

GENERALIZATION OF THE SHARPE RATIO AND THE ARBITRAGE-FREE PRICING OF HIGHER MOMENTS

Gaurav S. Amin^{*}

Harry M. Kat[#]

Working Paper

This version: May 14, 2002

Please address all correspondence to:

Harry M. Kat
ISMA Centre
The University of Reading
Whiteknights Park
Reading RG6 6BA
United Kingdom
Tel. +44-118-9316428
E-mail: h.kat@ismacentre.rdg.ac.uk

^{*}Ph.D student, ISMA Centre, University of Reading, [#]Associate Professor of Finance, ISMA Centre, University of Reading.

GENERALIZATION OF THE SHARPE RATIO AND THE ARBITRAGE-FREE PRICING OF HIGHER MOMENTS

ABSTRACT

We present an extension of the traditional Sharpe ratio to allow for the evaluation of non-normal return distributions. Combining earlier work in this area with stochastic simulation, we develop a procedure that allows for the construction of a benchmark for the evaluation of the performance of funds with a non-normal return distribution, while maintaining the operational ease of the Sharpe ratio. Similar to the latter, our procedure only requires the risk-free rate of interest rate, the distribution of the market index and an assumption about the type of return distributions to be evaluated. Unlike the Sharpe ratio, however, we are not restricted to normality but are able to handle any reasonable type of distribution. Since our benchmarking procedure is based on the no-arbitrage assumption, it also provides insight into the conditional arbitrage-free value of one distributional parameter in terms of another. We show that in case of the Johnson Su distribution the relationship between skewness and mean return is more or less flat. Skewness and median return on the other hand exhibit a strong negative relationship.

1. INTRODUCTION

Investment performance evaluation evolves around the construction of passive benchmarks. The latter can either be based on a fund's bottom line risk-return signature or on a fund's exposure to the relevant return generating factors. The Sharpe ratio (1966) and Jensen's alpha (1968) are the earliest and best-known examples of both these approaches. Over the past decennia, Jensen's original one factor model has undergone many extensions aimed at more fully capturing the return generating process. The Sharpe ratio has not seen a similar development. Recently, however, Amin and Kat (2001) employed a performance evaluation method that can be interpreted as a direct extension of the traditional Sharpe ratio. The basic motivation for the method is that the choice of evaluation method should depend on the nature of the return distribution of the funds to be evaluated. For those who tend to follow the factor model route, this will be an obvious point. However, this does not seem to be the case for users of the Sharpe ratio, who tend to apply it without much consideration for the actual risk-return profile of the funds that are being evaluated.

The Sharpe ratio of the market portfolio represents the set of return distributions that is obtained when statically combining the market portfolio with cash. With the market portfolio being highly diversified, this set of distributions offers the highest attainable expected return for every possible standard deviation. If a fund produces a return distribution with an expected return lower than that of a similar (in terms of standard deviation) distribution from this set, the fund is deemed to be inefficient and the other way around. Crucial in all this is the assumption that fund return distributions are fully described by their mean and standard deviation, i.e. are normal. If this is not the case, the Sharpe ratio is not the appropriate performance measure and is open to 'gaming'. Since the Sharpe ratio only looks at the mean and standard deviation of the return distribution one can easily create the illusion of superior performance by sacrificing the distribution's higher moments in exchange for a higher mean and/or lower standard deviation. Mechanically writing ordinary calls for example will raise the Sharpe ratio substantially.

Although the Sharpe ratio itself cannot, the (pragmatic) reasoning behind it can still be used when a fund's return distribution is non-normal. Instead of the highest

possible expected return on a portfolio with a given standard deviation, we should now look for the highest possible expected return on a portfolio that matches the standard deviation as well as all other relevant higher moments. This can be done on the basis of the individual moments but it is of course much more efficient to replicate the entire return distribution as a whole. This is the approach taken in Amin and Kat (2001). Since it makes no specific assumptions about fund return distributions, their method is extremely flexible. When a parameter is important enough to make a difference it will automatically be replicated so there is no need to explicitly decide which parameters to replicate.

One obvious advantage of the Sharpe ratio and an important reason behind its popularity with practitioners is that it makes performance evaluation operationally very simple. The highest attainable expected return for a given standard deviation can be calculated as the sum of the interest rate and the product of the Sharpe ratio of the market portfolio and the standard deviation in question. In other words, once we know the interest rate and the return distribution of the market portfolio we can evaluate an unlimited number of normally distributed funds. Since the method constructs tailor-made benchmarks on a fund-by-fund basis, it appears that the Amin and Kat (2001) method does not offer a similar convenient feature. This is not entirely true though. With the use of stochastic simulation, one can easily construct a general benchmark that can be used to evaluate all funds with a given type of return distribution. In the first part of this paper we will discuss how this is done and provide an example.

Since it looks for the highest possible expected return for a given set of parameters in an arbitrage-free world, an interesting by-product of the procedure discussed in this paper is that, given a particular type of distribution, it also allows us to draw conclusions about the arbitrage-free pricing of the distribution's parameters. Having determined the appropriate benchmark for a particular type of distribution, it is straightforward to tell how much an extra unit of standard deviation or skewness should be worth in terms of expected return for example. We will discuss this application in the second part of the paper.

The paper proceeds as follows. In section 2 we explain the step-by-step construction of the benchmark, while in section 3 we provide a concrete example based on Johnson's Su distribution. In section 4 we look at the arbitrage-free pricing of standard deviation and skewness in terms of the median return. Section 5 concludes.

2. CONSTRUCTION OF THE PERFORMANCE BENCHMARK

One of the most common errors made when evaluating investment performance is to choose the evaluation framework before carefully investigating the nature of the return distributions to be evaluated. Despite a high degree of non-normality, many authors have evaluated hedge fund performance for example using methods that require fund return distributions to be strictly normal. Without exception, they conclude that hedge funds offer investors tremendous benefits. When the whole distribution is taken into account, however, these benefits largely disappear. In Amin and Kat (2001) we discussed a procedure for performance evaluation that does not require any specific assumptions about the return distribution of the funds that are to be evaluated. The procedure was applied on a one-to-one basis to evaluate the performance of a number of hedge funds and hedge fund indices. When combined with some stochastic simulation, however, the same procedure can be used to create a general benchmark that can be used to evaluate any return distribution of a given type, irrespective of its particular parameter values. In what follows we discuss how to construct such a benchmark.

Since we want a benchmark that is capable of dealing with all stylised facts observed in fund returns, the first step in our evaluation procedure is to investigate what type of distribution best fits the available fund return data. After having done so, we use stochastic simulation methods to generate a variety of frequency distributions (of the above type) that cover all combinations of parameter values of practical interest. Suppose for example that, apart from the mean, the type of distribution that best described the available fund return data was fully defined by only two other parameters and that the parameter values estimated from fund returns ranged between 1 and 8 for the first and 5 and 10 for the second. In that case we would generate 48 (= 8×6) different frequency distributions with parameter values such as (1,5), (1,6),...

(2,5), (2,6), etc. and a mean equal to zero. The reason for the latter assumption will become clear later.

The third step is to turn the (let's assume monthly) return distributions thus generated into end-of-month payoff distributions by assuming we invest \$100 at the beginning of the month. The same is done for the market index. We do not make any assumptions about the distribution of fund returns but because we will use Black-Scholes (1973) type option pricing at a later stage, we explicitly assume the index to be normally distributed.

The fourth step is to construct a number of payoff functions that, when combined with the index distribution, reproduce the above payoff distributions. Since there are many functions that will map one distribution into the other, we make the additional assumption that each payoff function must be a path independent non-decreasing function of the index value at the end of the month.¹ Under the latter assumption, constructing the desired payoff function is quite straightforward. Suppose that the fund distribution told us there was a 10% probability of receiving a payoff lower than 100. We would then look up in the index distribution at which index value X there was a 90% probability of finding an index value higher than X . If we found $X=95$, the payoff function would be constructed such that when the index ended at 95 the payoff would be 100. Next, we would do the same for a probability of 20%. If the fund distribution told us there was a 20% probability of receiving a payoff lower than 110 and the index distribution said there was a 80% probability of finding an index value higher than 100, the payoff function would be constructed such that when the index ended at 100 the payoff would be 110. This procedure is repeated until we get to 100%.

The fifth step consists of finding the initial investment required by the self-financing dynamic trading strategy, trading the index and cash, that generates each payoff function in question. This is not different from what is usually done to price derivatives contracts. We can therefore use standard derivatives pricing technology and calculate the price of the derived payoff function as the discounted risk neutral expected payoff. Since the payoff is not a neat function, however, we do so using Monte Carlo simulation instead of a more elegant analytical or numerical technique.

The easiest approach is to assume we live in the world of Black and Scholes (1973). We can in that case generate end-of-month index values using the discretized and risk neutralized geometric Brownian motion given by

$$S(t + \delta t) = S(t) \exp\left((r - q - \frac{1}{2}\sigma^2)\delta t + \sigma \sqrt{\delta t} \varphi \right), \quad (1)$$

where $S(t)$ is the starting value of the index, r is the risk-free rate, q is the dividend yield on the index, σ is the index volatility, δt is the time step (one month) and φ is a random variable with a standard normal distribution. From the index values thus generated we subsequently calculate the corresponding payoffs, average them and discount the resulting average back to the present at the risk-free rate to give us the price of the payoff function in question.

Since we assume all return distributions to have a zero mean and assuming interest rates are non-zero, the above procedure will produce values that are below \$100, which implies that initially we should have specified a higher mean. To find out how much higher the mean should have been to get an initial value of \$100, the above procedure is repeated several times while raising the mean step by step. The mean that produces a payoff distribution that is valued at exactly \$100 is recorded as the highest attainable mean given the parameter values in question.

Since we now know the highest attainable mean for every possible combination of parameter values, we are ready to proceed with the actual performance evaluation. All this requires is to compare the mean returns of the funds to be evaluated with the highest attainable mean for parameter combinations equal to those of the latter funds. A mean higher than the highest attainable mean indicates superior performance and a lower mean inferior performance. Note that once the benchmark is constructed, the mechanics of the evaluation procedure are not different from a traditional Sharpe ratio analysis where funds' means are compared with those predicted by the capital market line. Like the Sharpe ratio, our benchmark only depends on the interest rate and the distribution of the index. Like the Sharpe ratio, we need to make an a priori assumption about the type of return distributions to be evaluated. Unlike the Sharpe

ratio, however, we are not forced to assume normality. Our procedure yields different benchmarks for different types of distributions.

3. AN EXAMPLE USING JOHNSON'S SU DISTRIBUTION

In this section we apply the above procedure assuming all fund returns follow a type of distribution known as a Johnson Su distribution. Johnson [1949, 1965] proposes a distribution, referred to as the Su distribution, with a distribution function given by

$$Y = \xi + \lambda \cdot \sinh\left[\frac{Z - \gamma}{\delta}\right], \quad Y > \xi, \quad (2)$$

where Z is a unit normal variable. The skewness and kurtosis of the distribution are determined by the shape parameters γ and δ . ξ and λ are the location and scale parameters respectively. Values for ξ and λ can be determined from the following expressions:

$$\Omega = \frac{\gamma}{\delta}$$

$$w = \exp(\delta^{-2})$$

$$E(Y) = -w^{\frac{1}{2}} \sinh \Omega$$

$$\sigma(Y) = \left[\frac{1}{2} (w - 1) (w \cosh 2\Omega + 1)^{\frac{1}{2}} \right]$$

$$\lambda = \frac{\sigma(Z)}{\sigma(Y)} \quad (3)$$

$$\xi = (\text{mean of } Z) - \lambda E(Y) \quad (4)$$

Assuming that all fund return distributions are Johnson Su distributions with zero excess kurtosis, we can evaluate these hypothetical funds' performance once we quantify the relationship between standard deviation, skewness and the highest possible expected return for this type of distribution.² To do so, we first have to generate a number of frequency distributions with varying standard deviations and skewness levels. Because the mean return is the variable that our procedure has to determine, we initially fix the mean of all distributions at 0%. Standard deviation is

varied from 1% to 5% (in 0.5 steps) and skewness from -1.5 to 1.5 (in 0.25 steps). Figure 1 and 2 provide two examples of frequency distributions (10,000 sample points) generated with a mean of 0%, a standard deviation of 1%, and a skewness of -1.5 and 1.5 respectively. The smooth curve is a normal distribution with the same mean and standard deviation

<< Insert Figure 1 and 2 >>

Given the above 117 (= 9 x 13) frequency distributions, the next step is to find out what these distributions' highest attainable means are given the interest rate and the distribution of the index. Given a distribution, using the methodology explained in the previous section and starting with a mean of 0%, we increase the mean in a stepwise fashion and re-price the distribution at every step. When the procedure prices the distribution at \$100 that value of the mean is recorded as the highest attainable mean return for the combination of standard deviation and skewness in question. After having done this for all 117 combinations of standard deviation and skewness, the results can conveniently be plotted in a 3-D graph. Assuming index returns are normally distributed with a mean return of 1.21% and a standard deviation of 4.14% (estimates from monthly S&P 500 returns over the period January 1990 - April 2000), the risk-free rate fixed at 5% and a dividend yield of 2.35%, the plane in figure 3 shows the highest attainable expected return as a function of standard deviation and skewness. All that is left at this point to finish the evaluation exercise is to plot the funds to be evaluated into mean-standard deviation-skewness space together with the benchmark plane. A fund that plots above the plane offers a superior risk-return profile, while funds that plot below offer inferior performance.

<< Insert Figure 3 >>

It is interesting to take a closer look at the shape of the benchmark in figure 3. In line with the idea that investors do not like standard deviation, we see a clear positive relationship between mean and standard deviation, irrespective of the level of skewness. Skewness on the other hand seems to have little impact on the mean. This goes very much against the generally accepted idea that investors prefer more

skewness over less and will therefore require a higher expected return when skewness deteriorates. It is important to note, however, that investor preferences are not really at stake here as the means in figure 3 are only indirectly determined by the marketplace. For different combinations of standard deviation and skewness, the plane in figure 3 shows the highest expected return that can be achieved following a dynamic trading strategy, trading the index and cash. The weak relationship between mean and skewness therefore is the result of the way the index is priced in the marketplace in combination with the dynamic trading strategy followed.

<< Insert Figure 4 >>

Although the benchmark does not require a higher expected return when skewness drops, this does not imply that skewness has no impact on the benchmark at all. When a distribution becomes more skewed, the mean will get dragged into the tail of the distribution. An increase in skewness will raise the mean, while a drop in skewness will reduce the mean. The fact that the means in figure 3 are more or less independent of the skewness levels therefore implies that the benchmark does produce lower (higher) prices for payoff distributions with lower (higher) skewness levels. It just does not show from the mean returns as the latter are heavily influenced by the skewness of the distribution. A less sensitive measure of location is formed by the median return. Actually, the difference between the mean and the median is often used as a rough measure of skewness. Figure 4 shows the benchmark in median-standard deviation-skewness space. From the graph we see a clear negative relationship between skewness and median return.

4. THE ARBITRAGE-FREE PRICING OF STANDARD DEVIATION AND SKEWNESS

An interesting by-product of the above benchmarking procedure is that it allows us to draw conclusions about the arbitrage-free pricing of a particular type of distribution's parameters. Having determined the appropriate benchmark for a particular type of distribution, it is straightforward to tell how much an extra unit of standard deviation

or skewness for example is worth in terms of mean or median return. Let's look at the example from the previous section again.

<< Insert Figure 5 and 6 >>

Under the same assumptions as before, figure 5 shows the highest attainable median return as a function of skewness for various levels of standard deviation. From the graph we see that in the absence of arbitrage there will be a negative relationship between skewness and median return, which becomes stronger as the standard deviation rises. With a standard deviation of 1%, the median drops 0.14% for every unit rise in skewness. With a standard deviation of 5% on the other hand, the median drops by 0.70%. For various levels of skewness, figure 6 shows the highest attainable median return as a function of standard deviation. The graph shows that the no-arbitrage relationship between standard deviation and median return is positive and grows stronger as skewness drops. For high skewness levels the relationship is almost flat. With skewness equal to 1.5 for example, the median return rises by only 0.04% for every 1% rise in standard deviation. With skewness equal to -1.5, however, it rises by 0.46%.

<< Insert Figure 7 >>

Obviously, when the distributions to be evaluated are normal, our procedure turns into a traditional Sharpe ratio analysis. To check this, we looked at the case of zero skewness. Figure 7 shows the relationship between the highest attainable mean return and the standard deviation, assuming the risk-free rate is at either 3,4, or 5%. This confirms that in the zero skewness case our procedure yields a benchmark that is equal to the standard capital market line.

5. CONCLUSION

In this paper we have discussed the extension of the traditional Sharpe ratio to allow for the evaluation of non-normal return distributions. Combining earlier work reported in Amin and Kat (2001) with stochastic simulation, we develop a procedure that allows us to construct a benchmark for the evaluation of the performance of funds with non-normal return distributions while maintaining the operational ease of the Sharpe ratio. Similar to the latter, our procedure only requires knowledge of the risk-free interest rate and the distribution of the market index, plus an assumption about the type of return distribution to be evaluated. Unlike the Sharpe ratio, however, we are not restricted to normality but are able to handle any reasonable type of return distribution.

Since our benchmarking procedure is based on the no-arbitrage assumption, it also provides insight into the arbitrage-free value of one distributional parameter in terms of another. Our results show that in case of the Johnson Su distribution the relationship between skewness and mean is more or less flat. Skewness and median on the other hand exhibit a strong negative relationship.

FOOTNOTES

1. As explained in more detail in Amin and Kat (2001), this assumption is based on the work of Cox and Leland (2000) and the payoff distribution pricing model of Dybvig (1988a, 1988b). It ensures that of all possible strategies that generate the desired payoff distribution we obtain the cheapest one.
2. Strictly speaking it is not possible to introduce more skewness without picking up some excess kurtosis as well. Table 1 contains the resulting kurtosis values given a skewness range from -1.5 to 1.5 . For ease of exposition, however, we will ignore kurtosis.

REFERENCES

Amin, G. and H. Kat (2001), Hedge Fund Performance 1990-2000: Do the Money Machines Really Add Value?, forthcoming *Journal of Financial and Quantitative Analysis*.

Black, F. and M. Scholes (1973), The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, Vol. 3, pp. 637-654.

Cox, J. and H. Leland (2000), On Dynamic Investment Strategies, *Journal of Economic Dynamics and Control*, Vol. 24, pp. 1859-1880.

Dybvig, P. (1988a), Distributional Analysis of Portfolio Choice, *Journal of Business*, Vol. 61, pp. 369-393.

Dybvig, P. (1988b), Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market, *Review of Financial Studies*, Vol. 1, pp. 67-88.

Jensen, M. (1968), The Performance of Mutual Funds in the Period 1945-1964, *Journal of Finance*, Vol. 23, pp. 389-416.

Johnson, N. (1949), Systems of Frequency Curves Generated by Methods of Translation, *Biometrika*, Vol 36, pp. 149-176

Johnson, N. (1965), Tables to Facilitate Su Frequency Curves, *Biometrika*, Vol. 52, pp. 547-558.

Sharpe, W. (1966), Mutual Fund Performance, *Journal of Business*, Vol. 39, pp. 119-138.

Table 1: Excess kurtosis accompanying different levels of skewness in Johnson Su distributions

Skewness	Excess Kurtosis
1.50	4.4
1.25	3.2
1.00	2.0
0.75	1.6
0.50	0.6
0.25	0.2
0.00	0.0
-0.25	0.2
-0.50	0.6
-0.75	1.6
-1.00	2.0
-1.25	3.2
-1.50	4.4

Figure 1: Example positively skewed Johnson Su distribution

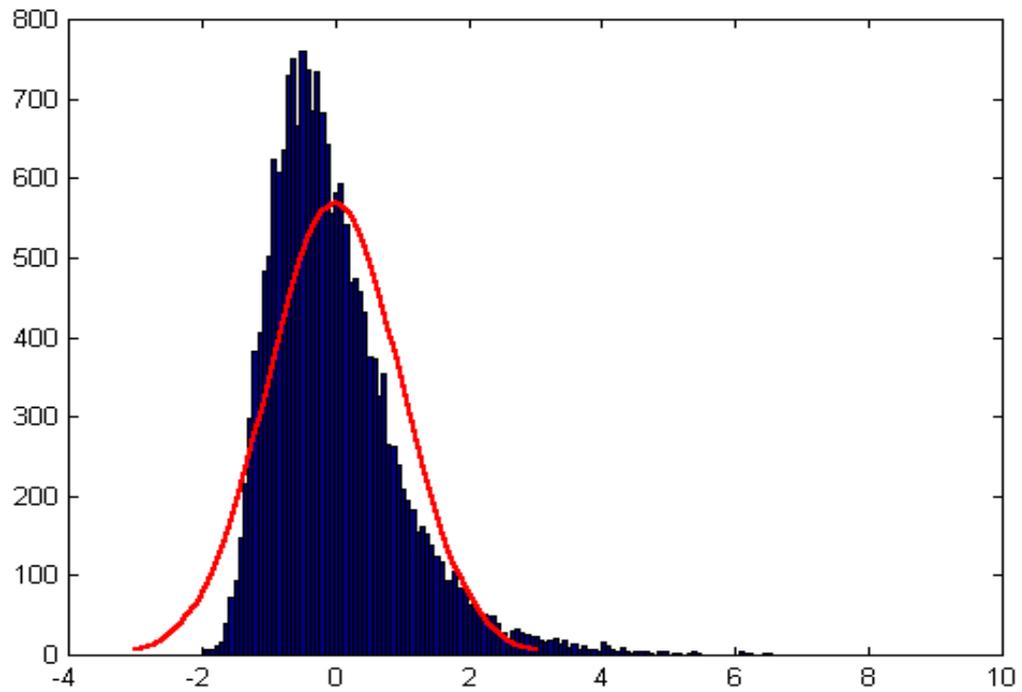


Figure 2: Example negatively skewed Johnson Su distribution

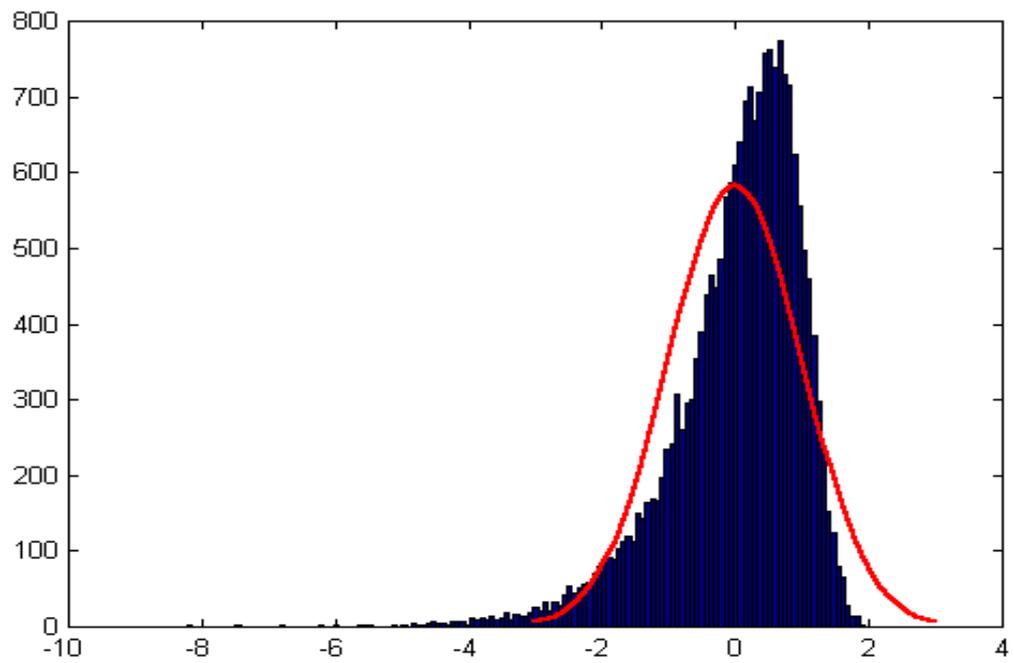


Figure 3: Highest attainable mean return for different combinations of standard deviation and skewness

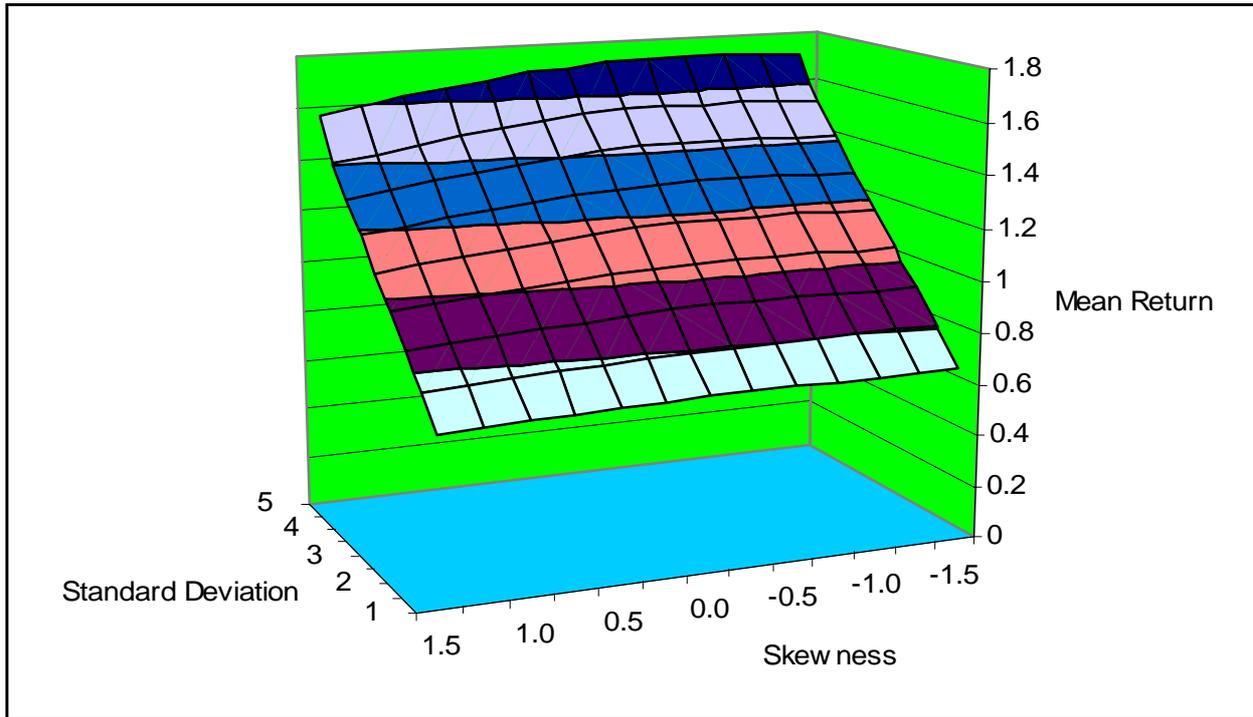


Figure 4: Highest attainable median return for different combinations of standard deviation and skewness

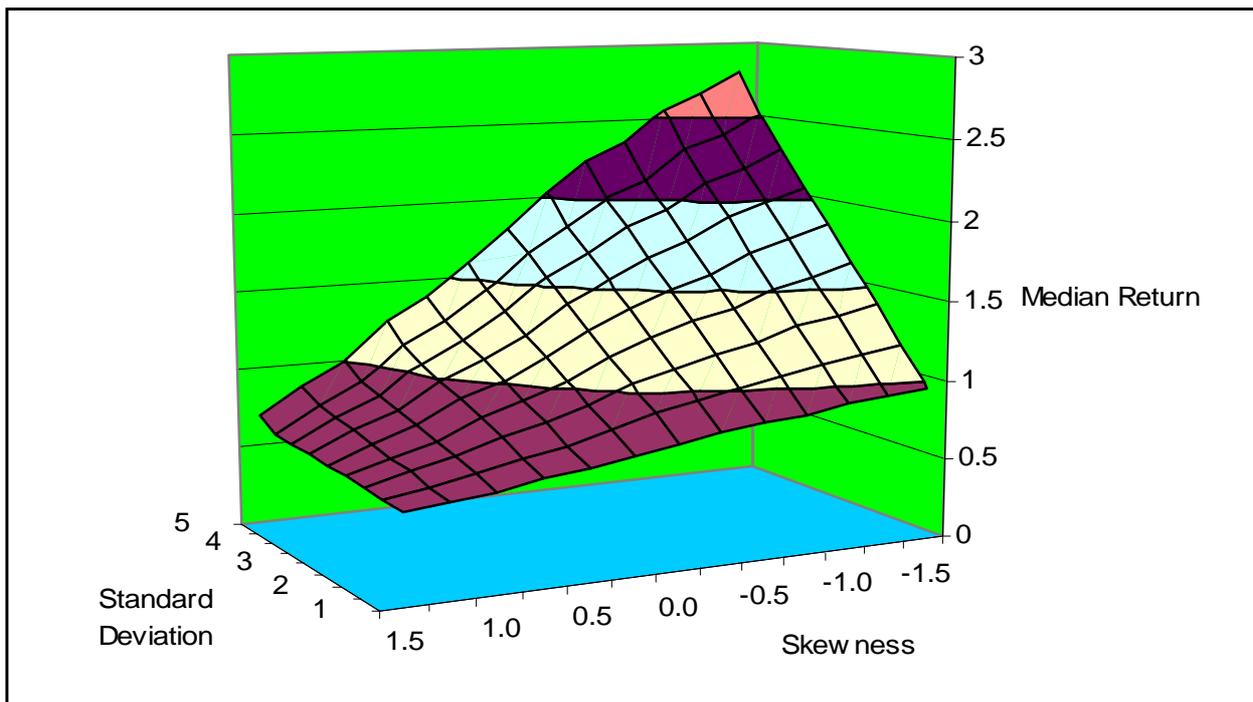


Figure 5: Highest attainable median return as function skewness for various levels of standard deviation

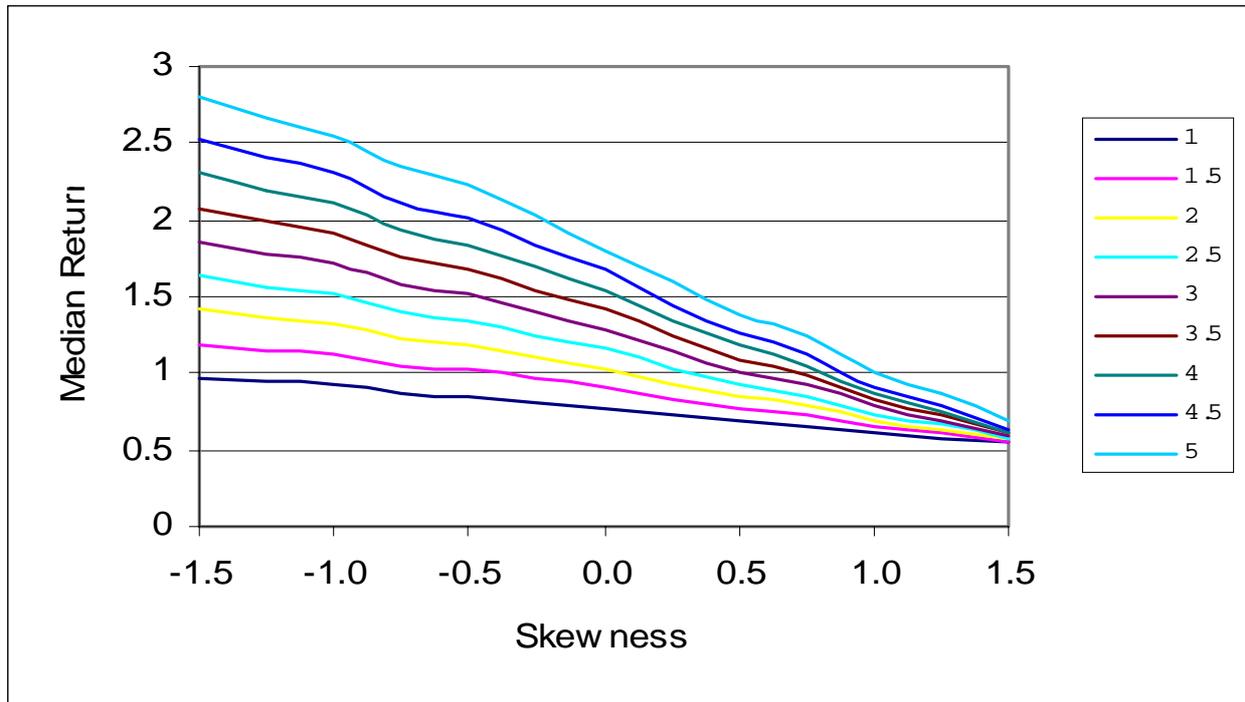


Figure 6: Highest attainable median return as function standard deviation for various levels of skewness

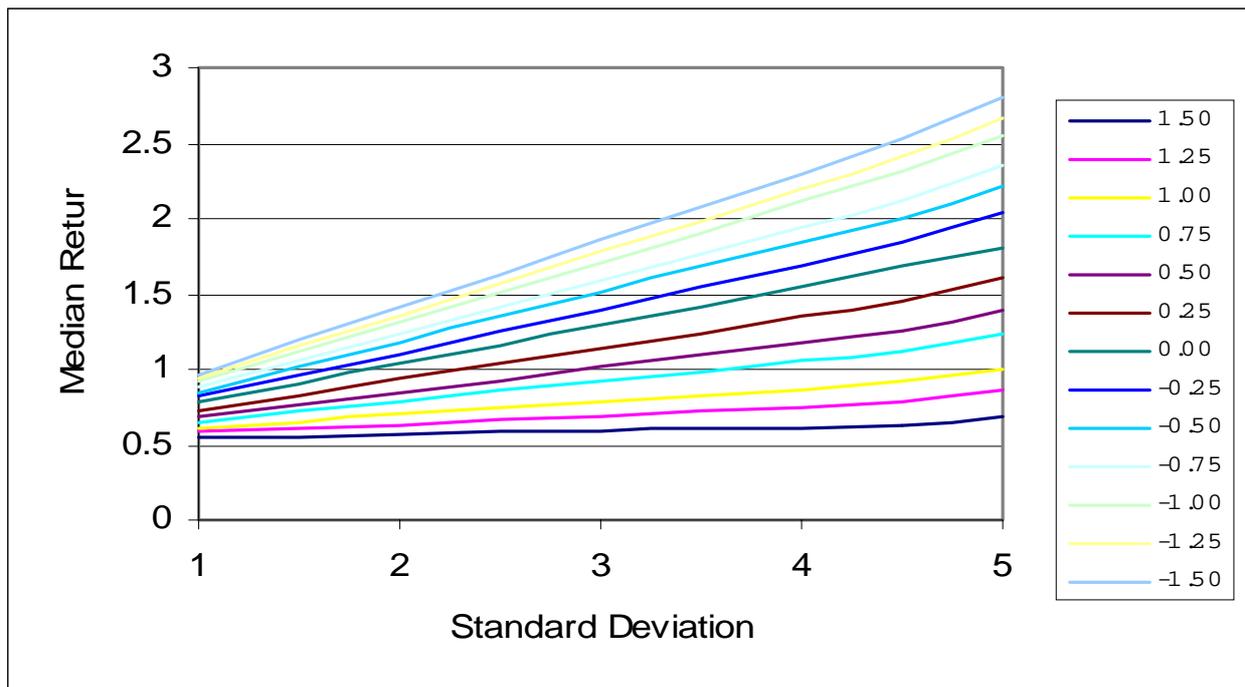


Figure 7: Highest attainable mean return as a function of standard deviation for 3 different interest rates and skewness fixed at zero

