ACTUARIAL RESEARCH CLEARING HOUSE
1997 VOL. 1

On Expert Use in Portfolio Management [1]

CHRISTIAN GENEST, Ph. D.
Professeur de statistique
Département de mathématiques et de statistique
Université Laval
Sainte-Foy (Québec)
Canada G1K 7P4

MICHEL GENDRON, Ph. D., A. S. A.
Directeur de la chaire en assurance
Département de finance et d'assurance
Université Laval
Sainte-Foy (Québec)
Canada G1K 7P4

This note is concerned with the evaluation of the common portfolio management strategy which consists of delegating to different managers portions of the assets to be invested. This practice is shown to be close to optimal under the realistic assumption that the private signals of the financial experts are positively correlated. This is accomplished by comparing this multi-manager approach to an optimal investment policy derived under full information, within the context of a standard normal returns / exponential utility market timing model.

Key Words: asymmetric information, opinion pool, portfolio management

1. Introduction

A common risk-spreading measure for large funds consists of delegating parts of the assets to several managers, each with a particular investment strategy or style, to produce an aggregate portfolio with diversification and characteristics that suit the owner's needs. As well established as this practice may be, there are financial analysts (e.g., Jeffrey, 1991) who contend that the desired diversification might be obtained more effectively and efficiently if a single manager were charged with an overall set of objectives. This approach could conceivably lead to a superior investment strategy if this manager's knowledge base encompassed that of the other agents. For, even leaving aside the fact that it is more costly to hire several managers than just one, an obvious weakness of the multi-manager approach is that it fails to exploit the communalities and dependencies between their information sets.

The purpose of this note is to measure the loss of efficiency encurred by standard multi-manager portfolio investment practices. This will be accomplished by comparing the expected utility of the latter approach with an investment strategy that would be optimal under full information. This analysis rests on a multivariate version of the standard normal returns / exponential utility market timing model. As will be shown, an approach involving a single, fully informed manager would indeed be uniformly superior, but only marginally so, under realistic conditions. This is essentially because standard practice turns out to be equivalent to the investment policy of a composite manager whose opinion would be a (weighted geometric) mean of the agents' beliefs.

[1]Partial funding in support of this work was provided by the Natural Sciences and Engineering Research Council of Canada, by the Social Sciences and Humanities Research Council of Canada, as well as by the Fonds pour la formation de chercheurs et l'aide à la recherche du Gouvernement du Québec.

A description of the differential information framework used in this note is given in Section 2, and the optimal investment strategy under full information is identified in Section 3. This ideal strategy is then compared with the multi-manager approach in Section 4, using the expected utility criterion. It will be seen, among other things, that the current practice which consists of delegating parts of a portfolio to different managers is close to optimal under the realistic assumption of positively-correlated signals. Section 5 concludes briefly.

2. The Model

The differential information framework used in this note is a straightforward generalisation of the standard normal returns / exponential utility market timing model exploited by Gendron & Genest (1990), among others, to evaluate the performance of a single portfolio manager facing investment constraints. In this one-period, partial equilibrium model, the market return in excess of the risk-free rate is assumed *ex ante* to be normally distributed with mean μ and variance σ^2, denoted

$$R_m \sim N(\mu, \sigma^2).$$

To portray the diversity of opinions, it is supposed that individual i in a set of $n \geq 2$ financial experts has the ability to observe a private signal Y_i which, given a realisation $R_m = r_m$ of the market return, also follows a normal distribution with mean r_m and variance σ_i^2. Possible dependence between these signals is accounted for by the assumption that conditionally on r_m, the column vector $Y = (Y_1, \ldots, Y_n)'$ follows a multivariate normal distribution with covariance matrix Σ, namely

$$Y | R_m = r_m \sim N(r_m e, \Sigma),$$

where e stands for an $n \times 1$ vector of ones.

Upon observing his/her signal $Y_i = y_i$, each expert uses Bayes' Theorem to update his/her subjective distribution for R_m. A standard calculation shows that his/her posterior distribution for R_m is

$$R_m | Y_i = y_i \sim N(\mu_i, \tau_i^2),$$

where $\mu_i = \tau_i^2(\mu/\sigma^2 + y_i/\sigma_i^2)$ and $1/\tau_i^2 = 1/\sigma^2 + 1/\sigma_i^2$. It is this posterior distribution which, upon being hired by an investor to manage his/her portfolio, expert i would use to determine the proportion X_i of assets that should be invested in the market. This proportion is obtained by maximizing the investor's end-of-period utility of portfolio return, $R_p = X_i R_m$.

Assuming an exponential utility function with risk aversion parameter $\theta > 0$, the solution to the maximization problem

$$X_i = \text{argmax } E(-e^{-\theta X_i R_m} | Y_i = y_i)$$

is unique and given by

$$X_i = \frac{E(R_m | Y_i = y_i)}{\theta var(R_m | Y_i = y_i)} = \frac{\mu}{\theta \sigma^2} + \frac{1}{\theta \sigma_i^2} y_i,$$

which shows that for expert i, the investment policy is a linear function of his/her private signal. In particular, this manager would shortsell the market if $X_i < 0$ and would borrow to invest if $X_i > 1$.

Now suppose that an investor hires $n \geq 2$ financial experts, each of whom is entrusted with a portion of the assets to be managed. Let $0 \leq w_i \leq 1$ represent the fraction of these funds allocated to manager $1 \leq i \leq n$, so that $\sum_{i=1}^{n} w_i = 1$ by definition. Care should be exerted to distinguish w_i from the proportion X_i that manager i invests in the market on behalf of the fund's owner. While X_i may be negative or greater than one to reflect the fact the manager may choose to shortshell or borrow to invest, his/her weight w_i is necessarily comprised between 0 and 1, as it merely represents the fraction of the investor's assets that this expert was entrusted with.

The market share of this multi-manager investment policy is given by

$$\bar{X} = \sum_{i=1}^{n} w_i X_i.$$

It amounts to investing in the market a weighted linear average of the proportions X_i of assets chosen by the managers. While this strategy stems from a portfolio delegation mechanism, it is akin to a linear combination of forecasts, a technique whose value has been demonstrated again and again in a variety of economic contexts (see Clemen, 1989, for a review). The main difficulty associated with this procedure is the determination of each expert's share of the funds under management (Winkler & Clemen, 1992). An appealing rule would be to allocate the assets proportionally to the level of "expertise" of the managers, in some sense of the word. In practice, absence of discriminating information in this regard would reasonably lead to equal weights. The issue of weight selection will be revisited in Section 4.

The investment strategy $\bar{X}$ can also be justified via an opinion pooling argument. Assuming that an investor could have access to the experts' posterior distributions $f_1, \ldots, f_n$ for R_m, he/she could then form his/her own opinion by taking a weighted geometric mean of these density functions. The resulting distribution,

$$f(r_m) = \prod_{i=1}^{n} f_i^{w_i}(r_m) / \int \prod_{i=1}^{n} f_i^{w_i}(r)dr,$$

is referred to as a logarithmic opinion pool in the statistical literature (Genest & Zidek, 1986). In the special case where each of the f_i's is normal with mean μ_i and variance τ_i^2 as defined above, it is easy to see that f is itself normal with mean $\bar{\mu}$ and variance $\bar{\tau}^2$ satisfying

$$\bar{\mu} = \bar{\tau}^2 \left(\sum_{i=1}^{n} w_i \mu_i / \tau_i^2 \right), 1/\bar{\tau}^2 = \sum_{i=1}^{n} w_i / \tau_i^2,$$

so that, given f, the proportion of assets that should be invested in the market is $\bar{X} = \bar{\mu}/\theta\bar{\tau}^2 = \sum_{i=1}^{n} w_i X_i$.

Weighted geometric averages of expert opinions are externally Bayesian in the sense of Madansky (1978). In short, this means that once f has been computed, an investor no longer needs to refer to the experts to update it when new, common knowledge evidence relevant to the market return becomes available. Because of the way in which f was obtained, its updating via Bayes' Theorem would be equivalent to the weighted geometric average of the experts' updated distributions. This property is valuable to investors in that recourse to the experts is superfluous when new public information is released.

3. The Full Information Investment Strategy

Suppose that an investor had direct access to the experts' private signals $Y_i = y_i$, or that he/she could infer them from their posterior distributions. Further assume that the investor knew the dependence structure between these signals, as embodied in the conditional distribution of the column vector $Y = (Y_1, ..., Y_n)'$ given $R_m = r_m$. If the prior distribution of the market return is common knowledge, he/she would then have in hand all the data required to make a decision based on the posterior distribution of R_m given $Y = y$. The solution to the maximization problem

$$X_{full} = \text{argmax} E(-e^{-\theta X_{full} R_m} | Y = y)$$

would then be optimal, based on the information sets of the experts consulted.

Given a value of the vector $Y = y$, the posterior distribution of R_m is normal with mean and variance

$$E(R_m | Y = y) = \mu + \sigma^2 e' \Psi^{-1} (y - \mu e)$$

and

$$var(R_m | y) = \sigma^2 \left(1 - \sigma^2 e' \Psi^{-1} e\right),$$

where $\Psi = \Sigma + \sigma^2 ee'$. Using the fact that $\Psi^{-1} = \Sigma^{-1} - \sigma^2 \Sigma^{-1} ee' \Sigma^{-1} / (1 + \sigma^2 e' \Sigma^{-1} e)$ whenever Σ^{-1} exists, it follows that

$$E(R_m | Y = y) = \frac{\mu + \sigma^2 e' \Sigma^{-1} y}{1 + \sigma^2 e' \Sigma^{-1} y}$$

and

$$var(R_m | Y = y) = \frac{\sigma^2}{1 + \sigma^2 e' \Sigma^{-1} e},$$

so that

$$X_{full} = \frac{\mu + \sigma^2 e' \Sigma^{-1} y}{\theta \sigma^2}$$

varies linearly with the signals, as is the case with $\bar{X}$. It is important to note, however, that the inverse of the correlation matrix Σ often contains *negative* elements.

To illustrate this remark, suppose that

$$\Sigma = \begin{pmatrix} 1 & 2 \\ 2 & 10 \end{pmatrix},$$

a situation where the experts' signals are highly correlated but of wildly varying precision. In that case, it is easy to verify that the weights of y_1 and y_2 in X_{full} are respectively equal to 8/6 and -1/6, so that

$$\theta X_{full} = \frac{\mu}{\sigma^2} + \frac{8}{6}y_1 - \frac{1}{6}y_2.$$

Clearly, this strategy could not be replicated by any choice of weights w_1, w_2 in

$$\theta \bar{X} = \frac{\mu}{\sigma^2} + w_1 y_1 + w_2 \frac{y_2}{10},$$

even negative ones!

This is hardly surprising, of course, given that investment policy X_{full} assumes direct access to the experts' private signals and complete knowledge of their dependence structure. In practice, however, such information is difficult to access and to process, if only because most financial experts are loath of formulating but their final investment policy. In fact, it could be argued that managers are usually unable to identify precisely their information basis, including what they believe to be private signals. While this makes inoperative the investment strategy X_{full} identified above, one may wonder how inefficient the current, multi-manager delegation investment practices may be, in relation to this ideal. This question is examined in the following section, with particular attention to the special cases where the experts' signals are independent or are equicorrelated.

4. Evaluation of the Standard Portfolio Delegation Strategy

The purpose of this section is to describe the loss entailed by the use of the standard diversification strategy $\bar{X}$, as compared to the full information investment policy X_{full} identified above. A natural measure of this loss is provided by the ratio, π, of the expected utility corresponding to the strategies X_{full} and $\bar{X}$, respectively. Since the utility function takes only negative values, one has

$$0 \le \pi = \frac{E(-e^{-\theta X_{full} R_m})}{E(-e^{-\theta \bar{X} R_m})} \le 1.$$

The expectations in this ratio can easily be computed, upon observing that both $\theta\bar{X}$ and θX_{full} can be expressed in the form $a'Y + b$ with $b = \mu/\sigma^2$ and an appropriate choice of $a = (a_1, ..., a_n)'$. Writing

$$E\left\{-e^{-(a'Y+b)R_m}\right\} = E\left\{-e^{-bR_m} E\left(e^{-a'Yr_m} | R_m = r_m\right)\right\}$$

and using the fact that the conditional distribution of $a'Y$ given $R_m = r_m$ is normal with mean $r_m a'e$ and variance $a'\Sigma a$, it is plain that

$$E(e^{-r_m a'Y} | R_m = r_m) = e^{-V(a) r_m^2}$$

with $V(a) = a'e - a'\Sigma a/2$, because the left-hand side is the moment generating function of the random variable $a'Y$ evaluated at the point $-r_m$. Simple integration then yields

$$E(-e^{-(a'Y+b)R_m}) = E(-e^{-V(a)R_m^2 - bR_m}) = \frac{-1}{\sqrt{1+2\sigma^2 V(a)}} e^{-\mu^2/2\sigma^2}$$

so long as $V(a) \geq -1/2\sigma^2$.

It is clear that $V(a)$ is a concave function whose maximum occurs when $a = \Sigma^{-1}e$, which corresponds to X_{full}, as expected. For investment strategy $\bar{X}$ with weight vector $w = (w_1, \ldots, w_n)'$, one has $a = \Delta^{-1}w$ with $\Delta = \text{diag}(\Sigma)$, from which it follows that

$$\pi^2 = \frac{1 + 2\sigma^2(e'\Delta^{-1}w - w'\Delta^{-1}\Sigma\Delta^{-1}w/2)}{1 + \sigma^2 e'\Sigma^{-1}e}.$$

The vector w^* that maximizes this expression yields the largest possible expected utility that can be achieved through the standard investment strategy $\bar{X}$. The behavior of π as a function of this optimal w^* and of Σ is examined below in two special cases of interest.

Case 1: Two experts

It can be assumed without loss of generality that $\sigma_1^2 \leq \sigma_2^2$. If ρ stands for the correlation between the signals Y_1 and Y_2 given $R_m = r_m$, the numerator of π^2 reduces to

$$1 + 2\sigma^2\left(\frac{w_1}{\sigma_1^2} + \frac{w_2}{\sigma_2^2}\right) - \sigma^2\left(\frac{w_1^2}{\sigma_1^2} + 2\frac{w_1 w_2 \rho}{\sigma_1\sigma_2} + \frac{w_2^2}{\sigma_2^2}\right),$$

so that the optimal weight for expert 1 is given by

$$w_1^* = \frac{\sigma_2^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2} \in [0, 1]$$

when $\rho \leq \sigma_1/\sigma_2$ and equals one otherwise. Accordingly, all the weight is given to the better of the two experts when their signals are highly positively correlated. Note that when $\rho = 0$, the weights w_i^* are proportional to the precision $h_i = 1/\sigma_i^2$ of the signals, a reasonable prescription which is not necessarily optimal when there are more than two experts, however.

A graph of the largest possible expected utility that can be achieved through the standard investment strategy $\bar{X}$ is depicted in Figure 1 as a function of ρ and σ_1^2 under the assumptions that $\sigma_2^2 = 4\sigma_1^2$ and that the variance σ^2 of the market return is 1. For comparison purposes, Figure 2 shows how the ratio π varies under the same conditions when both experts are given the same weight. It is clear from these pictures that for positively but not perfectly correlated signals, there is very little difference in terms of expected utility between the full information strategy X_{full} and the standard diversification policy $\bar{X}$, even with the unsophisticated equal weighing scheme $w = (1/2, 1/2)'$. As $|\rho| \to 1$, however, there is more and more to be gained from combining the experts' signals as prescribed by X_{full}. In the extreme case where $\rho = 1$, it is easy to check that $var(R_m | Y = y) = 0$, which implies that the aggregated signals reveal the exact value of R_m. To an investor, this is worth infinitely more than any diversification strategy.

Another observation derived from the graphs is the fact that as σ_1^2 increases, the influence of ρ is not felt so strongly. This is because there is not much insight on the value of R_m to be gained by combining imprecise signals, except when they are highly correlated. Similar conclusions can be reached from graphs of π as a function of ρ and σ_1^2 for cases where $\sigma_2^2 = k\sigma_1^2$ and $k > 1$. Such graphs (not displayed) also indicate that the role of ρ decreases as $k \to \infty$. The case $k = 1$, which corresponds to interchangeable experts, is treated next.

Case 2: Interchangeable Experts

Assume that conditional on the value of R_m, one has $var(Y_i) = \sigma_0^2$ and $corr(Y_i, Y_j) = \rho \geq -1/(n-1)$ for all $1 \leq i < j \leq n$. In such a case, standard diversification arguments suggest that equal weights should be given to each expert. It is easy to check that this allocation scheme yields the optimal, standard investment strategy $\bar{X}$. Substituting $w^* = e/n$ in the formula for π^2 and using the fact that there exists an explicit expression for the inverse of an equicorrelation matrix, one gets

$$\pi^2 = \frac{\sigma_0^2 + \sigma^2 \frac{(2n-1)-(n-1)\rho}{n}}{\sigma_0^2 + \sigma^2 \frac{n}{1+(n-1)\rho}}.$$

A graphical representation of π as a function of ρ and n is given in Figure 3. A simple calculation confirms that this function is both increasing in ρ and decreasing in n. It vanishes when $\rho = -1/(n-1)$ or $n = \infty$ and reaches 1 if either ρ or n equals 1. These results are conform to intuition. It should be observed in passing that π does not tend to zero when $\rho \to 1$, because $var(R_m|Y = y) = var(R_m|Y_i) = \sigma_0^2$ in the special case where the experts' signal are equicorrelated. Investment strategies $\bar{X}$ and X_{full} are thus equivalent under these extreme conditions. For $\rho < 1$, the decreasingness of π as a function of n can be explained by the fact that each expert brings in additional information which the optimal strategy can exploit more efficiently than the standard diversification scheme.

5. Discussion

There are two morals to the theoretical story told in this note. The first is that for realistic conditions of dependence between expert signals, the standard, multi-manager portfolio investment strategy which consists of delegating assets to different managers is reasonably efficient, even with unsophisticated weight allocation schemes. The second is that appropriate use of the latent dependencies in a set of expert opinions may sometimes yield dramatic gains in utility, particularly when expert signals are negatively correlated. However, the latter prospect appears somewhat flimsy, considering the large extent to which financial experts share the same information and methods of analysis, as highlighted for example by Figlewski and Urich (1983) in their investigation of composite predictions for the weekly change in the money supply.

The leverage effect of information pooling rests essentially on its ability to "borrow" from some of the experts' signals in order to "invest" in others. While this strategy is theoretically superior to the standard diversification strategy $\bar{X}$, it could only be implemented if the investor had direct access to these signals. This raises the issue of eliciting truthful information, as discussed by Bhattacharya and Pleiderer (1985) in a delegated portfolio management context, for example. This problem is complicated by the fact that in circumstances delineated by Admati and Pfleiderer (1990), portfolio managers may find it valuable to restrict the usage of their information, by adding noise in a direct sale or by pricing usage in their mutual funds. By avoiding the issue of signal elicitation, the standard diversification strategy $\bar{X}$ escapes most of these difficulties while being close to optimal in realistic conditions.

References

Admati, A. R. and P. Pfleiderer, "Direct and Indirect Sale of Information," *Econometrica*, 58 (1990), 901-928.

Bhattacharya, S. and P. Pfleiderer, "Delegated Portfolio Management," *Journal of Economic Theory*, 36 (1985), 1-25.

Clemen, R. T., "Combining forecasts: A review and annotated bibliography," *International Journal of Forecasting*, 5 (1989), 559-583.

Figlewski, S. and T. Urich, "Optimal Aggregation of Money Supply Forecasts: Accuracy, Profitability and Market Efficiency," *The Journal of Finance*, 28 (1983), 695-710.

Gendron, M. and C. Genest, "Performance Measurement under Asymmetric Information and Investment Constraints," *The Journal of Finance*, 45 (1990), 1655-1661.

Genest, C. and J. V. Zidek, "Combining Probability Distributions: A Critique and an Annotated Bibliography" (with discussion), *Statistical Science*, 1 (1986), 114-148.

Jeffrey, H. C., "Do clients need so many portfolio managers?" *The Journal of Portfolio Management*, 18 (1991), 13-19.

Madansky, A., "Externally Bayesian Groups," Unpublished manuscript, University of Chicago Graduate School of Business, 1978.

Winkler, R. L. and R. T. Clemen, "Sensitivity of weights in combining forecasts," *Operations Research*, 43 (1992), 609-614.

FIGURE CAPTIONS

Figure 1. Ratio π of expected utilities comparing the value of the standard investment strategy $\bar{X}$ with optimal choice of weights to that of the full information strategy X_{full} in the case of $n = 2$ experts. The ratio is plotted as a function of ρ and σ_1^2, representing respectively the correlation between $n = 2$ expert signals and the variance of the first signal, conditional on the value of the market return. The variance of the second expert's signal is set equal to $4\sigma_1^2$ and the variance of the market is assumed equal to one for illustration purposes.

Figure 2. Ratio π of expected utilities comparing the value of the standard investment strategy $\bar{X}$ with equal weights to that of the full information strategy X_{full} in the case of $n = 2$ experts. The ratio is plotted as a function of ρ and σ_1^2, representing respectively the correlation between $n = 2$ expert signals and the variance of the first signal, conditional on the value of the market return. The variance of the second expert's signal is set equal to $4\sigma_1^2$ and the variance of the market is assumed equal to one for illustration purposes.

Figure 3. Ratio π of expected utilities comparing the value of the standard investment strategy $\bar{X}$ with optimal choice of weights to that of the full information strategy X_{full}. in the case of n interchangeable experts. The ratio is plotted as a function of n and ρ, the correlation between the experts' signals, conditional on the value of the market return. The common variance of the experts' signals and the variance of the market are all set equal to one for illustration purposes.

FIGURE 1

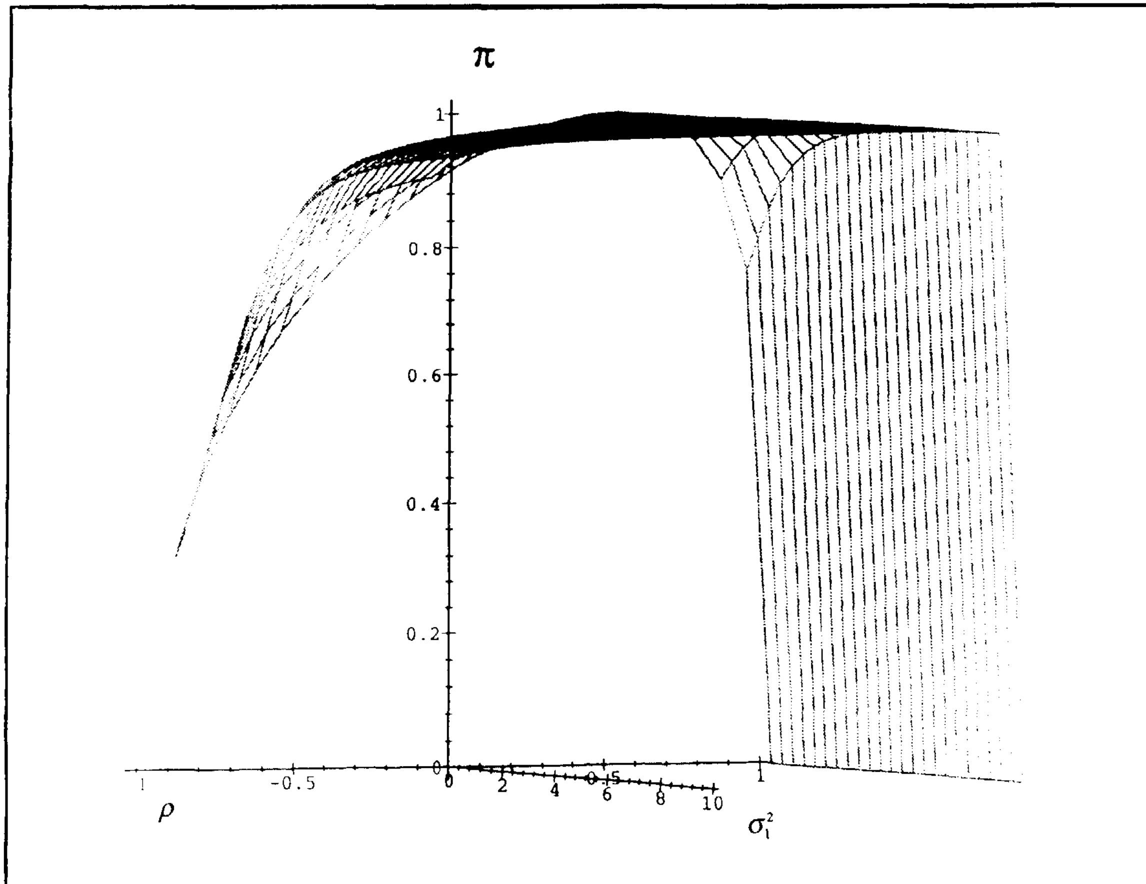

FIGURE 2