OMEGA

A New Tool for Financial Analysis

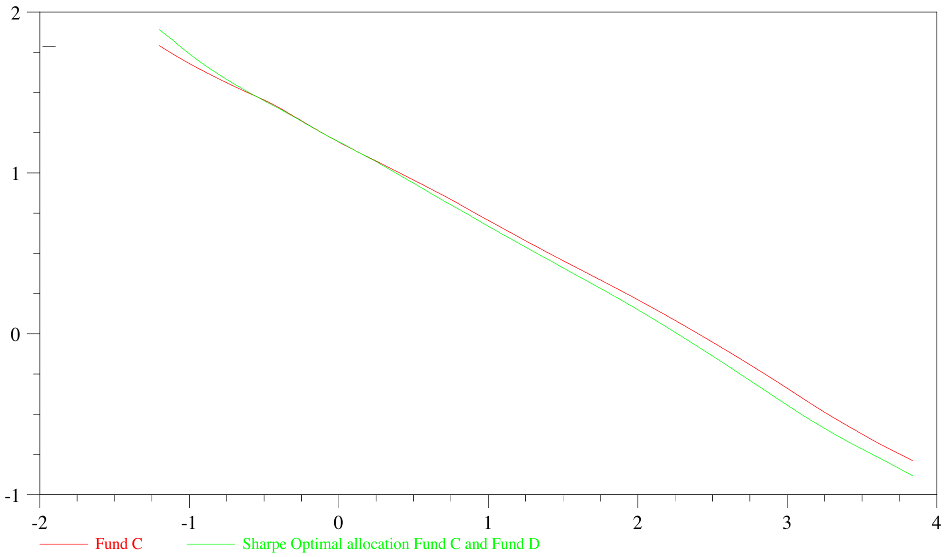

Fund C is a better bet than the Sharpe optimal combination of Fund C and Fund D for more than 70% of the observed range of returns. The result of reducing variance in this case is lower terminal values.

A New Tool for Financial Analysis

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 which use the 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. But by itself this is not enough. 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 is mathematically equivalent to the distribution itself–so it contains exactly the same information and, in particular, all of the information in tails. It is easy to interpret Omega curves; at any return level a higher value is preferable to a lower one and the flatter the curve the higher the risk.

We can see the impact of this by examining the returns from two distributions which have the same mean but are otherwise mirror images of each other: one exhibits negative skew, the other positive.

This is an example of the choice of returns from buying a lottery ticket (skewed heavily to the right by the small chance of a very large return) and selling a lottery ticket (skewed heavily to the left by the small chance of having to pay out a large win).

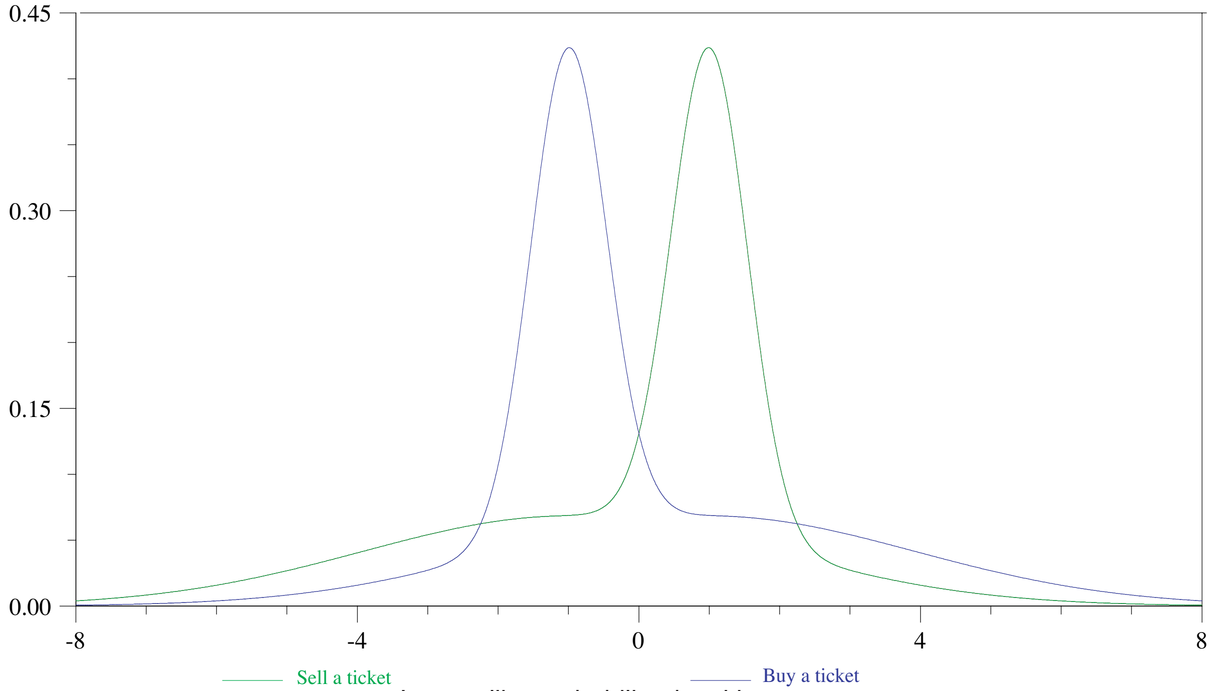

Lottery like probability densities

An individual faced with the choice of buying a ticket for $1 with a one in a million chance of winning a million dollars and the choice of selling the same ticket will invariably chose to buy rather than to sell. This choice cannot be justified on the basis of mean and variance because these are the same for both bets. The Omegas for this example are shown below. The only place at which they agree is the mean; for any other return level, the Omega for buying a ticket dominates the Omega for selling it.

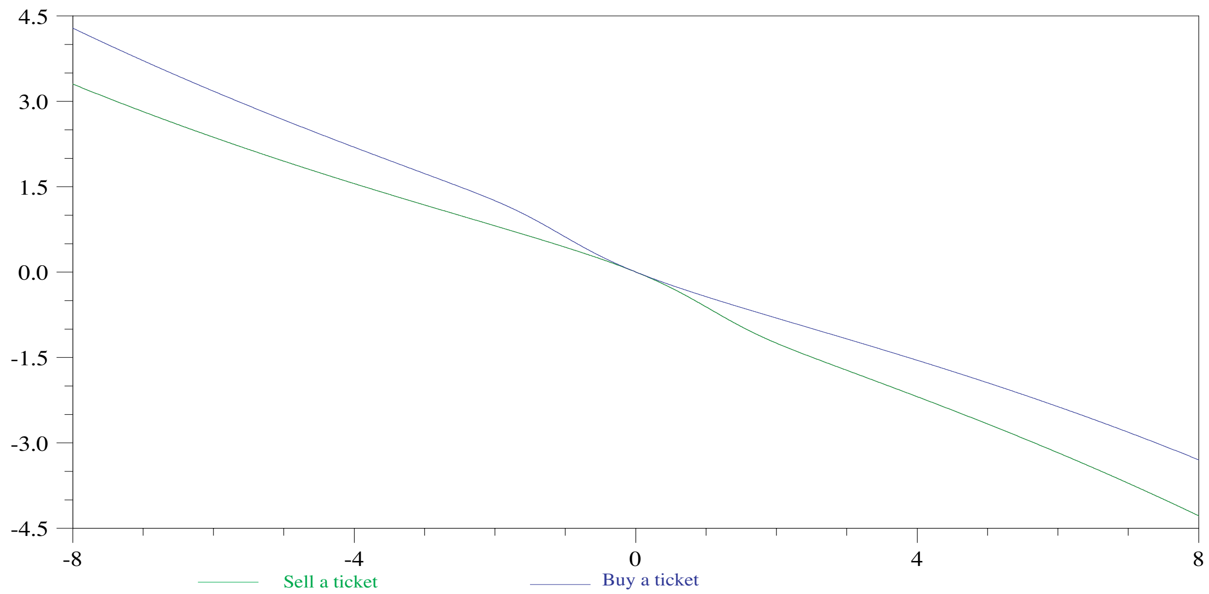

Log(Omega) for the lotteries

Portfolio Construction

While common sense is enough to distinguish between the choices of buying or selling a single lottery ticket, the construction of an investment portfolio, such as a fund of hedge funds, is less straightforward.

Because Omega incorporates all available information about the returns distribution, allocations made according to Omega preference take into account the effects of such features as skewness and fat tails which often outweigh the contribution of variance in hedge fund returns. This can produce dramatically higher terminal values than those obtained by optimising Sharpe ratios or other mean/ variance measures.

Omega analysis shows the limitations of using variance as a proxy for risk. Examples from hedge funds show that this can lead to combinations of assets which insure against large losses at a cost of abandoning the majority of upside potential.

Simulation

The analysis of Omegas for returns series lends itself naturally to simulations. The preferences obtained from Omega functions are easily interpreted,statistically significant and robust. By making use of bootstrap re-sampling, for example, one can easily evaluate the stability of predictions based on an observed set of returns data. Three to four years of monthly data can produce Omega functions which provide stable preferences between assets.

Scenario Analysis

Scenario analysis and stress testing are also straightforward to conduct and evaluate. The Omega functions provide easily interpreted analysis of the returns with the impact of historic market events added or with extreme gains or losses omitted. This allows a comprehensive assessment of the of theimpact of deviations from normality as well as an indication of the extent to which mean and variance or other moments capture the risk in an investment portfolio.

Higher Sharpe ratios can produce lower terminal values.

We illustrate the impact of Omega on portfolio construction with an example of allocations between a pair of hedge funds of similar style held within a fund of funds. The analysis is based on five years of monthly returns.

Fund A and Fund B were held in equal weights in the overall fund, with annual re-balancing. The terminal value of $1 invested in this way after 5 years was $1.55. If instead of maintaining the initial weights the annual re-balancing had been done to optimise the Sharpe ratio for the combination of Fund A and Fund B, the terminal value would have been reduced to $1.35.

By the end of the first four years, the Omegas for these funds show a clear dominance of Fund B over Fund A across most of the range of returns (which run from –17% to +23%). Switching from the equal allocation to 100% Fund A for the fifth year raises the terminal value to $1.70. The reason for the success of this strategy is simple. Fund A has a fat tail on the downside relative to Fund B and Fund B has 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 optimises the Sharpe ratio, using only mean and variance, does exactly the opposite, putting 61% of the portfolio into Fund A in July 2001.

After the fifth year, the Omegas for Fund A and Fund B are a close approximation to those for buying and selling a lottery ticket–the dominance of Fund A over Fund B is clear across the entire range of observed returns. This is illustrated in the diagram below.

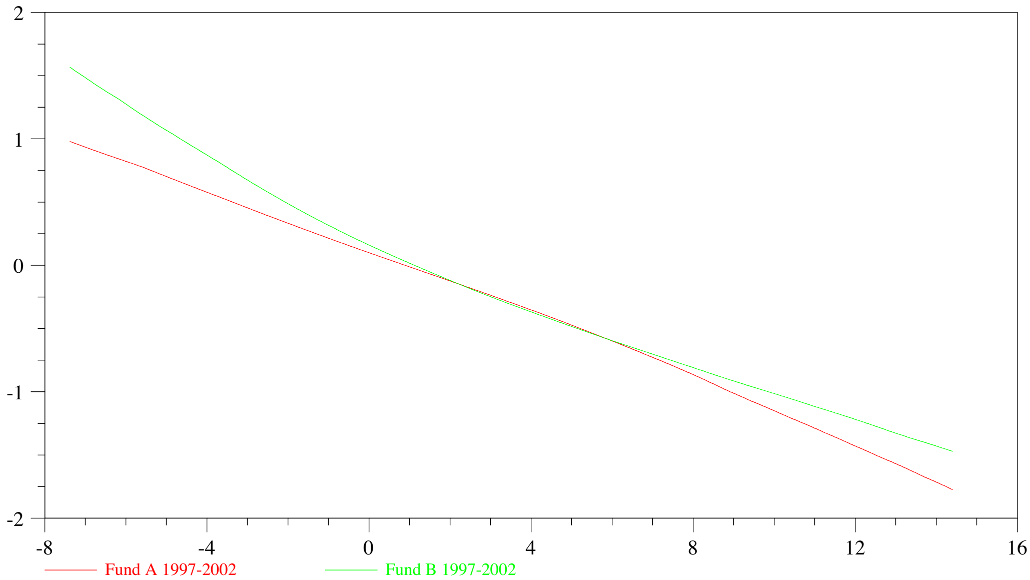

Nevertheless, based on the returns for the first five years, the combination of Fund A and Fund B that produces the optimum Sharpe ratio is 20% Fund A and 80% Fund B. As a result, one would expect the terminal value of the Sharpe optimal allocation to be substantially lower than that of the Omega allocation (100% Fund B) continued in future years.

Simulations of a further year of returns by bootstrap re-sampling from the observed returns distribution confirms this with an average six year terminal value for the Sharpe optimal allocation of $1.53 compared with $1.95 for the Omega allocation.

Reducing variance can be expensive

The motivation for the use of Sharpe ratios is one of indisputable value: to obtain the combination of assets which produces the highest possible return with the least possible risk. The use of variance as a proxy for risk, however, may reduce the benefits of this approach to an unexpected degree.

We illustrate this with a pair of hedge funds of similar style which were held in a fund of funds in equal weights. The Omega functions for three years of monthly returns for Fund C and Fund D make apparent a preference for Fund C, independent of return threshold. At this stage, the Omega analysis and mean/variance analysis reach the same conclusion: the best allocation is 100% FundC.

By the end of the fifth year however, changes to mean, variance and covariance move the optimal Sharpe ratio allocation down to 86% Fund C and 14% Fund D. The Omega functions, which incorporate all of the higher moments as well as the first two, show that Fund C still dominates Fund D by a wide margin across all return levels.

To understand the impact of this additional information we can compare the Omega function of Fund C with that of the Sharpe optimal combination of Fund C and Fund D. The latter reduces the standard deviation from 2.16 to 2.05 (and also reduces the mean return from 2.4 to 2.26).

The diagram below shows how costly it is to use variance as a proxy for risk: the portfolio with the optimal Sharpe ratio has an Omega which falls below that of Fund C more than one standard deviation below the portfolio mean return.

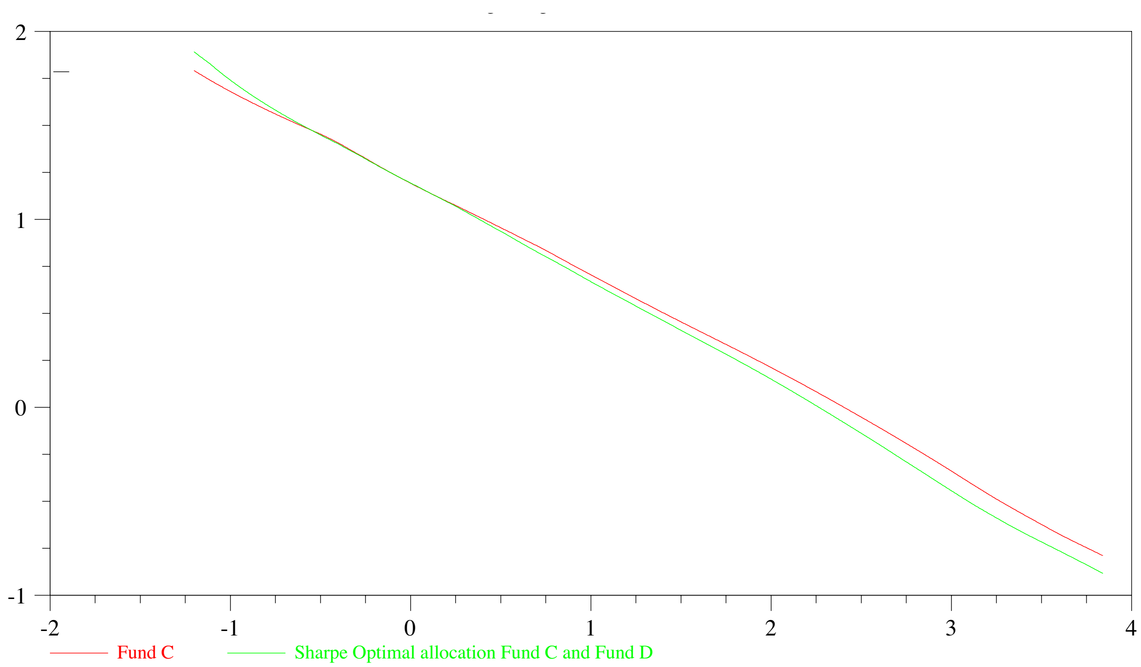

Fund C is a better bet than the Sharpe optimal combination of Fund C and Fund D for more than 70% of the observed range of returns. The result of reducing variance in this case is lower terminal values.

In fact, for any loss threshold above a return of 0.12% per month, holding Fund C is a better bet than holding the Sharpe Optimal portfolio. For a loss threshold below this level, the Sharpe Optimal portfolio is the better bet. But the support of the observed monthly returns distribution for the portfolio runs from –3.86% to 9.53% , so a fund manager who uses the optimal Sharpe ratio to reduce the standard deviation from 2.16 to 2.05 will have paid for this by reducing his upside potential over more than 70% of the observed returns.

Simulations of 10,000 sixth years from bootstrap re-sampling of the five year data sets reveal the cost of this strategy in lower terminal values. The average terminal value of $1 invested in Fund C for one additional year is $1.33 compared to only $1.14 for Fund D. As the Omega functions predict, the lowest terminal value (65 cents) was produced by Fund D while the lowest terminal value from Fund C was over $1.00.

Omega Risk Profiles

Funds CL and BH are hedge funds of the same investment style, contained within a fund of funds. Their Omegas, from monthly data to June 30 2002, reveal a global performance advantage for Fund CL over Fund BH .

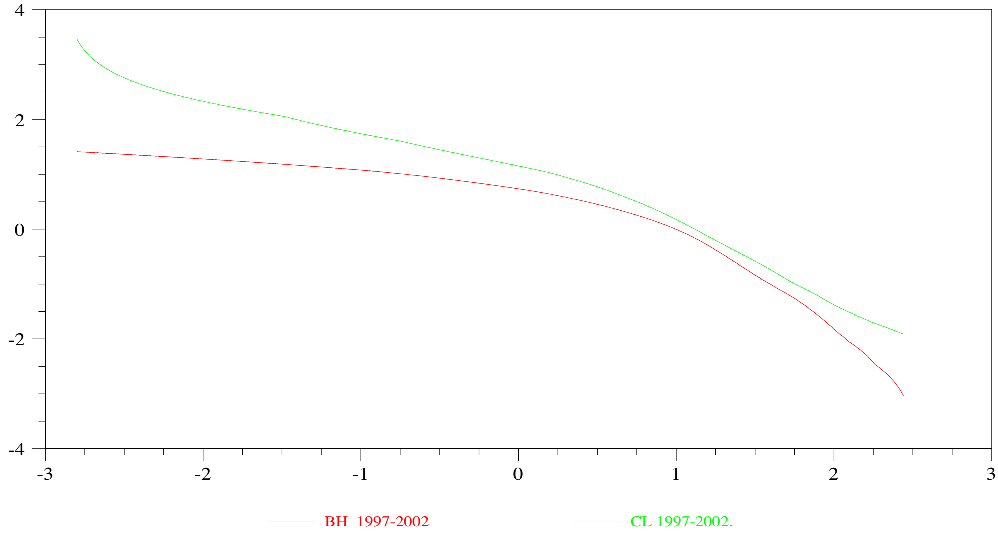

The slopes of the log(Omega) curves reveal the risk run by the managers of Funds BH and CL. The more negative the slope, the lower the risk. A manager whose curve is steep on the downside is producing lower downside risk. On the upside, a flatter curve provides more risk of outperformance. By this measure, Fund CL again dominates Fund BH. The BH manager is running large downside risk with very low prospect of achieving large returns.

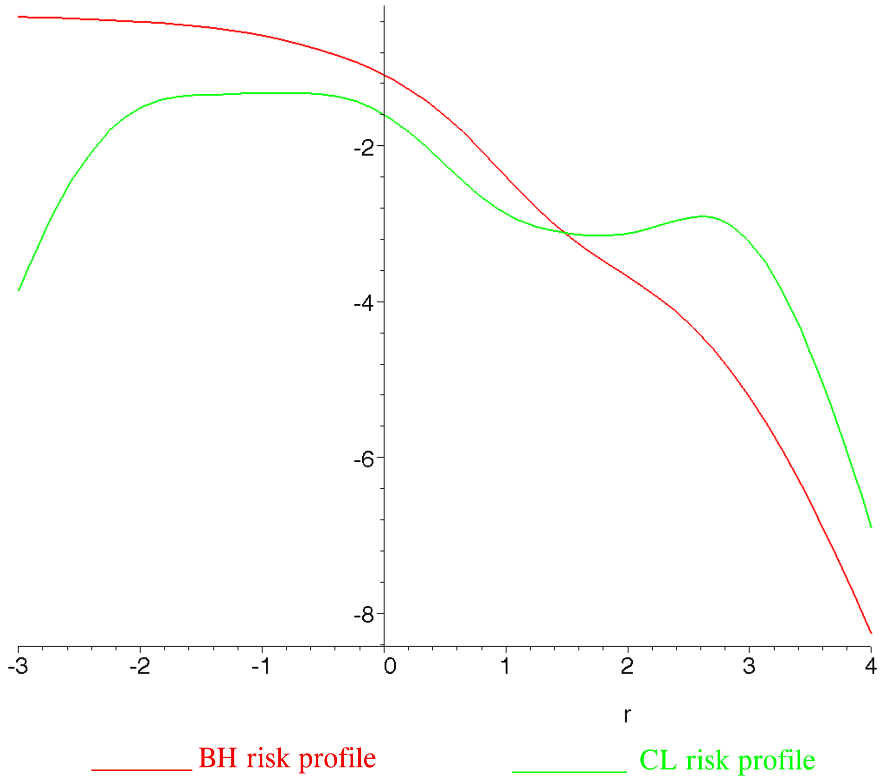

We can assess the effect of these risk profiles on terminal values by simulating additional one year periods by bootstrap sampling from the observed 5 year returns data to June 2002. The resulting distributions of terminal values for an investment of $1 for 10,000 simulated 12 month periods are shown below for BH in red and CL in green.

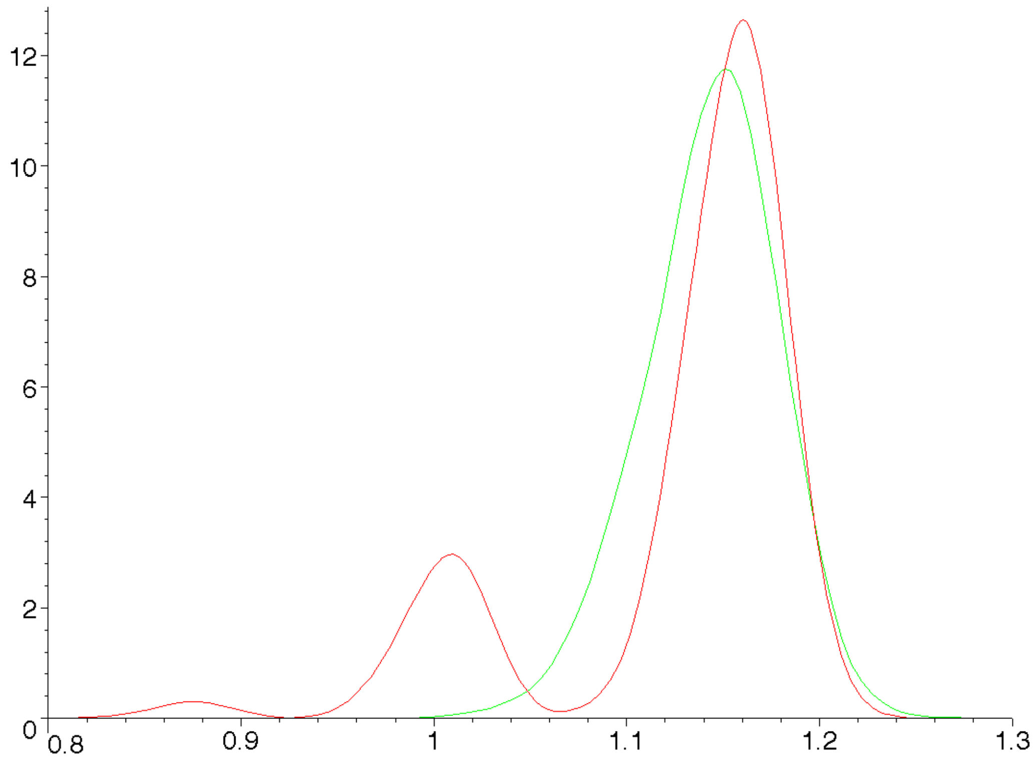

Distributions of terminal values for 10,000 simulated years. BH in red, CL in green.

The means are almost identical, at 1.13 for BH and 1.14 for CL. The difference is in the potential for large loss which the BH distribution of terminal values presents. Over 7.5% of the returns represent losses–two hundred and fifty times the frequency experienced with CL. The downside risk carried by BH is an example of the proverbial process of picking up pennies on a train track. In the fullness of time a large loss will be experienced. In the case of BH, this happened in September 2002 when the fund lost 53% and closed.

How much risk is hidden in the tails of a returns distribution?

The Omega curves for two pension fund managers (long only equity portfolios) are shown below together with their benchmark for three years ending in December 2002. As one would expect all three show significant losses, however both managers outperformed their benchmark both in average return and in terminal value. Manager 2 had the poorest risk profile of a group of 4 managers with the same mandate. He was the only one who ran higher downside risk than the benchmark. In spite of this high level of risk taking, he had the highest Sharpe ratio of the group. The additional risk would not only go unnoticed in conventional performance measures, it would be rewarded rather than penalised.

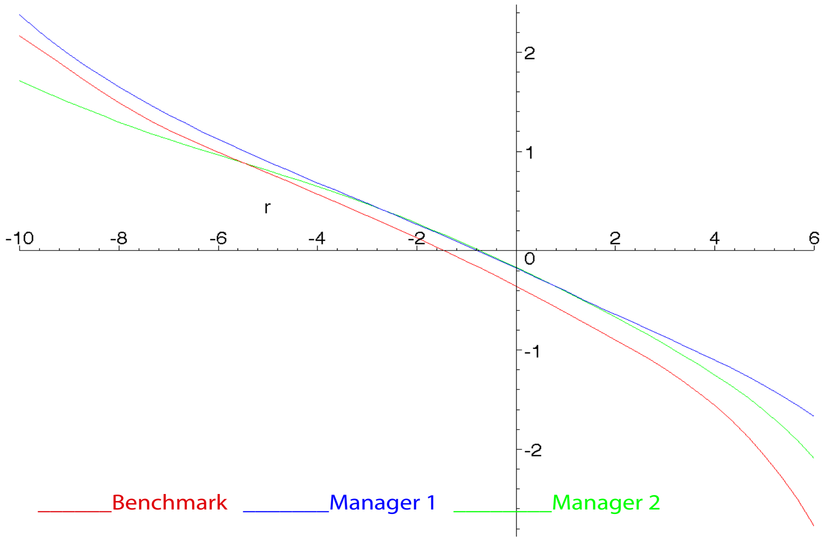

The next figure shows the risk profiles for the two managers together with the constant risk level implied by the normal approximation. This serves reasonably well only close to the mean return. It significantly understates the downside risk presented by manager 1 as well as overstating his upside potential. Both effects will prove costly for the plan sponsor if they go undetected.

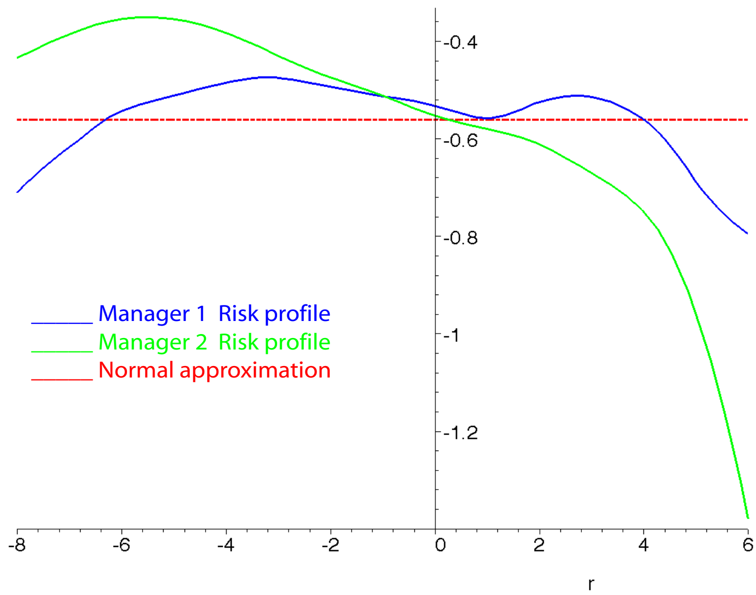

The next figure shows the Omega curves for two Funds of Hedge Funds, which are intended to be aggressive and conservative. As would be expected, the aggressive fund has a flatter Omega than the conservative. However the up and downside differences are noteworthy (as is the fact that the aggressive fund of funds returned less on average over the 3 year data set).

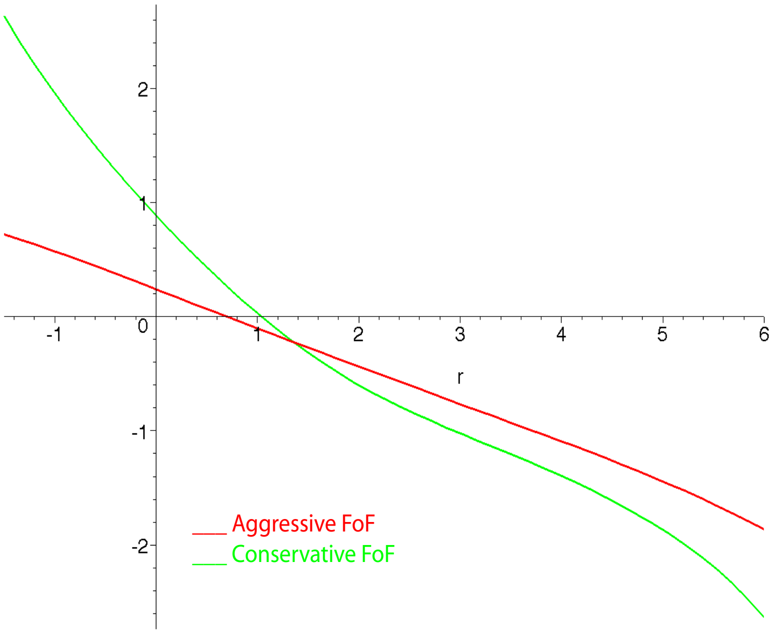

The figure shows that the additional downside risk of the aggressive fund relative to the conservative one is not matched by upside potential. (If it were, this diagram, which is on a log scale, would be symmetric around the mean). Simulations of additional 12 month periods confirm this effect. The aggressive fund is 200 times more likely to produce a loss than is the conservative fund. Its potential for large loss, relative to the conservative fund, is not offset by its increased potential for large gains.

The aggressive fund of funds has improved over the last year of returns as the figure below shows. The aggressive fund of funds has made improvements in its risk profile over the third year, perhaps indicating managers' adjustment to the bear market.

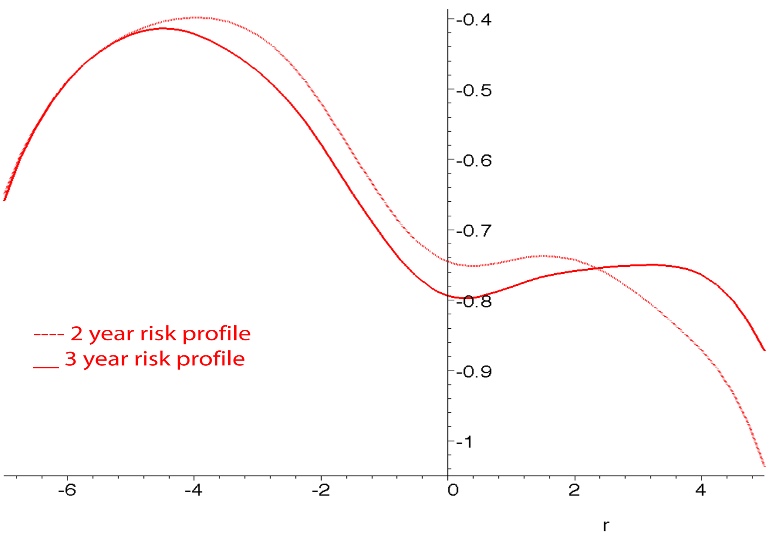

During the same period the conservative fund of funds has also improved its risk profile, so it appears likely that the conservative fund will continue to be the better bet.

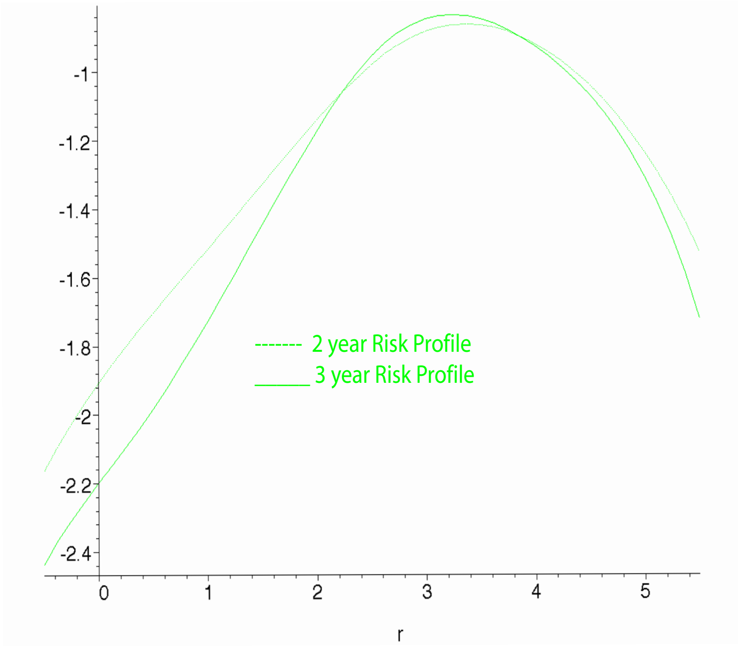

The next figure shows the extent to which a manager who is attempting to produce normally distributed returns, has succeeded. The fund strategy aims to control the first four moments and, by this measure is producing returns which are “close” to normal. The Omega risk profile shows that this metric is hiding both additional downside risk and even greater upside potential. This information is important both to the fund manager and his investors. With Omega profiles it becomes obvious at a level which no moment expansion can match.

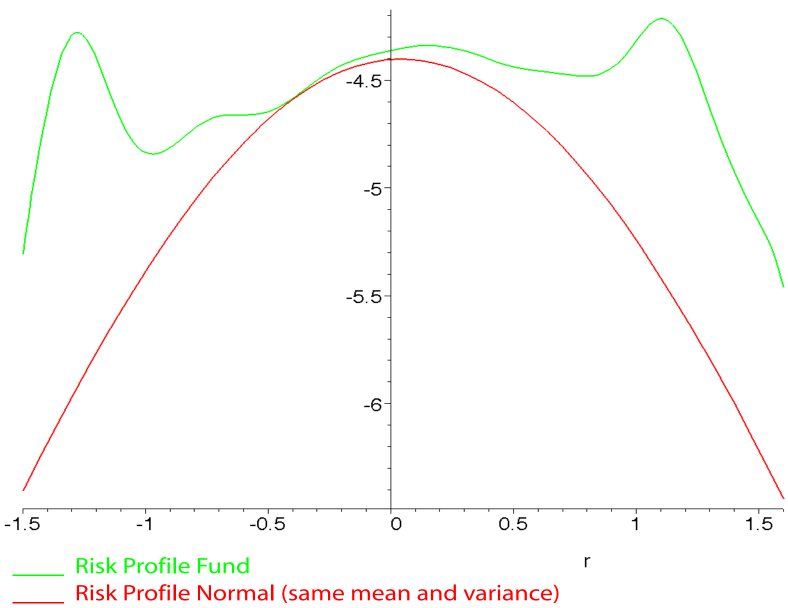

Omega Portfolio Optimisation

In this note we show how our simplest Omega based metric produces portfolios which are superior both to those created with conventional mean variance optimisation and to those produced by optimisation over the first four moments, mean, variance, skew and kurtosis. (The example uses monthly data from two hedge funds of different style for 5 years ending January 2003.)

In using Omega to produce optimal portfolios it is essential to create metrics which, like Omega itself, take all of the data in the returns distribution into account. In particular, any method which is based on the value of the Omega function at a particular threshold will produce 'pessimal' rather than optimal portfolios generically.

The reason for this can be seen clearly in Figure 1. Fund T and fund G have very different risk profiles: fund T has a much fatter tail on the downside while fund G has a fatter upside tail. The mean returns are similar with 0.92% per month for G and 0.97% per month for T. The value of the Omega function for T is greater than that of the Omega function for G for any return level between 0 and 1% per month, so any optimisation based on the value of Omega at any point in this interval will favour T over G. This is, of course, the wrong conclusion as the global picture makes it quite clear that G is superior to T over most of the observed range of returns.

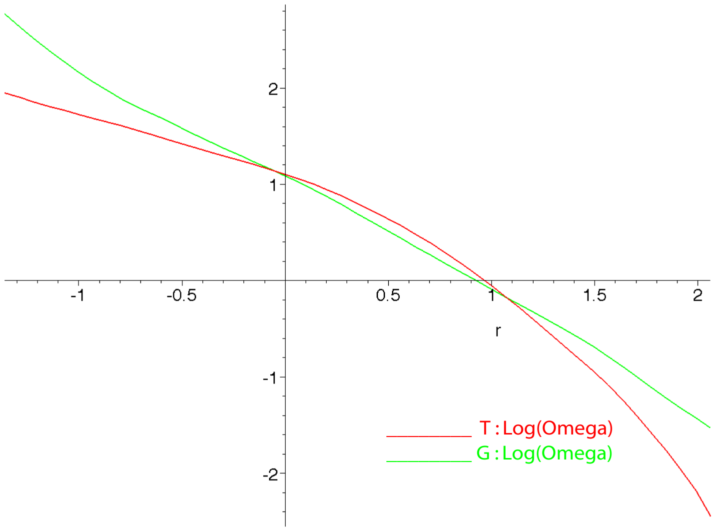

Figure 1 Log(Omega) for fund T and fund G.

Because there is not a strict dominance of G over T however, there *is* some benefit to diversification between these funds. In terms of the distributions, we may give up some of the upside apparent in G in order to produce a downside which is superior to that provided by either of T or G separately.

The conventional approach to combining the two funds is to make the allocation with the highest Sharpe ratio. By this measure, the optimal mix is 60% fund T and 40% fund G. The distribution which this produces has, by definition, the best return to standard deviation ratio of all possible portfolios of the two funds. As there is only a relatively small difference between the mean returns of the two funds the main effect is due to the reduction in variance to 0.45 for the Sharpe optimal portfolio compared with 0.65 and 0.83 for T and G respectively.

The Sharpe optimal allocation ignores all information aside from mean and variance, so in particular ignores the significant tail properties of funds G and T. A more sophisticated approach takes into account the third and fourth moments to produce a combination of T and G which is approximately normal. This 4-moment optimisation leads to a reduction in the allocation of T to approximately 50%.

While this approach is a substantial improvement on the Sharpe optimal allocation, it still fails to capture all of the available diversification benefits. To obtain these, we apply the simplest of our proprietary metrics to produce an Omega optimal portfolio. This is achieved with an allocation of 40% Fund T and 60% Fund G.

Figure 2 shows the extent of the improvement over the Sharpe optimal Omega which the 4 moment optimisation provides, together with the further gains which are achieved by moving the T allocation down to the Omega optimal level of 40%. The differences are relatively minor over returns from 0 to 1% but are dramatic for negative returns or for returns above 1.5% per month.

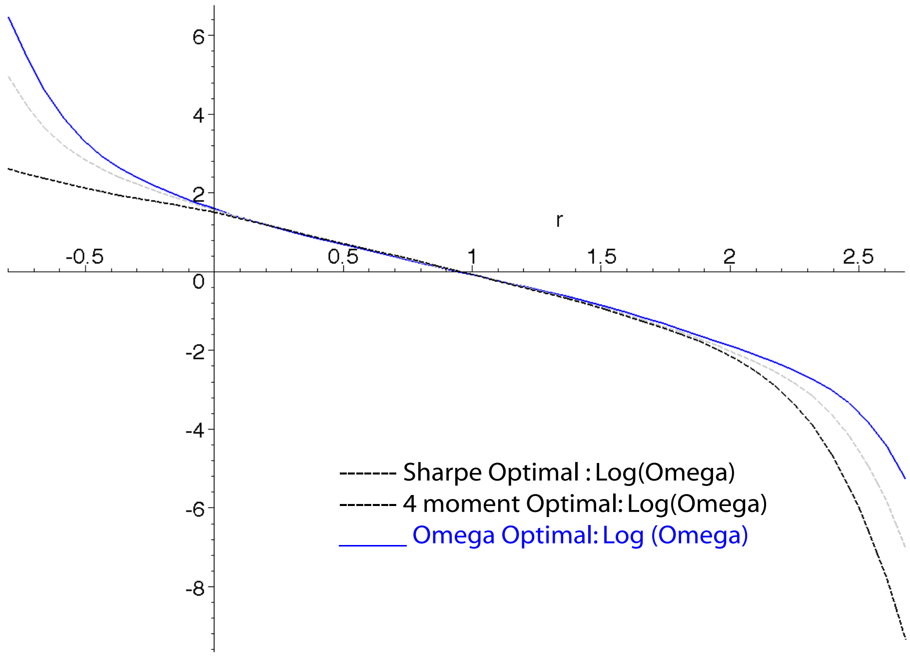

Figure 2 Log(Omega) for the Sharpe Optimal allocation of 60% T, the 4-moment optimal allocation of 50% T and the Omega optima allocation of 40% T.

In particular, as Figure 2 shows, adjusting for skew and kurtosis provides significantly more upside risk and less downside risk than the Sharpe optimal weighting, while lowering the proportion of T to the Omega optimal level provides an additional improvement of the same magnitude. The true diversification benefits available from combining these two funds are obscured by conventional mean-variance analysis and only partly revealed by the adjustment for third and fourth moments. By contrast, the benefits are immediate in terms of the Omega functions of the various portfolios.

There is, of course, no free lunch. Figure 3 shows the Omega optimal allocation of T and G against the Omega function for fund G. The improvement in the downside has been purchased by giving up some of the upside which fund G provides by itself. Note however that return level at which the Omega optimal allocation begins to trade off up for downside advantage is *above* the mean return of Fund G, as the close-up view provided by Figure 4 shows.

Finally, we note that in the 4 months of out of sample data to the end of May 2003, fund G outperformed fund T by over 11.5 %. Thus, in this period, both the Sharpe and 4-moment optimal combinations proved inferior to the Omega optimal allocation. This is indicative of the impact which Omega optimisation can have on terminal values.

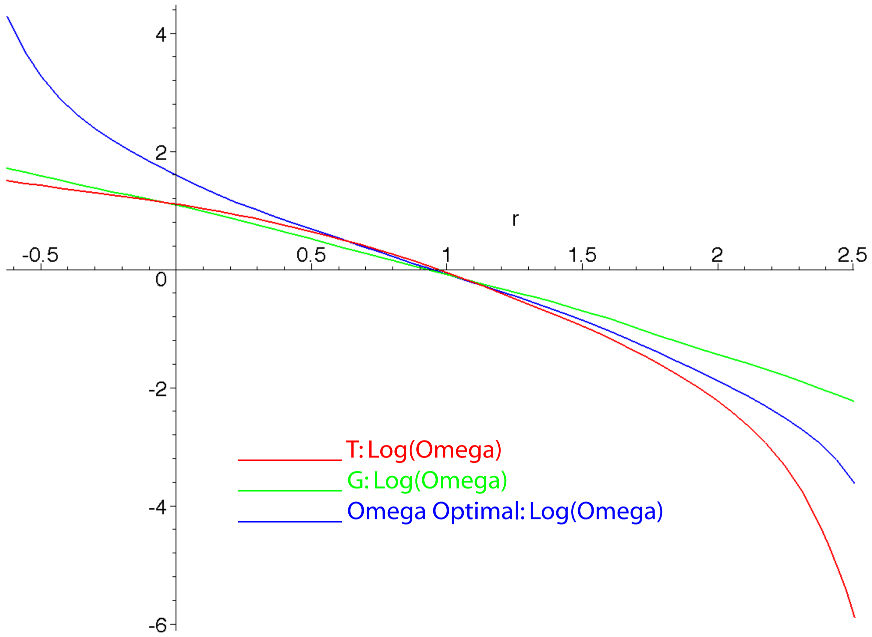

Figure 3 The Omega functions for the optimal combination of Fund T and Fund G together with those for Fund T and Fund G.

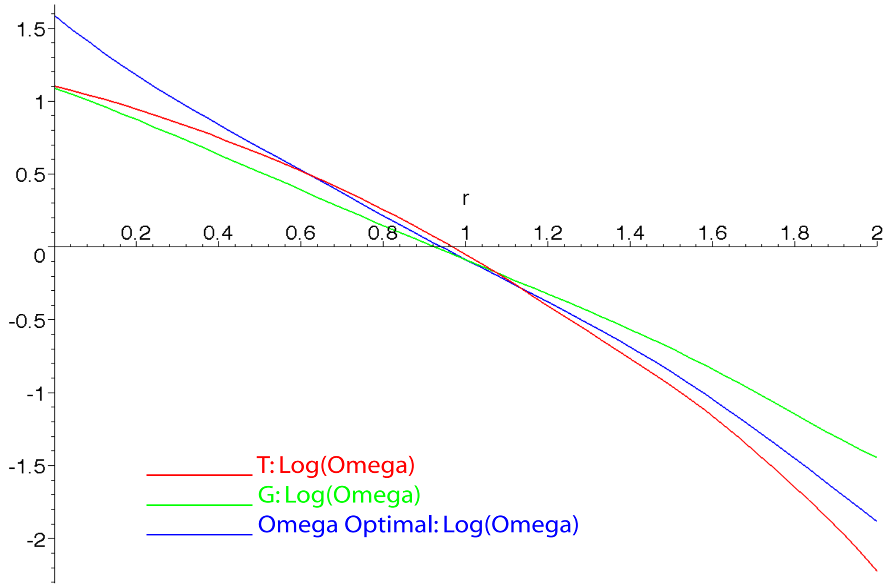

Figure 4 The Omega functions for the optimal combination of Fund T and Fund G together with those for Fund T and Fund G over returns from 0 to 2% per month. (The mean return has an Omega value of 1 so Log(Omega) is zero at the mean.)

TFDC's Omega Team

Dr. William Shadwick, Managing Director, is a prominent mathematician who was responsible for establishing the Fields Institute for Research in Mathematical Sciences before entering the finance industry in 1998. He was Assistant Director of Information Technology at Dresdner Kleinwort Benson where he was responsible for a number of projects involving technology transfer for derivatives pricing and hedging and for investment fund performance measurement. He was a member of the Steering Committee of the Financial Markets Group at the London School of Economics from 1998 to 2002. He has published key contributions to finance theory and practice and has made invited presentations of his work to leading industry groups and academic centres.

Dr. Ana Cascon, who was a tenured professor in the Applied Mathematics Department of the Federal University Fluminense, Rio de Janeiro, Brazil is a first rate mathematician who has held research and teaching positions in leading centres in France, Canada and Brazil, She is an expert in the use of symbolic computation techniques for mathematical modelling and is a co-author of TFDC's R&D papers on the extensions of Omega to multivariate distributions, portfolio optimisation and other proprietary trading trading applications.

Dr. Con Keating, a chemist and economist by training, has decades of top level market experience in the trading and research settings as well as in academic finance. He has held a number of high level consulting roles to investment banks, commercial banks, insurance and re-insurance companies, central banks and governments. He is a member of the Steering Committee of the Financial Econometrics Centre of the City University Business School. His econometric and trading experience provide a key component of our Omega R&D strategy.

Dr Bradley Shadwick, a physicist at the Lawrence Berekeley National Laboratory, is a one of the world's leading experts in numerical computation. He has extensive experience of applications of computational techniques in industrial settings and holds a patent for one of these. He is responsible for TFDC's information technology development and the software implementations of our proprietary computational methods for the use of Omega.