Omega Metrics

The 21st Century Standard for Performance and Risk Measurement

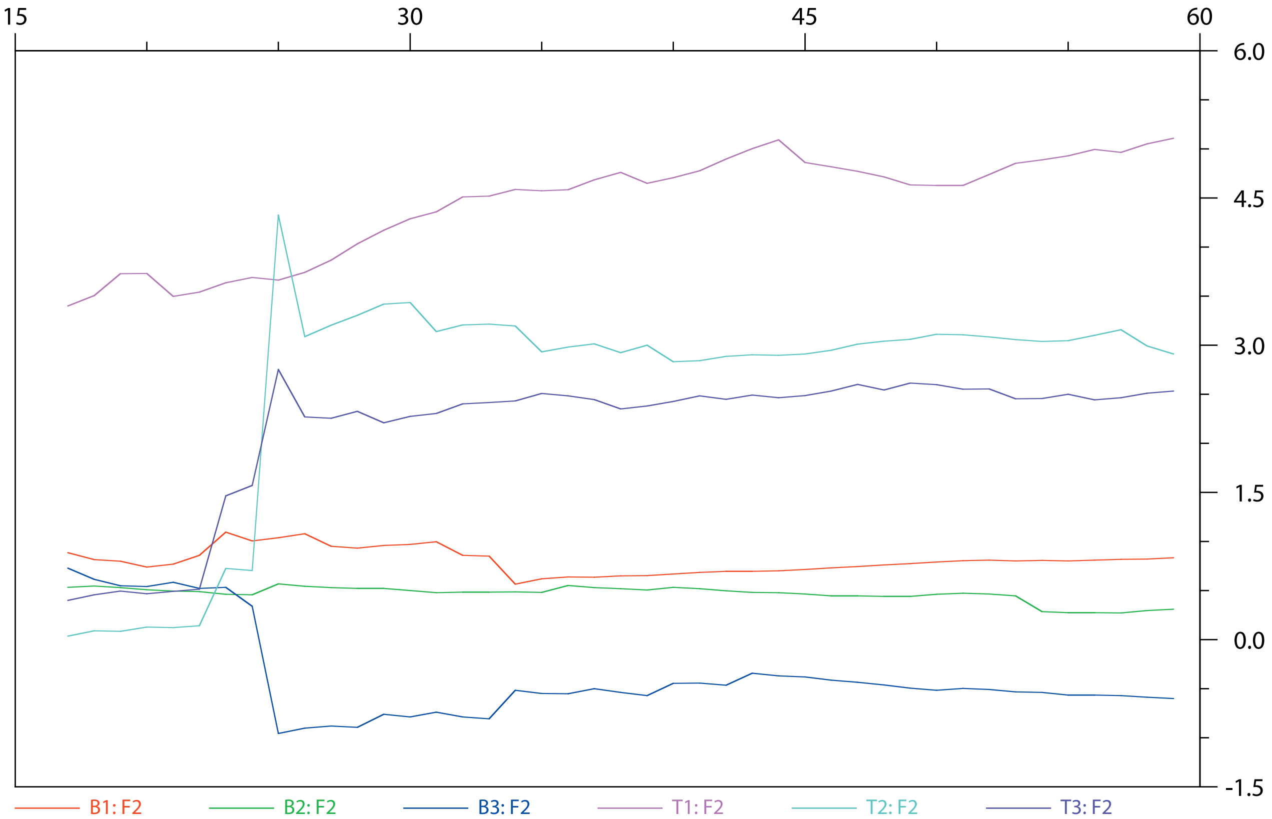

PMRA 2004

William.Shadwick@FinanceDevelopmentCentre.com

The Challenge of Asymmetric Returns

Hedge Funds, Funds of Hedge Funds (and even Hedge Fund Indices) have become standard for Institutional Investors, Pension Funds, and, increasingly, the retail investor.

These funds are *designed* to create asymmetric returns with upside dominating downside. (And can sometimes reverse this arrangement.)

As with any assets we must:

- assess the risk-reward trade offs
- make optimal allocations to portfolios
- re-balance or hedge portfolios effectively.

20th Century Technology

Markowitz Portfolio Optimisation (1953-56)

Sharpe Ratio Optimisation (1966)

Sortino Ratio (1991)

Ad hoc 3 and 4 moment methods (late 1990s)

The first two cannot cope with asymmetry unless it is insignificant in comparison with the effects of mean and variance.

This makes them suspect (and possibly dangerous) in any appliction to hedge funds (or even to corporate bond portfolios).

A Test Case for Handling Asymmetry: The Lottery

Buy Payoff: A loss of 90 cents 999,999 times in 1million
A gain of $999,999.10 one time in 1 million

Sell Payoff: A gain of 90 cents 999,999 times in 1million
A loss of $999,999.10 one time in 1 million

Why is this a test case?

- Common sense produces the ranking Buy over Sell (Ask the man in the street)
- Studies of Hedge Fund styles show that many Hedge Funds act as if they were selling out-of-the-money puts
- Hedge Funds are becoming an inescapable fact of life for investors from institutions to individuals.

If your optimiser or performance measure can't handle the lottery, you can't afford to trust it on assets with asymmetric returns.

William.Shadwick@FinanceDevelopmentCentre.com

How do the 20th Century Tools Stack Up?

Buy and Sell have the same mean (10 cents) and the same variance (999,999).

Buy and Sell are *indistinguishable* to mean/variance measures

The optimal mean/variance allocation is 50% of each.

This also reduces the skewness to zero and produces the minimal kurtosis, so 4 moment optimisers are no help.

This leaves the Sortino Ratio.

It correctly observes the heavy downside bias of the Sell and chooses Buy for MAR below the mean.

It fails the lottery test for MAR above the mean by reversing its preference and choosing Sell over Buy.

William.Shadwick@FinanceDevelopmentCentre.com

For MAR above the mean the Sortino ratio preference is for Sell over Buy.

This reverses the common-sense preference for Buy over Sell which it produces for MAR below the mean.

This is generic behaviour for Sortino ratios in Lotteries independent of mean and variance.

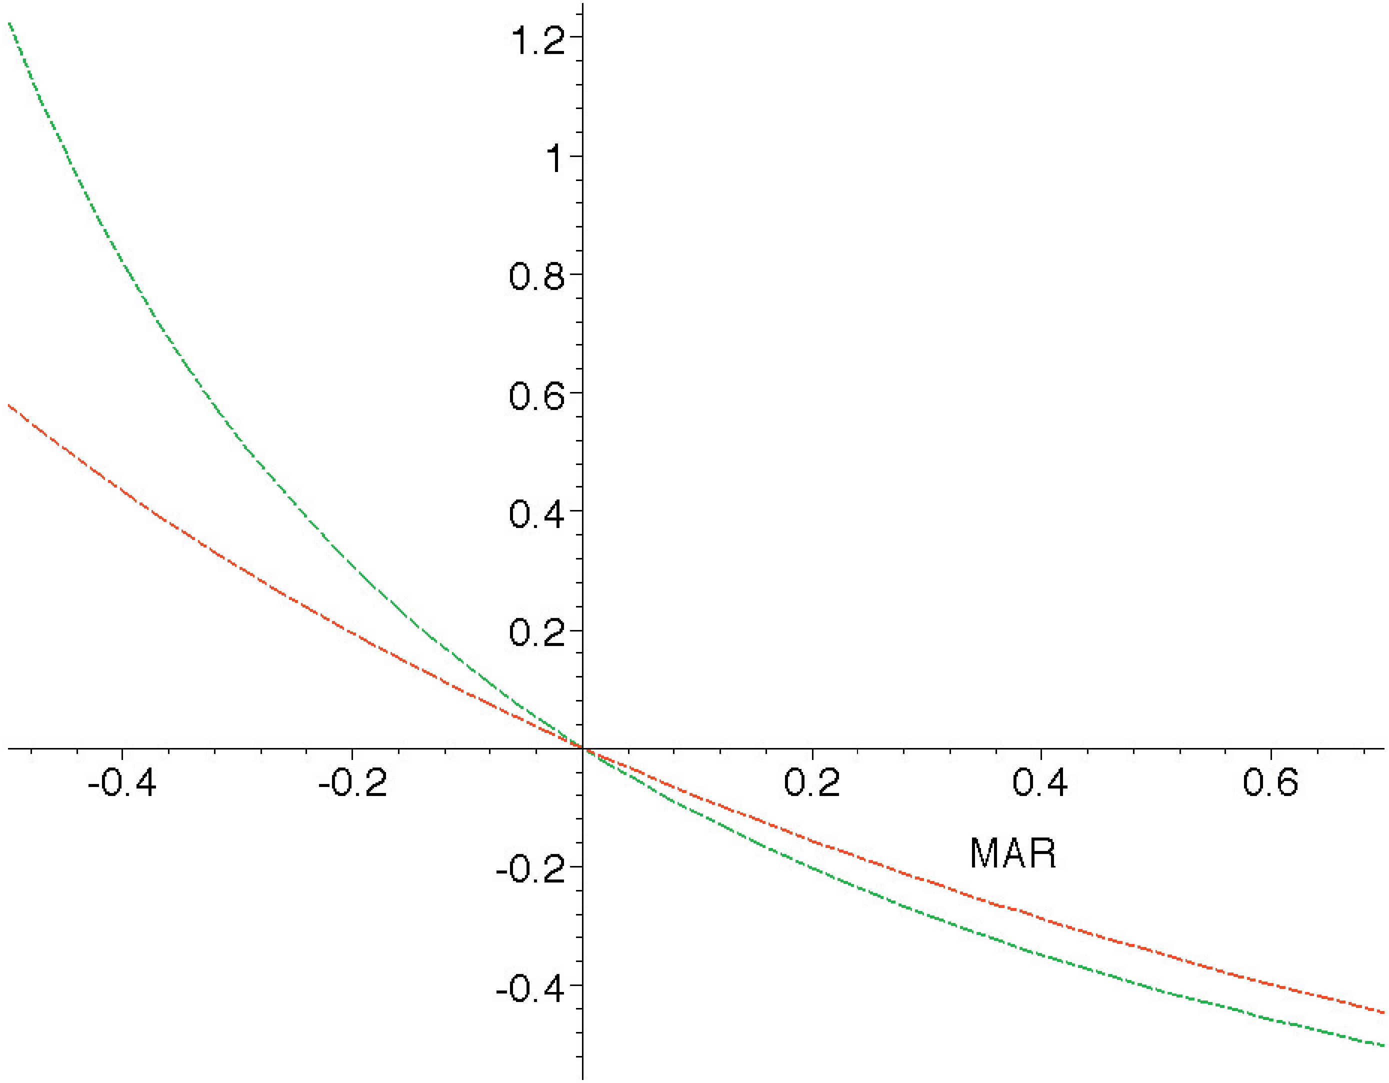

Sortino Ratio as a function of MAR for Buy (green) and Sell (red).

William.Shadwick@FinanceDevelopmentCentre.com

A Word About Utility Functions

If you have a utility function, use it.

An increasing utility function should be used to transform the random return variable r into another random variable U(r).

You still have to be able to compare distributions in ‘Utility space’.

Any risk averse utility function will lead to a preference for the Buy Lottery over the Sell lottery –the Buy will always have higher expected utility.

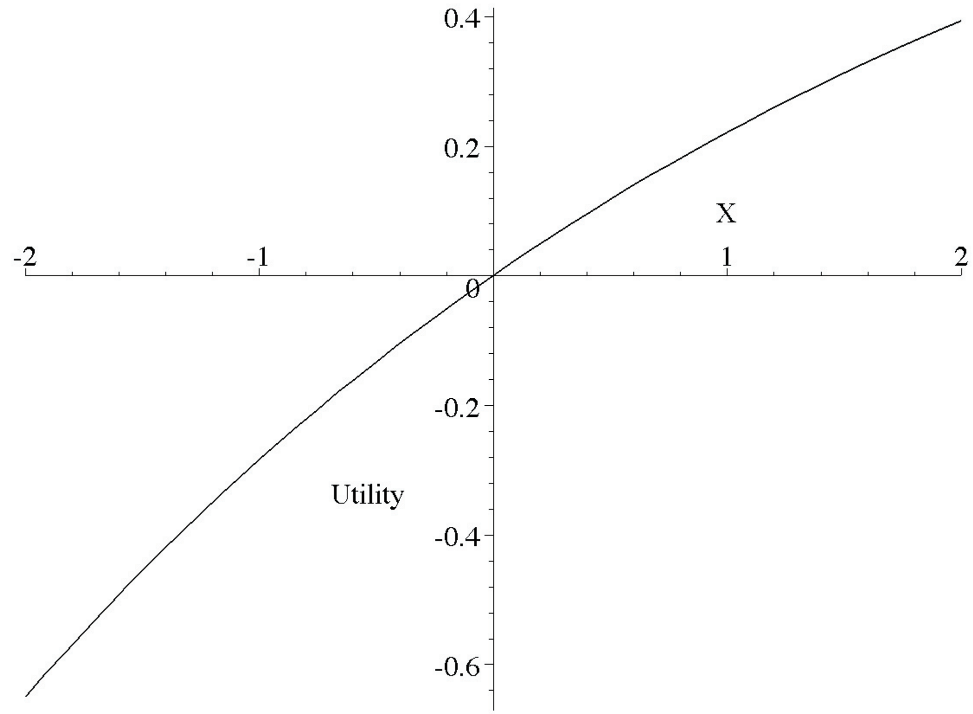

A (Mildly) Risk Averse Utility Function

Assume you are risk averse to some degree.

A symmetric distribution for r will produce a left biased one for U(r).

BUT

A sufficiently right biased distribution for R will produce right biased for U(R).

William.Shadwick@FinanceDevelopmentCentre.com

To rank investments we MUST rank distributions

A symmetric distribution for r will produce a left biased one for U(r) and a sufficiently right biased distribution for R will produce a right biased for distribution for U(R).

If the mean of U(r) and the mean of U(R) are the same, you are confronted with a Lottery-Like choice between U(r) and U(R).

In the upper panel we have symmetric and right biased distributions. In the lower panel the same distributions viewed through a risk averse utility function.

Unfortunately, they have the same mean.

Expected utility has run out of steam at this point. You have only *postponed* the problem, not solved it.

Now that you've used your utility function, use the Omega function of the new distribution.

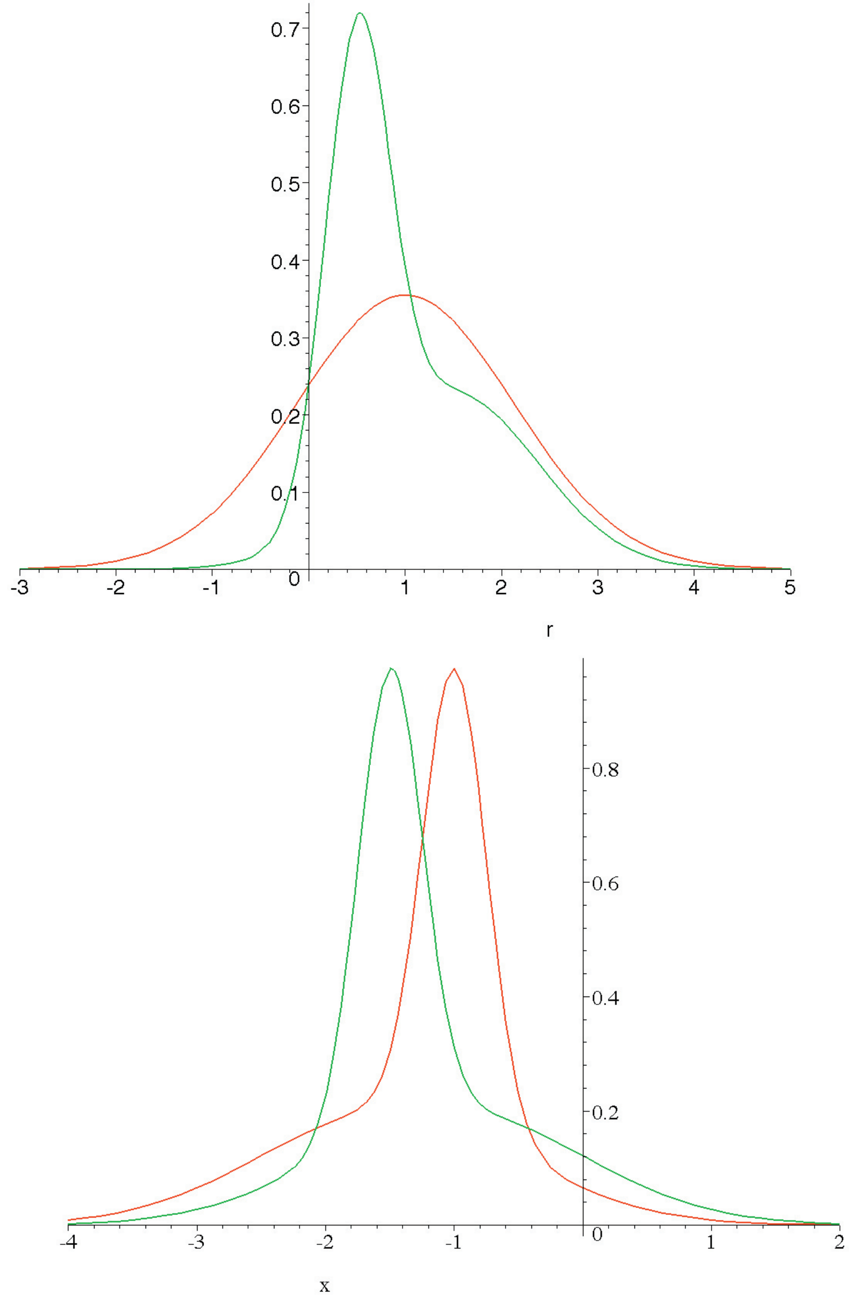

William.Shadwick@FinanceDevelopmentCentre.com

The 21st Century

Omega Functions make the Lottery decision correctly and easily.

The Buy Omega (green) dominates the Sell Omega (red).

This is generic for the lottery independent of mean or variance.

Omega functions are immediately informative about the quality of a bet on a given return level.

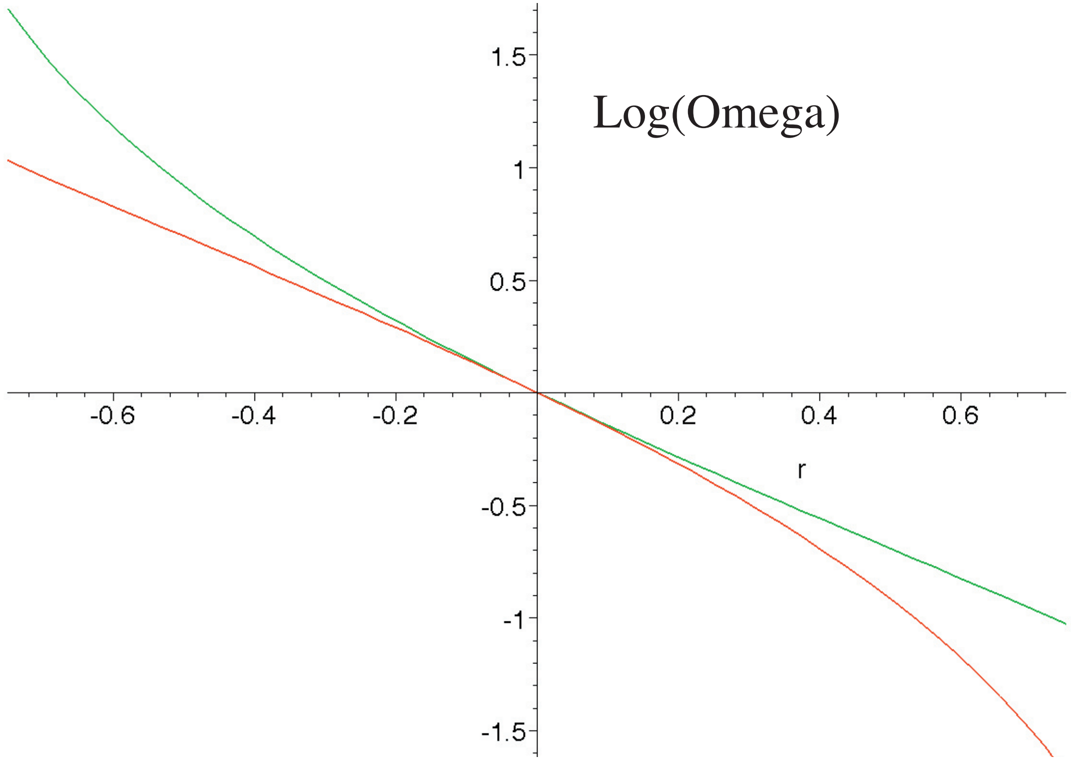

The larger the Omega value, the higher the quality of the bet.

Hedge Fund 'Lotteries, Omega functions and Sortino Ratios

Omega Functions often reveal 'lottery like' choices in hedge funds like FSell and FBuy.

The Sortino ratios treat these the same as the lottery Buy and Sell.

Who is right this time? Compare the Omega optimal portfolio with the Sharpe optimal (which is prefered by its Sortino ratio for MAR from 0.2 to 1.8% per month).

The portfolios are dramatically different:
Omega: 91%FBuy 9% FSell
Sharpe: 32%FBuy 68% FSell

The subsequent 12 month terminal values were:

Sharpe Optimal \$0.988
Omega Optimal \$1.013.

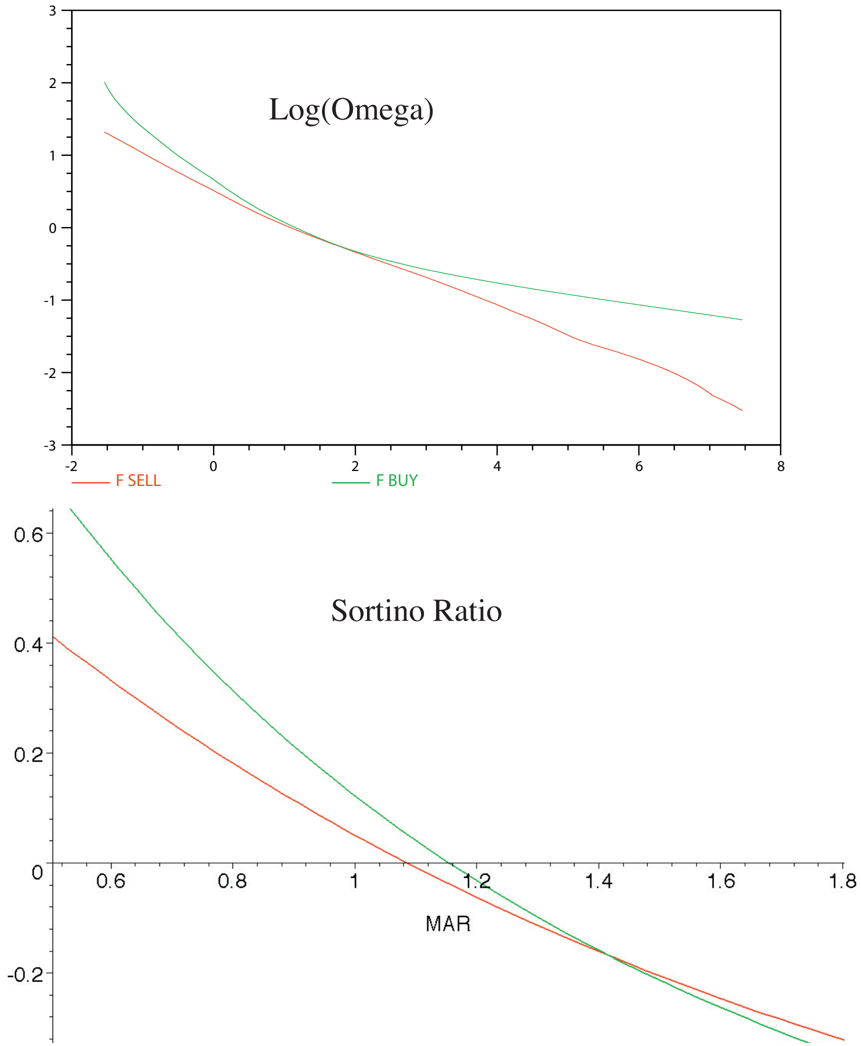

The Sortino Ratio (bottom panel) switches preference between FBuy and FSell for MAR 1.4% per month. The Omega preference (top panel) is constant for FBuy. FBuy in Green, FSell in Red in both.

William.Shadwick@FinanceDevelopmentCentre.com

Many managers are now reporting Omega function values.

This will lead to ambiguous, or worse, incorrect, conclusions!

The use of Omega function values at single or multiple points can miss critically important information. It is only the *entire* Omega function that carries the same information as the distribution itself.

It is easy to see the problem by looking at the Buy lottery with a succession of Sell lotteries with the same left bias but with increasing means.

Eventually, everyone will prefer the 'Sell' in spite of its left bias once its mean becomes big enough.

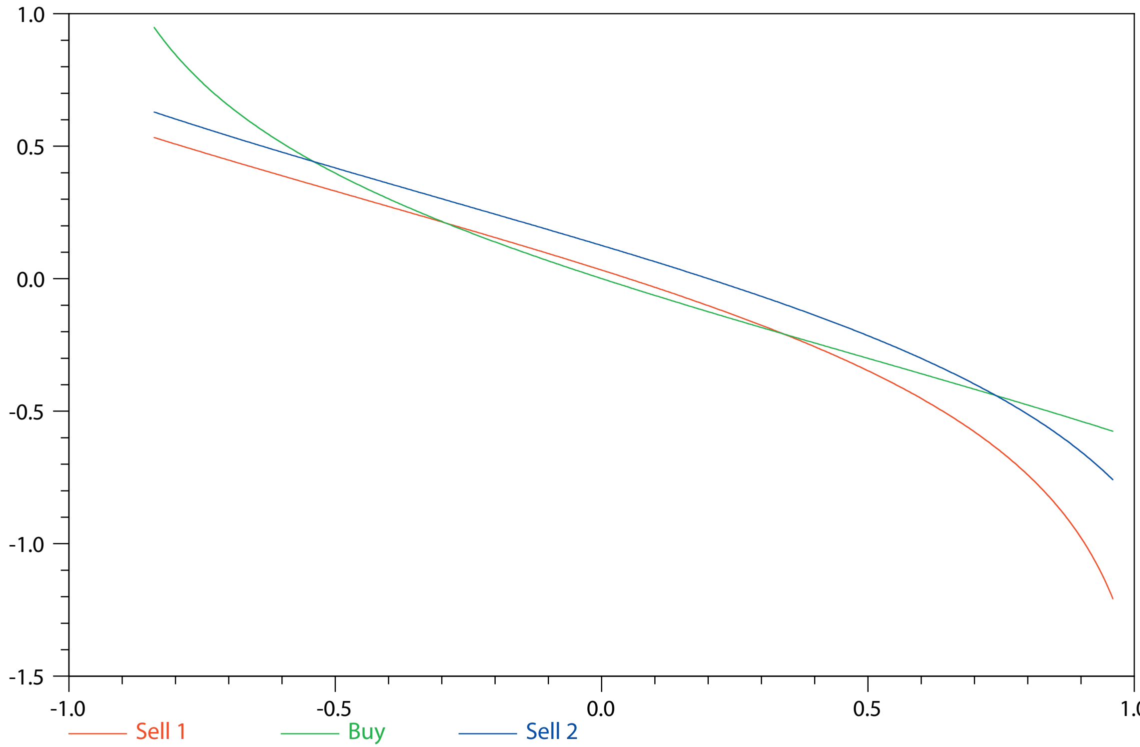

Ranks from Omega values at points can only lead to trouble.

William.Shadwick@FinanceDevelopmentCentre.com

To rank investments we MUST rank distributions

To rank distributions we have to produce a number for each.

For Normal distributions, only one number makes sense:
The Sharpe Ratio.

For asymmetric returns distributions, we need something new.

Expected values provide the most obvious route but we know this doesn't work when we only have 36 data points.

Fortunately there is another way:

William.Shadwick@FinanceDevelopmentCentre.com

Omega Metrics

- Evaluate risk and return to:

  Reward fat tails above the mean
  Penalise fat tails below the mean
  Reward higher mean
  Reduce to Sharpe ratio rankings for normal distributions

- Produce rankings that predict higher out-of-sample terminal values.
- May be customised to individual risk appetite
- May be combined with any utility function
- May be combined with market views or scenarios or calculated from historic data.

To rank investments we MUST rank distributions

Omega metrics start by transforming a distribution into its Omega function.

Omega metrics balance the right or left bias of the distribution, seen through its Omega function, against the contribution of the mean.

This is done in a way that reproduces Sharpe ratio rankings in the case of normal distributions.

It incorporates the effects of up and downside bias for asymmetric distributions, in analogy with the mean/variance efficient frontier.

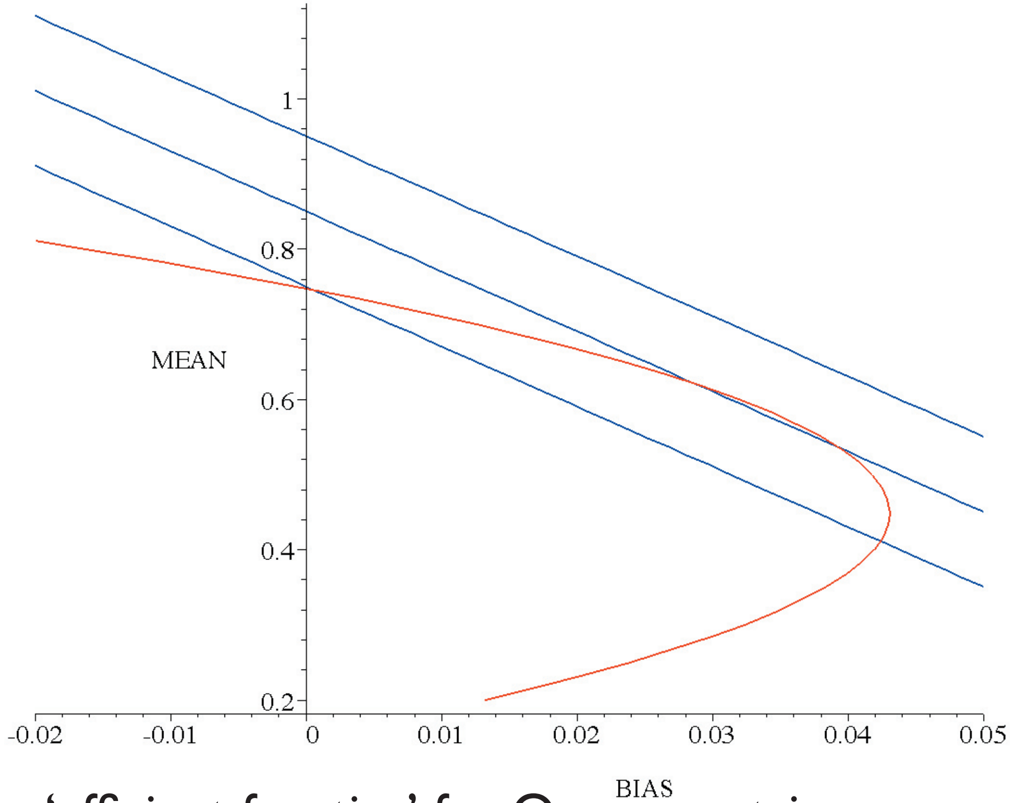

The 'efficient frontier' for Omega metrics.

Balancing 'mean' against 'bias' (which can be negative or positive) you want to be as far North-East as possible.

'Indifference curves' can be produced by linear or nonlinear trade-offs between mean and bias.

William.Shadwick@FinanceDevelopmentCentre.com

Higher Omega Scores Produce Higher Terminal Values

Omega Metrics allow you to rank individual assets, managers or funds based on their historic returns.

Omega F2 scores typically stabilise over 2 to 3 years of monthly data.

This effectively separates the top from the bottom performers.

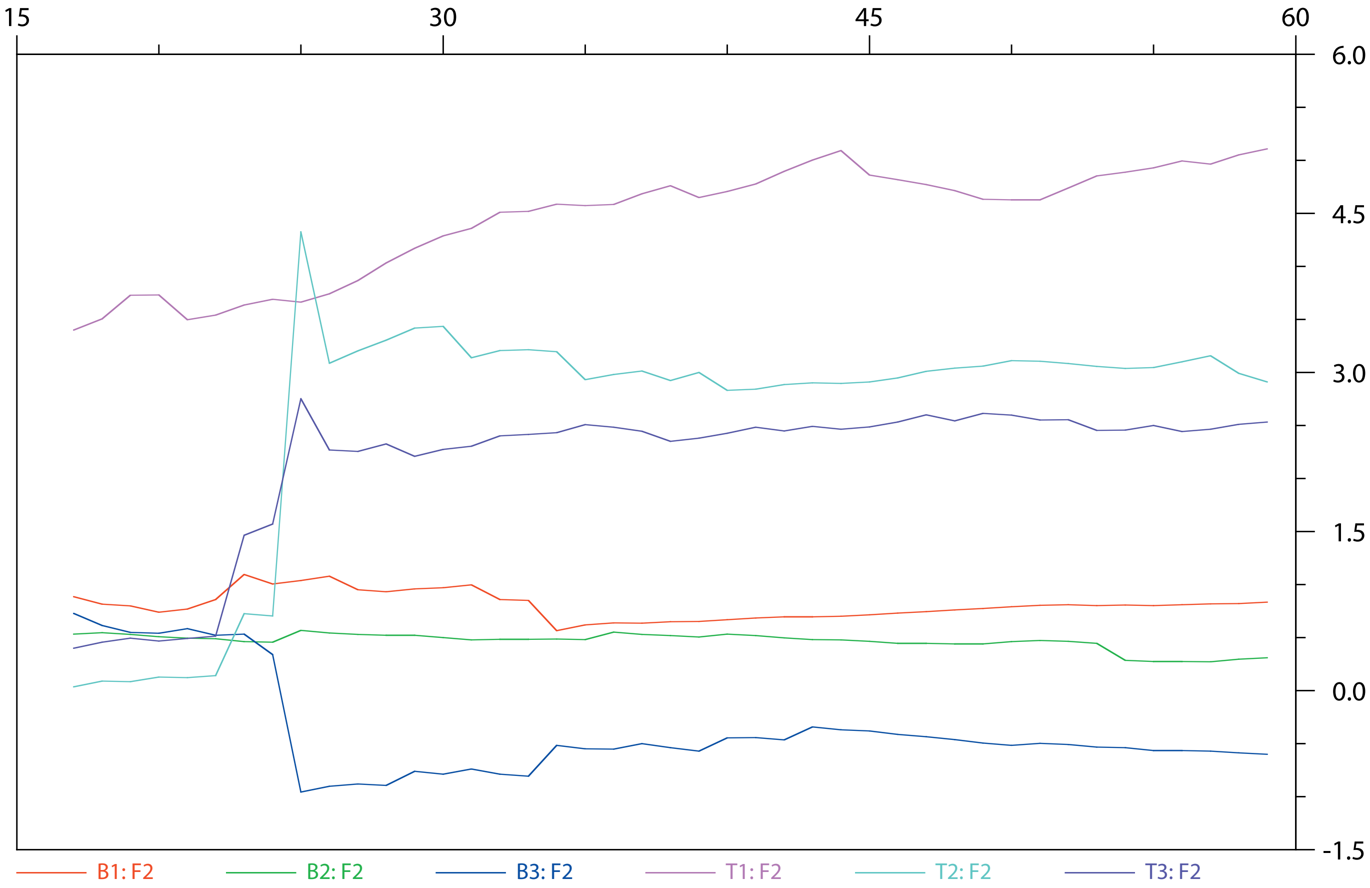

Three of the top 10 and three of the bottom 10 separated in 30 months

In a group of 30 managers the top 10 ranked group out-performed the bottom 10 producing an average of 5% more in the 12 months after the analysis.

William.Shadwick@FinanceDevelopmentCentre.com

Optimising Your Portfolio

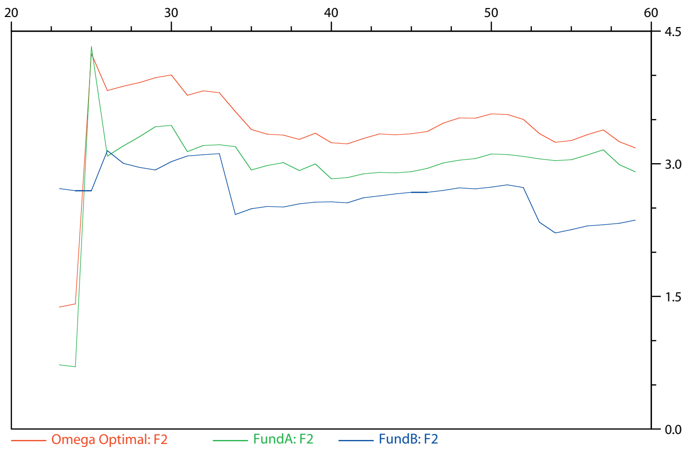

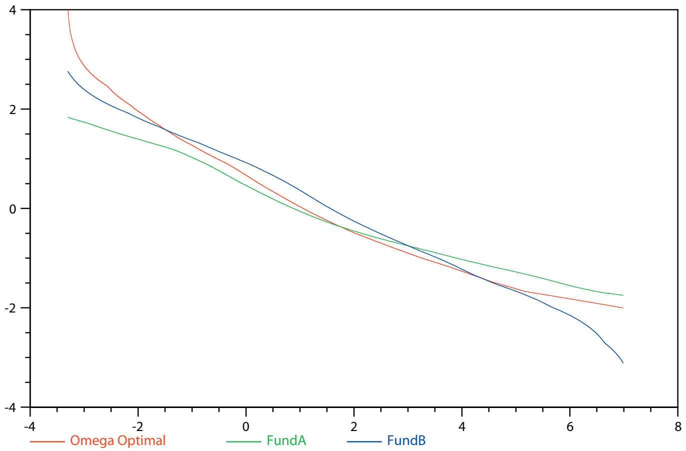

Omeg scores (top) and Omega functions (bottom) for Funds A and B and the Omega optimal combination. (60 Months of data)

An initial ranking of assets by their Omega scores shows you which should be selected or rejected for inclusion in a portfolio.

The allocation decision is the final, crucial step.

Omega optimisation produces the portfolio with the best possible Omega score based on the data available.

It incorporates dependence information from the joint distribution of returns, correctly accounting for the upside and downside bias of individual assets.

Omega Optimal Portfolios of assets with asymmetric returns routinely outperform those based on mean/variance measures.

In the 12 months after this analysis:

Omega Optimal Terminal Value: $1.32

Sharpe Optimal Terminal Value: $1.24

William.Shadwick@FinanceDevelopmentCentre.com

Hedging an Existing Portfolio

Using the same approach, Fund of Funds managers can hedge their portfolios.

The test of a hedge is its performance in the subsequent period.

CTA funds were considered as hedges for a portfolio of hedge funds.

The CTA fund chosen produced a dramatic reduction in the downside tail.

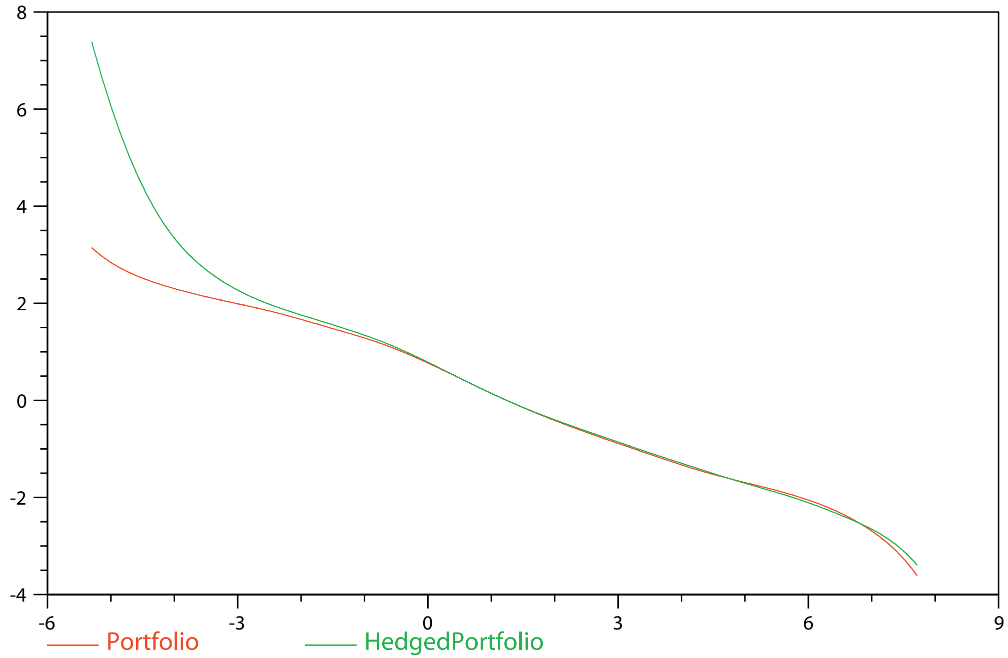

Log(Omega) for the hedged and unhedged portfolios

This was accompanied by a small increase in mean return–in effect an insurance policy with a negative premium.

In the subsequent 12 months, the unhedged portfolio would have returned $1.024 for every $1 invested. The hedged portfolio returned $1.059.

William.Shadwick@FinanceDevelopmentCentre.com

To rank investments we MUST rank distributions

The 21st Century Challenge is to rank asymmetric Distributions

The Tools of the 20th Century Are not Sufficient

Omega Metrics Provide a Simple, Effective Solution.

You Should Use Them.

William.Shadwick@FinanceDevelopmentCentre.com