A Universal Performance Measure

Con Keating*
William F. Shadwick

THE FINANCE DEVELOPMENT CENTRE
LONDON

Current Version: January 2002

Abstract

We propose and examine a new performance measure in the spirit of the downside, lower partial moment and gain-loss literatures. We indicate how this may be applied across a broad range of problems in financial analysis. The performance measure captures all of the higher moments of the returns distribution. The performance measure is applied to a range of hedge fund style or strategy indices.

* Corresponding Author: Con Keating E-mail: keating@risk.demon.co.uk

We should like to thank Jim Aspinwall, Philip Best, Chris Brooks, Ana Cascon, Jon Danielsson, Chris Golden, 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 the over-simplification that mean and variance fully describe the distribution of returns. It is a generally accepted stylised fact of empirical finance that few, if any, would now challenge, 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 easy to illustrate the importance of higher moments. For example the two distributions in Diagram 1.1 have the same mean and variance, however they differ in skew, kurtosis and all higher moments and represent vastly different processes. Though normality is now the standard model for returns distributions, some historic series such as daily Gilts and Austrian government bonds were markedly bimodal, a feature which arose from the account settlement practice of those markets.

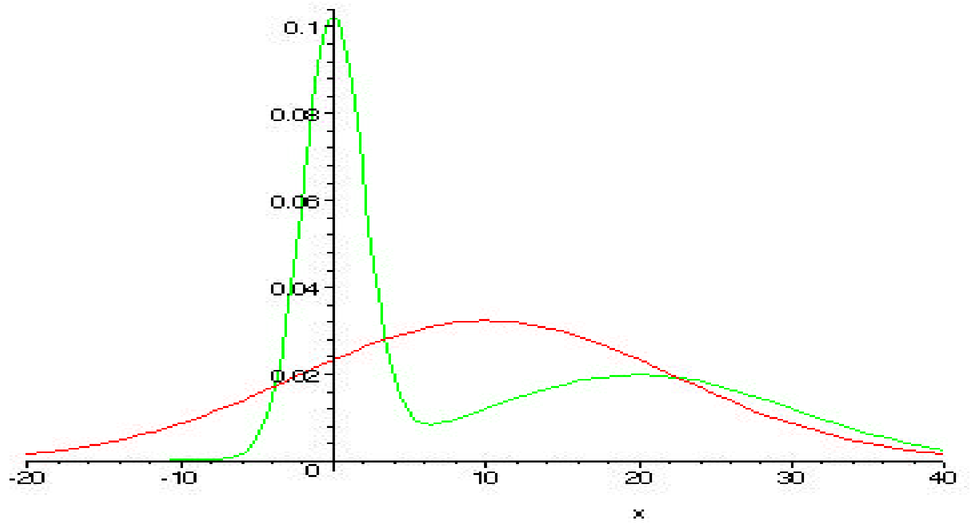

Diagram 1.1 Two distributions with mean 10 and variance 152.

While one would not now expect to see financial returns series for a single security following the bimodal distribution, portfolios can, of course, exhibit even more complex behaviour. It is the fact that we cannot know the precise distribution of real returns that makes the ability to deal with higher moments important. 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. The central problem is that the mean and variance toolkit is simply inadequate for such analysis. Some traditional asset "classes", such as corporate bonds, also require the consideration of higher moments in order to fully understand their performance characteristics.

By way of illustration, we shall show the effect of uncertainty and the higher moments of a returns distribution upon the value of a portfolio and demonstrate some of the shortcomings of traditional measures such as the Sharpe ratio.

In Table A below, in each case we are comparing an arbitrary 36 period set of returns. They have been chosen such that the expected value is the same for each but the most likely or terminal values differ by their higher moments. The difference between these terminal values and the expected value may be considered as effects due to uncertainty[1].

If we begin by preferring higher most likely terminal values, it is clear that we should rationally dislike the second and fourth moments[2] (variance and kurtosis) of the returns distributions in these examples as these moments unequivocally lower the most likely outcome. The third moment, skewness, increases the terminal, most likely value when positive and decreases it when negative.

The Sharpe ratio cited uses a risk-free rate of zero. Notice that it is the same in three instances even though terminal values differ markedly. Moreover in the case of negative skewness, the Sharpe ratio is better than in cases with higher most likely values[3].

These examples show that we should therefore expect, under typical market conditions, high returns[4] from securities or instruments which have high variance, high kurtosis and negative skewness as these all tend to increase risk. The impact of these higher moments is, of course, invisible to the Sharpe ratio.

Table A

	Certain	Case A	Case B	Positive Skew	Negative Skew	Kurtosis
Mean	10%	10%	10%	10%	10%	10%
Volatility	0	50.71	20.28	50.71	50.69	50.71
Skewness	0	0	0	0.0227	-0.0371	0
Excess Kurtosis	0	-2.12	-2.12	-2.065	-2.046	-2.011
Expected Value	30.91	30.91	30.91	30.91	30.91	30.91
Most Likely Value	30.91	0.4791	16.88	0.4863	0.4163	0.4539
Series (36)	10 certain	+60 –40 Equi	+30 –10 Equi	Complex	Complex	Complex
Sharpe (Zero)	Infinite	0.1972	0.4931	0.1972	0.1973	0.1972

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.

This paper introduces a performance evaluation measure, Γ, which captures the effects of all higher moments fully and which may be used to rank and evaluate manager performance. It avoids the problem of estimating individual moments by measuring their total impact, which is of course precisely what is of interest to practitioners. It distinguishes readily between distributions such as those in Diagram 1.1. Using returns data for actual financial time series for hedge funds, as we show later in applications, the additional information contained in the Gamma measure results in preference orderings that differ from the Sharpe ordering. Our approach also avoids the need for utility functions. In order to evaluate a collection of portfolios our performance evaluation function, Γ, will need just the simple decision rule that we prefer more to less, that we are not satiated[5].

Our performance measure is a natural feature of the returns distribution. In fact its construction from a returns distribution is entirely canonical, requiring no choices and admitting no ambiguity which is not already present in the data. As such it may be regarded as an extension of the notion of the cumulative distribution. It 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. The use of a function of returns rather than a single number to measure performance is essential as our examples show.

2. The Gamma Measure

We begin with an elementary heuristic. A direct analogy might be a simple bet. The investment situation differs from a standard gamble in that the "stake" is unknown at the outset. We wish to know what we stand to win if we win and what we stand to lose if we lose. In order to investigate this we need only specify the loss threshold L. This is the conditional expected return given loss. The return expectation is the conditional expected return given gain rather than the unconditional mean of the distribution. This is illustrated below, as diagram 2.1:

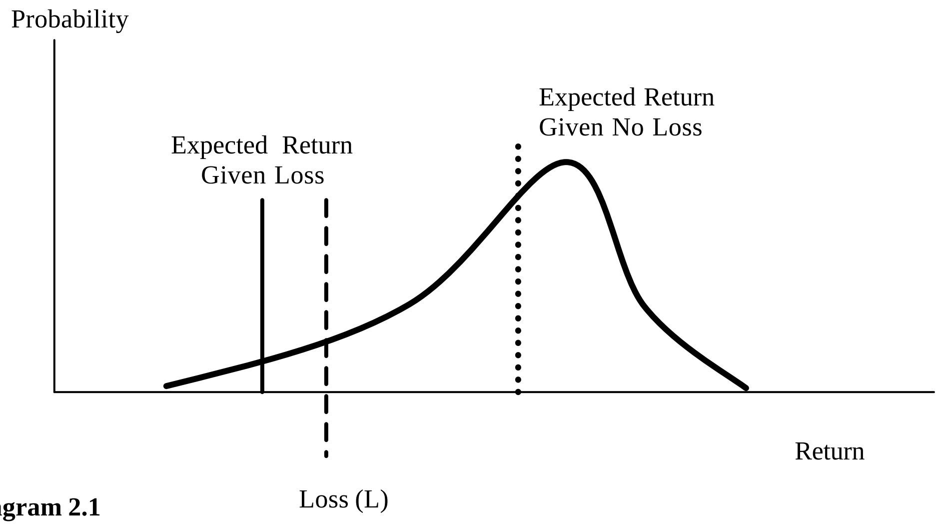

Diagram 2.1

The diagram above shows the conditional expected returns given loss and no loss for an arbitrary distribution of returns. The partitioning of the distribution by the loss line (L) may be around a zero return as would be implicit in the gambling analogy or it may be any other exogenously specified level. This may, for example, be the return from a benchmark index or an absolute rate of return such as that used in actuarial assumptions. It should be immediately recognised that this partitioning changes both expected gain and expected loss as it is varied.

The ratio of these two returns is directly analogous to the odds in a standard bet. If we now add consideration of the likelihood of each expectation, through a likelihood ratio, we have a measure of the quality of the bet taken, or in investment terms the portfolio performance. The likelihood ratio is the ratio of the areas to the right and left of the partitioning (L) in the diagram above. This statistic, which we shall refer to as Γ_S [6], is given by

$$\Gamma_S = \frac{E(r \mid r \geq L)}{E(r \mid r \leq L)} \frac{(1 - F(L))}{F(L)} \qquad [1]$$

where F is the cumulative distribution of the returns series. Graphically this statistic may be illustrated (Diagram 2.2) in terms of F as follows. The Γ_S statistic is the ratio of the crosshatched and striped areas.

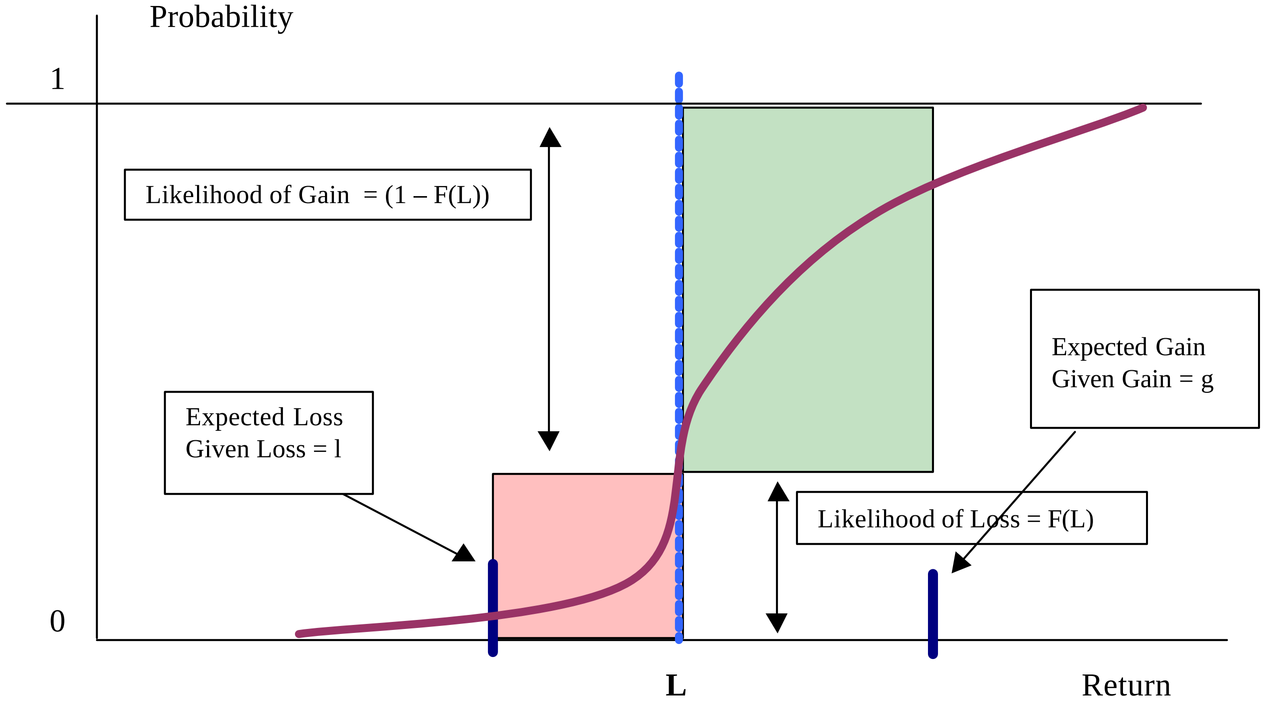

Diagram 2.2

This graphical representation makes obvious the inadequacy of using the conditional expectation in the formulation presented as Equation [1]. The Γ_S statistic when calculated as the ratio of the two areas omits some information and introduces some spurious elements. For instance we could equally well have used the probability weightings F(l) for the loss (as this is the probability of a return no greater than l) and (1-F(g)) for the gain (as this is the probability of a return no less than g). This would lead to the following diagram (diagram 2.3).

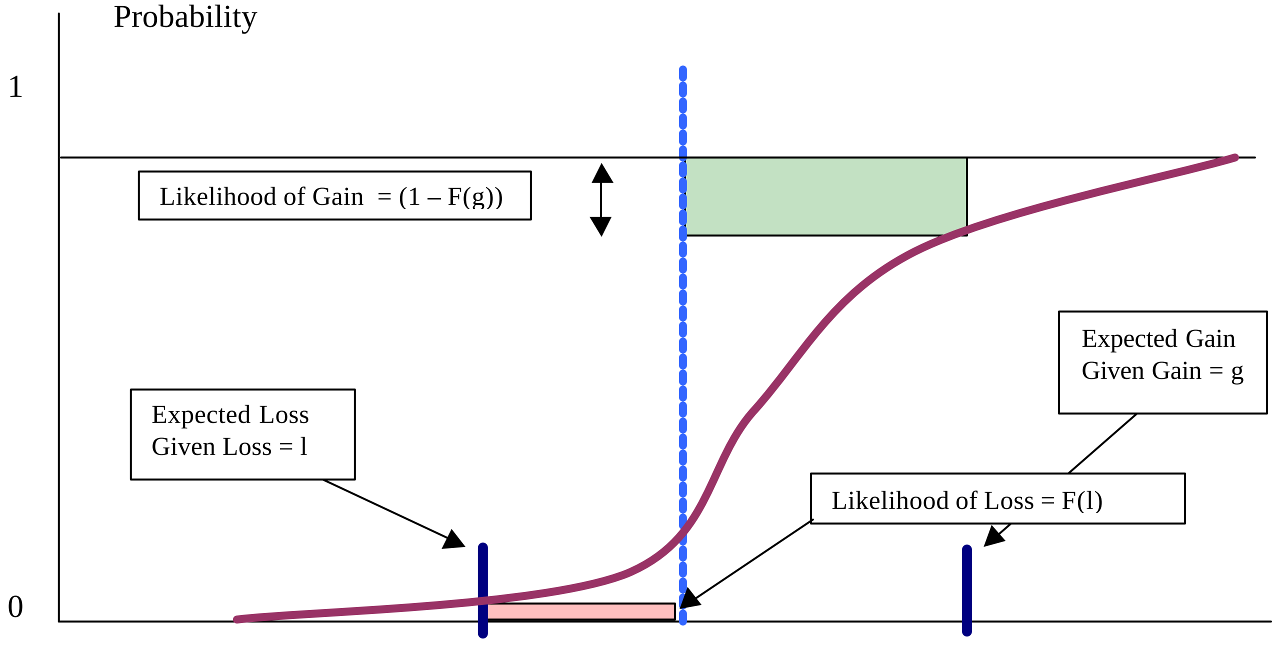

Diagram 2.3

However, we may take the concept a little further and correct for these failings as is illustrated below as diagram 2.4. We need only consider the limit in which the unit of gain or loss is allowed to tend to zero and sum the gains and losses with their appropriate weights.

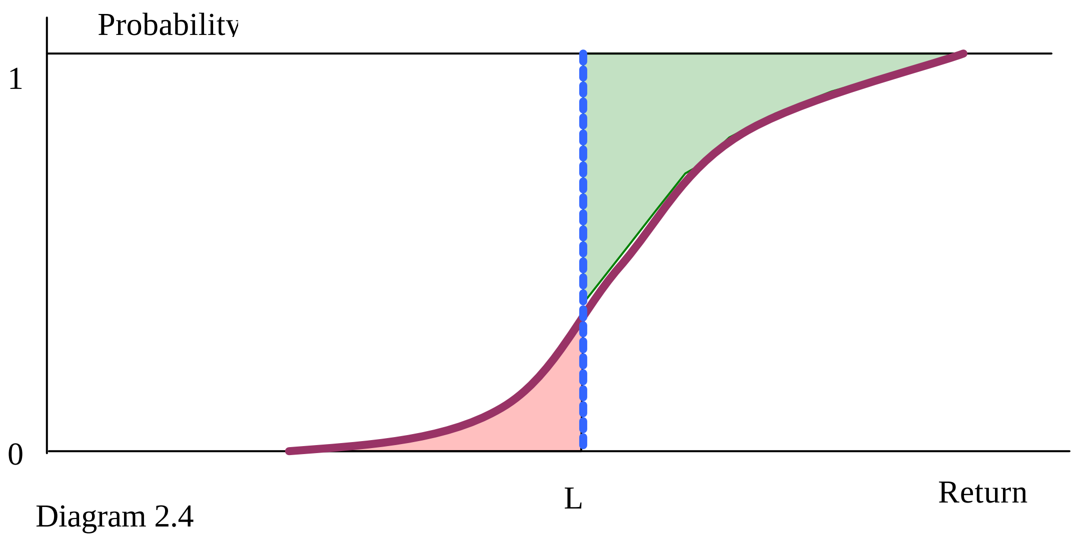

Diagram 2.4

The Gamma statistic will now be defined as the ratio of the two areas, striped and crosshatched. Mathematically this formulation of Gamma may be expressed as:

$$\Gamma = \frac{\int_L^b [1 - F(r)]\,dr}{\int_a^L F(r)dr} \qquad [2]$$

where the cumulative distribution F is defined on the interval (a,b). It is straightforward to see that this is the limiting case of the constructions above[7] .

All available information from the returns distribution, including higher moments, is contained in the cumulative distribution and hence is encoded in Γ. 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. The use of the statistic for choice is just a simple variation upon the use of stochastic dominance rules in decision theory[8]. In the case where the interval (a,b) is infinite there are distributions[9] for which the integrals above do not exist, but some elementary assumptions can exclude these and in practice existence presents no problems since we work with discrete return observations.

Assuming the convergence of the integrals, Γ is a natural feature of the underlying probability distribution. It is important to observe that, because the loss threshold can be any number, Γ is a function of the return level L. In fact, as we indicate below, Γ is a smooth monotone decreasing function from (a,b) to $(0,\infty)$. As we show in what follows, this function and its derivatives with respect to L have natural financial interpretations.

We will also show that regardless of the underlying distribution, Γ takes the value 1 at the distribution's mean m. Thus, as in the case of the Sharpe ratio, the single number $\Gamma(m)$ is inadequate to distinguish between distributions which have the same mean and variance but differ in their higher moments. With the risk threshold set at the mean, all bets are fair, a sidelight on the world of Sharpe and Markowitz and of course a martingale property. Nevertheless, the *function* $\Gamma(L)$ does encode the differences between such distributions as we illustrate in diagram 2.5 with the Gamma measures corresponding to the distributions of diagram 1.1 around their common mean of 10. Notice that as we move across the risk threshold our preference may change from one portfolio to another. This shift is a higher moment effect and not seen by the Sharpe measure. We provide some additional examples in applications and appendices later.

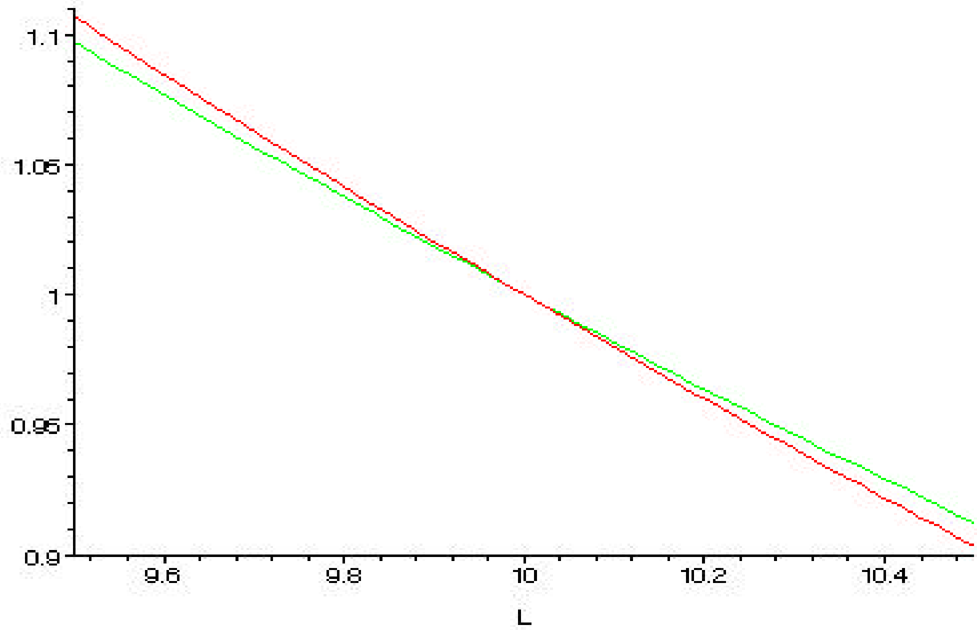

Diagram 2.5 The Gamma measures for the distributions of Diagram 1.1

There are many close relations to this ratio discussed in the downside and shortfall literatures. Perhaps the most common criticism of these downside-type measures is that the sample return set is small or even empty. It is commonplace to see parametric methods[10] imposed, such as the fitting of a three parameter log-normal distribution[11]. It is almost surely better to recognise that if a sample period contained no returns lower than the risk threshold, it was an unparalleled opportunity to borrow and profit[12]. The fundamental difficulty is that the sample set of return observations is insufficient to define the full stationary distribution[13]. This criticism of course holds true of any statistical analysis of portfolio returns.

Concerns over the existence and accuracy of estimated values of higher moments that render other approaches based upon the individual moments questionable are not relevant here. Although all of the information on higher moments is encoded in the formulation of Γ, it is obtained through the cumulative distribution and hence there is no need to know any of the individual moments in order to observe their effect in total.

We also observe that the definition of Γ is canonical given a returns distribution. As any invertible transformation of returns gives rise to an induced cumulative distribution, it also gives rise to an induced Gamma measure via this canonical construction. This means that one may transform the returns distribution at will prior to calculating the Gamma measure. For the purposes of comparison all that is necessary is that the same transformation be applied to all of the returns series in question. The effect of a transformation of the returns is to introduce a function of L rather than L itself as a risk threshold. In other words, we may introduce individual risk preferences by such a transformation. This may be thought of as providing alternative utility functions and has the effect of modifying means and variances as well as higher moments. We defer discussion of this point to a later paper.

3. Some elementary properties of the Gamma measure

We indicate some of the properties of the Gamma measure here. A more detailed exposition is left to another paper[14]. We begin by examining the sensitivity of Gamma to changes in the risk threshold L. Let $I_1 = \int_a^L F(r)dr$ and $I_2 = \int_L^b [1 - F(r)]dr$, so $\Gamma = \frac{I_2}{I_1}$.

We may differentiate this expression with respect to L to obtain $\frac{d\Gamma}{dL} = \frac{\frac{dI_2}{dL} I_1 - \frac{dI_1}{dL} I_2}{I_1^2}$, or more explicitly, $\frac{d\Gamma}{dL} = \frac{[F(L) - 1]I_1 - F(L)I_2}{I_1^2}$. In particular we see that $\frac{d\Gamma}{dL}$ is as smooth as F(L) and that $\frac{d\Gamma}{dL} < 0$ everywhere.

Thus $\Gamma(L)$ 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 $m = I_2(0) - I_1(0)$ and one may deduce from this and the definition of the Gamma measure that $\Gamma(m) = 1$.

Unless the cumulative distribution fails to be differentiable, the Gamma 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 Gamma measure as the underlying cumulative distribution varies according to various scenarios. This is also straightforward but we leave it to another paper[15].

The Gamma measure is an affine invariant of the returns distribution. That is, for any affine change of variable, $r \to j(r) = Ar + B$, there is an induced cumulative distribution and the Gamma measure for the induced distribution, $\hat{\Gamma}$ satisfies $\hat{\Gamma}(j(L)) = \Gamma(L)$. Conversely, if this relationship is satisfied by any change of variable j then $j(r) = Ar + B$.

The variation of Gamma with time is also meaningful as $\frac{d\Gamma}{dt}$ contains the full information set[16] 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. Clearly the integral of Gamma with respect to time is the normalised value of the security, firm or portfolio as a function of the risk threshold. The cumulative time evolution of Gamma for two index returns series evaluated at a risk threshold of zero and the equivalent Sharpe ratio are shown in the section Applications.

In common with most downside or lower partial moment measures, the components of Gamma (I_1, I_2) are sub-additive[17]. Strictly the ratio is, of course, dimensionless. The measure may also be directly related to other techniques in common use such as tracking error[18].

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[19]. With the Gamma function this dependence may be contained within a transformation of the "risk" threshold, *L*, which may also be thought of as a utility function. 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[20].

For the purposes of comparison, the only proviso which we need is that comparison is only valid between distributions at a common risk threshold. 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.

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 the Gamma measure 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[21]. 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[22]. 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[23] 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.

"Risk" Threshold – Zero

The Gamma and Sharpe ratios with a risk free of zero with the "risk" 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 Gamma ratio.

	Sharpe	Indicator		Gamma
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 Gamma measure 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, we should not only be considering the ordering by Gamma at a single threshold but rather the Gamma function over the range of returns, as is illustrated as Diagram 4.1 below.

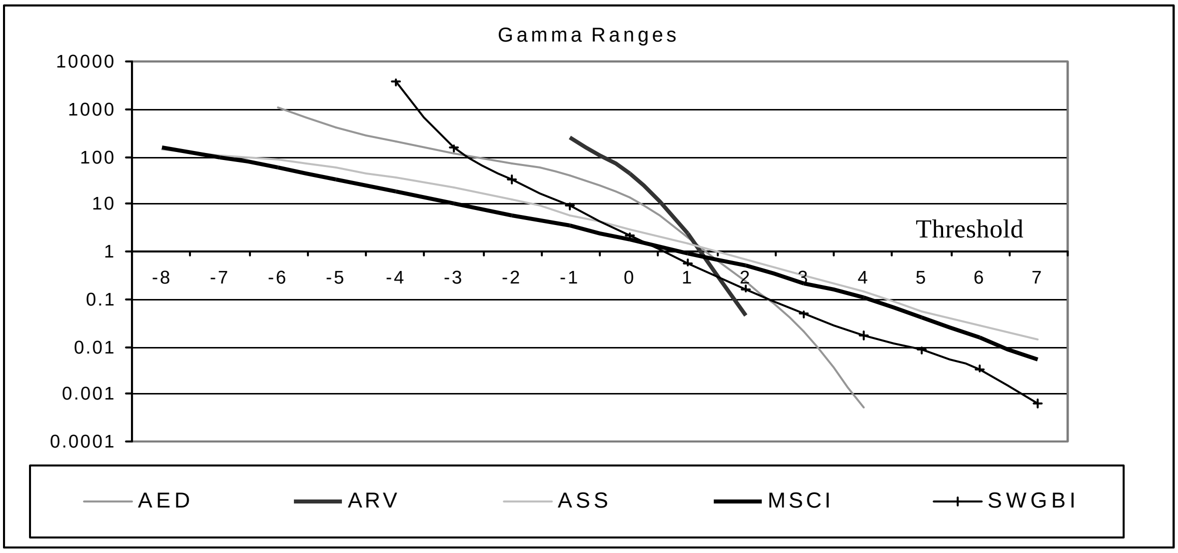

Diagram 4.1 Gamma as a Function of Return Threshold

For clarity, this shows only a selection of the index sequences over the range of returns experienced, on a logarithmic scale. Points where these curves, the Gamma measures, cross are indifference points for choices between particular portfolio pairings. In the broadest of terms, the steepness of the Gamma 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 Gamma measure 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 Gamma 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 the Gamma functions and a linear scale for the Sharpe ratio.

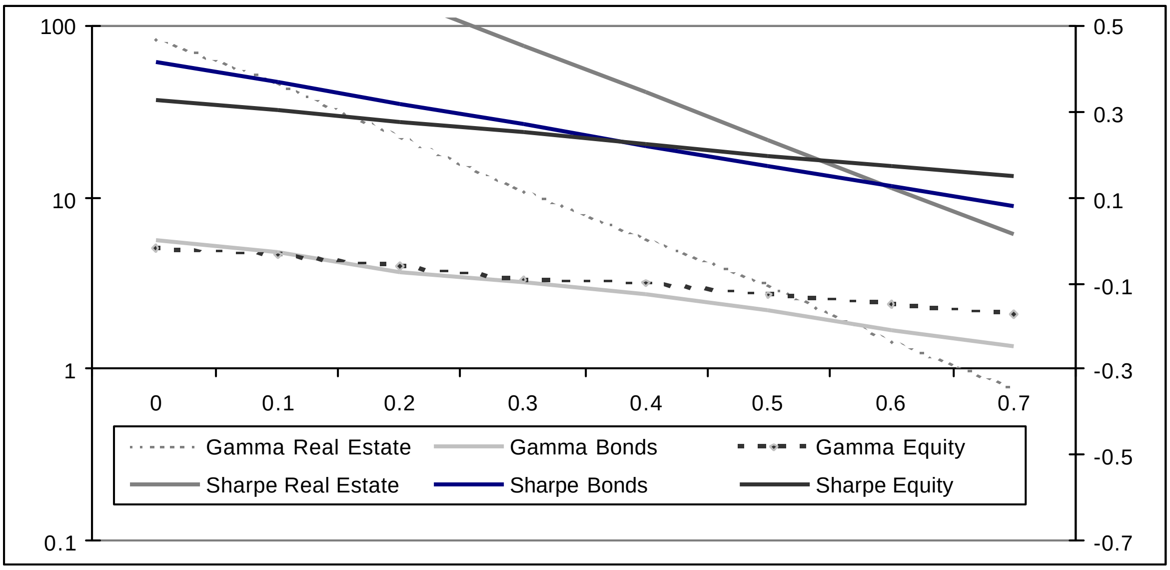

Diagram 4.2

Notice that none of the indifference points are coincident, again illustrating the importance of the higher moment information which the Gamma measure 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 again has the Sharpe ordering to the left and the Gamma ordering to the right.

	Gamma	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 Gamma 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.

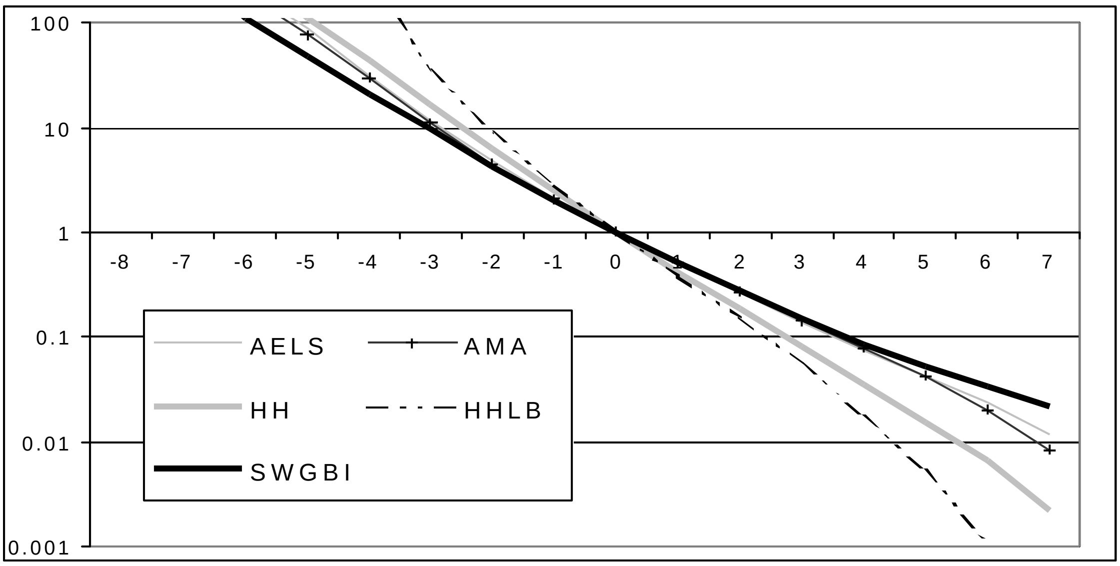

Diagram 4.3 Demeaned Gamma functions for a selection of MSCI relative indices

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[24] and by the Sharpe and Gamma measures at threshold zero.

	Tracking Error	Gamma	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 the more advanced Gamma measure 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	Gamma	Sharpe
Tracking Error	1	0.3088	0.3235
Gamma	0.3088	1	0.9559
Sharpe	0.3235	0.9559	1

Spearman's rank correlation coefficient :

	Tracking Error	Gamma	Sharpe
Tracking Error	1	0.4093	0.4289
Gamma	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 Gamma performance measures, it is hard not conclude that Tracking Error is a very poor performance measurement criterion or tool.

The final illustration is the time evolution of the cumulative Gamma 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.

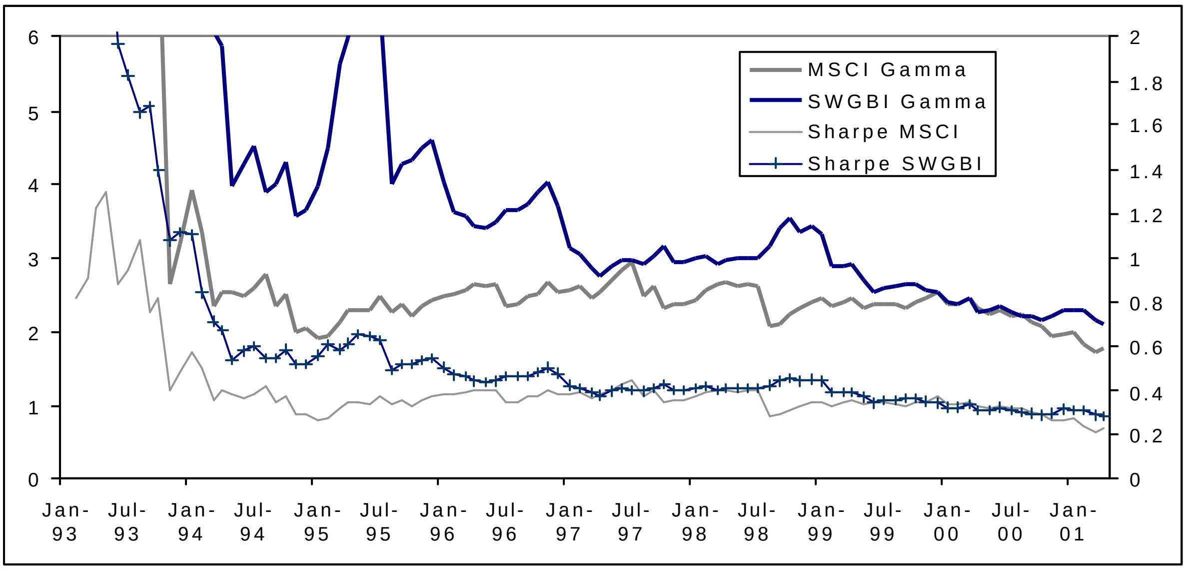

Diagram 4.4: A Comparison of the Cumulative Gamma and Sharpe Measures

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.

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 the Gamma measure captures. The presence of such effects in real data is, we trust, convincing evidence for the improvements in performance measurement which the Gamma measure provides.

We have also demonstrated, with examples from analytic probability distributions, that the Gamma measure is a powerful tool for the capture of higher moment effects. A number of most important basic properties of the Gamma measure have been stated. The affine invariance of Gamma 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 Gamma also provides for additional performance measures to be induced from more general transformations of the returns distribution and these may be interpreted as alternative utility functions encoding risk preferences or tolerances. The induced Gamma measure 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 Gamma and 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 Gamma would provide some useful insights.

The gain-loss literatures, such as Bernardo and Ledoit, already provide some insights as to how the Gamma measure might be used in asset pricing. An unpublished work of Agarwal and Naik extends that literature to optimal asset allocation[25].

The most obvious extension is to performance attribution. We might have followed Hicks and Marschak[26] in the observation that preferences are a function of all of the moments of a returns distribution and 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.

Insert tables 1- 4

Appendix A: Gamma 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:

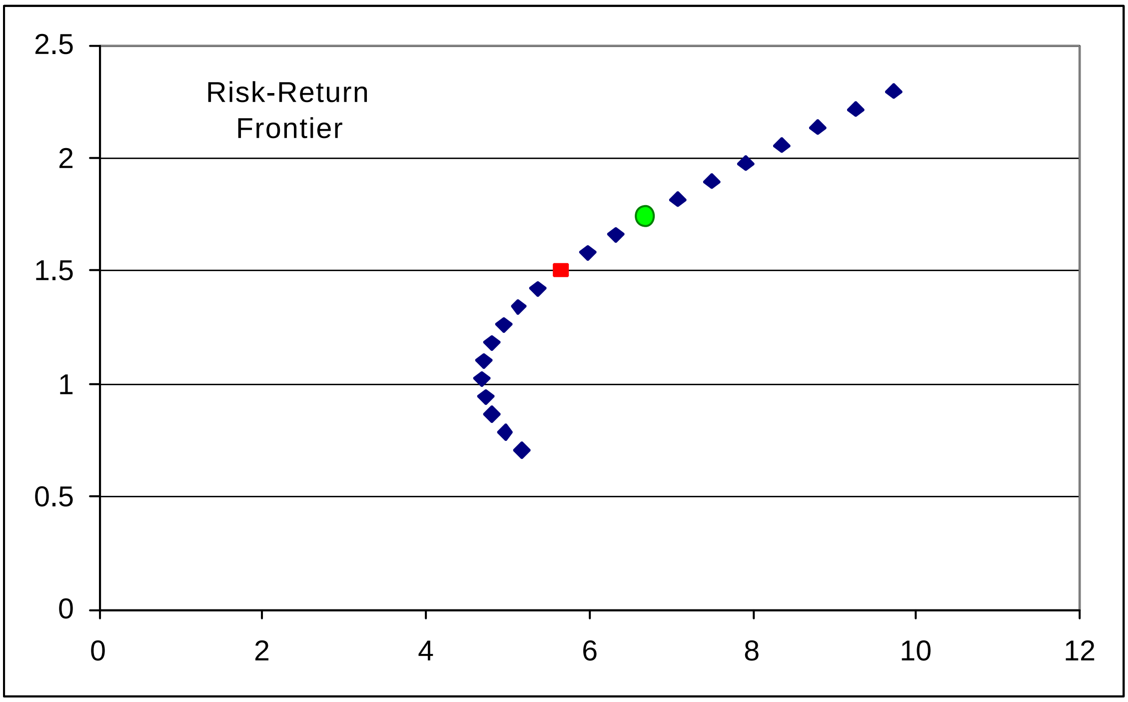

As is usual with these diagrams, the vertical axis is return (in %) and the horizontal standard deviation of returns (in %). The square is the optimal portfolio allocation given a risk-free return of 0%, a mixture of 50% A and 50% B. The Gamma optimal, the circle, is a significantly different asset allocation, (35% A and 65% B). We now present the Gamma analogue of this diagram:

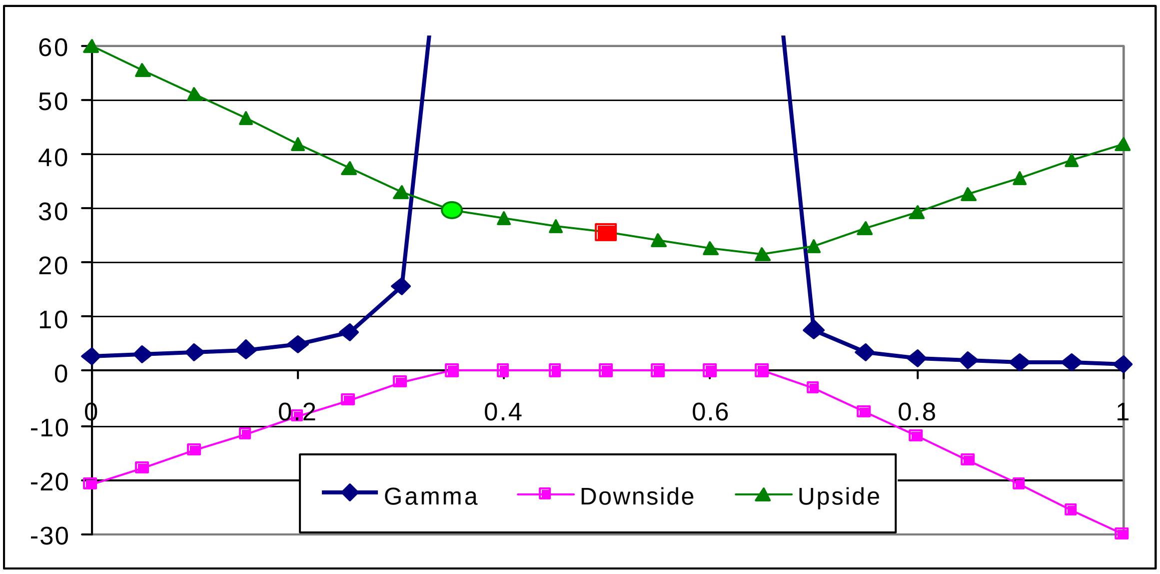

The vertical axis is the value of the upside or downside, I_2 or I_1 earlier. The horizontal axis is the proportion of asset A in the portfolio mix. The dotted triangular marked line is the upside value for differing mixes. The continuous square marked line is the downside value for differing mixes. The heavy lozenge marked line is the Gamma 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 the Gamma measure 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 Gamma measure for normal distributions as their variance changes at a common mean. Diagram B.1 shows Gamma for three normals, of mean zero, with standard deviations of 5,10 and 15, shown dotted, dashed and solid 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. The smaller the variance, the more negative the slope of the Gamma measure.

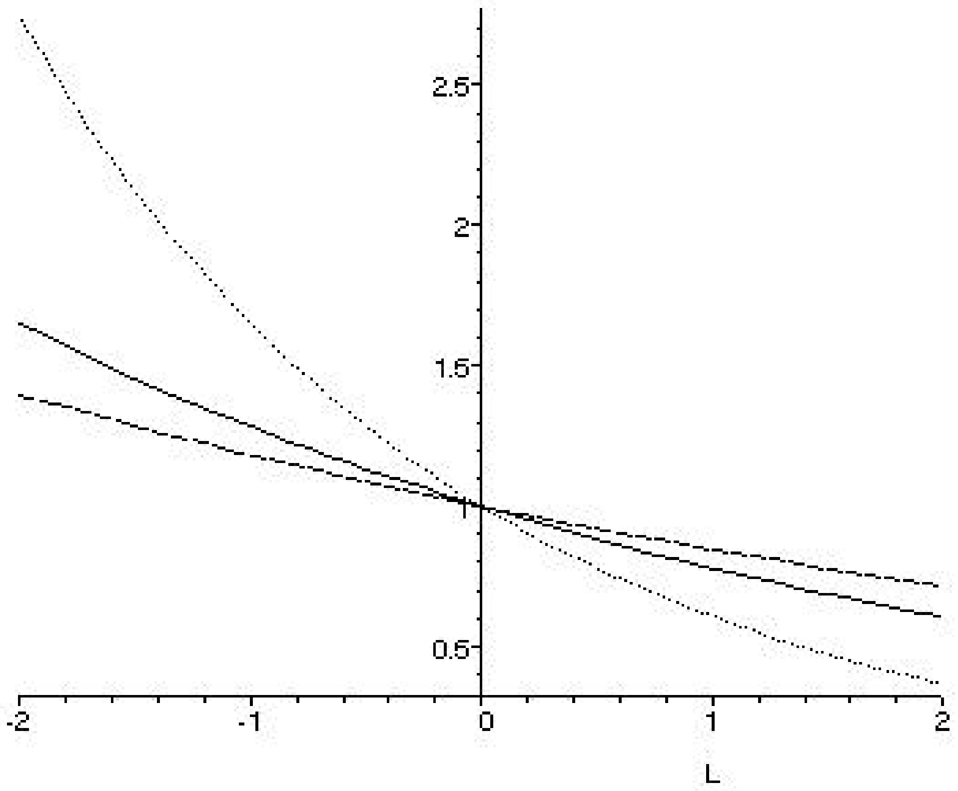

Diagram B.1: Gammas for normals of mean zero and variance 5,10 and 15

We next consider the effects of skew and kurtosis. More precisely, as the Gamma measure 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 Gamma measures for distributions with skew 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.

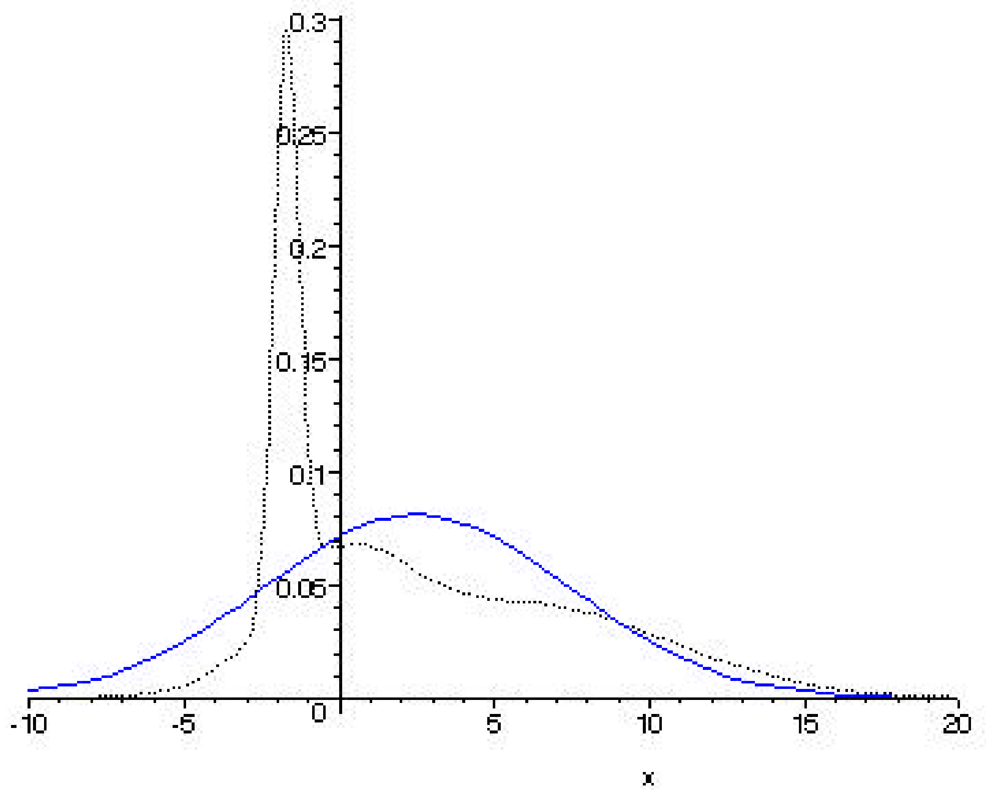

Diagram B.2: Two distributions with the same mean, variance and kurtosis

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 Gammas corresponding to the these distributions are shown in Diagrams B.3 and B.4. While there is a separation in the Gamma 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 Gamma with respect to L, illustrated in Diagrams B.5 and B.6 indicate.

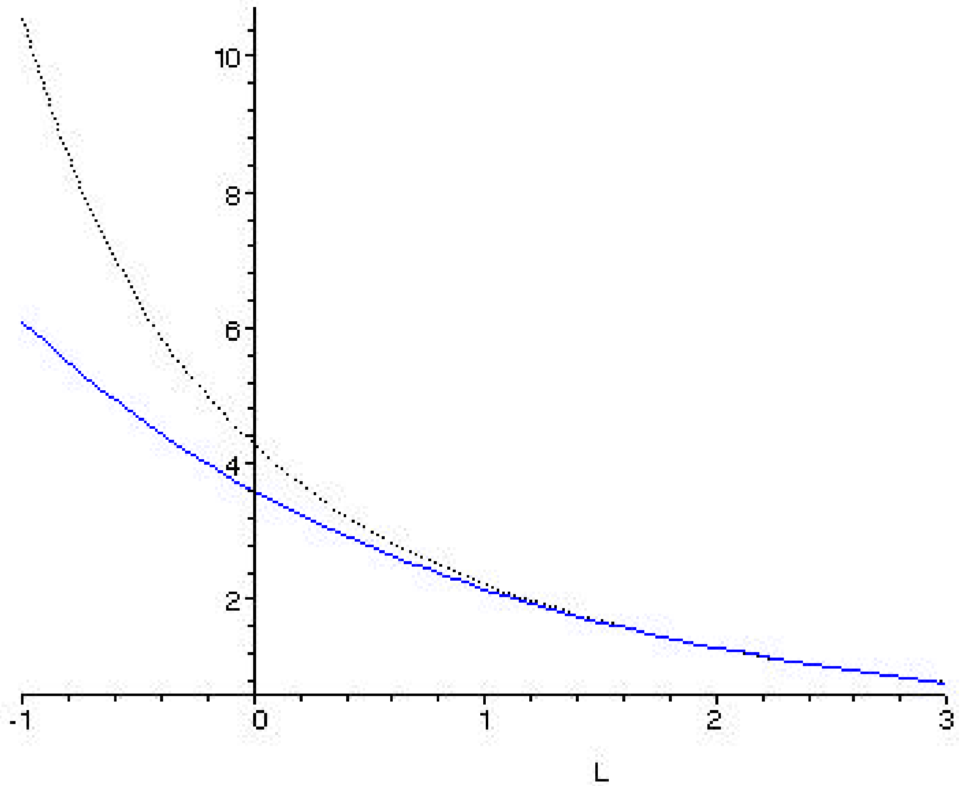

Diagram B.3: Gammas for the distributions in Diagram B.2

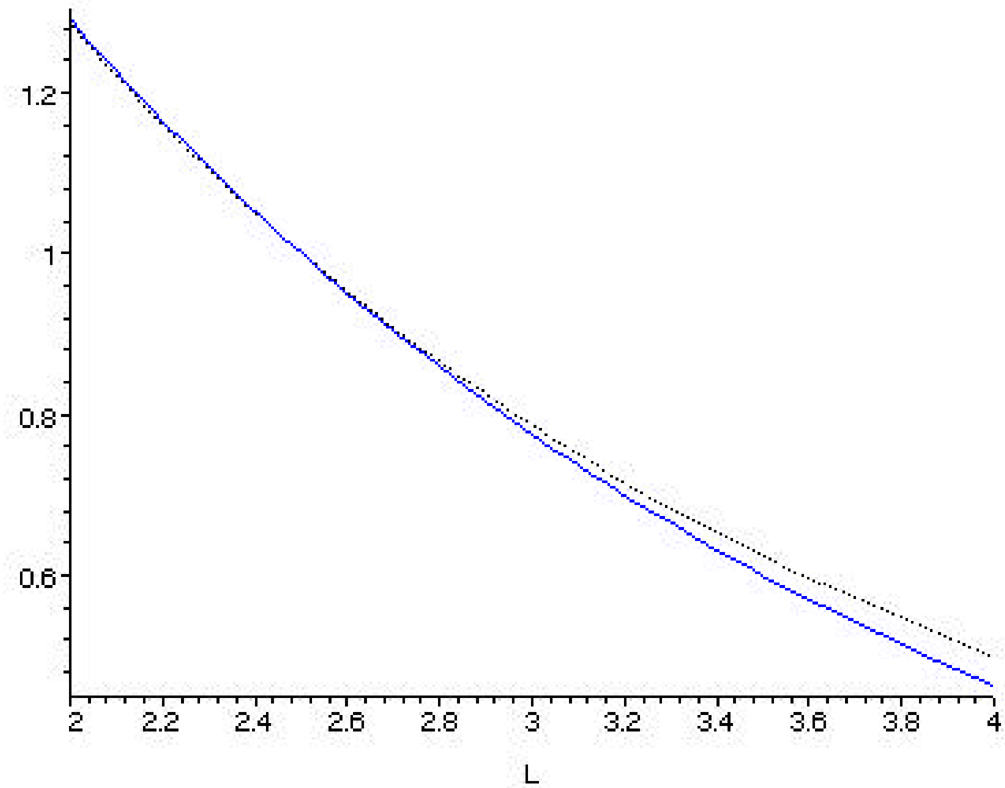

Diagram B.4: Gammas for the distributions in Diagram B.2

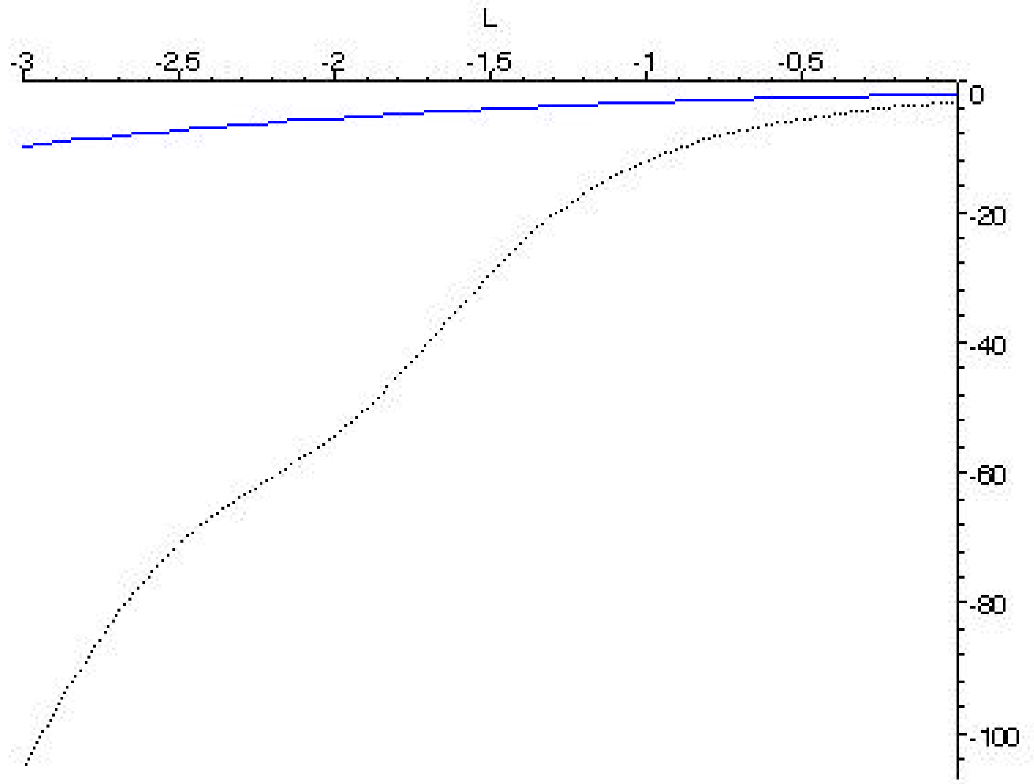

Diagram B.5: First derivatives of Gamma with respect to L

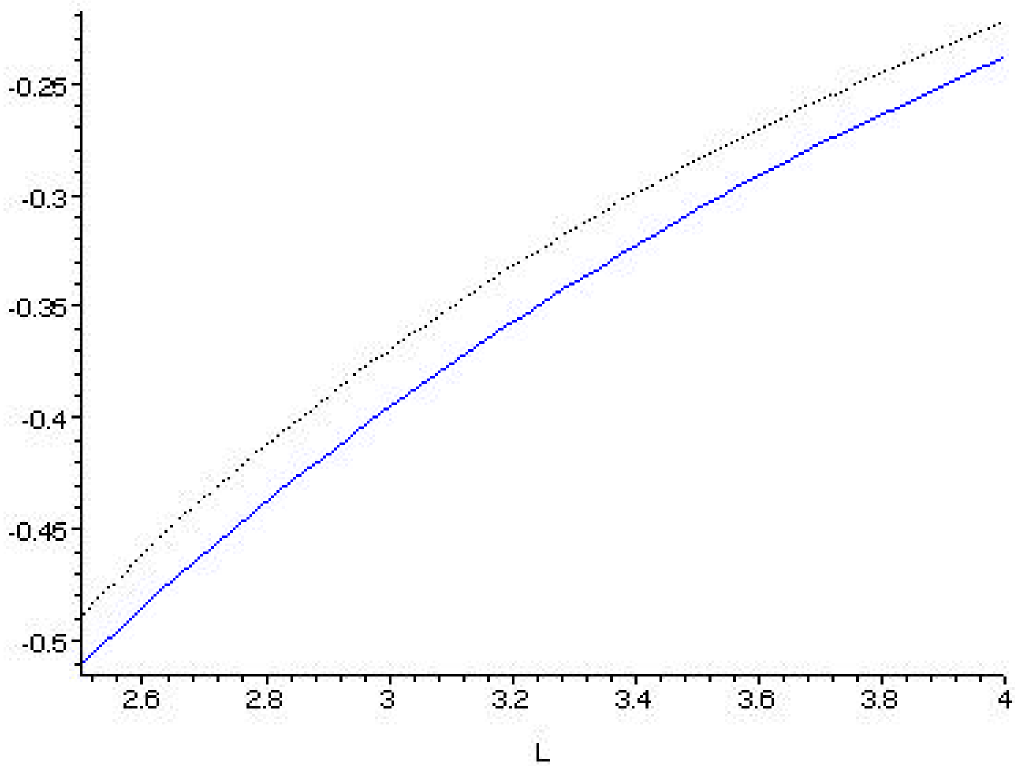

Diagram B.6: First derivatives of Gamma with respect to L

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

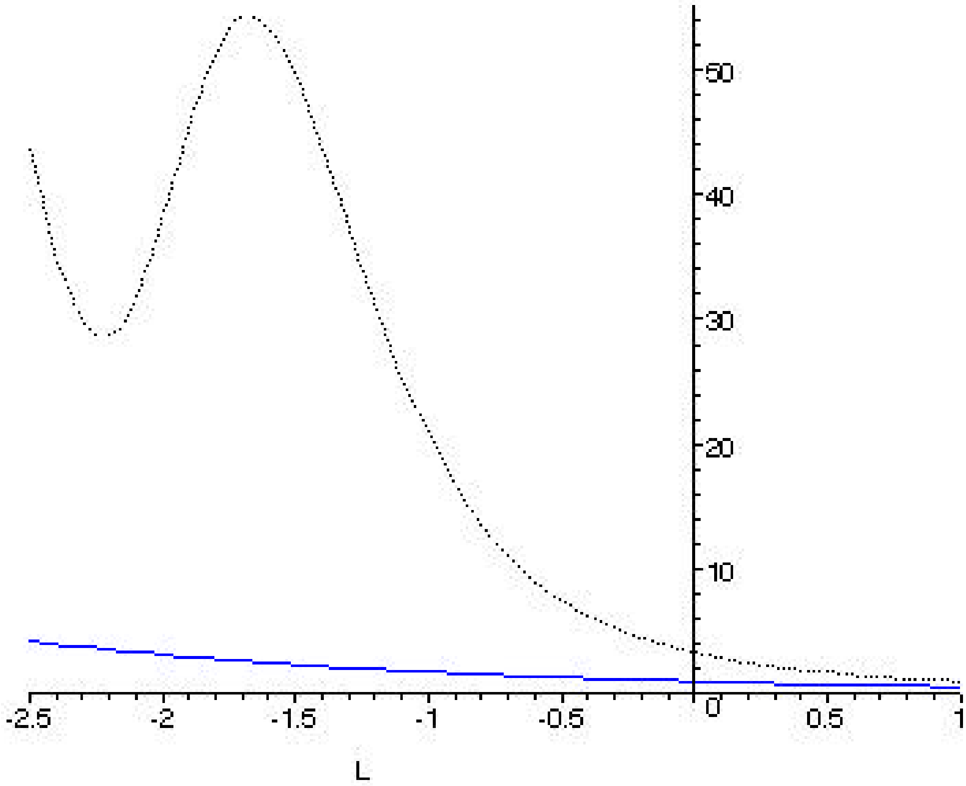

Diagram B.7: Second derivatives of Gamma

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

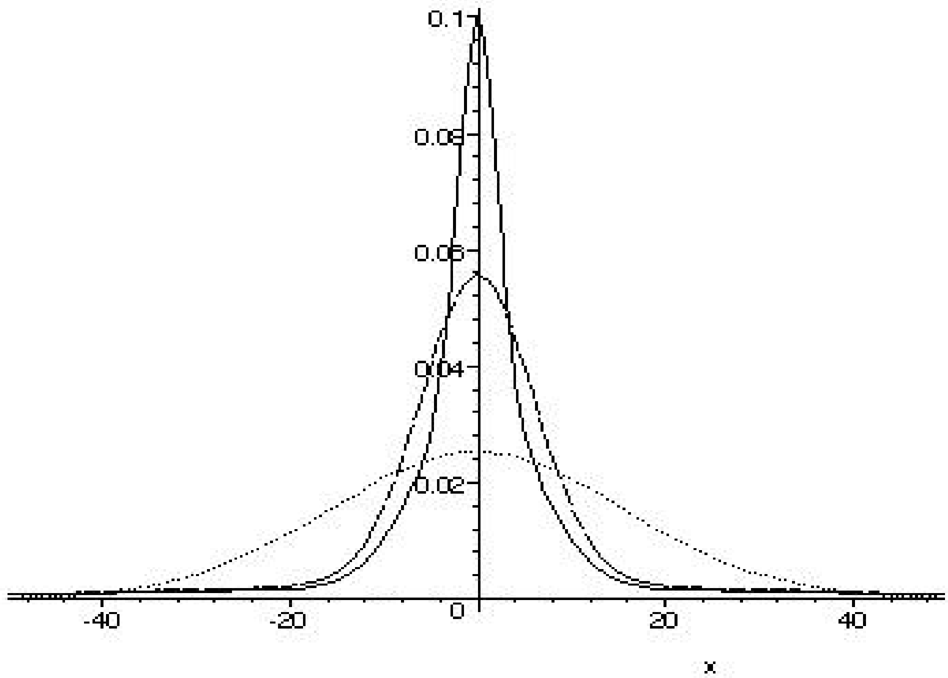

Diagram B.8: Symmetric distributions differing in kurtosis and higher even moments

As diagrams B 9,B.10 and B.11 show, although on very different scales, the Gamma 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 Gammas each cross the Gamma for the normal distribution in three places and themselves cross three times, at +/-20 and at their common mean of 0.

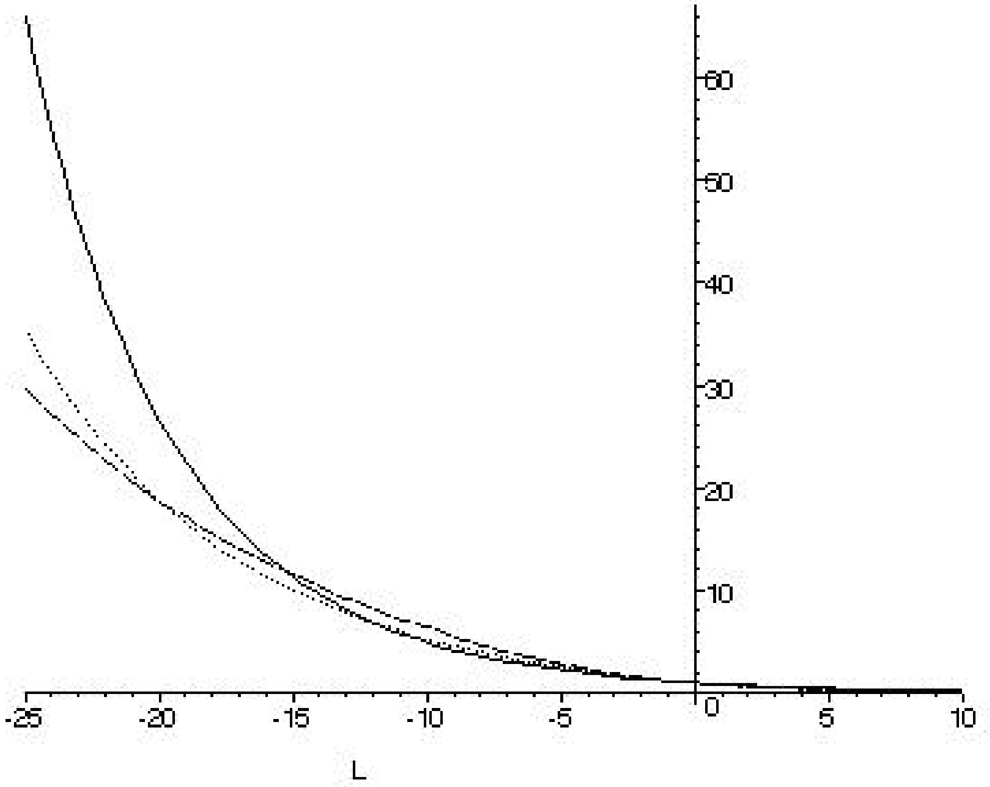

Diagram B.9: Gammas for the distributions of Diagram B.8

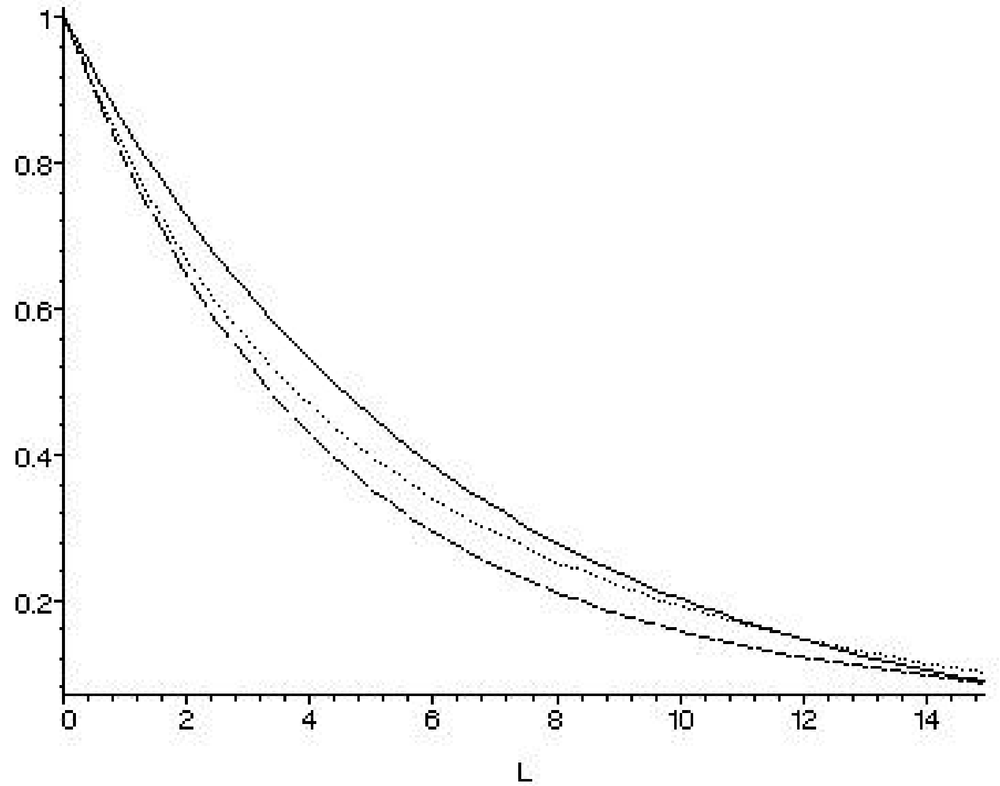

Diagram B.10: Gammas for the distributions of Diagram B.8

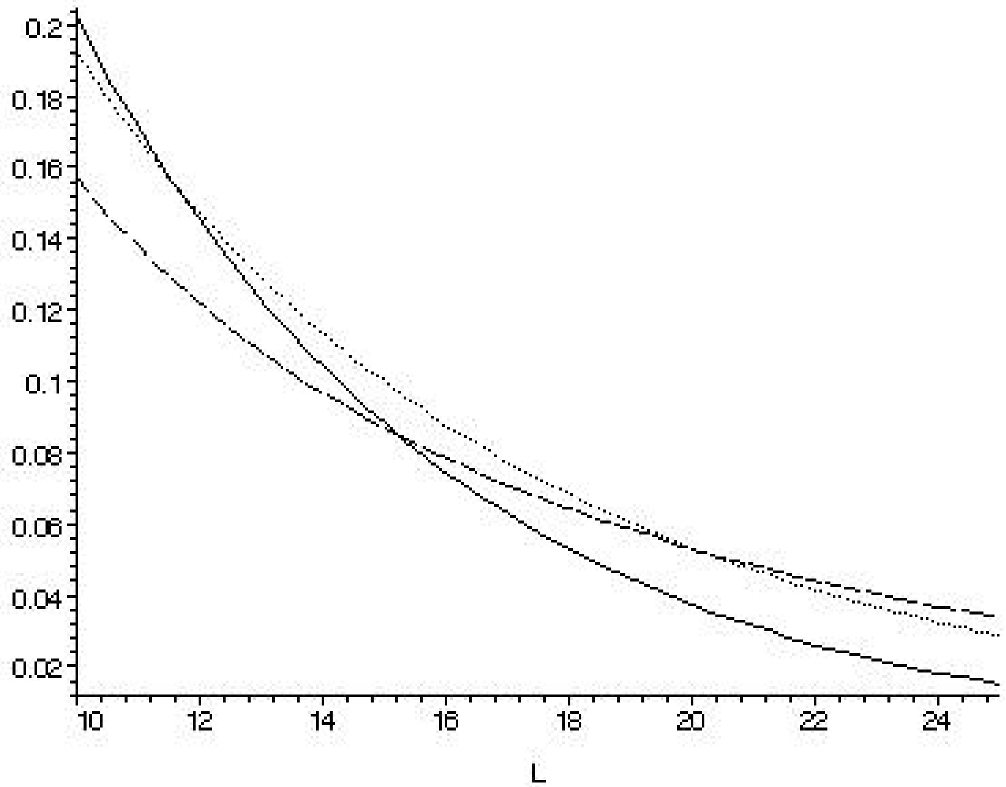

Diagram B.11: Gammas for the distributions of Diagram B.8

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.

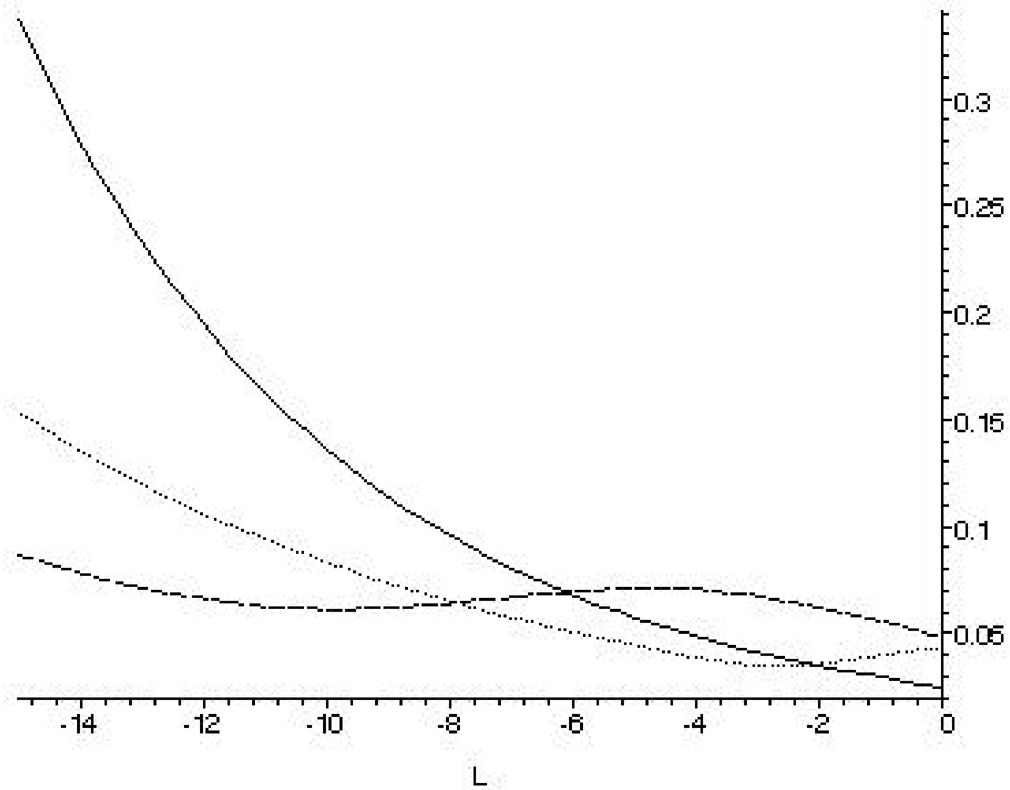

Diagram B.12: Second derivatives of Gamma with respect to L for the distributions of Diagram B.8.

---

1 The expected value of any of the series listed is $[1.1]^{36}$, while the most likely value in case A is $[1.6]^{18}[0.6]^{18}$ and $[1.3]^{18}[.9]^{18}$ for case B and so on. The row labelled Series is a succinct description of the series of returns which generates the descriptive statistics. For the skewed and kurtotic examples the series are irregular.

2 In point of fact, using the measure Gamma which we shall subsequently develop, we can illustrate that this is not generally true but depends upon the level of the "risk" threshold, which in this instance is a zero return. See appendix B.

3 The robustness of the Sharpe measure is in large part because the events which lead to skewness and asymmetry tend to be large relative to the average and when such an event occurs, the variance of the distribution also increases. Similarly events which change kurtosis will also be reflected in the variance, and kurtosis is also affected by a skewness event. In other words, the variance or standard deviation already captures some but by no means all of these effects. Disaggregating these effects, as we might wish if we were to specify a model which considers the effects of each central moment separately and the rates of substitution across them, is a non-trivial econometric exercise. It is also worth noting that the presence of asymmetry in the distributions of returns renders suspect the use of the product moment correlation as a measure of dependence as the distributions are no longer jointly elliptical.

4 See also: Scott R.C. and Horvath P.A. "On the Direction of Preference for Moments of Higher Order than the Variance." Journal of Finance Volume 35, Issue 4, pp 915 – 919 Sep 1980.

5 Quadratic utility does not satisfy the non-satiation property.

6 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

7 This follows from the definition of the Riemann integral.

8 See: V. Bawa and J. Lindenberg, "Capital Market Equilibrium in a Mean-Lower Partial Moment Framework" Journal of Financial Economics 5 pps 189-200 1977. See also appendix B for a cursory examination of the role of crossings.

9 A note on the convergence of the integrals:

Let p(x) be the underlying probability distribution so that the cumulative distribution function is given by $F(r) = \int_{-\infty}^{r} p(x)dx$. This is finite but it is possible that $I_1 = \int_{-\infty}^{L} F(r)dr$ will not converge. For example unless p(x) decays faster than $\frac{1}{x^2}$ as $x \to -\infty$, F(r) will asymptotically approach $\frac{1}{r}$ or worse as $r \to -\infty$ and similarly I_2 will diverge unless p(x) decays faster than $\frac{1}{x^2}$.

Even if the tail decay does not satisfy this condition it is of course possible that Γ will be defined in a Cauchy principal value sense – i.e. as

$$\lim_{A\to\infty} \frac{I_2(A)}{I_1(A)} \quad , \quad \text{where } I_1(A) = \int_{-A}^{L} F(r)dr \text{ and } I_2(A) = \int_{L}^{A} [1 - F(r)]dr \ .$$

10 There are also some serious statistical sampling theory issues here. For example, the sample standard deviation estimator is an estimator of sigma but not unbiased and moment estimators may be bounded by the sample size – a sample of 40 observations cannot produce a skewness of larger than 6.325 regardless of the true skewness of the distribution. For further detail see: Wallis J., Matalas N. and Slack J. "Just a Moment!" Water Resources Research, 10, pp 211- 219. 1974

11 The three parameter log-Normal uses mean, standard deviation and an extreme value. See Aitchinson J. and J. Brown, "The lognormal distribution" C.U.P. 1957

12 A Gamma value of 0 arises when there are no positive returns. A value of 1 for the Gamma statistic is the family of martingales and of course, an infinite value of Gamma is a riskless investment with a positive return, the free lunch of popular finance. This arises when the risk set is empty. In between these values, we have a range of quasi arbitrage opportunities, relatively high returns for a given level of risk. There are obvious applications to the pricing of non-replicable option portfolios. We can show that portfolios of relatively simple assets can give rise to very high and even infinite Gamma values (See appendix A).

13 It is also worth noting that the standard tests for stationarity operate only upon the mean and variance of the distribution.

14 A. Cascon, C Keating and W. Shadwick "Properties of the Gamma Measure" Preprint Finance Development Centre 2001

---

15 A Cascon, C. Keating and W Shadwick 2001 op cit

16 In the limit, $\frac{d\Gamma}{dt}$ is the innovation.

17 The failure of the sub-additivity test is perhaps the most damning criticism of Value at Risk. The practical consequence that the risk of a portfolio may be greater than the sum of its parts is sobering, if perhaps not significant in most applications.

18 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 Gamma statistic or some risk adjusted counterpart centred on the passive index give us direct measures of relative performance.

19 This is just the classic statement , $p = \frac{E(x)}{R^f} + \frac{\text{cov.}[b\, u'(c_{t+1}), x_{t+1}]}{u'(c_t)}$ 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.

20 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 a threshold function (L) which is not linear or constant over time will affect the descriptive statistics or properties of the returns distribution and consequently the Gamma functions.

21 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.

22 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.

23 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.

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

25 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

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