WILLIAM SHADWICK
Managing Director, Finance Development Centre, London, England

# From Alpha to Omega

## The new Omega function looks at all the distribution information.

**While it is a fact that returns from financial** instruments are not normally distributed, the standard analytic tools for investment portfolios are based simply on mean and variance.

Fat tails contain vital information about risk, for which variance is a poor proxy. In hedge funds, skewness and other tail effects normally dominate the information in their return variance. Portfolio allocations based on mean and variance can produce lower terminal values than allocations using information in the entire returns distribution.

The Omega function for a returns distribution is a new tool for financial analysis using all the information in the distribution. Comparing the Omega functions for two or more assets over a range of returns ranks their performance and risk profiles without estimating any moments. The evolution of a manager's Omega function over time provides a complete picture of performance and risk. Omega functions reveal information invisible to mean/variance measures and can lead to significant improvements in portfolio values.

The construction of the Omega function can be motivated by considering the quality of a bet on a return above a given level *r*, which we regard as a loss threshold. To do this we need to know how much we will win if we win and how much we will lose if we lose. We also need to know the probability of a win and a loss.

If F($x$) is the cumulative density of returns defined over the interval from *a* to *b*, then, by considering the sum of probability weighted gains and losses as the unit of gain and loss shrinks to zero, we are led to the ratio

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

as our measure of the quality of a bet on a return higher than *r*. The higher this value, the better the quality of our bet. We obtain the Omega function of the distribution *F* by considering all possible returns *r* between *a* and *b*.

The Omega function turns out to be mathematically equivalent to the distribution itself—so it contains exactly the same information and all of the information in tails. The derivation of Omega shows that a higher value is preferable to a lower one.

We apply Omega to two similar hedge funds from a YMG fund of funds. Fund A dominates Fund B across

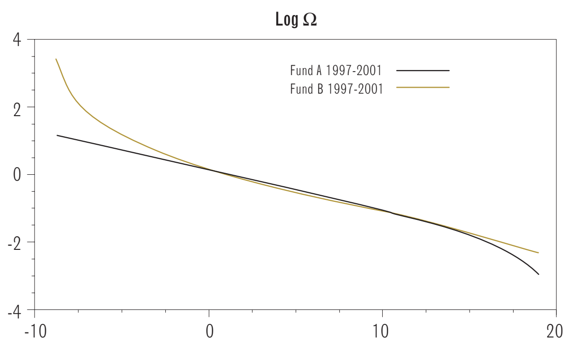


almost the entire range of monthly returns from 1997 to 2001 (which runs from -17 per cent to 23 per cent). Allocating $1 between these funds to maximize the Sharpe ratio of the portfolio would have produced a terminal value of $1.35 at the end of June 2002. Following an Omega allocation strategy would have produced a terminal value of $1.70 instead. The actual allocations were equal with a terminal value of $1.55.

The reason for the superiority of the Omega allocation is clear from the diagram. The flatter the Omega function, the fatter the tails of the distribution. Fund A has a fat tail on the downside relative to Fund B. Fund B has a fat tail on the upside relative to Fund A. A rational investor (independent of utility function) will prefer Fund B to Fund A. The allocation which optimizes the Sharpe ratio, using only mean and variance, does exactly the opposite, putting 61 per cent of the portfolio into Fund A in July 2001. ▮