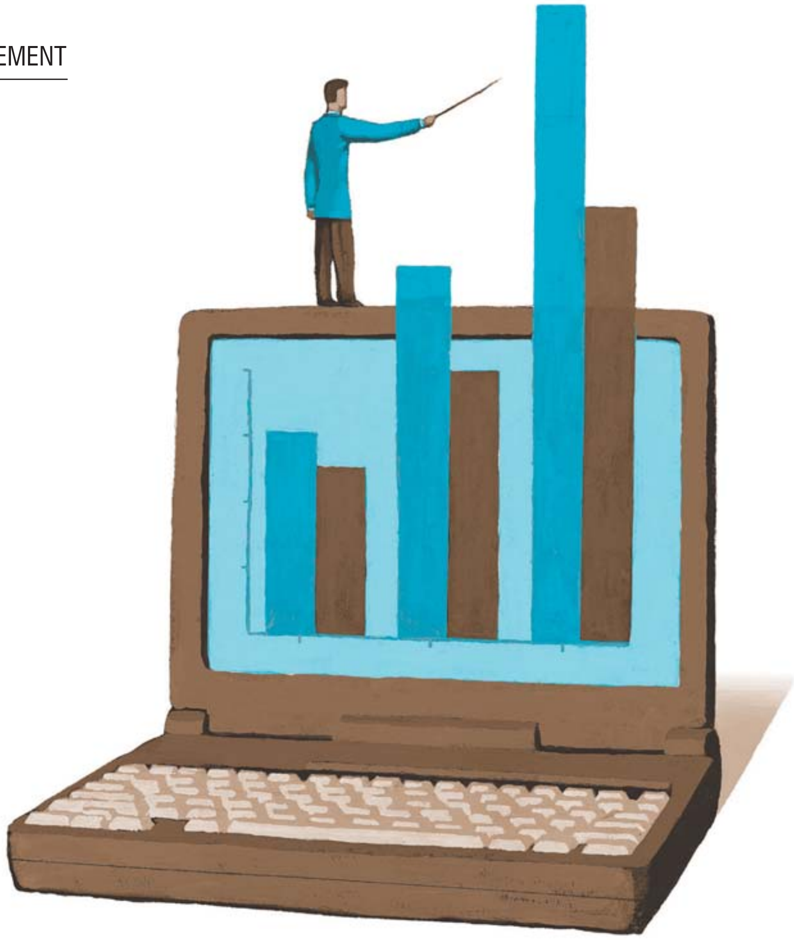

Omega sharpens portfolio performance evaluation

South African equity unit trust data confirm that the Omega ratio is a cut above the Sharpe ratio as a risk-adjusted performance measure.

by Ronel de Wet, Niel Krige and Eon Smit

Investing is a practice that is driven strongly by two factors. Not only are the return prospects important, but also the risks taken to achieve those returns. The two factors are opposite sides of the same coin; one should consider both in order to take sound investment decisions. Condensing risk and return into one useful risk-adjusted number is the key task of investment performance measurement.

The Sharpe ratio, a widely-used measure in financial portfolio performance evaluation, has been criticised lately for its simplistic definition of risk. One of the newer measures that is claimed to avoid the weaknesses of the Sharpe ratio is the Omega ratio. It follows in the footsteps of recent downside risk measurement models, like the Sortino ratio. Downside risk models depart from the conventional Sharpe ratio which assumes all volatility, upside as well as downside, as risk.

A research project at the University of Stellenbosch Business School (USB) examined the effectiveness of Omega as a tool for portfolio evaluation in the South African equity fund market, and found that it has definite advantages over the Sharpe ratio when funds have irregular return distributions.

The Sharpe ratio under scrutiny

Modern portfolio theory has its roots in a paper Markowitz wrote on portfolio selection and diversification in 1952. More papers by pioneers such as Lintner, Treynor, Sharpe and Mossin followed, which eventually led to the formation

THE SHARPE RATIO

The Sharpe ratio is calculated by dividing the average portfolio excess return, over a sample period, by the standard deviation of returns over the same period, as shown by the Sharpe formula:

$S = \{r(p) - r(f)\} \;/\; \sigma(p)$

r(p)	=	return on portfolio
r(f)	=	risk-free rate (or, alternatively, an agreed benchmark rate)
σ(p)	=	standard deviation of portfolio returns

The interpretation of this ratio is simply: the higher the ratio, the greater the risk-adjusted performance. A Sharpe ratio of 1 indicates a risk-adjusted return on investment that is directly proportional to the risk taken in achieving that return.

of the Capital Asset Pricing Model (CAPM) and the concept of the Capital Market Line (CML).

The Sharpe ratio, also known as the reward-to-volatility ratio, has become a long-standing benchmark in portfolio management evaluation. Its formula is based on the principle of the CML and captures the concept of total risk-adjusted performance (see box above).

The Sharpe ratio follows from Markowitz's mean-variance paradigm, which contends that the mean and standard deviation of the distribution of returns are sufficient statistics for evaluating the prospects of an investment portfolio. Central to this reasoning is the assumption that returns follow a normal distribution.

However, this assumption has drawn criticism from analysts. Recent international papers have highlighted that return distributions are not necessarily normal. The increasing popularity of options, futures and hedge funds in portfolios has led to probability distributions that are far from normal. Positive or negative skewness and fatter tails are commonly found in modern-day return distributions, and this phenomenon results in a Sharpe measure that inadequately accounts for all the risks or excess returns to which investors are exposed.

But the Sharpe ratio has other weaknesses too. For example, it gives rise to misleading rankings where the net returns are negative. Moreover, the building blocks of the Sharpe ratio – the expected return and volatility – are quantities that are statistically estimated from a sample of returns and are therefore subject to estimation error. Calculation of Sharpe ratios often occurs without disclosure of the statistical significance of the estimates. Yet the relative accuracy of these estimates should not merely be ignored when one draws a conclusion about the performance of a fund based on the Sharpe ratio, as is often the case.

How does the Omega ratio differ?

The Omega ratio makes no assumptions about the form of the return distribution. Omega is calculated by mathematically integrating all the characteristics of the return distribution without having to estimate the distribution or its parameters. Simply put, the formula calculates the ratio of the weighted probability of returns above a certain threshold to the weighted probability of returns below that threshold (see box **below**).

THE OMEGA RATIO

$$\Omega(r) = \frac{\int_r^b (1 - F(x))\,dx}{\int_a^r F(x)\,dx}$$

Where:

- F(x) is the cumulative distribution of returns
- (a,b) represents the interval of returns
- r is the threshold level set
- Ω(r) is the probability-weighted ratio of gains to losses relative to the threshold r

The Omega ratio involves partitioning returns into losses and gains above and below a return threshold and then considering the probability-weighted ratio of returns above and below the partitioning. The nature of the Omega ratio calculation is illustrated in the diagram below, with the horizontal axis as the return rate and the vertical axis the cumulative probability. The ratio Ω(r) can be visualised as area A divided by area B.

Cumulative return distribution F(x)

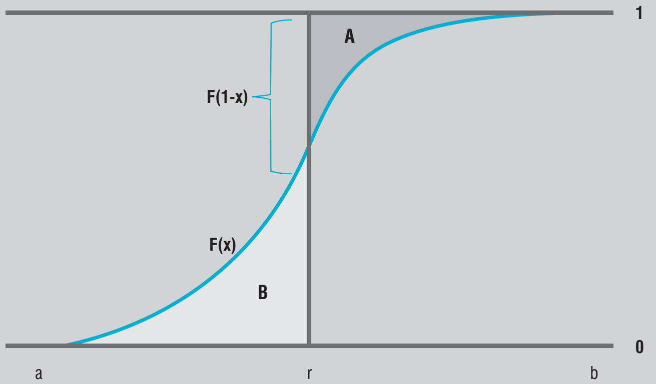

As can be seen, the formula uses the return distribution as it is. As a result, it encodes also the non-normal properties of skew or fat-tail distributions. If r coincides with the mean of the distribution, the ratio will be 1.

With acknowledgement to the paper *An introduction to Omega* by Keating and Shadwick, 2002, The Finance Development Centre Limited.

»

Rankings based on the Omega ratio are always possible, irrespective of the threshold. The Sharpe ratio, in contrast, may give an ambiguous ranking of funds earning negative returns. Another feature of the Omega ratio is that it is flexible with regard to the utility of the investor. The loss threshold is a function of the investor's preferences; the Sharpe measure does not consider this.

The principles behind the Omega ratio expose the limitations in using variance as a proxy for risk

Where returns are normally distributed, or are close to normality, the Omega results will correspond to the Sharpe ratio. In the event of irregular return distributions, however, the Omega ratio will take additional distribution information into account. Whereas the Omega ratio will distinguish between the higher upside potential or higher risk underlying skew distributions – and as such will rank them differently – the Sharpe ratio cannot make this distinction.

The principles behind the Omega ratio expose the limitations in using variance as a proxy for risk. By simply equating variance to risk, the composition of a portfolio may be insured against losses at the cost of the upside potential. Theoretically, therefore, the Omega measure adds another dimension to risk-adjusted performance evaluation, and it makes sense that it be incorporated in the portfolio evaluation process.

Putting Omega and Sharpe to the test

The USB study set out to investigate the difference between applying the Sharpe and Omega ratios in the South African investment environment, with the focus on equity funds.

Return distribution characteristics

As the Omega ratio becomes relevant where return distributions deviate from normality, the research started by examining the return distributions of the South African equity market. The All-Share, Gold, Industrial and Financial indices – which represent the general equity market – were subjected to goodness-of-fit tests in respect of different theoretical distributions, including the normal distribution.

Data over a ten-year period showed that the All-Share and Financial indices had sharper than normal peaks and fatter tails, thereby indicating that the *normal* distribution is not an accurate description of their respective distributions. Moreover, the All-Share, Financial and Industrial indices were negatively skewed, while the Gold index was positively skewed. A deeper analysis revealed that the *extreme* value and the *three-parameter lognormal* distributions described the Gold index well, and the *logistic* distribution described the Financial and Industrial indices well. Interestingly, the combination of these indices resulted in an All-Share index that is best described by the *stable* distribution.

For the five-year period, however, return data showed much more symmetry and less inclination to fatter tails. The *normal* distribution is a reasonably accurate approximation for returns over this period.

Comparing Omega and Sharpe rankings

To compare the rankings of the Omega ratio with those of the Sharpe ratio, monthly return data of general equity unit trusts over various periods were obtained. Since daily returns of unit trusts are not readily available in South Africa, the investigation had to be conducted with monthly data. Five-, three- and one-year-period data were used.

The funds were then ranked by applying the risk-free rate, which was determined by averaging the 91-day Treasury bill rate over the three periods respectively. Analyses were performed based on all funds that had monthly returns for the full period being ranked.

a. Five-year period

For the five-year period, very little difference was found between the Omega-based and the Sharpe-based rankings. The threshold rate applied was 0,75% (monthly return). A rank order correlation test revealed that the similarity was statistically significant at the 1% level.

b. Three-year period

The analysis based on three-year data showed that, overall, the rankings were very similar. As in the five-year case, a rank order correlation test found a significant relationship between the rankings at the 1% level of significance. There were, however, three funds whose rankings differed by more than 10 places out of the total of 53 funds ranked.

To gain insight into the reason for the difference in rankings, two funds – Investec Equity Fund A and Old Mutual Investor's Fund R – were selected for further investigation. The Investec fund was ranked higher by Sharpe, but lower by Omega. The descriptive statistics shown in the table **below** form the rationale for the Sharpe ranking:

Descriptive statistics: Investec vs Old Mutual

	Investec Equity Fund A	OM Investor's Fund R
Mean	0.0361	0.0316
Standard deviation	0.0389	0.0381
Kurtosis	-0.7802	-1.0250
Skewness	-0.1671	0.2877
Minimum	-0.0342	-0.0331
Maximum	0.1040	0.1089

The higher ranking of the Investec fund based on the Sharpe ratio is understandable on account of the higher mean achieved, even though some of this benefit is offset by the higher standard deviation (risk) of the Investec fund. The Old Mutual fund has a positive skewness that could indicate a heavier positive tail in the distribution.

The difference in Omega ranking is illustrated by the **following** table:

Analysis of Omega value: Investec vs Old Mutual

	Investec Equity Fund A	OM Investor's Fund R
Probability-weighted average gains	0.03523	0.02973
Probability-weighted average loss	0.00562	0.0046
Omega value (PWA gains/PWA loss)	6.2	6.4

This table **(p15)** shows that the Old Mutual fund enjoys a higher Omega ranking not as a result of a heavier positive tail but because of a lower probability-weighted loss. The Old Mutual fund's probability-weighted average gain was 15.6 % less than that of the Investec fund, which was more than offset by a 18.1 % lower probability-weighted average loss. This results in the higher Omega ratio for the Old Mutual fund.

c. One-year period

The analysis of one-year data once again indicated a similar ranking of funds. There were, however, eight funds whose rankings differed by more than 15 places – of which five funds differed by more than 20 places – out of the total of 78 funds ranked. To investigate this difference in more detail, a comparison was done between the Stanlib R fund and the Sanlam A fund. The descriptive statistics are shown in the box **below:**

Descriptive statistics: Stanlib R vs Sanlam A

	Stanlib Fund R	Sanlam Fund A
Mean	0.0398	0.0405
Standard deviation	0.0443	0.0412
Kurtosis	0.1859	-0.7064
Skewness	0.5535	-0.6791
Minimum	-0.0318	-0.0362
Maximum	0.1235	0.0857

The Omega-based calculations are shown **as follows**:

Analysis of Omega value: Stanlib R vs Sanlam A

	Stanlib Fund R	**Sanlam Fund A**
Probability-weighted average gains	0.03775	0.0405
Probability-weighted average loss	0.00352	0.00567
Omega value (PWA gains/PWA loss)	10.7	7.1

The table **above** shows that the comparatively low probability-weighted average loss of the Stanlib fund gives rise to its higher Omega ranking.

This particular analysis highlights some of the key advantages of the Omega ratio over the Sharpe ratio. Sharpe penalises the Stanlib fund based on the larger standard deviation, which it interprets as risk. But the larger standard deviation of Stanlib is largely attributable to the positive returns being generated by the fund; it earned a maximum of 12.35 % versus the 8.57 % of Sanlam. This penalty is further accentuated by the Sharpe ratio that ignores the lower probability-weighted average loss of 0.00352 of the Stanlib fund, as opposed to the 0.00567 of the Sanlam fund.

Sharpe penalises the Stanlib fund based on the larger standard deviation, which it interprets as risk

Omega a finer tool for evaluating performance over shorter periods

The Omega measure enables one to incorporate the characteristics of return distributions into performance measurement without having to determine their shape. The importance of this extra sophistication has been demonstrated by the research. By applying the Omega ratio to empirical data, it was shown that in cases where the return distributions departed from normal it ranked the fund quite differently from its Sharpe ranking. Specifically, the study showed that Omega places a finer focus on the true risk profile underlying the return distribution of the particular fund.

The research also revealed that the equity unit trust environment is relatively homogeneous over the medium term, as shown by the normally distributed returns over a five-year period. Therefore the Omega and Sharpe rankings were very similar. The similarity of rankings in this case is precisely how Omega functions. The more return distributions approximate normality, the more evaluations based on Omega will approximate Sharpe rankings.

Over the shorter period, where certain return distributions deviated more noticeably from normality, Omega provided its richer insight. The added advantage of Omega is that when one evaluates performance over shorter periods, there are fewer returns on which to estimate a distribution. The practice to assume normality is therefore even more questionable for shorter investigation periods. What gives Omega the edge is that it does not estimate the distribution; it simply uses each bit of return information to make its calculation.

Ronel de Wet

Ronel de Wet, Niel Krige and Eon Smit published this research under the title *An investigation into performance rankings of the Omega ratio vs the Sharpe ratio applied to South African general equity unit trusts* in *Studies in Economics and Econometrics*, 32(2), 2008, and as *An investigation into the return distributions of the South African general equity universe in Management Dynamics: Journal of the Southern African Institute for Management Scientists*, 16(3), 2007. The articles were based on De Wet's MBA research project *The evaluation of Omega as an effective tool for portfolio evaluation in the South African context*, supervised by Prof Niel Krige and Prof Eon Smit, which was presented to the University of Stellenbosch Business School (USB) in December 2006.

Write to us: Do you believe that the Omega ratio should be used more extensively by the unit trust industry to report risk-adjusted investment performance?
lab@usb.ac.za