FINANCIAL OPTIONS RESEARCH CENTRE

University of Warwick

A Generalization of the Sharpe Ratio and its Applications to Valuation Bounds and Risk Measures

Stewart Hodges

April 1998

Financial Options Research Centre
Warwick Business School
University of Warwick
Coventry
CV4 7AL
Phone (0)24 76 524118

FORC Preprint: 1998/88

A Generalization of the Sharpe Ratio and its Applications to Valuation Bounds and Risk Measures

Stewart Hodges

The Sharpe Ratio is a commonly used measure of the reward for risk given by investments. It has also been used recently by Cochrane and Saá Requejo to obtain "no good deal" bounds for derivative instruments in incomplete markets. Unfortunately, because it is based on the mean and variance alone, the Sharpe Ratio is incapable of providing satisfactory rankings of non-Gaussian distributions.

The paper describes these problems and develops a modified measure which overcomes them. It demonstrates how the new measure provides a more robust criterion for obtaining valuation bounds in incomplete markets, and that this approach can be employed in a wide variety of situations. The Generalized Sharpe Ratio can also be used to derive risk measures (e.g. for Value at Risk purposes) which are coherent in the sense described by Artzner, Delbaen, Eber and Heath.

A Generalization of the Sharpe Ratio and its Applications to Valuation Bounds and Risk Measures

1 Introduction

The Sharpe Ratio of the expected (or realized average) excess return divided by standard deviation of return is a convenient measure for assessing reward relative to risk. It was first introduced in 1965 and has been more recently discussed in Sharpe, 1994. Not surprisingly, it is commonly used in the financial community, particularly in the context of investment management. However, as is well known, the Sharpe Ratio (and indeed any measure based solely on the mean and variance) is incapable of providing satisfactory rankings of opportunities in the presence of arbitrary non-normal distributions. This disadvantage has become an increasingly awkward problem as the acceptance of derivatives has made it possible to engineer virtually any distribution at all.

The purpose of this paper is threefold. First it provides a generalization of the Sharpe Ratio which solves these problems. Second, it shows how this new measure provides an appropriate framework for deriving valuation bounds for derivatives in incomplete markets. Finally, it provides some further characterizations of these bounds and discusses their possible role as measures of risk (i.e. related to Value at Risk concepts).

The structure of the paper is as follows. First, an example is provided to remind the reader of the nature of the problems which arise with the conventional Sharpe Ratio. Next, the generalized measure is developed. In section 4, this is then applied to show how in an incomplete market the values of options (or other derivatives) can be bounded relative to the attractiveness of the market opportunities which would exist at various prices. This part of the paper relates to and extends existing work by Cochrane and Saa Requejo, 1996. The computation of these valuation bounds is discussed and an example is provided. Bounds can be developed in dynamic and static frameworks, including where the incompleteness arises from trading costs or jumps. Further, heuristics can also be used to obtain conservative bounds where exact computation of the embedded optimal control problem is too onerous. Section 5, develops further properties of our bounds. It is shown that they are closely related to the class of "coherent measures of risk" described by Artzner, Delbaen, Eber and Heath, 1997. In particular, used as risk measures, our bounds have all but one of the attractive properties advocated in that paper. Finally, the paper concludes with a general discussion of the nature of this measure and its advantages and shortcomings applied to performance measurement.

2 A Sharpe Ratio Paradox

It is well known that for general distributions the Sharpe Ratio leads to unsatisfactory "paradoxes" which make it unsuitable for ranking investment opportunities. The following simple example illustrates the nature of the problem. Table 1 provides two probability distributions, A and B, of percentage excess return over some horizon.

Table 1

Sharpe Ratio Paradox

Distribution A

X	-25	-15	-5	5	15	25	35
Pr	0.01	0.04	0.25	0.40	0.25	0.04	0.01

Mean: 5.00 **SD: 10.00** **SR: 0.500** **GSR: 0.498**

Distribution B

X	-25	-15	-5	5	15	25		45
Pr	0.01	0.04	0.25	0.40	0.25	0.04		0.01

Mean: 5.10 **SD: 10.34** **SR: 0.493** **GSR: 0.500**

Clearly distribution B should be preferred to distribution A: the outcome of +35 has been shifted to +45 which is a clear improvement under the weak assumption of non-satiation. Unfortunately, for our calculations, this shift increases the standard deviation by a larger percentage than it increases the mean of the distribution, so their ratio, the Sharpe Ratio, actually falls. Distribution B has a Sharpe Ratio of only 0.493, as against 0.500 for the original and inferior distribution A.

There is no mystery as to why this arises. For general distributions, quadratic utility is the only utility function consistent with making decisions based on mean and variance alone. Unfortunately, of course, quadratic utility functions display negative marginal utility for sufficiently high wealth levels. The rankings provided by the Sharpe Ratios are supported implicitly by a quadratic utility function which gives negative marginal utility to the higher outcomes.

Our generalization will avoid this problem. For normal distributions the conventional definition will be retained. Neither of the above distributions are normal: their Generalized Sharpe Ratios (GSR's) are 0.498 for distribution A and 0.500 for distribution B, providing a ranking which is consistent with decision making based on maximizing expected utility (and with first order stochastic dominance considerations). For the example shown above, the adjustment is fairly minor, however this consistency is absolutely essential for the applications we have in mind for valuing derivatives, calculating risk measures and measuring performance. The Generalized Sharpe ratio is developed and defined in the next section.

3 The Generalized Sharpe Ratio (GSR)

The generalization is based on two principles. First, we preserve the usual Sharpe Ratio for normal distributions, but for other distributions we provide a generalization which is free of the paradoxes associated with the Sharpe Ratio. Second, the generalization is based on the expected utility to investors with constant absolute risk aversion. Other generalizations are possible, but this one has certain advantages. In particular, the measure obtained (like the original Sharpe Ratio) does not depend on any particular assumption of risk aversion or wealth level, but instead, simply describes a measure of the extent of market opportunities (or realizations).

Consider an investor who maximises $E[U(\tilde{w})]$ with particular $U = -e^{-\lambda w}$, (for some $\lambda > 0$ which governs the absolute risk aversion). For simplicity let us assume that the investor starts with initial wealth, $w_0 = 0$, so no investment gives $U = -1$. However, favourable investment opportunities can lead to an optimal expected utility U^* anywhere within [-1, 0). We will examine first the well known case where the investor faces a single opportunity whose outcomes after time T are normally distributed with (forward annualized) mean excess return μ and standard deviation σ.

In this situation the Sharpe Ratio is of course $SR = \mu/\sigma$. We will redefine it in terms of the optimal expected utility so that it gives the usual measure for this case of Normal distributions, and something more consistent for other distributions.

Our investor's forward investment opportunity for date T has future outcomes distributed as $N(\mu T, \sigma^2 T)$. To maximize $E[U(w)]$ with $U = -e^{-\lambda w}$ the choice problem is:

$$\underset{x}{Max}\ E[U] = -\exp\{-\lambda(\mu x T - \tfrac{1}{2}\lambda\sigma^2 T x^2)\}$$

The first order condition is $\mu T - \lambda\sigma^2 T x = 0$, so $x = \dfrac{\mu}{\lambda\sigma^2}$.

Substituting into the expression for $E[U]$, we find

$$U^* = \underset{x}{Max}\ E[U] = -\exp\{-\tfrac{1}{2}\frac{\mu^2}{\sigma^2}T\}$$

From this equation we can down write an expression for the Sharpe Ratio μ/σ in terms of the optimal utility U^* and use this to define a Generalized Sharpe Ratio, to be used instead of the usual measure whenever distributions are non-normal.

The Sharpe Ratio $\dfrac{\mu}{\sigma}$ is obtained as $\sqrt{\dfrac{-2}{T}ln(-U^*)}$.

We are now in a position to define the generalized measure:

Definition Generalized Sharpe Ratio

We define the Generalised Sharpe Ratio (GSR) as a measure of market opportunities as:

$$GSR = \sqrt{\frac{-2}{T}ln(-U^*)}$$

where U^* is the optimal expected utility obtainable by an investor with $U = -e^{-w}$, and initial wealth, $w_0 = 0$. T is the length of the investment horizon in years.

Discussion

It is worth noting that the optimal utility U^* is independent of the original investor's risk aversion λ, so we may as well define $\lambda = 1$ as in the above definition. The level of risk aversion is irrelevant since the investor is considered free to choose an desired level of investment (or sale). We have already proved that in the case of an investment opportunity with a Normal distribution of outcomes, the Generalized Sharpe Ratio (GSR) equates to the usual definition of the Sharpe Ratio (SR) as μ/σ.

The new measure provides rankings which are always consistent with stochastic dominance considerations (since the utility function has derivatives which alternate in sign). For the definition of the Sharpe Ratio we know that the equity premium is of the order of 0.5 (e.g. 8% risk premium over 16% standard deviation). With Normal distributions a Sharpe Ratio of 3 is almost a pure arbitrage with virtually no probability of loss. The following simple example demonstrates some behaviour of the GSR measure.

An Example:

Consider the GSR provided by the forward gamble of +£1 with probability p and -£1 with probability q:

$$U = -pe^{-x} - qe^{x}$$

$$\therefore \frac{dU}{dx} = pe^{-x} - qe^{x} = 0$$

$$\therefore \ln\left(\frac{p}{q}\right) - x = x \text{ and } x = \frac{1}{2}\ln\left(\frac{p}{q}\right) = \ln\sqrt{\frac{p}{q}}$$

$$\therefore U^* = -p\sqrt{\frac{q}{p}} - q\sqrt{\frac{p}{q}} = -2\sqrt{pq}$$

$\text{GSR} = \sqrt{-2\ln(-U^*)} = \sqrt{-2\ln\left(2\sqrt{pq}\right)}$ (ignoring the time dimension of the gamble).

In comparison the Sharpe Ratio for this situation is given as

$$SR = \frac{|p-q|}{2\sqrt{pq}}$$

Figure 1 shows a graph of GSR against p. For p close to 0.5 this is a small bet and the GSR tracks very closely to the usual SR. For $p = 0$ and $p = 1$ the GSR curves upwards asymptotically towards infinity, as a sure thing at the wrong price would be a pure arbitrage. However, is notable how close to these extreme values most of the curvature occurs. The GSR reaches 1 at $p = 0.90$ and 0.10, it reaches 2 at $p = 0.995$ and 0.005, while for GSR = 3 we require $p = 0.99997$ or 0.00003. It should be appreciated that these values are for a single gamble. A *repeated* gamble of this kind would be much more attractive for p close to 0.5.

4 Valuation Bounds in Incomplete Markets

Cochrane and Saa Requejo, 1996, introduced the idea of using the Sharpe Ratio to characterize valuation bounds on derivatives in incomplete markets. They used a pricing function with linear marginal utility, which was subsequently constrained to take zero values wherever it would otherwise have been negative, in order to avoid the problems we have seen associated with the conventional Sharpe Ratio. Our approach provides a neater solution to the same problem. First, the general philosophy is described, and then the details of the method are developed, with an example.

Even where exact replication of derivatives is impossible, we may have an idea that one price for a given contingent claim is "cheap" and that another is "dear". We introduce a formulation whereby bounds on option prices are established according to the reward for risk opportunities which they introduce. Under exact replication there is a unique market price, C_0, for a given claim. Even at a price, $C_0-\varepsilon$, an infinite reward for risk opportunity (theoretically) exists to buy the claim and hedge it for an exact profit of ε. Similarly at $C_0+\varepsilon$ an infinite reward for risk opportunity exists for selling the claim and hedging it.

We propose a valuation approach in which optimal choices are made under the standard expected utility of wealth paradigm. The enhancement of utility through being able to buy or sell the claim leads to specific bounds for the option price relative to various reward for risk levels.

In particular we may compute price levels which give particular values of the Generalized Sharpe Ratio (such as ½ or 1) from either buying or selling the option, as well as no-arbitrage bounds and a "no gain" level at which there is no incentive to either buy or sell.

We describe the approach first in the context of a statically hedged position, and where the analysis is done under a risk neutral measure. We then subsequently describe extensions to dynamic hedging and to incorporate market risk premia for assets other than the one of interest.

In general we proceed in two stages:

A: Find the best dynamic strategy with hedging securities alone. This gives a benchmark GSR which will be bettered when the claim is added.

If we model the hedging securities under any risk neutral measure this benchmark GSR will be exactly zero and this whole stage can be skipped.

B: For each possible price C_0 for the claim, find the optimal strategy and expected utility from either buying or selling a fixed fraction y of the claim and dynamically hedging it using the available securities. Interpret this as the GSR for that price C_0. Finally, invert to obtain option valuation bounds as a function of the available Sharpe Ratio. For simplicity, all prices can be taken as forward prices.

Valuation Bounds Under Static Hedging

We solve the choice problem for an investor who maximizes $E[U(w)]$ with $U = -e^{-w}$.

The investor buys y units of the contingent claim, and hedges with x units of the underlying:

$$\underset{x,y}{\text{Maximize}}\, E[U] = -E[e^{-x(S_T - S_0) - y(C_T - C_0)}]$$

The value of the expected utility provides a measure of the "goodness" of any particular C_0. We will generalize to dynamic hedging later. More specifically we can define a lower bound, *LB*, and an upper bound, *UB*, as the prices at which we can respectively buy and hedge or sell and hedge so as to achieve a specified GSR or U^* level. Thus LB and UB are defined by the following equations:

$$\underset{y>0,x}{\text{Maximize}}\, E[U] = -E[e^{-x(S_T-S_0)-y(C_T-LB)}] = U^*, \text{ defines } LB.$$

$$\underset{y>0,x}{\text{Maximize}}\, E[U] = -E[e^{-x(S_T-S_0)-y(UB-C_T)}] = U^*, \text{ defines } UB.$$

Computational Considerations

Whereas to compute the GSR for a general distribution simply involves a one-dimensional search for the optimal scale in which to invest, the GSR which can be obtained from a derivative or portfolio of derivatives depends on how the portfolio is hedged. In the single asset static hedge situation we have to optimize over x, the hedge quantity, and y, the quantity of the derivative bought or sold. The first order conditions for the optimization are:

$$E[(S_T - S_0)\exp\{-x(S_T - S_0) - y\,C_T\}] = 0,$$
$$E[(C_T - LB)\exp\{-x(S_T - S_0) - y\,C_T\}] = 0.$$

If we first choose a value of y, then these equations can be solved sequentially. The first equation involves a one-dimensional search for x. Since the first derivative is available, it can be accelerated by the Newton-Raphson method. Given x and the assumed value of y, the second equation provides the corresponding LB value as:

$$LB = \frac{E[C_T \exp\{-x(S_T - S_0) - y\,C_T\}]}{E[\exp\{-x(S_T - S_0) - y\,C_T\}]}.$$

Finally, substituting back into the original maximization we obtain U^* and GSR.

An Example

Figures 2 and 3 show the nature of these bounds for a single call option with a strike price of 95 and one year to expiry. The assumptions made are that we live in a (risk neutral) Black-Scholes world with a zero interest rate, and that the asset has volatility equal to 0.15. The option is only hedged at the beginning and not dynamically rebalanced. Figure 2 shows how the GSR depends on price of the option when the asset price is 100. At any option price below 5 a pure arbitrage exists, and the GSR curves up asymptotically towards this. The no-arbitrage upper bound is a long way away at 100. Figure 3 shows valuation bounds for at variety of asset prices, and under different assumptions. The two outer curves show prices which give a GSR of ½ if the option is bought or sold outright and not hedged at all. This is done simply by restricting $x = 0$. The curves inside these show upper and lower bounds based on a single static hedge to give a GSR of ½. The graph also shows the Black-Scholes values (the centre curve) and the no-arbitrage lower bound (piece-wise linear). Note that this is violated by the no-hedge lower bound but not by the one using a static hedge.

Risk Premia

The calculations illustrated above were made under the assumption that the risk premium on the underlying asset is zero. This assumption is unnecessary. We can use the following two-stage procedure. First optimize over x to obtain a benchmark level of expected utility, U_0. For a given incremental GSR, the required U^* value is given as

$$U^* = U_0 \exp\{-½\frac{\mu^2}{\sigma^2}T\}$$

The proof of this result is straightforward.

Dynamic Hedging

The approach generalizes to dynamic hedging situations in incomplete markets.
As before we solve the choice problem for an investor who maximizes $E[U(w)]$ with $U = -e^{-w}$.
The investor buys y units of the contingent claim, and hedges dynamically with x units of the underlying:

$$\underset{x,y}{\text{Maximize}}\, E[U] = -E[e^{-\left\{\int_0^T x_t \,.dS_t + y(C_T - C_0)\right\}}]$$

Again the value of the expected utility provides a measure of the "goodness" of any particular C_0. In principle x_t and dS_t may be vectors, corresponding to dynamic hedging with several instruments. The lower bound, LB, and an upper bound, UB, are still the prices at which we can respectively buy and hedge or sell and hedge so as to achieve a specified GSR or U^* level. The only difference from the previous static case is that we have a stochastic optimal control problem to solve. The nature of this problem and its numerical complexity depends on the specific nature of the processes involved and the reason for market incompleteness.

In the case of Black-Scholes assumptions but with rebalancing at discrete dates the problem amounts to a dynamic programming formulation with only a single state variable. This is quite easily solved, and it is clear that the approach described here generalizes the Black-Scholes replication method. We will obtain upper and lower bounds which tend towards the Black-Scholes limit as the revision dates become closer together. There are a few other relevant processes which also lead to single state variable dynamic programming formulations. One important case is the jump-diffusion process, which will not provide a unique value as trading becomes continuous.

Other situations may result in problems with two or more state variable. For example, introducing stochastic volatility introduces a second state variable. Transactions costs are a further significant cause of incompleteness. Work by Hodges and Neuberger, 1989, already introduced a formulation of this problem under negative exponential utility and it is only a question of reinterpreting the results of that paper (and subsequent developments to it) in the light of the framework presented here.

In computing numerical solutions a certain amount of care is required, as the exponential form of the utility function makes it quite easy to get overflow or underflow kinds of computational errors. Nevertheless, the formulation is intrinsically robust in the sense that because our $E[U]$ is bounded, the methodology is not subject to generalized St. Petersburg paradoxes.

Finally, there are many models of hedging in incomplete markets where the computation of exact optimal hedging strategies is too daunting to be feasible. Nevertheless, the framework introduced here is helpful even in these situations. Whatever heuristics may be employed to hedge with, the hedge performance can be simulated and conservative valuation bounds can be calculated from the resulting distributions. The better the heuristic you can come up with, the tighter the valuation bound you are entitled to apply!

5 Properties of the Measures

The bounds we can compute relative to a given GSR, have a close relationship to the idea of "coherent risk measures" described in a recent paper by Artzener, Delbaen, Eber and Heath, 1997. That paper advocates that for purposes such as calculating margin requirements or Value at Risk for bank capital requirements there are various desirable properties that we would like to have. They state these properties in an axiomatic fashion and derive a characterization of risk measures which possess them. In this section, we will describe how our bounds also qualify a risk measures which satisfy all but one of their axioms. We argue that there is no need to satisfy the remaining one, and we also look at the resulting characterization.

First, our lower bound, *LB*, was derived on the principle that if we could buy at that price we could achieve a specified GSR or U^* level. If our U^* level is chosen sufficiently high enough, then *LB* becomes an appropriate candidate for a measure of how bad our portfolio might become. In fact, for our choice of utility function, we can demonstrate a probabilistic interpretation for U^* which gives us a bound on the quantile of the distribution. It is easily shown that

$$\Pr\{C_T < LB\} \leq -U^*$$

For other utility functions too, we could in principle obtain similar risk measures. To do so, the lower bound would be defined by the equation:

$$\underset{y>0}{Max}\ E[U(w_0 + y(\tilde{C}_T - LB)] = U^*$$

We can describe and demonstrate most of the properties advocated by Artzener et al, 1997, in terms of this slightly broader category of risk measures than those generated by negative exponential utility:

"Linearity": $B[\alpha\,\tilde{C}] = \alpha\, B[\tilde{C}]$, and

$$B[\beta + \tilde{C}] = \beta + B[\tilde{C}]$$

The meaning of these "linearity" properties, which hold for both upper and lower bounds, *B*, is as follows. First that valuations can be scaled linearly. If we double all the cash flows of the claim we double any valuation bound computed for it. This follows for our bounds as a simple consequence of the freedom to choose (linearly) the quantity to buy or sell, and also to hedge with. Thus if y^* solves the problem for C, then y = y/α solves it for αC, and it follows immediately that $LB(\alpha C) = \alpha\, LB(C)$.

The second property is that packaging with a riskless cash flow is transparent. Again, the algebra to establish this is straight forward. Suppose

$$E[U(w_0 + y^*(\tilde{C}_T - LB)] = U^* \text{ , then}$$
$$E[U(w_0 + y^*(\beta + \tilde{C}_T - \beta - LB\ )] = U^*$$

showing that $LB(\beta{+}C) = \beta + LB(C)$.

Monotonicity:

$$\tilde{C} \le \tilde{D} \Rightarrow LB[\tilde{C}] \le LB[\tilde{D}]$$
$$-\tilde{D} \le -\tilde{C} \Rightarrow UB[\tilde{C}] \le UB[\tilde{D}]$$

This property is also established rather easily. Suppose that L is the lower bound of C which is stochastically dominated by D, in some sense of stochastic dominance which is satisfied by our utility function. Then the expected utility given by D at the y level and bound L established for C is at least equal to U^* and possibly higher. In other words:

$$E[U(w_0 + y_C{}^*(\tilde{D}_T - L)] \ge U^*.$$

Since marginal utility is positive it follows that LB[C] $\le LB$[D].

Subadditivity:

$$LB[\tilde{C}] + LB[\tilde{D}] \le LB[\tilde{C} + \tilde{D}]$$
$$UB[\tilde{C} + \tilde{D}] \le UB[\tilde{C}] + UB[\tilde{D}]$$

The idea here is that the bounds obtained from portfolios are tighter than the sums of the bounds from their components. Although this is a simple and intuitive notion, the proof is slightly longer than for the other properties. Suppose:

$$E[U(w_C)] = E[U(w_0 + y_C{}^*(\tilde{C}_T - LB_C)] = U^*, \text{ and}$$
$$E[U(w_D)] = E[U(w_0 + y_D{}^*(\tilde{D}_T - LB_D)] = U^*$$

We know that U is concave, so $E[U(\theta_C w_C + \theta_D w_D)] \ge U^*$,
for any θ_C, $\theta_D \ge 0$ with $\theta_C + \theta_D = 1$.
The result can now be established by choosing $\theta_C = y_D/(y_C + y_D)$ and using the monotonicity property.

Characterization

Artzener et al draw attention to a feature of risk measures which satisfy their axioms, that (in our notation)

$$LB[\tilde{C}] = \inf_{P \in \rho} E_P[\tilde{C}]$$

(In their case attention is confined to losses only, and the expectation is only over losses, also since they considering "maximum loss scenarios" they have a supremum.)

This infimum over a set of probability measures also arises naturally in our analysis. We have already seen that the first order condition for the quantity y provides the equation:

$$LB = \frac{E[C_T \exp\{-y\,C_T\}]}{E[\exp\{-y\,C_T\}]}.$$

As usual, the marginal utilities provide a transformation of the objective probability measure to a new "risk neutral" one. We can show that under our model the lower bound is obtained as the infimum of $E_P[C]$ over probability measures which take this form and arise from situations where $E[U] \leq U^*$. The intuition for this is that for low levels of expected utility y is close to zero and LB is close to $E[C]$. We can only get measures which take LB further into the tail of the distribution at higher values of both y and $E[U]$.

Use for Performance Evaluation

The fact that *ex post* distributions differ from *ex ante* ones means that no single measure provides a completely satisfactory measure of performance, and the Generalized Sharpe Ratio is not advocated here as a complete answer. Nevertheless, it is immediately apparant that it avoids some of the inconsistencies which can occur with the usual version. It is also robust to "gaming" in a continuous time world where the price of risk μ/σ is constant. In such a world an investor with negative exponential utility will always bet the same amount on the market and will obtain a Normal distribution of return. In this case the best attainable Generalized Sharpe Ratio (*ex ante*) is equal to the conventional *ex ante* Sharpe Ratio.

6 Summary and Conclusions

This paper has shown three things. First it has provided a generalization of the Sharpe Ratio which gives coherent rankings of probability distributions which are consistent with stochastic dominance rankings. Second, it has demonstrated how this new measure provides an appropriate framework for deriving valuation bounds for derivatives in incomplete markets. Finally, it has given some further characterizations of these bounds and discussed their appropriateness as measures of risk (i.e. related to Value at Risk concepts).

The Sharpe Ratio is much used in the context of performance measurement, and it is therefore sensible to ask whether the Generalized Sharpe Ratio is equally applicable. Certainly, it avoids the immediate problems which may arise when outcomes are clearly come from non-normal distributions and comparison in a mean-variance framework is inappropriate. However, it must be pointed out that the GSR is no panacea. In most cases where large discrepancies from Normality are encountered empirical work we have no guarantee that we have sampled the entire distribution; even if we have a very large there may be no reason to expect the distribution to be stationary. It therefore seems that the GSR is likely to be a much more powerful and appropriate tool for the analysis of *ex-ante* distributions than for the analysis of *ex-post* returns.

References

Artzner P, F Delbaen, J Eber, D Heath, 1997, "Definition of Coherent Measures of Risk", Working Paper, March 1997, Cornell University.

Bouchard J-P, and D Sornette, 1994, "The Black-Scholes Option Pricing Problem in Mathematical Finance: Generalization and Extensions For a Large Class of Stochastic Processes", *Journal de Physique I, France*, 4, 863-881.

Cochrane J H and J Saá Requejo, 1996, "Beyond Arbitrage: 'Good-Deal' Asset Price Bounds in Incomplete Markets" Working Paper February 1996, Graduate School of Business, University of Chicago.

Hodges S D and A Neuberger, 1989, "Optimal Replication of Contingent Claims Under Transactions Costs", *Review of Futures Markets*, 8, 222-239.

Sharpe, W F, 1994, "The Sharpe Ratio", *Journal of Portfolio Management*, 21, 49-59.

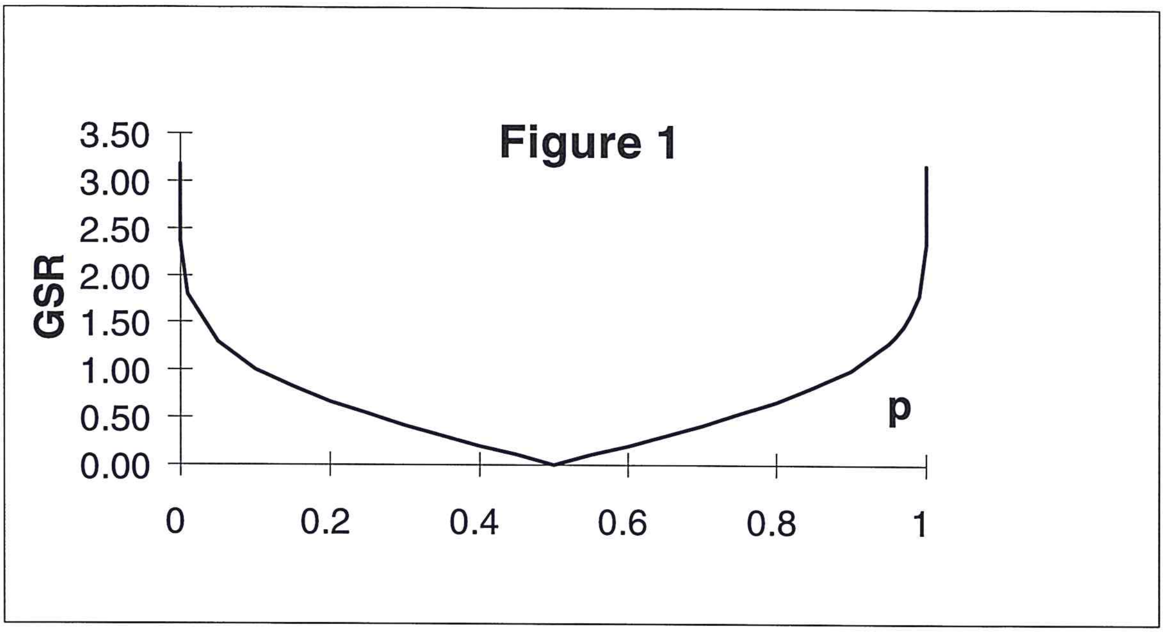
Figure 1
GSR
3.50
3.00
2.50
2.00
1.50
1.00
0.50
0.00
p
0
0.2
0.4
0.6
0.8
1

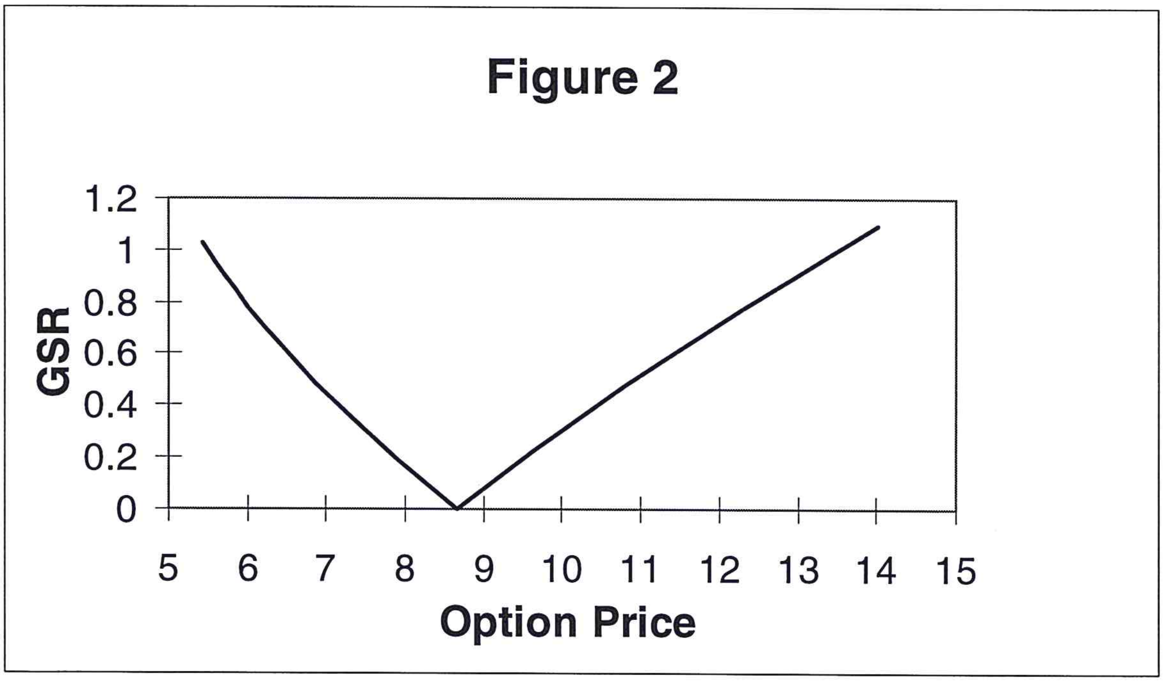
Figure 2
GSR
1.2
1
0.8
0.6
0.4
0.2
0
5
6
7
8
9
10
11
12
13
14
15
Option Price

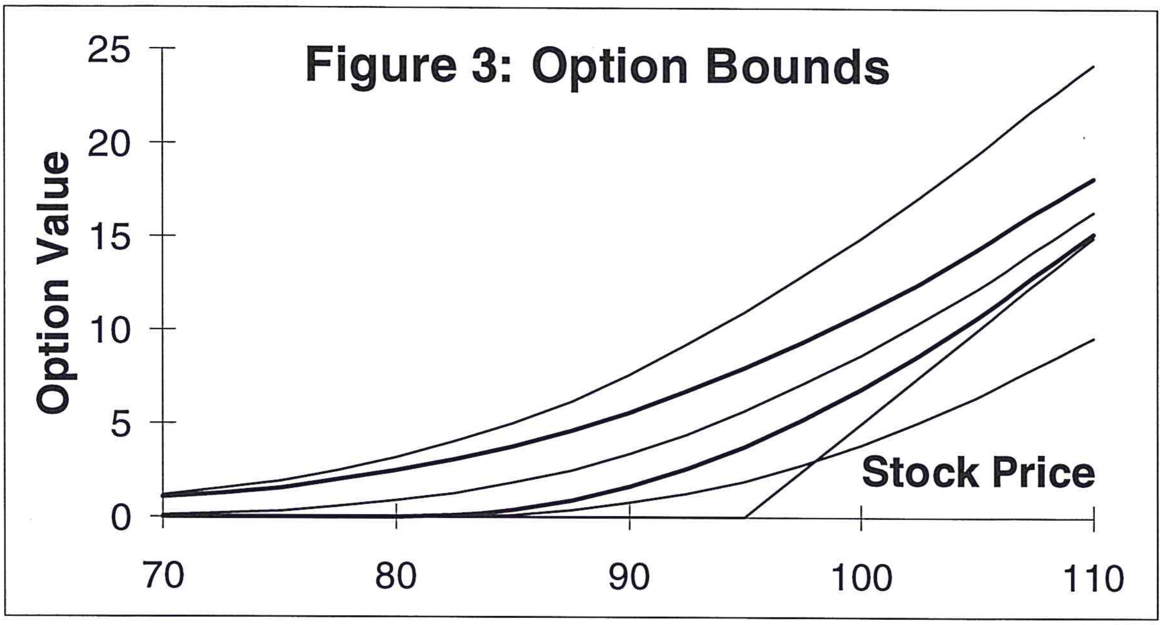
Figure 3: Option Bounds
Option Value
Stock Price
25
20
15
10
5
0
70
80
90
100
110