A Universal Performance Measure

Con Keating*
William F. Shadwick

THE FINANCE DEVELOPMENT CENTRE
LONDON

Current Version: May 2002

Abstract

We present a new approach to analysing returns distributions, the Omega function, which may be used as a natural performance measure. Analysis based on Omega is in the spirit of the downside, lower partial moment and gain-loss literatures. The Omega function captures all of the higher moment information in the returns distribution and also incorporates sensitivity to return levels. We indicate how this may be applied across a broad range of problems in financial analysis and apply it to a range of hedge fund style or strategy indices.

* Corresponding Author: Con Keating E-mail: Con.Keating@FinanceDevelopmentCentre.Com

We should like to thank Jim Aspinwall, Philip Best, Chris Brooks, Ana Cascon, Jon Danielsson, Jim Golden, Drago Indjic, Thomas Klepsch, Narayan Naik and two anonymous referees for helpful discussions and comments. All errors are, of course, our own.

A Universal Performance Measure

1. Introduction

Many of the difficulties we encounter in performance measurement and attribution are rooted in two over-simplifications. The first is that mean and variance fully describe the distribution of returns. The second is that the risk-reward characteristics of a portfolio may be described without reference to any return level aside from the mean return. It is a generally accepted fact of empirical finance that returns from investments are not distributed normally. Thus in addition to mean and variance, higher moments are required for a complete description. It is likewise clear that a return at the level of the mean may be regarded as a gain by one investor and as a loss by another and that the "risk" of a return far above the mean has a different impact than that of one far below the mean.

In this paper we introduce a performance evaluation measure, Ω, which accomplishes the task of incorporating all of the higher moments of a returns distribution. It provides a full characterisation of the risk reward characteristics of the distribution in a way which is intuitively appealing and easily calculated. Instead of estimating any individual moments it measures their total impact, which is of course precisely what is of interest to practitioners. It also provides a risk-reward evaluation of a returns distribution which incorporates the beneficial impact of gains as well as the detrimental effect of losses, relative to any individual's loss threshold.

Omega is a natural feature of the returns distribution. In fact it is, in a mathematically precise sense, equivalent to the returns distribution. Thus its construction from a returns distribution is entirely canonical, requiring no choices and introducing no ambiguity not already present in the returns data. Omega is a function that may be evaluated at any value in the range of possible returns, so that it allows performance comparisons with respect to any 'risk' threshold in this range.

Omega may be used to rank manager performance, without the need to introduce utility functions[1]. In order to rank a collection of portfolios our performance function, Ω, will need just the simple decision rule that we prefer more to less, that we are not satiated.

In addition to providing corrections to mean-variance measures by taking higher moment information into account, Omega also takes into account the level of return against which a given outcome will be viewed as a gain or loss. Even in the case where returns are normally distributed, this provides additional information which mean and variance alone do not encode. This can lead to significantly different portfolio optimisations than are produced by traditional mean-variance analysis.

Before defining Omega, we provide some model returns distributions which illustrate the limitations of analysis which uses only the mean and variance. We show that the question of preference for one portfolio over another must vary with the level of return which is considered. We also show that the information included in a mean-variance approximation can be far less significant than the information which is neglected.

First we consider two assets with normally distributed returns.

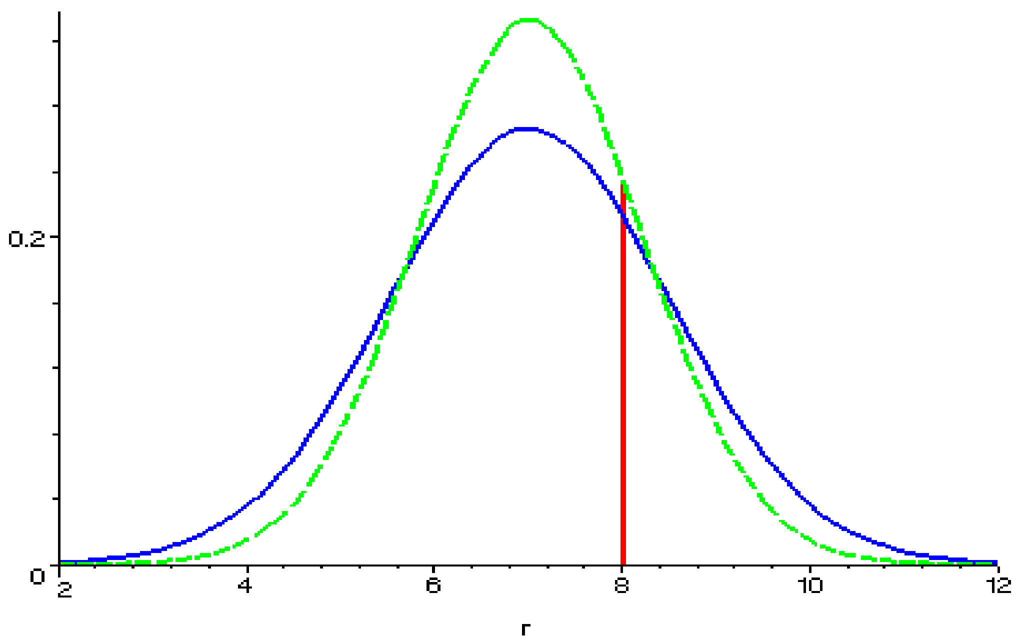

Diagram 1.1 Asset A (dashed) with mean 7 and variance 1.44, asset B (solid) with mean 7 and variance 2.25 and a loss threshold at r= 8.

Asset A has a mean of 7 and variance of 1.44, asset B has the same mean and a variance of 2.25. The Sharpe ratio says that we should prefer asset A to asset B. This ranking minimises the potential for loss, but of course it also minimises the potential for gain. If we use only mean and variance our rankings are biased in that they implicitly regard the potential for large gain and large loss as equally undesirable. We consider the position of an investor who regards a return below 8 as a loss and one above 8 as 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 not merely the mean return and the variance of returns. It is apparent from Diagram 1.1 that from this point of view, asset B is more attractive than asset A. For asset B about 25% of returns are above this investor's loss threshold, as compared with only 20% for asset A. The ratios of likelihood of gain to loss are 0.338 for asset B and 0.254 for asset A. This approach is not merely reversing the usual ranking in terms of variance however. For an investor whose loss threshold is at 6 rather than 8, the same analysis shows that asset A would be preferable to asset B. The ratios of likelihood of gain to loss relative to a threshold of 6 are 3.94 for asset A and 2.96 for asset B. As we will show in what follows, the Omegas for these two assets will show a change of preference at the mean, with asset A preferable to asset B for all return levels below the mean and asset B preferable to asset A for all return levels above the mean.

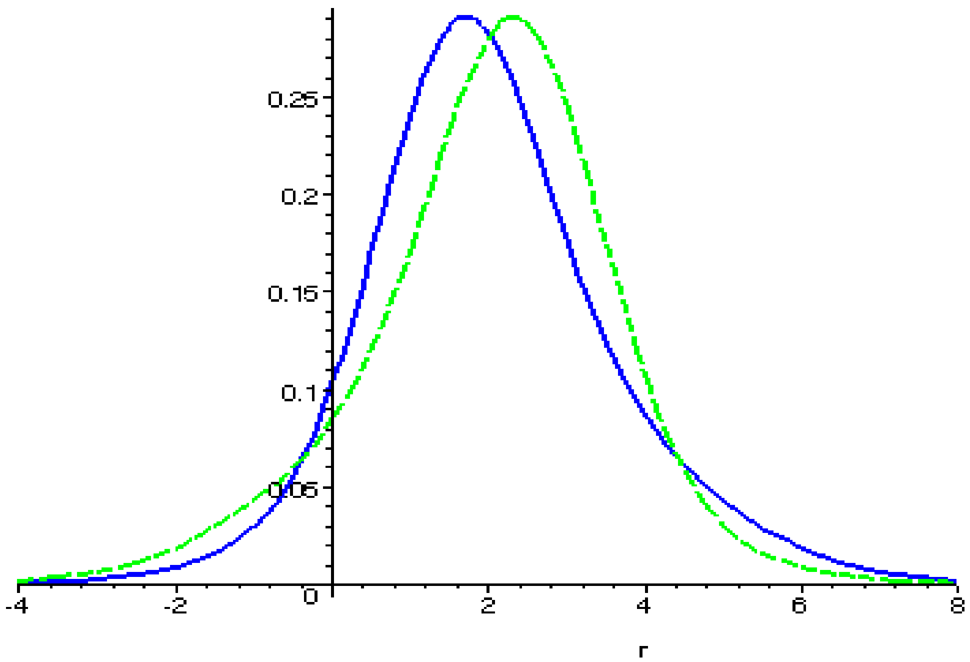

Diagram 1.2 Returns distributions for assets C (dashed) and D (solid) with mean of 2 and standard deviation of 1.6 .

Next, we consider two model returns distributions which have both have mean 2 and standard deviation of 1.6. These are illustrated in Diagram 1.2. All of their even moments are the same and all of their odd moments are equal in magnitude but opposite in sign. For example, asset C has skew of –0.398 while asset D has skew of 0.398 and both have kurtosis of 3.84.

From the point of view of mean variance analysis however, these two assets are indistinguishable. [2] The likelihood of a return at any level below the mean is greater for asset C than for asset D while any return above the mean is more likely with asset D. In particular, catastrophic loss is much more likely with asset C than asset D. The probability of a return 4 standard deviations below the mean is 5 times as high for asset C as it is for asset D. We will see later that the omegas for these two assets agree at the mean but that for any other return level asset D is preferable to asset C. As all of the even moments are equal, this preference is a feature of the difference in the odd moments.

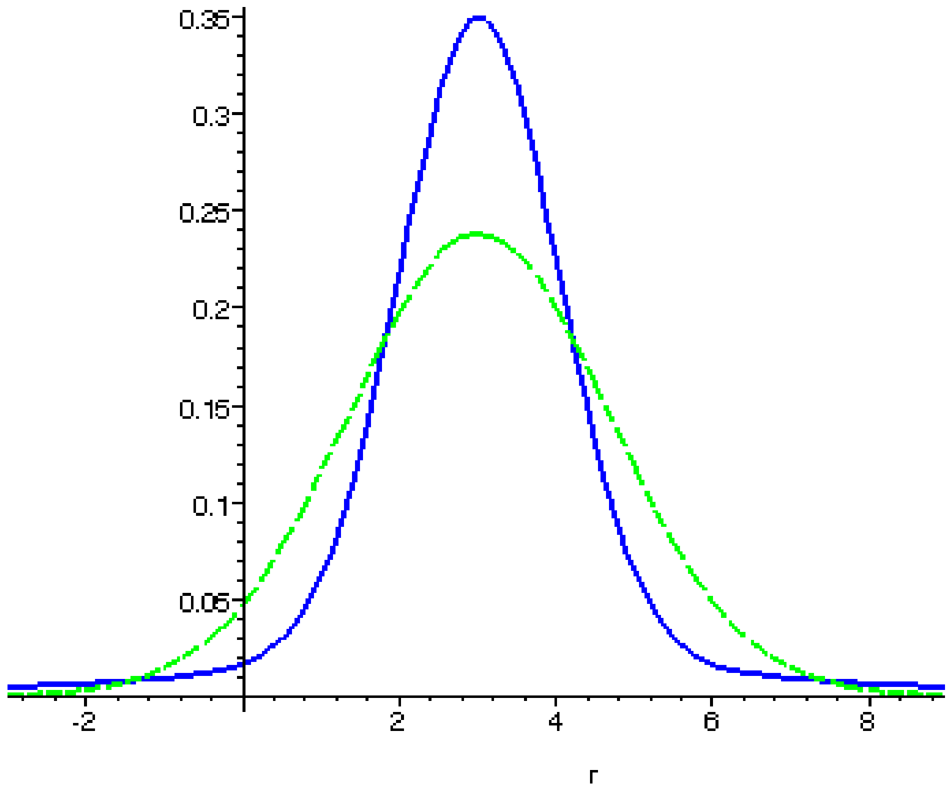

Diagram 1.3 Returns distributions for assets E(dashed) and F (solid).

Finally we consider two model returns distributions which again have the same mean and variance but which have dramatically different capacities for large losses and large gains. Asset E has normally distributed returns with a mean return of 3 and a variance of 2.76. The returns for asset F are symmetrically distributed and it has the same mean and standard deviation as asset E. Asset F exhibits a vastly different propensity both for catastrophic loss and for large gain than does asset E. For example the likelihood of a return more than 4 standard deviations below the mean is 137 times as likely for asset F than for asset E. Because the returns are symmetrically distributed, a return more than 4 standard deviations above the mean is also 137 times as likely with asset F than with asset E. On the other hand, in the range of returns closer to the mean it is less obvious how to rank the two. The expected return conditional on a return less than the mean is 1.88 for asset F and only 1.66 for asset E. The expected return conditional on a return greater than the mean however, is 4.12 for asset F and 4.34 for asset E so that preferences for asset E and F clearly depend on the range of returns under consideration. Because the two assets have the same mean and variance, it must be the case that these effects are due to higher (even order) moments. The kurtosis of asset F is 9.6 (i.e. excess kurtosis is 6.6) and, in fact, all of the even moments are higher than those of the normal distribution.

Practitioners are well aware that the impact of higher moments can be significant. Over the last decade, we have seen many attempts to analyse and attribute returns from portfolios and securities which by design seek to capture asymmetries in investment returns. This has resulted in the current fashion for style analysis. Much of this work has focussed upon “hedge” funds. Some traditional asset “classes”, such as corporate bonds, also require the consideration of higher moments in order to fully understand their performance characteristics. The central problem is that, as in our examples above, the mean and variance toolkit is simply inadequate for such analysis.

We note that there is a very substantial body of work that seeks to extend the mean-variance framework of modern finance to encompass higher moments. The theoretical difficulties within that literature arise from the need to specify the form of a utility function and the substitution across moments. In addition, there is a serious obstacle to incorporating the effects of higher moments in performance measurement, as data are often both sparse and noisy. This means that estimation of the moments is error prone and any attempt to attribute performance characteristics to them individually is therefore difficult if not impossible to do reliably.

What is more, it is extremely difficult to establish that an effect is caused by, say the third moment as opposed to all moments of order three or higher. While the model portfolios in our examples appear to conform to the notion that positive skew is desirable and negative skew undesirable, this is not by itself true in general. In fact, asset C and D have the same even moments of all orders and opposite odd moments of all orders, they do not simply differ in skewness. Assets E and F have all odd moments equal to zero but differ substantially not only in their kurtosis but in all of their even moments. In these cases, moments of high order have a substantial impact on the risk-reward characteristics of the distributions.

This suggests very strongly that any approach which depends on systematically extending econometric analysis based on individual moments is doomed to fail.

2. The Omega Function

We begin with an analogy with a simple bet. The first consideration is how much we stand to win if we win and how much we stand to lose if we lose. By itself however, this information is not enough to judge the quality of our bet. To make a meaningful comparison of the potential gains and losses, we need to weight them appropriately by their probabilities of occurrence.

The investment situation differs from this only in that the 'stake' is unknown at the outset and the loss threshold L must be specified exogenously. That the level which constitutes 'loss' varies by individual and application of the investment is clear. The loss threshold might be the rate of inflation for an investment providing a pensioner's income or the return on a benchmark index for a 'growth' fund. We illustrate this in Diagram 2.1 with a model returns distribution relative to a benchmark index and a return threshold set at 70 basis points above the benchmark return.

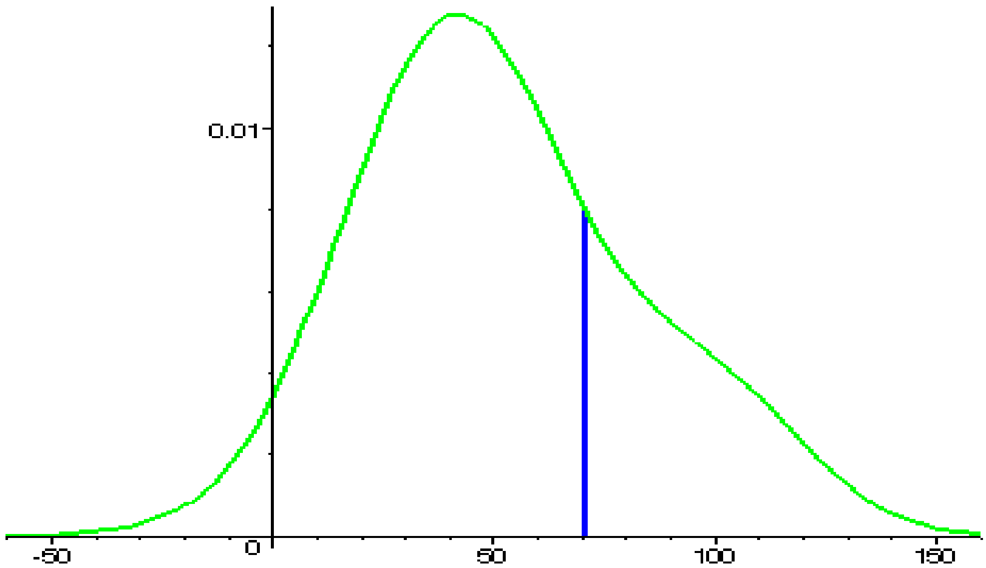

Diagram 2.1 The returns distribution with the threshold set at 70 basis points.

Once we have specified the return level L however, we may make a probability weighted comparison of gains and losses relative to it, in the same way we did with the simple bet. The expected gain, given a return greater than L, is just the amount by which the conditional expectation $E(r | r \geq L)$ exceeds the threshold. The expected loss, given a return less than L, is the amount by which the conditional expectation

$E(r|\, r \leq L)$ falls below the threshold. Thus the expected gain and loss relative to the threshold r = L, are $g = E(r|\, r \geq L) - L$ and $l = L - E(r|\, r \leq L)$ respectively. We illustrate this in Diagram 2.2 with the cumulative distribution for the returns of Diagram 2.1.

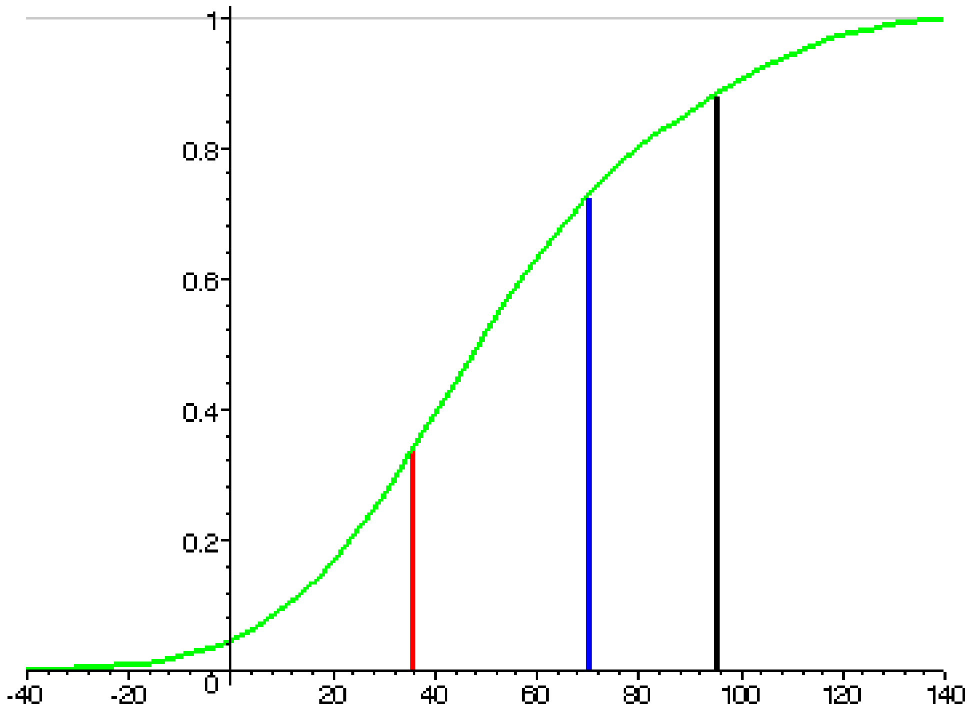

Diagram 2.2 The cumulative distribution with the expected returns given a return above and below 70 basis points.

As before, these levels by themselves are not informative. In order to compare potential gains and losses meaningfully, we need to weight them by the appropriate probabilities. If $F(r)$ is the cumulative distribution function for the returns, then the probability of a return less than our threshold is $F(L)$. The probability of a return above the threshold is $1 - F(L)$. Thus the ratio $\frac{g \times (1 - F(L))}{l \times F(L)}$ is a measure of the quality of our investment 'bet'. As is illustrated in Diagram 2.3, this is the ratio of the area of the upper rectangle to the area of the lower.

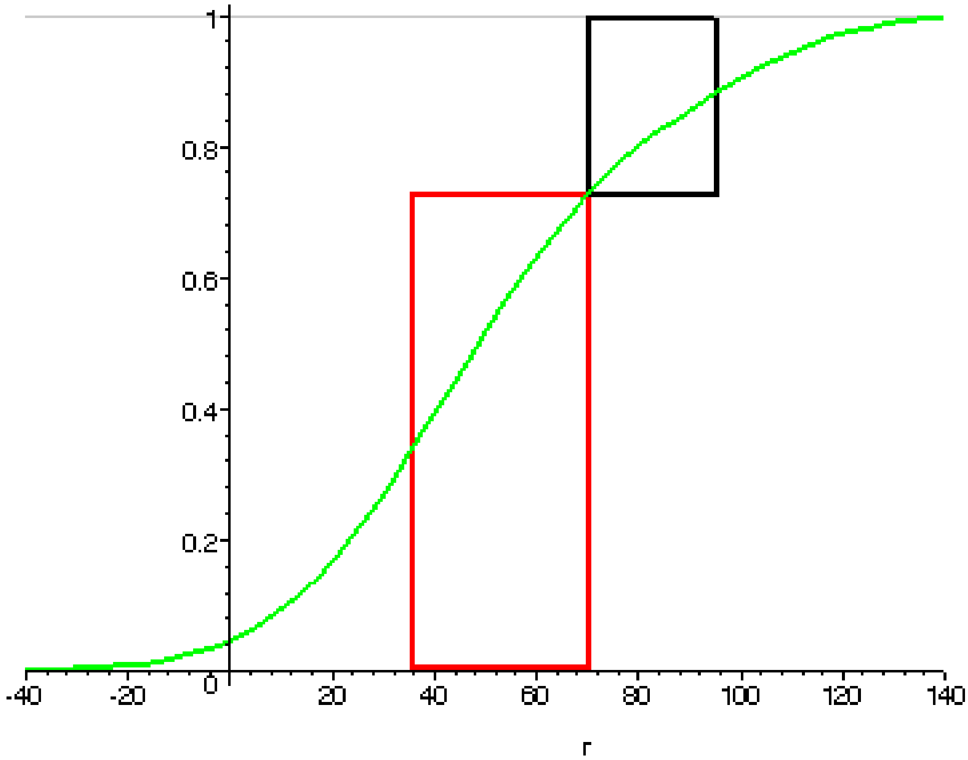

Diagram 2.3 Probability weighted gains and losses relative to the threshold of 70 basis points. The ratio of the area of the upper box to the area of the lower box is an estimate of the quality of a bet on a return greater than 70 basis points.

This ratio considers only one particular gain and loss possibility. However, if (a,b) is the possible range of returns then gains and losses of any amount in this interval can occur with some probability. To take this into account, we may generalise our initial comparison of gains and losses by considering a sequence of gains and losses and summing these with their appropriate probability weights. This approach leads to a unique limiting case as we allow the unit of gain or loss to become progressively smaller. We illustrate this idea in Diagrams 2.4 and 2.5 below.

As we allow our unit of gain or loss to shrink to zero and sum the probability weighted gains, we obtain $I_2(70) = \int_{70}^{b} [1 - F(x)]\, dx$ in the limit. Similarly, the sum of the probability weighted losses leads to $I_1(70) = \int_{a}^{70} F(x) dx$.

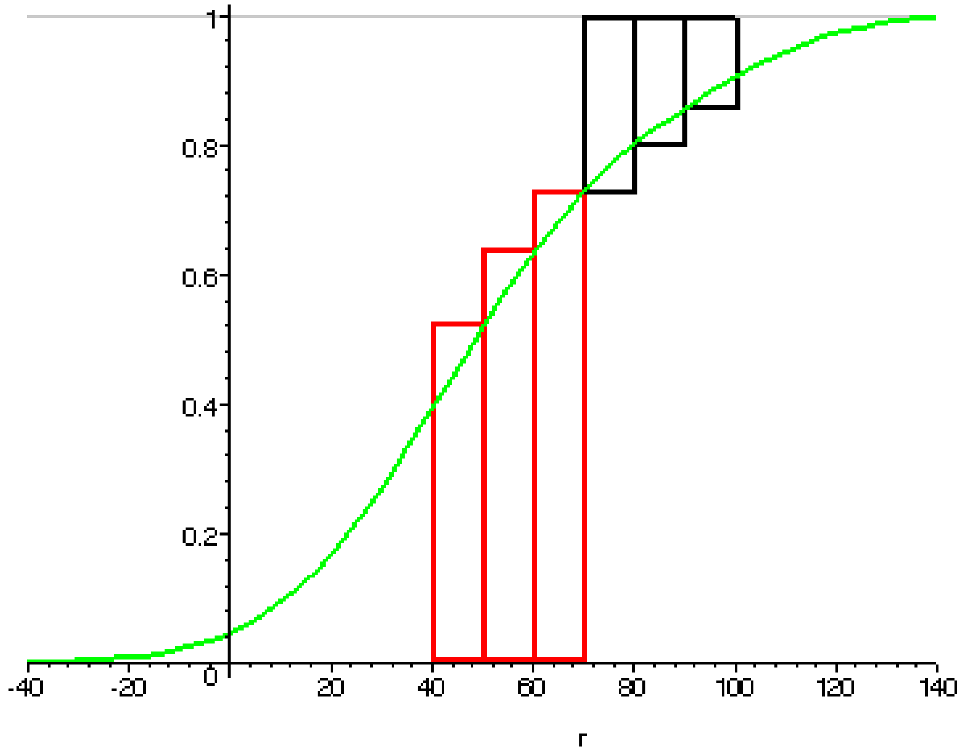

Diagram 2.4. Reducing the unit of gain and loss refines the estimate of the quality of a bet on a return above the level of 70 basis points.

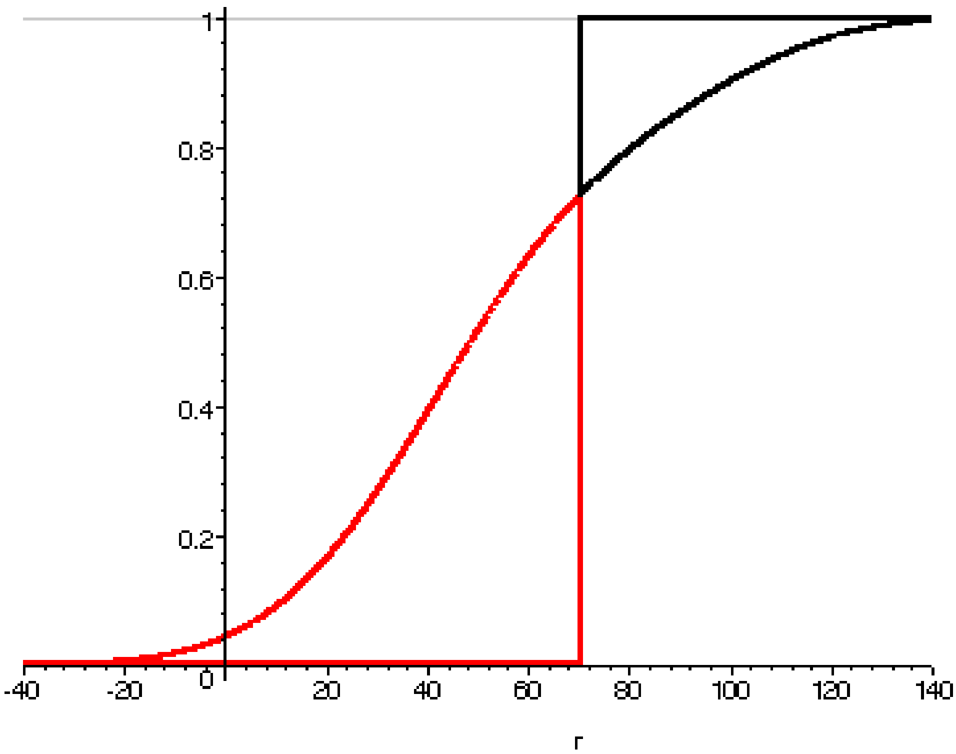

Diagram 2.5 The limit as the unit of gain and loss shrinks to zero. The ratio of the upper area to the lower is Ω(70).

The ratio of these two is a measure of the quality of our investment 'bet' relative to the return threshold r = 70.

The construction which led us to this ratio can be repeated for any return threshold r and we call the resulting function Omega:

$$\Omega(r) := \frac{\int_r^b [1 - F(x)]\, dx}{\int_a^r F(x) dx} .$$

Thus, in our example, the probability weighted ratio of gains to losses relative to a return of 70 basis points above the benchmark mean return is Ω(70).

Like the cumulative distribution for the returns, Ω is a function of return level. In fact, in a mathematically precise sense, Ω is equivalent to the returns distribution. As a result, available information from the returns distribution, including higher moments, is encoded in Ω. We shall illustrate this point later with a range of examples.

No parametric assumptions are needed and no constraints are placed upon the form of the distribution aside from the requirement that the integrals $I_1(r) = \int_a^r F(x)dx$ and $I_2(r) = \int_r^b 1 - F(x)dx$ exist. For analytic distributions defined over infinite intervals this is easy to deal with. In practice, existence presents no problems since we work with discrete return observations.

Assuming the convergence of the integrals, Ω is a natural feature of the underlying probability distribution. As we indicate below, Ω is a smooth monotone decreasing function from (a,b) to $(0,\infty)$. We will also show that, independent of the returns distribution, Ω takes the value 1 at the distribution's mean μ.

The function Ω(r) allows us to compare returns for different assets and to rank them on the basis of the magnitude of their Omegas. The rankings will depend on the interval of returns under consideration and will incorporate all higher moment effects.

We conclude this section with the Omega functions for the model asset distributions introduced in section1. (Because of the large variation in the values of Omega over the range of possible returns, we have presented the natural logarithm of Omega rather than Omega.)

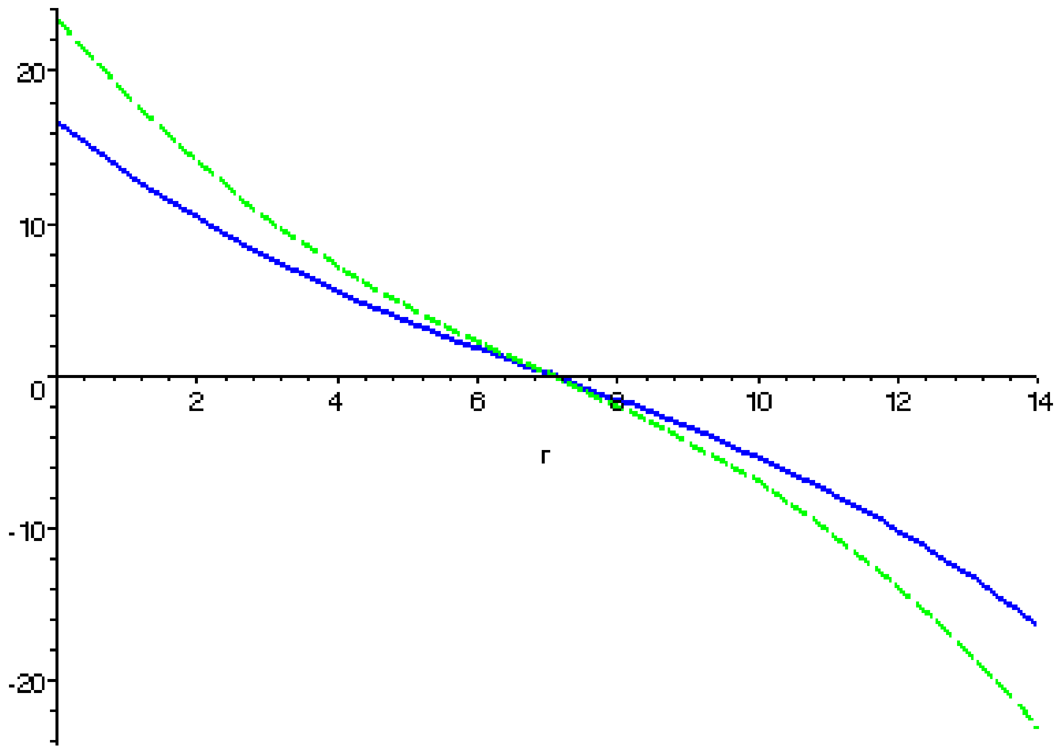

Diagram 2.6 log(Ω (r)) for assets A(dashed) and B (solid). Asset A is preferable at any return level below the mean and asset B is preferable above the mean.

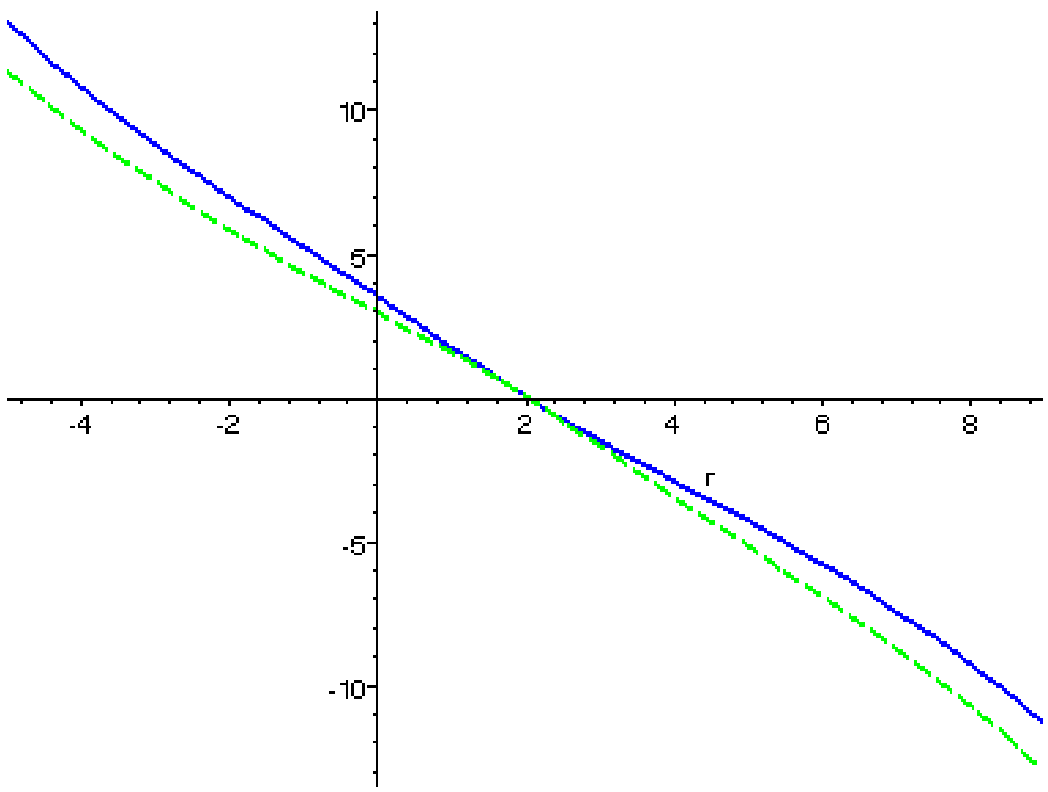

Diagram 2.7 log(Ω(r)) for assets C(dashed) and D (solid). Asset D is preferable to asset C everywhere except at the mean.

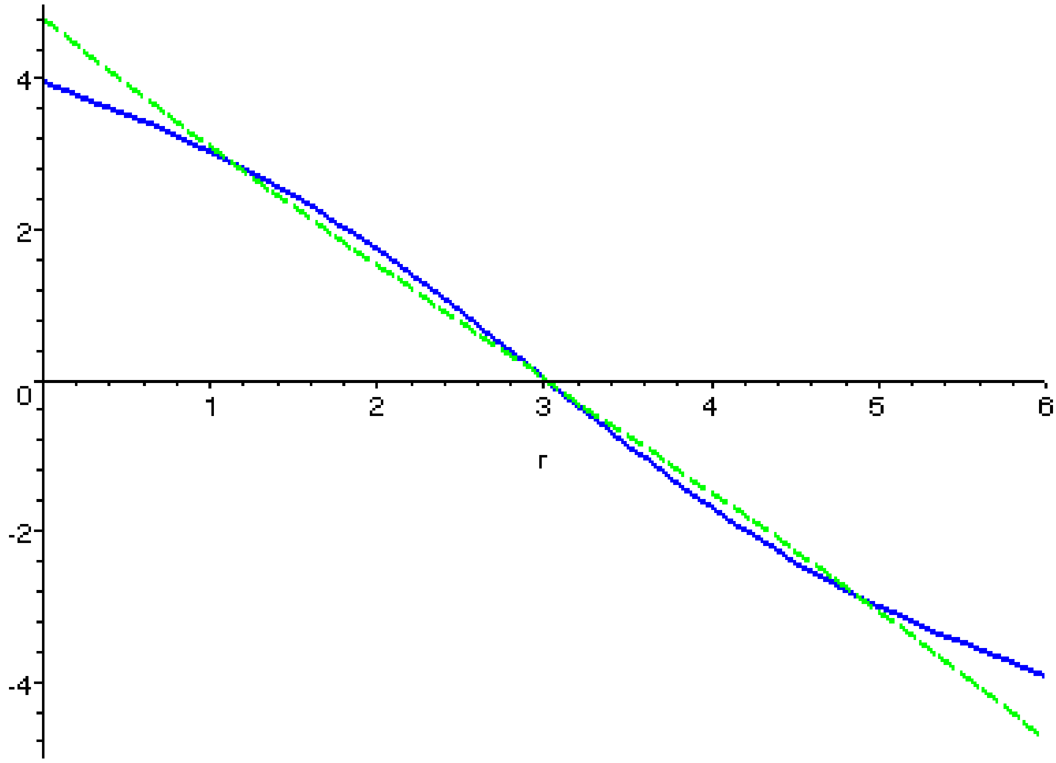

Diagram 2.8 logΩ(r) for assets E (dashed) and F (solid). Asset E is preferable for all returns below 1.4 while asset F is preferable for all returns above 4.6. Between 1.4 and the mean return, asset F is preferable to asset E. This preference is reversed between the mean and 4.6.

3. Some elementary properties of the Omega measure

We indicate some of the properties of the Omega function here. A more detailed exposition is left to another paper[3]. We begin by examining the sensitivity of Ω to changes in the return level. We have $\Omega(r) = \frac{I_2(r)}{I_1(r)}$ where $I_1(r) = \int_a^r F(x)dx$ and $I_2(r) = \int_r^b [1-F(x)]dx$. We may differentiate this expression with respect to r to obtain $\frac{d\Omega}{dr} = \frac{\frac{dI_2}{dr} I_1 - \frac{dI_1}{dr} I_2}{I_1^2}$, or more explicitly, $\frac{d\Omega}{dr} = \frac{[F(r)-1]I_1 - F(r)I_2}{I_1^2}$. In particular we see that $\frac{d\Omega}{dr}$ is as smooth as *F(r)* and that $\frac{d\Omega}{dr} < 0$ everywhere.

Thus Ω is a smooth monotone decreasing function from (a,b) onto $(0,\infty)$ from which it follows that it takes the value 1 precisely once. It is a consequence of the definitions of I_1 and I_2 that the mean satisfies $\mu = I_2(0) - I_1(0)$ and one may deduce from this, and the definition of Ω, that $\Omega(\mu) = 1$.

Unless the cumulative distribution fails to be differentiable, the Omega measure has derivatives of at least second order and these may also be used to distinguish between distributions which differ in their higher moments, as we illustrate in appendix B.

We may also consider the variation of the Omega measure as the underlying cumulative distribution varies according to various scenarios. This is also straightforward but we leave it to another paper[4].

The Omega measure is an affine invariant of the returns distribution. That is, for any affine change of variable, $r \to \varphi(r) = Ar + B$, with *A>0,* there is an induced cumulative distribution and the Omega measure for the induced distribution, $\hat{\Omega}$ satisfies $\hat{\Omega}(\varphi(r)) = \Omega(r)$. Conversely, if this relationship is satisfied by any change of variable φ then $\varphi(r) = Ar + B$. For affine changes of variable with *A<0*, the relation is $\hat{\Omega}(\varphi(r)) = \dfrac{1}{\Omega(r)}$.

The variation of Omega with time is also meaningful as $\dfrac{d\Omega}{dt}$ contains the full information set of the series, such as auto-correlations of all orders, and when used to compare with another security or portfolio of similar periodicity, must also contain the cross-relations. The time evolution of Ω for two index returns series evaluated at a risk threshold of zero and the equivalent Sharpe ratio are shown in the section Applications.

Omega may also be directly related to other techniques in common use such as tracking error[5]. In the applications which follow, we also report the hedge fund indices relative to the return on the MSCI index for the purpose of illustration.

As we have pointed out, in order simply to rank the performance of two or more assets, no utility function is required, however we may apply any. For example: Any utility function which is monotone may be used as a transformation of the returns (or equivalently to induce a transformation of the cumulative distribution). The resulting Omegas will produce changes of preference at different points except in the case where the change of variable is affine.

4 Applications

Both initial applications are to a set of portfolio returns for a range of hedge fund style indices from two vendors and two traditional comparisons, MSCI and SWGBI. The data is monthly for the period beginning January 1993 and ending April 2001, 100 data points for each series. The data was presented blind and nothing is known of the portfolios beyond their name descriptions. The descriptive statistics are presented as Table 01. It is evident that these distributions are far from normally distributed but the Jarque Bera statistics are not reported for brevity. A pseudo-Sharpe ratio, where the risk-free rate is zero, is presented. These values range from 1.70 to 0.23. It is interesting to note that by the Sharpe measure, and also by Omega, which includes higher moment effects, the bond and equity indices are the worst performing portfolios. Auto-correlation values for the series and the squared series were also computed. Positive (0.20 – 0.55) first order auto-correlation was evident but perhaps more interestingly similar pairs such as ACSA – HCSA return differing values - 0.40 and 0.55 respectively in this case[6]. The squared series showed no evidence of statistically significant auto-correlation.

The overwhelming impression is that, though these are competing suppliers of index data series, there is a considerable disparity between them. There is an important caution for anyone comparing a portfolio with its style index here. The style index chosen is critical. If we compare the obvious pairings and consider the hypothesis that each of the two sequences is a sampling from some common distribution, the hypothesis is rejected (at 5%) for all pairs except AMA-HMA, where the difference appears to be a simple bias[7]. It is interesting to note that the government bond index exhibits small positive skewness as might be expected from the effect of convexity.

Somewhat more surprisingly the world equity index exhibits negative skewness and proved difficult for the traditional fund manager[8] to outperform over this period.

The (product moment) correlations of the data sets were calculated and are presented as table 02. The point to note here is that the equity and world government bond indices were in line with historic relations at 0.24. The hedge fund strategies showed negative correlation with the world government bond index. Perhaps the most alarming feature of this matrix is that the correlations with other strategies from the same data supplier were typically higher than correlations across pairs.

A principal components analysis was also conducted. Table B, below, lists the first six values and their cumulative explanatory power. The surprise here, particularly so as there are pairs within the data, is that the explanatory power of the eigenvalues diminishes very slowly.

Eigenvalues	1	2	3	4	5	6
Value	9.3869	1.6709	1.3582	1.0966	0.7814	0.7344
% of variability	52.15	9.28	7.55	6.09	4.34	4.08
Cumulative %	52.15	61.43	68.98	75.07	79.41	83.49

Table B

The correlation to factors matrix is given as table 03. The points to note in this are that the world government bond index is essentially uncorrelated to the first factor, while with the exception of index HMN, all others are strongly positively so. The second factor has equity, global bonds and HMN responding positively strongly while ACSA responds strongly negatively. In the table some of the more intriguing relations have been printed in bold type. For example, indices AM and HM share a response to factor 4 in the 10% - 25% explanatory range. The overwhelming impression from this factor matrix is that noise is dominant.

For the purposes of comparison, the only further proviso which we need is that comparison is valid across any range of returns.

Return Threshold – Zero

The Omega and Sharpe ratios with a risk free of zero with the return threshold set at zero are tabulated below, where the left hand sequence lists the indices ranked by their Sharpe ratio and the right hand lists the indices ranked by their Omega values, $\Omega(0)$.

	Sharpe	Indicator		$\Omega(0)$
ARV	1.699	1	ARV	43.98
ACSA	1.338	0	AELS	20.33
AMA	1.245	0	ACSA	19.19
AELS	1.214	0	AMA	18.04
HMA	1.040	0	AED	12.80
AED	0.964	0	HMA	11.94
HCSA	0.741	0	ADIST	6.19
ADIST	0.703	0	HED	6.04
HED	0.652	0	HDIST	5.73
HDIST	0.608	0	HCSA	5.69
HMN	0.544	0	HH	4.16
HH	0.533	0	HMN	3.79
HM	0.475	1	HM	3.79
AM	0.459	1	AM	3.27
ASS	0.399	1	ASS	2.85
HHLB	0.284	0	SWGBI	2.11
SWGBI	0.281	0	HHLB	2.06
MSCI	0.232	1	MSCI	1.78

The column marked indicator takes the value 1 when the Sharpe ratio agrees with the Omega ranking as to rank order. There are just five points of agreement, a clear indication of the importance of higher moment effects. The Kendall and Spearman rank correlations are 0.89 and 0.97.

Of course, to obtain a full picture of relative performance we should not only be considering the ordering by Omega at the single threshold $r = 0$, but rather the function $\Omega(r)$ over the range of returns, as is illustrated as Diagram 4.1 below.

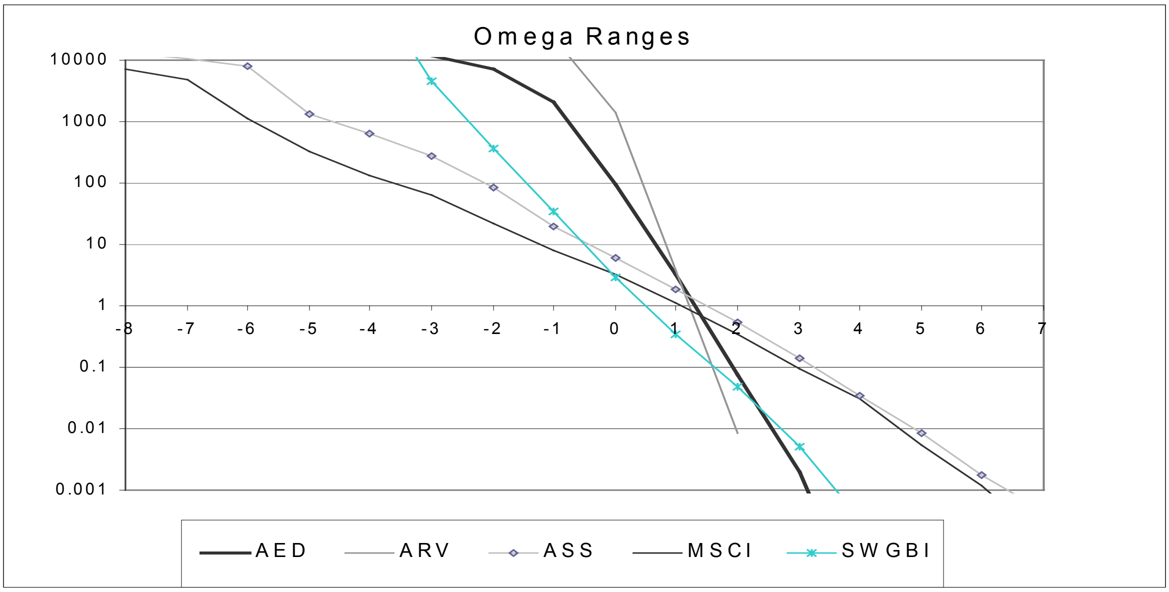

Diagram 4.1 Omega as a Function of Return Threshold.

For clarity, this shows only a selection of the index sequences over the range of returns experienced, on a logarithmic scale. Points where the Omegas cross are indifference points for choices between particular portfolio pairings. In the broadest of terms, the steepness of the Omega function is a measure of its risk. The steeper, the less risky. Though not shown here, the majority of the hedge fund indices are steeper than the SWGBI and MSCI, in the manner of ARV and AED above. The index ASS above seems to offer true value, being most consistently better than the MSCI. Above its mean, a steeply sloped Ω also implies a very limited potential for further gain.

At threshold –1, a high risk tolerance, the preference ordering is ARV, AED, SWGBI, ASS, MSCI while at threshold +2, the preference ordering is ASS, MSCI, AED, SWGBI, ARV.

Comparison of the indifference points based upon the Sharpe measure, where the risk free rate changes, and the Omega function is also possible as is illustrated in Diagram 4.2 below. The data-set here consists of a UK equity index, an international bond index and a UK property index. This diagram has a log scale for Ω and a linear scale for the Sharpe ratio.

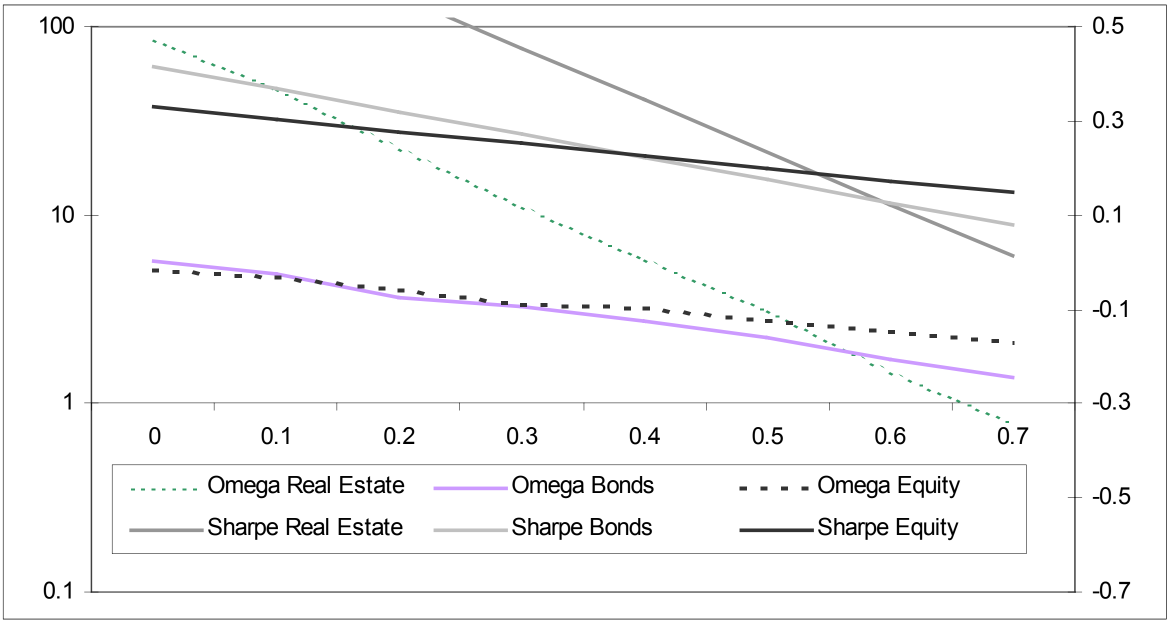

Diagram 4.2 Omega and Sharpe ratios (for threshold varying risk free)

Notice that none of the indifference points are coincident, again illustrating the importance of the corrections due to higher moment information which Omega incorporates.

"Risk" Threshold - MSCI returns

The second application uses the return from the MSCI as its "risk" threshold. The descriptive statistics for these distributions are appended as table 04. The tabulation below has the Sharpe ordering to the right and the Ω ordering to the left.

	Ω	Indicator		Sharpe
HH	1.58	1	HH	0.17
HM	1.43	0	ASS	0.14
ASS	1.41	0	HM	0.13
AMA	1.39	1	AMA	0.13
AED	1.36	1	AED	0.12
AM	1.35	1	AM	0.12
ACSA	1.32	1	ACSA	0.11
ADIST	1.29	1	ADIST	0.10
ARV	1.26	0	HED	0.09
HED	1.26	0	ARV	0.09
AELS	1.18	0	HHLB	0.07
HHLB	1.18	0	AELS	0.06

HMA	1.10	1	HMA	0.04
HDIST	1.03	1	HDIST	0.01
HCSA	1.03	1	HCSA	0.01
HMN	0.84	1	HMN	-0.07
SWGBI	0.79	1	SWGBI	-0.09

There is once more some disagreement between the two rank orderings, due to higher moment effects.

A selection of the Ω functions for these MSCI relative portfolios is shown below as Diagram 4.3. In order to facilitate comparison with Tracking Error type measures, we present these demeaned or normalised.

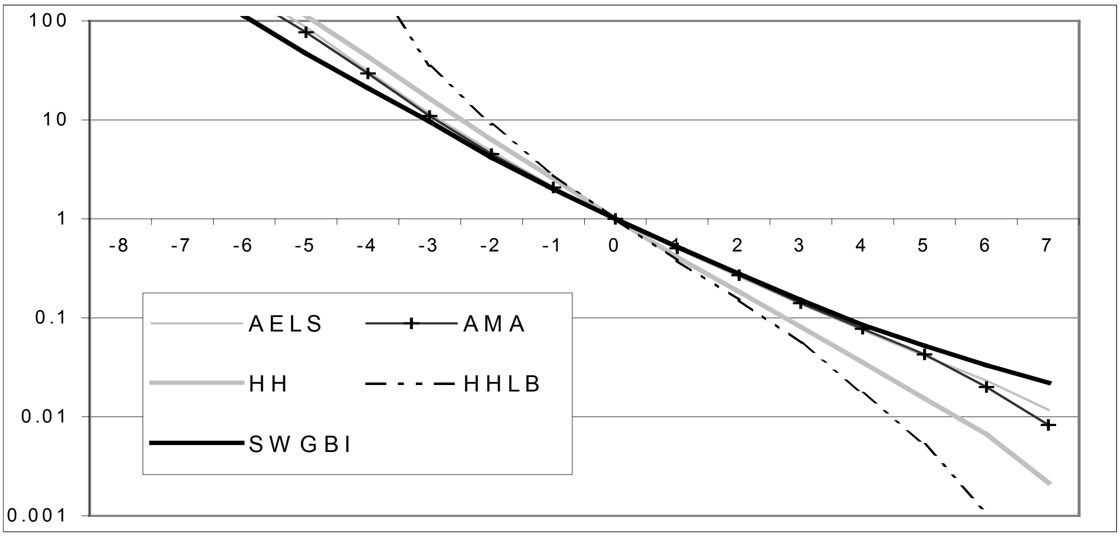

Diagram 4.3 **Demeaned Omegas for a selection of MSCI relative indices.**

Tracking Error

As Tracking Error is now a popular measure of portfolio performance we shall also report the rank ordering of these MSCI relative portfolios by Tracking Error[9] and by the Sharpe measure and Omega at threshold zero.

	Tracking Error	Ω(0)	Sharpe
HHLB	1	12	11
HH	2	1	1
HED	3	10	9
AED	4	5	5
ADIST	5	8	8
HM	6	2	3
AELS	7	11	12
HDIST	8	14	14
AMA	9	4	4
HMA	10	13	13
AM	11	6	6
ARV	12	9	10
ASS	13	3	2

HCSA	14	15	15
ACSA	15	7	7
HMN	16	16	16
SWGBI	17	17	17

Notice that the only points of agreement between the rank ordering by Tracking Error by Sharpe and by Omega, which incorporates all the higher moment effects, are for the two poorest performing portfolios HMN and SWGBI. We report the Kendall and Spearman rank correlation statistics below:

Kendall's rank correlation coefficient :

	Tracking Error	Ω	Sharpe
Tracking Error	1	0.3088	0.3235
Ω	0.3088	1	0.9559
Sharpe	0.3235	0.9559	1

Spearman's rank correlation coefficient :

	Tracking Error	Ω	Sharpe
Tracking Error	1	0.4093	0.4289
Ω	0.4093	1	0.9926
Sharpe	0.4289	0.9926	1

Here, though the principal point to note is only that Tracking Error is poorly correlated with either Sharpe or Ω performance measures, it is hard not to conclude that Tracking Error is a very poor performance measurement criterion or tool.

The final illustration is the time evolution of the cumulative Omega measures of the MSCI and SWGBI, at risk threshold zero, and for comparison their Sharpe analogues at zero risk free, which is shown below as Diagram 4.4. Here it is evident that the Sharpe measure captures much, but no means all, of the value evolution.

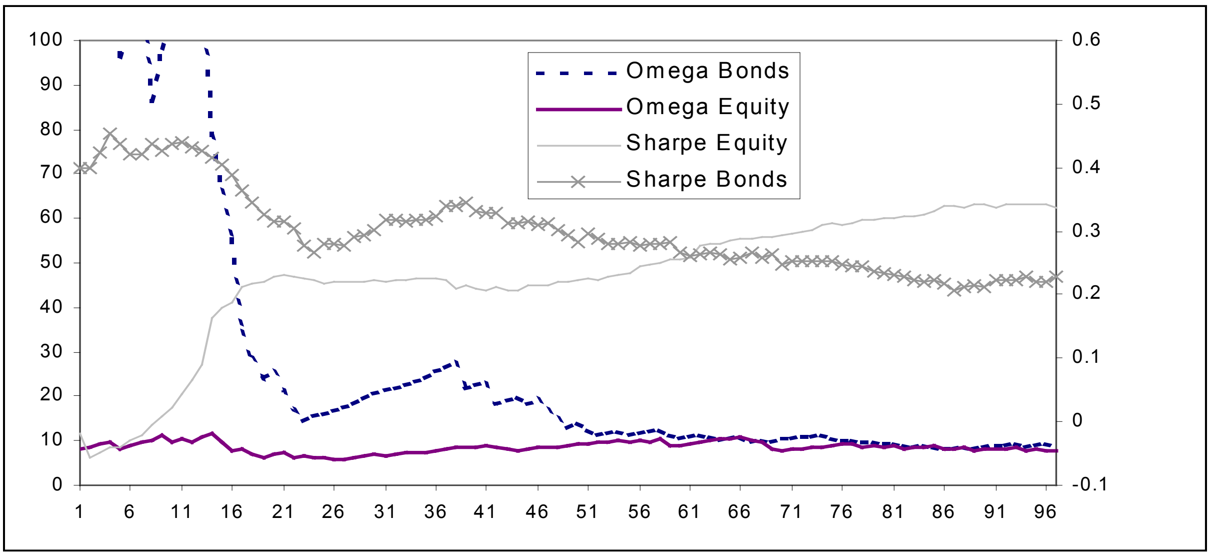

Diagram 4.4: A Comparison of the Cumulative Omega and Sharpe Measures

Table 01	ACSA	ADIST	AED	AELS	AM	AMA	ARV	ASS	HCSA	HIST	HH	HMN	HHLB	HED	HM	HMA	MSCI	SWGBI
Number of observations	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Missing values	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Sum of weights	101	101	101	101	101	101	101	101	101	101	101	101	101	101	101	101	101	101
Minimum	-2.65	-7.07	-6.71	**-2.49**	-4.15	-5.28	-1.95	**-15.08**	-4.95	-7.91	-7.92	-2.44	-10.82	-7.82	-6.62	-5.50	-13.45	-4.07
Maximum	3.32	5.25	4.07	3.50	9.21	3.25	2.62	**11.17**	3.33	5.76	10.07	4.14	8.53	8.59	9.84	**2.91**	8.91	7.42
Mean	1.32	1.24	1.32	1.14	1.37	1.37	1.25	1.44	0.94	0.94	**1.38**	0.63	1.06	1.20	1.36	1.03	0.90	**0.55**
Standard dev.	**0.98**	1.76	1.36	0.93	2.91	1.09	0.73	3.57	1.28	1.56	2.61	1.16	3.76	1.85	2.88	1.00	**3.89**	1.97
1st quartile	**0.92**	0.15	0.54	0.60	-0.66	0.94	0.75	-0.79	0.52	0.25	-0.44	0.04	-0.94	0.06	-0.40	0.60	**-1.45**	-0.73
Median	1.39	1.28	1.50	1.12	1.20	**1.51**	1.34	1.47	1.21	1.07	1.40	0.65	1.44	1.41	0.84	1.17	1.51	**0.36**
3rd quartile	2.00	2.29	2.00	1.79	3.28	1.97	1.77	3.59	1.65	1.70	2.77	**1.33**	**3.54**	2.33	3.35	1.67	3.31	1.77
Kurtosis	2.11	4.30	11.10	1.40	-0.40	**12.66**	2.86	3.98	5.35	11.63	2.18	**0.39**	0.85	6.69	0.69	17.06	0.99	1.17
Skewness	-0.87	-0.97	-1.96	-0.45	0.32	-2.43	-1.01	-0.83	-1.73	-1.84	0.16	-0.16	-0.75	-0.48	**0.45**	**-2.91**	-0.72	0.43
Std dev of mean	0.10	0.17	0.14	0.09	0.29	0.11	0.07	0.36	0.13	0.16	0.26	0.12	0.37	0.18	0.29	0.10	0.39	0.20
Zero Sharpe	1.35	0.71	0.97	1.22	0.47	1.25	**1.71**	0.40	0.74	0.60	0.53	0.54	0.28	0.65	0.47	1.03	**0.23**	0.28

Table02	ACSA	ADIST	AED	AELS	AM	AMA	ARV	ASS	HCSA	HDIST	HH	HMN	HHLB	HED	HM	HMA	MSCI	SWGBI
ACSA	1.00	0.64	0.67	0.55	0.48	0.57	0.83	0.44	0.60	0.58	0.46	0.47	0.30	0.50	0.37	0.41	0.24	-0.10
ADIST	0.64	1.00	0.93	0.57	0.58	0.74	0.72	0.61	0.43	0.68	0.65	0.33	0.53	0.61	0.58	0.56	0.47	-0.05
AED	0.67	0.93	1.00	0.63	0.53	0.91	0.72	0.63	0.52	0.74	0.68	0.35	0.56	0.69	0.55	0.71	0.49	-0.07
AELS	0.55	0.57	0.63	1.00	0.44	0.58	0.80	0.53	0.38	0.43	0.64	0.52	0.50	0.48	0.45	0.46	0.48	0.06
AM	0.48	0.58	0.53	0.44	1.00	0.35	0.55	0.54	0.30	0.32	0.60	0.12	0.46	0.40	0.68	0.26	0.45	0.12
AMA	0.57	0.74	0.91	0.58	0.35	1.00	0.59	0.54	0.51	0.65	0.57	0.34	0.46	0.67	0.46	0.77	0.42	-0.05
ARV	0.83	0.72	0.72	0.80	0.55	0.59	1.00	0.53	0.51	0.62	0.57	0.55	0.43	0.57	0.51	0.45	0.36	-0.05
ASS	0.44	0.61	0.63	0.53	0.54	0.54	0.53	1.00	0.47	0.52	0.59	0.05	0.53	0.52	0.56	0.45	0.50	-0.18
HCSA	0.60	0.43	0.52	0.38	0.30	0.51	0.51	0.47	1.00	0.53	0.46	0.33	0.37	0.52	0.38	0.38	0.25	-0.05
HDIST	0.58	0.68	0.74	0.43	0.32	0.65	0.62	0.52	0.53	1.00	0.52	0.23	0.50	0.69	0.48	0.59	0.41	-0.10
HH	0.46	0.65	0.68	0.64	0.60	0.57	0.57	0.59	0.46	0.52	1.00	0.32	0.86	0.68	0.63	0.48	0.68	0.07
HMN	0.47	0.33	0.35	0.52	0.12	0.34	0.55	0.05	0.33	0.23	0.32	1.00	0.23	0.30	0.21	0.24	0.20	0.15
HHLB	0.30	0.53	0.56	0.50	0.46	0.46	0.43	0.53	0.37	0.50	0.86	0.23	1.00	0.65	0.53	0.45	0.79	0.08
HED	0.50	0.61	0.69	0.48	0.40	0.67	0.57	0.52	0.52	0.69	0.68	0.30	0.65	1.00	0.60	0.61	0.53	-0.02
HM	0.37	0.58	0.55	0.45	0.68	0.46	0.51	0.56	0.38	0.48	0.63	0.21	0.53	0.60	1.00	0.35	0.49	0.04
HMA	0.41	0.56	0.71	0.46	0.26	0.77	0.45	0.45	0.38	0.59	0.48	0.24	0.45	0.61	0.35	1.00	0.41	-0.04
MSCI	0.24	0.47	0.49	0.48	0.45	0.42	0.36	0.50	0.25	0.41	0.68	0.20	0.79	0.53	0.49	0.41	1.00	0.24
SWGBI	**-0.10**	**-0.05**	**-0.07**	**0.06**	**0.12**	**-0.05**	**-0.05**	**-0.18**	**-0.05**	**-0.10**	**0.07**	**0.15**	**0.08**	**-0.02**	**0.04**	**-0.04**	**0.24**	**1.00**

Table 03	factor 1	factor 2	factor 3	factor 4	factor 5	factor 6	factor 7	factor 8	factor 9	factor 10	factor 11	factor 12	factor 13	factor 14	factor 15	factor 16	factor 17	factor 18
ACSA	0.73	**-0.44**	0.20	0.23	0.00	0.12	0.06	-0.22	-0.08	0.20	0.10	0.06	0.08	0.21	-0.02	-0.03	-0.07	0.00
ADIST	0.86	-0.10	-0.07	0.07	0.25	-0.10	-0.08	-0.17	-0.15	-0.24	-0.03	0.05	-0.05	-0.05	0.20	-0.05	-0.04	-0.04
AED	0.91	-0.16	-0.12	-0.09	0.22	-0.09	0.00	-0.06	-0.15	-0.13	-0.06	0.00	0.04	-0.02	-0.01	0.00	0.00	0.09
AELS	0.75	-0.05	0.34	0.06	-0.10	**-0.36**	0.20	0.13	0.22	0.03	-0.20	-0.10	0.08	-0.07	0.00	-0.03	-0.10	0.00
AM	0.64	0.30	0.07	**0.54**	0.28	0.04	-0.03	0.04	-0.16	0.20	0.08	-0.02	-0.07	-0.18	-0.10	0.03	-0.02	-0.01
AMA	0.82	-0.22	-0.15	-0.30	0.22	-0.06	0.05	0.14	-0.12	-0.08	-0.05	0.01	0.18	0.03	-0.18	0.08	0.03	-0.04
ARV	0.82	-0.29	0.28	0.23	-0.01	-0.12	0.02	-0.13	0.16	0.10	-0.03	-0.02	0.03	-0.02	0.09	0.08	0.16	0.00
ASS	0.73	0.11	-0.32	0.27	-0.05	-0.06	0.35	0.12	0.17	-0.16	0.11	0.22	-0.14	0.07	-0.04	0.01	0.01	0.00
HCSA	0.63	-0.26	0.00	0.02	-0.27	**0.57**	0.27	0.15	-0.12	-0.06	-0.03	-0.11	0.03	-0.09	0.07	-0.01	0.00	0.00
HDIST	0.77	-0.19	-0.23	-0.15	0.02	0.19	-0.13	-0.29	0.30	-0.07	0.02	-0.16	-0.14	-0.05	-0.13	-0.02	-0.02	-0.01
HH	0.83	0.35	0.04	0.00	-0.20	-0.05	-0.03	-0.01	-0.18	0.05	-0.22	-0.02	-0.11	0.09	-0.08	-0.17	0.07	-0.01
HMN	0.44	-0.29	**0.69**	-0.16	-0.23	-0.08	-0.21	0.15	-0.08	-0.15	0.19	0.07	-0.13	-0.04	-0.06	0.01	0.00	0.00
HHLB	0.73	**0.49**	-0.05	-0.16	-0.31	-0.05	-0.02	-0.14	-0.13	0.03	-0.06	-0.04	-0.08	0.06	0.04	0.20	-0.04	0.00
HED	0.80	0.05	-0.14	-0.20	-0.12	0.19	-0.25	0.04	0.14	0.17	-0.10	0.31	0.10	-0.11	0.03	-0.02	-0.02	0.00
HM	0.70	0.29	-0.04	**0.31**	0.08	0.16	-0.34	0.30	0.15	-0.11	0.03	-0.13	0.08	0.15	0.05	0.01	-0.01	0.01
HMA	0.69	-0.14	-0.22	-0.45	0.19	-0.11	0.05	0.23	0.00	0.28	0.14	-0.09	-0.16	0.03	0.10	-0.02	0.00	0.00
MSCI	0.65	**0.58**	0.04	-0.18	-0.13	-0.09	0.11	-0.15	0.02	-0.03	0.31	-0.05	0.20	-0.05	0.01	-0.08	0.02	0.00
SWGBI	**-0.01**	**0.45**	**0.61**	-0.28	**0.43**	**0.31**	0.17	-0.05	0.11	-0.03	-0.08	0.06	-0.05	0.05	0.01	0.01	0.00	0.00

Table04	ACSA	ADIST	AED	AELS	AM	AMA	ARV	ASS	HCSA	HDIST	HH	HMN	HHLB	HED	HM	HMA	SWGBI
Observations	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Missing values	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Sum of weights	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Minimum	-9.42	-8.92	-7.84	-7.79	-8.01	-7.03	-9.79	-7.32	-9.6	-9.75	-6.75	-10.42	-4.35	-8.16	-10.39	-7.78	-8.4
Maximum	10.8	11.42	10.24	10.96	13.35	9.98	11.5	9.76	11.52	10.39	9.01	12.07	6.76	8.64	10.74	8.89	17
Mean	0.409	0.332	0.407	0.230	0.427	0.458	0.334	0.521	0.040	0.045	0.486	-0.273	0.162	0.304	0.467	0.133	-0.351
Standard dev.	3.80	3.459	3.451	3.561	3.667	3.589	3.710	3.754	3.800	3.565	2.860	3.857	2.494	3.326	3.538	3.621	3.945
1st quartile	-2.44	-2.04	-1.98	-2.13	-2.29	-2.20	-2.27	-2.08	-2.64	-2.34	-0.97	-2.74	-1.68	-2.2	-1.6	-2.74	-2.99
Median	-0.14	0.165	0.16	-0.13	0.575	0.275	-0.065	-0.05	-0.1	0.27	0.29	-0.74	0.06	0.11	0.59	-0.14	-0.88
3rd quartile	2.80	2.485	2.737	2.28	2.49	2.83	2.54	2.79	1.97	2.25	1.90	2.37	1.63	2.66	2.22	2.48	1.91
Kurtosis	0.253	0.619	0.256	0.407	1.277	-0.019	0.492	-0.330	0.683	0.389	0.536	0.675	-0.344	-0.092	1.251	-0.242	2.931
Skewness	0.50	0.446	0.480	0.595	0.390	0.491	0.466	0.160	0.499	0.189	0.285	0.495	0.377	0.127	-0.016	0.326	0.986
Std dev of mean	0.38	0.345	0.345	0.356	0.366	0.358	0.371	0.375	0.380	0.356	0.286	0.385	0.249	0.332	0.353	0.362	0.394

Descriptive Statistics: Indices relative to MSCI

Conclusions and Further Work

We have introduced a simple measure of performance which is both natural from the standpoint of probability and statistics and heuristically appealing in its financial interpretation. It is defined in the most basic of terms but captures all higher moment information in a distribution of returns. It is broadly in the spirit of the downside and related literature, but could also be related to the stochastic dominance and decision literature. There is existing work, by Liang Zou, which approaches downside type measures axiomatically[10].

We have applied this to a set of hedge fund index returns. We accept that these returns sequences have survivor and other biases present but for pedagogic purposes they suffice. The results, based on the simplest of decision rules, namely that we prefer more to less, show a markedly different order of preference from more traditional measures such as the Sharpe ratio or tracking error. In the case of the Sharpe ratio, this difference arises from the additional higher moment information which Omega captures. The presence of such effects in real data is, we trust, convincing evidence for the improvements in performance measurement which the Omega function provides.

We have also demonstrated, with examples from analytic probability distributions, that Omega is a powerful tool for the capture of higher moment effects. A number of the most important basic properties of Omega have been stated. The affine invariance of Ω allows comparisons to be made in a way that is independent of scaling and translations of the underlying returns or equivalently, of the risk threshold.

The canonical nature of Ω also provides for additional performance measures to be induced from more general transformations of the returns distribution. This is equivalent to the use of alternative utility functions encoding risk preferences or tolerances. The induced Omega function may be used to provide a consistent performance ranking for each such risk adjustment. This is a subject for further investigation.

In this paper we have not considered in any detail, either time serial behaviour or portfolio optimisation; these will be the subject of further, later papers. A more advanced analysis of the properties of Ω and some of its statistical characteristics is the subject of Cascon, Keating and Shadwick 2001.

The most obvious further requirement is for a more advanced technique for the estimation of stationarity which explicitly considers the higher moments. The ideal would provide some indication of the likelihood of stationarity based upon prior arrivals, but this is a non-trivial affair. It seems likely that a frequency domain analysis of Ω would provide some useful insights.

The gain-loss literatures, such as Bernardo and Ledoit[11], already provide some insights as to how Omega might be used in asset pricing. An unpublished work of Agarwal and Naik extends that literature to optimal asset allocation[12].

The overwhelming lesson from modern finance is that the state of the economy in which a return is received is a prime determinant of its value[13]. It is comparatively trivial to revert to the more familiar ground of a fixed risk free rate or even the returns from some passive index as a proxy for the wealth and consumption capacity[14], as we have illustrated in the index-relative case earlier. This introduction is limited to the simple univariate case while for most pricing and portfolio applications consideration of the multivariate case is required. This is the subject of further, later papers.

The most obvious extension is to performance attribution. We might have followed Hicks and Marschak[15] in the observation that preferences are a function of all of the moments of a returns distribution and we have demonstrated earlier why that might be rational choice. We might then have simply noted that these are a function of the moment generating function of the returns distribution and in turn the characteristic and cumulative density functions. The obvious extension from there is to the frequency domain and examination of the spectra of the returns series, where the limiting requirement of most techniques applied is only covariance stationarity.

Perhaps the most interesting hypothesis to be investigated is that the activity in funds where higher moments are significant is directly related to the sale of liquidity and that will take us into the monetary economics, and possibly the causal identification of the factors driving performance.

One further avenue for investigation is that of behavioural finance, where, with the higher moments now accounted for, we might investigate by way of the penalty function the nature of some of the irrationality they claim.

Appendix A: Omega and Markowitz Frontiers

For two returns series, A and B, which are negatively correlated, we show the frontier achieved by weights varying from [0% A, 100% B] to [100% A, 0% B] in 5% steps: As is usual with these diagrams, the vertical axis is return (in %) and the horizontal standard deviation of returns (in %). The square is the optimal portfolio allocation given a risk-free return of 0%, a mixture of 50% A and 50% B. The Ω optimal, the circle, is a significantly different asset allocation, (35% A and 65% B).

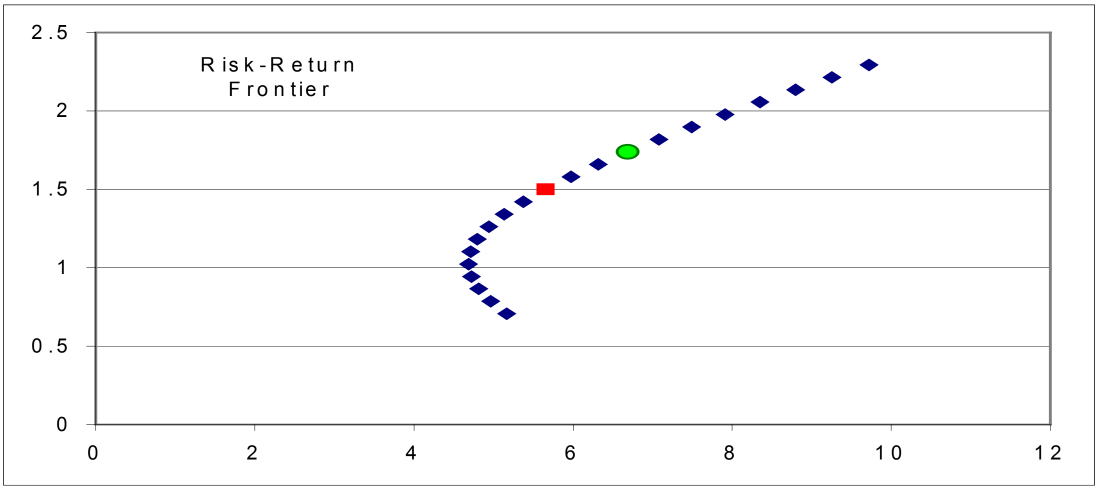

Diagram A1 Risk –Return Frontier

We now present the Ω analogue of this diagram:

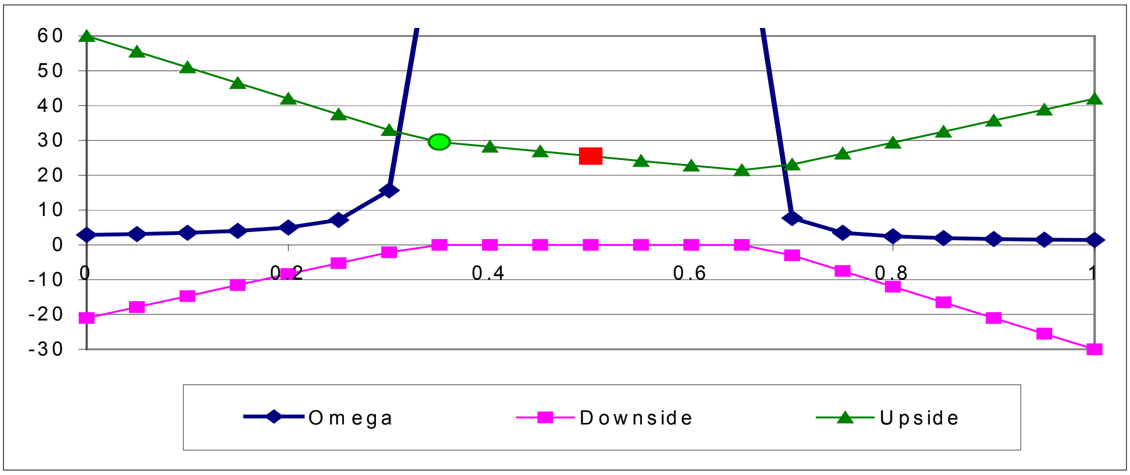

Diagram A2 Omega at 0 threshold

The vertical axis is the value of the upside or downside, I_2 or I_1 earlier. The horizontal axis is the proportion of asset A in the portfolio mix. The dotted triangular marked line is the upside value for differing mixes. The continuous square marked line is the downside value for differing mixes. The heavy lozenge marked line is the Ω function. Note that this is infinite at and between portfolio mixes of 35% and 65% A. In this range the effects of diversification mean that the portfolio has no downside, and the downside function is zero in this range. Using our preference for more rather than less, we should therefore prefer the portfolio 35% A and 65% portfolio B, marked as a circle, rather than the Markowitz optimal portfolio of 50% A and 50% B, marked as a square. We are in fact choosing among "free lunches" as there is no downside risk present.

Notice also that the traditional risk-return framework does not suggest at any point that the portfolio has no (downside) risk. The risk as measured by the standard deviation of returns is always positive.

Asset allocation and portfolio optimisation will be the subject of another later paper.

Appendix B

We illustrate the effect of higher moments on Ω in this appendix. For this purpose we have used analytic distributions constructed from linear combinations of normal distributions.

We first consider the behaviour of the Omega function for normal distributions as their variance changes at a common mean. Diagram B.1 shows Ω for three normals, of mean zero, with standard deviations of 5,10 and 15, shown dotted, solid and dashed respectively. The reversal of preferences across the mean is the effect of variance. On the upside, increased variance provides more chance of gain, while on the downside it provides , symmetrically, more chance of loss. Thus, the smaller the variance, the steeper the slope of the Omega function.

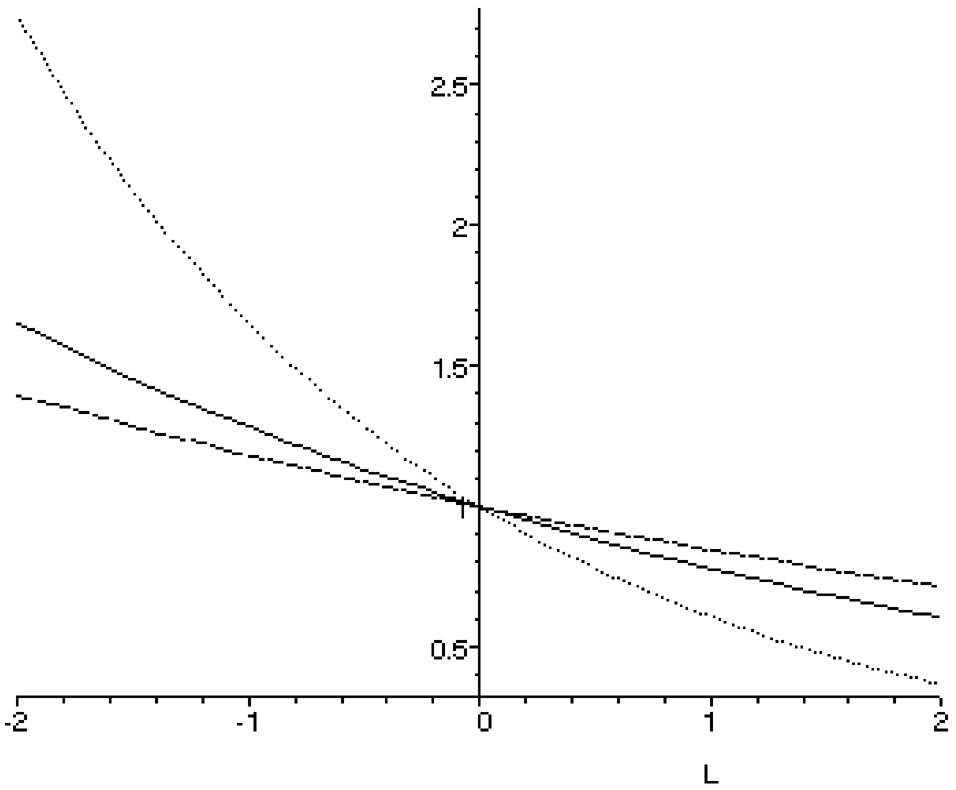

Diagram B.1: Ωs for normals of mean zero and variance 5,10 and 15

We next consider the effects of skew and kurtosis. More precisely, as the Omega function responds to the effects of all moments, we can illustrate the effects of third and higher moments and fourth and higher moments.

First we make a comparison of the Omega measures for distributions whose kurtosis is the same as a normal with the same mean and variance. This illustrates the impact of skew and moments of fifth and higher orders.

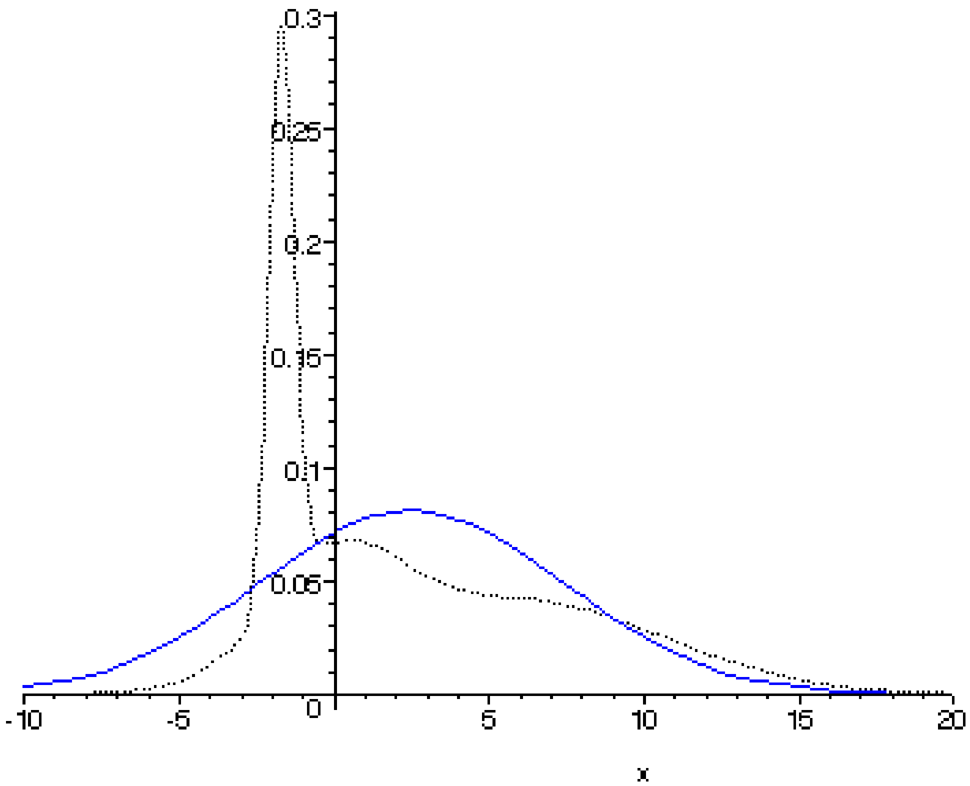

Diagram B.2: Two distributions with the same mean, variance and kurtosis

The skewed distribution in diagram B.2 has the same mean, 2.5, variance, 24 and kurtosis, 3 as the normal distribution shown by the solid curve. The skewness is 0.86. Thus this distribution differs from the normal only in its skewness and fifth and higher moments.

The Ωs corresponding to the these distributions are shown in Diagrams B.3 and B.4. While there is a separation in the Ω curves away from the mean, the differences around the mean are small. The shapes of the curves differ strongly even here however as the derivatives of Ω with respect to L, illustrated in Diagrams B.5 and B.6 indicate.

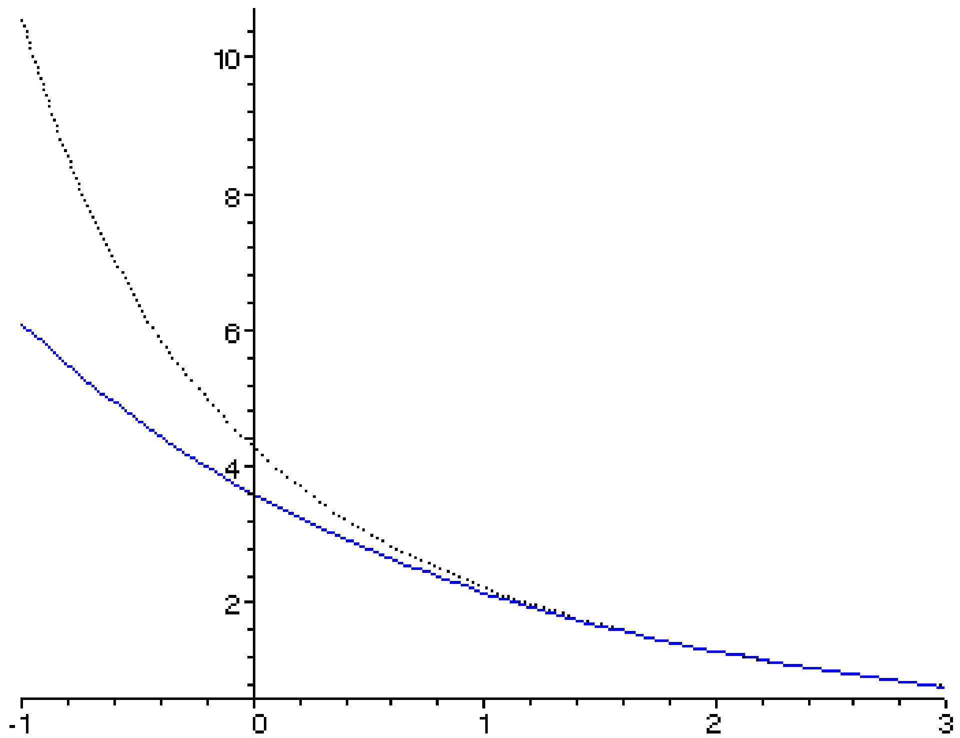

Diagram B.3: Ωs as functions of r for the distributions in Diagram B.2

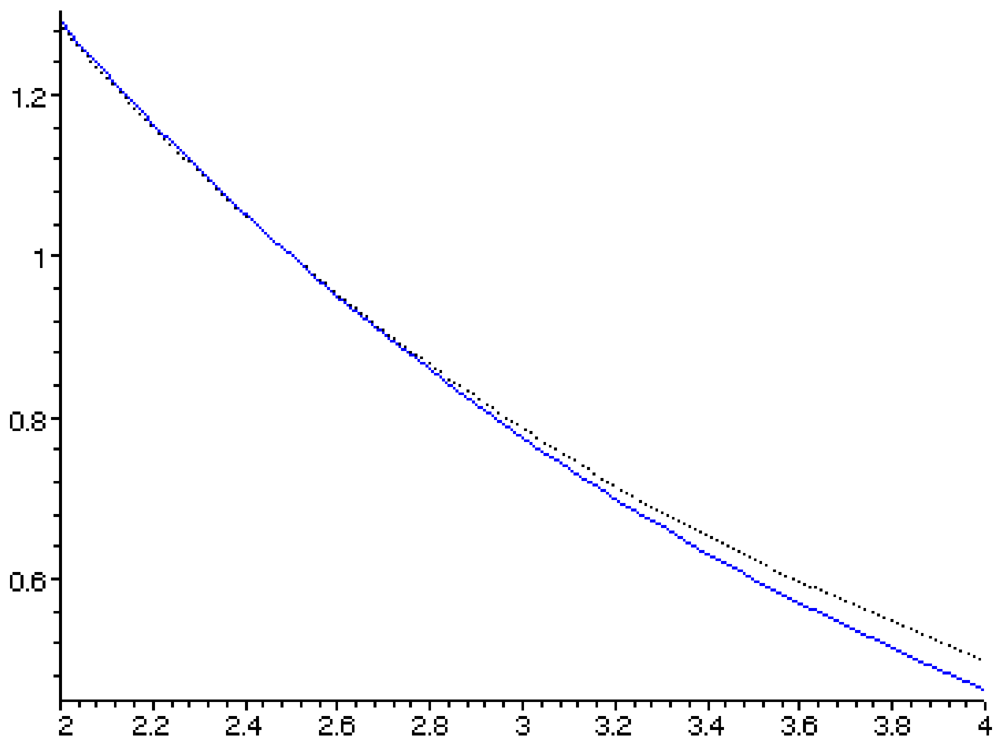

Diagram B.4: Ωs as functions of r for the distributions in Diagram B.2

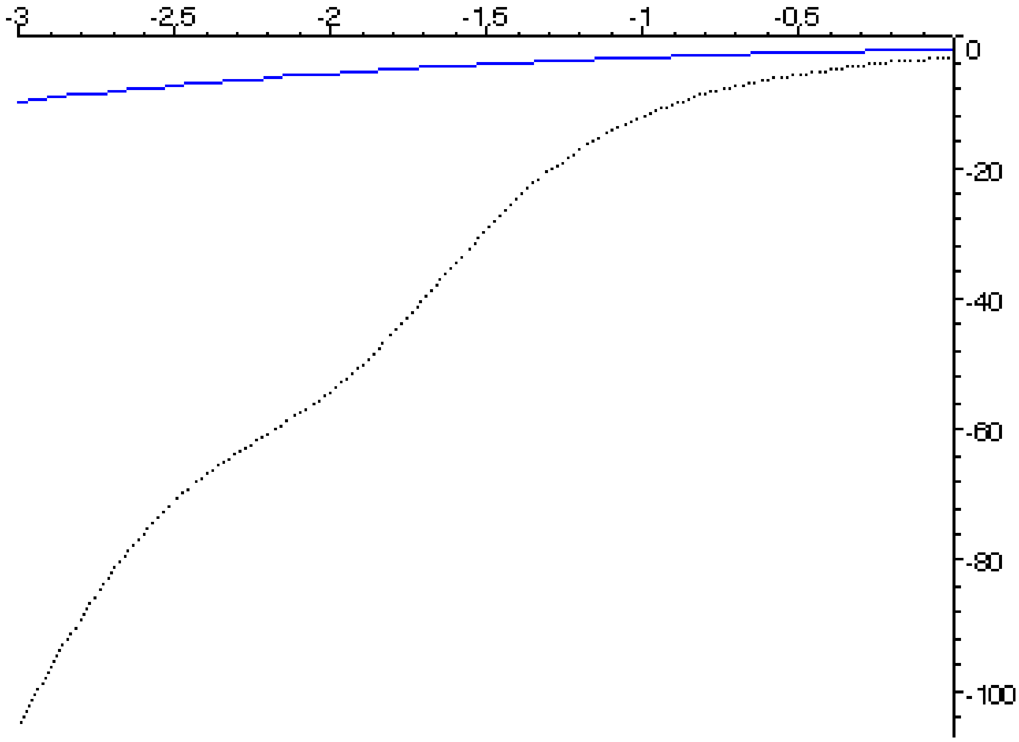

Diagram B.5: First derivatives of Ω with respect to r

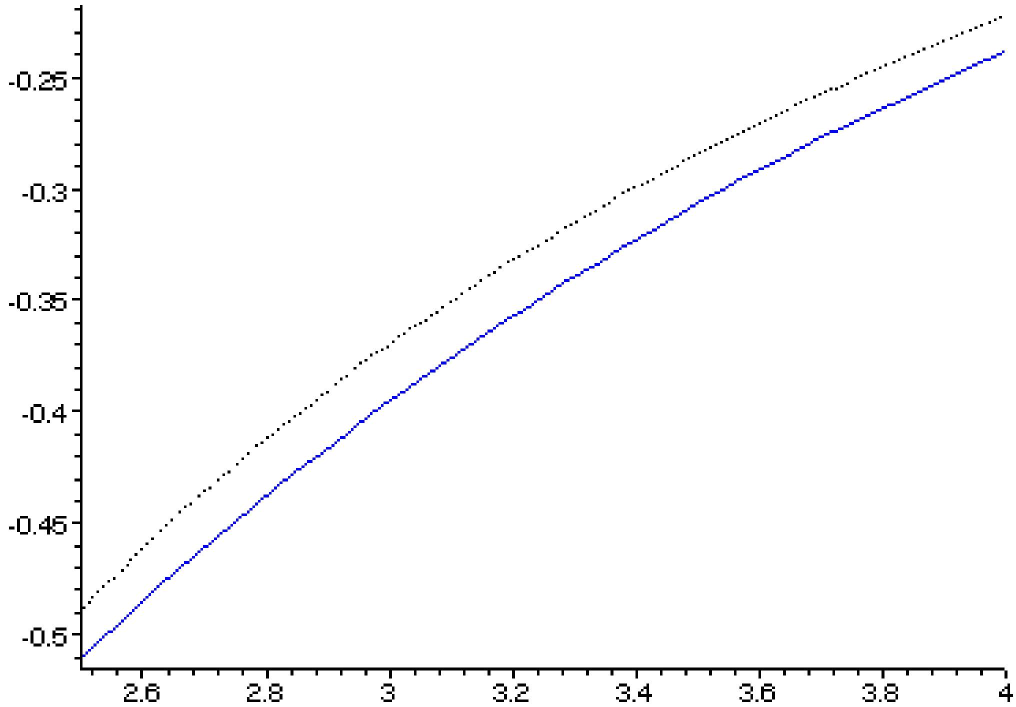

Diagram B.6: First derivatives of Ω with respect to r.

The second derivative behaviour is even more markedly different for the skewed distribution.

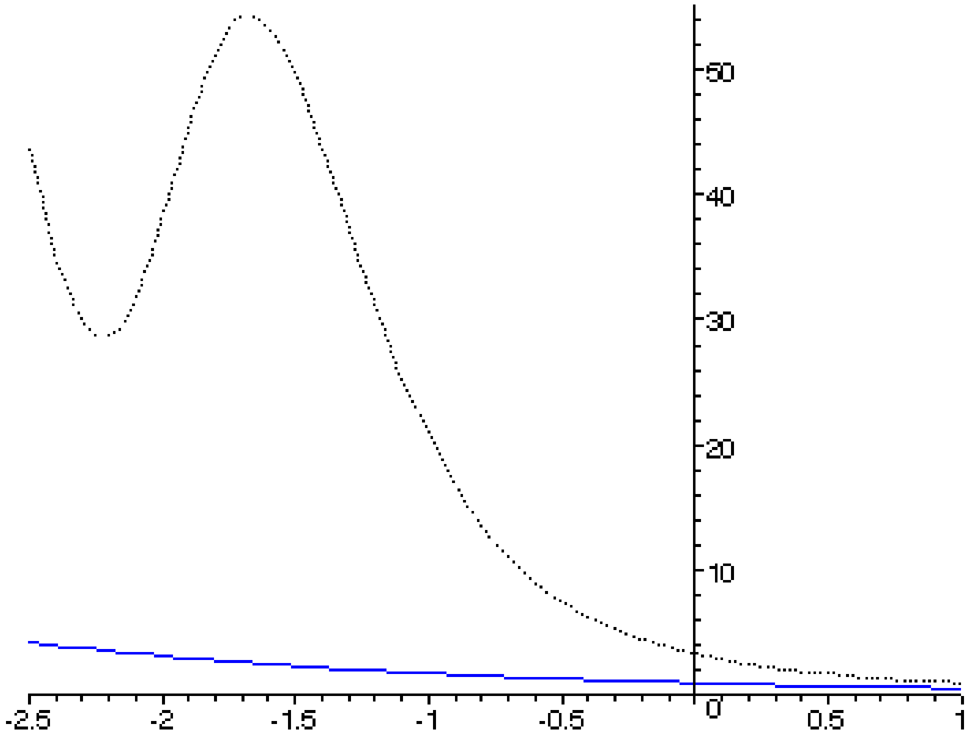

Diagram B.7: Second derivatives of Ω with respect to r,

Finally we make a comparison of the Ωs for three symmetric distributions with the same mean and variance and kurtosis of 3, 5.9 and 11.8. The differences in the Ωs are therefore produced by kurtosis and by sixth and higher even moments.

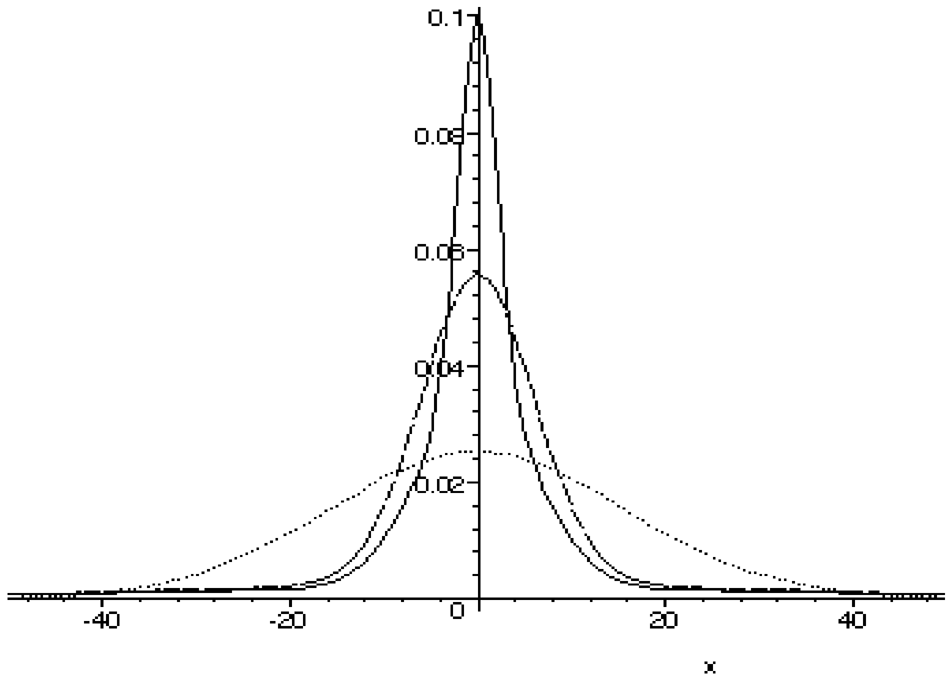

Diagram B.8: Symmetric distributions differing in kurtosis and higher even moments

As diagrams B 9,B.10 and B.11 show, although on very different scales, the Omega measure displays significant differences due to these higher moment effects. The asymptotic rankings are as indicated in B.9 and B.11, corresponding to the higher probability of large losses and gains increasing with kurtosis. The higher kurtosis Ωs each cross the Ω for the normal distribution in three places and themselves cross three times, at +/-20 and at their common mean of 0.

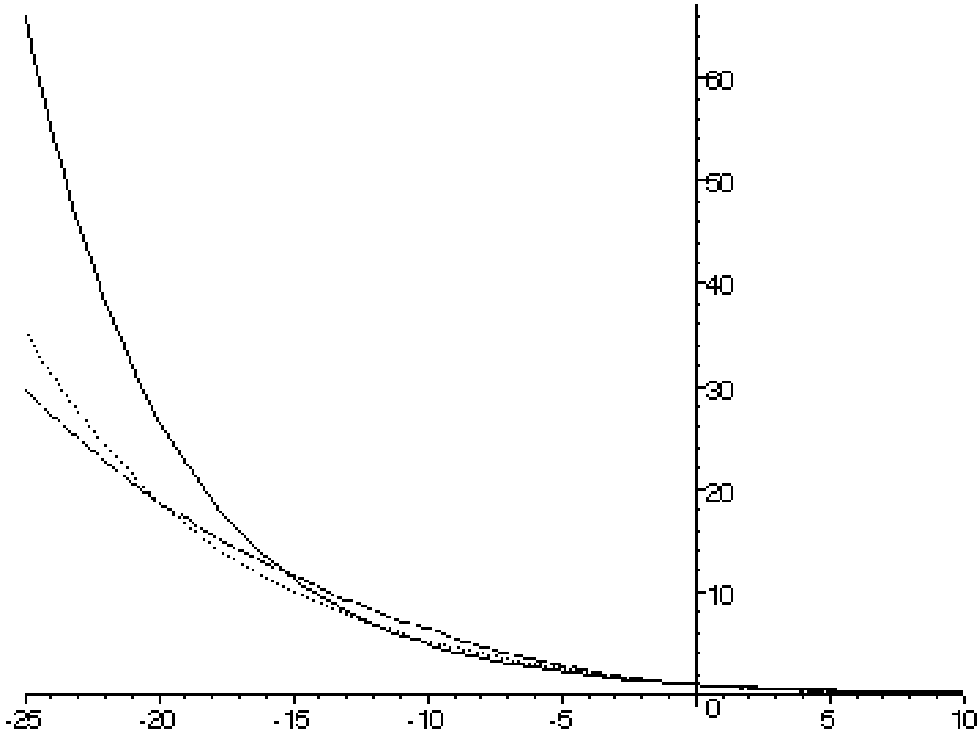

Diagram B.9: Ωs as functions of r for the distributions of Diagram B.8

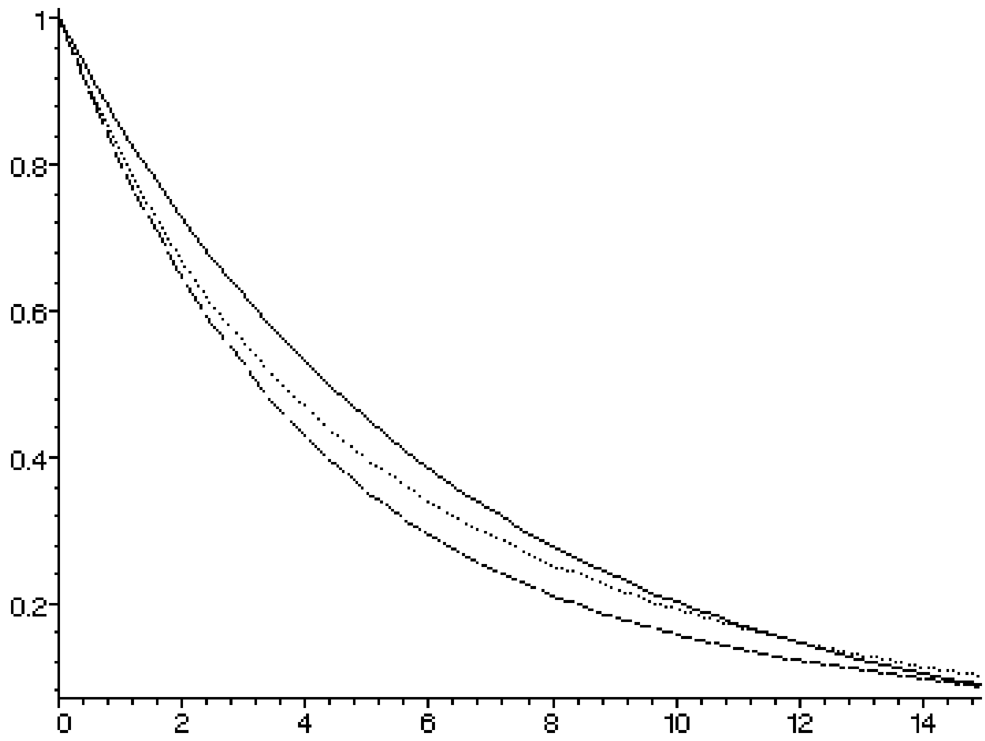

Diagram B.10: Ωs as functions of r for the distributions of Diagram B.8

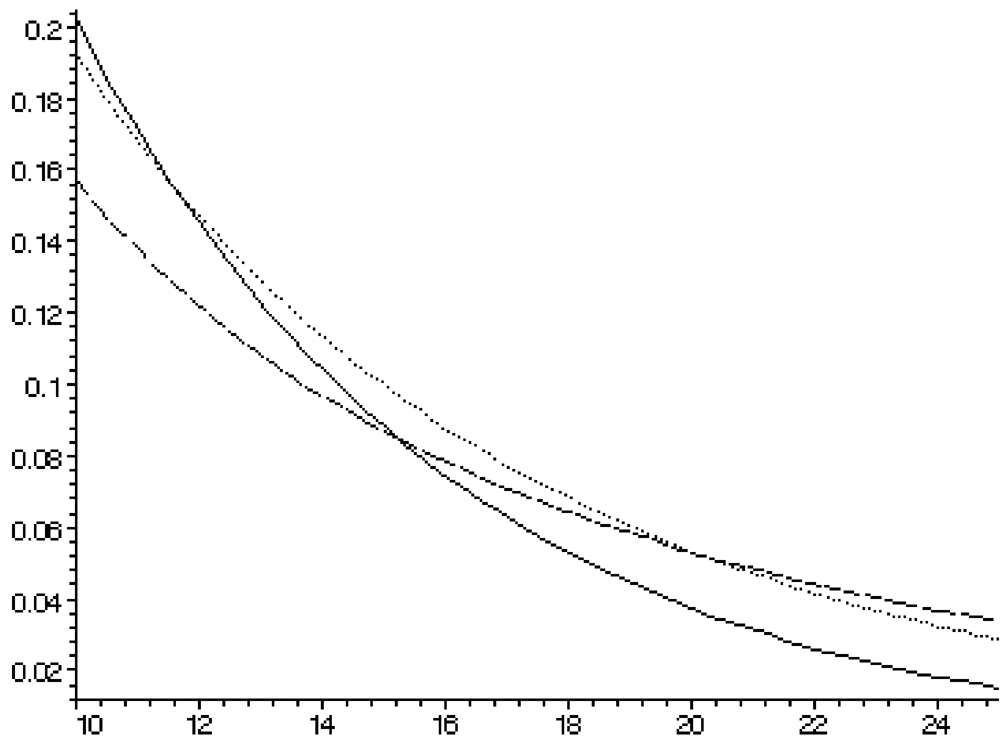

Diagram B.11: Ωs as functions of r for the distributions of Diagram B.8

The second derivative behaviour here may be contrasted with that in Diagram B.7. Although the two cases have large differences in variance, the effect seen here is primarily due to skew and higher odd moments.

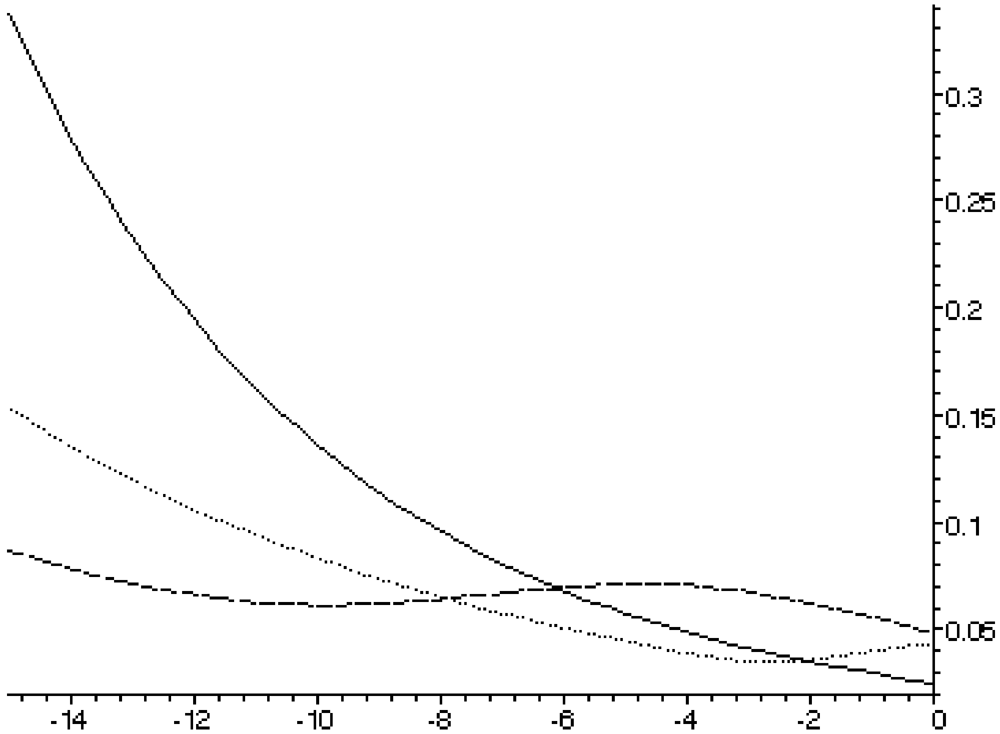

Diagram B.12: Second derivatives of Ω with respect to r for the distributions of Diagram B.8.

[1] The absence of a utility function here may, at first sight, be disconcerting. However, in order to *rank* portfolios over an interval of possible returns, all that is needed is a comparison of the magnitudes of their omegas over that interval. For example, if asset A's Omega is larger than asset B's over an interval, we should prefer asset A in that range of returns. It is this application which we explore in the present paper. If we wish to quantify the difference between the two assets on the other hand, we must introduce additional structure which can, for example, decide how much better is an omega of 2 than an omega of 1.5. A utility function is the obvious way to do this.

[2] This example is a less dramatic variant on the choice between buying a lottery ticket for £1 with a one in a million chance of winning £1million and selling a lottery ticket for £1 with a one in a million chance of having to pay out £1million. We have, to date, found no one who regards these two as equally attractive.

[3] A. Cascon, C Keating and W. Shadwick "Properties of the Omega Measure" Preprint Finance Development Centre 2001

[4] A Cascon, C. Keating and W Shadwick 2001 op cit

[5] The commonly used tracking error measure, the square root of the variance of the difference portfolio, assumes a common location or mean for the portfolio and its comparator index. More generally this should be the square root of the sum of the mean squared plus the variance and even then we need the value of the mean or at least its sign to know whether the manager is adding or subtracting value relative to the passive. The Omega measure or some risk adjusted counterpart centred on the passive index give us direct measures of relative performance.

[6] For a fuller discussion of hedge fund indices see: Brooks C. and Kat H. "The Statistical Properties of Hedge Fund Index Returns and Their Implications for Investors" Working Paper, ISMA Centre, University of Reading October 2001.

[7] Five tests were utilised: A Stochastic Dominance Test, Wilcoxon-Mann-Whitney, Student's T, a known variance Z, and Fisher's F. Full details are available from the authors on request.

[8] There is some evidence that traditional fund managers are biased towards positive skewness, which would explain, beyond dealing costs, their average apparent inability to outperform benchmark indices. The intuition here may be that they buy fewer of the high risk investments that are present in the market. See for example: F.D. Arditti "Another Look at Mutual Fund Performance" Journal of Financial and Quantitative Analysis, June 1971.

[9] The tracking error calculated here is the popular version, the standard deviation of the difference portfolio.

[10] Liang Zou "Dichotomous Theory of Choice under Risk I: Basic Framework and Axiomatic Foundation" Working Paper, Tinbergen Institute, University of Amsterdam 2002
see also www.tinbergen.nl/discussionpapers/00050.pdf and www.tinbergen.nl/discussionpapers/00108.pdf

[11] Bernardo and Ledoit (see below) consider a gain / loss function, which is defined as the expected positive excess returns divided by the expected negative excess returns under some risk adjusted probability measure. This suffers from the difficulty that their ratio $\frac{E^*[\tilde{x}^+]}{E^*[\tilde{x}^-]}$ is not continuous and that the derivatives do not therefore exist. This is not a problem for the measure that we shall settle upon.
A.E. Bernardo and O. Ledoit "Gain, Loss and Asset Pricing" Journal of Political Economy, 2000 Vol. 8 No 1 pp 144 – 172

[12] V. Agarwal and N. Naik, "Does Gain-Loss Analysis Outperform Mean-Variance Analysis? Evidence from Portfolios of Hedge Funds and Passive Strategies." Unpublished manuscript – London Business School November 1999

[13] This is just the classic statement , $p = \frac{E(x)}{R^f} + \frac{\text{cov}.[\beta u'(c_{t+1}), x_{t+1}]}{u'(c_t)}$ where $u'(c)$ is the marginal utility, that an asset's price is lowered if it covaries positively with consumption and of course the theoretical basis for insurance.

[14] This is in the spirit of Brown and Gibbons. See: S. Brown and M. Gibbons " A Simple Econometric Approach for Utility Based Asset Pricing Models" Journal of Finance 40, pp 359-382 1985. It should be recognised that if a transformation of returns is not affine or is not constant over time this will affect the descriptive statistics or properties of the returns distribution and consequently the Omega functions.

[15] Jacob Marschak "Money and The Theory of Assets" Econometrica 6 1938.
J.R. Hicks "Value and Capital" O.U.P. 1946