Risk, Reward and Performance Attribution

The financial markets now devote much time, effort and expense to the analysis of risk; apparently more so than return. By way of illustration, a Google search of the internet on the words "Risk Finance" returned 129 million articles while a search using "Return Finance" showed 83.3 million articles[1]. The standard model of financial analysis, of course, holds that the relation between risk and return is linear[2]; this is often interpreted as more risk (equals) more return. Of course if return and risk were related in this way, the analysis of risk would be functionally equivalent to the analysis of return. This chapter will consider some basic relations between risk and return and some of the consequences for risk, regulatory and investment management.

Unfortunately definitions of risk adequate for analytic use are elusive. The Concise Oxford Dictionary defines risk as "a chance or possibility of danger, loss, injury or other adverse consequence". Whether we might really wish to consider a chance of danger, a conditioned conditional, as risk is debatable but the element of interest is chance or possibility. This definition is incomplete in the sense that it is just a likelihood, but demonstrates that as a minimum we need to motivate the concept of probability[3].

Time is of the essence in this or, more precisely, the irreversibility of time; if time were reversible then we could go back and alter our historic actions until a satisfactory outcome arose and risk would lose any meaning. It is also obvious that we only observe or experience the single outcome. Uncertainty or a measure on it, probability, along with dynamics, arise as a fundamental property of nature in open (non-deterministic) systems,

---

1 A number of my reviewers have repeated this exercise and report numbers which differ greatly from those reported here. This is an illustration of an aspect of uncertainty not covered by this chapter; that of precision in knowledge. This can, of course, be very significant in non-linear systems or more generally in any process which exhibits sensitive dependence upon initial conditions.

2 Some have asserted that non-linearity would imply the presence of arbitrage opportunities. In fact any arbitrage opportunity available would be probabilistic rather than deterministic.

3 In this essay probability may, with some abuse, be thought of as either relative frequency or arising from the axiomatic framework due to Kolmogorov. For an intriguing discussion of the history of mathematical thought on probability, and a novel, game-theoretic approach to probability and finance the reader is referred to: Glenn Shafer & Vladimir Vovk *"Probability & Finance"*Wiley 2001 ISBN 0-471-40226-5. It should be noted though that this stops far short of an adequate discussion of the more recent algorithmic complexity approaches.

usually far from equilibrium; such systems also have the pleasing property that they may give rise to events of entropy production and self-organisation, and the complexity and richness that we observe in our existence. The relation between risk and dynamics is often overlooked. It is not at all uncommon to see comparative statics[4] presented as being informative about risk, when the reality is that these studies are at best incomplete. Risk may perhaps be more productively thought of as the dynamics of a process under uncertainty.

Time is rather more than our existential medium; it orders the events which we refer to as gain or loss. In financial application, as in nature, risk, probability and dynamics are intrinsically related; this is the source of the more complete definition of risk as the product of likelihood and consequence of an event. This description gives rise immediately to concern over a risk-management technique much favoured by regulators, the stress test or scenario modelling as it is sometimes known; here only the consequence of an event is considered but not the likelihood of occurrence.

For the purpose of this essay, it is sufficient to consider risk as the undesired subset of uncertainty and leave, for now, the undesired partition subjective; this reduces the problem to one of gain or loss events, and their related likelihoods. It of course presents problems for anyone wishing to speak of the riskiness of a return, since we can now only consider such concepts in terms of gain and loss as ensembles in defined time. This may for simplicity be thought of as another incarnation of the single outcome observed.

Moving to the specific situation of financial risk, the effect of uncertainty on realised returns is well-known; it drives a wedge between arithmetic and geometric returns. Mathematically this relation is[5]:

$$R_{\text{Geometric}} = R_{\text{Arithmetic}} - \frac{1}{2}\sigma^2 + \cdots \quad , \qquad [1]$$

---

4 Comparative statics is the analysis of the effect of changes in exogenous parameters or variables upon the endogenous variables within a model. It was introduced by J.R. Hicks in his 1939 *"Value and Capital"* and expounded somewhat more fully by Paul Samuelson in *"Foundations of Economic Analysis"*(1947). It is a technique widely used in micro-economic analysis.

5 This is a Mercator or Taylor series expansion.

where σ^2 is the variance of the return process. This formulation is known as the certainty equivalent premium (CEP), and is obviously non-linear in nature. This is illustrated below for two series of random samples from Normal distributions. These samples are drawn from N(.05, .10) and N(.05, .20) and are shown together with a deterministic or certain 0.05 return. The descriptive statistics for these samples are N(0.049,0.202) and N(0.052,0.104). The gross regularities evident do not therefore arise from sampling difficulties. The scale is geometric resulting in the certain five per cent return producing a straight line.

Figure 1: Value evolution of samples from N(.05,.10) and N(.05,.20) returns together with fixed 0.05 return.

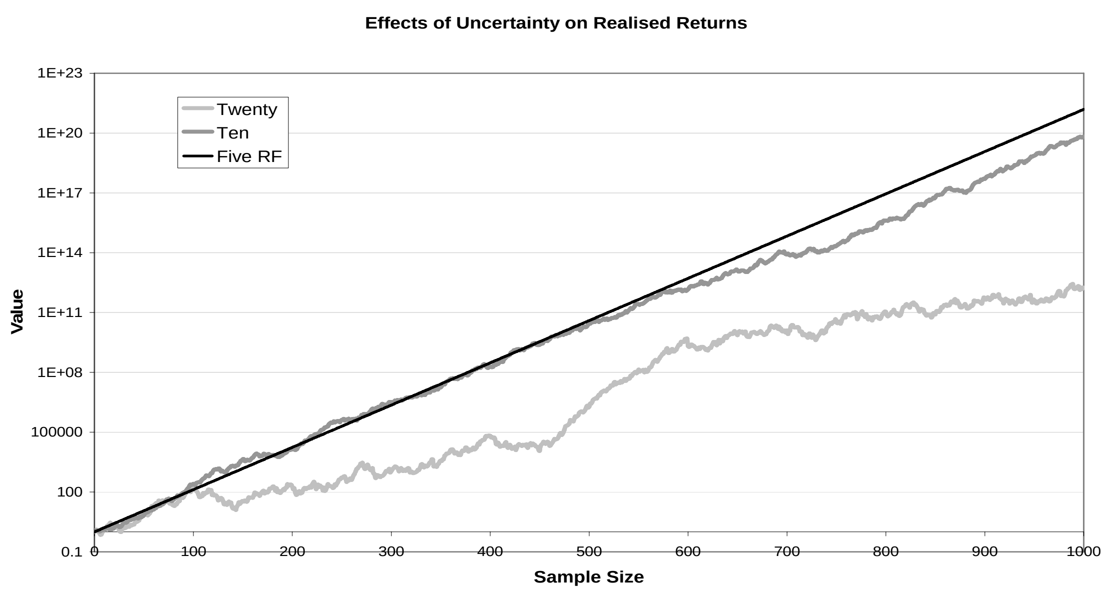

We see here that the effect of increasing uncertainty is to lower future wealth. In fact the arithmetic mean of the N(.05,.10) sample lies, perhaps surprisingly, above that of the certain wealth. The non-linearity of the relation between uncertainty as measured by standard deviation and the certainty equivalent premium is illustrated below:

Figure 2: Certainty Equivalent Premium and Volatility

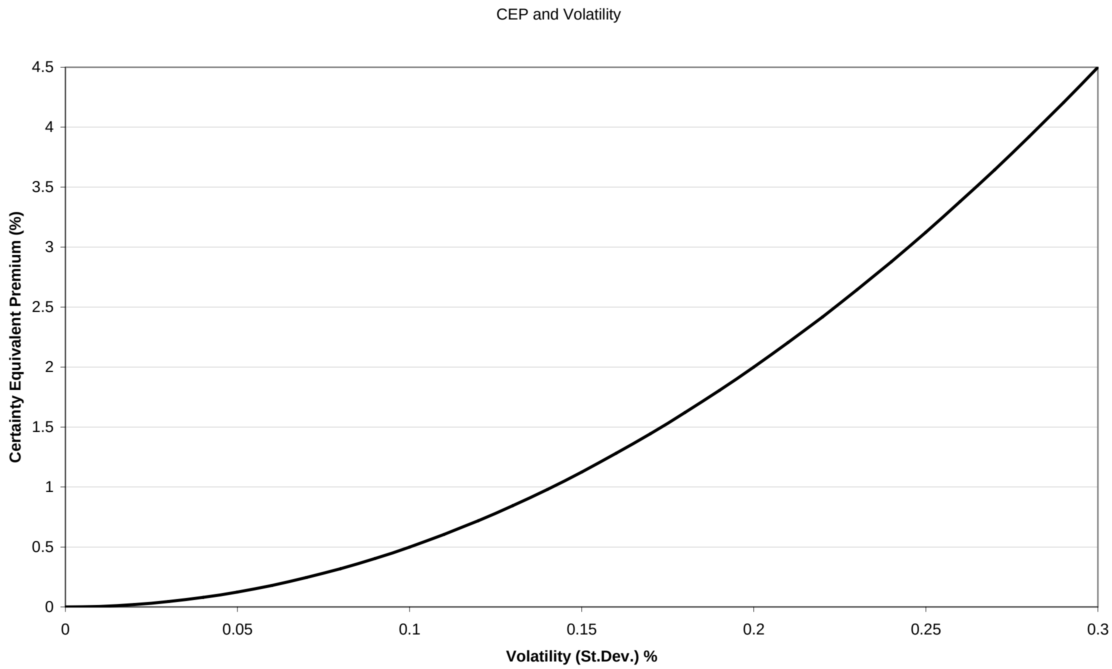

This raises a fundamental difficulty for performance attribution based upon a linear view of risk and return. The characteristic of a risk free interest or discount rate is precisely that it has coincident arithmetic and geometric returns; this is a deterministic rate and of course cannot exist in reality. It follows that if we are to go forward and consider the existence or magnitude of risk premia (RP), we need to be looking to the realised geometric returns, and volatility or risk as variation around those. Mathematically this relation is:

$$RP = \left(\frac{V_T}{V_0}\right)^{T^{-1}} - \left(1 + r_{det}\right), \qquad [2]$$

with V representing value and r_{det} the risk free or deterministic interest rate. Again we have the single observation problem; in a single period we have just one observation and the arithmetic and geometric returns cannot differ. Risk lacks meaning.

There is also the question of the uncertainty effects of symmetric and asymmetric shocks. We illustrate the effects of large positive shocks by taking the earlier N(.05,.10) data and applying a single .75% event to that returns series in period 100. We adjust the balance of the series such that they retain a common return. There is a small difference in volatilities but sufficiently small that it can be ignored.

Figure 3: Effects of Asymmetry

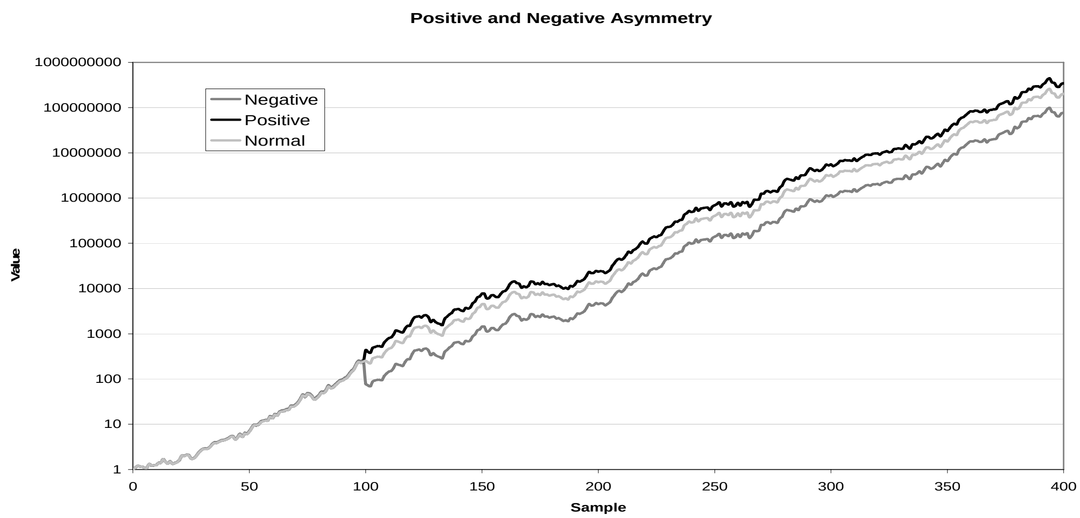

It is immediately obvious that asymmetric effects have rather different characteristics to symmetric. Where symmetric effects tend to fade away rather rapidly, asymmetric are marked by their persistence. The literature on risk management often contains rather inchoate cautions with regard to asymmetry; this illustrates a clear effect on wealth.

There is now good evidence to believe that asymmetric shocks to capital markets are far more persistent and have more meaningful longer term effects upon an economy than symmetric[6]. This is also related to an important distinction between realised and unrealised losses or gains. A loss or gain realised is asymmetric due to the finality of realisation; perhaps the most obvious form of this is an insured event. By contrast an unrealised gain or loss holds the potential for change, and symmetric outcomes, at future times. One of the weaknesses of elementary modern finance is that it does not distinguish between realised and unrealised gains or losses[7].

---

[6] Bloom N., 2007 "The Impact of Uncertainty Shocks, A Firm-Level Estimation and a 9/11 Simulation". www.stanford.edu/~nbloom/ImpactUncertaintyShocks.pdf It is far from obvious why one form of capital market shock should cause a slowdown in the real economy but there are a number of plausible hypotheses; these range from psychological wealth effects to the persistence of the lower collateral value of the capital market asset after a substantial loss.

[7] This is rather more a tax effect, though these can be relevant. If we think of capital markets as representing an American option on liquidity, we can see that we are destroying value by exercise.

For completeness, as this is a topic much discussed among practitioners, we will illustrate the effects of kurtosis or "fat tails"[8]. To facilitate this we take draws from a Uniform (-.15,+.25), the descriptive statistics for this sample and the Normal (.05,.10) data earlier are shown in table 1:

	Uniform	Normal
Mean	0.051	0.052
St. Dev.	0.116	0.104
Skew	-0.003	-0.033
Kurt	-1.242	-0.203

Table 1: Descriptive statistics kurtosis illustration

There is a caution relevant for anyone who seeks to use moment expansions, such as the Taylor series based upon the sample statistics of a dataset, such as those shown above. A single event influences the value of all these moment estimates; this means that to estimate the moments correctly it is necessary to estimate them jointly. This is something that we have never seen done in practitioner research. It can also lead to spurious results in analytic inference; for example a single event can raise the estimates of volatility, skewness and kurtosis making it difficult to attribute causality to any one of them.

Figure 4: Kurtosis Effects

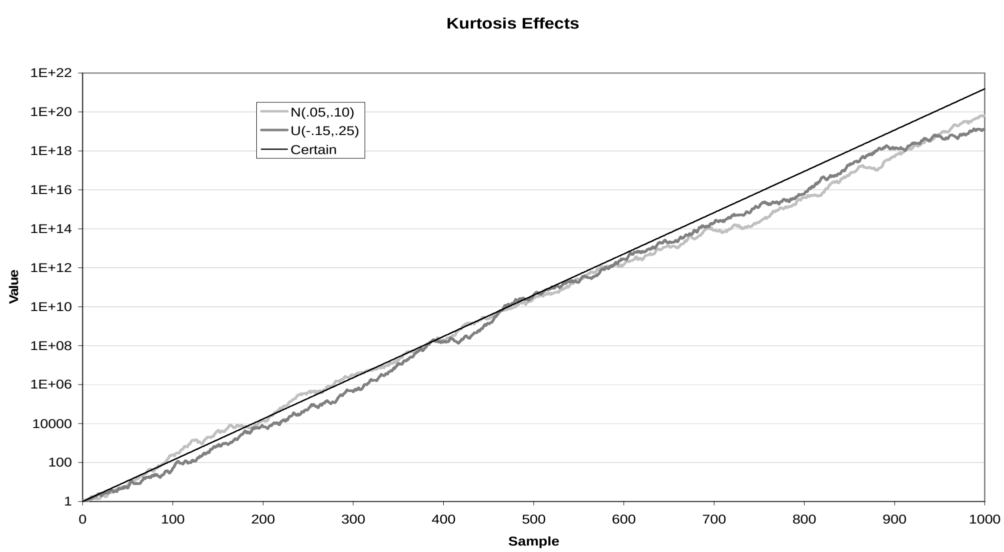

[8] A little caution is needed with respect to the interpretation of excess kurtosis as "fat tails". It is actually a rather trivial exercise to generate distributions which have excess kurtosis but are thin tailed relative to a Normal.

Though kurtosis is present in the illustration at levels which would be considered significant by practitioners, its effect is not as evident as the asymmetric effects shown earlier. There is perhaps some general tendency to under-perform relative to the Normal but this is not compelling.

Having seen the effects on wealth of symmetry and asymmetry in distributions of returns, we shall next consider the relation between gain and loss, if you will return and risk. If we first consider a uniform distribution and partition this at the mean of the distribution, it is evident that the relationship between gain and loss is linear in the sense that for every loss there is an equal and opposite gain. Somewhat more formally we can rotate losses or gains about the mean axis and cover the other. This does not however mean that we can map gains to losses in a one to one manner, merely that on average we expect a loss or gain of similar magnitude – these gains or losses are the conditional expected gains or losses. In the case of the Uniform these return values are -.05% and .15%. The slope of this gain-loss line when the partition is at the mean is unity. The process is a martingale, a fair bet. If we partition the gain-loss distribution about some other level, the relationship remains linear though now the slope of the line differs. If partitioned below the mean the slope is greater than unity and above less. An interpretation of this is that as we become less risk averse through the lower partition between gain and loss we stand to gain more. More risk equals more return in a linear fashion. If we were to partition the distribution at say the 25% percentile, (-.05%), the conditional expected loss is -.10% and the gain +.10% but there is three times the likelihood of a gain than a loss. The slope of the gain-loss function is now three.

About the mean of a uniform distribution, of course, losses and gains are equi-probable. In fact for any symmetric distribution partitioned about the mean the gain-loss function is linear. This is the world of Markowitz and Sharpe and much of modern financial theory, but it carries with it a consequence. If this linear world holds with respect to means, what are we to do when we have assets with differing means?

Let us consider the work-horse distribution of modern finance, the normal distribution. Again this is symmetric, so relative about the mean we have linearity but as we move the partition away from this level, we observe a much more complex situation. Firstly the gain-loss pay-off function relative to the 25% level is now 3.55 instead of the 3.00 calculated earlier, but more importantly the mapping from loss to gain has acquired non-linearity. This can be illustrated well in a Monte Carlo simulation. We draw at random with replacement 100 losses and 100 gains from the 25% partitioned Uniform sample and plot these against one another as in figure 5.

Figure 5: Gain-Loss relationship for Uniform distribution.

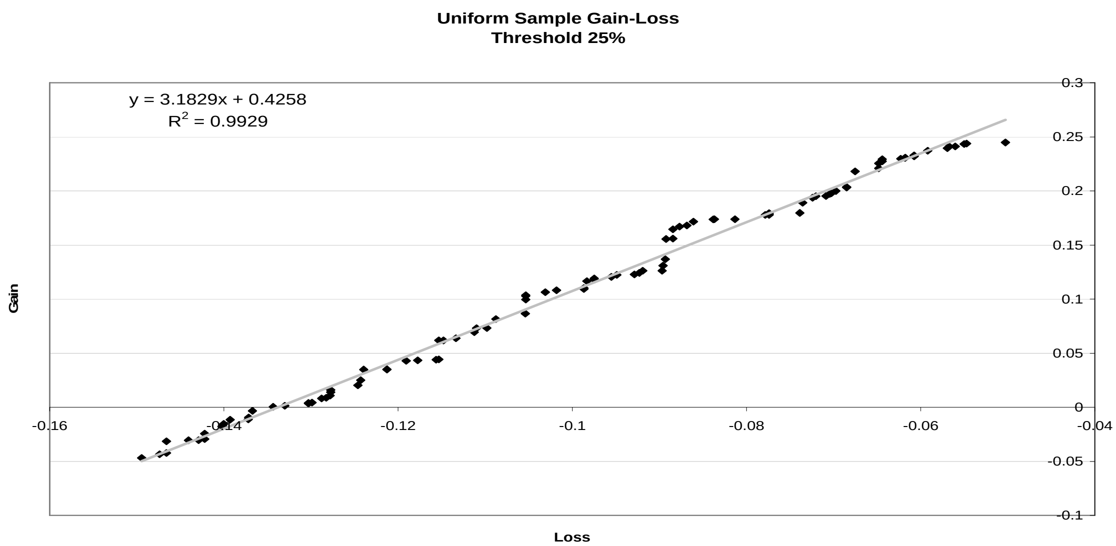

We have fitted a linear regression line to this and show the equation and R-squared for it on the figure. It is clear that the Uniform gain-loss relationship is well described as linear and that due to sampling error there is some departure from the expected threefold relationship. Next we perform the same actions for the Normal (0.05,.10) dataset used earlier and reproduce this as figure 6.

Figure 6: Gain-Loss Function for Normal with threshold 25%

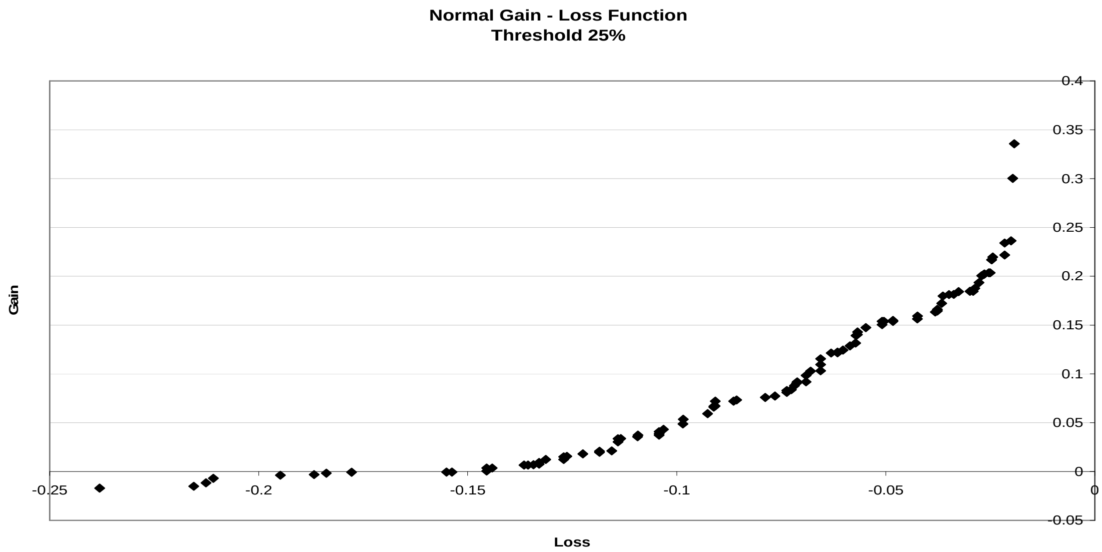

It is clear that the relationship between gain and loss is markedly non-linear and further that the interpretation of the 3.55 slope coefficient cited earlier is not obvious.

The practitioner and academic finance literatures are replete with cautions concerning asymmetric distributions, most notably concerning options. We choose to examine this with a sample from a Beta (5,50) distribution. The histogram of samples and density function for this is shown as figure 7. The skewness of this distribution is in fact quite modest at 0.69 even though the asymmetry is clearly visible.

Figure 7: Beta (5,50) distribution and sample histogram

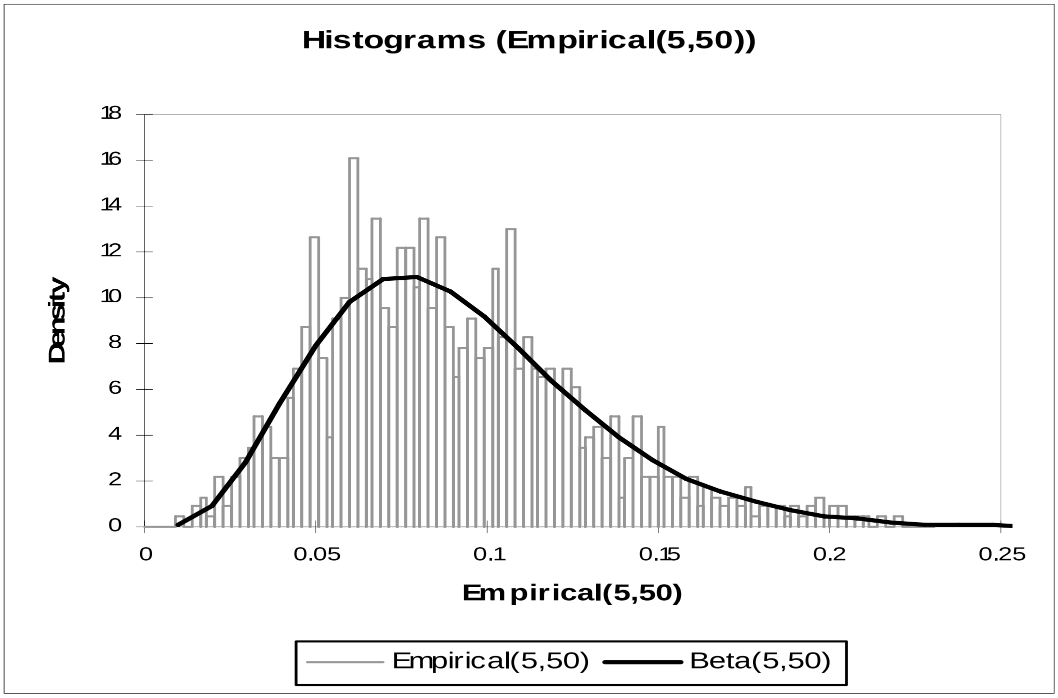

For this sample distribution 56% lies below the mean of 0.09, and 44% above. If we repeat the earlier sampling process considering gain and loss but here relative to the mean, then we observe a marked non-linearity which becomes even more pronounced for the 25% threshold. These two situations are illustrated as figure 8.

Figure 8: Gain-loss functions for mean and 25% threshold of sample from Beta(5,50) distribution.

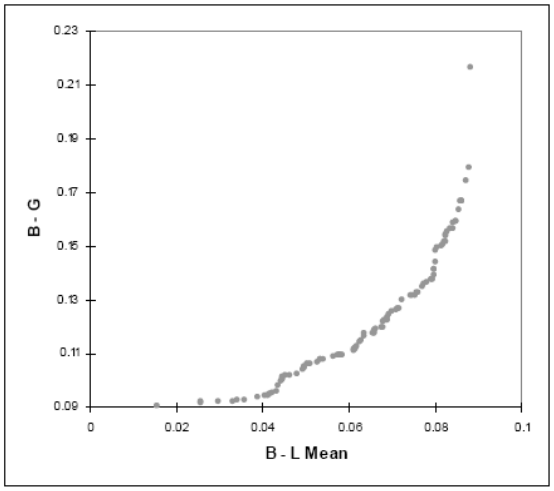

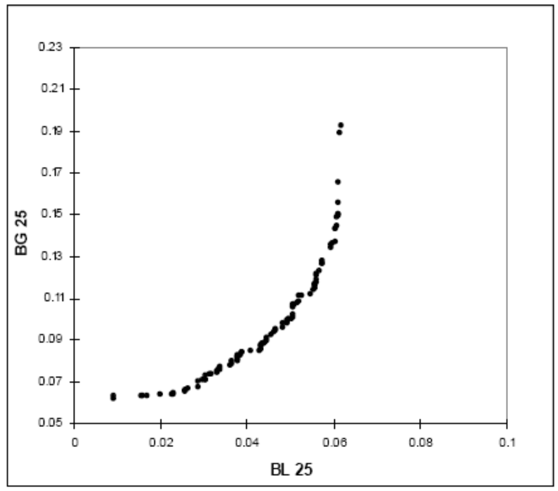

This analysis has considered gain and loss relative to a threshold with the relevant measure being return. Before considering alternate measures of risk to this, it is worth investigating briefly how diversification might function in this set-up. In standard modern finance diversification of a portfolio across assets allows the variance to diminish faster than return increasing the return per unit of risk. Perhaps a better way to understand what is going on is to consider the effect of diversification of the realised geometric return, the wealth or value function. We will use an equally weighted (rebalanced) portfolio of two assets which are themselves just samples from Normal (.05,0.15)and a Normal (.07,0.25) distributions and for further comparison a portfolio which was initially equally weighted but not rebalanced. This descriptive statistics for these assets and portfolios are shown as Table 2.

Table 2

	N(7,25)	N(5,15)	Rebalanced	Initial
Mean	0.0682	0.0499	0.0591	0.0546
St.Dev.	0.251	0.149	0.147	0.179
Skew	-0.245	-0.032	-0.097	-0.188
Kurtosis	0.045	-0.029	-0.186	0.588

The evolution of the value function for these assets and portfolios is shown below as figure 9.

Figure 9: Diversification and Strategy Effects

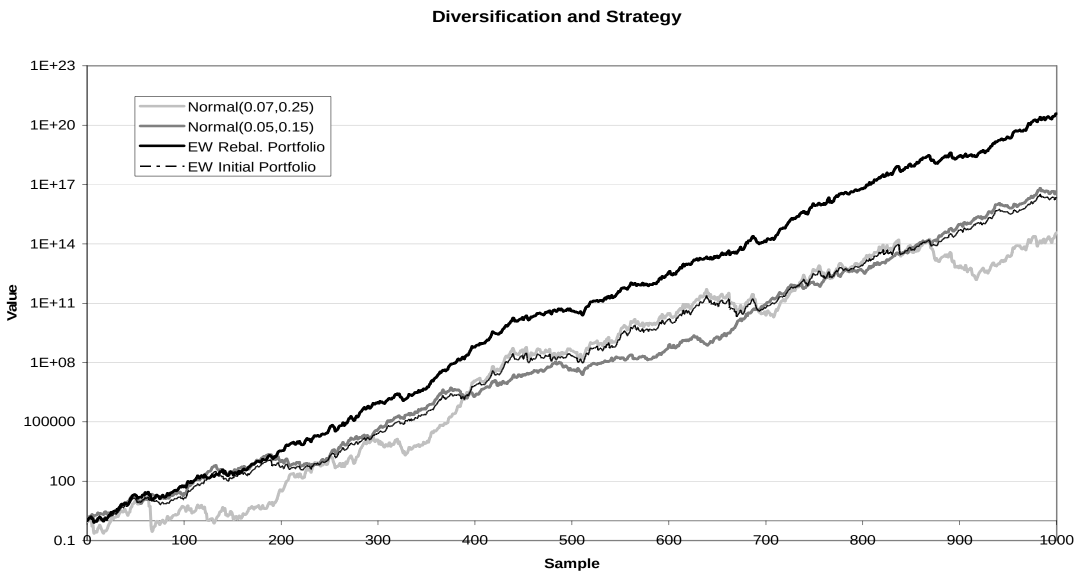

The effects are quite startling with the diversified rebalanced portfolio performing consistently above either of its constituent assets. If we examine the descriptive statistics of these assets and portfolio we observe that the arithmetic mean return of the rebalanced portfolio lies half-way between its constituent assets but the volatility is actually slightly lower than the volatility of its least volatile constituent. The performance evident here is (almost) entirely attributable to the lowering of uncertainty. The initially equally weighted portfolio by contrast tends to track the higher returning of constituent assets at any time, but its arithmetic mean is now significantly below the mid-point of the constituent asset returns and its volatility markedly higher than the rebalanced portfolio. This is actually quite a well known problem associated with indices, usually referred to as chasing the winners; in somewhat more technical terms the asset allocation arising from asset performance is not stationary. These weight variations are illustrated in figure 10.

Figure 10: Constituent asset weights for an initially equally weighted portfolio.

Constituent Weights Initial EW Portfolio

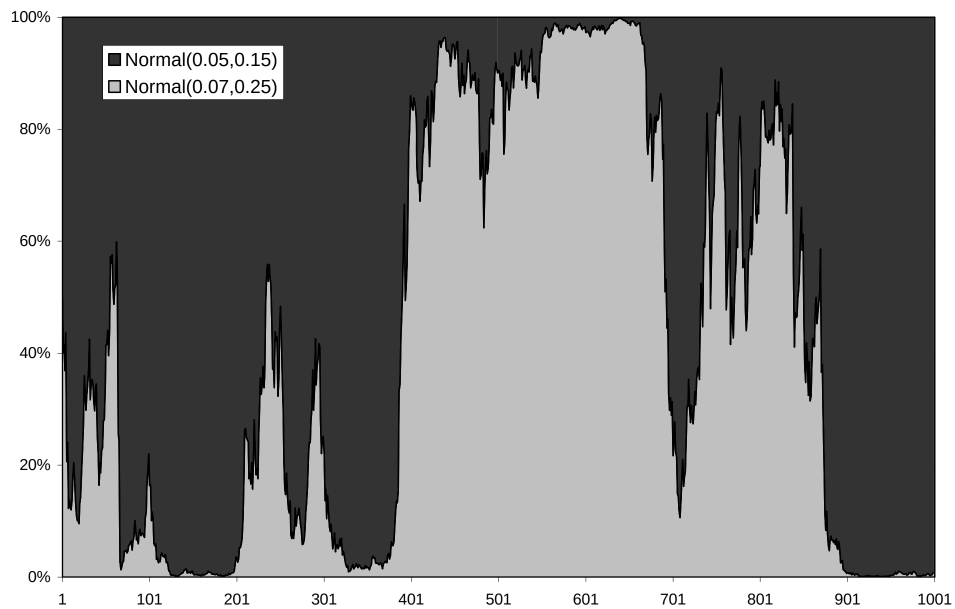

There has been considerable interest in creating portfolios which exploit variations on this problem in indices; most use fundamental financial analysis[9] rather than market capitalisation as the weight decision criterion. Another way of viewing the problem is that it is variation of the market timing problem for active managers, where it can be demonstrated that a fund manager wishing to outperform by selection of one or another asset finds it very hard indeed to outperform the simple rebalanced portfolio. For those seeking to perform performance attribution relative to a benchmark, there is also a problem in that this benchmark is itself not stationary. Incidentally this non-stationarity has consequences for the famous Beebower, Brinson and Hood study[10] which found that asset allocation could account for 93.6% of the variability of returns, which is widely misquoted among practitioners as asset allocation explains 93.6% of the returns of portfolios. This extension might perhaps be justified if the relation between variability and return were linear. The most important lesson here though is that portfolio strategies now matter. It is also evident that some of the standard tools of financial analysis are of little help. By way of illustration we show as figure 11, the first two principal components of the returns series for the two assets, the rebalanced and the initial portfolio.

---

[9] The best known name in this field is Robert Arnott who has published widely on the subject, perhaps most notably in the Financial Analysts Journal.

[10] Beebower, Brinson & Hood, "Determinants of Portfolio Performance" FAJ 1986

Figure 11: First two principal components of returns series

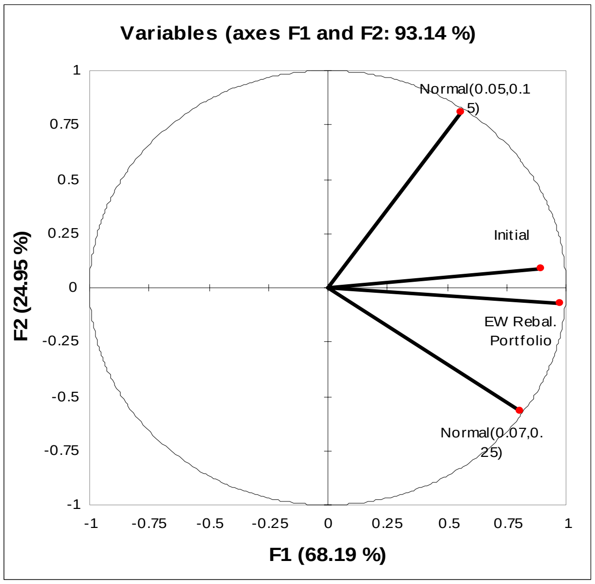

The two portfolios look very similar in this decomposition which explains 93.2% of overall variance. The two assets can be seen to be very close to independent; they are approximately orthogonal.

We have often heard it said by academics and practitioners that portfolio returns should be presented in a risk-adjusted manner, which would require the use of some risk measure. For most the default risk measure is the standard deviation or volatility of the returns series; this is implicit in the information ratio and the Sharpe ratio. These differ only in the presence of the "risk-free rate" in the Sharpe formulation. This rate, however, serves a significant purpose; it resolves the difficulty of scale invariance of the information ratio. By this we mean that, by the information ratio, two series which are simply rescaled variants of one another in respect of say their mean and variance are equivalent, but when we look to their geometric means we see that the series differ with the larger scaled version being superior. Given the significance of leverage, or rescaling of returns, in hedge funds, this is a material concern.

Early in this essay, we made the point that risk and dynamics are integrally related, which raises the question why the market standard measure of risk has become the standard deviation of the series distribution when in mechanics and engineering we would move directly to the derivatives of the process, rather than higher moments. In fixed income analysis we do precisely this with the dollar duration and convexity measures. There is, in fact, a mathematical duality between the moments and the derivatives. This becomes evident when we look to the conditional expected loss (relative to the mean) of a normal distribution and notice that this is $-\sigma\sqrt{2\pi}$, or that the derivative of the Omega function at the mean is $\frac{-\sqrt{2\pi}}{\sigma}$. The conditional expected loss has several other possible interpretations; it is the pay-off to an option with strike at the mean or alternately we might think that it is a natural choice of quantile for value at risk.

There are endless debates in the practitioner literature as to the deficiencies of variance based measures; this is mirrored in portfolio allocation optimisation models which account for higher moments. It is perhaps helpful to consider these, but better is simply to look to the symmetry or asymmetry effects – to consider the situation holistically. This is what led us to the Omega function[11] and indeed to tests for symmetry and entirely new, but very natural, classes of distribution.

The Omega function is defined as: $\Omega(r) = \frac{\int_r^{\infty} 1 - F(x)\,dx}{\int_{-\infty}^{r} F(x)\,dx}$ with obvious notation. There are many alternate formulations, in terms of for example upper and lower partial moments. Its financial application is evident that this is the ratio of a call to a put option with strike at the partition or threshold. Interestingly it shows that put-call parity arises as a property of probability distributions rather than the economics of Black and Scholes. The function is a pay-off ratio for all strikes and as such offers insight into the shape of the expected utility function implicit to a portfolio or asset. The attraction of the function is that it contains all information which is contained in the distribution; this means that it contains information concerning all moments. Clearly it does not contain any time serial or dynamic effects other than in the sense that their historic effects are present in the

---

11 Further references and papers are available from the author.

distribution. The appropriate risk measure is the proportional rate of change of the function, and this is local and global. Comparison of Omega functions requires comparison of the entire function, which may on occasion, most notably when there are double crossings of the Omega functions of security returns series, prove difficult by visual inspection. In such situations it is necessary to resort to a metric which maps the function to a value on the real line; there are in fact infinitely many possible such mappings but perhaps the most obvious, with some abuse of notation, is:

$$M = \int_a^b \Omega(x) f(x) dx \qquad [2]$$

In discrete form this is simply the sum of the likelihood adjusted pay-offs. There are some minor sampling problems to be overcome with this concerning the distribution range end-points in order to implement it. There have been many attempts to optimise portfolios based upon the Omega function; most have used single thresholds, a few have used two or three thresholds. It is comparative trivial to optimise using metrics and in fact these are closer to the tradition of mean variance optimisation in the sense that they consider the situation globally.

Before moving on to consider benchmark relative performance attribution, we should make the point that if we are considering assets or liabilities then we must recognise that the first order measure of risk on this as would be obvious from the dynamic perspective, its proportional rate of change, is better known as return – which of course motivates the entire gain-loss analysis.

The past few years have seen great interest develop in a new form of benchmark relative performance, but instead of an index or asset portfolio, liabilities are used; this is Liability Driven Investment (LDI). The performance of a portfolio relative to a benchmark is well trodden territory, with the usual measure the Tracking Error. Unfortunately this is usually defined as the standard deviation of the difference portfolio - that is asset returns minus benchmark returns. This is adequate if there is no expected return, positive or negative, as was the case in the origin of the measure - index portfolio replication. However, in active management it becomes necessary to include the difference in expected return also and the correct formula[12] becomes:

$$T.E. = \sqrt{\mu_{diff}^2 + \sigma_{diff}^2}, \text{ with obvious notation.}$$

A consequence of this is that there are bounds to the returns which can be achieved for a given tracking error constraint; impossible mandates have been observed. A 3% return cannot be achieved with a 3% tracking error constraint other than deterministically.

To consider liability driven investment, we need first to understand how an asset differs from its corresponding liability. This is not a question of translation along the real line, a simple relocation of the distribution, but rather of rotation around the zero axis; the effect of this is to reverse the signs of symmetry. A left skewed asset is right skewed as a liability. If preferences over magnitudes of gains and losses do differ in level, this simple fact may have quite profound consequences for valuation. More importantly we can see this problem of matching assets to liabilities as analogous to the lottery problem used extensively for pedagogic purposes in earlier Omega papers. Asymmetric payoffs such as purchased and sold lotteries are in fact only coincident in value, or Omega function, at their mean; everywhere else they differ.

The next problem with LDI arises from the fact that this operates in current values. Assets are marked to market while liabilities are discounted present values of future cash-flow estimates. The use of market prices for assets is stochastic; in theory these prices consider all future states of the uncertain world, but liabilities, for example in pension application, use a single rate, albeit market determined. FRS 17, the UK accounting standard specifies the AA corporate bond rate. This is clearly deterministic and the earlier discussion of the certainty equivalent premium should make it obvious that this mixed attribute nature of the valuation standard introduces bias and error. If we move to use a yield curve, then the problem remains; this is still deterministic though of higher order than the simple one rate for all.

---

[12] Derivation available on request from the author.

There is also the question of risk premia. It is true that £100 of bonds is equivalent to £100 of equities. The future cash-flows from the equities are, however, higher than those from bonds at the horizons of liability cash-flows – the probability measures under which the cash-flows from equities and the cash-flows from bonds are discounted differ. If we consider the dynamics of the situation we also observe a term structure of volatility[13], where as holding period increases bond investments prove more volatile than at short horizons while the converse is true for equities. Perhaps more than anything this appears to be an instance where the nomenclature is influencing the analysis; if holding an asset for a longer period increases the return achieved[14] , or lower the volatility, is it really sensible to describe the difference as a "risk premium"?

It seems that there is much misguided research in the field of asset and liability management; the problem appears to be much misunderstood. The problem is actually one of coverage and is perhaps best simply and diagrammatically illustrated.

13 The existence of a term structure of volatility means that, in particular, there is one approach to asset allocation which seems suspect, risk budgeting; this seems to have arisen as returns are themselves rather difficult to forecast. Given that returns, that is $E(r)$, are problematic to forecast, it is difficult to understand why risk, that is $\sqrt{E(r^2)}$, should be any more predictable, other of course than the trivial it lies only on the positive half-line.

14 It is worth noting that if we consider markets as American options on liquidity, we would predict such effects. There really is need for a coherent theory of liquidity in the context of asset prices and markets – Holmstrom and Tirole's *"Liquidity Asset Pricing Model"* might perhaps be seen as a first step in this direction.

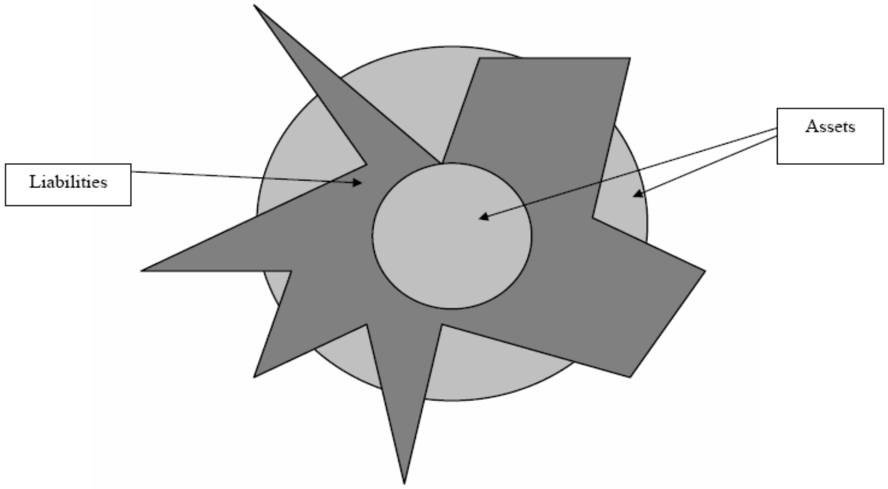

The diagram above shows liabilities in dark grey; the area is total cost. Obviously liabilities in reality are multi-dimensional, rather than this simple two dimensional surface. Much effort is currently being expended in defining on defining the boundaries of liabilities. Much can also be futile, for example measuring the perimeter of such a complex shape tells us little about its contained area[15]. Assets are shown in this diagram as a simple model, a circle. Clearly the inner, small circle is insufficient for most purposes, the larger circle is perhaps sufficient though it leaves some extremities uncovered and covers other areas unnecessarily, and of course it is possible to have the liabilities entirely contained within a very large circle, though that of course would have many areas of redundancy. Perhaps an attraction of this simple circle model is that here the circumference can tell us precisely the area or total cost contained within it. This does make obvious that if we are to consider coverage at minimal cost we must start with assets and liabilities which share a common location. This simple diagrammatic representation should though make evident the fact that there are trade-offs between cost and coverage that are far from simple.

---

[15] This is conceptually related to the Holder exponent (H) and Mandelbrot's concept of box dimension. The latter is the number of square boxes of dimension $dt \cdot dt$ required to cover the object – so an object of area A has dimension 2 , while a continuous line has dimension 1, and so on.

If we are to use models for asset behaviour, in asset allocation or performance attribution, we must recognise that the degree of complexity with which they are specified also limits the extent to which we can use them for the analysis of their dynamics or risk.

In many instances we are using attribution analysis to inform and guide management action, which may itself be problematic. The overwhelming majority of financial asset models assume that risk is exogenous, for example all prescribed regulatory evaluations. Whether we carry an umbrella against rain is an illustration of good management of an exogenous risk; carrying the umbrella does not affect the likelihood of rain though it does mitigate the consequence. These are games against nature but financial markets are far from completely such games. They are, in fact, mixed games, partly against nature but predominantly against others. Our management actions and strategies do affect the outcomes and in particular strategies matter.

Having questioned the applicability of the linear model, it is worth considering the alpha-beta distinction currently in favour in asset allocation; these arguments are usually presented as why pay active management fees, that is alpha generation fees, when what is delivered is market related, or beta, and cheap to obtain through replication strategies. In terms of the economics the distinction between these is determined by the extent to which higher asset prices affect future consumption prices – and this is usually considered small. If, by contrast, we consider the true situation faced by an open pension scheme we notice immediately that future contributions can be large relative to the currently held portfolio, the endowment. In this situation the pension scheme is a consumer of future investments. If we want to consume a hamburger every day of our lives, do we want the market price of hamburgers to rise or fall? The objective function of the pension scheme is no longer the unconditional maximise wealth at the immediate future if this has a cost because of higher asset market prices at the time of future investments. However returns independent of the market, one definition of alpha, are always attractive. This is an aspect which is independent of the linearity of the model.

Perhaps we would all do well to recall the econometrician George Box's observation: "All models are wrong, but some are useful".