# A Primer on Omega Metrics Tail Risk Analysis

*A New Risk Management Tool*

Effective Risk Management depends on the ability to differentiate between acceptable and unacceptable downside exposure or 'tail risk'. Hedge fund returns typically have characteristics which make this essential task difficult, if not impossible, to carry out with standard statistical tools. The assumption that returns are well modelled as samples from a normal distribution may lead to serious under-estimation of the potential for large losses. Managing hedge fund risk requires tools which do not depend on assumptions of normality and which give robust signals even from short data samples.

Omega Metrics Tail Risk Analysis makes use of new statistics which provide a 'fingerprint' of a distribution from as few as 18 months of history. This allows both peer group comparison and 'style checks' against the fingerprints of hedge fund indices for the same investment strategy. Historic tracking of the tail risk of funds highlights increased or decreased risk which may be correlated with market events or with changes in trading strategy or effectiveness.

*Omega Metrics Tail Risk Analysis is A Practical Way to Measure and Monitor Tail Risk*

Tail Risk is present in any investment portfolio. To manage this risk, it is essential to be able to measure it and to be aware of any significant changes in it over the life of the investment. Omega Metrics Tail Risk Analysis takes an investment's tail risk fingerprint and tracks its variation over time. A risk manager may use this to assess the degree to which a fund's tail risk is changing with market conditions or as a result of changes in trading strategy, for better or worse. It may also be used to identify under or out performance relative to competing strategies or to flag risk which is inconsistent with that exhibited by a peer group or proxy such as an investable hedge fund index or other benchmark. As our analysis of the Amaranth Fund shows (page 3), this can provide significant advance warning of changes in the risk taken on by a manager.

*Four Categories of Tail Risk*

It is common practice, even in dealing with hedge funds, to use units of standard deviation–as in 'a 2 sigma loss'. The probability of such a return is assumed to be the probability of the same event in normally distributed returns–in this case about once every 43 months. In the absence of any information to the contrary, if an assumption must be made about a 'population distribution', the assumption that returns are normal is the natural one. The problem is that returns might well be better approximated as draws from a distribution with 'fat tails' in which the probability of a 2 sigma loss is three or even ten times greater than it would be for a normal distribution. Something that had been expected only three times in a decade might occur more than once a year as a result. Knowing when the tail risk exhibited by a manager falls inside or outside what could be expected from a normal distribution is therefore critical to successful risk management.

*Omega Metrics Tail Risk* is measured by the value of the C-S Character (see Appendix) of the historic returns distribution. This statistic, which provides an extremely robust picture of tail risk, even from small data samples, is the basis for our classification of tail risk into four categories: Normal, Moderate, High and Extreme.

Tail risk which is statistically indistinguishable from that of normally distributed returns is classified as 'Normal'. Returns distributions which exhibit this sort of tail risk are the ones in which standard mean/variance techniques are least likely to fail to adequately assess risk. Portfolios of Normal tail risk assets, and well diversified portfolios in general, tend to exhibit Normal tail risk. Well constructed hedge fund and fund of hedge fund indices may fall into this category over extended periods. Developed equity markets may also spend long periods in the Normal tail risk category but routinely exhibit more tail risk than is likely in normal distributions.

The next category, Moderate, is close to Normal but includes Tail Risk Index values which are either sufficiently high or low that they would not be found persistently in normally distributed data. Hedge fund and fund of hedge fund indices, together with developed equity markets, usually remain within the Moderate category when they depart from Normal. Assets with tail risk in this category often exhibit significantly higher probabilities of extreme events than would be likely for normally distributed returns. Because returns in the Moderate category are less likely to be normal, the use of mean/variance analysis on such assets may introduce serious errors.

Hedge fund style indices such as Emerging Markets can exhibit higher than Moderate tail risk, as can equity markets and hedge fund indices in certain market conditions. The High tail risk category includes returns distributions for which the probability of a given loss may exceed that in a normal distribution by a factor of ten or even one hundred. High tail risk may be accompanied by high absolute or risk-adjusted returns, and may be sought by some investors while it would be avoided by others. In either case, it is crucial to know if High tail risk is a feature of an investment. Any assumptions about downside exposure which depend on returns being approximately normal are likely to fail catastrophically for assets which exhibit High tail risk.

In unusual market conditions, such as the Summer of 1998, hedge fund indices of all styles can exhibit tail risk which exceeds the High category and, for short periods, reaches a level which we have classified as Extreme. This level is so high that it is consistent with sampling from distributions which have wild variability. The threshold of the Extreme tail risk range is visible in Figure 1 which shows a time series of Tail Risk Index values.

*Examples of Tail Risk For Hedge Fund and Market Indices*

As might be expected given the diversification over many individual hedge funds and also over hedge fund styles, the Credit Suisse Tremont Investable Index and the HFR Fund of Funds Indices exhibit close to Normal tail risk from January 2004 to November 2006. The S&P 500 Index generally has Moderate tail risk while the CS/T Emerging Markets Index has tail risk trending upward steadily, and remaining in the High category since mid 2005.

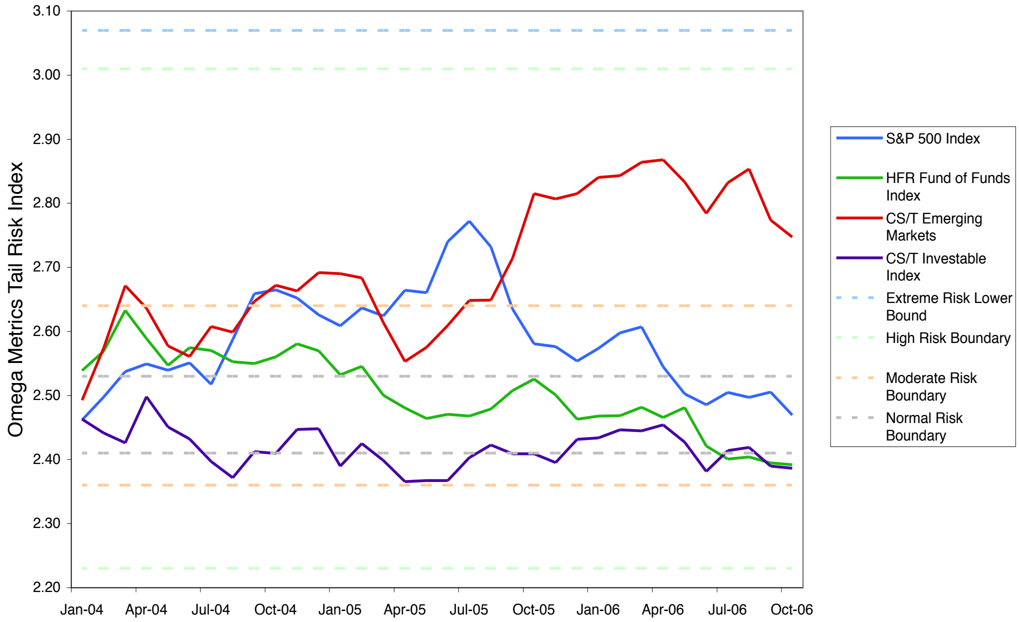


Figure 1. Tail Risk Category Bands. For most of the period from January 2004 to November 2006, hedge funds have shown less tail risk than the S&P 500 Index. Since Q4 2003, the Emerging Markets style index tail risk has trended upward, reaching the High risk category in September 2005.

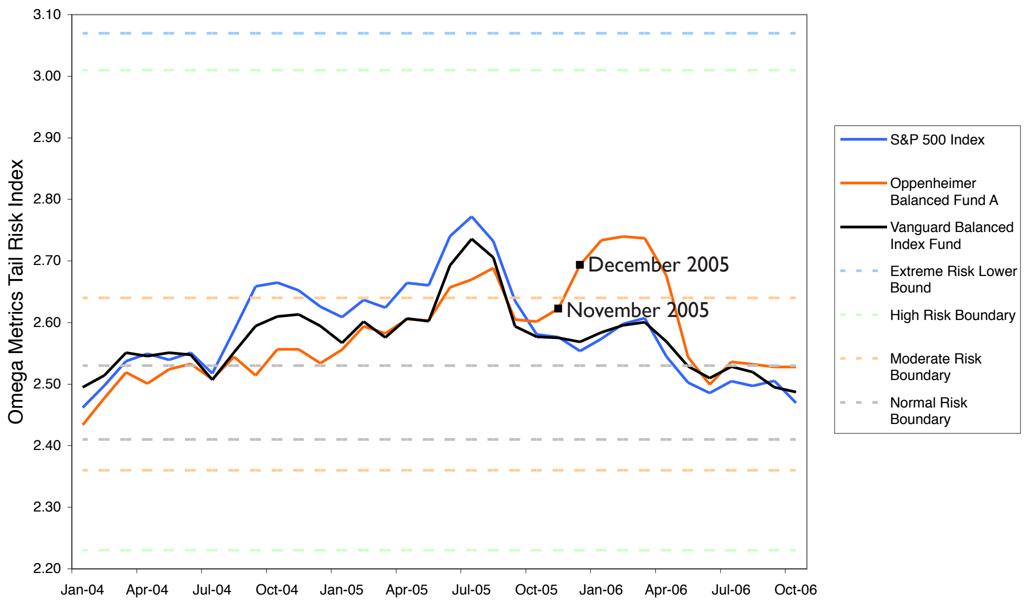


Figure 2. Tail Risk history for two Balanced Equity/Bond Mutual Funds and for the S&P 500 Index. The Tail Risk Indices for the funds and the S&P 500 Index move in lock step until November/December 2005.

*The Amaranth Fund*

The severe losses of the Amaranth Fund in September 2006 have prompted a number of retrospective analyses. The data history for Amaranth shows anomalous returns in April and May of 2006. Depending on lock-up provisions, warnings of trouble by early to mid May, when the April results were first known, might still not have provided sufficient motivation to secure an exit before the September 2006 collapse of the fund. Early exit penalties also complicated the decision.

Omega Metrics Tail Risk Analysis provides risk managers with a robust tool for detecting increases in tail risk at an early stage. In the case of the Amaranth Fund, this warning came more than 6 months before the April 2006 return. As Figure 3 shows, the Amaranth Fund had exhibited Normal tail risk since June 2002 but by April 2005 had risen to Moderate. Tail risk then began to rise steadily in the third quarter of 2005 until, by September 2005, the fund entered the High tail risk category for the first time since 2001. This coincided with the higher profits and volatility from the increased concentration on gas trading. Whether or not the Amaranth Fund investors were aware of the significance of this shift in strategy, the increased tail risk would have been very apparent to a risk manager using Omega Metrics Tail Risk Analysis.

With the April 2006 return, the Fund was well above the tail risk of even the Emerging Markets style index and was close to the *upper* boundary for High tail risk. In May it entered the Extreme tail risk category from which it emerged only briefly in August 2006 before going off the scale of Figure 3 in September 2006.

A risk manager who was aware of the increased tail risk carried by the Amaranth Fund since September 2005 would likely have been advocating exit after April 2006, even with an early exit penalty.

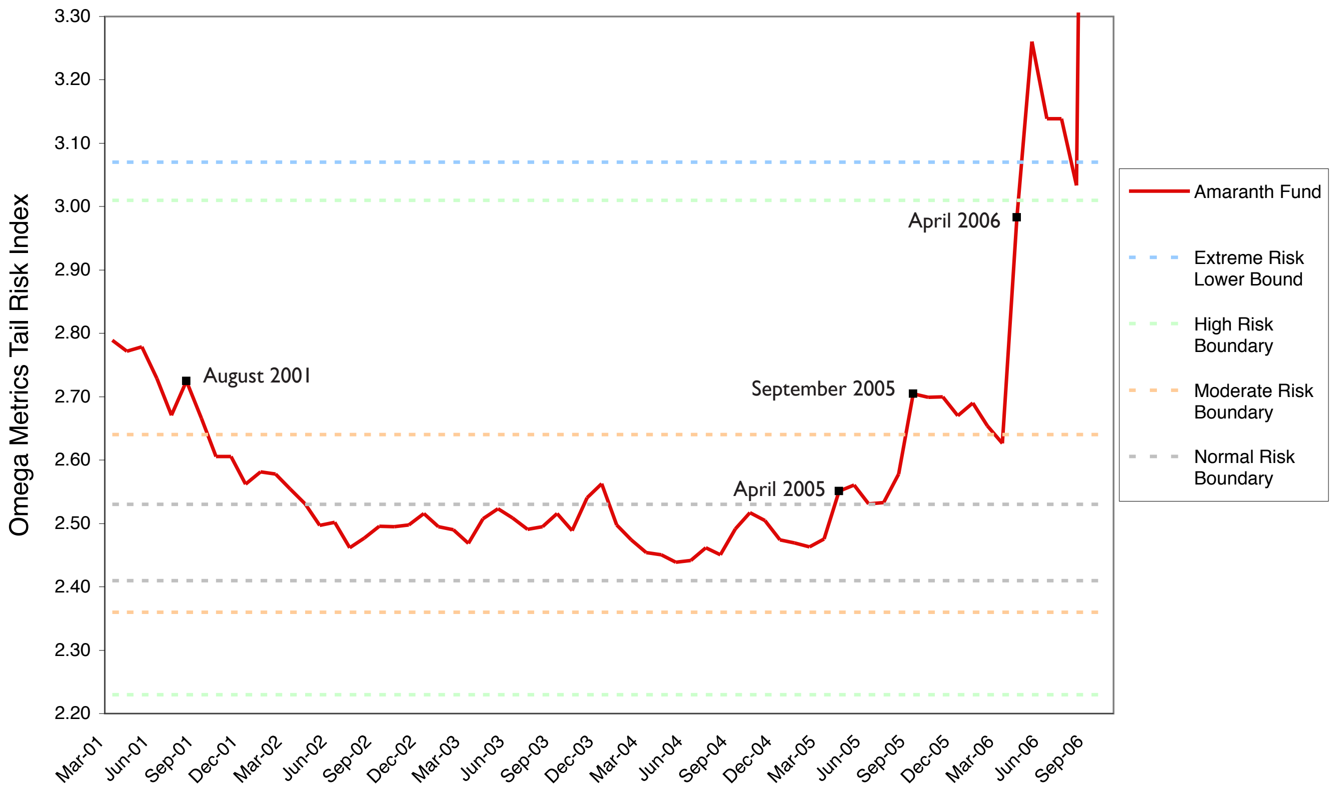


Figure 3. Tail Risk history for the Amaranth Fund. After largely Normal tail risk in 2002 through March 2005, tail risk rapidly rose to a level last observed in 2001. By September 2005 the Fund had moved to High tail risk, 6 months in advance of the anomalous returns in April and May 2006.

*Statistical Analysis of Sample CS-Characters*

The returns from an investment fund are not random although they are subject to random events. This means that the use of a mathematical model in which returns are treated as random variables and statistical analysis is employed is the only practical approach to quantitative risk management. If this is to be effective, it is essential that the statistical tools employed be well suited to the data which is available.

In the case of hedge funds the data is monthly returns. Regardless of how much returns history is available, it is rarely the case that market conditions or external factors remain constant. Whatever period is used, it will necessarily be short and there will always be sampling error as a result, even if the fiction that returns are independent draws from a fixed population distribution were correct.

The C-S Character is defined for any distribution which has a finite standard deviation and it has remarkable sampling properties. The Omega Metrics Tail Risk Index levels are the result of extensive analysis of the sampling distributions of the C-S Characters. These provide very good accuracy even from small samples of the sort routinely available from hedge funds. With draws of only 48 points from a normal distribution for example, over 80% will have C-S Characters which are within plus or minus 5% of the true value.

One might also try to use kurtosis to measure tail risk, however its sampling distribution shows that kurtosis is much too variable to be relied on in typical hedge fund data sets. For example, in samples of 48 points from a normal distribution, the error at the 80% confidence level is 25% of the true value. For distributions with fatter tails than normal, estimates of kurtosis from small samples such as this are so volatile as to be completely useless. This is a very serious objection to the creation of 'synthetic hedge funds' based on matching mean, variance, skew and kurtosis.

Figure 4 shows the extreme variability of kurtosis compared with the C-S Character for data drawn from a normal distribution.

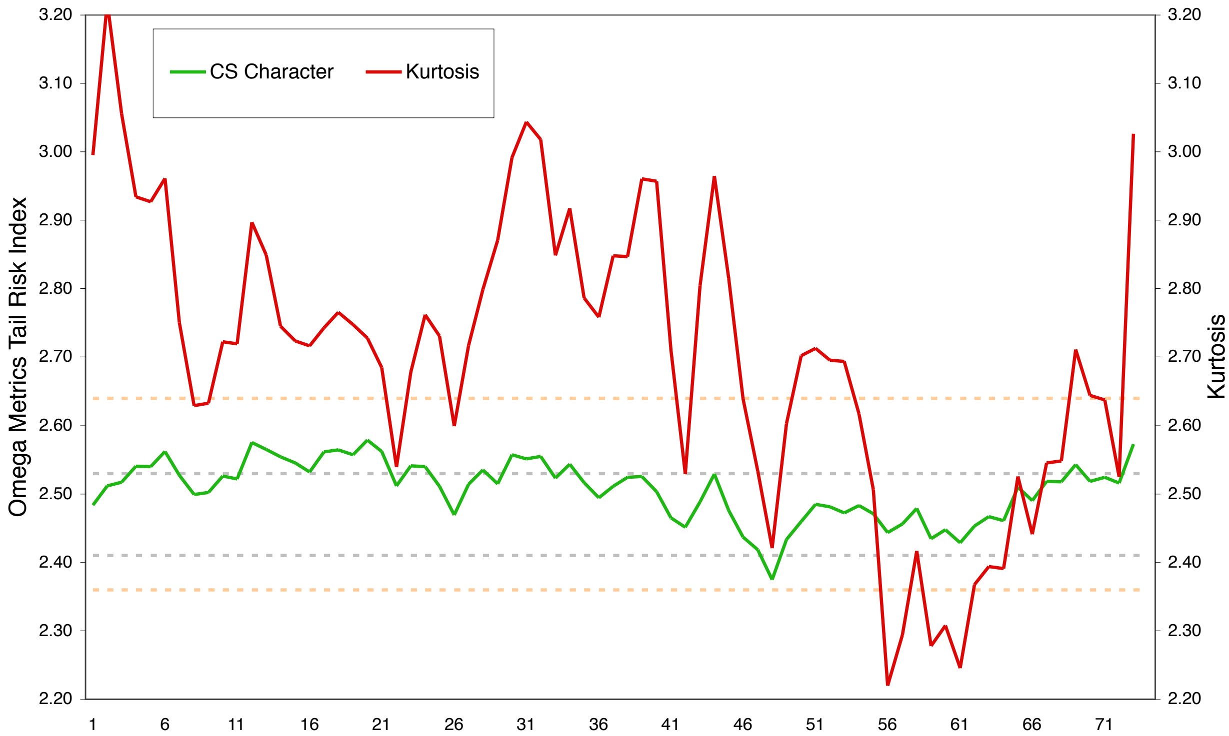


Figure 4. Time series of C-S Characters and Kurtosis for samples from a normal distribution. The variability of the C-S Character within the Normal and Moderate tail risk bands is typical: less than one third of the values are in the Moderate bands and over two thirds are within the Normal band. While the tail risk bands are irrelevant for Kurtosis, the diagram shows the extreme variablity of this statistic compared with the C-S character. In this sample, the Character and Kurtosis have correlation of 0.75.

*To Subscribe to Omega Tail Risk Analysis or for more information contact:* TailRisk@OmegaAnalysis.com

Appendix The C-S Characters and Tail Risk

The C-S Characters of a distribution (Reference 1) are statistics which are independent of scale and location, so for example, every normal distribution, regardless of its mean or variance, has the same C-S Characters. In this case the 'population value' of the first C-S character is approximately 2.51, just as the kurtosis of any normal distribution, independent of mean or variance, is 3.

The first C-S Character is defined as the ratio of the standard deviation of a distribution to its 'om', a statistic which is defined in terms of the slope of the Omega function at the distribution's mean. The om is also equal to half of the mean absolute deviation.[1]

The om is a direct measure of the degree of concentration of returns around the mean–so when two distributions have the same standard deviation, the one with the smaller om will be more concentrated around the mean. For a fixed standard deviation, the smaller the om the greater degree of concentration about the mean and the higher the C-S Character. Financial returns distributions whose C-S Characters are above 2.51 are generally more concentrated about the mean *but also have fatter tails* than a normal distribution. Figure A.1 shows an example with analytic distributions formed by mixtures of normals, for which the results may be calculated exactly. It shows the standard normal distribution (mean 0, standard deviation 1) in green and another distribution with the same mean and standard deviation in blue. For the normal, just over 68% of the probability is concentrated between -1 and +1 – i.e. within one standard deviation of the mean. For the blue distribution this rises to 77%. The trade-off is a higher probability of events far from the mean as the fatter blue tails in Figure A.1 indicate. For example, the probability of a return more than 3 standard deviations below the mean is 0.8% –six times as large as for a normal. The C-S Character of the blue distribution is 2.84, significantly higher than the Character for the normal distribution.

Conversely, a lower degree of concentration about the mean corresponds to lower C-S character and is accompanied by tails which are thinner than those of a normal distribution. A population value of the C-S Character below 2.51 is the statistical indicator of this behaviour. Figure A.2. again shows the standard normal distribution in green. The blue distribution has only 64% of its probability within one standard deviation of the mean. Its C-S Character is 2.34.

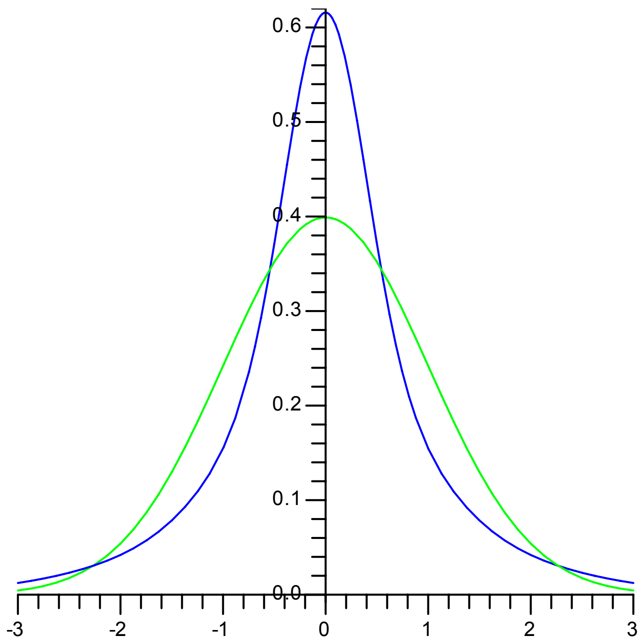


Figure A.1 The standard normal distribution (green) and a distribution with higher than Normal Tail Risk Index of 2.84 (blue).

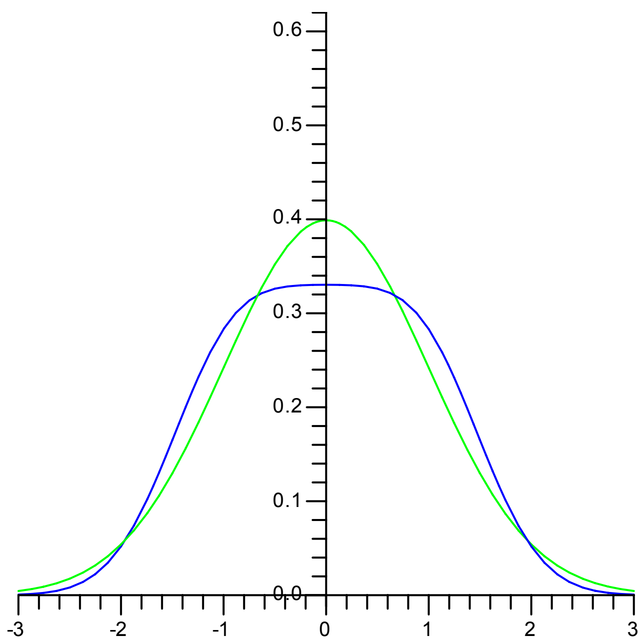


Figure A.2 The standard normal distribution (green) and a distribution with lower than Normal Tail Risk Index of 2.34 (blue).

If your analysis depends on the assumption that returns are normal, both higher than Normal and lower than Normal tail risk values can lead to serious errors. If the C-S Character is above 2.51, this arises because events far from the mean may be much more likely than they are in the normal case. In the case of a Character value below 2.51 it may happen that a string of returns at the lowest end of those historically observed and a string of returns around the mean are equally likely. For the distribution in blue in Figure A.2, a run of losses between 0.75 and 1 standard deviations below the mean is almost as likely as a run of the same length around the mean. For the normal it is only two thirds as likely.

The Extreme tail risk category arises from distributions which can exhibit wild variability. Figures A.3 and A.4 show an example which combines features of those in Figures A.1 and A.2. This distribution has fatter tails than the blue distribution in Figure A.1 but has less of its probability mass between -1 and +1 than the blue distribution in Figure A.2.

The Extreme tail risk in this case is due to infinite standard deviation. The distribution's om is finite however, so we have a natural way to compare it with a normal distribution even though we can't do this through matching standard deviations. The comparisons in Figures A.3 and A.4 are made by matching the om of the red distribution to the om of the standard normal distribution. This is, in a precise mathematical sense, the best basis for comparing the two distributions. While units of standard deviation are meaningless for the red distribution, units of om make perfect sense. The probability mass within 2.5 oms of the mean for the normal distribution is approximately 68% while for the red distribution it is only 43%. The probability of a return 7 oms below the mean for the normal distribution is only 0.3% while it is 13%– over 40 times greater for the red distribution.

While the population value of the C-S Character for this distribution is infinite, samples of any size drawn from this distribution will have finite standard deviation and hence finite sample C-S Character. The Extreme tail risk threshold is defined by the requirement that there be a significant chance that it would be observed in samples from a distribution such as that shown in red in Figures A.3 and A.4. Samples from such a distribution may exhibit swings of arbitrary magnitude. An investment whose returns behave in this way must be sized with extreme care to avoid large losses.

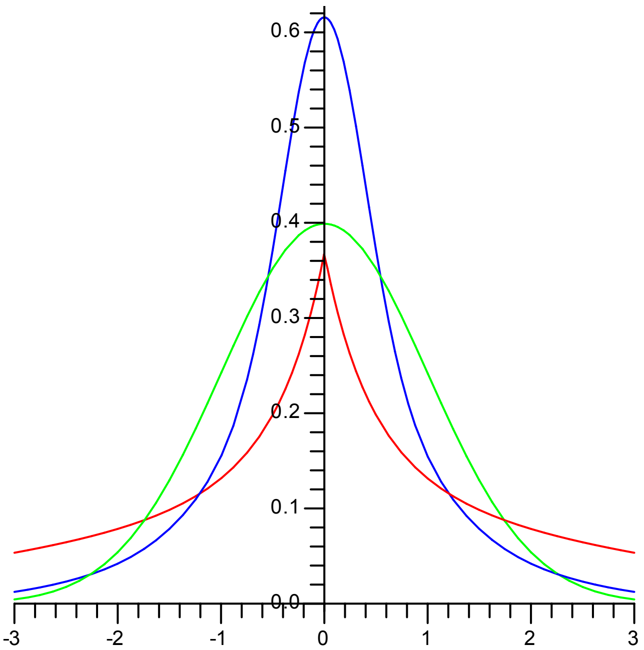


Figure A.3 The Distribution in red has fatter tails than either the normal or the blue distribution.

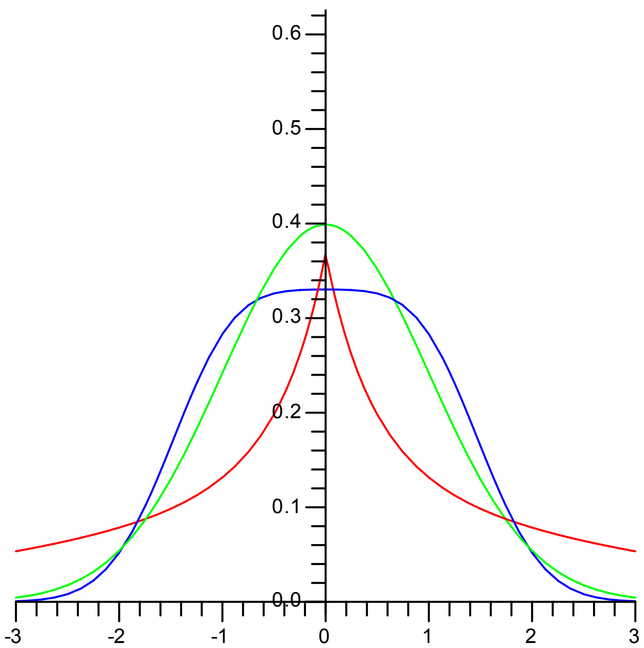


Figure A.4 The Distribution in red has less of its probability mass between -1 and +1 than either the normal or the blue distribution.

[1]The om of a distribution is given in terms of its Omega function by $\omega = \dfrac{-1}{\Omega'(\mu)}$. This statistic and those created from it have many remarkable properties (see References 1&2).

The exact value of the first C-S Character for a normal distribution is $\sqrt{2\pi}$ .The second C-S Character is defined as the ratio of the distribution's second L moment to its om. This statistic is defined for any distribution which has a finite mean. For a normal distribution it is $\sqrt{2}$ or approximately 1.41.

Reference 1: *The C-S Character and Limitations of the Sharpe Ratio in Financial Data*. Ana Cascon and William F. Shadwick, Journal of Investment Consulting, Vol 8, Number 1, Summer 2006, pp 36-52.
Reference 2: *The Standard Dispersion and its Applications to Portfolio Management*. Ana Cascon and William F. Shadwick, to appear in The Journal of Investment Consulting, 2007.