INTRODUCING THE NORMAL EQUIVALENT

TO EVALUAUTE HEDGE FUND

PERFORMANCE

By

MATTHEW EDWARD HOOD

A dissertation submitted in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

WASHINGTON STATE UNIVERSITY
Program in Business Administration

MAY 2006

To the Faculty of Washington State University:

The members of the Committee appointed to examine the dissertation of MATTHEW EDWARD HOOD find it satisfactory and recommend that it be accepted.

Chair

ACKNOWLEDGMENTS

A great deal of effort and sacrifice has gone into completing my degree and dissertation. Certainly not all by me. My family, friends, cohorts, and faculty deserve more appreciation and recognition than will be shown them. My wife, Susan, has supported our family and grown to meet new challenges. I love her. My mother and father gave me every opportunity to succeed; I hope they are proud of what I have done with them. I have asked for the help of my committee and especially my committee chair, John Nofsinger, on short notice many times and they always came through for me. I received good advice from everyone around me and have always looked forward to seeing those that I associate with. I am reminded of a scripture in the Old Testament where Samuel sets a stone to remember the help of the Lord in their victory. He names it Ebenezer, *stone of help*. I'll keep a stone from the Palouse to remember that I could not have done this by myself.

INTRODUCING THE NORMAL EQUIVALENT

TO EVALUAUTE HEDGE FUND

PERFORMANCE

Abstract

by Matthew Edward Hood, Ph.D.
Washington State University
May 2006

Chair: John R. Nofsinger

I develop normal equivalents to assess the impact of non-normal return distributions on investor utility using the power utility function. From Taylor series approximations and simulated return data generated from both a non-central *t* distribution and normal distribution, I find that the statistical rejection of normality is not strong enough to determine that a sample distribution is economically meaningful. A sample of 373 hedge funds, of which 86% reject normality using the Jarque-Bera test, with complete monthly data from 1996 to 2005 is analyzed with the normal equivalents. As a stand alone investment, only 4% of them have return distributions that would cause an investor to be willing to trade ten basis points (per month) for the normal distribution. It does not appear that managers are taking positions that are utility detracting, which would have negative skewness and excess kurtosis, because there are a nearly equal number being utility enhancing (7) as utility detracting (9). The Sharpe ratio's use of the standard deviation to measure risk is generally robust; 90% of the funds had Sharpe ratios that were between 3.9% underestimated and 4.5% overestimated. None of the Sharpe ratios

computed with normal equivalents were statistically different than the standard Sharpe ratio. Overwhelmingly, this sample made a positive contribution to a portfolio of stocks and T-bills, even though funds with different strategies impacted the allocations of stocks and T-bills differently. Also, even though 78% of these constructed portfolios reject normality, none of them have non-normality that would substantially impact the utility of an investor. The Sharpe ratio is overestimated for these portfolios by between 0.7% and 2.6% (all positive due to the negative skewness of the portfolios). Hedge fund rankings based on the Sharpe ratio are reasonable for this group if considered as a member of a portfolio, but analysts using it to compare funds as a stand alone investment should be more cautious.

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

CHAPTER ONE: OPENING

This dissertation introduces normal equivalents as a new measure to determine the economic significance of the shape of the distribution in Chapter Two. Standardizing is done by equating the average utility of the sample to the expected utility of normally distributed returns. The power utility function is used because it is the most common choice for academics that matches utility theory. The most common alternative to the power utility function is the exponential utility function. The “Utility Functions” Section of Chapter Five describes the small differences between the two. The imposition of the normal distribution is done because it is a good approximation for how many investment’s returns are actually distributed and is a valid assumption for the Sharpe ratio. The uniform distribution is presented in the Section of Chapter Five titled “Alternative Distributions” as a possible alternative because it is familiar and the stability of the higher moments. The normal equivalents allow for a non-normality premium that determines how much the distribution matters to an investor. It also allows for an adjustment to the standard deviation to incorporate the non-normality into a better approximation of risk in the Sharpe ratio.

As a new measure, it is important to determine its ability to provide meaningful results and determine the general cases for which the non-normality matters to investors and analysts. I explore the normal equivalents with the mathematics of the expected value of Taylor series expansions around the mean and simulations of a non-central *t* distribution in Chapter Three. Also in Chapter Three are simulations from the normal distribution to find the frequency for which meaningful non-normality emerges.

This dissertation applies the normal equivalents to a sample of hedge funds and to a portfolio that includes a hedge fund with a stock index and one-month T-bill in Chapter Four.

Hedge funds are noted for their rejection of normality and wide range in their operating behaviors and return distributions. Hedge funds are also well known for the inherent difficulties in getting data from a group of managers who are not required to report their returns and may have incentives to start or stop reporting that creates more biases in the data than found in standard investments. The literature addressing this problem is summarized in the "Data Biases" Section of Chapter Five. After examining the non-normality of hedge funds as a stand alone investment, I consider the non-normality of a portfolio that includes a hedge fund. Creating a common portfolio for funds with a wide range of goals and activities would introduce too great of a restriction on an investor. Therefore, the portfolios are customized to each hedge fund for the proper allocation of stocks and T-bills so as to maximize expected utility. The "Hedge Funds in a Portfolio" Section of Chapter Five explains the procedure for constructing and optimizing the portfolio and the value that hedge funds contribute to the portfolio.

This dissertation finishes with concluding comments in Chapter Six.

CHAPTER TWO: INTRODUCING THE NORMAL EQUIVALENT

Several recent papers identify non-normality in the returns of hedge funds and propose alternative measures to account for it. This renews the debate over investors' preferences for higher moments and the effect of higher moments on the Sharpe ratio, the industry standard measure for portfolio performance. This chapter introduces normal equivalents to describe non-normality and suggests a modification to the Sharpe ratio if considerable non-normality is present. Normal equivalents are uniquely able to quantify the non-normality in the data by placing its effect on utility into the mean or standard deviation. Instead of a statistical determination to decide if a distribution rejects normality, normal equivalents decide if a distribution's non-normality is economically meaningful to investors.

Performance Measurement

A traditional and common performance yardstick is the Sharpe ratio, a measure of reward to variability proposed by Sharpe (1966). To compute the *ex post* measure, the fund's net excess return is computed for each time period. The average of the net returns is divided by the standard deviation.

$$SR = \frac{\bar{r}}{s_r}$$

This ratio measures the efficiency of the fund – a fund with a higher ratio produces more reward per unit of risk. An important assumption of the Sharpe ratio is that the allocation to the fund is determined by the investor according to his level of risk aversion. Suppose, in *ex ante* terms, Fund A has an expected net return of 5% with a standard deviation of 20% and Fund B has an expected net return of 1% with a standard deviation of 4% yielding identical Sharpe ratios of 0.25 (See Figure 2-1). Investors would judge both fund managers as equally effective at managing risk. However, these individual Sharpe ratios do not determine the proper allocation into the funds. Investors make the allocation decision based on their own risk-reward tradeoff from utility.

The asymptotic variance of the Sharpe ratio (Equation 2-1), derived by Lo (2002) for any independent and identical distribution, is substantial for a metric so heavily used to evaluate performance. The variance also increases as the Sharpe ratio increases.

$$\mathrm{V}(SR) \overset{a}{=} 1 + \frac{SR^2}{2} \tag{2-1}$$

The 95% confidence interval for the true Sharpe ratio from return data with *T* draws from an *iid* normal distribution is:

$$SR \pm \frac{1.96 \times \sqrt{1 + \frac{SR^2}{2}}}{\sqrt{T}}$$

Therefore a Sharpe ratio of 0.25 determined from 120 observations could be expected to come from an *iid* normal distribution with an *ex ante* Sharpe ratio of between 0.08 and 0.43. Figure 2-2 shows the upper and lower bounds of the 95% confidence region for Sharpe ratios ranging from 0.1 to 2.0.

Sharpe's 1966 paper used yearly observations for ten years in the creation of the ratio. The general practice today is to use monthly observations for five years then multiply by the square root of twelve to annualize the ratio:

$$\bar{r}_{yearly} \approx \bar{r}_{monthly} \times 12$$
$$s_{yearly} \approx s_{monthly} \times \sqrt{12}$$
$$SR_{monthly} = \frac{\bar{r}_{monthly}}{s_{monthly}}$$
$$SR_{yearly} = \frac{\bar{r}_{yearly}}{s_{yearly}} \approx \frac{\bar{r}_{monthly} \times 12}{s_{monthly} \times \sqrt{12}}$$

$$SR_{yearly} \approx SR_{monthly} \times \sqrt{12} \qquad (2\text{-}2)$$

The Sharpe ratio requires several simple assumptions to be appropriate for performance measurement. It requires one interest rate for borrowing and lending (the risk free rate), positive net returns, no autocorrelation, and either the funds have normally distributed returns or investors have quadratic utility functions. In practice, the borrowing and lending rates are not the same, so some risky investments might appear relatively more attractive using one rate and less attractive using the other. Thus making it difficult for investors to be able to allocate their investments properly. In addition, the Sharpe ratio is not always positive (a negative Sharpe ratio appears to improve with greater risk because the denominator is increased).

The third assumption of no autocorrelation has recently been researched and found to be overly restrictive for hedge funds. As noted by Sharpe (1994), returns may have serial correlation. After Lo (2002) derives the distribution of the Sharpe ratio under the assumption of independent and identically distributed returns he shows that with autocorrelation present, and unaccounted for, the estimation of the Sharpe ratio is biased. A simple autoregressive process, such as an AR1 – this period's return is dependent to some extent on last period's return, can be accounted for by multiplying by a factor other than the square root of twelve – recall Equation 2-2. This factor will be less than the square root of twelve for positive autocorrelation and larger for negative autocorrelation. Since hedge funds often have positive autocorrelation, Lo found that their Sharpe ratios are often overstated and that rankings “can change dramatically.” This is an indication that the autocorrelations also vary dramatically as well.

Lastly, the assumption that funds either have normally distributed returns or that investors have quadratic utility functions has prompted even more research. Fung and Hsieh (1999b) state that hedge funds’ returns are not normally distributed. Kouwenberg (2003) finds that 735 of 1,399 hedge funds, over a six-year period, reject normality at the 5% level and most of these funds have negative skewness and excess kurtosis. Amin and Kat (2003) reject normality for 66 of 77 hedge funds with ten-years of monthly data, noting their skewness in particular. Other than hedge funds, Pdersen and Satchell (2000) find that more than half of the emerging market funds they analyzed reject normality because of their skewness and kurtosis. Liang (2004) finds substantial autocorrelation in portfolios of commodity trading advisors. Clearly, many types of portfolios do not have normally distributed returns.

With all of these findings of non-normality, this fourth assumption that either returns are normally distributed or that investors do not care about higher moments often hangs on the utility function of hedge fund investors. Scott and Horvath (1980) show that investors like odd moments, such as skewness, and dislike even moments, such as kurtosis (both negative skewness and excess kurtosis are utility detracting features for a return distribution). Even with these preferences, several papers use Taylor series approximations to demonstrate that the loss due to the higher moments is trivial. Prominent among these are Levy and Markowitz (1979) and Hlawitschka (1994). Levy and Markowitz (1979) show that the second order Taylor series approximation had a very strong correlation (often greater than 99%) with the actual utility for a range of exponential and power utility functions using mutual funds, stocks, and small random portfolios of stocks. Hlawitschka (1994) substantiates the findings that the second order Taylor series approximations are excellent for stocks. It also concludes that the second order Taylor series approximations are adequate for portfolios of call options and preferable to higher order approximations. Fung and Hsieh (1999b) extended this work towards hedge funds, concluding that the second order Taylor series approximations of utility "produce rankings which are nearly correct" for a broad range of risk aversion levels in the power and exponential utility functions. However, they also conclude that the Sharpe ratio works well only for high levels of risk aversion.

Therefore, the Sharpe ratio is based on assumptions that are not true for many investments. Because second order Taylor series approximations are excellent approximations for utility, the Sharpe ratio is probably robust to non-normality. But the Sharpe ratio is a measure of efficiency, reward per unit of risk, not utility. It is not meant to determine which fund, 100% invested, would yield the highest utility. Again, consider the Sharpe ratio array in Figure 2-1. Suppose three funds are available, Funds A and B both have a Sharpe ratio of 0.25 and Fund C has a Sharpe ratio of 0.20. Funds A and B are more efficient than Fund C, yet Fund C garners more utility because its choice of risk and reward is more preferable. The efficiency of A and B are not directly transferable to utility, so neither is the Sharpe ratio.

Most of the recent papers exposing flaws in using the Sharpe ratio for evaluating non-normal portfolios propose alternative measures. Amin and Kat (2003) propose an efficiency test, which determines the cost of a dynamic trading strategy (with a normal distribution and cash) that would generate the same payoff distribution. They continue on by determining the optimal allocation with the rest invested in the S&P 500, allowing for a broadly framed measure. Kouwenberg (2003) uses a weighted least squares regression to compute the intercept (Jensen's alpha). Using weighted least squares instead of ordinary least squares to compute the intercept allows some observations to be weighted more heavily than others. He uses the marginal utility of the gross return which gives the worst returns the largest weight. (Utility increases with return but at a decreasing rate, so marginal utility approaches zero.) This measure puts Jensen's alpha more in line with utility but is narrowly framed and comparisons must be made for stand alone investments only. Sharma (2004) chooses the certainty equivalent for evaluating hedge funds. A comparison of certainty equivalents is identical to comparisons of utility. This measure is also narrowly framed and is used only for comparing funds on a stand alone basis.

The Sharpe ratio that these papers seek to replace with their own measure is aimed at efficiency. High Sharpe ratio funds are more efficient and hence more valuable additions to a portfolio. These should not be mistaken as the funds that provide investors the most utility. The normal equivalents proposed here are used to understand and quantify the shape of the distribution and determine if the risk measure of the Sharpe ratio is actually measuring the risk to utility. If not, then replace it with one that does.

Normal Equivalents

The certainty equivalent, *CE*, for any risky investment is the return, if known for certain, that would leave the individual indifferent between the risky investment and the certain return. There is a direct relationship, both monotonic and positive, between the certainty equivalent and the utility. The *CE* is equal to the expected return for risk neutral investors and bounded from above by the expected return for risk averse investors. The risk premium, *RP*, is the difference between the expected return and the *CE* and shows the individual's preferences about the risky investment. The larger the *RP* the more the investor is willing to sacrifice, in terms of expected return, to avoid the risk.

The certainty equivalent and risk premium are dependent upon the return distribution of the investment, the starting conditions, and the individual's preferences. Work using utility to evaluate performance generally uses an assumption for starting conditions. If the individual's attitude towards risk is dependent upon starting wealth, such starting wealth must be assumed, therefore the choice of a constant relative risk aversion utility function, CRRA, is prevalent. Such a function makes the starting wealth unimportant as only returns relative to the starting wealth concern the individual. Mean-variance utility functions and power utility functions are the most prominent of the CRRA family of utility functions. Power utility functions are more realistic and more popular among academics because they are more consistent with utility theory. (For a discussion of constant absolute risk aversion, CARA, and the exponential utility function, see the "Utility Functions" Section in Chapter Five.) The power utility function is defined as:

$$U = \frac{R^{1-\gamma}}{1-\gamma}$$

The power utility function has the desirable property of diminishing marginal utility. For example, an investor is indifferent between a gamble that has equal chances at a 10% loss and an 11% gain and a certain return of 0% if the risk aversion is one. A 10% loss has more impact on utility than a 10% gain does. As risk aversion increases, the amount the investor needs to balance the chance of losing 10% increases as well. At a risk aversion of five, a gain of 20% is needed to offset a loss of 10%.

With the current debate over the effect of the distribution of returns for atypical investments, like hedge funds which usually feature negative skewness and excess kurtosis, a measure of the effect of the distribution on utility is important. Three normal equivalent measures are introduced here to compare the realized return distribution for any risky asset to the normal distribution. The normal distribution is used for comparison because it is a common assumption for returns and is relatively accurate for many assets. The normal distribution is a common assumption for performance measures, like the Sharpe ratio and Jensen's alpha, because of its familiar and desirable properties and reasonable approximations. However, other distributions could be used following the same procedures. The "Alternative Distributions" Section in Chapter Five shows how the uniform distribution would be used and also why other popular distributions, such as the log-normal are unworkable because of their higher moments. The normal equivalents use utility to find an equivalent normal distribution, they do not assume returns are normally distributed.

Investments into assets for which these measures would be computed cannot lose more than their initial investment. This would lead to censoring observations below zero, at which is a total loss of the investment. Censoring in this fashion replaces any observation below zero with zero. (The alternative to censoring is truncating, where any observation below zero is replaced with another draw from the distribution. Censoring the return distribution of a fund is equivalent to giving an investor a total loss when the fund goes bankrupt and is consistent with limited liability. Truncating the distribution is equivalent to giving the investor a 'do over' if the fund goes bankrupt. Censoring is more consistent with the experience of investors than truncating.) For our purposes, the censoring is not done at a total loss, but at a loss of 99%. This makes the functions better behaved because the power utility function is undefined at zero. Another popular practice is to place only ninety-nine percent of the portfolio to go to the risky asset, leaving one percent in a risk-free investment to effectively censor the returns. The probability distribution function of a normal distribution with a reasonable mean and variance approaches zero faster than the utility function approaches negative infinity so the mathematical difference is negligible for the improved analytic approximations.

In the formulas used to estimate the normal equivalents, R is a normally distributed random variable censored from below at 1%. The parameters μ and σ^2 represent the mean and variance from the underlying uncensored normal distribution. The mean and variance of the sample returns from the data are represented as $\overline{R}$ and s^2, and the average utility as $\overline{U}$.

$$\overline{U} = \frac{\sum_{t=1}^{T} U(R_t)}{T}$$

To determine the gain or loss in the utility due to the non-normality of the distribution, the expected utility from the censored normal returns with the mean and variance equal to that of the sample is calculated – the normal equivalent-utility (NE_U). The name is derived from its definition – it is the expected utility that would be derived from normally distributed returns with the same mean and variance as the sample. It can be compared directly to the average utility from the data. If the NE_U is greater than the average utility of the sample, $\overline{U}$, then the sample return distribution is worse than the normal distribution and the sample's higher moments are utility detracting. The NE_U is bounded from above by the utility of the average return, since the normal distribution involves risk it must provide less utility than a distribution without risk. The derivation of the NE_U for the power utility function is:

$$NE_U = \mathrm{U}\left(R \mid R \sim \mathrm{N}\left(\overline{R}, s^2\right)\right)$$

$$NE_U = \int_{-\infty}^{0.01} \frac{0.01^{1-\gamma}}{1-\gamma} \times \frac{e^{-\frac{[R-\overline{R}]^2}{2s^2}}}{\sqrt{2\pi s^2}} \times dR + \int_{0.01}^{+\infty} \frac{R^{1-\gamma}}{1-\gamma} \times \frac{e^{-\frac{[R-\overline{R}]^2}{2s^2}}}{\sqrt{2\pi s^2}} \times dR$$

$$NE_U < \mathrm{U}\left(\overline{R}\right)$$

To determine the relative importance of the distribution compared to the importance of the variability in returns, the normal equivalent-mean (NE_μ) is defined as the expected return such that a normally distributed return with the same standard deviation as the sample would have an expected utility equal to the average from the sample. The NE_μ holds the utility and the standard deviation of the sample constant allowing for the mean to compensate for the move from the data's actual distribution to the normal distribution. Although it cannot be computed analytically, the integrals can be approximated through iteration. The equation necessary to solve for the NE_μ using the power utility function is:

$$\overline{U} = \mathrm{U}\left(R \mid R \sim \mathrm{N}\left(NE_\mu, s^2\right)\right)$$

$$\overline{U} = \int_{-\infty}^{0.01} \frac{0.01^{1-\gamma}}{1-\gamma} \times \frac{e^{-\frac{[R-NE_\mu]^2}{2s^2}}}{\sqrt{2\pi s^2}} \times dR + \int_{0.01}^{+\infty} \frac{R^{1-\gamma}}{1-\gamma} \times \frac{e^{-\frac{[R-NE_\mu]^2}{2s^2}}}{\sqrt{2\pi s^2}} \times dR$$

The NE_μ can be compared to the sample mean and to the certainty equivalent. If the NE_μ is greater than the sample mean then the sample distribution is preferred to the normal distribution and is utility enhancing. The NE_μ is the expected return that would leave an investor indifferent between a normal distribution with a mean return of NE_μ and the sample's distribution with a mean return of $\overline{R}$ (both distributions having the same variance). Therefore, if the NE_μ is greater than the $\overline{R}$, the investor would prefer the sample distribution to the normal distribution if both had the same mean and variance.

The NE_μ, like the sample mean, will never be less than the certainty equivalent for risk averse investors. The non-normality premium, *NNP*, is defined as the difference between the mean of the sample and the NE_μ. While the risk premium must be positive because risk averse investors find risk to be utility detracting, the non-normality premium may be positive or negative because investors may find the distribution of the sample to be better or worse than the normal distribution. A positive *NNP* will be the result of a distribution that is utility *detracting* and will be no larger than the *RP* because the premium for a normal distribution must be less than the premium for a risk free investment.

$$\overline{R} = NE_{\mu} + NNP$$

The normal equivalent-standard deviation, NE_{σ}, is the standard deviation from normally distributed returns with the same mean as the sample that will yield the average utility of the sample. It is useful for comparing the risk to utility instead of the risk to return. The standard measure of risk in finance is standard deviation. For the standard deviation to be a valid measure of risk across assets, either higher moments are unimportant to investors (which is not likely) or the higher moments are all the same (which is also not likely). This normal equivalent puts all of the risk into a single measure and can be used for comparing asset returns with different distributions because they are standardized into the normal distribution. If the NE_{σ} is greater than the sample standard deviation then investors must be compensated (by less variance) for assuming the unattractive return distribution of the sample. The sample's distribution is utility detracting. A negative NE_{σ} is nonsensical and cannot occur because the normal distribution is not preferable to a guaranteed return. The derivation of the NE_{σ} for the power utility function is:

$$\overline{U} = \mathrm{U}\left(R \mid R \sim \mathrm{N}\left(\overline{R}, NE_{\sigma}^{\ 2}\right)\right)$$

$$\overline{U} = \int_{-\infty}^{0.01} \frac{0.01^{1-\gamma}}{1-\gamma} \times \frac{e^{-\frac{\left[R-\overline{R}\right]^2}{2NE_{\sigma}^{\ 2}}}}{\sqrt{2\pi NE_{\sigma}^{2}}} \times dR + \int_{0.01}^{+\infty} \frac{R^{1-\gamma}}{1-\gamma} \times \frac{e^{-\frac{\left[R-\overline{R}\right]^2}{2NE_{\sigma}^{\ 2}}}}{\sqrt{2\pi NE_{\sigma}^{2}}} \times dR$$

Many of the standard portfolio measures, such as the Sharpe ratio, are derived using standard deviation as the measure of risk. If the NE_{σ} is a more appropriate measure for risk, then the same formula could be used by replacing the sample standard deviation with the NE_{σ}. The

Sharpe ratio that is modified by normal equivalents, SR_{NE}, can also be defined using the ratio of the sample standard deviation to the NE_σ.

$$SR_{NE} = SR \times \frac{s}{NE_\sigma} \tag{2-3}$$

Lastly, the utility adjustment to the standard deviation, *UAS*, is defined as one less than the ratio of the NE_σ and the standard deviation. When the normal equivalent-standard deviation is larger than the sample standard deviation (the distribution is utility detracting) the *UAS* will be positive and when the normal equivalent-standard deviation is smaller than the sample standard deviation the *UAS* will be negative. Or further, the *UAS* can be used to derive the modified Sharpe ratio in terms of the *UAS* by combining it with Equation 2-3 to produce Equation 2-4.

$$UAS = \frac{NE_\sigma}{s} - 1$$

$$SR_{NE} = \frac{SR}{1 + UAS} \tag{2-4}$$

A substantial utility adjustment to the standard deviation might be necessary in order to have a statistically significant impact on the Sharpe ratio. The *UAS* that would cause a significant difference decreases as the Sharpe ratio increases, shown in Equation 2-5 which is derived from Equations 2-1 and 2-4. The necessary *UAS* for a sample of 120 observations for a statistically significant difference at the five percent level in the Sharpe ratio is shown graphically in Figure 2-3.

$$SR_\alpha = SR + N(\alpha) \times \sqrt{1 + \frac{SR^2}{2}} \div \sqrt{T}$$

$$SR_\alpha = \frac{SR}{1 + UAS_\alpha}$$

$$\frac{SR}{1+UAS_{\alpha}} = SR + N(\alpha) \times \sqrt{1+\frac{SR^2}{2}} \div \sqrt{T}$$

$$UAS_{\alpha} = \frac{-N(\alpha) \times \sqrt{1+\frac{SR^2}{2}}}{SR\sqrt{T} + N(\alpha) \times \sqrt{1+\frac{SR^2}{2}}} \qquad (2\text{-}5)$$

The normal equivalents can place the effects on utility of the non-normality of the return distribution into either the mean or standard deviation. Placing it into the mean would be especially appropriate for an investor to determine if the skewness and kurtosis would alter their opinion of the investment. Placing it into the standard deviation would be especially appropriate for performance evaluation to determine if the risk is being properly identified. The normal equivalents are not designed to determine portfolio allocation nor performance evaluation, but to describe the impact of the distribution on investor utility or the risk measurement of the Sharpe ratio.

The primary means to determine the impact of the higher moments on investors in the following chapters will be the non-normality premium. A non-normality premium of ten basis points, calculated on a monthly basis, is defined as a distribution with higher moments that are so utility detracting that an investor should be aware of its extra risks. A non normality premium of minus ten basis points is defined as a distribution with utility enhancing higher moments that an investor would value. The primary means to determine if the standard deviation is an appropriate measure of risk for the Sharpe ratio will be the utility adjustment to the standard deviation. If it shows a positive or negative bias of more than ten percent, the standard deviation is not accurately portraying the risks. Further, if the range of the utility adjusted standard deviations is large, this is an indication that the Sharpe ratio is not appropriate for ranking funds. Although a ten percent difference between the standard deviation and the normal equivalent-standard deviation is used in this dissertation; this difference in the Sharpe ratio is not statistically significant as shown in Figure 2-3. However, it is plausible that analysts' opinions of performance might change with as small a difference as ten percent and quite likely that the difference required for statistical significance is not commonly known or used.

CHAPTER THREE: EXPLORING THE NORMAL EQUIVALENT

A normal equivalent imposes the utility gained or lost from the non-normality of the return distribution into either the mean or standard deviation. This allows a researcher to describe the value of the distribution in useful terms. The normal equivalent-mean allows for the computation of a non-normality premium, the difference between the sample mean and the mean that would yield the same utility if the distribution were normal. A positive non-normality premium is the result of a sample distribution that is utility detracting – worse for the investor than the normal distribution. Large non-normality premiums (in absolute terms) indicate a return distribution that is valued substantively different than the normal distribution.

The non-normality premium describes the effect of the sample distribution in the simplest terms. The normal equivalent-standard deviation is the standard deviation that normally distributed returns would require in order to attain the same utility as the sample. The utility adjustment to the standard deviation is one less than the ratio of the normal equivalent-standard deviation to the sample standard deviation (see Equation 2-4). If this adjustment is positive then the sample distribution that is utility detracting – normally distributed returns with the same mean as the sample could achieve the same utility as the sample with a larger standard deviation. If this adjustment is negative, then the sample distribution is utility enhancing – it would take a smaller standard deviation, accompanied with the sample mean, for a normal distribution to achieve the same utility. Utility detracting higher moments, negative skewness and excess kurtosis, represent an increased risk; specifically, a greater chance at a large loss relative to the normal distribution. If the goal is to incorporate that risk into the Sharpe ratio, using the utility adjustment to the standard deviation is the most appropriate way to accomplish this task. Normal equivalents allow for a description of the sample's return distribution and determination of the robustness of the Sharpe ratio.

Taylor Series Expansion

Taylor series expansions have led several researchers to conclude that higher moments, such as skewness and kurtosis, are not important to the utility of investors of stocks, mutual funds, and hedge funds. Levy and Markowitz (1979) find that second order Taylor series approximations of utility had a very strong correlation with the actual utility for mutual funds, stocks, and small portfolios of stocks. Hlawitschka (1994) substantiates these findings for stocks and portfolios of call options and then further argues that higher order approximations were less correlated with the actual utility. Fung and Hsieh (1999b) extend these finding for hedge funds. This chapter uses fourth order Taylor series expansions around the mean to find levels of skewness and kurtosis that are important to investors and analysts using the Sharpe ratio. This is not done to invalidate their work, but to discover when such higher moments are helpful.

The Taylor series approximation is a useful tool for approximating a difficult function around a specific point with a simpler formula. The marginal impact of skewness and kurtosis compared to the mean and standard deviation is determined using a fourth order Taylor series approximation. Further, it is used to find the combination of risk aversion, standard deviation, skewness, and kurtosis necessary for a measurable impact on utility or the Sharpe ratio to emerge. To do this, the expected utility must be calculated for the fourth order Taylor series approximation of the utility function around the mean. The derivation below is done with a generic power utility function. The risk aversion parameter, γ, is positive and larger parameters indicate more risk averse investors. If the risk aversion parameter is set to one, the function reduces to the familiar log utility function.

$$\mathrm{U}(R) = \frac{R^{1-\gamma}}{1-\gamma}$$

$$\mathrm{U}(R) \approx \frac{\mu^{1-\gamma}}{1-\gamma} + \frac{1}{\mu^{\gamma}} \times [R-\mu] - \frac{\gamma}{2\mu^{\gamma+1}} \times [R-\mu]^2 + \frac{\gamma[\gamma+1]}{6\mu^{\gamma+2}} \times [R-\mu]^3 - \frac{\gamma[\gamma+1][\gamma+2]}{24\mu^{\gamma+3}} \times [R-\mu]^4$$

After taking the expected value of the Taylor series approximation and substituting the mean, variance, skewness, and kurtosis into the equation we can take a total derivative. The true skewness and kurtosis are identified by η_1 and η_2 and sample estimates by h_1 and h_2. (Skewness and kurtosis can also be represented as γ_1 and γ_2, but η_1 and η_2 were chosen to avoid confusion with the most common symbol of risk aversion, γ. The sample skewness and kurtosis statistics are represented as h_1 and h_2 for consistency.)

$$\mathrm{E}(\mathrm{U}(R)) \approx \frac{\mu^{1-\gamma}}{1-\gamma} - \frac{\gamma\sigma^2}{2\mu^{\gamma+1}} + \frac{\gamma[\gamma+1]\sigma^3\eta_1}{6\mu^{\gamma+2}} - \frac{\gamma[\gamma+1][\gamma+2]\sigma^4\eta_2}{24\mu^{\gamma+3}}$$

$$\begin{aligned} d\,\mathrm{E}(\mathrm{U}(R)) \approx {} & \frac{1}{\mu}d\mu + \frac{\gamma[\gamma+1]\sigma^2}{2\mu^{\gamma+2}}d\mu - \frac{\gamma\sigma}{\mu^{\gamma+1}}d\sigma - \frac{\gamma[\gamma+1][\gamma+2]\sigma^3\eta_1}{6\mu^{\gamma+3}}d\mu \\ & + \frac{\gamma[\gamma+1]\sigma^2\eta_1}{2\mu^{\gamma+2}}d\sigma + \frac{\gamma[\gamma+1]\sigma^3}{6\mu^{\gamma+2}}d\eta_1 \\ & + \frac{\gamma[\gamma+1][\gamma+2][\gamma+3]\sigma^4\eta_2}{24\mu^{\gamma+4}}d\mu - \frac{\gamma[\gamma+1][\gamma+2]\sigma^3\eta_2}{6\mu^{\gamma+3}}d\sigma \\ & - \frac{\gamma[\gamma+1][\gamma+2]\sigma^4}{24\mu^{\gamma+3}}d\eta_2 \end{aligned}$$

Using the total derivatives, the marginal impact of skewness and kurtosis can be compared to that of the mean and standard deviation.

$$-\frac{d\mu}{d\eta_1} \approx \frac{4\mu^2\sigma^3}{\frac{24\mu^4}{\gamma[\gamma+1]} + 12\mu^2\sigma^2 - 4[\gamma+2]\mu\sigma^3\eta_1 + [\gamma+2][\gamma+3]\sigma^4\eta_2}$$

$$-\frac{d\sigma}{d\eta_1} \approx \frac{-[\gamma+1]\mu\sigma^2}{6\mu^2 - 3[\gamma+1]\mu\sigma\eta_1 + [\gamma+1][\gamma+2]\sigma^2\eta_2}$$

$$-\frac{d\mu}{d\eta_2} \approx \frac{-[\gamma+2]\mu\sigma^4}{\frac{24\mu^4}{\gamma[\gamma+1]} + 12\mu^2\sigma^2 - 4[\gamma+2]\mu\sigma^3\eta_1 + [\gamma+2][\gamma+3]\sigma^4\eta_2}$$

$$-\frac{d\sigma}{d\eta_2} \approx \frac{[\gamma+1][\gamma+2]\sigma^3}{24\mu^2 - 12[\gamma+1]\mu\sigma\eta_1 + 4[\gamma+1][\gamma+2]\sigma^2\eta_2}$$

To understand the marginal impact of skewness and kurtosis it is helpful to replace the symbols with potential values. The monthly mean will be assumed to be 1%, the monthly standard deviation 5%, and the skewness and kurtosis set to that of the normal distribution (skewness = 0 and kurtosis = 3). Work in utility theory and investments reveal that the risk aversion coefficient for a power utility function is likely small. Mehra and Prescott (1985) believe the plausible extreme for the coefficient is ten. Sharma (2004) cites evidence from Ait-Sahalia, Parker, and Yogo (2001) who estimate the risk aversion coefficient at 3.2 and Osband (2002) who suggests two to four. This dissertation uses a range of one to five with three as a starting point.

The marginal impact of skewness and kurtosis relative to the marginal impact of the mean and standard deviation are generally small and only sensitive to changes in the risk aversion coefficient and standard deviation (Table 3-1). The table assumes a net mean monthly return of 1%, skewness of 0, and kurtosis of 3. With a risk aversion coefficient of 3 and standard deviation of 5%, an increase of 1 in the skewness would have the same impact on utility as a increase in the mean of 2.41 basis points (per month). Relating this marginal change instead to the standard deviation, its impact would be the same as decreasing the standard deviation by 16.11 basis points. This decrease in the standard deviation would increase the Sharpe ratio by

3.3% (before annualizing). A sizeable change in the skewness has a small but measurable effect on the utility of an investor relative to the mean and standard deviation. The marginal impact of kurtosis is smaller but the kurtosis can be expected to vary over a larger range. An increase in the kurtosis of 1 would have the same impact on utility as a 0.15 basis point reduction in the mean or 1.00 basis point increase in the standard deviation. This increase in the standard deviation would decrease the Sharpe ratio by 0.2%. The kurtosis would have to change by about sixteen to make the same impact as a change of one in the skewness.

The starting conditions imposed on the mean, skewness, and kurtosis has a small impact on the marginals. Altering the starting conditions of the mean, skewness, and kurtosis had very little effect. Using 0% for the mean net return increased the marginal impacts of the skewness and kurtosis relative to the mean and standard deviation, but by less than two-tenths of a basis point. Decreasing the skewness to -2 left the marginal impacts relative to the mean unchanged but reduced the size of the marginal impacts relative to the standard deviation to -13.50 for the skewness and 0.84 for the kurtosis. Increasing the kurtosis to 25 had almost the same impact as decreasing the skewness to -2. The marginal impacts relative to the mean were unchanged and the size of the marginal impacts relative to the standard deviation were reduced to -13.70 for the skewness and 0.85 for the kurtosis.

Altering the starting conditions of the standard deviation and risk aversion parameter are important. Increasing the monthly standard deviation to 9% increases the size of the marginal utility of skewness relative to the marginal utility of the mean to 13.61 basis points per month and to the marginal utility of the standard deviation to -49.53 basis points. At these levels, an investor's concern for skewness is well founded and a skewness of just -1 would need to be offset by an increase of about 1.8% per year in the mean or decrease of about 1.7% per year in the standard deviation. The marginal utility of kurtosis rises faster but starts at a lower level. An increase of 1 in the kurtosis would be equivalent to a 1.52 basis point decrease in the monthly mean or 5.52 basis point increase in the monthly standard deviation. These numbers are still small, but a large difference in the kurtosis of different funds is possible and would be important.

Cutting the standard deviation to 3% similarly reduces the impact of the higher moments. An increase of 1 in the skewness is then equivalent to a 0.53 basis point improvement in the mean or 5.89 basis point improvement in the standard deviation. A decrease of 1 in the kurtosis is reduced to a 0.02 basis point improvement in the mean or a 0.22 basis point improvement in the standard deviation.

The standard deviation measures the spread or size of the distribution. If the standard deviation is relatively small, then all risk premiums will be small as there is little variation to be concerned about. If the standard deviation is large, then things that affect the risk (skewness and kurtosis) are more important. For investments that have standard deviations less than 5% the higher moments are not very important. But if a fund has a standard deviation of 9% or more, the higher moments are essential to the investor and a proper risk measure for the Sharpe ratio.

The risk aversion coefficient has a similar, but smaller, effect on the marginal change in the mean and standard deviation necessary to achieve a marginal change to the skewness or kurtosis. Decreasing the risk aversion coefficient to one (with a 5% standard deviation) causes a 0.41 basis point improvement in the mean or 8.19 basis point improvement in the standard deviation to be equivalent to an increase in the skewness of 1. It also causes a 0.02 basis points reduction in the mean or 0.30 basis point increase in the standard deviation to be equivalent to a increase in the kurtosis of 1. Investors that are less risk averse than is perhaps typical have less need to concern themselves with the higher moments of their investments' return distribution.

Increasing the risk aversion coefficient to five causes a 5.90 basis point improvement in the mean or 23.54 basis point improvement in the standard deviation to be necessary to be equivalent to an increase in the skewness of 1. It also causes a 0.51 basis point reduction in the mean or 2.04 basis point increase in the standard deviation to be necessary to be equivalent to an increase in the kurtosis of 1. Investors that are more risk averse than is perhaps typical have more need to concern themselves with the higher moments of their investments' return distribution.

Two caveats emerge from this analysis. First, since skewness and excess kurtosis often coexist and can have sizeable ranges, the effect of the higher moments could correctly be perceived to be higher than what has been shown. Second, a combination of a larger than average standard deviation with an investor who is relatively more risk averse should cause more concern than either one alone.

The expected value of the fourth order Taylor series approximation around the mean of the power utility function can reveal the moments and risk aversion coefficient necessary to achieve any desired non-normality premium. For example, Harvey and Siddique (2000) argue that investors may prefer portfolios with high positive skewness. For the purposes of this study, a non-normality premium of more than 10 basis points or less than -10 basis points per month is considered a return distribution that is notably different to investors than the normal distribution. Equations 3-1 and 3-2 below must yield the same utility for a given non-normality premium, *NNP*, for the appropriate skewness and kurtosis. Equation 3-3 then solves for the skewness, given the other five parameters. Allowing the kurtosis to vary, indifference curves can be graphed in two dimensions. Each point on a line has the same non-normality premium with the same mean and standard deviation and, therefore, the same utility. The investor would be indifferent between any two points on the same line.

$$\mathrm{E}(\mathrm{U}(R)) \approx \frac{\mu^{1-\gamma}}{1-\gamma} - \frac{\gamma\sigma^2}{2\mu^{\gamma+1}} + \frac{\gamma[\gamma+1]\sigma^3\eta_1}{6\mu^{\gamma+2}} - \frac{\gamma[\gamma+1][\gamma+2]\sigma^4\eta_2}{24\mu^{\gamma+3}} \tag{3-1}$$

$$\mathrm{E}(\mathrm{U}(R)) \approx \frac{[\mu - NNP]^{1-\gamma}}{1-\gamma} - \frac{\gamma\sigma^2}{2[\mu - NNP]^{\gamma+1}} - \frac{\gamma[\gamma+1][\gamma+2]\sigma^4}{8[\mu - NNP]^{\gamma+3}} \tag{3-2}$$

$$\begin{aligned}\eta_1 = {} & \frac{[\gamma+2]\sigma}{4\mu}\eta_2 + \frac{3\mu}{[\gamma+1]\sigma} + \frac{6\mu^{\gamma+2}[\mu - NNP]^{1-\gamma}}{[1-\gamma]\gamma[\gamma+1]\sigma^3} - \frac{6\mu^3}{[1-\gamma]\gamma[\gamma+1]\sigma^3} \\ & - \frac{3\mu^{\gamma+2}}{[\gamma+1][\mu - NNP]^{\gamma+1}\sigma} - \frac{3[\gamma+2]\mu^{\gamma+2}\sigma}{4[\mu - NNP]^{\gamma+3}}\end{aligned} \tag{3-3}$$

The abnormality necessary for a distribution to be 10 basis points worse than the normal distribution when the risk aversion coefficient is three, the net mean return is 1%, and the standard deviation is 5% is extreme. The lines of Figure 3-1 represent indifference curves—each point on the line would yield the investor the same utility. In Panel A, if there is no excess kurtosis the skewness must be less than -4.15. Even when the kurtosis is 60 the skewness must be less than -0.62. A distribution that is 10 basis points better than the normal requires skewness to be more than 4.14 when there is no excess kurtosis. Skewness of 7.66 is needed when the kurtosis reaches 60.

The non-normality premium is very sensitive to the standard deviation as was previously demonstrated with the marginal utilities. Panel B shows the skewness and kurtosis necessary to achieve a non-normality premium of 10 with the same mean and risk aversion coefficient as before, 1% and 3 respectively, yet allows for different standard deviations. The dotted line is the same as in the previous graph, representing a standard deviation of 5%. Increasing the standard deviation to 7% and 9% decreases the need for as much negative skewness and excess kurtosis. At 9%, the skewness need only reach -0.74 with no excess kurtosis to achieve a non-normality premium of at least ten basis points and if the kurtosis is more than 10, a non-normality premium of at least ten basis points can occur even with positive skewness. Decreasing the standard deviation to 3% increases the need for negative skewness and excess kurtosis far beyond the graph.

The non-normality premium was also shown to be very sensitive to the risk aversion coefficient. Panel C shows the skewness and kurtosis necessary to achieve a non-normality premium of 10 with the same mean and standard deviation as before, 1% and 5% respectively, yet allows for different risk aversion coefficients. The dotted line is the same as in the previous graph, representing a risk aversion level of three. Increasing the risk aversion to four and five decreases the amount of negative skewness and excess kurtosis necessary. At five, the skewness need only reach -1.70 with no excess kurtosis to achieve a non-normality premium of at least ten basis points and if the kurtosis is more than 24, a non-normality premium of at least ten basis points can occur even with positive skewness. Decreasing the risk aversion parameter to one increases the need for negative skewness and excess kurtosis beyond the graph. In fact, the skewness needs to be less than -20 in order to achieve a non-normality premium of 10. This should be expected, as risk aversion increases then the premiums to avoid risk will increase.

The expected value of the Taylor series approximation around the mean of a power utility function shows that for most portfolios, where the standard deviation of returns is small, significant non-normality is necessary to have a substantial impact on investor's utility. Rejecting normality is not enough to conclude that the higher moments have a noticeable effect on utility unless the standard deviation and/or the risk aversion are high. In fact, if the standard deviation and risk aversion are both slightly higher than average, a fund need not reject normality in order for the abnormality to have a noticeable impact on the utility. A distribution with 1%

mean net monthly returns, 9% standard deviation, skewness of -0.06, and kurtosis of 3 will result in a non-normality premium of ten basis points if the risk aversion coefficient is set to 5. Deciding to reject normality is not equivalent to determining that the higher moments have a meaningful impact on the utility to investors.

This experiment can be repeated to examine the robustness of the Sharpe ratio to non-normality with the utility adjustment to the standard deviation, *UAS*. The Sharpe ratio places the standard deviation in the denominator as a measure of risk ignoring the higher moments. If the higher moments show risk, by increased probabilities of very bad outcomes (negative skewness and excess kurtosis), then the Sharpe ratio is misleading and biased upwards. For the purposes of this study, a *UAS* of more than 10% or less than -10% per month is considered a return distribution that indicates risk may not be adequately explained by the standard deviation. Equation 3-4 must be equal to Equation 3-1 for a given *UAS*, with the other parameters known: mean, standard deviation, skewness, kurtosis, and risk aversion. Equation 3-5 finds the skewness necessary to balance any kurtosis.

$$\mathrm{E}(\mathrm{U}(R)) \approx \frac{\mu^{1-\gamma}}{1-\gamma} - \frac{\gamma\sigma^2[1+UAS]^2}{2\mu^{\gamma+1}} - \frac{\gamma[\gamma+1][\gamma+2]\sigma^4[1+UAS]^4}{8\mu^{\gamma+3}} \tag{3-4}$$

$$\eta_1 = \frac{[\gamma+2]\sigma}{4\mu}\eta_2 + \frac{3\mu}{[\gamma+1]\sigma} - \frac{3\mu[1+UAS]^2}{[\gamma+1]\sigma} - \frac{3[\gamma+2]\sigma[1+UAS]^4}{4\mu} \tag{3-5}$$

The abnormality necessary for a distribution to have 10% more risk than the normal distribution when the risk aversion coefficient is three, the net mean return is 1%, and the standard deviation is 5% is still large. The lines of Figure 3-2 are indifference curves – each point represents the same amount of risk to utility, and because they have the same mean, they produce the same utility. In Panel A, if there is no excess kurtosis the skewness must be less than -3.27 but when the kurtosis is 60 the skewness may even be 0.26. A distribution that has

10% less risk than the normal distribution requires substantial skewness and kurtosis – the skewness must be more than 2.94 when there is no excess kurtosis and 6.47 when the kurtosis reaches 60.

The utility adjustment to the standard deviation is very sensitive to the standard deviation as was previously demonstrated with the marginal utilities. Panel B of Figure 3-2 shows the skewness and kurtosis necessary to achieve a *UAS* of 10% with the same mean and risk aversion coefficient as before, 1% and three respectively, yet allows for different standard deviations. The dotted line is the same as in the previous graph, representing a standard deviation of 5%. Increasing the standard deviation to 7% and 9% decreases the need for negative skewness and excess kurtosis. At 9%, the skewness need only reach -1.92 with no excess kurtosis to achieve a *UAS* of at least 10% and if the kurtosis is more than 20, a *UAS* of at least 10% can occur even with positive skewness. Decreasing the standard deviation to 3% increases the need for negative skewness and excess kurtosis. In fact, the skewness needs to be less than -3.24 in order to achieve a non-normality premium of 10 even with a kurtosis of 60. The changes to the *UAS* are smaller compared to that of the *NNP* because the basis points to adjust the standard deviation must increase.

The utility adjustment to the standard deviation was also shown to be very sensitive to the risk aversion coefficient. Panel C shows the skewness and kurtosis necessary to achieve a *UAS* of 10% with the same mean and standard deviation as before, 1% and 5% respectively, yet allows for different risk aversion coefficients. The dotted line is the same as in the previous graph, representing a risk aversion level of three. Increasing the risk aversion to four and five decreases the amount of negative skewness and excess kurtosis necessary. At five, the skewness need only reach -2.24 with no excess kurtosis to achieve a *UAS* of 10% and if the kurtosis is more than 30, a *UAS* of at least 10% can occur even with positive skewness. Decreasing the risk aversion parameter to one increases the need for negative skewness and excess kurtosis; the skewness needs to be less than -6.41 in order to achieve a *UAS* of 10% when no excess kurtosis is present.

The more risk averse the investor is assumed to be, the less robust the assumption of normality is for the measure of risk. The expected value of the Taylor series approximation around the mean of a power utility function shows that for funds where the standard deviation of returns is small, or investors are not assumed to be strongly risk averse, the assumption of normality is relatively robust. Rejecting normality is not enough to conclude that the higher moments have a noticeable effect on risk unless the standard deviation and/or the risk aversion are high. However, if the higher moments do reject normality and the standard deviation is sizeable and/or the investor is assumed to be strongly risk averse, the Sharpe ratio's use of the standard deviation to measure risk is not well founded and should be replaced.

The largest caveat for the results from Taylor series expansions around the mean is that fourth order approximations are no better than second order approximations at estimating utility or ordering from utility. Unreported results substantiated Hlawitschka's (1994) claims to this effect. Although the third and fourth terms are small in absolute magnitude for funds in the desired range, they can strongly affect the approximation to utility without making improvements. The results from the Taylor series approximations should not be used alone to determine if a fund has non-normality present to affect investors or analysts rankings.

Non-Central *t* Distribution Simulations

To further investigate the normal equivalents we present their usefulness using simulated data. In this section we show that the normal equivalents are useable, fit with the Taylor series approximation results, and provide insight into the effect of the distribution on investors and performance analysis. The distribution used to simulate returns needs to have some control over the mean, standard deviation, skewness, and kurtosis in order to analyze the effects of the higher moments over the desired range. This lends itself to the selection of a censored non-central *t* distribution to accomplish this purpose (see Harvey and Siddique, 1999). A *t* distribution is symmetric with excess kurtosis. The kurtosis decreases as the degrees of freedom increases, approaching three (no excess kurtosis relative to the normal) as the degrees of freedom approaches infinity. From the work with Taylor Series approximation it is clear there is more of a need to produce a wide range in the kurtosis than producing platykurtosis (kurtosis less than three).

The *t* distribution is one of several that can take on a non-centrality parameter for different statistical needs. The non-centrality parameter specifically changes the skewness of the distribution. A negative non-centrality parameter produces a left or negatively skewed distribution, a positive non-centrality parameter produces a right or positively skewed distribution, and a non-centrality parameter of 0 replicates the symmetric *t* distribution. A non-central *t* random variable, with degrees of freedom ν and non-centrality δ, is the result of a standard normal random variable plus the non-centrality parameter, divided by the square root of a chi-square over its degrees of freedom:

$$t = \frac{z + \delta}{\sqrt{\chi^2 / \nu}}$$

The probability density function for the non-central *t* distribution is shown in Equation 3-6 (Evans, Hastings, and Peacock, 2000). The gamma function is defined below in Equation 3-7, which can be reduced for whole numbers (Mittelhammer, 1996) to Equation 3-8.

$$f(t) = \frac{\nu^{\nu/2} e^{-\delta^2/2}}{\Gamma\left(\frac{\nu}{2}\right) \times \sqrt{\pi}\left[\nu + t^2\right]^{[\nu+1/2]}} \times \sum_{i=0}^{\infty}\left[\Gamma\left(\frac{\nu+i+1}{2}\right) \times \left[\frac{\sqrt{2}t\delta}{\sqrt{\nu+t^2}}\right]^i \times \frac{1}{i!}\right] \tag{3-6}$$

$$\Gamma(a) = \int_0^{\infty} b^{a-1} e^{-b} db \tag{3-7}$$

$$\Gamma(a) = [a-1]! \tag{3-8}$$

The mean and variance for the non-central t distribution are defined below with c_{11} and c_{20} used to simplify the formulas (Hogben, Pinkham, and Wilk (1961)).

$$\begin{aligned}
\mathrm{E}(t) &= c_{11}\delta \\
\mathrm{E}(t - \mathrm{E}(t))^2 &= c_{20}\left[1+\delta^2\right] - c_{11}^2\delta^2 \\
c_{11} &= \sqrt{\frac{\nu}{2}} \times \frac{\Gamma\left(\frac{\nu-1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)} \\
c_{20} &= \frac{\nu}{\nu-2}
\end{aligned}$$

The skewness and kurtosis for the non-central t distribution are defined below with c_{11}, c_{20}, c_{31}, c_{33}, c_{40}, c_{42}, and c_{44} used to simplify the formulas.

$$\begin{aligned}
\eta_1 &= \frac{c_{33}\delta^3 + c_{31}\delta}{\left[c_{20}\left[1+\delta^2\right] - c_{11}^2\delta^2\right]^{3/2}} \\
\eta_2 &= \frac{c_{44}\delta^4 + c_{42}\delta^2 + c_{40}}{\left[c_{20}\left[1+\delta^2\right] - c_{11}^2\delta^2\right]^2} \\
c_{31} &= \frac{3c_{11}c_{20}}{\nu-3} \\
c_{33} &= 2c_{11}^3 - \frac{[2\nu-7]c_{11}c_{20}}{\nu-3} \\
c_{40} &= \frac{3\nu c_{20}}{\nu-4} \\
c_{42} &= 2c_{40} - 2[\nu-1]c_{11}c_{31} \\
c_{44} &= \frac{2[\nu-5]c_{11}c_{31}}{3} + \frac{c_{40}}{3} - 3c_{11}^4
\end{aligned}$$

A non-central t distribution can be shifted to any desired mean and scaled to any desired standard deviation without impacting the higher moments. The desired mean and standard deviation are identified in this paper as μ and σ and are attained by Equation 3-9. For comparison purposes, Figure 3-3 shows the cumulative (Panel A) and probability (Panel B) density functions of a non-central t distribution with δ=-1 and v=5 compared to the normal when both distributions have μ=1% and σ=5% to approximate monthly returns. The density functions shown here are atypical of the statistics literature. The t distribution is not usually shown using its standard deviation as the scale and when compared to the normal, the standard deviations are not the same. This problem is exacerbated for the non-central t distribution because its standard deviation is even larger and its mean is not adjusted either. It is important, however, for this study to show the difference in the distribution with the same mean and standard deviation as the normal.

$$x = \frac{\left[t_{\delta,v} - \mathrm{E}\left(t_{\delta,v}\right)\right]\sigma}{\sqrt{\mathrm{E}\left(t_{\delta,v} - \mathrm{E}\left(t_{\delta,v}\right)\right)^2}} + \mu \qquad (3\text{-}9)$$

Panel A shows that the probability of getting a return less than -6.8% or greater than +15.6% is exactly the same for the non-central t distribution as the normal distribution. Due to the negative skewness of the non-central t distribution the thicker tail emerges sooner on the left-hand side of the distribution. The probability density function of the non-central t distribution lies above that of the normal distribution above 18.0% and below -10.5%. It also has a higher peak; it lies above it between -0.8% and 6.3%.

An investment that is distributed as a non-central t distribution has a greater chance of a small gain, a smaller chance at a small loss or medium sized gain, and a larger chance of a medium or large loss or large gain. This investment may have only a 36% chance of losing money compared to 42% for a normally distributed investment with the same mean and standard deviation but when it does lose, it is more likely to lose big.

The higher moments are illustrated for different choices for the δ and v parameters in Table 3-2. The standard case for this study, and those producing the probability and cumulative density functions, has a skewness of -1.27 and a kurtosis of 13.32. Increasing the degrees of freedom causes the distribution to more closely approximate the normal (η_1=0 and η_2=3). As the non-centrality parameter goes to zero the distribution more closely approximates the normal distribution. Each of these two parameters, though, influences the higher moments differently. The non-centrality parameter affects the skewness more and the degrees of freedom parameter affects the kurtosis more. This is helpful for the simulations as some control over both of the higher moments is desirable.

The normal distribution and the non-central *t* distribution each range over the entire number line but can be censored or truncated to better approximate the reality facing an investor and allow utility functions to find a solution. For this application all returns below a loss of 99% (a gross return of 1%) are censored and replaced with a loss of 99%. Truncating would replace any observation below a loss of 99% with a new random draw from the distribution. Censoring is preferred in this case because the distribution is intended to simulate the possible returns for an investment and if the management company does lose more than their entire funds under management the result for an investor would not be a 'do-over' but would be a near complete loss.

For a normal distribution much less than one observation in a trillion would need to be censored. However, the fat lower tail of a *t* distribution with δ=-1 and v=5 means that 1 observation in 263,025 will need to be censored. Table 3-3 shows the effect of changing the standard deviation, non-centrality parameter, or degrees of freedom. As the standard deviation increases, the odds of censoring increase (the odds of censoring are 1:17,006 if the standard deviation is 9%). As the non-centrality parameter increases, the odds of censoring decrease (the odds of censoring are 1:438,943,140 if the non-centrality parameter is 2). Lastly, as the degrees of freedom increases, the odds of censoring decrease (the odds of censoring are 1:2,112,141,300 if the degrees of freedom parameter is 25).

Some of the effect of censoring the non-central *t* distribution can be describing by the moments of a non-central *t* distribution that is censored from below at -99%. Censoring increases the mean by 0.01 basis points and decreases the standard deviation by 0.27 basis points. The result of censoring has little effect on the first two moments, but is larger on the higher moments – it increases the skewness from -1.27 to -1.22 and decreases the kurtosis from 13.32 to 10.59. Since observations will be censored from below, the mean and skewness are higher. And since observations in the tail are censored, the standard deviation and kurtosis are lower.

The computation of the moments from the censored non-central *t* distribution follows from the probability density function, Equation 3-6. The expected value of the censored non-central *t* distribution raised to the power *j* is used as a reasonable approximation to find the moments.

$$\mathrm{E}(t_C^j) = F(-0.99) \times [-0.99]^j + \int_{-0.99}^{2.00} f(t) \times t^j dt + [1 - F(2.00)] \times [2.00]^j$$

$$
\begin{aligned}
\mu_C &= \mathrm{E}(t_C) \\
\sigma_C^2 &= \mathrm{E}\left(t_C^2\right) - \mathrm{E}(t_C)^2 \\
\eta_{C,1} &= \frac{\mathrm{E}\left(t_C^3\right) - 3\,\mathrm{E}\left(t_C^2\right)\mathrm{E}(t_C) + 2\,\mathrm{E}(t_C)^3}{\sigma_C^3} \\
\eta_{C,2} &= \frac{\mathrm{E}\left(t_C^4\right) - 4\,\mathrm{E}\left(t_C^3\right)\mathrm{E}(t_C) + 6\,\mathrm{E}\left(t_C^2\right)^2 - 3\,\mathrm{E}(t_C)^4}{\sigma_C^4}
\end{aligned}
$$

The effect of censoring decreases when the parameters lead to a more normal or less negatively skewed distribution because fewer observations are censored. Censoring is necessary for the mathematics and has a small and measurable impact on the moments of the sampling distribution.

The non-normality premium and utility adjustment to the standard deviation will be found for 120 observations (representing ten years of monthly observations) for 10,000 funds with data created from a censored non-central *t* distribution. Throughout the simulations the mean net monthly return is 1% and net returns below -99% are replaced with -99%. To show the effects of the standard deviation, skewness, and kurtosis on the non-normality premium and utility adjustment to the standard deviation, the parameters that influence them will be allowed to vary over a set range. The standard deviation will range from 5% to 9%, the non-centrality parameter will range from -2 to +2, and the degrees of freedom parameter will range from 5 to 25. Additionally, the risk aversion coefficient will range from 1 to 5 to show the effect of risk aversion on the non-normality premium and utility-adjusted standard deviation.

A non-normality premium of 10 basis points (or less than –10 basis points) will be considered as representing a fund with non-normality that is important to the investor. If an investor is willing to sacrifice 10 basis points per month (about 1.25% per year) in the mean return because of the non-normality then it has a noticeable effect on the investor. The next section shows that a *NNP* of 10 basis points is highly unlikely if the underlying distribution is normal. So, a *NNP* of 10 basis points per month is not only economically large, but would be considered outside statistical confidence bounds. A utility adjustment to the standard deviation, *UAS*, of more than 10% (or less than -10%) shows that the standard deviation is not an adequate measure for risk (damaging yardsticks such as the Sharpe ratio). The next section also examines *UAS*.

These determinations are very different than concluding that a fund does not have a normal distribution. This dissertation pursues the question of an investor's perception (using utility) of the non-normality of a distribution, not whether the distribution is normal or not.

Twenty runs of 10,000 funds consisting of simulated monthly returns for ten years were completed to find the percentage of funds that have non-normal returns that critically affect an investor (Table 3-4). Using the standard for this paper, μ=1%, σ=5%, δ=-1, ν=5, and γ=3, the percentage of funds that had a non-normality premium larger than 10 basis points ranged from 7.20% to 7.64%, with more than 99% of these having a positive premium (the sample distribution is utility detracting).

Panel A shows the impact of the risk aversion coefficient. Relatively less risk averse investors (γ=1) would be concerned with the non-normality in 0.56% of simulated returns, all with a positive premium. Relatively more risk averse investors (γ=5) would be concerned with the non-normality in 26.12% of simulated returns, with more than 99% of these having a positive premium. An increase in the risk aversion not only increases the weight in both ends, but also shifts weight to a larger (positive) premium because the sampling distribution has negative skewness (-1.3) and excess kurtosis (13.3).

Panel B shows the impact of the non-centrality parameter. When δ=-2, the distribution is more skewed (-2.1) and has even more kurtosis (20.8), which causes the number of funds that have economically meaningful non-normality to rise to 13.39%, all of which have positive non-normality premiums. When the distribution is symmetric, yet still has excess kurtosis (9.0), the number of funds with economically meaningful non-normality falls to 2.38% with 75% of these funds having positive non-normality premiums. When δ=2, the distribution is skewed to the right (+2.1) and is leptokurtic (20.8), 5.51% of the funds had economically meaningful non-normality and all of them had utility enhancing distributions. In this case, the utility enhancing positive skewness overcomes the utility detracting kurtosis and the net effect is only half as strong as when both are utility detracting. Overall, as the non-centrality parameter increases the funds have a more negative non-normality premium.

Panel C shows the impact of the degrees of freedom parameter. When there are 15 degrees of freedom only 0.06% of the funds showed non-normality that would be nontrivial to an investor. And just 0.01% of the simulated funds passed the cut-off when there were 25 degrees of freedom. The degrees of freedom parameter diminishes the skewness while reducing the kurtosis. With 25 degrees of freedom the skewness (-0.1) and the kurtosis (3.3) are nearly that of the normal. As should be expected, approaching the normal distribution diminishes the non-normality premium.

Panel D shows the impact of the standard deviation. When the standard deviation is increased to 7%, 28.77% of the funds showed economically meaningful non-normality, of which more than 99% had positive non-normality premiums. At 9%, 55.96% of the funds have non-normality premiums larger than ten (and still more than 99% of these were positive). A larger standard deviation increases the risk premium of a fund, which makes the non-normality relatively more important. This is shown by the pull towards more extreme non-normality premiums, particularly on the utility detracting side which the sampling distribution favors.

Strongly risk-averse investors need to exercise more caution when approaching funds with non-normal returns. All investors should exercise caution when the skewness and kurtosis are very high. A large standard deviation might cause investors some concern by itself, however, it also makes the higher moments more important as well.

Twenty runs of 10,000 funds consisting of simulated monthly returns for ten years were also completed to find the percentage of funds for which the utility adjustment to the standard deviation is more than 10% different than the sample standard deviation (Table 3-5). Using the base-case for this paper, μ=1%, σ=5%, δ=-1, ν=5, and γ=3, the percentage of funds that had more than a 10% difference ranged from 6.54% to 7.18%, with more than 99% of these having a utility adjustment to the standard deviation greater than the sample standard deviation (the sample distribution is utility detracting). Funds such as these have risk that is not well identified by the sample standard deviation. Hence, performance measurements based on the sample standard deviation, such as the Sharpe ratio, will produce ratings and rankings that are inconsistent with the risk to utility. The Sharpe ratio is calculated by dividing reward (the sample mean less the risk free rate) by the risk (the sample standard deviation).

The risk aversion coefficient for the utility function increases the importance of losses relative to gains. Panel A shows that if the risk aversion is assumed to be small (γ=1), then adjusting the standard deviation to account for utility will make little difference to the Sharpe ratio. The Sharpe ratio would be overestimated by 10% (because the risk would be underestimated by 10%) in only 1.45% of the funds simulated from a censored non-central *t* distribution with a mean net return of 1%, standard deviation of returns of 5%, non-centrality parameter of -1, and 5 degrees of freedom. However, the Sharpe ratio would not be underestimated for any of the 10,000 simulated funds by more than 10%. The Sharpe ratio would be overestimated by 10% in 16.91% of the funds when the risk aversion coefficient is 5, and would be underestimated by 10% in 0.05% of the funds. Increasing risk aversion makes the Sharpe ratio less reliable for rating funds because its measure of risk (standard deviation) is less accurate.

The non-centrality parameter affects the skewness and kurtosis of the non-central *t* distribution. Panel B shows that if the non-centrality parameter is decreased (δ=-2), then adjusting the standard deviation to account for utility is even more necessary as 13.50% of the funds' Sharpe ratios would be overestimated by more than 10% compared to the normal equivalent's measure. (Again, in none of the 10,000 cases would the Sharpe ratio be underestimated by more than 10%.) A symmetric *t* distribution caused 1.60% of the funds to have an overestimated Sharpe ratio and 0.30% of the funds to have an underestimated Sharpe ratio. More funds would be overestimated than underestimated because although the skewness is gone, excess kurtosis still persists. A positive non-centrality parameter of 2 causes the skewness to be 2.1 and the kurtosis to be 20.8. In this case, 0.01% of the funds would have an overestimated Sharpe ratio and 4.70% of the funds would have an underestimated Sharpe ratio. Negative skewness is more influential than positive skewness for the risk to utility because it has an extremely fat lower tail where marginal utility is the greatest.

The non-centrality parameter affects the skewness and kurtosis of the non-central *t* distribution. Panel C shows that if the degrees of freedom are increased (making the distribution more normal), fewer funds would produce utility adjustments to the standard deviation that are more than 10% different than the sample standard deviation. With 15 degrees of freedom (η_1=-0.2 and η_2=3.6) only 0.02%, and none of the simulated funds with 20 and 25 degrees of freedom, had utility adjustments to the standard deviations that were very different than the sample standard deviation. As expected, the closer the sampling distribution is to the normal, the better the standard deviation is at identifying risk.

While the standard deviation does not affect risk aversion nor does it affect the skewness or kurtosis it does have a strong impact on an investor's utility and the larger it is the more relevant the lower tail becomes. Panel D shows that as the standard deviation increases, the probability of a fund having a utility adjustment to the standard deviation more than 10% greater than the sample standard deviation grows. For the simulated funds, a standard deviation of 5% leads to 6.93% of the funds needing an adjustment, at 7% it is 15.51%, and at 9% it is 25.51%. The probability of a fund having a utility adjustment to the standard deviation more than 10% less than the sample standard deviation was consistently less than 0.05%. The effectiveness of standard deviation to measure risk is diminished as the standard deviation increases.

The results of the ratio of the utility adjustment to the standard deviation to the sample standard deviation using the 20 simulations from a censored non-central *t* distribution parallel that of the non-normality premium. The ability of the Sharpe ratio to judge performance is undermined by substantial abnormality (skewness and kurtosis). Also, the more risk averse the investor, the more risk is misidentified by the standard deviation. And finally, the higher the standard deviation, the more important the higher moments become to an investor, which the Sharpe ratio ignores.

In this section, the normal equivalents were applied to several sets of simulated data from a censored non-central *t* distribution to show the importance of the normality assumption for investors and performance measurement. The rejection of normality did not mean that investors were concerned with normality nor did it lead to poor analysis with the Sharpe ratio. However, considering that negative skewness and excess kurtosis may be prevalent for specific types of funds such as hedge funds, the non-normality premium improves investor's understanding of the fund. The utility adjustment to the Sharpe ratio standardizes the distribution to make a more fair rating and further a better comparison because funds may have utility enhancing or utility detracting return distributions.

Normal Distribution Simulations

In this section, the likelihood of normally distributed returns achieving economically meaningful levels of non-normality is pursued with evidence from the normal equivalents. The normal equivalents cannot be used to test for normality. But are simulations from normally distributed returns capable of showing non-normality economically meaningful to investors and analysts? Using simulations from a normal distribution will also be helpful in determining the consistency of the measures.

To answer these questions, returns are simulated for ten years of monthly data from the normal distribution using identically and independently draws from set parameters. Following the convention of the paper, the net mean monthly return is assumed to be 1% and although different simulation runs will have different standard deviations, the focus will be around 5%. Furthermore, the risk aversion parameter will range from one to five and each simulation will have 10,000 funds. Observations that fall below a loss of 99% (a gross return of 1%) are replaced with a loss of 99%.

The non-normality premiums generated by 10,000 simulated runs of 120 months of identically and independently distributed normal returns (1% mean net return, 5% standard deviation, and a risk aversion level of three) are very close to zero. Table 3-6 records that only one percent of them are below -1.6 basis points (utility enhancing) and only one percent of them are above 1.1 basis points (utility detracting). These non-normality premiums are very small, rare were the cases that a simulated fund would even require a 1.6 basis point improvement in the mean in order to compensate investors for getting an exact normal distribution instead of what they realized. Equally rare were the cases in which an investor would be willing to accept a 1.1 basis point reduction in the mean in order to get the normal distribution.

Panel A shows that as the risk aversion increases, so too does the spread in the percentiles. A risk aversion of one is low enough that one percent of funds have a non-normality premium of -0.3 basis points or lower and another one percent have non-normality premiums of 0.1 basis points or higher. In fact, at a risk aversion level of one, 90% of the simulated funds had non-normality premiums between -0.2 and 0.0. Increasing the risk aversion to five, one percent of funds were able to generate non-normality premiums as low as -3.6 and another one percent generated non-normality premiums as high as 3.2. Still these premiums are very small and 90% of the funds were between -2.7 and 1.7. Risk aversion causes investors to avoid risks, the concavity of the utility function causes an increase in the weight on the utility function for large losses.

Panel B shows that the non-normality premiums are greatly influenced by the standard deviation. An increase in overall risk increases the probability in the tails, which is particularly distasteful to risk averse investors. If the standard deviation is decreased to 3%, then the first percentile is at -0.4 and the ninety-ninth percentile is at 0.1 basis points. If the standard deviation is increased to 9%, then the first percentile is at -9.1 and the ninety-ninth percentile is 8.7. A nine basis point non-normality premium may appear very small, however these are monthly returns so an investor that must be compensated thusly for assuming the normal distribution is equivalent to an investor requiring a 13.2% normally distributed return to achieve the same utility as a 12.0% return from a utility enhancing distribution in annual terms.

Panel C shows that as the sample size increases the non-normality premiums shrink. This should be expected because the larger the sample size, the closer the fit will be to the normal distribution. The first percentile is -2.5 and the ninety-ninth percentile is 1.4 basis points with sixty months of simulated data. These same percentiles are -1.0 and 0.8 with two hundred and forty months. Note also that throughout all these panels the non-normality premiums tend to run stronger on the negative side, representing distributions that are preferable to the normal. This declines as the sample size increases. It would seem that a discrete distribution is slightly better than a continuous distribution. This comes because a continuous distribution will typically give a sliver of weight to outcomes that are much worse than what appears in a discrete distribution and much better than what appears in a discrete distribution and that the lower tail has a larger importance on utility than the upper tail. Larger sample sizes decrease this difference because the shape of the simulated distribution will more closely resemble the normal.

Not surprisingly, Table 3-7 shows that the utility adjustment to the standard deviation tracks closely with these results. One percent of funds had a utility adjustment of less than -2.0% and another one percent had an adjustment of more than 1.3%. A utility adjustment of -2.0% indicates that the risk of the simulated fund was equivalent to a standard deviation of 4.9%, not 5% because the distribution was preferable to the normal. The results of Panel A show that a risk aversion level of one will cause the first and ninety-ninth percentiles to be -1.2% and 0.4% but that a risk aversion level of five causes them to be -2.8% and 2.3%. Relative risk aversion increases the need for risk measures that incorporate the higher moments.

Panel B in Table 3-7 has a substantial difference than Panel B in Table 3-6. In Table 3-6, the percentiles move outward dramatically as the standard deviation increases. In Table 3-7, the percentiles move out at a slower rate because the divisor is increasing. The increased damage to risk is partially compensated by the fact that a ten basis point adjustment to the standard deviation has a far smaller impact on the Sharpe ratio when the standard deviation is 9% as opposed to 3%. Yet still the range increases, meaning that the higher moments' importance to the Sharpe ratio increases with increased standard deviations. When the standard deviation is 3% the first percentile is -1.4% and when it is 9% it is -3.4%. When the standard deviation is 3% the ninety-ninth percentile is 0.6% and when it is 9% it is 3.2%. Also interesting is that the range on each side becomes more similar with the increased standard deviation.

Panel C shows the sample size has the same impact on the utility adjustment to the standard deviation as it does the non-normality premium. The percentiles are closer to zero and more symmetric as the sample size increases. The first percentile is -3.0% when five years of data are simulated and -1.3% when twenty years of data are simulated. The ninety-ninth percentile is 1.7% when five years of data are simulated and 1.0% when twenty years of data are simulated.

The simulations with normal distributions show that identically and independently draws from normally distributed returns generally lead to very trivial non-normality premiums or adjustments to the standard deviation and Sharpe ratio. However, more risk aversion or higher standard deviations increase the likelihood of a sample that would statistically reject normality also producing non-normal affects to utility. Small samples also increase this likelihood. Even a sample as large as five years with a standard deviation of nine percent will occasionally reject normality statistically, and do so in an economically meaningful way.

Chapter Summary

Throughout this dissertation the assumption of a ten basis point non-normality premium or a ten percent difference between the normal equivalent-standard deviation and the standard deviation plays a large impact on the conclusions. A ten basis point premium amounts to a difference of about 1.35% per year and a ten percent difference in the standard deviation would amount to a Sharpe ratio that is off by ten percent, whether in monthly terms or annualized. This chapter reveals the skewness and kurtosis, in combination with the risk aversion and standard deviation, which is necessary to achieve these levels. From the previous chapter, a statistically significant difference to the Sharpe ratio will require a greater change in the risk than ten percent and a smaller number of simulated funds would reach this level. The boundaries are set at where investors and analysts might perceive a difference, whether it is statistically significant or not.

CHAPTER FOUR: APPLYING THE NORMAL EQUIVALENT

Hedge funds provide an interesting source for the application of the normal equivalents because their returns are not normally distributed. Several recent papers have centered on how their return distributions impact portfolio evaluation measures and possible corrections (see Amin and Kat, 2003, Capocci and Hübner, 2004, Fung and Hsieh, 1999b, Kouwenberg, 2003, and Liang, 2000). But failing a normality test is not a determination of whether the return distribution is undesirable; some non-normal return attributes may be valued by investors. I analyze hedge fund return distributions in the framework of investor utility, thus they are able to describe which characteristics are desirable to investors and which are not.

The non-normal returns of hedge funds may also affect performance evaluation. It is common in the investment community to evaluate hedge fund performance using the Sharpe ratio. However, the Sharpe ratio uses only the first two moments of the distribution (mean and standard deviation) in its computation. It assumes that either the distribution is normal or that investors are not impacted by higher moments. As most hedge funds returns are not normal, I assess the robustness of the Sharpe ratio for fund evaluation.

This chapter uses normal equivalent methods to examine the return distribution when an investor has a power utility function. Normal equivalents find the mean or standard deviation necessary for any return distribution to have the same utility to the investor as the normal distribution with the other parameter unchanged. The normal equivalents permit a description and quantification of the shape of a distribution's affect on utility and create a standard for comparing funds with different distributions on an equal footing. In addition, they can be used to illustrate what non-normal characteristics are desired by investors and which are undesired.

While hedge funds are typically evaluated in isolation, investors should also assess their desirability in a portfolio context. Therefore, this chapter also examines the return distribution of a portfolio consisting of the risk-free rate, stocks, and a hedge fund. First, ten percent of the portfolio is allocated to a hedge fund. Then, utility is used to optimize the allocation to stocks and the risk-free rate. The return distribution of this optimal portfolio is examined and compared to a utility maximizing portfolio with no hedge fund allocation. Investors are found to benefit from investing in hedge funds regardless of their risk aversion and the magnitude and coverage of the higher moments is greatly reduced.

Hedge Funds

Recent estimates of the hedge fund industry place the number of funds at 6,000 with $400 billion in capital (Capocci and Hübner, 2004). This is small relative to the mutual fund industry, which has more than $6 trillion in capital, yet attention to the industry exceeds its size due to its effect on the markets, remarkable returns, and managers' fame. The curiosity into hedge funds is further heightened because they are exotic, secretive and forbidden to the common man.

In 1949, Alfred Winslow Jones wrote "Fashions in Forecasting" for *Fortune* magazine on technical analysis and market timers. From his research he developed the idea for a market-neutral, or hedged, portfolio and began pursuing it by establishing a partnership even before the article was published. The fund's originality was not limited to hedging; the use of hedging allowed for leveraging large positions with less capital and the fund had an incentive fee of 20% of profits in lieu of a fixed fee. Seventeen years later Mr. Jones was featured in an article in *Fortune* magazine and other managers began their own funds with some similar features (Brown, Goetzmann, and Ibbotson, 1999). Today's hedge funds are a heterogeneous group. Many hedge funds do not hedge and their degree of leverage may vary, but most maintain the incentive fees.

Bookstaber (2003) contends that the only suitable description for all hedge funds is what they are not – a traditional investment fund.

Hedge funds differ from mutual funds because the managers choose to operate in a different realm of government regulation (Fung and Hsieh 1999a). They forfeit the ability to accept non-accredited investors' money and the ability to solicit in order to avoid government regulations of transparency, soundness, liquidity, and agency. In addition, hedge funds consider their trading strategies and holdings to be proprietary information and that public knowledge would trigger copycats and frontrunners, drastically reducing their ability to generate a profit. Hedge funds use derivative contracts to conduct arbitrage, isolate risks, and leverage their portfolio. Because their investment strategies may be illiquid, they use lock-ups and only allow investor purchases and redemptions with advance notice. Mutual funds, by their use of standard stock and bond investments, are generally assumed to have returns with normal distributions. Because hedge funds use investments with non-linear payoffs and a wide array of return distributions, their returns will almost certainly deviate from normality.

Because they do not follow the same regulations as mutual funds, hedge funds are not allowed to solicit new investors and they cannot advertise to the public. This outlaws not only newspaper, magazine, radio, and television advertisements, but also unrestricted Internet postings. The number of investors allowed into a hedge fund is limited by law, usually to less than 100, of which only a few may be non-accredited investors. To be accredited an individual must have a net worth of more than $1,000,000 or an income greater than $200,000 in each of the past two years and a trust or organization must have more than $5,000,000 in assets.

Four papers draw the conclusion that the majority of hedge funds' returns are not normally distributed. Fung and Hsieh (1999b) state that hedge funds' returns are not normal.

Kouwenberg (2003) finds that 735 of 1,399 hedge funds, over a six-year period, reject normality at the five percent level and most of these funds have negative skewness and excess kurtosis. Amin and Kat (2003) reject normality for 66 of 77 hedge funds with ten-years of monthly data, noting their skewness in particular. Sharma (2004) rejects normality for 418 in his sample of 788 hedge funds with five-years of monthly returns, noting the average of the kurtosis is 6.02, 3.02 larger than the normal distribution.

The Sharpe ratio, the industry standard performance measure, relies on either normally distributed returns or quadratic utility to make a ratio of reward per unit of risk. Reward is defined as the mean of the returns and risk is defined as the standard deviation of the returns. If returns are normally distributed the mean and standard deviation are sufficient statistics for the entire distribution and the standard deviation is sufficient for measuring risk – regardless of the investor's utility function. Since many hedge funds reject normality, the Sharpe ratio must rely on investors having quadratic utility functions to rank funds. Scott and Horvath (1980), however, show that investors like odd moments, such as skewness, and dislike even moments, such as kurtosis (negative skewness and excess kurtosis are utility detracting features for a return distribution). Even with these preferences, several authors use Taylor Series approximations to demonstrate that the loss due to the higher moments is trivial. These are discussed in Chapter Three and include papers by: Levy and Markowitz (1979), Hlawitschka (1994), and Fung and Hsieh (1999b).

Other authors seek modifications to account for non-normality. Kouwenberg (2003) modifies the regression to compute Jensen's alpha by using a weighted least squares (WLS) model. The weights are the marginal utility of the return using the power utility function with the risk aversion parameter set to six. The WLS model causes returns in the left tail to be more heavily weighted than those in the right because this utility function has the desirable property of decreasing marginal utility. Comparing the WLS alpha to the ordinary least squares he makes the determination that the difference is "rather small … suggest[ing] that the non-normality of hedge fund returns does not have a major and uniform impact on the portfolios of passive investors." (page 372)

Sharma (2004) uses the certainty equivalent, to evaluate the performance for a set of hedge funds. Sharma also uses the power utility function and sets the risk aversion parameter at four. The certainty equivalent allows him to rank the hedge funds by utility and find the correlation between the Sharpe ratio and the certainty equivalent to be 0.46. Among the advantages of using the certainty equivalent to evaluate hedge fund performance is that it adequately adjusts for the higher moments. Another advantage is that it punishes funds for increased leverage. The Sharpe ratio is invariant to leverage; leverage increases the mean and standard deviation at the same rate. The certainty equivalent's sensitivity to leverage can also be a disadvantage for performance evaluation. An investor can be assumed to be diversified and allocate wealth differently depending on its risk. A fundamental assumption of the Sharpe ratio is that two funds with the same Sharpe ratio, yet different standard deviations, are equally efficient and are equally attractive (Figure 2-1). Evaluating performance by the certainty equivalent does not maintain that assumption.

Many hedge funds have non-normal returns, yet the question remains as to what extent the Sharpe ratio's attempt to measure efficiency is undermined by its use of standard deviation to measure the risk to utility. To this end, normal equivalents are used to equate the non-normality with a change in the mean and standard deviation. This will allow a determination of the importance of the normality assumption on investors and the Sharpe ratio.

Data

This paper focuses on the distribution of hedge funds and the effects of the distribution on the utility of investors. This is done on both a stand-alone investment into a hedge fund and on a portfolio that also includes the stock market and risk free rate. The descriptive statistics will mimic the popular press: the means and standard deviations are net returns (as opposed to gross) and net of fees. The correlation and beta statistics will incorporate the market and one-month T-bill rates of return. The market rate of return is CRSP's monthly value weighted returns and the one-month T-bill rates are from Ibbotson and Associates, Inc., both made available at Kenneth French's website at: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. The returns are in real terms, using the Consumer Price Index (available at the Bureau of Labor Statistics website: http://www.bls.gov/cpi/home.htm) to adjust for inflation.

Data was collected on 373 funds with complete monthly return data from 1996 to 2005 available from hedgefund.net. Efforts were made to combine funds of the same fund family that provided identical or nearly identical returns. A majority of the funds (211) were U.S. onshore hedge funds, 131 were offshore hedge funds, and 31 had both onshore and offshore accounts available. All funds were examined to ensure the returns were denoted in dollar terms to avoid problems with exchange rates. The descriptive statistics of another 32 funds that did not have complete reporting are shown in Table 5-1. Eleven of these funds reported their December 2005 return in March 2006 after the analysis was done. Four more funds did not report December 2005 returns by April 2006. Finally, an additional seventeen funds from a previous study of 230 onshore hedge funds and commodity trading advisor funds (1995-2004) dropped out in 2005.

The funds report their general strategies among several options but are not limited to a single category. It is valuable to note some of the most popular strategies. Table 4-1: Panel A shows that a sizeable portion of them are Fund of Funds – hedge funds that invest in other hedge funds (less than 140 because a few chose more than one Fund of Funds category). Another large group identified themselves with their relation to the equity markets. Of these, 7 have a short bias and 15 are market neutral and only 16 identified themselves as Long Only. For other strategies, 55 identified themselves as Convertible Arbitrage, Capital Structure Arbitrage, Fixed Income Arbitrage, Merger/Risk Arbitrage, Statistical Arbitrage, or Other Arbitrage. Another 34 identified themselves as Value and 25 identified themselves as seeking Distressed Securities – usually in conjunction with one of the previously listed strategies, or Event Driven or Special Situations.

Hedge funds are noted for their abilities to generate above average returns with small standard deviations and little correlation with the general stock market. Panel B shows that the average monthly return for these 373 hedge funds was 1.01% per month, 15 basis points better than the stock market averaged over this time period. The average standard deviation was 3.36%, much less than the market (4.66%). In fact, 36% of the funds had both a higher mean and less variance than the market. Figure 4-1 shows the range in the means and standard deviations for the sample. The average correlation with the stock market was 42%, although this ranged from -90% to +91%. (Enthusiasm about hedge fund returns should be tempered because of such biases as backfilling or survivorship that are prevalent, see the "Data Biases" Section of Chapter Five.)

As should be expected, the average correlation with the market among those that categorize themselves as short bias hedge funds was -56%. Also in line with expectations, the average correlation with the market was 7% for market neutral equity hedge funds. Continuing with this trend, the average correlation rose to 65% for long only hedge funds. Interestingly, of the funds that selected long/short equity for their category had an average correlation of 41%. It appears they are more long than short from this correlation.

Lo (2002) showed that autocorrelation affects the annualization of the Sharpe ratio. The Sharpe ratio is annualized from monthly returns by multiplying by the square root of twelve because it is assumed that the yearly mean return is twelve times larger than the monthly and that the yearly standard deviation is only the square root of twelve times larger than the monthly standard deviation (Equation 2-2). However, the yearly standard deviation is more than the square root of twelve times larger if the fund has positive autocorrelation. Funds with a positive autocorrelation have Sharpe ratios that are upwardly biased. (Negative autocorrelation has the opposite effect). The autocorrelation of these funds averaged 17%, ranging from -37% to +96%. Considering that the market's autocorrelation was only 4% and 301 of the hedge funds had a larger autocorrelation than this, it might be easy to decide that autocorrelation is a large problem and that perhaps some of them are smoothing their returns. However, keeping in mind that many of these funds attempt to hedge and that the risk free rate had an autocorrelation of 96%, such numbers are reasonable. Yet this still leaves the Sharpe ratios biased upwards for most hedge funds compared to mutual funds which more closely track the stock market and not reliable at ranking hedge funds against each other.

Lastly, and of primary importance for this paper, 86% of the funds reject normality at the 95% level using the Jarque-Bera test (Jarque and Bera, 1987) in Equation 4-1. Figure 4-2 plots the skewness and kurtosis of the 373 funds with a preface of how many funds fit into each of the eight possible configurations. In total, 53% of the sample had negative skewness and the left tail was very long, extending out to -8.6 (The mean skewness was negative and the skewness of the skewness was very negative). While it has been noted that hedge funds often have negative skewness, the kurtosis was a much larger factor in rejecting normality. Only eleven of the funds had less kurtosis than the normal distribution and only one of them rejected normality. The kurtosis ranged as high as 89.7 although the median was only 5.8 (which is still sufficient to reject normality). A scatterplot of the skewness and kurtosis looks like a crooked smile because there is a 93% correlation between the absolute value of the skewness and the kurtosis – funds that were not symmetric also had excess kurtosis. The smile is crooked because the left side is longer – a few funds have very strongly utility detracting features.

$$JB = T\left[\frac{h_1^{\ 2}}{6} + \frac{[h_2 - 3]^2}{24}\right] \tag{4-1}$$

$$JB \sim \chi_2^2$$

How does this sample compare with that used by recent researchers? Capocci and Hübner (2004) have a sample of 2796 funds by combining the Zurich Hedge Fund Universe and HFR databases from 1994 to 2000. The mean return is 1.29% and the standard deviation is 5.56%. Sharma (2004) has a sample of 787 individual (not including fund of funds) HFR funds from 1997 to 2001. The mean return is 6.53% and the standard deviation is 16.55% (annual and subtracting the risk free rate). The mean skewness and kurtosis (calculated monthly) are -0.14 and 6.02. The skewness ranges from -7.18 to 5.78 and the kurtosis ranges from 2.14 to 54.41. 53% of his sample reject normality at the 5% level using the Jarque-Bera test (Equation 4-1). Kouwenberg (2003) uses the Zurich Hedge Fund Universe database from 1995 to 2000. The mean return is 13.7% and the standard deviation is 14.2% (annual). An estimate of the average skewness and kurtosis from his reported results are -0.46 and 6.21 (about 52% reject normality – again using the Jarque-Bera test of Equation 4-1). The sample funds from this chapter generally have a lower mean and standard deviation with similar higher moments. The increased rejection rate of normality comes in part from the longer time frame.

Results

Although 86% of the funds had return distributions that were statistically different than the normal, only 4% of the funds were economically different than the normal (with the risk aversion parameter set to three). Breaking down Table 4-2 shows that seven funds had non-normality premiums of -10 or less (the lowest being -26.1). These funds had such utility enhancing higher moments that an investor would be willing to assume a normal return distribution only if they were compensated by more than a ten basis point higher monthly expected return. While all seven had the utility detracting feature of excess kurtosis, the benefits of their positive skewness was overwhelming. Nine funds had non-normality premiums of +10 or more (the highest being 116.3). These funds had such utility detracting higher moments that an investor would be willing to sacrifice ten basis points, or more, per month in expected return in order to have the normal distribution. The extreme example came from a fund practicing both a short bias and an arbitrage strategy that lost 60% in a single month.

An interesting occurrence in the nine utility detracting funds is that one of the funds actually had positive skewness (0.2) and a kurtosis that was smaller than the mean (7.1). In this case, the fund had such a large standard deviation (14.4%), the second largest of all 373 funds, that the detrimental impact of the excess kurtosis was enough to cause a non-normality premium of 42.9 basis points. A large standard deviation results in a magnified non-normality premium because the risk premium is so large. It should be noted that the largest standard deviation was 15.5% and that fund had a skewness and kurtosis of 1.1 and 11.9, respectively, bringing the non-normality premium to 4.4.

Most of the non-normality premiums were relatively close to zero because most of the funds had such small standard deviations that there was little risk to begin with. Of the 245 funds with a standard deviation of less than 4%, only 2 had a *NNP* of more than 5, while 22 of the 50 funds with a standard deviation of more than 6% had a *NNP* of greater than 5 or less than -5. In order for the higher moments to matter, there must be a hefty amount of risk shown in the standard deviation as well.

The non-normality premium is also sensitive to the risk aversion parameter. An investor that is relatively less risk averse would have a smaller risk aversion parameter. By decreasing the risk aversion parameter to one, no funds have an economically meaningful non-normality. Increasing the risk aversion parameter to five brings the number of funds with economically meaningful non-normality to 31 (eight percent of the sample). The application to hedge fund data matches the explorations of Taylor series approximations and simulations with non-central *t* distributions; the non-normality is more important as the portfolio standard deviations and investor sensitivity to risk increases.

The non-normality premiums are computed without regard to the risk-free rate, but because the normal equivalent-standard deviation and the utility adjustment to the standard deviation are meant to correspond to the Sharpe ratio, they are computed net of the risk free rate. Using the risk aversion parameter of three, eleven funds (three percent) have a standard deviation that is not sufficiently capturing risk. Two of these funds have such utility enhancing higher moments that the standard deviation overestimates their risk, and understates their Sharpe ratio by ten percent (See Table 4-3). These two funds had a positive skewness of more than 3.5 which dampens the possibility of a large loss that would devastate an investor's utility. The other nine funds have their risks underestimated with the standard deviation, thereby overestimating their Sharpe ratio by ten percent.

A large non-normality premium can only emerge if the standard deviation is large enough to sufficiently magnify all risks – including risks coming from the higher moments. This is diminished for the utility adjustment to the standard deviation because it is in percentage terms of the standard deviation. Although diminished, it still exists because the impact on utility increases faster than the standard deviation. There is a 31% correlation between the standard deviation and the absolute value of the *UAS* (down from 44% for the absolute value of the *NNP*).

Just as the non-normality premiums move to the tails with increased risk aversion, so too do the utility adjustment to the standard deviations. As shown in Table 4-3, five funds have *UAS*s below -10% and seventeen funds have *UAS*s above +10% when the risk aversion parameter is increased to five. Decreasing the risk aversion parameter to one cuts the number of funds with overestimated Sharpe ratios to three and left no funds with underestimated Sharpe ratios.

The determination of the appropriate risk aversion level is very important to determining if the Sharpe ratio is adequate and that risks in the higher moments have a negligible impact on utility. Even though only three percent of the sample had mis-measured Sharpe ratios, the use of Sharpe ratios as a selection tool among funds more than doubles the difficulty. This is because Sharpe ratios can, and do, miss on both sides. While this leaves them unbiased (the mean and median utility adjustment to the standard deviation are 0.3% and -0.1%, respectively), it is entirely possible that a comparison of two Sharpe ratios will be in error by ten percent even if neither fund has such a large error by itself. The non-extreme 90% of the funds have *UAS* that range from -3.9% to 4.5% (See Table 4-4: Panel C).

None of the 373 funds have sufficient non-normality that the Sharpe ratio is statistically different. Figure 4-3 shows that none of the points even approach the 95% confidence interval; the largest p-value for a test of a significant difference between the Sharpe ratio and the Sharpe ratio-normal equivalent is 58%. Even if the risk aversion parameter is increased to five none of the funds had statistically different Sharpe ratios.

Further evidence that the funds' Sharpe ratios were not strongly affected by the inclusion of their non-normality into their risk measure comes from the use of the Spearman rank correlation. The non-parametric Spearman rank correlation (shown below) is preferred to the parametric Pearson correlation when there might be a nonlinear correlation. Each of the hedge fund's Sharpe ratio is ranked and compared to their ranking using the normal equivalent-standard deviation in place of the standard deviation for the Sharpe ratio. The correlation is 99.77% and ranges from 99.53% to 99.95% for the different risk aversion parameters (see Table 4-5).

$$\rho = 1 - \frac{6\sum_{i=1}^{N} \text{Difference}_{\text{i}}^{2}}{N\left[N^{2} - 1\right]} \tag{4-2}$$

The strategies of the hedge funds may have little impact on the non-normality premiums and the utility adjustment to the standard deviation. Tables 4-6 and 4-7 reveal the minimum, median, and maximum non-normality premiums and utility adjustment to the standard deviations for the different categories. Besides identifying the most extreme observation as an on-shore fund practicing both a short bias and arbitrage strategy no determinations can be made. This is because the median, which would show the non-normality tendencies of the funds are nearly all zero. The most extreme median is with the short bias funds, which is also the smallest group and therefore least reliable.

After examining the non-normality of hedge funds as a stand alone investment, the non-normality of a portfolio that includes a hedge fund is pursued to better match their impact on investors. Supposing investors who are optimally balanced between the stock market and T-bills were to place ten percent of their investment into a hedge fund and then optimally rebalance. The Section in Chapter Five titled "Hedge Funds in a Portfolio" describes the process of optimal balanced portfolios based on utility, the impact the hedge funds have on portfolio allocations, the value of a hedge fund, and the changes in the risk and return of the portfolio with their inclusion. This chapter continues with the impact including a hedge fund has on the non-normality.

Normality tests for the portfolios with ten percent allocation to hedge funds shows that 79% of them (optimized for a risk aversion level of three and a ten percent investment into the individual hedge fund) statistically reject normality. None of the portfolios had positive skewness and just eight had less kurtosis than the normal. All of the portfolios have skewness and kurtosis that is tightly packed, the strongest skewness is -1.01 and coincided with a kurtosis of 5.57 (see Table 4-8). The optimal portfolio without hedge funds had a skewness of -0.62 and a kurtosis of 3.41 (excess kurtosis of 0.41) which is enough to reject normality. The ten percent share to the hedge fund is not large enough to seriously impact the higher moments.

The spread in the higher moments is affected by the size of the allocation into the hedge fund. The greater the allocation, the greater the spread in the higher moments as the funds are less similar. With an allocation of 25%, the skewness approaches two and the kurtosis passes ten in the most extreme case. Counter-intuitively, the skewness shifts more negative and the range in the kurtosis increases as the optimal allocation includes a smaller allocation to stocks and a greater T-bill allocation. By themselves, the stocks have more negative skewness and larger kurtosis, so this is an unexpected result. The median skewness is -0.63 and the median kurtosis is 3.48. A difference of only -0.01 in the skewness and +0.07 in the kurtosis. Examining the medians of the higher moments of portfolios including hedge funds by their strategy shows little difference between the types of funds. Each strategy produces medians very near to that of the overall median.

The non-normality premiums for the portfolios including hedge funds have a very tight range (Table 4-6: Panels A and B). The inner ninety percent of the non-normality premiums for these hedge funds as a stand alone investment is from -3.4 to +2.9. The inner ninety percent for the portfolios including a hedge fund is from +0.4 to +0.7. The decreased range is similar across all risk aversion levels. In fact, the highest non-normality premiums are with the risk aversion level of three. This is because the riskiest hedge funds are compensated by a larger investment into the risk free investment which reduces the non-normality of the higher moments. A similar impact is shown in Table 4-7: Panels C and D with the utility adjustment to the standard deviation. The inner ninety percent ranges from -3.9% to 4.5% for hedge funds by themselves and only from 1.1% to 1.7% for a portfolio including hedge funds. The range is very small, yet a slight bias creeps in. This is because for many of the funds, the standard deviation is larger with an allocation into the stock market and the stock market has a slightly utility detracting distribution. This is also evidenced in the reduced *UAS* for higher risk aversion levels – with lower allocations into the stock market. Table 4-6 and Table 4-7 show that there is no particular strategy that has *NNP*s or *UAS*s that are unusual either by themselves or in a portfolio setting.

Chapter Summary

Chapter four applies the normal equivalents to a sample of hedge funds from hedgefund.net that reported monthly data for every month for ten years. 86% of the funds reject normality as a stand alone investment and 79% of portfolios with a ten percent allocation into a fund reject normality. Hedge funds are a fruitful source for a discussion on the impact of non-normality because there is a wide range in the means, standard deviations, skewnesses, and kurtoses. Despite the large number of funds that statistically reject normality, only 4% of the funds have enough to cause an investor to be willing to trade ten basis points for normally distributed returns and only 3% of them cause the Sharpe ratio to be off by ten percent. Additionally, none of the funds have distributions strong enough that a reasonable portfolio with will have economically meaningful non-normality.

However, the risk aversion level of three is perhaps on the lower end and a more risk averse parameter would better approximate an investor's utility function. If that is the case, several more funds will have economically meaningful return distributions. Also, these are surviving funds that may have been lucky rather than good and represent the better than average return distributions, further cautioning investors into a hedge fund with an unknown future. In summary, these hedge funds had favorable returns and rarely generate economically meaningful non-normality, but more risk averse investors may want to exercise more caution than these numbers suggest and remember that this sample may be atypical.

CHAPTER FIVE: ROBUSTNESS

Utility Functions

The power utility function features constant relative risk aversion, CRRA. This means that the increase in utility is the same for a given percentage return regardless of the starting wealth. (If an investor's wealth increases 10%, her utility increases by the same factor whether she started with \$100 or \$100,000.) This makes the starting wealth unimportant in the derivation of the investor's opinion of the gross return offered by a risky investment. The equations below define the power utility function in terms of wealth, *W*, and gross returns, *R*. The risk aversion parameter, γ, must be positive for risk averse investors. The limiting case of γ=1 is log-utility. The mathematics that follows shows how the starting wealth is unimportant to utility.

$$U = \begin{cases} \ln(W) & \gamma = 1 \\ \dfrac{W^{1-\gamma}}{1-\gamma} & \gamma \neq 1 \end{cases}$$

$$W_1 = W_0 R$$

$$U_1 = \frac{W_1^{1-\gamma}}{1-\gamma} = \frac{[W_0 R]^{1-\gamma}}{1-\gamma}$$

$$\frac{U_1}{U_0} = \frac{\dfrac{[W_0 R]^{1-\gamma}}{1-\gamma}}{\dfrac{W_0^{1-\gamma}}{1-\gamma}} = R^{1-\gamma}$$

$$U = \begin{cases} \ln(R) & \gamma = 1 \\ \dfrac{R^{1-\gamma}}{1-\gamma} & \gamma \neq 1 \end{cases}$$

The exponential utility function features constant absolute risk aversion, CARA. This means that the increase in utility is the same for a given dollar return, regardless of the starting wealth. (If an investor's wealth increases by \$10, his utility increases by the same amount whether he started with \$100 or \$100,000.) This makes the starting wealth important in the derivation of the investor's opinion of the gross return offered by a risky investment. The exponential utility function is below, in terms of wealth and still using γ as the risk aversion parameter, which again must be positive for risk averse investors. The mathematics that follows shows how the starting wealth is important to utility for gross returns.

$$U = -e^{-\gamma W}$$

$$\begin{aligned} W_1 &= W_0 R \\ U_1 &= -e^{-\gamma W_1} = -e^{-\gamma W_0 R} \\ \frac{U_1}{U_0} &= \frac{-e^{-\gamma W_0 R}}{-e^{-\gamma W_0}} = e^{-\gamma W[R-1]} \end{aligned}$$

The papers that use the exponential utility function set the starting wealth to one to cause it to drop out for the ending period wealth as well.

$$U_0 = -e^{-\gamma}$$

$$U_1 = -e^{-\gamma R}$$

Constant absolute and relative risk aversion are often demonstrated by taking the first and second derivatives of the utility function for the Arrow-Pratt measures. The equations for absolute risk aversion, *ARA*, and relative risk aversion, *RRA*, are shown in the equations below.

$$ARA = -\frac{U''(R)}{U'(R)}$$

$$RRA = -\frac{U''(R)}{U'(R)} R$$

The Arrow-Pratt measures for the power utility function and exponential utility functions (setting starting wealth to one) are derived and shown below.

$$U = \frac{R^{1-\gamma}}{1-\gamma}$$
$$U' = R^{-\gamma}$$
$$U'' = -\gamma R^{-\gamma-1}$$

$$ARA = -\frac{-\gamma R^{-\gamma-1}}{R^{-\gamma}} = \gamma R^{-1}$$

$$RRA = -\frac{-\gamma R^{-\gamma-1}}{R^{-\gamma}} R = \gamma$$

$$U = -e^{-\gamma R}$$
$$U' = \gamma e^{-\gamma R}$$
$$U'' = -\gamma^2 e^{-\gamma R}$$

$$ARA = -\frac{-\gamma^2 e^{-\gamma R}}{\gamma e^{-\gamma R}} = \gamma$$

$$RRA = -\frac{-\gamma R^{-\gamma-1}}{R^{-\gamma}} R = \gamma R$$

It is important to note that these results hold for the power utility function even if the risk aversion parameter is set to unity and the limiting case is necessary. Of even more importance, both utility functions have well-behaved first and second derivatives. Economic theory requires that utility increases with returns, which requires that the first derivative is positive. Further work shows that utility should increase at a decreasing rate, which requires that the second derivative is negative. However, evidence and behavioral finance theory do not support utility functions that display increasing relative risk aversion, IRRA. IRRA has investors become relatively more risk averse as wealth increase and is a feature of the exponential utility function. The CRRA feature of the power utility function is more desirable and useful for normal equivalents.

Typically utility functions are shown graphically by plotting the utility for different returns. While this is helpful for showing the effect of utility in absolute terms for the utility functions it does not allow for a comparison of how utility is affected between different utility functions. Perhaps the best way to show the relative impact on utility from net returns to wealth for both the power utility function and exponential utility function is with the return necessary to compensate for a fifty percent chance at a loss of a given percent. This is derived below for a loss of ten percent and shown for losses of one to twelve percent graphically for risk aversion parameters of one, three, and five in Figure 5-1.

The derivations begin with Equation 5-1 to solve for the different utility functions, in net returns, *r*.

$$0.5 \times \mathrm{U}(-10\%) + 0.5 \times \mathrm{U}(r) = \mathrm{U}(0\%) \tag{5-1}$$

First, the power utility function with the risk aversion parameter set to one which approaches the log-utility function in the limit:

$$U = \ln(1 + r)$$

Step 1-1: Plug the utility function into the derivation equation.

$$0.5 \times \ln(1 - 0.10) + 0.5 \times \ln(1 + r) = \ln(1.00)$$

Step 1-2: Multiply both sides by two and note that the log of 1 is 0.

$$\ln(0.90) + \ln(1 + r) = 0$$

Step 1-3: Rearrange terms and raise both sides to an exponent.

$$e^{\ln(1+r)} = e^{-\ln(0.90)}$$

Step 1-4: Note that the exponent of a log of a variable is equal to the variable and that the log of a variable is equal to the inverse of the log of the reciprocal of the variable.

$$1 + r = \frac{1}{0.90}$$

Step 1-5: Solve for r.

$$r = \frac{1}{0.90} - 1 \approx 11.11\%$$

Second, the power utility function with a positive risk aversion parameter that is not equal to one:

$$U = \frac{[1 + r]^{1-\gamma}}{1 - \gamma}$$

Step 2-1: Plug the utility function into the derivation equation.

$$0.5 \times \frac{[1 - 0.10]^{1-\gamma}}{1 - \gamma} + 0.5 \times \frac{[1 + r]^{1-\gamma}}{1 - \gamma} = \frac{[1]^{1-\gamma}}{1 - \gamma}$$

Step 2-2: Multiply both sides by two times 1-γ and note that one to any power is one.

$$0.90^{1-\gamma} + [1 + r]^{1-\gamma} = 2$$

Step 2-3: Rearrange terms and raise both sides to the reciprocal of 1- γ.

$$1 + r = \left[2 - 0.90^{1-\gamma}\right]^{\frac{1}{1-\gamma}}$$

Step 2-4: Subtract one from both sides, now given the risk aversion parameter the net return can be solved.

$$r = \left[2 - 0.90^{1-\gamma}\right]^{\frac{1}{1-\gamma}} - 1$$

Third, the exponential utility function with a positive risk aversion parameter:

$$U = -e^{-\gamma R}$$

Step 3-1: Plug the utility function into the derivation equation.

$$0.5 \times -e^{-\gamma[1-0.10]} + 0.5 \times -e^{-\gamma[1+r]} = -e^{-\gamma}$$

Step 3-2: Multiply both sides by two and rearrange terms.

$$e^{-\gamma[1+r]} = 2e^{-\gamma} - e^{-0.90\gamma}$$

Step 3-3: Take the log of both sides and note that the log of an exponent of a variable is equal to the variable.

$$-\gamma[1+r] = \ln\left(2e^{-\gamma} - e^{-0.90\gamma}\right)$$

Step 3-4: Divide both sides by negative γ and subtract one, now given the risk aversion parameter, the net return can be solved.

$$r = -\frac{\ln\left(2e^{-\gamma} - e^{-0.90\gamma}\right)}{\gamma} - 1$$

The derivations above were done with a fifty percent chance at a loss of 10% and a certain alternative of a net return of zero percent. The net return necessary for the remaining fifty percent in order for the utility of the risky investment and the certain return is determined by *r*. This could be done for any loss, probability of loss, and certain return.

Figure 5-1 plots different losses, ranging from 12% to 1%, on the x-axis with the necessary gains, ranging from 1% to 35%, for both the power and exponential utility functions with the risk aversion parameter set to 1, 3, and 5. This graph shows the effect of the risk aversion parameter and utility function on the impact of net returns on utility. As the risk aversion parameter increases, risk aversion increases. For the power utility function with the risk aversion parameter set to one a fifty percent chance of a 10% loss and a fifty percent chance at an 11.11% gain will yield an investor the same utility as a certain net return of 0%. Under the same possible loss and certain return, a fifty percent chance at a 20.40% return is necessary to provide the same utility when the risk aversion parameter is set to five. Relative risk aversion is shown by the compensating potential gain.

Figure 5-1 also shows that the curvature of the utility functions. A loss that is twice is large will require a gain that is more than twice as large in order to compensate the utility. This curvature increases as risk aversion increases. Also, the exponential utility function shows slightly more relative risk aversion than the power utility function. The difference between the power utility function and the exponential utility function are quite small. In fact, the difference is not noticeable until the risk aversion and loss are relatively large for this graph. Since the differences are minute, the normal equivalents would be redundant.

Alternative Distributions

The normal distribution is favored for this dissertation, but alternate distributions would satisfy the necessary conditions. The second best choice is the continuous uniform distribution because. The uniform distribution is very familiar to academics and is the simplest to draw and describe. The uniform distribution gives every observation in the domain an equal probability – making the probability density function a horizontal line and the cumulative density function a diagonal line starting at zero for the lowest possible observation and ending at one for the highest possible observation. Below are the density functions for the continuous uniform distribution; following standard practice for this distribution the lower and upper limits are *a* and *b* and the random variable is *R*, gross returns.

$$f(R) = \begin{cases} \dfrac{1}{b-a} & a \le R \le b \\ 0 & \text{otherwise} \end{cases}$$

$$F(R) = \begin{cases} 0 & R < a \\ \dfrac{R-a}{b-a} & a \le R \le b \\ 1 & R > b \end{cases}$$

The uniform distribution is not only a good alternative because it has recognizable features and is easy to interpret but is also fully defined by the first two moments. In order for a distribution to create equivalents, the first two moments must be sufficient statistics. It also satisfies the assumptions of the Sharpe ratio, namely that the returns have a uniform distribution the utility function can be any with constant relative risk aversion and does not need to be mean-variance. However, it is not generally a good approximation to the return distribution found in the market. The mean and standard deviation for the uniform distribution are derived below.

$$\mu = \int_a^b \frac{R}{b-a} dR$$
$$\mu = \frac{b^2}{2[b-a]} - \frac{a^2}{2[b-a]}$$

$$\mu = \frac{b+a}{2}$$

$$\mathrm{E}(R^2) = \int_a^b \frac{R^2}{b-a} dR$$
$$\mathrm{E}(R^2) = \frac{b^3}{3[b-a]} - \frac{a^3}{3[b-a]}$$
$$\mathrm{E}(R^2) = \frac{b^2 + ab + a^2}{3}$$

$$\sigma^2 = \mathrm{E}(R^2) - \mu^2$$
$$\sigma^2 = \frac{b^2 + ab + a^2}{3} - \frac{[b+a]^2}{4}$$
$$\sigma^2 = \frac{4b^2 + 4ab + 4a^2 - 3a^2 - 6ab - 3b^2}{12}$$
$$\sigma^2 = \frac{b^2 - 2ab + a^2}{12}$$

$$\sigma = \frac{b-a}{2\sqrt{3}}$$

$$\mu = \frac{b+a}{2} \Rightarrow b = 2\mu - a$$
$$\sigma = \frac{b-a}{2\sqrt{3}} = \frac{2\mu - 2a}{2\sqrt{3}} \Rightarrow a = \mu - \sqrt{3}\sigma$$
$$\therefore b = \mu + \sqrt{3}\sigma$$

$$f(R) = \begin{cases} \frac{1}{2\sqrt{3}\sigma} & \mu - \sqrt{3}\sigma \le R \le \mu + \sqrt{3}\sigma \\ 0 & \text{otherwise} \end{cases}$$

A desirable property of the uniform distribution emerges in that censoring would not be necessary because the uniform distribution has a lower bound already that should be restrictive enough that the utility function will be well defined over the entire range (standard deviations larger than 0.50 being very unlikely). The calculation ease of the probability density function creates yet another desirable feature – the expected utility is integrable. The expected utility of a uniform return distribution with the power utility function is derived below.

$$\mathrm{E}(\mathrm{U}(R)) = \int f(R) \times U(R) \times dR$$
$$\mathrm{E}(\mathrm{U}(R)) = \int_{\mu-\sqrt{3}\sigma}^{\mu+\sqrt{3}\sigma} \frac{1}{2\sqrt{3}\sigma} \times \frac{R^{1-\gamma}}{1-\gamma} \times dR$$

$$\mathrm{E}(\mathrm{U}(R)) = \frac{\left[\mu + \sqrt{3}\sigma\right]^{-\gamma} - \left[\mu - \sqrt{3}\sigma\right]^{-\gamma}}{2\sqrt{3}\sigma}$$

The uniform equivalent-mean is created by replacing the average utility from the sample for the expected utility and the sample standard deviation for σ, and solving for μ. Similarly, the uniform equivalent-standard deviation is created by replacing the average utility from the sample for the expected utility and the sample mean for μ, and solving for σ. As an example, the easiest case when the risk aversion parameter is one is derived completely.

$$\overline{U} = \frac{\left[UE_{\mu} + \sqrt{3}s\right]^{-\gamma} - \left[UE_{\mu} - \sqrt{3}s\right]^{-\gamma}}{2\sqrt{3}s}$$

$$\gamma = 1$$

$$\Rightarrow \overline{U} = \frac{\left[UE_\mu + \sqrt{3}s\right]^{-1} - \left[UE_\mu - \sqrt{3}s\right]^{-1}}{2\sqrt{3}s}$$

$$2\sqrt{3}s\overline{U} = \frac{UE_\mu - \sqrt{3}s - UE_\mu - \sqrt{3}s}{\left[UE_\mu + \sqrt{3}s\right]\left[UE_\mu - \sqrt{3}s\right]}$$

$$2\sqrt{3}s\overline{U} \times \left[UE_\mu{}^2 - 3s^2\right] = -2\sqrt{3}s$$

$$UE_\mu = \sqrt{3s^2 - \overline{U}^{-1}}$$

$$\overline{U} = \frac{\left[\mu + \sqrt{3}UE_\sigma\right]^{-\gamma} - \left[\mu - \sqrt{3}UE_\sigma\right]^{-\gamma}}{2\sqrt{3}UE_\sigma}$$

$$\gamma = 1$$

$$\Rightarrow \overline{U} = \frac{\left[\overline{R} + \sqrt{3}UE_\sigma\right]^{-1} - \left[\overline{R} - \sqrt{3}UE_\sigma\right]^{-1}}{2\sqrt{3}UE_\sigma}$$

$$2\sqrt{3}UE_\sigma\overline{U} = \frac{\overline{R} - \sqrt{3}UE_\sigma - \overline{R} - \sqrt{3}UE_\sigma}{\left[\overline{R} + \sqrt{3}UE_\sigma\right]\left[\overline{R} - \sqrt{3}UE_\sigma\right]}$$

$$2\sqrt{3}UE_\sigma\overline{U} \times \left[\overline{R}^2 - 3UE_\sigma{}^2\right] = -2\sqrt{3}UE_\sigma$$

$$UE_\sigma = \sqrt{\frac{\overline{R}^2\overline{U} + 1}{3\overline{U}}}$$

The uniform distribution is as effective as the normal distribution in that it is recognizable and sufficient for the assumptions of the Sharpe ratio. It can be more practical because its probability density function is integrable. However, the commonality and better approximation to historical returns of the normal distribution outweigh the easier mathematics of the uniform.

Several other distributions have been assumed for returns, most notably the log-normal distribution (Equation 5-2) that would fail to satisfy the requirements for a standard. The log-normal distribution has been used for the returns for common stocks, assuming that stock prices are normally distributed stock returns are log-normal. Besides the theoretical background, it is bounded from below at 0, and would need no censoring. Also, the mean and variance are sufficient statistics for the log-normal distribution, which is important for computing equivalents.

However, the skewness and kurtosis are not constant and depend on the mean and variance (Evans, Hastings, and Peacock (2000)), which makes it inappropriate for a standard and not satisfactory for the Sharpe ratio because it would require a quadratic utility function.

$$f(R)=\frac{e^{-\frac{[\ln(R)-\mu]^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma R} \tag{5-2}$$

Normal equivalents are created instead of uniform equivalents because the loss due to censoring and non-integrability is not as important as the gain due to the common assumption of normal returns. The normal distribution is used as a standard over other possible return distributions because the mean and standard deviation are sufficient statistics for the distribution, it satisfies the assumptions of the Sharpe ratio, it is symmetric, and its higher moments are not dependent upon the mean and variance.

Data Biases

The difference in the characteristics of surviving and disappearing funds has been noted to be a problem for both mutual funds and hedge funds. The survivorship bias is the difference between the mean return of a sample with just surviving funds and a sample with all funds. Literature cited by Liang (2000) find the survivorship bias to be about one percent per year for mutual funds, and hedge funds can be expected to have a greater bias because they can choose higher risk strategies, face less public scrutiny, and are a fledging alternative investment vehicle set. Liang's work estimates the survivorship bias at a little more than two percent per year. Brown, Goetzmann, and Ibbotson (1999) estimate the survivorship bias to be three percent per year and Ackermann, McEnally, and Ravenscraft (1999) estimate the survivorship bias to be nearly non existent.

The spread in the survivorship bias results can be explained by the differences in the samples. Brown, Goetzmann, and Ibbotson (1999) use a yearly directory of offshore hedge funds from 1989 to 1995. Ackermann, McEnally, and Ravenscraft (1999) combine data from Hedge Fund Research, HFR, and the Zurich Hedge Fund Universe (then called Managed Account Reports, Inc.) from 1988 to 1995. Liang (2000) combines HFR and TASS Management Limited, TASS. Liang contends that the different databases have different rates of survivorship, which explains the variety in the results (HFR data have fewer dissolved funds). Other possible sources of the spread and inconsistency might be the sample – certain years may have greater dropout rates, and the timing – the length of time a fund must remain in the sample to be considered a survivor.

Six sources of bias are outlined in Ackermann, McEnally, and Ravenscraft (1999), to which I add a seventh that may arise from data collected from hedgefund.net.

1) Survivor – The bias that occurs from considering only the funds that are alive and active in the database at the end of the sample period.
2) Termination – The bias that occurs from not considering funds that have been terminated during the sample period.
3) Self-Selection – The bias that occurs from not considering funds that have elected to no longer report to the database.
4) Liquidation – The bias that occurs from not considering the last stages of a fund before it disappears.
5) Backfilling – The bias that occurs from accepting a fund's historical returns before it began reporting to the database.

6) Multi-Period Sampling – The bias that occurs from not considering the returns to funds that entered the database during the sample period.

7) Late Reporting – The bias that occurs from not considering the returns to funds that report their returns in an untimely fashion.

The data from the hedgefund.net database is subject to all of these biases. Some are particularly strong because of the ten year sample period necessary to estimate skewness and kurtosis. An unknowable, but presumably large, number of funds have dropped out or been added to the database during these ten years. Fund return data is updated to the service, so returns for firms that have dropped out are lost and returns are not coded to indicate when they were reported, so the backfilling is also unknown. Lastly, 373 unique funds were identified to have begun by January 1996 for the sample period and reported complete data through December 2005. Table 5-1 reports the descriptive statistics of another 15 funds had full reports through November 2005, but had not reported December's return as of the first weekend of March 2006. Eleven of these funds had reported their returns by the first weekend of April and four had not. The descriptive statistics of these funds have similar characteristics of the 373 funds in the analysis. Although no sizeable difference appears from the funds that reported late, it would not be a leap to consider that late reporting funds might have had the worst returns (reluctant to report) or most unusual returns (difficulty in valuing illiquid investments).

Besides the difference in the mean, Ackermann, McEnally, and Ravenscraft (1999) and Kouwenberg (2003) find that the disappearing funds are more volatile – use greater leverage, have a higher standard deviation, and a worse minimum monthly return. Brown, Goetzmann, and Ibbotson (1999) find that few funds and fund managers last three years in the directory and their further work (Brown, Goetzmann, and Park (2001)) find that few fund managers that disappear from the directory ever returned.

Hedge fund features may also play a role in the decision to terminate that does not occur in mutual funds. Hedge funds are allowed to have an asymmetric fee schedule, the most common is a twenty percent fee for profits above a certain level. If a fund fails to achieve that level for several time periods, the managers must catch-up in order to receive the bonus and the level becomes a high watermark. A hedge fund manager that feels the probability to reach the high watermark is small is more apt to dissolve the fund.

There are many features of hedge funds that can create a biased sample, not only in the means, but also in the standard deviation and higher moments. It is believed that due to the necessary time period for the data that the results here are atypical of hedge funds in general, which should be expected to have a lower mean and greater risk, perhaps including utility detracting higher moments.

Hedge Funds in a Portfolio

As a stand-alone investment, this sample of hedge funds seem favorable and though they reject normality statistically, few of them have enough non-normality to impact investors' utility and analysts' performance measurement. In this section I seek to further the understanding of their risks by combining them with the market portfolio and the risk free rate. Hedge funds have a wide variety of styles and putting each hedge fund into the same portfolio setting may not be fair. Therefore, a portfolio that maximizes utility is created. This allows each fund to have a different effect on the allocation into the market and risk-free rate to reflect their unique characteristics.

The portfolios consist of a stock index and one-month T-bills, chosen so as to maximize utility. The rates of returns for the stock index and one-month T-bills comes from Kenneth French's website: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. The market rate of return is CRSP's monthly value weighted returns and the one-month T-bill rates are from Ibbotson and Associates, Inc. Additionally, the returns are in real terms, using the Consumer Price Index (available at the Bureau of Labor Statistics website: http://www.bls.gov/cpi/home.htm) to adjust for inflation. The returns of the 120 months from 1996 to 2005 are assumed to be equally likely occurrences for the next month and adjusting for inflation so as to capture real utility. (See Table 5-2 for the descriptive statistics of these three variables.) The power utility function with a risk aversion parameter of three is selected to best represent the sensitivities to risk and return for potential investors into hedge funds. The allocation between the stock index and the one month T-bills is allowed only to range from zero to one – no borrowing and no short selling. With these conditions, the appropriate weighting decision between stocks and T-bills is demonstrated below.

It begins by maximizing the expected utility of the real gross returns to the portfolio, R, given the risk aversion parameter of γ. The real gross return is determined by the allocation to the stock market, w; and the net returns to the stock market, r_m; one-month T-bill, r_f; and inflation rate, p.

$$\max_{w} \mathrm{E}(\mathrm{U}(R)) = \frac{R^{1-\gamma}}{1-\gamma}$$
$$st. 0 \le w \le 1$$
$$R = \frac{1 + wr_m + [1-w]r_f}{1+p}$$

Since each of the 120 outcomes is equally likely, maximizing the expected value is the same as maximizing the average. Recognizing the denominator as a scalar it can be relocated.

$$\max_{w}\left(\frac{\sum_{t=1}^{T}\left[\frac{1+wr_{mt}+[1-w]r_{ft}}{1+p_t}\right]^{1-\gamma}}{[1-\gamma]T}\right)$$

The optimal weights for risk aversion parameters from one to five was found to the nearest hundredth of a percent. These optimal weights and the descriptive statistics for these portfolios are reported in Table 5-3 and to show the behavior of the functions, the certainty equivalents are plotted in Figure 5-2.

The optimal weight for the stock market index meets the no borrowing boundary condition for the less risk averse parameters of one and two. From there it falls relatively quickly to 81.45% for a risk aversion level of three and all the way to 49.06% for a risk aversion level of five. This causes the mean and standard deviation to fall as well. The non-normality decreases as the portfolio relies less and less on the stock market, although always rejecting normality at the five percent level using the Jarque-Bera test (Equation 4-1). The shape of the distribution improves from a utility standpoint as the skewness increases and the kurtosis decreases, however the change in utility is very small. Increasing the risk aversion parameter decreases the certainty equivalent because more risk averse investors would accept a smaller guaranteed return than their less risk averse counterparts. The mean decreases faster than the certainty equivalent from three to five, so the actual risk premium decreases. (Meeting the boundary condition at less risk averse levels make comparisons difficult at one and two.)

With these standards in place, I seek to determine if adding a specified amount into a hedge fund increases utility and if it does so by increasing the return or decreasing the risk. Further, I seek to determine what affect its inclusion has on the optimal allocation into the stock index and the one-month T-bills. I do this by using a sample of 373 hedge funds with complete monthly data for this time period available on the website of hedgefund.net. (See Table 4-1 for descriptive statistics of these funds.)

The additional work to make optimal portfolios to determine if the hedge funds add value appears necessary from the effects of their inclusion (Table 5-4) with the stock index and one month T-bills. Beginning with a mix of 81.45% in stocks and 18.55% in T-bills, there was a wide range in the appropriate allocations with a 10% share in a hedge fund. The resulting mix ranges from 66% stocks and 24% T-bills to 90% stocks and 0% T-bills. (There are still restraints placed on the system by assuming that the hedge funds are designed for investors across all risk aversion levels and with the same 10% allocation.)

The most interesting reallocation occurs for funds with a short bias. Six of these seven funds counterbalance the stock market such that the investor would select a portfolio without T-bills. The fifteen market neutral equity funds have a median allocation of 81% stocks and 9% T-bills and the sixteen long only funds have a median allocation of 74% stocks and 16% T-bills. Market neutral equity funds appear to substitute for the T-bills and long only funds appear to substitute for the stock market.

The fifty-five funds seeking arbitrage opportunities have similar reallocations as the market neutral equity funds, substituting for T-bills. A category which seems to contain the most variety is the value hedge funds. Some strongly decrease the demand for stocks and some strongly decrease the demand for T-bills. Hedge funds with a strategy seeking distressed securities also had a wide range, but appear more synchronous with the market because none of them increase the demand for stocks.

These funds are nearly all valuable additions to a portfolio (Table 5-5). Fully 96% of the funds increase utility and the average gain in utility is equivalent to a 53 basis point increase in the certainty equivalent. This increased utility comes about with increased expected returns, up 55 basis points, which compensates for the increased risk; the risk premium is up 1.5 basis points (the standard deviation increased 12 basis points). The percentage of funds that increase utility is consistent for allocations as small as 5% and as large as 25% into a particular hedge fund. However, their average improvement to utility increases with a larger weight. By the largest allocation, the average certainty equivalent is 125 basis points higher, including a 4.5 basis point improvement in the risk premium (the standard deviation is 46 basis points smaller). Less risk averse investors would find fewer attractive funds in this list because some of their returns are not as great as the market. However, more risk averse investors do not find a larger amount of funds more attractive than do those of the standard risk aversion. These investors start with an allocation with less risk and less return and have a nearly identical change in the risk and return by including a hedge fund.

All types of funds improve utility. The largest improvement in the certainty equivalent comes from those with a short bias (95 basis points). These funds increase the expected return by an average of 93 basis points and reduce risk by an average of 2 basis points, the only group to reduce risk. The funds that invest in other hedge funds tend to underperform their peers, with none of the three groups able to reach the average increase in the certainty equivalent. On a clear comparison however, with Fund of Funds – Market Neutral versus Market Neutral Equity funds, they are able to increase the mean by 9 more basis points and increase risk by 1 less basis point, but with a combined total of 25 funds, this very small difference is also not statistically significant.

Hedge funds that remain in a database (see the previous section) are valuable additions to a portfolio of stocks and T-bills; generally with increased returns with little change in the risk.

The funds have a variety of effects on the optimal allocations of stocks and T-bills and short bias hedge funds can even eliminate the need for T-bills.

CHAPTER SIX: CONCLUSIONS

This dissertation introduces normal equivalents to allow funds with different higher moments to be compared together on an equal footing. This is done by forcing a sample return distribution's utility onto a normal return distribution and determining the equivalent mean or standard deviation. Since normal equivalents standardize the distribution, it also identifies both the impact the distribution has on utility of an investor and the degree to which the standard deviation accurately portrays risk.

The non-normality premium emerges to identify the cost of the non-normal distribution, just as the risk premium identifies the cost of the risk. Unlike the risk premium, the non-normality premium may be negative as well as positive because a distribution may be utility enhancing or utility detracting, whereas risk is always utility detracting for risk averse investors. The utility adjustment to the standard deviation is the percentage by which the standard deviation must be adjusted in order to represent the risk-to-utility instead of the risk-to-return. As the Sharpe ratio depends on the standard deviation to accurately portray risk-to-utility, the utility adjustment to the standard deviation shows the amount the Sharpe ratio is over or under estimated by using the standard deviation.

The statistics generated in this paper cannot be determined to be significant or insignificant because they show the importance to investors utility and perception of risk. The adjustment to the standard deviation necessary for a statistically significant change in the Sharpe ratio is stronger than the basis used here. Therefore assumptions are made for analysis and meant to correspond to perceptions where the performance evaluation tool has a large variance that may make it inefficient yet it remains the most popular. First, that a ten basis point non-normality premium determined on a monthly basis is economically meaningful to investors.

Second, that a Sharpe ratio that is off by ten percent makes its use for evaluating performance doubtful. Simulations suggest that the likelihood of a normal distribution producing such non-normal returns to be very small. Another critical assumption is that the power utility function with a risk aversion level of three is the most appropriate function which appears to be on the lower end of the realistic spectrum (although easily inside the popular range).

With these purposes and assumptions, I use the expected value of Taylor series expansions around the mean to determine the relative importance of skewness and kurtosis relative to the mean and standard deviation. I then determine levels of skewness and kurtosis that will cause concern to investors and invalidate the Sharpe ratio. Further exploration is done by using a non-central *t* distribution to create returns that are not normal with a range of standard deviations, skewnesses, and kurtoses. Lastly, the normal equivalents are applied to a sample of hedge funds because they are known for their statistical rejection of normality. The hedge funds are evaluated as stand alone investments and as incorporated into a portfolio of stocks and one-month T-bills.

From this work I conclude that the non-normality of a distribution does not have a meaningful impact on utility unless the skewness is less than -2 or greater than +2 for common standard deviations for this level of risk aversion. The kurtosis is less important to compute because its impact is smaller and there is such a strong correlation between the kurtosis and the absolute value of the skewness that computing only one is enough. The sample of hedge funds routinely rejected normality but rarely was the non-normality enough to have a meaningful impact on utility. This seemed to be because the standard deviations were often so small as to make all the risks small. A portfolio of stocks and the risk-free rate that also includes a small position in a hedge fund will see no noticeable impact on its higher moments and have very reasonable estimates as to the Sharpe ratio in a portfolio setting but less so for a stand alone investment because of funds with under and over estimated ratios. Rankings based on the Sharpe ratio were almost the same whether it was the standard formula or the normal equivalent version.

Normal equivalents were designed to test the economic relevance of the non-normality of hedge funds. Other financial investments such as options and currency futures might also have non-normality and inconsistency in their higher moments and benefit from the normal equivalents. And lastly, it is hoped that the normal equivalents derived, explored, and applied here will spawn other equivalents to improve understanding and check robustness when differences are assumed away.

BIBLIOGRAPHY

Ackermann, Carl, Richard McEnally, and David Ravenscraft. 1999. “The performance of hedge funds: Risk, return, and incentives.” *Journal of Finance*, vol. 54, no. 3 (June): 833-874.

Aït-Sahalia, Yacine, Jonathan Parker, and Motohiro Yogo. 2004. "Luxury goods and the equity premium." *Journal of Finance*, vol. 59, no. 6 (December): 2959-3003.

Amin, Gaurav and Harry Kat. 2003. “Hedge fund performance 1990-2000: Do the “Money Machines” really add value?” *Journal of Financial and Quantitative Analysis*, vol. 38, no. 2 (June): 251-274.

Bookstaber, Richard. 2003. “Hedge fund existential.” *Financial Analysts Journal*, vol. 59, no. 5 (September/October): 19-23.

Brown, Stephen, William Goetzmann, and Roger Ibbotson. 1999. “Offshore hedge funds: Survival and performance, 1985-95.” *Journal of Business*, vol. 72, no. 1: 91-117.

Brown, Stephen, William Goetzmann, and James Park. 2001. “Careers and survival: Competition and risk in the hedge fund and CTA industry. *Journal of Finance*, vol. 56, no. 5 (October): 1869-1876.

Capocci, Daniel and Georges Hübner. 2004. “Analysis of hedge fund performance.” *Journal of Empirical Finance*, vol. 11: 55-89.

Evans, Merran, Nicholas Hastings, and Brian Peacock. 2000. *Statistical Distributions*, third edition. New York: John Wiley & Sons, Inc.

Fung, William and David Hsieh. 1999a. “A primer on hedge funds.” *Journal of Empirical Finance*, vol. 6: 309-331.

Fung, William and David Hsieh. 1999b. “Is mean-variance analysis applicable to hedge funds?” *Economic Letters*, vol. 62: 53-58.

Harvey, Campbell and Akhtar Siddique. 1999. “Autoregressive Conditional Skewness.” *Journal of Financial and Quantitative Analysis*, vol. 34, no. 4 (December): 465-487.

Harvey, Campbell and Akhtar Siddique. 2000. “Conditional Skewness in Asset Pricing Tests.” *Journal of Finance*, vol. 55, no. 3 (June): 1263-1295.

Hlawitschka, Walter. 1994. "The empirical nature of Taylor-series approximations to expected utility." *American Economic Review*, vol. 84, no. 3 (June): 713-719.

Hogben, D., R.S. Pinkham, and M. B. Wilk (1961). “The moments of the non-central t-distribution.” *Biometrika*, vol. 48, no. 3/4 (December): 465-468.

Jarque, Carlos and Anil Bera. 1987. “A test for normality of observations and regression residuals.” *International Statistics Review*, vol. 55: 163-172.

Kouwenberg, Roy. 2003. “Do hedge funds add value to a passive portfolio? Correcting for non-normal returns and disappearing funds.” *Journal of Asset Management*, vol. 3, no. 4 (March): 361-382.

Levy, Haim and Harry Markowitz. 1979. "Approximating expected utility by a function of mean and variance." *American Economic Review*, vol. 69, no. 3 (June): 308-317.

Liang, Bing. 2000. “Hedge funds: The living and the dead.” *Journal of Financial and Quantitative Analysis*, vol. 35, no. 3 (September): 309-326.

Liang, Bing. 2004. “On the Performance of Alternative Investments: CTAs, Hedge Funds, and Funds-of-Funds.” *Journal of Investment Management*, vol. 2 (4^{th} Quarter): 76-93.

Lo, Andrew. 2002. “The statistics of Sharpe ratios.” *Financial Analysts Journal*, vol. 58, no. 4 (July/August): 36-52.

Mehra, Rajnish and Prescott, Edward. 1985. "The equity premium: a puzzle." *Journal of Monetary Economics*, vol. 15, no. 2 (March): 145-161.

Mittelhammer, Ron. 1996. *Mathematical Statistics For Economics and Business*. New York: Springer-Verlag New York Inc.

Osband, Kent. 2002. *Iceberg Risk: An Adventure in Portfolio Theory*. New York: Texere LLC.

Pedersen, Christian and Stephen Satchell. 2000. "Small Sample Analysis of Performance Measures in the Asymmetric Response Model." *Journal of Financial Quantitative Analysis*, vol. 35, no. 3 (September) 425-450.

Scott, Robert and Philip Horvath. 1980. "On the direction of preference for moments of higher order than the variance." *Journal of Finance*, vol. 35, no. 4 (September) 915-919.

Sharma, Milind. 2004. "AIRAP – Alternative RAPMs for alternative investments." *Journal of Investment Management*, vol. 2, no. 4 (Fourth Quarter).

Sharpe, William. 1966. "Mutual fund performance." *Journal of Business*, vol. 39, no. 1 (January): 119-138.

Sharpe, William. 1994. "The Sharpe ratio." *Journal of Portfolio Management*, vol. 21, no. 1 (Fall): 49-58.

TABLES AND FIGURES

Table 3-1: Marginal Utilities with Respect to the Mean, Standard Deviation, Skewness and Kurtosis

The marginal utility to skewness and kurtosis relative to the marginal utility to the mean and standard deviation by the standard deviation and risk aversion. The first column for each standard deviation is for the marginal utility for the mean and the second column is for the marginal utility for the standard deviation. The first row for every risk aversion is the marginal utility for the skewness and the second row is for the marginal utility for the kurtosis. The base level of the mean is 1% and it is 0 and 3 for the skewness and kurtosis, respectively. All are in basis points. Thus, for a risk aversion coefficient of three and a standard deviation of 5% an increase of 1 in the skewness has the same impact on utility as an increase of 2.41 basis points in the mean or a decrease of 16.11 basis points in the standard deviation.

		Standard Deviation							
	Marginal Utility of	**3%**		**5%**		**7%**		**9%**	
		Marginal Utility of		Marginal Utility of		Marginal Utility of		Marginal Utility of	
Risk Aversion		**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**
1	η_1	0.09	-2.96	0.41	-8.19	1.12	-15.94	2.36	-26.11
	η_2	0.00	0.07	-0.02	0.30	-0.06	0.83	-0.16	1.75
2	η_1	0.26	-4.43	1.22	-12.20	3.31	-23.58	6.97	-38.28
	η_2	-0.01	0.13	-0.06	0.60	-0.23	1.63	-0.62	3.41
3	η_1	0.53	-5.89	2.41	-16.11	6.53	-30.86	13.61	-49.53
	η_2	-0.02	0.22	-0.15	1.00	-0.57	2.67	-1.52	5.52
4	η_1	0.87	-7.33	3.98	-19.90	10.67	-37.71	21.93	-59.72
	η_2	-0.04	0.33	-0.30	1.48	-1.11	3.92	-2.93	7.98
5	η_1	1.31	-8.75	5.90	-23.54	15.61	-44.07	31.56	-68.74
	η_2	-0.07	0.45	-0.51	2.04	-1.89	5.35	-4.92	10.72

Table 3-2: Higher Moments of the Non-Central *t* Distribution

The skewness and kurtosis are reported for the non-central *t* distribution for various non-centrality parameters and degrees of freedom. The skewness and kurtosis are represented as η_1 and η_2, respectively. The non-centrality and degrees of freedom parameters are represented as δ and ν, respectively. The shaded statistics are those that come from the standard in this paper, δ=-1 and ν=5.

	δ=-2		**δ=-1**		**δ=0**		**δ=1**		**δ=2**	
	η_1	η_2	η_1	η_2	η_1	η_2	η_1	η_2	η_1	η_2
ν= 5	-2.08	20.84	-1.27	13.32	0.00	9.00	1.27	13.32	2.08	20.84
ν=10	-0.72	4.83	-0.40	4.25	0.00	4.00	0.40	4.25	0.72	4.83
ν=15	-0.45	3.84	-0.24	3.63	0.00	3.55	0.24	3.63	0.45	3.84
ν=20	-0.33	3.53	-0.17	3.42	0.00	3.38	0.17	3.42	0.33	3.53
ν=25	-0.26	3.38	-0.13	3.31	0.00	3.29	0.13	3.31	0.26	3.38

Table 3-3: Censoring the Non-Central *t* Distribution

The odds of censoring the non-central *t* distribution when censoring at -0.99 when the mean is 0.01. For the first column the odds are for a non-centrality parameter of -1 and 5 degrees of freedom for various standard deviations. For the second column the odds are for a standard deviation of 5% and 5 degrees of freedom for various non-centrality parameters. For the third column the odds are for a standard deviation of 5% and non-centrality parameter of -1 for various degrees of freedoms.

δ=-1 and ν=5		σ=5% and ν=5		σ=5% and δ=-1	
σ=5%	263,025	δ=-2	128,278	ν= 5	263,025
σ=6%	111,003	δ=-1	263,025	ν=10	261,549,280
σ=7%	53,975	δ=0	1,228,719	ν=15	2,076,298,100
σ=8%	29,117	δ=1	17,704,201	ν=20	2,111,961,600
σ=9%	17,006	δ=2	438,943,140	ν=25	2,112,141,300

Table 3-4: Simulations and the Non-Normality Premium

The percentage of funds with non-normality premiums listed is reported from simulations using a censored non-central *t* distribution. The observations are censored from below at -0.99, the 1% mean and reported standard deviations are from the uncensored distribution. 10,000 funds are simulated.

	Less than -10 basis points	Between -10 and -5 basis points	Between -5 and 0 basis points	Between 0 and +5 basis points	Between +5 and +10 basis points	More than +10 basis points
Panel A: By Risk Aversion, σ=5%, δ=−1, ν=5						
γ=1	0%	0%	15%	84%	1%	1%
γ=2	0%	0%	10%	84%	4%	2%
γ=3	0%	0%	8%	74%	10%	7%
γ=4	0%	0%	7%	61%	17%	15%
γ=5	0%	0%	6%	46%	22%	26%
Panel B: By Non-Centrality Parameter, σ=5%, ν=5, γ=3						
δ=−2	0%	0%	0%	70%	16%	13%
δ=−1	0%	0%	8%	74%	10%	8%
δ=0	1%	1%	53%	40%	2%	2%
δ=+1	3%	8%	85%	4%	0%	0%
δ=+2	6%	14%	80%	0%	0%	0%
Panel C: By Degrees of Freedom, σ=5%, δ=−1, γ=3						
ν=5	0%	0%	8%	75%	10%	7%
ν=10	0%	0%	26%	71%	2%	1%
ν=15	0%	0%	38%	62%	0%	0%
ν=20	0%	0%	47%	53%	0%	0%
ν=25	0%	0%	51%	49%	0%	0%
Panel D: By Standard Deviation, δ=−1, ν=5, γ=3						
σ=5%	0%	0%	8%	75%	10%	8%
σ=6%	0%	0%	7%	60%	17%	16%
σ=7%	0%	0%	6%	43%	22%	29%
σ=8%	0%	0%	5%	29%	22%	44%
σ=9%	0%	1%	5%	19%	19%	56%

Table 3-5: Simulations and the Utility Adjustment to the Standard Deviation

The ratio of the standard deviation to the normal equivalent-standard deviation, *UAS*, from simulations using a censored non-central *t* distribution. The observations are censored from below -0.99, the 1% mean and reported standard deviations are from the uncensored distribution. 10,000 funds are simulated.

	More than 10% smaller	**Between 5% and 10% smaller**	**Between 0% and 5% smaller**	**Between 0% and 5% larger**	**Between 5% and 10% larger**	**More than 10% larger**
Panel A: By Risk Aversion, σ=5%, δ=−1, ν=5						
γ=1	0%	0%	15%	80%	4%	1%
γ=2	0%	0%	10%	77%	9%	4%
γ=3	0%	0%	8%	71%	13%	7%
γ=4	0%	0%	7%	62%	19%	11%
γ=5	0%	0%	6%	53%	24%	17%
Panel B: By Non-Centrality Parameter, σ=5%, ν=5, γ=3						
δ=−2	0%	0%	0%	61%	25%	14%
δ=−1	0%	0%	8%	71%	14%	7%
δ=0	0%	2%	54%	39%	3%	2%
δ=+1	2%	12%	82%	4%	0%	0%
δ=+2	5%	27%	68%	0%	0%	0%
Panel C: By Degrees of Freedom, σ=5%, δ=−1, γ=3						
ν=5	0%	0%	8%	71%	14%	7%
ν=10	0%	0%	26%	72%	2%	0%
ν=15	0%	0%	37%	62%	1%	0%
ν=20	0%	0%	47%	53%	0%	0%
ν=25	0%	0%	51%	49%	0%	0%
Panel D: By Standard Deviation, δ=−1, ν=5, γ=3						
σ=5%	0%	0%	8%	71%	14%	7%
σ=6%	0%	0%	7%	64%	18%	11%
σ=7%	0%	0%	6%	57%	21%	16%
σ=8%	0%	0%	5%	50%	25%	20%
σ=9%	0%	0%	5%	43%	27%	26%

Table 3-6: Non-Normality Premiums by Percentiles for a Normal Distribution of Returns

NNP are computed for each of 10,000 simulated funds. The *NNP*s are reported at the 1st, 5th, 10th, 90th, 95th, and 99th percentiles. The simulations are conducted using the normal distribution with the base-case of 1% mean, 5% standard deviation, risk aversion of 3, and 120 months of returns. These parameters are varied in the panels.

Panel A: By The Risk Aversion Parameter, σ=5%, T=120

	1%	**5%**	**10%**	**90%**	**95%**	**99%**
γ=1	-0.3	-0.2	-0.2	0.0	0.0	0.1
γ=2	-0.8	-0.6	-0.5	0.1	0.2	0.4
γ=3	-1.6	-1.2	-0.9	0.3	0.5	1.1
γ=4	-2.5	-1.9	-1.5	0.7	1.0	2.0
γ=5	-3.6	-2.7	-2.2	1.1	1.7	3.2

Panel B: By Standard Deviation, γ=3, T=120

	1%	**5%**	**10%**	**90%**	**95%**	**99%**
σ=3%	-0.4	-0.3	-0.2	0.0	0.0	0.1
σ=5%	-1.6	-1.2	-0.9	0.3	0.5	1.1
σ=7%	-4.3	-3.1	-2.5	1.3	2.0	3.5
σ=9%	-9.1	-6.4	-5.2	3.2	4.8	8.7

Panel C: By Sample Size, σ=5%, γ=3

	1%	**5%**	**10%**	**90%**	**95%**	**99%**
T=60	-2.5	-1.8	-1.5	0.2	0.6	1.4
T=120	-1.6	-1.2	-0.9	0.3	0.5	1.1
T=240	-1.0	-0.7	-0.6	0.3	0.5	0.8

Table 3-7: Utility Adjustments to the Standard Deviation by Percentiles for a Normal Distribution of Returns

UAS are computed for each of 10,000 simulated funds. The *UAS*s are reported at the 1st, 5th, 10th, 90th, 95th, and 99th percentiles. The simulations are conducted using the normal distribution with the base-case of 1% mean, 5% standard deviation, risk aversion of 3, and 120 months of returns. These parameters are varied in the panels.

Panel A: By The Risk Aversion Parameter, σ=5%, T=120

	1%	**5%**	**10%**	**90%**	**95%**	**99%**
γ=1	-1.2	-1.0	-0.8	0.0	0.1	0.4
γ=2	-1.6	-1.2	-1.0	0.2	0.4	0.8
γ=3	-2.0	-1.5	-1.3	0.4	0.7	1.3
γ=4	-2.4	-1.8	-1.5	0.7	1.0	1.8
γ=5	-2.8	-2.1	-1.7	0.9	1.3	2.3

Panel B: By Standard Deviation, γ=3, T=120

	1%	**5%**	**10%**	**90%**	**95%**	**99%**
σ=3%	-1.4	-1.1	-0.9	0.1	0.2	0.6
σ=5%	-2.0	-1.5	-1.3	0.4	0.7	1.3
σ=7%	-2.7	-2.0	-1.7	0.9	1.3	2.2
σ=9%	-3.4	-2.5	-2.1	1.2	1.9	3.2

Panel C: By Sample Size, σ=5%, γ=3

	1%	**5%**	**10%**	**90%**	**95%**	**99%**
T=60	-3.0	-2.3	-2.0	0.3	0.8	1.7
T=120	-2.0	-1.5	-1.3	0.4	0.7	1.3
T=240	-1.3	-1.0	-0.8	0.4	0.6	1.0

Table 4-1: Descriptive Statistics

Descriptive statistics for the 373 hedge funds reporting complete monthly data from 1996 to 2005 available to hedgefund.net. The returns are absolute and net of fees.

Panel A: By Category

	Number of Funds	Mean	Standard Deviation	Skewness	Kurtosis	Reject Normality	Auto Correlation	Market Correlation
All Funds	373	1.01	3.36	-0.19	8.60	86%	17%	37%
Location:								
Onshore	211	1.05	3.52	-0.24	8.46	85%	16%	40%
Offshore	131	0.95	3.20	-0.08	8.44	87%	17%	35%
Both	31	0.96	2.96	-0.38	10.21	87%	26%	25%
Fund of Funds:								
Single Strategy	29	0.87	2.66	-0.10	8.27	97%	20%	47%
Multi Strategy	101	0.88	2.20	-0.40	9.26	89%	21%	40%
Market Neutral	10	0.84	0.97	-0.07	9.08	90%	34%	29%
Equity Positions:								
Short Bias	7	0.85	6.86	0.04	8.83	71%	3%	-56%
Market Neutral Equity	15	0.72	2.00	0.03	6.55	87%	10%	7%
Long Only	16	1.16	6.59	-0.31	6.79	94%	8%	65%
Long/Short Equity	97	1.14	4.63	0.32	6.80	82%	12%	41%
Other:								
Arbitrage	55	0.83	1.97	-0.89	10.82	85%	22%	21%
Value	34	1.24	4.96	0.09	6.56	91%	11%	44%
Distressed	25	1.13	2.99	-0.58	9.20	92%	23%	38%

Panel B: All Funds

	Mean	Standard Deviation	Skewness	Kurtosis	Auto Correlation	Market Correlation
Mean	1.01	3.36	-0.19	8.60	17%	37%
Standard Deviation	0.42	2.41	1.46	9.38	16%	30%
Minimum	-0.38	0.20	-8.61	1.55	-37%	-90%
Median	0.94	2.74	-0.06	5.81	17%	42%
Maximum	3.78	15.51	3.93	89.65	96%	91%

Table 4-2: Non-Normality Premiums

Non-normality premiums (in basis points) for the 373 hedge funds with complete data from 1996 to 2005.

Panel A: Non-Normality Premiums by Risk Aversion

		Non-Normality Premiums						
		Less than -10	Between -5 and -10	Between -1 and -5	Between -1 and +1	Between +1 and +5	Between +5 and +10	More than +10
Risk Aversion	γ=1	0	3	10	350	7	3	0
	γ=2	2	7	29	315	12	2	6
	γ=3	7	5	41	284	19	8	9
	γ=4	11	8	50	249	34	3	18
	γ=5	13	16	46	230	38	12	18

Panel B: Non-Normality Premiums by Standard Deviations

		Non-Normality Premiums							
		Less than -10	Between -5 and -10	Between -1 and -5	Between -1 and +1	Between +1 and +5	Between +5 and +10	More than +10	Total
Standard Deviation	2%	0	0	0	134	0	0	0	134
	4%	0	0	7	111	4	2	0	124
	6%	0	3	19	31	10	1	1	65
	8%	1	0	14	6	4	1	3	29
	10%	4	2	1	2	0	2	3	14
	Above	2	0	0	0	1	2	2	7
	Total	7	5	41	284	19	8	9	373

Table 4-3: Utility Adjustments to the Standard Deviations

Utility adjustments to the standard deviations (in percentages) for the 373 hedge funds with complete data from 1996 to 2005.

Panel A: Utility Adjustments to the Standard Deviation by Risk Aversion

Risk Aversion	Utility Adjusted-Standard Deviation: Less than -10	Between -5 and -10	Between -1 and -5	Between -1 and +1	Between +1 and +5	Between +5 and +10	More than +10
$\gamma=1$	0	4	67	273	20	6	3
$\gamma=2$	1	6	89	217	50	2	8
$\gamma=3$	2	9	92	179	73	9	9
$\gamma=4$	4	11	99	154	81	12	12
$\gamma=5$	5	17	101	123	99	11	17

Panel B: Utility Adjustments to the Standard Deviation by Standard Deviations

Standard Deviation	Utility Adjusted-Standard Deviation: Less than -10	Between -5 and -10	Between -1 and -5	Between -1 and +1	Between +1 and +5	Between +5 and +10	More than +10	Total
2%	0	0	15	93	26	0	0	134
4%	0	1	42	53	21	3	4	124
6%	1	6	19	18	18	1	2	65
8%	0	1	10	11	4	2	1	29
10%	1	1	4	3	2	1	2	14
Above	0	0	2	1	2	2	0	7
Total	2	9	92	179	73	9	9	373

Table 4-4: Percentiles by Risk Aversion

Percentiles of the non-normality premiums and utility adjusted-standard deviation for the 373 hedge funds with complete data from 1996 to 2005 as stand alone investments and as included in a portfolio of T-bills and the stock market. All done for risk aversion levels of one to five.

Panel A: Non-Normality Premiums (in basis points), Stand Alone Investment

Non-Normality Premium
Stand Alone Investment
Percentiles

Risk Aversion	1%	5%	10%	50%	90%	95%	99%
γ=1	-3.8	-0.7	-0.4	0.0	0.1	0.3	2.9
γ=2	-7.7	-1.9	-1.0	0.0	0.4	1.1	12.5
γ=3	-13.6	-3.4	-1.6	0.0	1.0	2.9	36.3
γ=4	-20.9	-5.3	-2.5	0.0	1.8	5.5	82.5
γ=5	-28.8	-7.4	-3.5	0.0	3.2	9.1	162.0

Panel B: Non-Normality Premiums (in basis points), Portfolio Investment

Non-Normality Premium
Portfolio Investment
Percentiles

Risk Aversion	1%	5%	10%	50%	90%	95%	99%
γ=1	0.0	0.0	0.0	0.1	0.1	0.1	0.2
γ=2	0.2	0.3	0.3	0.4	0.4	0.5	0.7
γ=3	0.4	0.4	0.5	0.5	0.6	0.7	0.8
γ=4	0.2	0.3	0.3	0.4	0.4	0.5	0.6
γ=5	0.0	0.2	0.2	0.3	0.3	0.4	0.5

{table continued on next page}

Panel C: Utility Adjustment to the Standard Deviation (in percentages), Stand Alone Investment

Utility Adjustment-Standard Deviation
Stand Alone Investment
Percentiles

Risk Aversion		1%	5%	10%	50%	90%	95%	99%
	γ=1	-5.6	-2.4	-1.6	-0.3	0.8	1.8	9.1
	γ=2	-7.7	-3.1	-2.2	-0.2	1.5	3.1	15.3
	γ=3	-9.3	-3.9	-2.7	-0.1	2.2	4.5	22.6
	γ=4	-10.7	-4.6	-3.1	0.0	3.2	6.1	30.8
	γ=5	-11.8	-5.2	-3.6	0.0	4.0	7.8	40.2

Panel D: Utility Adjustment to the Standard Deviation (in percentages), Portfolio Investment

Utility Adjustment-Standard Deviation
Portfolio Investment
Percentiles

Risk Aversion		1%	5%	10%	50%	90%	95%	99%
	γ=1	0.3	0.4	0.5	0.6	0.7	0.7	0.9
	γ=2	0.7	0.9	1.0	1.1	1.3	1.4	1.7
	γ=3	1.0	1.1	1.2	1.4	1.5	1.7	2.0
	γ=4	0.7	0.9	1.0	1.3	1.5	1.6	2.1
	γ=5	0.4	0.8	0.9	1.2	1.4	1.6	2.3

Table 4-5: Spearman Rank Correlations

The Spearman rank correlation between the Sharpe ratio using the standard deviation to measure risk and the normal equivalent-standard deviation to measure risk for the 373 hedge funds with data from 1996-2005 available from hedgefund.net.

Risk Aversion	Spearman Correlation
1	99.95%
2	99.88%
3	99.77%
4	99.66%
5	99.53%

Table 4-6: Non-Normality Premiums by Strategy

The non-normality premiums (in basis points) for the 373 hedge funds with complete data from 1996 to 2005 as stand alone investments and as included in a portfolio of T-bills and the stock market. The range for each strategy is shown by the minimum, median, and maximum non-normality premiums. All are computed with a risk aversion level of three.

		Non-Normality Premiums					
		Stand-Alone Investment			Portfolio Investment		
	Funds	Minimum	Median	Maximum	Minimum	Median	Maximum
All Funds	373	-26.1	0.0	116.3	0.2	0.5	1.0
Location:							
Onshore	211	-25.1	0.0	116.3	0.2	0.5	1.0
Offshore	131	-26.1	0.0	33.8	0.4	0.5	0.8
Both	31	-9.8	0.0	61.6	0.4	0.5	0.9
Fund of Funds:							
Single Strategy	29	-1.4	0.0	1.5	0.5	0.5	0.7
Multi Strategy	101	-2.7	0.0	68.2	0.4	0.5	1.0
Market Neutral	10	-0.1	0.0	0.0	0.5	0.5	0.6
Equity Positions:							
Short Bias	7	-3.6	-2.1	116.3	0.4	0.4	0.7
Market Neutral Equity	15	-0.6	0.0	0.2	0.4	0.5	0.6
Long Only	16	-16.3	0.8	9.1	0.4	0.5	0.8
Long/Short Equity	97	-12.6	-0.1	13.8	0.2	0.5	0.7
Other:							
Arbitrage	55	-2.5	0.0	116.3	0.5	0.5	0.7
Value	34	-16.3	-0.2	6.6	0.4	0.5	0.7
Distressed	25	-4.0	0.0	13.8	0.5	0.6	0.7

Table 4-7: Utility Adjustment to the Standard Deviations by Strategy

The utility adjustment to the standard deviations (in percentages) for the 373 hedge funds with complete data from 1996 to 2005 as stand alone investments and as included in a portfolio of T-bills and the stock market. The range for each strategy is shown by the minimum, median, and maximum utility adjustment to the standard deviations. All are computed with a risk aversion level of three.

$$UAS = \frac{NE_{\sigma}}{s} - 1$$

		Utility Adjustment Standard Deviation					
		Stand-Alone Investment			Portfolio Investment		
	Funds	Minimum	Median	Maximum	Minimum	Median	Maximum
All Funds	373	-13.1	-0.1	43.4	0.7	1.4	2.6
Location:							
Onshore	211	-13.1	0.0	43.4	0.7	1.4	2.6
Offshore	131	-10.1	-0.3	16.3	1.0	1.4	2.0
Both	31	-9.7	0.0	27.6	1.1	1.4	2.3
Fund of Funds:							
Single Strategy	29	-3.3	-0.5	3.9	1.2	1.4	1.7
Multi Strategy	101	-4.4	0.0	39.2	1.0	1.4	2.6
Market Neutral	10	-2.6	-0.2	1.0	1.3	1.4	1.4
Equity Positions:							
Short Bias	7	-2.8	-1.6	43.4	0.9	1.2	1.7
Market Neutral Equity	15	-3.9	-0.4	2.0	1.1	1.3	1.4
Long Only	16	-4.2	1.3	14.6	0.9	1.4	2.0
Long/Short Equity	97	-10.1	-0.4	6.3	0.7	1.3	1.8
Other:							
Arbitrage	55	-3.8	0.1	43.4	1.2	1.4	1.7
Value	34	-6.4	-0.5	4.8	0.9	1.4	1.8
Distressed	25	-4.8	0.7	7.0	1.4	1.5	1.8

Table 4-8: Hedge Fund Effects on Higher Moments

The effects of allocating ten percent into a hedge fund on the autocorrelation, skewness, and kurtosis of a portfolio.

Panel A: All funds, risk aversion level of three

		Autocorrelation			Skewness			Kurtosis		
	Funds	Min	Med	Max	Min	Med	Max	Min	Med	Max
All Funds										
5% Investment	373	3%	5%	8%	-0.80	-0.63	-0.50	3.15	3.45	4.26
10% Investment	373	1%	5%	12%	-1.01	-0.63	-0.36	2.94	3.48	5.57
15% Investment	373	-2%	6%	15%	-1.25	-0.63	-0.21	2.79	3.52	7.34
20% Investment	373	-6%	6%	18%	-1.53	-0.63	-0.06	2.69	3.56	9.50
25% Investment	373	-9%	6%	22%	-1.84	-0.63	0.19	2.53	3.61	11.98

Panel B: All funds, ten percent investment into the hedge fund

		Autocorrelation			Skewness			Kurtosis		
	Funds	Min	Med	Max	Min	Med	Max	Min	Med	Max
All Funds										
$\gamma = 1$	373	1%	5%	10%	-0.96	-0.64	-0.42	3.02	3.49	5.22
$\gamma = 2$	373	1%	5%	10%	-0.96	-0.64	-0.42	3.02	3.49	5.22
$\gamma = 3$	373	1%	5%	12%	-1.01	-0.63	-0.36	2.94	3.48	5.57
$\gamma = 4$	373	-1%	6%	14%	-1.12	-0.60	-0.23	2.79	3.44	6.51
$\gamma = 5$	373	-2%	7%	17%	-1.24	-0.57	-0.11	2.69	3.40	7.55

{Table continued on next page}

Panel C: Ten percent investment into the hedge fund, risk aversion set to three.

		Autocorrelation			Skewness			Kurtosis		
	Funds	Min	Med	Max	Min	Med	Max	Min	Med	Max
All Funds	373	1%	5%	12%	-1.01	-0.63	-0.36	2.94	3.48	5.57
Location:										
Onshore	211	1%	5%	12%	-1.01	-0.63	-0.36	2.94	3.49	5.57
Offshore	131	1%	5%	11%	-0.83	-0.63	-0.47	3.06	3.46	4.89
Both	31	4%	5%	7%	-0.90	-0.64	-0.53	3.23	3.47	5.15
Fund of Funds:										
Single Strategy	29	4%	5%	7%	-0.72	-0.64	-0.56	3.20	3.50	4.17
Multi Strategy	101	3%	5%	7%	-1.01	-0.64	-0.51	3.02	3.50	5.57
Market Neutral	10	5%	5%	6%	-0.66	-0.64	-0.60	3.33	3.49	3.56
Equity Positions:										
Short Bias	7	1%	4%	6%	-0.74	-0.59	-0.48	3.08	3.35	3.73
Market Neutral Equity	15	4%	5%	6%	-0.65	-0.62	-0.53	3.14	3.41	3.54
Long Only	16	4%	7%	9%	-0.83	-0.63	-0.44	3.24	3.55	4.50
Long/Short Equity	97	2%	6%	11%	-0.76	-0.61	-0.36	2.94	3.42	4.29
Other:										
Arbitrage	55	1%	5%	7%	-0.74	-0.64	-0.58	3.27	3.49	3.90
Value	34	2%	6%	12%	-0.76	-0.63	-0.44	3.01	3.52	4.29
Distressed	25	5%	6%	11%	-0.78	-0.67	-0.63	3.44	3.59	4.24

Table 5-1: Descriptive Statistics of Lost Funds

Descriptive statistics for 32 hedge funds that did not report complete information by March, 2006 for returns from 1996 to 2005. Data collected from hedgefund.net. The returns are absolute and net of fees. Eleven funds reported their December 2005 return after the first weekend of March 2006 but before the first weekend of April 2006. Another four funds did not report their December 2005 return by the first weekend of April 2006. Finally, seventeen more funds from an earlier version of this paper – onshore hedge funds and CTAs with complete data from 1995 to 2004, stopped reporting in 2005.

	Funds	Mean	Standard Deviation	Skewness	Kurtosis	Auto Correlation	Market Correlation
Late Reports	11	0.92	3.91	-0.21	8.45	13%	38%
Dropped Out After November 2005	4	0.75	1.14	0.49	5.95	20%	27%
Previous Draft (Dropped out in 2005)	17	0.87	4.64	0.33	9.14	10%	4%

Table 5-2: Descriptive Statistics

Descriptive statistics for the stock market, one-month T-bills, and the inflation rate (from the CPI) from 1996 to 2005.

	Stock Market	One month T-bills	Inflation Rate
Mean	0.86%	0.30%	0.21%
Standard Deviation	4.66%	0.15%	0.29%
Skewness	-0.69	-0.27	0.02
Kurtosis	3.53	1.57	4.33
Autocorrelation	4%	96%	33%
Correlation with:			
Stock Market	100%	1%	-10%
One month T-bills	1%	100%	2%
Inflation Rate	-10%	2%	100%

Table 5-3: Portfolio Allocation

The optimal portfolio allocation for the stock market and one-month T-bills, taking the inflation rate (from the CPI) into account. The optimal allocation is determined by considering each month from 1996 to 2005 as equally likely. The descriptive statistics of these portfolios are given in nominal terms.

Risk Aversion	1	2	3	4	5
Stock Allocation	100.00%	100.00%	81.45%	61.26%	49.06%
T-bills Allocation	0.00%	0.00%	18.55%	38.74%	50.94%
Mean	0.65%	0.65%	0.55%	0.43%	0.36%
Standard Deviation	4.70%	4.70%	3.83%	2.90%	2.33%
Skewness	-0.64	-0.64	-0.62	-0.59	-0.56
Kurtosis	3.45	3.45	3.41	3.35	3.29
Normality p-value	1%	1%	1%	2%	4%
Autocorrelation	5%	5%	5%	5%	6%
Certainty Equivalent	0.54%	0.42%	0.32%	0.26%	0.23%
Risk Premium	0.11%	0.22%	0.23%	0.17%	0.14%

Table 5-4: Portfolio Allocation with the Hedge Fund

The optimal portfolio allocation for the stock market and one-month T-bills, taking the inflation rate (from the CPI) into account and given a ten percent allocation to a hedge fund. The optimal allocation is determined by considering each month from 1996 to 2005 as equally likely. The minimum, median, and maximum allocation to stocks and T-bills are reported.

Panel A: All funds, risk aversion level of three

		Stock Index Allocation			T-Bill Allocation		
All Funds	Funds	Min	Med	Max	Min	Med	Max
5% Investment	373	74%	80%	88%	7%	15%	21%
10% Investment	373	66%	79%	90%	0%	11%	24%
15% Investment	373	58%	78%	85%	0%	7%	27%
20% Investment	373	51%	77%	80%	0%	3%	29%
25% Investment	373	43%	75%	75%	0%	0%	32%

Panel B: All funds, ten percent investment into the hedge fund

		Stock Index Allocation			T-Bill Allocation		
All Funds	Funds	Min	Med	Max	Min	Med	Max
$\gamma = 1$	373	90%	90%	90%	0%	0%	0%
$\gamma = 2$	373	90%	90%	90%	0%	0%	0%
$\gamma = 3$	373	66%	79%	90%	0%	11%	24%
$\gamma = 4$	373	46%	59%	74%	16%	31%	44%
$\gamma = 5$	373	34%	47%	62%	28%	43%	56%

{Table continued on next page}

Panel C: Ten percent investment into the hedge fund, risk aversion set to three.

		Stock Index Allocation			T-Bill Allocation		
	Funds	Min	Med	Max	Min	Med	Max
All Funds	373	66%	79%	90%	0%	11%	24%
Location:							
Onshore	211	66%	79%	90%	0%	11%	24%
Offshore	131	67%	80%	83%	7%	10%	23%
Both	31	69%	80%	90%	0%	10%	21%
Fund of Funds:							
Single Strategy	29	74%	79%	81%	9%	11%	16%
Multi Strategy	101	69%	80%	83%	7%	10%	21%
Market Neutral	10	80%	81%	81%	9%	9%	10%
Equity Positions:							
Short Bias	7	79%	90%	90%	0%	0%	11%
Market Neutral Equity	15	78%	81%	85%	5%	9%	12%
Long Only	16	67%	74%	79%	11%	16%	23%
Long/Short Equity	97	68%	77%	90%	0%	13%	22%
Other:							
Arbitrage	55	76%	81%	84%	6%	9%	14%
Value	34	69%	76%	90%	0%	14%	21%
Distressed	25	68%	79%	81%	9%	11%	22%

Table 5-5: Effects of Hedge Fund

The effects of allocating ten percent into a hedge fund. The percentage of funds that increase utility, the average changes in the certainty equivalent, risk premium, mean, and standard deviation are in basis points and are in nominal terms.

Panel A: All funds, risk aversion level of three

	Funds	Increased Utility	Change in *CE*	Change in *RP*	Change in Mean	Change in St.Dev.
All Funds						
5% Investment	373	96%	27.03	0.36	27.39	2.52
10% Investment	373	96%	53.13	1.48	54.60	11.72
15% Investment	373	96%	78.22	2.60	80.82	20.64
20% Investment	373	96%	102.20	2.72	104.92	19.94
25% Investment	373	95%	124.71	-4.55	120.16	-45.92

Panel B: All funds, ten percent investment into the hedge fund

	Funds	Increased Utility	Change in *CE*	Change in *RP*	Change in Mean	Change in St.Dev.
All Funds						
$\gamma = 1$	373	84%	29.63	-14.69	14.94	-321.40
$\gamma = 2$	373	93%	44.80	-29.86	14.94	-321.40
$\gamma = 3$	373	96%	53.13	1.48	54.60	11.72
$\gamma = 4$	373	96%	52.52	2.22	54.74	17.93
$\gamma = 5$	373	96%	51.91	2.83	54.74	22.94

{Table continued on next page}

Panel C: Ten percent investment into the hedge fund, risk aversion set to three.

	Funds	Increased Utility	Change in *CE*	Change in *RP*	Change in Mean	Change in St. Dev.
All Funds	373	96%	53.13	1.48	54.60	11.72
Location:						
Onshore	211	96%	55.94	1.56	57.50	12.60
Offshore	131	95%	47.99	1.65	49.64	13.21
Both	31	97%	55.69	0.17	55.87	-0.62
Fund of Funds:						
Single Strategy	29	100%	40.49	0.67	41.16	4.82
Multi Strategy	101	95%	46.61	0.51	47.13	2.77
Market Neutral	10	100%	50.45	0.08	50.54	0.28
Equity Positions:						
Short Bias	7	86%	95.01	-2.19	92.83	-14.47
Market Neutral Equity	15	93%	40.32	0.90	41.21	9.62
Long Only	16	88%	33.51	4.85	38.36	37.35
Long/Short Equity	97	94%	58.64	2.58	61.22	23.65
Other:						
Arbitrage	55	98%	47.59	0.66	48.26	4.37
Value	34	100%	65.27	2.69	67.96	21.99
Distressed	25	100%	65.43	0.23	65.66	-4.17

Figure 2-1: The Sharpe Ratio and Utility

The Sharpe ratio measures the reward per unit of risk. Funds A and B have different means and standard deviations but produce the same Sharpe ratio — they are equally efficient. Fund C has a lower Sharpe ratio yet has a higher utility. If the investor needed to choose only one fund, the best choice would be C. However, if the investor were allowed to lend or borrow at the risk free rate, Point D could be achieved by Fund A or Fund B. Point D has less risk and a higher return than does Fund C.

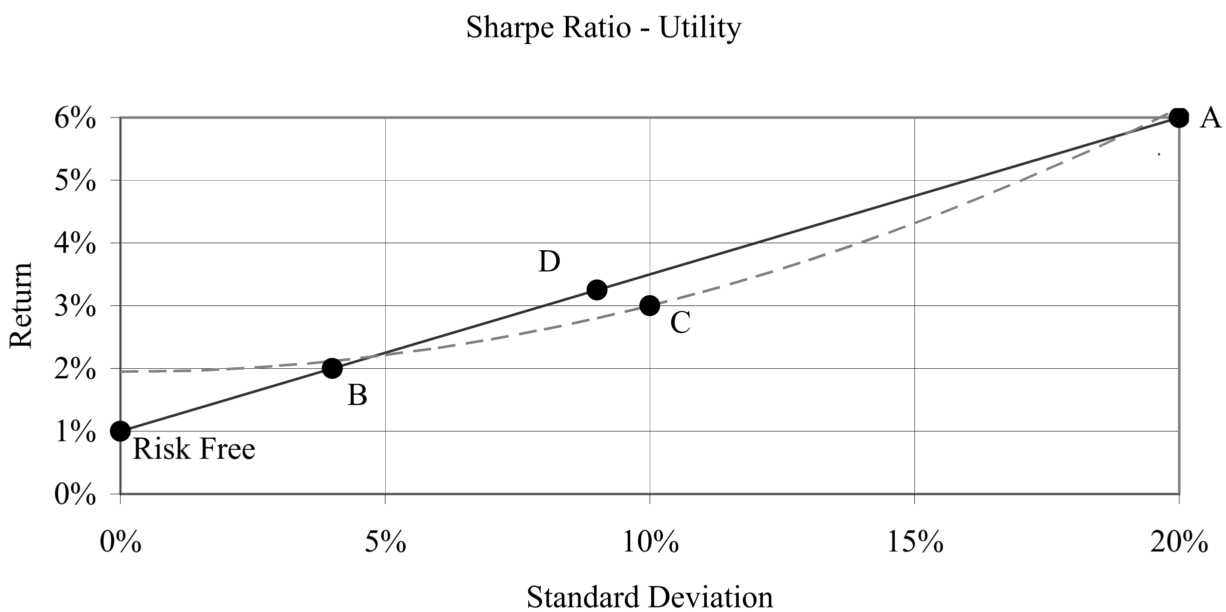

Figure 2-2: The Confidence Region of the Sharpe Ratio

The Sharpe ratio measures the reward per unit of risk. The variance of the Sharpe ratio increases as the estimate of the Sharpe ratio increases. The 95% confidence region with 120 observations for point estimates ranging from 0.1 to 2.0. The point estimates are shown as open circles, with the upper and lower bound as solid lines.

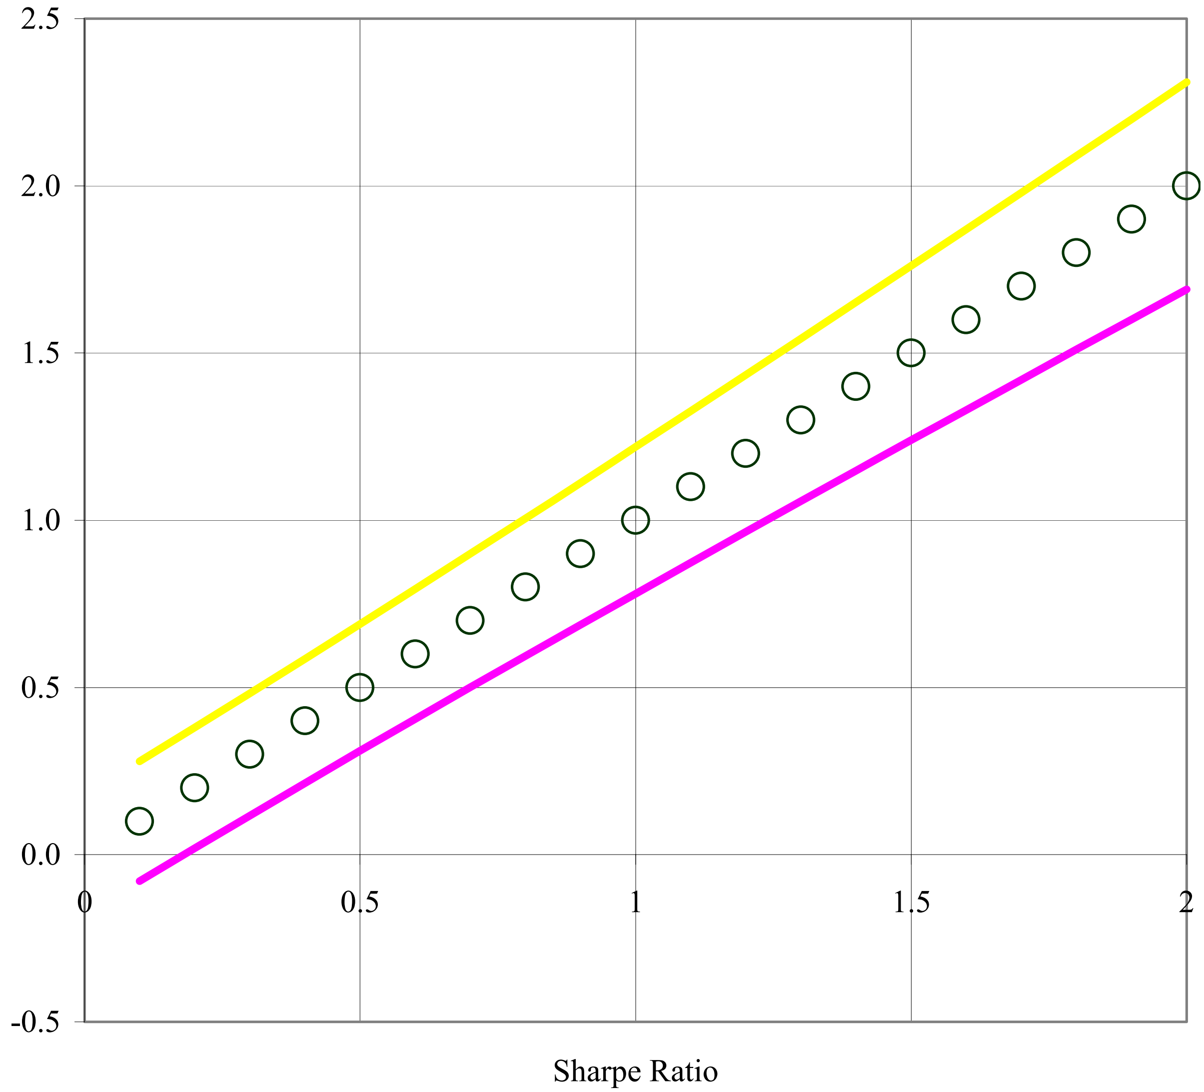

Figure 2-3: The Confidence Region of the Utility Adjustment to the Sharpe Ratio

The Sharpe ratio measures the reward per unit of risk. In order to produce a statistically significant difference in the Sharpe ratio, at the five percent level and with 120 observations, the utility adjustment to the Sharpe ratio must lie outside of these bounds. Points above the upper bound represent utility detracting sample distributions and show that the standard Sharpe ratio is overestimated.

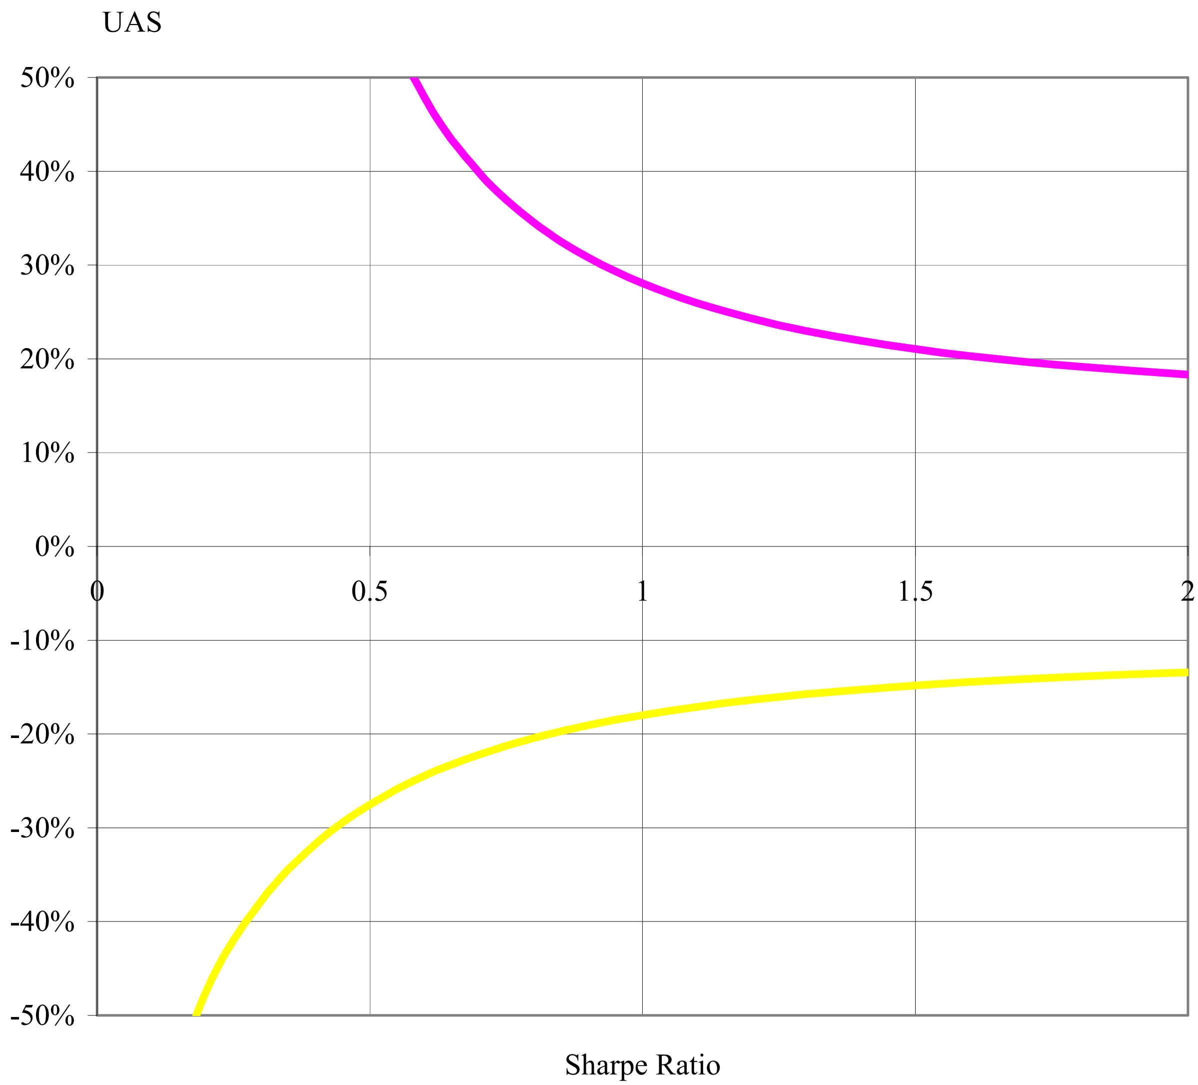

Figure 3-1: Indifference Curves from Taylor Series Approximations

The non-normality premiums as determined by the expected value of fourth order Taylor series expansions around the mean of 1%. The power utility function is used with the risk aversion parameter set to three. The monthly mean and standard deviation are 1% and 5%, respectively. Each point on a line yields the investor the same utility, these are indifference curves.

Panel A: The combination of skewness and kurtosis necessary to achieve given non-normality premiums of –5, 0, 5, and 10 basis points. Risk aversion parameter set to 3, standard deviation set to 5%.

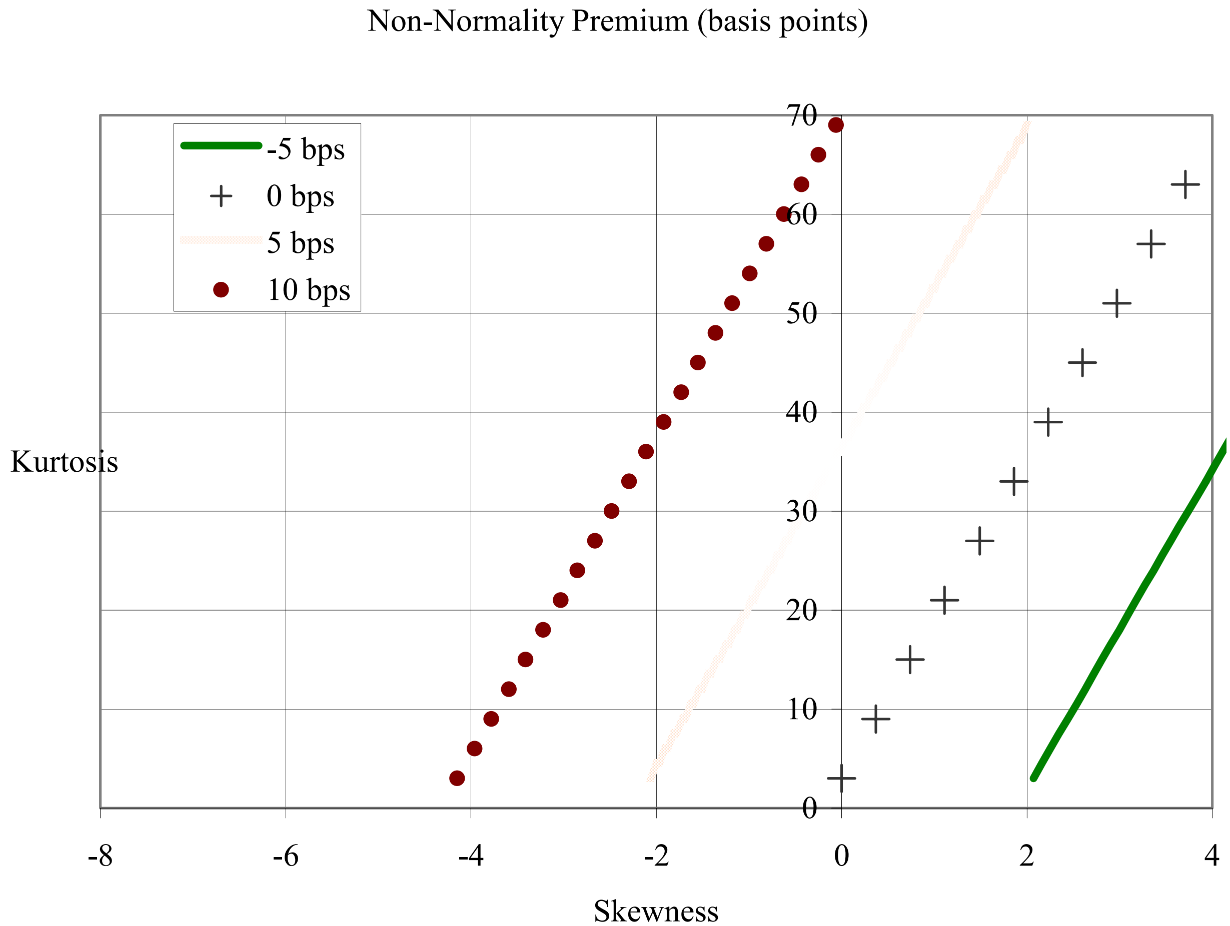

{figure continued on next page}

Panel B: The combination of skewness and kurtosis necessary to achieve a given non-normality premium of ten basis points given 5%, 7%, and 9% standard deviations. Risk aversion parameter set to three.

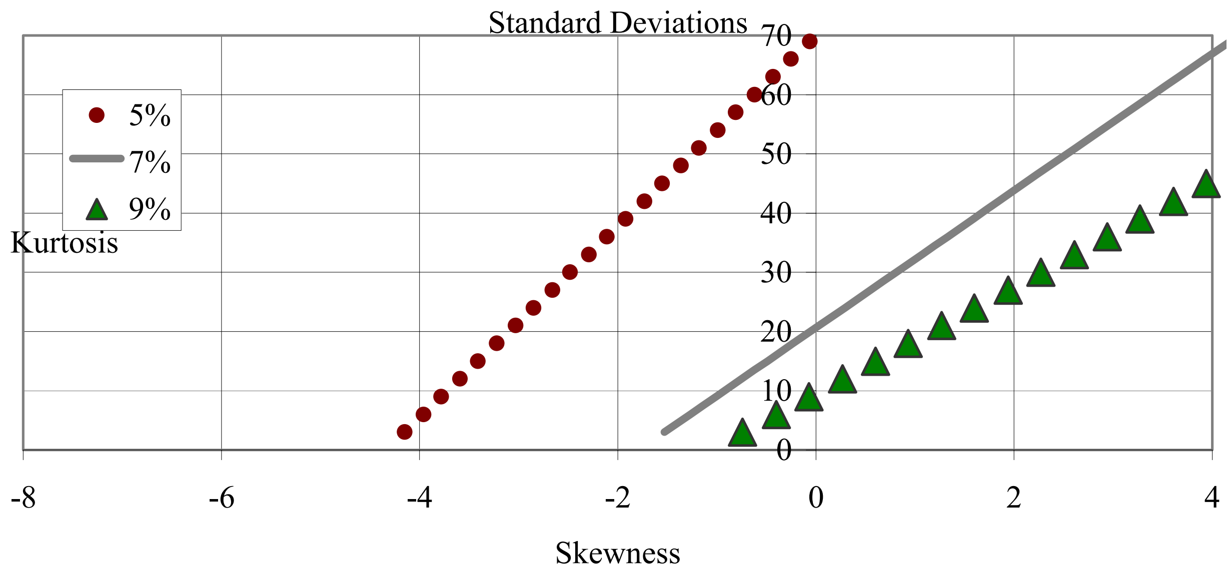

Panel C: The combination of skewness and kurtosis necessary to achieve given a non-normality premium of ten basis points given risk aversion coefficients of 2, 3, 4, and 5. Standard deviation set to 5%.

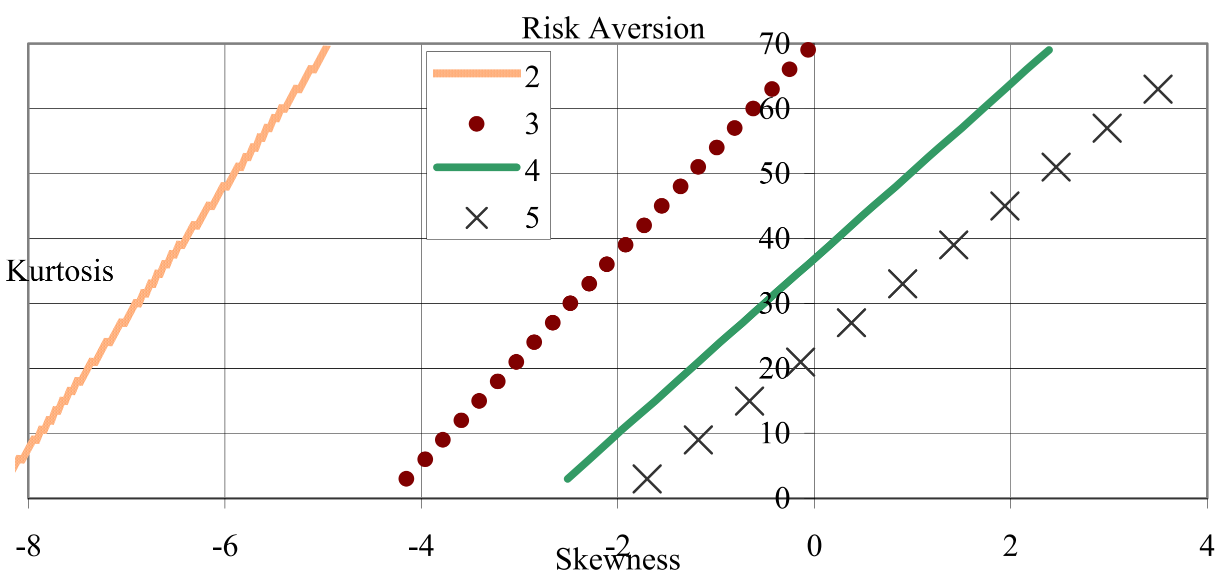

Figure 3-2: Indifference curves for the Utility Adjustment to the Standard Deviation (UAS)

Each point along the line is an indifference curve, representing the same amount of risk to utility for the level of kurtosis and skewness in the distribution with the same mean.

Panel A: Varying the *UAS* from 10% to –10%

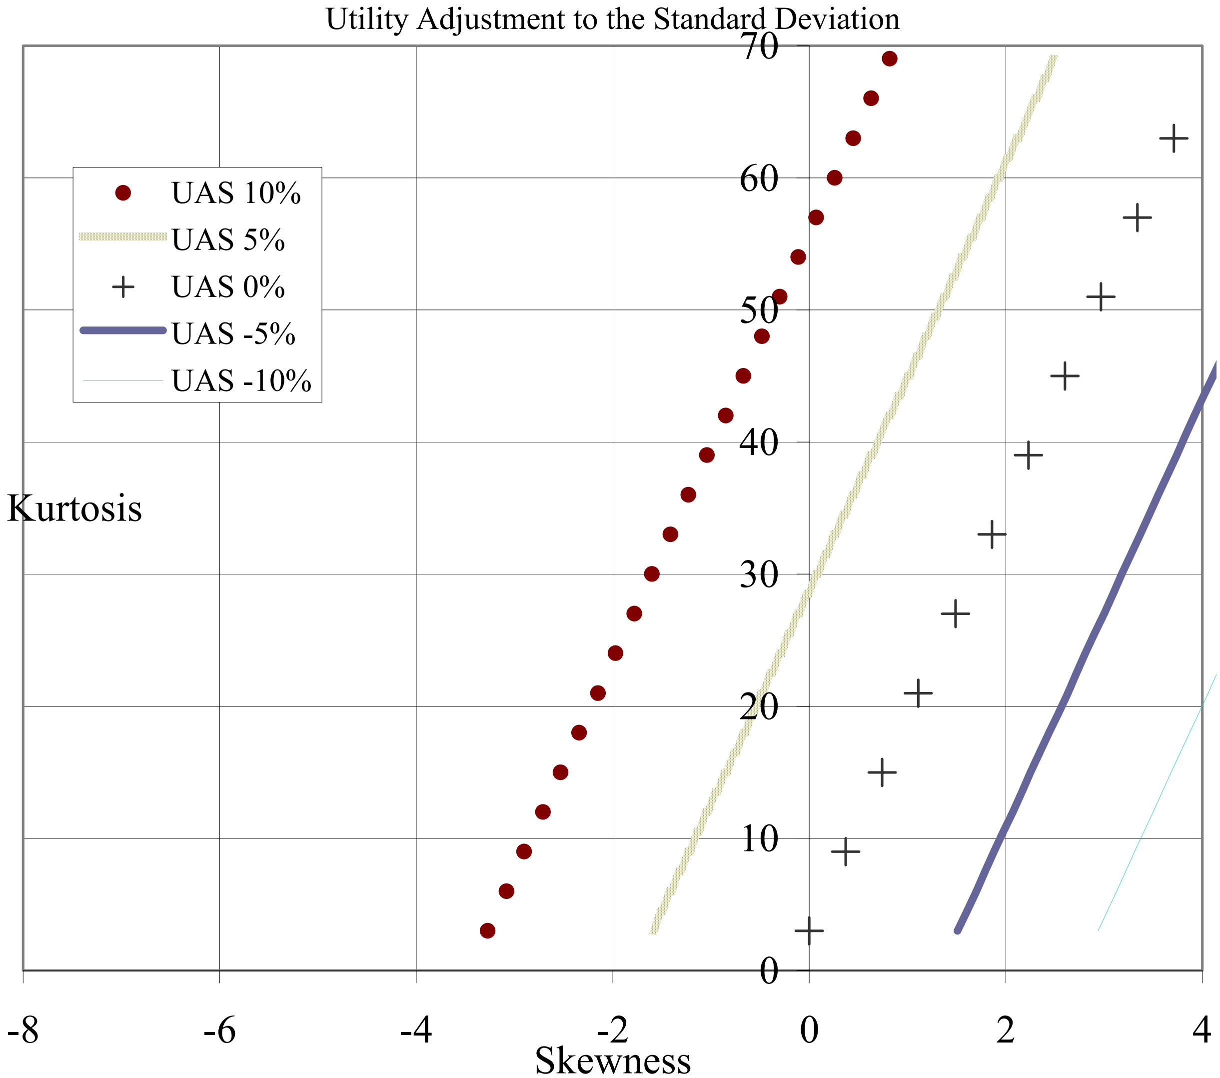

{Figure continues on next page}

Panel B: Varying the standard deviation from 3% to 9% per month

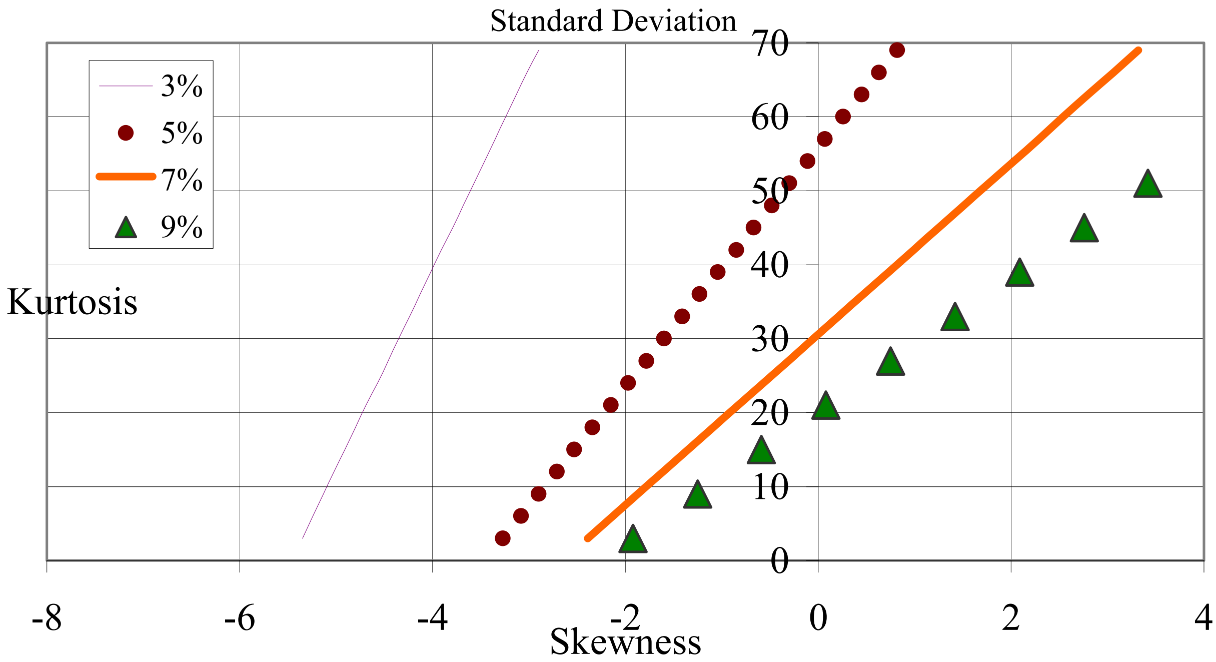

Panel C: Varying the risk aversion parameter from 1 to 5

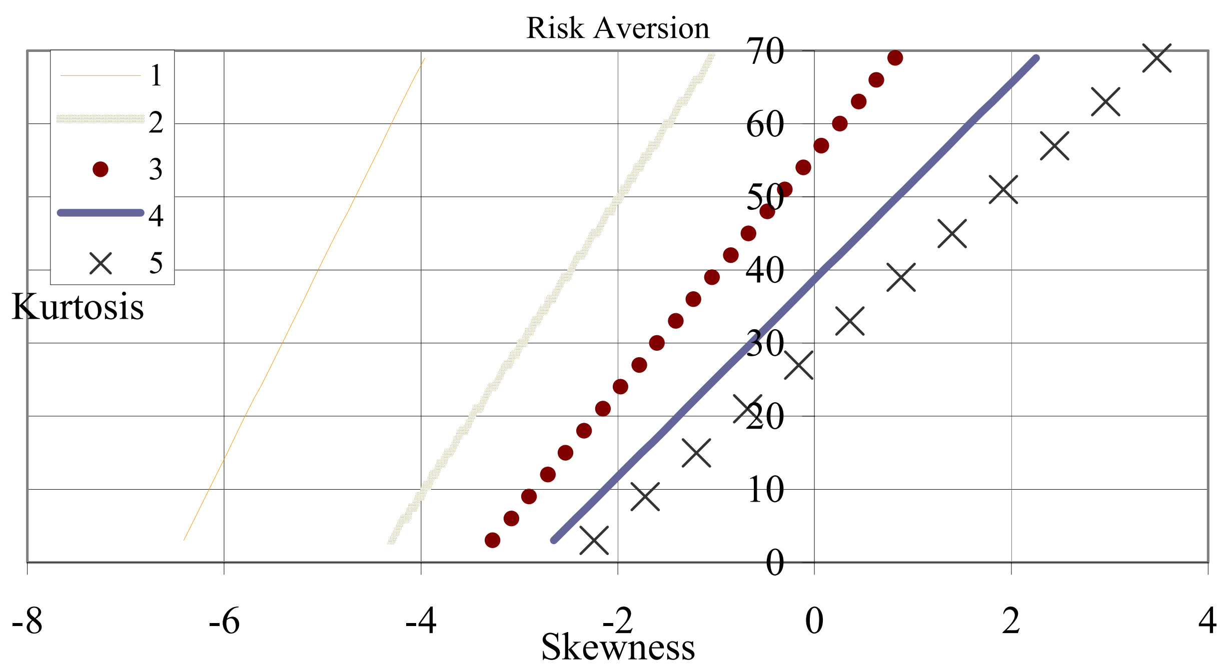

Figure 3-3: Non-Central *t* Distribution

The non-central *t* distribution is compared to the normal distribution. Both distributions have the same 1% mean and 5% standard deviation.

Panel A: Cumulative Density Functions

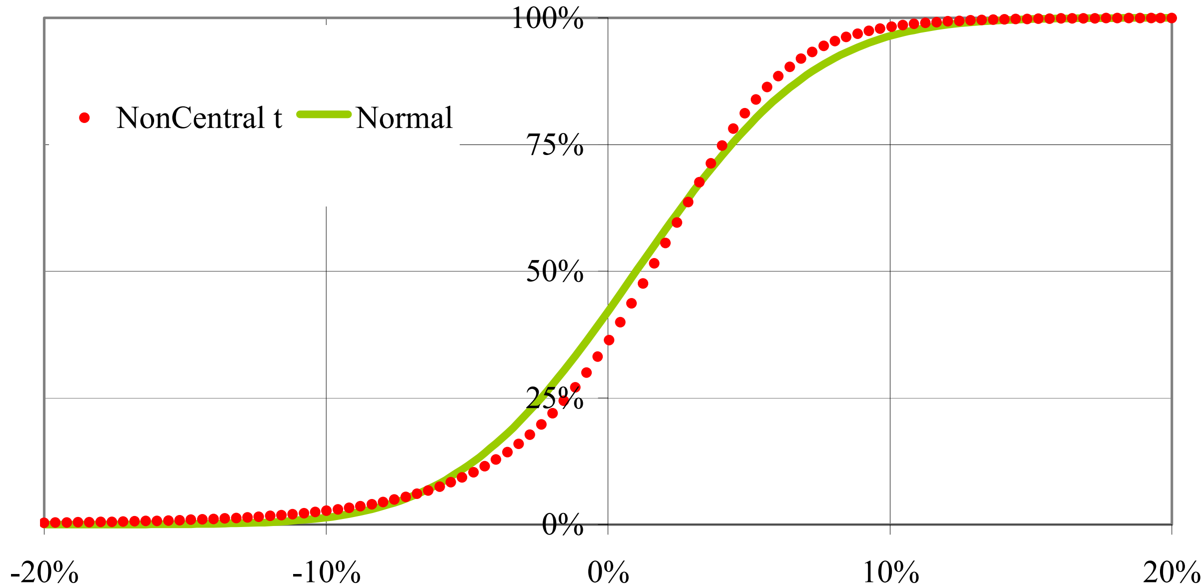

Panel B: Probability Density Functions

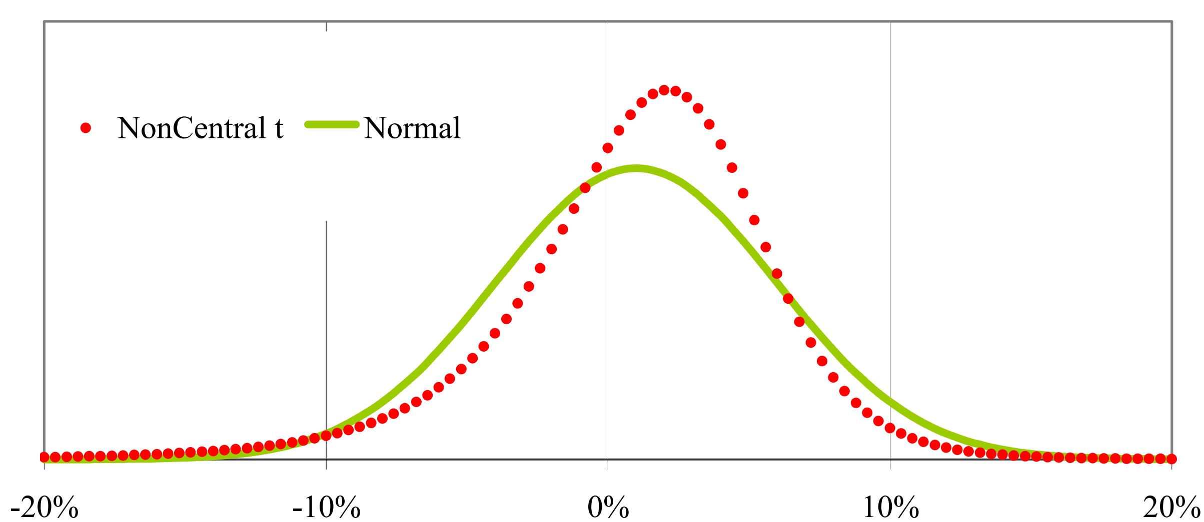

Figure 4-1: Hedge Funds – Risk and Return

Means and standard deviations for the 373 hedge funds reporting complete monthly data from 1996 to 2005 available at hedgefund.net. The returns are absolute, net of fees, and in percentages.

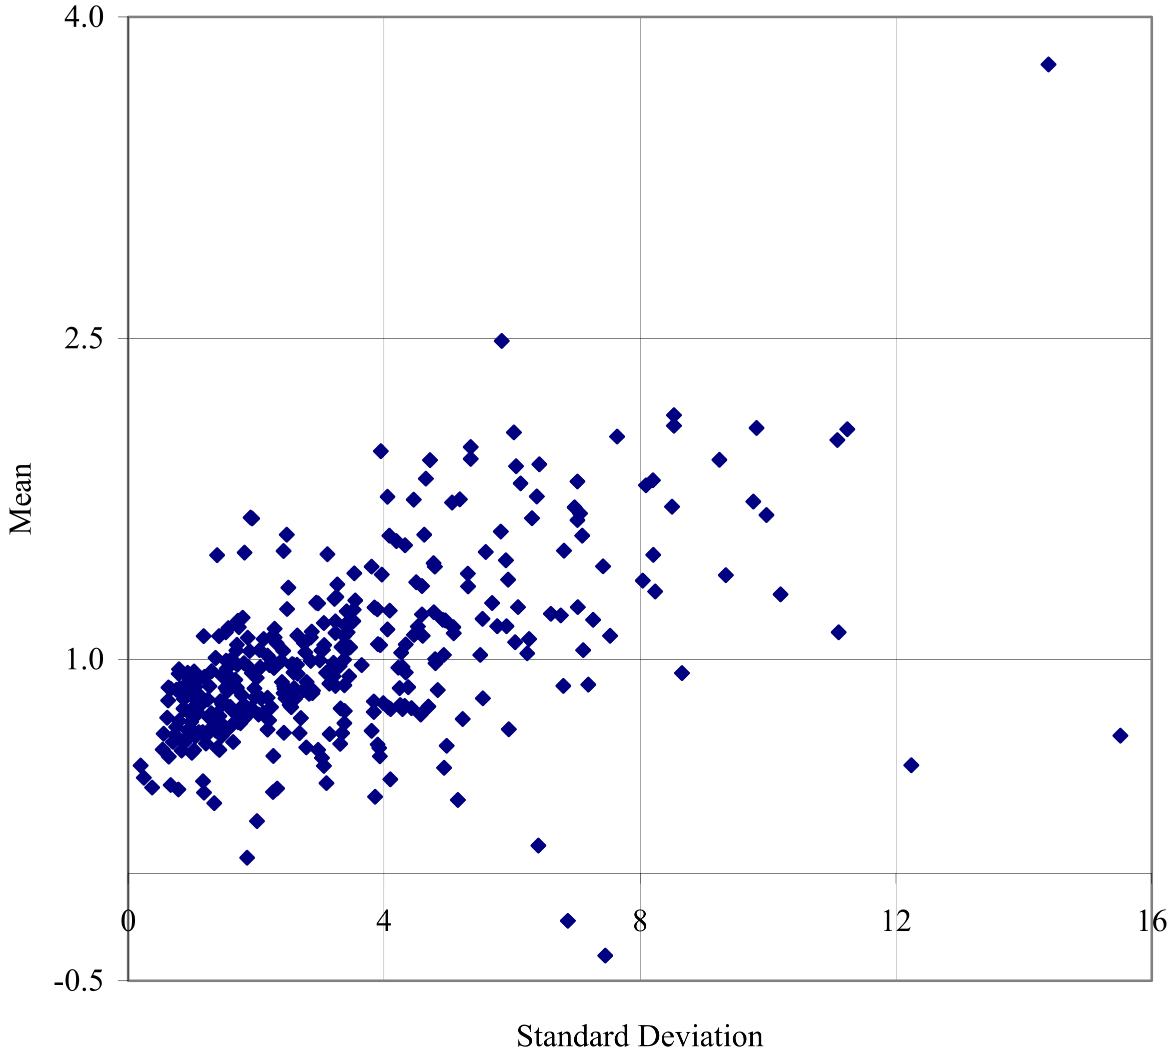

Figure 4-2: Hedge Funds – Skewness and Kurtosis

Skewness and kurtosis for the 373 hedge funds reporting complete monthly data from 1996 to 2005 available at hedgefund.net. The returns are absolute and net of fees.

Panel A: The number of funds that fall in the eight potential regions depending on if they reject normality or not, positive or negative skewness, excess kurtosis or no excess kurtosis.

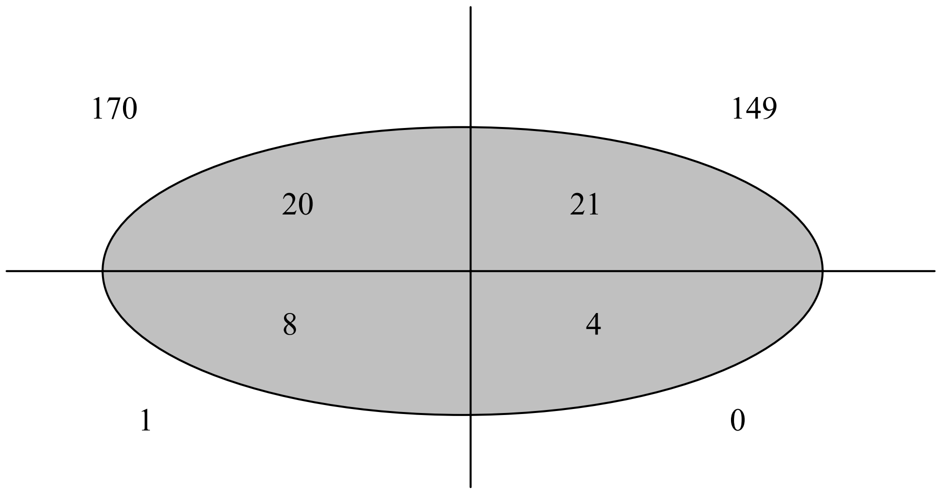

Panel B: The skewness and kurtosis of the 373 funds.

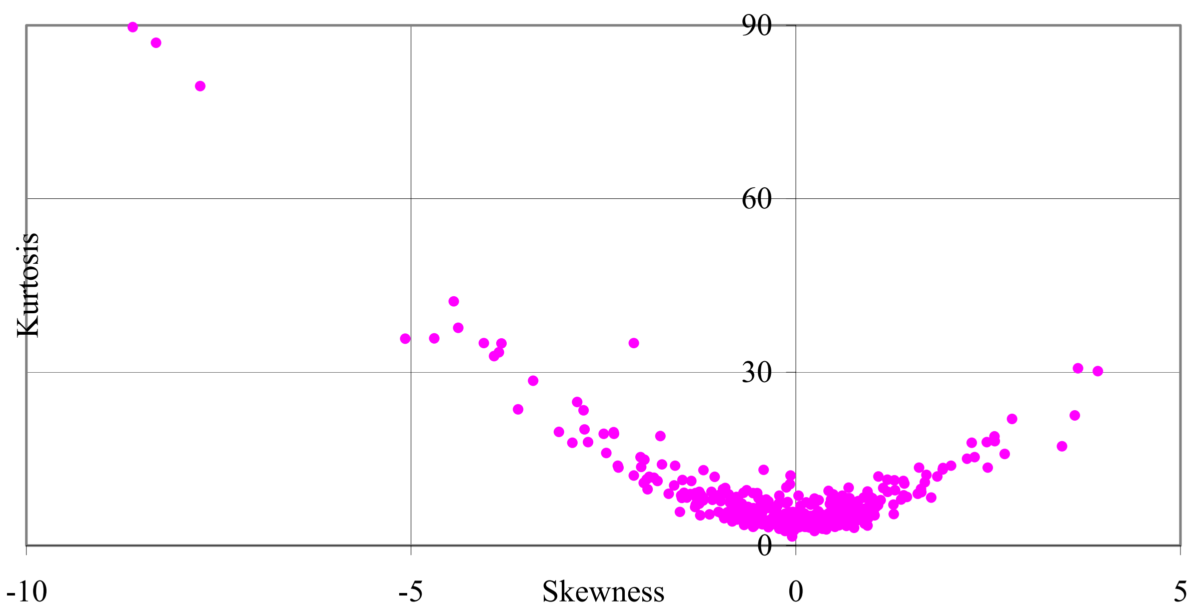

Figure 4-3: UAS and the Sharpe ratio

The utility adjustment to the standard deviation plotted against the Sharpe ratio for the 373 funds with complete monthly data from 1996-2005 made available by hedgefund.net. The solid lines represent the necessary *UAS* to generate a statistically significant Sharpe ratio. Sharpe ratios are in monthly terms. A positive *UAS* comes from a utility detracting distribution with the Sharpe ratio overestimated.

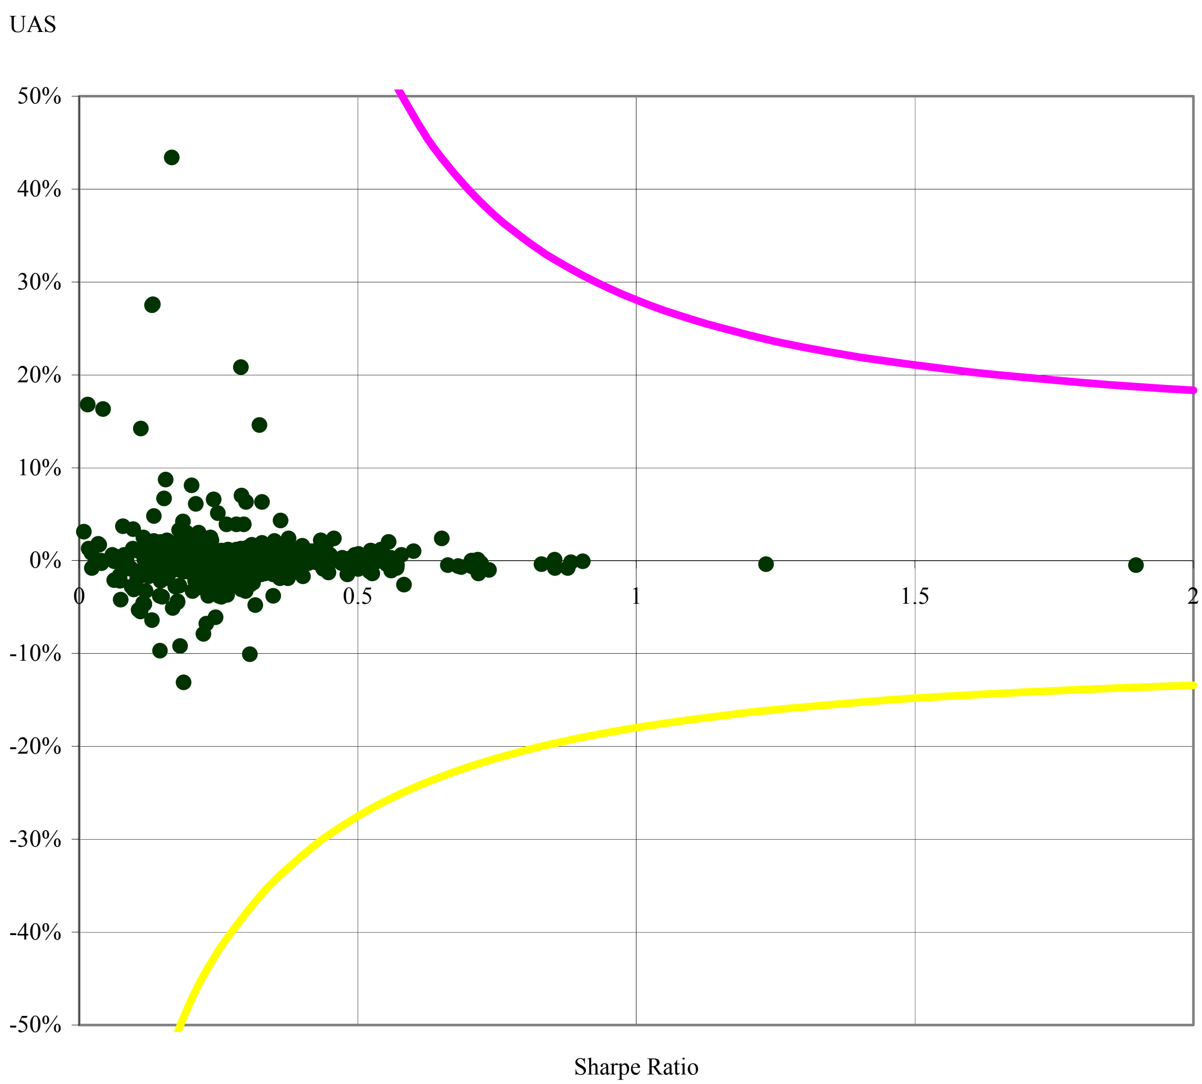

Figure 5-1: Compensating Returns

The amount of possible gain and loss, with equal probabilities, that would leave an investor indifferent between the risky gamble and a certain return of zero percent. The graph is plotted for the constant relative risk aversion power utility function and the constant absolute risk aversion exponential utility function with risk aversion levels of one, three, and five.

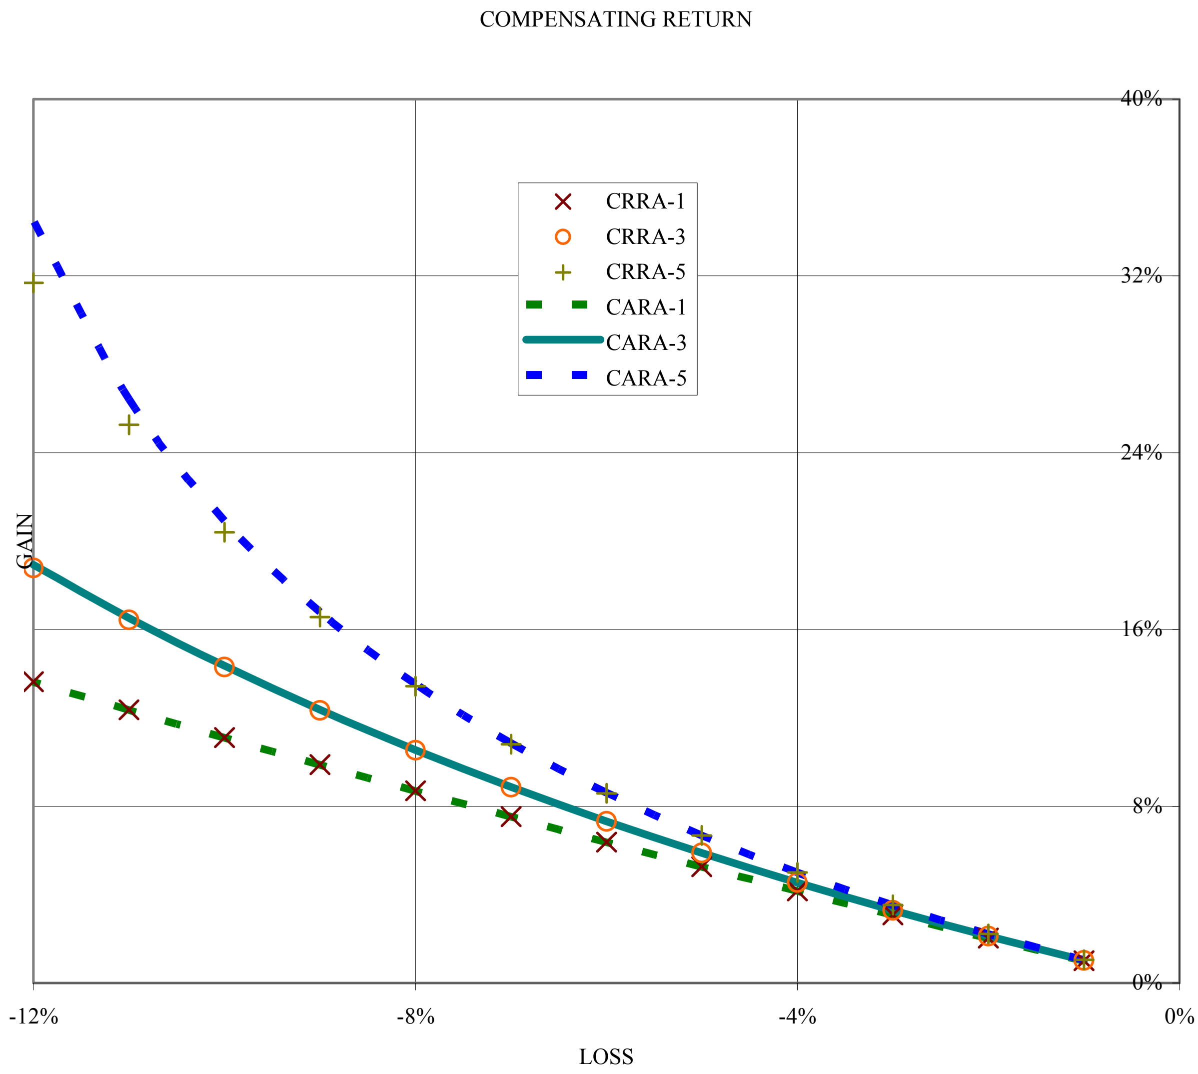

Figure 5-2: Optimal Portfolio Allocation

The certainty equivalents by their allocation into the stock market (the remainder in T-bills) depending upon the risk aversion level. The returns are in real terms by discounting by the Consumer Price Index.