


An Introduction to Omega[®]

Con Keating and William F. Shadwick

These distributions have the same mean and variance.



Are you indifferent to their risk-reward characteristics?

From Alpha to Omega

Comprehensive Performance Measures

While everyone knows that mean and variance cannot capture all of the risk and reward features in a financial returns distribution, except in the case where returns are normally distributed, performance measurement traditionally relies on tools which are based on mean and variance. This has been a matter of practicality as econometric attempts to incorporate higher moment effects suffer both from complexity of added assumptions and apparently insuperable difficulties in their calibration and application due to sparse and noisy data.

A measure, known as Omega, which employs all the information contained within the returns series was introduced in a recent paper¹. It can be used to rank and evaluate portfolios unequivocally. All that is known about the risk and return of a portfolio is contained within this measure. With tongue in cheek, it might be considered a Sharper ratio, or the successor to Jensen's alpha.

The approach is based upon new insights and developments in mathematical techniques, which facilitate the analysis of (returns) distributions. In the simplest of terms, as is illustrated in Diagram 1, it involves partitioning returns into loss and gain above and below a return threshold and then considering the probability weighted ratio of returns above and below the partitioning.




Diagram 1

The cumulative distribution F for Asset A, which has a mean return of 5. The loss threshold is at $r=7$. I_2 is the area above the graph of F and to the right of 7. I_1 is the area under the graph of F and to the left of 7. Omega for Asset A at $r=7$ is the ratio of probability weighted gains, I_2 , to probability weighted losses, I_1 .

THE FINANCE DEVELOPMENT CENTRE LIMITED

By considering this Omega ratio at all values of the returns threshold, we obtain a function which is characteristic of the particular asset or portfolio. We illustrate this in Diagram 2.




Diagram 2

Omega for Asset A as a function of returns from $r=-2$ to $r=15$. Omega is strictly decreasing as a function of r and takes the value 1 at Asset A's mean return of 5.

The evaluation statistic Omega has a precise mathematical definition as:

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

where (a,b) is the interval of returns and F is the cumulative distribution of returns. It is in other words the ratio of the two areas shown in Diagram 2 with a loss threshold set at the return level r . For any return level r , the number $\Omega(r)$ is the probability weighted ratio of gains to losses, relative to the threshold r .

The Omega function possesses many pleasing mathematical propertiesⁱⁱ that can be intuitively and directly interpreted in financial terms. As is illustrated above, Omega takes the value 1 when r is the mean return. An important feature of Omega is that it is not plagued by sampling uncertainty, unlike standard statistical estimators—as it is calculated directly from the observed distribution and requires no estimates. This function is, in a rigorous mathematical sense, equivalent to the returns distribution itself, rather than simply being an approximation to it. It therefore omits none of the information in the distribution and is as statistically significant as the returns series itself.

As a result, Omega is ideally suited to the needs of financial performance measurement where what is of interest to the practitioner is the risk and reward

THE FINANCE DEVELOPMENT CENTRE LIMITED

characteristics of the returns series. This is the combined effect of all of its moments, rather than the individual effects of any of them—which is precisely what Omega provides.

Now to use Omega in a practical setting, all that is needed is a simple decision rule that we prefer more to less. No assumptions about risk preferences or utility are necessary though any may be accommodated. The Omega function may be thought of as the canonical risk-return characteristic function of the asset or portfolio.

In use, Omega will usually show markedly different rankings of funds, portfolios or assets from those derived using Sharpe ratios, Alphas or Value at Risk, precisely because of the additional information it employs. In the cases where higher moments are of little significance, it agrees with traditional measures while avoiding the need to estimate means or variances. In those cases where higher moments do matter—and when they do their effects can have a significant financial impact—it provides the crucial corrections to these simpler approximations. It also makes evident that at different levels of returns, or market conditions, the best allocation among assets may change.

In many respects, Omega can be thought of as a pay-off function, a form of bet where we are considering simultaneously both the odds of the horse and its observed likelihood of winning. Omega provides, for each return level, a probability adjusted ratio of gains to losses, relative to that return. This means that at a given return level, using the simple rule of preferring more to less, an asset with a higher value of Omega is preferable to one with a lower value. We illustrate the use of the Omega function in a choice between two assets. Diagram 3 which shows their Omegas as functions of the return level r .




Diagram 3

THE FINANCE DEVELOPMENT CENTRE LIMITED

Notice that the mean return of asset A (the point at which Ω_A is equal to 1) is higher than that of Asset B. The point I, where the two Omegas are equal, is an indifference point between A and B. At return levels below this point, using just the “we prefer more to less” criterion, we prefer Asset B, while above we prefer Asset A. This phenomenon of changes of preference, crossings of the Omega functions, is commonplace and multiple crossings can occur for the same pair of assets. The additional information built into Omega, can lead to changes in rational preferences which cannot be predicted using only mean and variance.

We remark that in the example, Asset A is riskier than Asset B in the sense that it has a higher probability of extreme losses and gains. This aspect of risk is encoded in the slope of the Omega function: the steeper it is, the less the possibility of extreme returns. A global choice in this example involves an investment trade-off between the relative safety of Asset B compared to Asset A and the reduced potential this carries for large gains.

We may use the Omega functions calculated over a selected sequence of times to investigate the persistence or skill in a manager’s performance. Some preliminary work suggests that there is far more persistence than academic studies have indicated previously.

The Omega function can also be used in portfolio construction, where markedly different weights from those derived under the standard mean variance analysis of Markowitz are obtained. In fact those “efficient” portfolios can be shown to be a limited special case approximation within the more general Omega framework.

Omega, when applied to benchmark relative portfolios, provides a framework in which truly meaningful tracking error analysis can be carried out, a significant expansion of existing capabilities.

All things considered, Omega looks set to become a primary tool for anyone concerned with asset allocation or performance evaluation. Particularly those concerned with alternative investments, leveraged investment or derivatives strategies. The first of a new generation of tools adapted for real risk reward evaluation.

Ω Notes: Ω and Normal distributions

Here we consider the simplest application of all, to returns distributions which are normal. The approach to ranking such distributions via the Sharpe ratio involves an implicit choice to consider the possibility of a return above the mean and a return below the mean as equally 'risky'. For two normal distributions with the same mean, the Sharpe ratio favours the one with the lower variance, as this minimizes the potential for losses. Of course it also minimizes the potential for gains. Thus, the use of variance as a proxy for risk considers the downside as more significant than the upside, even in the case where these are equally likely.

Here we consider two assets, A and B, which both have a mean return of 2 and have standard deviations of 3 and 6 respectively. Their probability densities are shown in Diagram 1. In terms of their Sharpe ratios, A is preferable to B. If we were to rank these assets in terms of their potential for gains however, the rankings would be reversed.

Consider the ranking which an investor who requires a return of 3 or higher to avoid a shortfall might make. From this point of view a return below 3 is a loss, while one above 3 is a gain. To assess the relative attractiveness of assets A and B such an investor must be concerned with the relative likelihood of gain or loss. It is apparent from Diagram 1 that this is greater for asset B than for asset A.




Diagram 1. The distributions for assets A and B with a loss threshold of 3

This is due to the fact that the distribution for B has substantially more mass to the right of 3 than the distribution A. For asset B, about 43% of the returns are above 3 while for asset A the proportion drops to 37%. The ratios of the likelihood of gain to loss are 0.77 for B and 0.59 for A.

From this point of view, our rank order is the reverse of the ranking by Sharpe ratios. We are not simply reversing the Sharpe ratio bias however. If the investor's loss

THE FINANCE DEVELOPMENT CENTRE LIMITED

threshold were placed at a return of 1 rather than 3, the same process would lead to a preference for A over B. The ratios of the likelihood of gain to loss with the loss threshold set at a return of 1 are 1.71 for A and 1.31 for B. Clearly, at any loss threshold above the mean the preference will be for B over A, while for any loss threshold below the mean the preference will be reversed. With the loss threshold set at the mean, both assets produce a ratio of 1.

Like this simple process, the use of Ω treats the potential for gains and losses on an equal footing and provides rankings of A and B which depend on a loss threshold. The function $\Omega(r)$ compares probability weighted gains to losses relative to the return level r . As a result, rankings will vary with r . Diagram 2 shows the Ω s for assets A and B. For any value of r greater than the common mean of 2, the probability weighted gains to losses are higher for asset B than for asset A. For any value of r less than the mean, the rankings are reversed. The relative advantage of asset A to asset B declines smoothly as r approaches their common mean of 2 and thereafter the relative advantage of B to A increases steadily.




Diagram 2. Omega for assets A B as a function of return level r

Ω Notes: The Omega of a Sharpe Optimal Portfolio

The mean-variance approach to performance measurement and portfolio optimization is based on an approximation of normality in returns. In this note we show that even in the case of two assets with normally distributed returns, a portfolio which maximizes the Sharpe ratio will be sub-optimal over a significant range of returns. This is a manifestation of the inherent bias in regarding losses and gains as equally 'risky'. As Omega rankings change with the level of returns, an Omega optimal portfolio's composition will vary over different ranges of returns. This extra flexibility can be very important as we show here. A portfolio composition which is independent of returns level and is optimal on the downside, as the Sharpe optimal portfolio is, must be sacrificing considerable upside potential.

We consider two assets, A and B which have independent, normally distributed returns with means and standard deviations of 6 and 4 and 7 and 3 respectively.

We let a denote the weight of asset A and $1-a$ the weight of asset B in the portfolio.

The Sharpe ratio for the portfolio is then $SR(a) = \frac{6a + 7(1-a)}{\sqrt{16a^2 + 9(1-a)^2}}$, which has its

maximum value at about $a=0.68$. The Sharpe optimal portfolio has normally distributed returns with a mean and standard deviation of 6.7 and 2.4. The portfolio distribution is shown in Diagram 1, with the distributions for assets A and B.




Diagram 1. The distributions for asset A, asset B and the Sharpe optimal portfolio.

In Diagram 2 we show the Omegas for asset A, asset B and the Sharpe optimal portfolio for returns below the level of 5. The Sharpe optimal portfolio clearly dominates both asset A and asset B over this range.

Diagram 3 shows that for returns levels above about 5.4, the Sharpe optimal portfolio

THE FINANCE DEVELOPMENT CENTRE LIMITED

has a lower value of Omega than a portfolio consisting only of asset A. Diagram 4 shows that the Sharpe optimal portfolio also has a lower value of Omega than asset B for returns above 7.8.

For returns above 5.4 and below 10 we obtain a higher value of Omega by holding asset A. For returns above 10, holding only asset A continues to be preferable to the Sharpe optimal portfolio but holding only asset B is preferable to both these options, as one sees in Diagram 5. It is apparent that in this example the Sharpe optimal portfolio is sacrificing a considerable amount of the available upside. Fully 69% of the returns from the Sharpe optimal portfolio and from asset A are above 5.4. Over 55% of the returns from asset B are above this level.




Diagram 2. The Omegas for asset A , asset B and the Sharpe optimal portfolio as functions of the return level r .

The crossing in the Omegas for asset A and the Sharpe optimal portfolio identifies the point at which a portfolio consisting of 100% asset A has better risk-reward characteristics than the Sharpe optimal portfolio. Holding asset A, together with a put option with a strike at the return level of this crossing is an obvious strategy for a risk averse investor. Strategies for any risk preference may be obtained by optimising Omega over the appropriate range of returns, as is illustrated in Diagram 6.

THE FINANCE DEVELOPMENT CENTRE LIMITED




Diagram 3. Above 5.4 the Sharpe optimal portfolio has a lower value of Omega than asset A.




Diagram 4. The Sharpe optimal portfolio has a lower value of Omega than either asset A or asset B for returns above 7.8.

THE FINANCE DEVELOPMENT CENTRE LIMITED




Diagram 5. The Omega optimal portfolio is 100% asset B for returns above 10.





Diagram 6. The Omega optimal portfolio (black) for returns between 4.4 and 5.4 is 80% asset A, 20% asset B.

THE FINANCE DEVELOPMENT CENTRE LIMITED

Ω Notes: Does negative skew and higher than normal kurtosis mean more risk?

This is the distribution formed by combining 3 normal distributions with means of 0, 78.5 and -76 respectively and standard deviations of 11.2, 20.8 and 20.8. Their weights in the combination are 62%, 7% and 31%. The mean and standard deviation of the resulting distribution are -18 and 46. It has skew of -0.05 and kurtosis of 3.17 – about 6% higher than a normal distribution's kurtosis of 3. These are both widely regarded as signs of higher than normal risk.

We show the distribution (A) and a normal distribution with the same mean and variance (B).




In spite of the indications from skew and kurtosis, it is the normal distribution which has the heavier tails on both the up and downsides. A $2\text{-}\sigma$ gain is almost 1.8 times as likely from distribution A as from the normal however a $3\text{-}\sigma$ gain is only 0.85 times as likely. At the $4\text{-}\sigma$ level the gain is 35 times more likely from the normal.

The downside is more alarming. The probability of a $1\text{-}\sigma$ loss is about 1.4 times higher than for the normal however a $2\text{-}\sigma$ loss is only 0.7 times as likely. The normal is almost 80 times as likely to produce a $3\text{-}\sigma$ loss and over 100,000 times more likely to produce a $4\text{-}\sigma$ loss.


Both the large loss and large gain regimes are produced by moments of order 5 and higher which dominate the effects of skew and kurtosis. It is not possible to estimate moments of such orders from real financial data.

The Omegas for these two distributions, capture this information completely, with no need to compute moments of any order. The crossings indicate a change of preference.

THE FINANCE DEVELOPMENT CENTRE LIMITED




This is the leftmost crossing of the Omegas. The combination has less downside and a higher value of Omega.



This is the rightmost crossing of the Omegas. The normal has more upside.

Ω Notes: How many moments do you need to describe tail behaviour?

This is a distribution formed by combining three normal distributions whose means are -5 , 0 and 5 with standard deviations of 0.5 , 6.5 and 0.5 respectively. The respective weights are 25% , 50% and 25% . It is shown below with a normal distribution with the same mean (0) and variance (5.8)




The distribution A with a normal B of the same mean and variance.

The kurtosis of this distribution is 2.65 or about 88% of the normal value of 3 . Although lower kurtosis is often regarded as indicating lower risk, this distribution has heavier tails than a normal with the same mean and variance.


The distribution is symmetric so the odd moments are, like those of the normal, all zero. The 6^{th} moment is identical with that of the normal to within 2 parts in 1,000 and the eighth moment only differs from that of the normal by 24% . It is only with the tenth moment that a more substantial deviation from the normal appears. The tenth moment is 55% greater for the portfolio than for the normal.

The dominant effects producing the heavy tail behaviour therefore come from moments of 8 , 10 and higher. These simply cannot be estimated from real data.

THE FINANCE DEVELOPMENT CENTRE LIMITED



The Omegas for distributions A and B around their common mean of 0. Crossings in Omegas indicate a change in preference.



The large loss regime and the left-most preference change (2σ is 11.6). Distribution A has a lower Omega to the left of -12.2 as it has the higher catastrophic loss potential. Distribution A has almost 3 times the likelihood of a $4\text{-}\sigma$ loss or gain than the normal with the same mean and variance.

ⁱ Con Keating and William F. Shadwick, *A Universal Performance Measure* © The Finance Development Centre 2002.

ⁱⁱ Ana Cascon, Con Keating and William F. Shadwick, *The Mathematics of the Omega Measure* © The Finance Development Centre 2002.

An Introduction to Omega © The Finance Development Centre, February 2002