W McLeod and G van Vuuren*

Interpreting the Sharpe ratio when excess returns are negative

1. INTRODUCTION

Fund performance evaluation has evolved from the simple comparison of portfolio returns, to the comparison of excess portfolio returns (returns relative to a risk-free rate), to the highly acclaimed Sharpe Ratio in which excess portfolio returns per unit of risk taken to achieve those returns are compared. Of these developments, the Sharpe Ratio enjoys arguably the greatest success and most widespread implementation. It has undergone several refinements and augmentations in its 37-year life, but the basic concept has survived remarkably intact and with few modifications.

Whilst one of the aims of good fund management remains the active selection of those funds with the largest Sharpe Ratio, the aim of this article is to introduce an alternative interpretation and definition of this procedure. The standard view suggests that superior funds are those having larger Sharpe Ratios, i.e. the higher excess returns per unit of risk. We show that, in fact, the choice should be based upon the maximum probability of outperforming the risk free rate – a subtle but important shift in emphasis.

Section 2 presents an overview of some of the issues that have shaped the evolution of the 'standard' Sharpe Ratio. These developments necessarily include criticisms and suggested improvements as well as adaptations of the Sharpe Ratio as markets have changed and new products and procedures have evolved. This Section is by no means exhaustive and aims to provide only a general background to what has become an extensive body of research.

Section 3 provides a brief overview of the mathematical definition (and graphical explanation) of the standard interpretation of the Sharpe Ratio (*ex post* and *ex ante*). This formulation works well when excess returns are, or are expected to be positive, but a flaw arises in this interpretation when negative excess returns are used. Whilst it could be argued that investors would not invest in funds offering negative excess returns (*ex ante*) and hence the Sharpe Ratio would provide no useful information in these situations, we argue (section 4) that, with our new interpretation, the Sharpe Ratio may be, and should continue to be used even when excess returns are *expected* to be negative (and should certainly continue to be used for *positive* excess returns).

Section 4 introduces a formal, probabilistic description of an alternative interpretation of the Sharpe Ratio and provides supporting graphical and mathematical explanations to substantiate the concept. We show that when the new interpretation is applied to funds, it generates a consistent ranking of these funds under all potential excess return scenarios, not only when excess returns are, or are expected to be positive. Some examples are provided to illustrate the consistency of this new interpretation and section 5 concludes the article.

2. THE SHARPE RATIO

The Sharpe Ratio was first introduced by Sharpe (1966) as a measure for the comparison of mutual fund performance. Originally called the *reward-to-variability ratio*, it quickly gained widespread popular acceptance and today enjoys almost ubiquitous implementation in the financial world. Despite this, or because of it, a number of generalisations and improvements to this statistic have subsequently been suggested within the portfolio management literature.

The Sharpe Ratio concept rests upon the Markowitz mean-variance paradigm, which assumes (a) that the distribution of one-period portfolio returns is normal and (b) that the mean and the standard deviation of the distribution are necessary and sufficient statistics for evaluating the risk-adjusted performance of an investment portfolio. Sharpe (1966) originally intended the ratio to be used *ex ante*, i.e. as a performance measure forecast, but it has subsequently been widely implemented as an *ex post* measure as well – to record and rank portfolio (and/or fund manager) performance. Sharpe (1994) explored this issue of *ex post* versus *ex ante* values and contended that historic data, used for the computation of most performance measures, is justified on the basis of predicted relationships. Practical implementations use *ex post* results while theoretical discussions focus largely on *ex ante* values. Implicitly or explicitly then, historic results are generally assumed to have at least some predictive ability. Sharpe did, however, point out that the use of unadjusted historic (*ex post*) Sharpe Ratios as surrogates for unbiased predictions of *ex ante* ratios "remained subject to serious question" (Sharpe, 1994). Despite such caveats and in the absence of significantly better performance measures, Sharpe contended that "there is much to recommend a measure that at least takes into account both risk and expected return over any alternative that

*Respectively from Quantitative Risk Group, Old Mutual Asset Managers, PO Box 878, Cape Town 8000, Republic of South Africa and Market Risk Group, Standard Bank London Ltd, Cannon Bridge House, 25 Dowgate Hill, London EC4R 2SB, United Kingdom.
Email: wmcleod@omam.com and
gary.vanvuuren@standardbank.com

focuses only on the latter."

In the same article, Sharpe examined the important issue that the Sharpe Ratio does not accommodate fund return correlations. The *ex ante* Sharpe Ratio takes into account both the expected differential return and the associated risk, while the *ex post* version takes into account both the average differential return and the associated variability. Neither of these, however, incorporates information about the correlation of a fund or strategy with other assets, liabilities, or previous realizations of its own return (but see Dowd, 2000). For this reason, the Sharpe Ratio may need to be supplemented in certain applications.

Earlier work of Sharpe (1975, 1992) focused attention on the use of asset class factor models as a means of assessing fund manager performance. Sharpe recommended grouping portfolio exposures into various 'major' asset classes and then gauging that component's performance relative to a similarly composed benchmark fund. This process assists in the determination of any value added by portfolio managers and helps "make order out of chaos".

Dowd (1999) highlighted the importance of dealing with net rather than gross portfolio exposures and suggested implementing a value at risk (VaR) approach to risk-return analysis, rather than the Sharpe Ratio. This methodology (used for investment, hedging and general portfolio management decisions) is particularly useful when making hedge decisions by helping to avoid a number of problems that easily arise using the traditional approach to hedging and are not addressed by the Sharpe Ratio.

Dowd (2000) drew attention to the limiting assumption that returns of the assets used for the calculation of the Sharpe Ratio are normally distributed (but see Černý, 2003). The assumption that individual asset returns are uncorrelated with the returns to an institution's existing portfolio was also addressed. Dowd's solution to both of these limitations was to construct Sharpe Ratios for each of the alternative portfolios under consideration rather than the alternative investments themselves. Dowd's suggestion (or 'new Sharpe rule') is valid regardless of the correlations of prospective new assets with the existing portfolio. These correlations can take any value and are implicitly permitted in the construction of the portfolio Sharpe Ratio.

Harding (2002) criticised the definition of 'risk' in the construction of the Sharpe Ratio. He argued that the return of an asset or portfolio is a "definite and meaningful quantity" whilst risk is not. Risk may be easily calculated as the standard deviation of a set of return data, but it is not obvious that the meaning of 'risk' will be the same for all time series. For a standard deviation to be a meaningful statistic, it must be generated from a stationary and parametric process and it is not at all always true that the return series is thus defined.

In the same article, Harding (2002) considered the problem that large positive returns are penalised in the calculation of the Sharpe Ratio. The removal of the highest returns from a data series could actually increase the Sharpe ratio by reducing the 'risk' – a case of *reductio ad absurdum*. An associated measure (the Sortino Ratio), which considers only the negative, semi-standard deviation in the denominator is sometimes considered an improvement, but the assumptions of stationarity and parametricity remain.

Lo (2002) argued that the constituents of the Sharpe ratio are unknown quantities that must be estimated statistically and are, therefore, subject to estimation error. The degree of accuracy in the measurement of Sharpe Ratios must therefore be addressed. Lo derived explicit expressions for the statistical distribution of the Sharpe Ratio using standard asymptotic theory under several sets of assumptions for the return-generating process – independently and identically distributed returns, stationary returns, and with time aggregation. He showed that monthly Sharpe ratios could not be annualized using 'square root of time' rule except under very special circumstances, and derived the correct method of conversion in the general case of stationary returns.

Goetzman, Ingersoll, Spiegel and Welch (2002) showed that the Sharpe Ratio and other related reward-to-risk measures may be manipulated with option-like strategies. They derived the general conditions for achieving the maximum expected Sharpe ratio, static rules for achieving the maximum Sharpe Ratio with two or more options, as well as a continuum of derivative contracts. Dynamic rules for increasing the Sharpe Ratio were also derived.

Černý (2003) found that the Sharpe Ratio is closely related to quadratic utility and argued that there is a one-to-one relationship between the maximum quadratic utility attainable in the market and the Sharpe Ratio. Černý extended the definition of the Sharpe Ratio from quadratic utility to an entire family of CRRA (Constant Relative Risk Aversion) utility functions and restated the equilibrium restrictions in terms of 'Generalised Sharpe Ratios' (first introduced by Dowd, 1999). Unlike the standard Sharpe Ratio, the Generalised Sharpe Ratio provides a consistent ranking of investment opportunities even when asset returns are non-normal. Arbitrage adjusted Sharpe Ratios are also discussed in Černý (2004), to be published at a later date.

These are some of the many articles in the literature which address supplementary – rather than alternative – measures to the Sharpe Ratio of evaluating fund performance. They all aim to augment or improve upon Sharpe's original construct, but as testimony to its robust formulation, it still survives virtually unchanged.

3. FUND SELECTION USING THE SHARPE RATIO

It is useful to step back at this point and briefly examine the mathematical description of the *ex ante* and *ex post* Sharpe Ratios.
The *ex ante* Sharpe Ratio may be expressed mathematically as

$$SR = \frac{E(R_P) - R_F}{\sigma_P} \qquad \ldots (1)$$

where

SR is the Sharpe Ratio,

$E(R_P)$ is the fund's expected return,

R_F is the risk free rate and

σ_P is the fund's expected standard deviation. The quantity $E(R_P) - R_F$ is the excess return.

The *ex post* Sharpe Ratio uses *historical* returns, and becomes

$$SR = \frac{\bar{R}_P - R_F}{S_P}.$$

Here, the expected returns, $E(R_P)$, have been replaced by the average of the *historical* return series, $\bar{R}_P$, and σ_p has been replaced by the standard deviation of the *historical* return series, S_p. Note that Sharpe (1966, 1992) often uses the return series average in this way. Essentially, a sample of realised monthly returns are used to approximate the true distribution. Consequently, the probability of outperforming the risk free rate, based on a distribution of returns similar to the realised series, may be calculated.

Assume two funds, A and B. If Fund A's Sharpe Ratio were larger than Fund B's, Fund A would be expected to be the better-performing fund (note that this argument applies equally well to both *ex post* and *ex ante* measures of risk and reward). For example:

Risk free rate = 10%	Fund A	Fund B
Expected return [$E(R_P)$]	20%	25%
Risk (σ_P)	5%	10%
Sharpe Ratio	2,0	1,5

Ranking these funds using the standard interpretation classifies (correctly) Fund A as being more desirable than Fund B since $SR_A > SR_B$. The risk and excess return profiles of the two funds are shown below in Figure 1. The fund that has the steeper gradient drawn from the risk-free rate is the fund with the larger Sharpe Ratio.

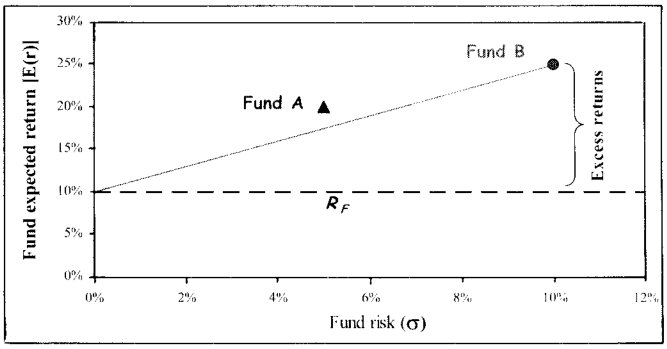

Figure 1: Risk and excess return profiles (relative to the risk-free rate) for Funds A and B. The dashed line shows the level of the risk-free rate.

There are, of course, no guarantees of positive excess returns for any asset or portfolio. When these excess returns are negative, problems arise for the standard Sharpe Ratio interpretation. Consider first, the case where the funds have the parameters as shown in the table below. The expected returns for both funds are identical and both are slightly smaller than the risk-free rate.

Risk free rate = 10%	Fund C	Fund D
Expected return [$E(R_P)$]	9,9%	9,9%
Risk (σ_P)	15%	10%
Sharpe Ratio	-0,007	-0,010

Choosing the fund with the maximum Sharpe Ratio indicates that Fund C is preferable to Fund D. In standard parlance, Fund C, then, with the higher Sharpe Ratio, is the fund with "the higher expected return per unit of risk" – but the funds have identical returns and Fund C has the higher risk! The standard interpretation of the Sharpe Ratio is clearly at odds with this ranking and requires some modification.

Now consider the case when the funds have the following parameters:

Risk free rate = 10%	Fund E	Fund G
Expected return [$E(R_P)$]	−12%	−8%
Risk (σ_P)	15%	10%
Sharpe Ratio	−1,467	−1,800

Again, choosing the fund with the maximum Sharpe Ratio indicates that Fund E is preferable to Fund G, or, the standard definition is that Fund E has "the higher return per unit of risk". But examination of Fund E's Sharpe Ratio shows that it comprises a *lower* excess return and a *higher* risk than Fund G. This clearly presents a problem for a naïve interpretation of the standard Sharpe Ratio.

4. PROBABILISTIC INTERPRETATION

The definition of maximising the Sharpe Ratio within a risk-versus-reward paradigm fails when excess returns (*ex post* or *ex ante*) are zero or negative. A broader definition, encompassing situations in which excess returns may take on any value, must be sought.

Closer examination of the Sharpe Ratio reveals the following probabilistic equivalence[1]:

$$\begin{aligned}
\text{Max Pr}\,(R_P \geq R_F) &\Leftrightarrow \text{Max Pr}\left(\frac{R_P - E(R_P)}{\sigma_P} \geq \frac{R_F - E(R_P)}{\sigma_P}\right) \\
&\Leftrightarrow \text{Max Pr}\left(z \geq \frac{R_F - E(R_P)}{\sigma_P}\right) \quad \text{where } z \sim N(0,1) \\
&\Leftrightarrow \text{Min}\left(\frac{R_F - E(R_P)}{\sigma_P}\right) \quad \text{or, alternatively} \\
&\Leftrightarrow \text{Max}\left(\frac{E(R_P) - R_F}{\sigma_P}\right) \\
&= \text{Max}(\text{Sharpe Ratio}).
\end{aligned}$$

This, then, represents an alternative interpretation for the goal of selecting funds with the highest Sharpe Ratio. Of any fund set, selecting the best-performing fund traditionally means selecting that fund with the maximum Sharpe Ratio. We have shown that this selection process is equivalent to selecting that fund with the maximum probability of outperforming the risk-free rate. A numerical example, using the parameters from Funds E and G above, illustrates the point. The probability of Fund E achieving a greater excess return than the risk-free rate is:

$$\begin{aligned}
\text{Probability}_{R_E > R_F} &= \Pr\left(z \geq \frac{10\% - (-12\%)}{15\%}\right) \\
&= \Pr(z \geq 1.467) \\
&= 0.071,
\end{aligned}$$

whilst the probability of Fund G achieving an excess return greater than the risk-free rate is:

$$\begin{aligned}
\text{Probability}_{R_G > R_F} &= \Pr\left(z \geq \frac{10\% - (-8\%)}{10\%}\right) \\
&= \Pr(z \geq 1.800) \\
&= 0.004.
\end{aligned}$$

Fund E, therefore, has a marginally better chance of outperforming the risk free rate than Fund G (or Fund E has a slightly higher probability of obtaining a positive excess return than Fund G).

This example is shown graphically below in Figures 2a and b in which the risk associated with respective funds is shown as a normal distribution around the excess return for each particular fund. Figure 2a shows the case when funds have small, negative, identical excess returns, but different risks (clearly showing why Fund C is superior to Fund D when employing the new definition. Fund C has a higher probability of outperforming the risk-free rate as evidenced by the larger proportionate area in the tail of the probability distribution above the risk-free rate). Figure 2b shows this situation for unequal, negative, excess returns.

It is not immediately obvious that Fund E is superior to Fund G since Fund E has a lower excess return coupled with a greater level of risk. It may be argued, however, that by holding the constituents of the portfolio in the same proportion except for cash, which could be increased, one could quite easily 'move up' the line that joins the risk free rate (cash) to Fund E in Figure 2b. This would then have created a portfolio which has a higher expected return and a smaller risk than Fund G: this new 'adjusted' Fund E, then, is certainly superior.

One might claim that one would not even consider Fund E in its "raw" state, but one requires caution for such a view. If one makes an *ex post* comparison between Fund E and Fund G, one could not adjust or reconstruct Fund E. If one's objective were to outperform the risk free rate (cash) then Fund E would certainly have provided a better chance of achieving this objective than Fund G.

If one were making an *ex ante* comparison between the two funds then one could certainly move up the straight line that joins the risk free rate (cash) to Fund E. When moving up this straight line one maintains the same gradient and consequently the same Sharpe Ratio. In reality, however, constraints are imposed on cash-holding: for example, an equity portfolio would certainly not consist of 50% cash, let alone 100% cash.

A fund manager's primary objective might be to maximise the probability of achieving a return greater than the risk free rate (cash), but a secondary objective – to outperform the risk free rate by a specified percentage – might also be considered. If this were the case then the further a portfolio's profile is moved up the line joining the risk free rate to Fund E (maintaining the same Sharpe ratio) the smaller the probability of that portfolio outperforming the risk free rate by the specified percentage.

[1]There are many ratios in use in fund management and assessment, for example the Sortino (Pederson, 2002) and Information (Grinold, 1989, Goodwin, 1998) Ratios. Since quotients of excess returns and risk – in various guises – are often involved, the same 'new interpretation', applied here to the Sharpe Ratio, also pertains to these measures.

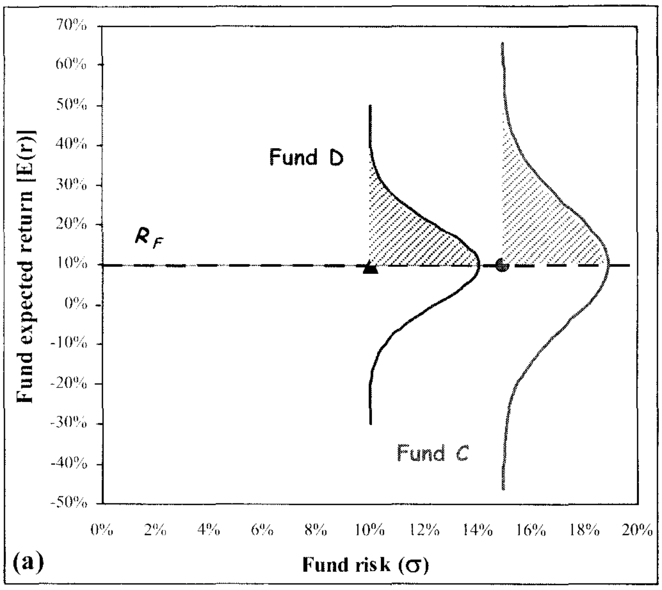

Figure 2a: Low excess returns (close to 0%) and corresponding return distributions with the dashed line showing the level of the risk free rate. The area in the upper tail of each distribution – *above* the risk free rate (shown shaded) – is the probability of the portfolio return being greater than the risk free rate of return.

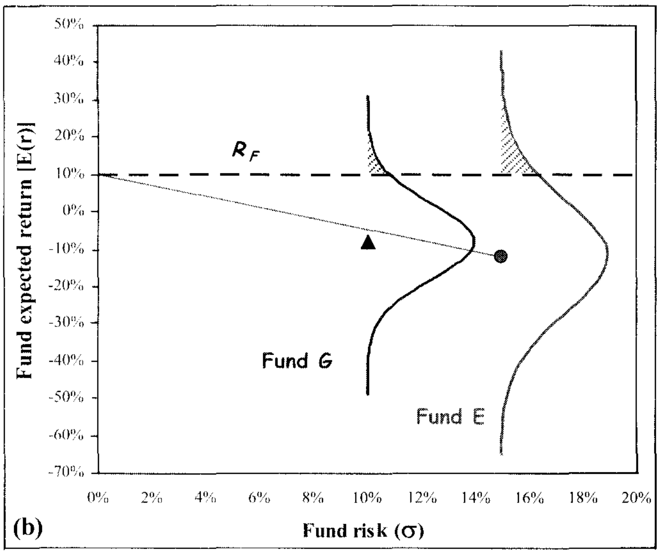

Figure 2b: Negative excess returns and corresponding return distributions. Again, the area in the upper tail of each distribution – that area *above* the risk free rate (shown shaded) – is the probability of the portfolio return being greater than the risk free rate of return.

5. CONCLUSION

We have explored the widely accepted paradigm that the fund with the largest Sharpe Ratio is that one which has the largest excess return per unit of risk. Whilst this may hold true for positive excess returns, it is not true for negative excess returns. We showed that superior fund selection may be realized using the basic goal of maximising the Sharpe Ratio and that this is valid for *all* values of excess returns. The interpretation of that goal, however, requires a shift in emphasis, namely: the fund with the maximum Sharpe ratio is that fund with the highest probability of outperforming the risk-free rate, not necessarily the fund with the largest excess return per unit of risk. We examined and proved this claim using probabilistic mathematics. It will be good news for the investment community everywhere that the robust Sharpe Ratio has survived yet another inspection with only a minor interpretation shift.

REFERENCES

Černý A. 2003. Generalised Sharpe ratios and asset pricing in incomplete markets. *European Finance Review*, 7:191-233.

Černý A. 2004. Mathematical techniques in finance: tools for incomplete markets. *Princeton University Press*. Release date: January, 2004.

Dowd K. 1999. A value at risk approach to risk-return analysis. *Journal of Portfolio Management*, 25(4):60–67.

Dowd K. 2000. Adjusting for risk: An improved Sharpe ratio. *International Review of Economics and Finance*, 9(3):209-222.

Goetzman WJ, Ingersoll M, Spiegel P and Welch I. 2002. Sharpening Sharpe ratios. *Yale ICF Working Paper* No. 02-08; AFA 2003 Washington, DC Meetings. http://ssrn.com/abstract=302815.

Goodwin TH. 1998. The information ratio. *Financial Analysts Journal*, 54(4):34–43.

Grinold RC. 1989. The fundamental law of active management. *Journal of Portfolio Management*, 15(3):30–37.

Harding DA. 2002. Critique of the Sharpe ratio. *Winton Capital Management*, (www.wintoncapital.com/sharpe.htm), London.

Lo AW. 2002. The statistics of Sharpe ratios, *Financial Analysts Journal*, 58(4):36–52.

Pedersen CS and Satchell SE. 2002. On the foundation of performance measures under asymmetric returns. *Quantitative Finance*, 2(June 2002):217-223.

Sharpe WF. 1966. Mutual fund performance. *Journal of Business*, Supplement on Security Prices, 39:119–138.

Sharpe WF. 1975. Adjusting for risk in portfolio performance measurement. *Journal of Portfolio Management*, Winter:29–34.

Sharpe WF. 1992. Asset allocation: Management style and performance measurement. *Journal of Portfolio Management*, Winter:7–19.

Sharpe WF. 1994. The Sharpe ratio. *The Journal of Portfolio Management*, Fall:49–58.