OPTIMAL CONTRACTS FOR SECURITY ANALYSTS AND PORTFOLIO MANAGERS

by

Richard E. Kihlstrom

(16-86)

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

Optimal Contracts for Security Analysts and Portfolio Managers

by

Richard E. Kihlstrom
University of Pennsylvania

September, 1985
(Revised August, 1986)

Support for this research from the National Science Foundation is gratefully acknowledged. The author is also grateful to Bill Rogerson and Rick Lambert for helpful discussions. He has also benefitted from the remarks of participants at the 1985 Nothwestern University summer workshop on Industrial Organization.

I. INTRODUCTION

In this paper we consider the problem of how a security analyst should be compensated. We view the analyst as a supplier of information about securities and we assume that the analyst is better informed than his client about the reliability of the information he provides. We consider two basic and fundamentally different cases which, using standard terminology, can be referred to as the cases of moral hazard and adverse selection. The moral hazard case is one in which the quality of the information provided by the analyst is determined by the amount of effort he expends in acquiring it. The analyst knows the information's reliability because he knows how much effort he has invested in its acquisition. The analyst's client doesn't know how hard the analyst has worked and he, therefore, doesn't know the reliability of the information supplied by the analyst. The case of adverse selection arises when the reliability of the analyst's information is determined by his ability which is unknown to the client.

Each of these cases is studied using standard approaches developed in the literatures of moral hazard and adverse selection. In each case, the standard approach must, however, be revised to account for the special features of the security analyst problem. The paper begins by adapting the approach, developed in the work of Mirrlees [1974 and 1976], Holmstrom [1979], Shavell [1979] and Grossman-Hart

[1983], to moral hazard in the principal-agent problem to provide a framework for analyzing the incentive implications of alternative copensation schemes for security analysts. The analysis of adverse selection uses the "revelation principal" discussed in Myerson [1979]. This analysis is also closely related to Stiglitz [1977].

The situation considered is one in which a securities analyst provides an investor with information about the return on a risky asset which we interpret as the market portfolio. Thus, we assume that there are only two assets, the risky asset which is the subject of the analyst's investigation and a second asset which yields a riskless return. For simplicity we assume that only two events are possible; the return on the risky asset can be high or low. The analyst's information is interpreted as a prediction of the return. This is, in effect, Merton's [1981] market timing case. The reliability of the analyst's information is measured by the conditional (on the true return) probability of an error. In this context, two types of errors are possible: the analyst can incorrectly predict a low return or he can err in predicting a high return. Our measure of reliability is meaningful if the conditional probabilities of these two different types of errors are equal. We limit our attention to this case.

In this situation, a moral hazard problem arises if the reliability of the information acquired for the investor by the analyst depends on the amount of effort expended by the analyst and if the investor is unable to monitor the analyst's efforts. When effort is unobservable, the investor is unable to write an enforceable contract linking the analyst's efforts to his compensation. The only enforceable contracts are those relating the analyst's compensation to observables. In the situation considered, the only relevant observables are the return predicted by the analyst and the return actually earned on the asset. We study the form of the optimal contract that satisfies the constraint that the analyst's compensation can only be linked to these two observables. The optimal contract satisfying this constraint is then compared to the contract that would be optimal if effort could be observed and if the analyst's compensation could therefore be made contingent on his expenditure of effort. The latter contract is referred to as the "first-best contract", while the former is called the "second best contract".

Our approach is to begin the study of this problem by posing it in its simplest form. The investor is assumed to be risk neutral as he would be if he were well diversified and if all risks were diversifiable. If the investor is interpreted as an institutional investor it is, indeed, reasonable to assume that he is well diversified. In order to insure that a moral hazard problem exists, we assume that the analyst is risk averse. In such a setting, the optimal first-best contract would have the investor bear all risks. Thus, the optimal contract would involve the payment of a non-contingent fee that would compensate the analyst for his effort. The optimal amount of effort would be that which maximized the amount by which the investor's expected return exceeded the fee required to compensate the analyst. Since effort is costly to the analyst, he can only be induced to expend effort if his compensation is tied to his level of effort. When effort is observable, the contract can link effort to compensation directly by specifying that the analyst will be paid only if he expends a predetermined level of effort. When effort is unobservable, the link between effort and compensation must be established indirectly by forcing the analyst to bear some of the risk of inaccurate predictions. This is achieved if the analyst's compensation is specified by a contingent fee schedule that, in effect, penalizes the analyst for mistakes.

In general, the investor can use any one of a number of different fee schedules to give the analyst the incentive to choose a particular level of effort. Since he is risk neutral, the investor will prefer the schedule that achieves the specified effort level at minimimum expected cost. Because the analyst is risk averse, the expected fee will be minimized by the contract that achieves the desired effort level at minimum risk to the analyst. When the conditional probabilities of each type of prediction error are equal, the analyst bears no unnecessary risk if the fee paid when there is an error is independent of the nature of the error and if the fee paid when the prediction is accurate is also independent of the correctly predicted return. Thus, under the optimal second-best contract, the analyst receives the same compensation when he wrongly predicts a high return as he does when he errs by predicting a low return. Furthermore, the optimal second-best contract also insures that the analyst receives the same compensation whenever he is correct.

It should be noted that when the optimal contract is used, it is not the return to the investor's portfolio that determines the analyst's compensation. As noted above, the analyst will receive exactly the same fee when he correctly predicts a high return as he does when he correctly predicts a low return. In general, the investor will not earn the same return in these two different cases. On the contrary, the investor will typically earn more on his portfolio when a high return is correctly predicted than he will when the analyst correctly predicts a low return.

The moral hazard problem just discussed will be replaced by an adverse selection problem if the analyst's ability to provide reliable information is unknown to the investor. The simplest context in which to study adverse selection is one in which the reliability of the information provided by the analyst is predetermined and unrelated to the amount of effort he expends. The specific model considered assumes that two levels of reliability are possible. On the one hand, the analyst may be able to provide information that is reliable enough to be useful. On the other hand his information may be essentially useless. The usefulness of the analyst's information is known to the analyst but not to his client.

The investor can effectively deal with the problem created by his ignorance of the reliability of the analyst's prediction by offering two different contracts: one designed to be more attractive when the analyst's information is reliable and the other designed to be chosen by an analyst with useless information. In the language of the adverse selection literature, these contracts are referred to as "separating" contracts and the analyst's choice of a contract is a self-selection process that reveals the quality of his information to the investor.

The investor would, of course, prefer not to pay for useless information. If, however, the analyst's contract choice can be used to infer the reliability of his information, the investor can make effective use of that knowledge when making his portfolio choice. Specifically, when the contract designed for an analyst with useless information is chosen, the investor can ignore the information when he makes his portfolio choice. When the other contract is chosen, the investor can infer that the analyst's information is reliable and use it as the basis for his porfolio choice. Thus, the investor is willing to pay the analyst with unreliable information simply to induce him to reveal the quality of the information.

We study the form of the optimal separating contracts which are those that minimize the expected fee paid by the risk neutral investor. The investor's expected fee will be smallest if he does not have to compensate the risk averse analyst for risk bearing. One of the available contracts should, therefore, be a noncontingent contract that guarantees a fixed fee. The other available contract should shift a minimal fraction of the prediction error risk to the analyst by offering a fee that is contingent on the prediction's accuracy. Specifically, a correct prediction should be rewarded with a fee in excess of that paid by the fixed fee contract, and an error should, similarly, be penalized with a fee that is lower than the one paid by the noncontingent contract.

The two contracts should be designed to be equally attractive to the analyst with unreliable information. We can then assume that he chooses the fixed fee contract. When the analyst with unreliable information is indifferent between the two contracts, the analyst with reliable information, who has a higher probability of being correct, will prefer the contingent contract. An analyst who chooses the contingent contract over the noncontingent contract reveals that his information is reliable.

The contingent contract must, of course, promise the analyst a risk premium, an expected fee in excess of the fixed fee paid by the noncontingent contract. The risk premium will be minimized if the risk of error imposed on the analyst by the contingent contract is the minimal amount necessary to achieve separation. When the two types of prediction error are equally probable (conditional on the true returns) separation will be obtained with minimal risk to the analyst, and, hence, at minimum expected cost to the investor, if the fee paid when there is an error is independent of the nature of the error and if the fee paid when the prediction is accurate is also independent of the return predicted and achieved. Thus, when the separating contracts used to solve the adverse selection problem are optimally designed, the contingent contract offered for analysts with reliable information has the same features as the second best contract that solves the moral hazard problem.

The alternative to separating contracts is to offer a single noncontingent fee contract (a "pooling" contract in the language of the adverse selection literature) which will be attractive to the analyst regardless of the reliablity of his information. If such a contract is offered, the self-selection process will not operate and the investor will be unable to infer the reliability of the analyst's information from his contract choice. Our analysis demonstrates that a pooling contract is dominated by the best separating contract.

The security analyst model discussed in this paper can be reinterpreted as a description of the relationship between an investor and a portfolio manager who uses the information provided by a prediction of the risky asset's return to choose a portfolio for the investor. If the investor can also observe the prediction, he can write a contract that specifies how the portfolio should be related to the observed prediction. In this case, our analysis can be directly reinterpreted as yielding the optimal contract for the portfolio manager.

In the portfolio manager reinterpretation, it is more realistic to assume that the portfolio manager is better informed than the investor whose portfolio is being managed. Thus, it is more interesting to consider the case in which the investor cannot directly observe the analyst's prediction and cannot, therefore, directly specify how the portfolio should be related to the prediction. In this case it is still reasonable to assume that the investor can observe the portfolio chosen by the manager. The investor can, therefore, make the portfolio manager's compensation depend on the portfolio choice. If, in making the choice of a portfolio, the portfolio manager uses the information provided by the prediction, his portfolio choice will act as an informative signal of the prediction. By using this signal as a proxy for the prediction, it is apparently possible for the investor to write contracts relating the portfolio manager's compensation to his portfolio choice and thereby acheive any outcome that could be achieved by specifying how the portfolio should be related to each possible prediction when the prediction is actually observed by the investor. There is, of course, a potential problem inherent in the use of such contracts. Specifically, the fact that the portfolio manager's fee is determined by the portfolio he chooses, may create an incentive for the portfolio manager to use his information inefficiently when he makes the portfolio choice. In certain circumstances, the potential problem fails to arise and the security analyst's optimal contract can be used to used to set the portfolio manager's contingent compensation schedule. These circumstances are described below.

The portfolio manager reinterpretation explicitly described in this paper is limited to the case of moral hazard. Although it is reasonable to expect that a similar reinterpretation can be made for the case of adverse selection, we have not yet treated that case.

Before concluding this introduction, it should be emphasized that the analysis discussed here is explicitly partial equilibrium. In particular, it ignores the important issues raised in the literature on rational expectations. Thus, we ignore the possibility that the value of information may be reduced by the fact that uninformed investor's may infer informed investors' information by observing market statistics. One partial justification for ignoring this problem is provided by the model of Grossman and Sitglitz [1980] in which the existence of "noise" prevents information from being revealed by prices. Ultimately, however, this issue will have to be faced more squarely by extending the analysis to a general equilibrium setting in which the consequences of information revealing price movements can be explicitly dealt with. Although we do not describe this extension, its formalization will presumably be based on an model similar to that of Grossman and Stiglitz. Two recent contributions of Admati and Pfleiderer [1985a and 1985b] also contain results that should be useful in this connection.

The moral hazard problem studied in this paper is formally closely related to the principal agent problems studied by Lambert [1986] and Rogerson [1985]. In their models, as in the model of the present paper, the agent acquires information for the principal. Bhattacharya and Pfleiderer [1985] also study the relationship between a portfolio manager and his client. Allen [1985a and 1985b] also deals with issues related to those treated in the present paper.

The paper is organized as follows. The moral hazard model is described in the next section. The portfolio manager reinterpretation follows in Section 3. Section 4 describes the adverse selection case. Section 5 summarizes our results.

II. MORAL HAZARD AND OPTIMAL CONTRACTS FOR SECURITY ANALYSTS

The safe asset is assumed to earn a zero rate of return. The random rate of return, $\tilde{r}$, paid by the risky asset can take on two values, $r_1>0$ and $r_2<0$. For both the analyst and the investor, the <u>a priori</u> probability that $\tilde{r}$ equals r_k is π_k, for $k = 1$ and 2.

The security analyst observes a prediction, $\tilde{x}$, of $\tilde{r}$. An observation that $\tilde{x}$ equals x_k, is interpreted as a prediction that $\tilde{r}$ will equal r_k. The conditional probability mass function of $\tilde{x}$ given that $\tilde{r} = r$ is denoted

$f(x|r)$.

We assume that

$$f(x_k|r_j) = \begin{cases} \theta & , \text{ if } j = k \\ 1 - \theta, & \text{if } j \neq k. \end{cases}$$

where θ is a measure of the prediction's reliability that is determined by the effort expended by the analyst. If the analyst expends no effort, $\theta = 1/2$ and observations of $\tilde{x}$ provide no information about $\tilde{r}$. As the analyst's efforts increase, θ rises and $\tilde{x}$ becomes a more reliable predictor of $\tilde{r}$. In fact, it is easy to demonstrate that, when θ rises, $\tilde{x}$ becomes more informative about $\tilde{r}$ in the sense of Blackwell [1951, 1953]. The utility cost to the analyst of the expenditure of effort required to acheive a prediction of reliability θ is denoted by $C(\theta)$. We normalize by assuming that $C(1/2) = 0$. The fact that effort must rise if θ is to be increased implies that $C'(\theta) > 0$. We will also assume that $C''(\theta) > 0$ and that $C(1) = \infty$. When $\theta = 1$, the prediction is perfectly accurate. Thus, the last condition imposed on C implies that a perfect prediction is prohibitively expensive.

The moral hazard problem is introduced by the fact that the θ chosen by the analyst is unobserved by the investor. The investor does, however, observe the prediction $\tilde{x}$ and he chooses his portfolio knowing the prediction. The investor's total initial wealth is denoted by I. A part, $\bar{a}$, of I is invested in the portfolio. The amount invested in the risky asset is denoted by a. We assume that

$$\bar{a} \geq a \geq 0.$$

Thus, in addition to being unable to invest more than $\bar{a}$ in the risky asset the investor cannot sell the risky asset short. Since the investor is risk neutral, his optimal strategy will be to choose $a(x) = \bar{a}$ if $E[\tilde{r} \mid x] > 0$, and to choose $a(x) = 0$ if $E[\tilde{r} \mid x] < 0$.

If the unconditional expectation

$$E\tilde{r} = r_1\pi_1 + r_2\pi_2$$

is zero, then, for all $\theta>1/2$,

$$E[\tilde{r} \mid x_1]$$
$$=[r_1\theta\pi_1+r_2(1-\theta)\pi_2]/[\theta\pi_1+(1-\theta)\pi_2]$$

will be positive and $a(x_1)$ will, hence, equal $\bar{a}$. In this case, we will also have, for all $\theta>1/2$,

$$E[\tilde{r} \mid x_2]$$
$$=[r_1(1-\theta)\pi_1+r_2\theta\pi_2]/[(1-\theta)\pi_1+\theta\pi_2]$$
$$< 0$$

and, hence, $a(x_2) = 0$. When the unconditional expectation $E\tilde{r}$ is positive (negative), there will exist a threshold level of θ, that we will refer to as θ^*, such that both $E[\tilde{r}|x_1]$ and $E[\tilde{r}|x_2]$ will be positive (negative) and both $a(x_1)$ and $a(x_2)$ will equal $\bar{a}$ (0) when $\theta<\theta^*$. When the unconditional expectation $E\tilde{r}$ is unequal to zero, and $\theta > \theta^*$, $E[\tilde{r}|x_1]$ will be positive and $a(x_1)$ will equal $\bar{a}$ while $E[\tilde{r}|x_2]$ will be negative and $a(x_2)$ will equal zero. (When $E\tilde{r}=0$, θ^* can be interpreted as one-half.)

When $\theta>\theta^*$, the expected return, $R(\theta)$, to using the optimal strategy $a(x_1)=\bar{a}$, $a(x_2)=0$ is

$$\bar{a}(1+[r_1\theta\pi_1+r_2(1-\theta)\pi_2]).$$

When $\theta<\theta^*$, the expected return, $R(\theta)$, to using the optimal strategy $a(x_1)=a(x_2)=\bar{a}$, if $E\tilde{r}>0$ and $a(x_1)=a(x_2)=0$ if $E\tilde{r}<0$ is

$$\bar{a}(1+E\tilde{r})=\bar{a}(1+r_1\pi_1+r_2\pi_2).$$

The graph of $R(\theta)$ is shown in Figure 1, for the cases $E\tilde{r}=0$ and $E\tilde{r}\neq 0$. Notice that when $E\tilde{r}=0$ and when $E\tilde{r}\neq 0$ and $\theta>\theta^*$, $R(\theta)$ is linear with slope

$$\bar{a}[r_1\pi_1-r_2\pi_2]$$

which is positive since $r_1>0$ and $r_2<0$. Also note that when $E\tilde{r}\neq 0$ and $\theta<\theta^*$, $R(\theta)$ is constant at $R(\theta)=R(1/2)$, where $R(1/2)$ is, of course, the value of no information. In this example, θ can be interpreted as the amount of information about $\tilde{r}$ obtained by observing

$\tilde{x}$. $R(\theta)$ can, therefore be interpreted as the value of information. Thus, the nonconcavity in $R(\theta)$ that arises when $E\tilde{r}\neq 0$ can be interpreted as the Radner-Stiglitz [1984] nonconcavity in the value of information. Note that there is no nonconcavity when $E\tilde{r}=0$. In this case, the assumptions that imply the Radner-Stiglitz nonconcavity fail to hold. A more complete discussion of the Radner-Stiglitz nonconcavity in an example that is formally equivalent to the present example is contained in Kihlstrom [1984].

The security analyst's utility function depends on wealth as well as on θ. Specifically, the analyst's utility function is

$$u(m)-C(\theta)$$

where m is his monetary income. The properties of C have already been discussed. The function u is assumed to be such that $u'>0$ and $u''<0$. Thus, as mentioned in the introduction, the analyst is risk averse with respect to income. In designing a compensation scheme for the analyst, the investor is constrained by the fact that the analyst can achieve a utility level $\bar{u}$ if he rejects the contract offered by the investor. This would, indeed, be the only constraint imposed on the investor's choice if the analyst's θ choice could be observed and the analyst's compensation could, therefore, be made contingent on θ.

Before proceeding to the discussion of moral hazard and the optimal second-best contract, we will first discuss the optimal first-best contract. This will provide us with a standard of comparison by which the second-best contract can be judged. We first note that although the analyst's compensation could be made contingent on both $\tilde{r}$ and $\tilde{x}$, this would only impose risks on the risk-averse analyst. These risks can be more efficiently borne by the risk neutral investor. Thus, the optimal first-best contract guarantees the analyst a fixed fee. Specifically, for any θ that might be specified by the contract, the analyst will be compensated with an amount $m(\theta)$ that is just sufficient to induce him to expend the effort required to provide a prediction of reliabilty θ. Formally, $m(\theta)$ is the solution to the constraint

(1) $u(m(\theta))-C(\theta)=\bar{u}$.

If the contract failed to satisfy this constraint, the analyst would reject the contract in favor of the other available option that guarantees him the utility level $\bar{u}$. By implicitly differentiating (1), we can obtain an expression for $m'(\theta)$, the increase in the analyst's fee required to induce him to raise θ. The expression that results from this implicit differentiation is

(2) $m'(\theta)=C'(\theta)/u'(m(\theta))$.

When the analyst chooses θ and the investor pays the analyst $m(\theta)$, the investor's final wealth is

(3) $\quad I+\bar{a}[r_1\theta\pi_1+r_2(1-\theta)\pi_2]-m(\theta).$

The θ that maximizes (3) is that which the analyst is instructed to choose by the first-best contract. We will use θ_1 to denote the first- best level of θ. Since

$$m''(\theta) = \frac{[C''(\theta)u'(m(\theta))-C'(\theta)m'(\theta)u''(m(\theta))]}{[u'(m(\theta))]^2} > 0,$$

(3) is a strictly concave function of θ and the first-best θ, θ_1, is that which solves the first-order condition

(4) $\quad \bar{a}[r_1\pi_1-r_2\pi_2]=m'(\theta)=C'(\theta)/u'(m(\theta)).$[1]

When θ cannot be observed by the investor, he cannot write an enforceable contract that instructs the analyst to expend the effort required to yield a prediction of reliability θ_1. In this case, the analyst will be guided by his own self interest when he chooses θ. If the analyst were paid a noncontingent contract in such a situation, increases in θ would cause him to bear additional costs but would yield him no additional benefits. Formally, the analyst's marginal utility of effort would be $-C'(\theta)<0$, the derivative of $u(m)-C(\theta)$ with respect to θ. He would, therefore have no incentive to provide any effort and the investor could expect to receive a completely uninformative prediction. It is possible, however, for the investor to devise a compensation scheme that causes increases in θ to result in benefits to the analyst. These benefits provide an incentive for the analyst to bear the cost of the effort required to provide informative predictions. These benefits can only be introduced, however, by making the analyst's compensation contingent on $\tilde{r}$ and $\tilde{x}$, and thereby forcing him to bear some risk.

Formally, a contingent fee schedule will be represented by a vector $M=(m_{11},m_{21},m_{12},m_{22})$ where m_{jk} is the analyst's fee when $\tilde{x}=x_j$ and $\tilde{r}=r_k$. If such a fee schedule is to induce the analyst to choose a reliability level θ, it must satisfy the constraint

(5) $U(M,\theta)-C(\theta)=\bar{u}$,

where

(6) $U(M,\theta)=[u(m_{11})\theta+u(m_{21})(1-\theta)]\pi_1$
$+[u(m_{12})(1-\theta)+u(m_{22})\theta]\pi_2$.

The constraint (5) generalizes (1) to the case of a contingent fee schedule. It, therefore, plays the same role as (1); it guarantees that the analyst will not reject the contract to pursue his alternative option.

Note that the constraint (5) can be rewritten as

(5') $[\bar{u}_C-\bar{u}_E]\theta+\bar{u}_E-C(\theta)=\bar{u}$

where

$\bar{u}_C=u(m_{11})\pi_1+u(m_{22})\pi_2$

and

$\bar{u}_E=u(m_{21})\pi_1+u(m_{12})\pi_2$.

Note also that $\bar{u}_C$ is the conditional expected utility of the analyst when he knows that his prediction will be correct. Similarly, $\bar{u}_E$ is the conditional expected utility of the analyst when he knows that his prediction will be wrong.

The optimal second-best contract must satisfy another constraint in addition to (5). It must also provide the analyst with an incentive to expend effort. Thus, the compensation schedule must be designed to guarantee that the analyst benefits from increases in θ. This will be accomplished if the derivative

$$(7)\quad U_{\theta}(M,\theta)=[u(m_{11})-u(m_{21})]\pi_1 + [u(m_{22})-u(m_{12})]\pi_2, = [\bar{u}_C-\bar{u}_E]$$

which measures the marginal utility to the analyst of a change in θ, is positive. This derivative will indeed be positive if m_{11} exceeds m_{21} and m_{22} exceeds m_{12}. In other words, the analyst will derive positive benefits from increases in θ if he is punished for his mistakes in the sense that he receives a higher fee when he correctly predicts r_1 than when he incorrectly predicts r_2 and he also receives a higher fee when he correctly predicts r_2 than when he incorrectly predicts r_1. In fact, the derivative will be positive if errors are punished in the weaker sense that $\bar{u}_C$, the conditional expected utility of the analyst when he knows that his prediction will be correct, exceeds $\bar{u}_E$, the conditional expected utility of the analyst when he knows that his prediction will be wrong.

The θ chosen by the analyst when he is faced with the fee schedule M will be that which maximizes his net utility

$$(8)\quad U(M,\theta)-C(\theta).$$

This θ will be denoted by $\theta(M)$. Since

(9) $U_{\theta\theta}(M,\theta)-C''(\theta)=-C''(\theta)<0$,

(8) is a concave function of θ, and $\theta(M)$ will be the solution to the first order condition

(10) $U_{\theta}(M,\theta)-C'(\theta)=0$.

This is the second constraint that must be satisfied by the optimal contract. Note that the constraint (10) can be rewritten as

(10') $[\bar{u}_C-\bar{u}_E]-C'(\theta)=0$.

Using $\bar{m}$ to denote the expected fee

$$[m_{11}\theta+m_{21}(1-\theta)]\pi_1$$
$$+[m_{12}(1-\theta)+m_{22}\theta]\pi_2$$

associated with the contingent fee schedule M, the second-best contract, $\langle\theta_2,M_2\rangle$, can now be defined formally as the $\langle\theta,M\rangle$ that solves the problem of maximizing the investor's expected final wealth

(11) $W=I+\bar{a}[r_1\theta\pi_1+r_2(1-\theta)\pi_2]-\bar{m}$

subject to the two constraints

(5') $[\bar{u}_C-\bar{u}_E]\theta+\bar{u}_E-C(\theta)=\bar{u}$

and

(10') $[\bar{u}_C-\bar{u}_E]-C'(\theta)=0$.

It is clear from these forms of the constraints that, if we start with a desired level of θ and a fee schedule M that, together with θ, satisfies the constraints and if we then alter the schedule so as to leave $\bar{u}_C$ and $\bar{u}_E$ unchanged, then, together with the same θ, the altered fee schedule will also satisfy the constraints. Thus, for each level of $\bar{u}_C$ we need only consider contingent fees m_{11} and m_{22} which minimize the conditional expected fee

$$m_{11}\pi_1+m_{22}\pi_2$$

while maintaining the equality

$$u(m_{11})\pi_1+u(m_{22})\pi_2=\bar{u}_C.$$

Since u is strictly concave, this is accomplished by making m_{11} equal to m_{22}. Similarly, for any level of $\bar{u}_E$, we want to choose m_{21} and m_{12} to minimize the conditional expected fee

$$m_{21}\pi_1+m_{12}\pi_2$$

while maintaining the equality

$$u(m_{21})\pi_1+u(m_{12})\pi_2=\bar{u}_E.$$

The solution to this problem is to make m_{12} equal to m_{21}.

The fee paid the analyst when his prediction is correct will be denoted by m_C and the fee paid when the prediction is incorrect will be represented by m_E. Thus, the optimal fee schedule will be such that

(12) $\quad m_{11}=m_{22}=m_C$

and

(13) $\quad m_{12}=m_{21}=m_E.$

We noted earlier, that the analyst will choose not to expend any effort unless the marginal benefit of e is positive. When the fee schedule satisfies (12) and (13), the marginal benefit to the analyst of a e increase is

(7′) $\quad U_e(M,e)=u(m_C)-u(m_E).$

This will be positive if and only if m_C exceeds m_E. Thus, the optimal second-best contract will result in reliable predictions if and only if the analyst is punished for mistakes in the sense that m_C, the fee paid by the contract when the prediction is correct, is higher than m_E, the fee paid by the contract when the prediction is wrong.

The mathematical problem has now been reduced to one of finding m_C, m_E and θ to maximize

$$(11') \quad I+\bar{a}[r_1\theta\pi_1+r(1-\theta)\pi_2]-[m_C\theta+m_E(1-\theta)]$$

subject to the two constraints

$$(5'') \quad u(m_C)\theta+u(m_E)(1-\theta)-C(\theta)=\bar{u},$$

and

$$(10'') \quad \theta=\underset{\theta}{\operatorname{argmax}}\{u(m_C)\theta+u(m_E)(1-\theta)-C(\theta)\}$$

Since, as we have noted earlier,

$$u(m_C)\theta+u(m_E)(1-\theta)-C(\theta)$$

is a strictly concave function of θ, the constraint (6''') can be replaced by the first order condition

$$(10''') \quad [u(m_C)-u(m_E)]-C'(\theta)=0.$$

This reduced problem is formally equivalent to the standard moral hazard problem as posed and studied by Marshall [1976]. The reduced problem is also a special case of the more general moral hazard problem studied by Holmstrom [1979], Shavell [1979] and Grossman-Hart [1983]. It should be emphasized, however, that, in its original form, the investor(principal)-analyst(agent) problem posed and studied above is not a special case of the problem studied by Marshall, Holmstrom, Shavell or Grossman-Hart.

In our discussion of the reduced form problem, we will find it convenient to use the approach of Grossman-Hart and to compute, for each θ, the cost to the investor of the contract that induces the analyst to expend the effort required to supply predictions of reliability θ. Because the investor is risk neutral, this cost is simply the expected fee

(14) $\bar{m}(\theta)=m_C(\theta)\theta+m_E(\theta)(1-\theta)$

where $\langle m_C(\theta),m_E(\theta)\rangle$ is the contingent contract implicitly defined as the solution of the contraints (5'') and (10''').

We can now demonstrate that moral hazard raises the cost of reliablity by comparing $\bar{m}(\theta)$ with $m(\theta)$, the guaranteed fee that satisfies the constraint

(1) $u(m(\theta))-C(\theta) = \bar{u}$.

When there is no moral hazard problem, $m(\theta)$ is the cost of a prediction of reliability θ. Since $m(\theta)$ satisfies (1), since $\langle m_C(\theta),m_E(\theta)\rangle$ satisfies (5'') and since the analyst is risk averse, $m(\theta)$ is the certainty equivalent of the random fee schedule $\langle m_C(\theta),m_E(\theta)\rangle$ and $m(\theta)$ is, therefore, exceeded by $\bar{m}(\theta)$, the expected value of $m_C(\theta)$ and $m_E(\theta)$. Figure 2 describes the relationship between $\bar{m}(\theta)$ and $m(\theta)$.

Moral hazard raises the cost of reliability because it necessitates the use of contingent contracts that reward accuracy and penalize errors. The amount by which the expected fee $\bar{m}(\theta)$ exceeds the fixed fee $m(\theta)$ is the risk premium that compensates the analyst for the risks imposed by the contingent contract. If there is no moral hazard, the risk premium is unnecessary because the investor can simply pay the analyst the noncontingent fee $m(\theta)$ that compensates him for his effort.

Using the cost function $\bar{m}(\theta)$, the formal problem facing the investor can be restated as one of choosing θ to maximize his expected final wealth

(15) $I+\bar{a}[r_1\theta\pi_1+r(1-\theta)\pi_2]-\bar{m}(\theta)$.

Notice the difference between this problem and the one faced by the investor when θ is observable and moral hazard is, therefore, absent. Recall that, when there is no moral hazard, the final wealth maximized by the investor is

(3) $I+\bar{a}[r_1\theta\pi_1+r_2(1-\theta)\pi_2]-m(\theta)$.

where $m(\theta)$ is the guaranteed fee that we have just shown to be smaller than $\bar{m}(\theta)$, the expected fee subtracted in computing (15). It is clear from expressions (3) and (15) that, although moral hazard raises reliability's cost, it has no effect on benefits. In both of these expressions for the investor's final wealth, the benefits to the investor of a prediction of reliability θ are measured by the term

$$\bar{a}[r_1\theta\pi_1+r(1-\theta)\pi_2].$$

The necessary first order condition satisfied at θ_2, the solution to the problem of maximizing (15), is,

(16) $\bar{a}[r_1\pi_1-r\pi_2]-\bar{m}'(\theta) = 0$,

where

(17) $\bar{m}'(\theta)=m_C(\theta)-m_E(\theta)$

$+m_C'(\theta)\theta+m_E'(\theta)(1-\theta)$

and where $m_C'(\theta)$ and $m_E'(\theta)$ are obtained by implicit differentation of the two constraints (5'') and (10''').[2]

If $m'(\theta)$, the marginal cost of reliability without moral hazard, could be compared to $\bar{m}'(\theta)$, the marginal cost of reliability with moral hazard, the first order conditons (4) and (16) could be used to determine moral hazard's effect on the reliability of the information supplied by the analyst. In general, this marginal cost comparison yields ambiguous results. It is nevertheless possible to use (17) to study the effect of moral hazard on the marginal cost of reliability. Note that without moral hazard, the marginal cost of reliability is simply the analyst's marginal cost of effort. With moral hazard, the marginal cost of reliability also includes an increase in the risk premium. The increase in the risk premium is required because the analyst will have an incentive to raise reliability only if the contingent contract is altered to make him bear more of the risk of error.

The expression for $\bar{m}'(\theta)$ given in (17) permits us to be specific about the determinants of the marginal cost of reliability with moral hazard. Note, in particular, that reliability increases raise the expected cost of the contingent contract $\langle m_C(\theta),m_E(\theta)\rangle$ in two ways that are measured by the two terms in (17). The first term,

(17.a) $m_C(\theta)-m_E(\theta)$,

is the premium the contract pays for accuracy. An increase in reliability raises the probability that this premium will be paid and, hence, raises the expected cost of the contract.

The second term in (17),

(17.b) $m_C'(\theta)\theta+m_E'(\theta)(1-\theta)$,

measures the expected cost increase that occurs when the fee schedule $\langle m_C(\theta),m_E(\theta)\rangle$ is adjusted to give the analyst the incentive to raise the reliability of his prediction. Implicit differentiation (which is carried out in the appendix) of the constraints (5'') and (10''') implies that

(18) $m_E'(\theta)=-\theta C''(\theta)/u'(m_E(\theta))$

and

(19) $m_C'(\theta)=(1-\theta)C''(\theta)/u'(m_C(\theta))$.

Thus, $m_C(\theta)$ must rise and $m_E(\theta)$ must fall if the contract change is to induce the analyst to increase θ. The increase in $m_C(\theta)$, which can be interpreted as an increase in the reward for accuracy, raises the expected cost of the contract but the fall in $m_E(\theta)$, which can be viewed as an increase in the penalty for errors, results in an expected cost decrease. When (18) and (19) are substituted in (17.b), we obtain an expression for the net effect of the contract change on the expected cost of the contract. The result is

$$(17.b') \quad m_C'(\theta)\theta+m_E'(\theta)(1-\theta)$$
$$=\theta(1-\theta)C''(\theta)\{[1/u'(m_C)]-[1/u'(m_E)]\},$$

which is positive because the increase in $m_C(\theta)$ outweighs the decrease in $m_E(\theta)$. Thus, the expected cost of the contract increases when $m_C(\theta)$ increases and $m_E(\theta)$ decreases so as to create an incentive for a more reliable prediction.

In summary, the marginal cost of reliability is positive for two reasons. On the one hand, more reliable predictions are simply more likely to be correct and to earn the analyst a premium for accuracy. On the other hand, although the increased premium required to provide the incentive for the increase in reliability is achieved by a simultaneous increase in both the reward for a correct prediction and the penalty for an error, the increase in the reward must outweigh the increase in the penalty. The increase in the expected fee that results as the sum of these two effects can be viewed as an increase in the risk premium required not only to cover the marginal cost of effort, but also to compensate the analyst for the increase in risk he faces because the penalty and the reward have both increased.

Using the expression in (17.b'), the expression (17) for the marginal cost of reliability can be rewritten as

$$(17') \quad \bar{m}'(\theta)=[m_C-m_E]$$
$$+\theta(1-\theta)C''(\theta)\{[1/u'(m_C)]-[1/u'(m_E)]\}.^{\ge}$$

We can also substitute (17′) in the first order-condition (16) that determines θ_2. The result is

(16′) $\bar{a}[r_1\pi_1-r\pi_2]$

$$= [m_C(\theta)-m_E(\theta)]$$

$$+\theta(1-\theta)C''(\theta)\{[1/u'(m_C)]-[1/u'(m_E)]\}.$$

As noted earlier, there is, in general, no reason for the expression for $\bar{m}'(\theta)$ in (17′) to exceed the expression for $m'(\theta)$ in (2). If $m'(\theta)$ did exceed $m'(\theta)$, then the first order conditions (4) and (16′) could be used to prove that, in accordance with intuition, moral hazard reduces the analyst's effort and the reliability of his prediction.

III. MORAL HAZARD AND OPTIMAL CONTRACTS FOR PORTFOLIO MANAGERS

The preceding analysis determines how a security analyst's compensation should be related to his performance. The same analysis can, in certain circumstances, be reinterpreted to yield optimal contracts for portfolio managers. The analysis can, in fact, be applied without alteration to cases in which the investor observes the prediction, $\tilde{x}$. In such cases, the investor is able to write contracts that specify how the prediction, the portfolio choice and the manager's compensation should be related. Thus, the contract that is optimal for the security analyst is feasible and, indeed, optimal for the portfolio manager.

It is more realistic, however, to consider cases in which the investor cannot observe $\tilde{x}$ and cannot, therefore, write enforceable contracts that directly specify, for each possible prediction, how the portfolio should be chosen and how the manager should be compensated. Even if the investor cannot observe $\tilde{x}$, he usually can observe the portfolio choice. He can, therefore, write a contract that relates the manager's compensation to the portfolio. If the manager uses $\tilde{x}$ as the basis for his portfolio choice, the portfolio can be used as a proxy for $\tilde{x}$ in structuring the contract that determines the manager's compensation.

A potential problem arises when the manager's compensation is determined by his portfolio choice. Specifically, he may have an incentive to use the information provided by the prediction inefficiently when choosing the portfolio. As we now demonstrate, there are cases where this potential problem does not arise and where the security analyst's optimal contract can be rewritten to provide an optimal contract for the portfolio manager. We begin this demonstration by describing how the terms of the security analyst's optimal contract can be revised so as to make it feasible.

Recall that the security analyst's optimal contract $\langle\hat{\theta},\hat{m}_C,\hat{m}_E\rangle$ maximizes (11′) subject to (5′′) and (10′′′). When the investor cannot observe $\tilde{x}$, he can use the portfolio choice as a proxy for $\tilde{x}$ by announcing that he will pay the portfolio manager m_C if $a=\bar{a}$ and $\tilde{r}=r_1$ or if $a=0$ and $\tilde{r}=r_2$ and that the portfolio manager's fee will be m_E if $a=\bar{a}$ and $\tilde{r}=r_2$ or if $a=0$ and $\tilde{r}=r_1$. This revised contract will be formally referred to as the <u>revised security analyst contract</u>.

We can now ask how the manager's portfolio choice will be related to his prediction when he is faced with the revised security analyst contract. For any possible θ choice above θ^*, consider first the case in which $\tilde{x}$ has been observed to equal x_1. If the portfolio manager chooses $a=\bar{a}$ as he should in this case, his expected utility will be

(20) $u(m_C)\theta\pi_1+u(m_E)(1-\theta)\pi_2.$

If, however, the portfolio manager chooses $a=0$, his expected utility will equal

(21) $u(m_E)\theta\pi_1+u(m_C)(1-\theta)\pi_2$.

The portfolio manager will, therefore, choose the appropriate portfolio if (20) exceeds or equals (21). It is easily verified that (20) does, indeed, exceed (21) if and only if

(22) $[u(m_C)-u(m_E)][\theta-\pi_2]\geq 0$.

Since m_C is larger than m_E, (22) is equivalent to

(23) $\theta\geq\pi_2$.

When $\tilde{x}=x_2$, the portfolio manager will choose a=0, as he should, if the expected utility

(24) $u(m_E)(1-\theta)\pi_1+u(m_C)\theta\pi_2$

obtained by choosing a=0 exceeds the expected utility

(25) $u(m_C)(1-\theta)\pi_1+u(m_E)\theta\pi_2$.

obtained by choosing $a=\bar{a}$. Again, since m_C exceeds m_E, (24) will exceed (25) if and only if

(26) $\theta\geq\pi_1$.

The preceding remarks imply that, when the portfolio manager is faced with the revised security analyst contract, he will make appropriate use of the information provided by predictions of reliability greater than $\max\{\pi_1,\pi_2\}$ and he will make no use of predictions with reliability levels below $\max\{\pi_1,\pi_2\}$. If, for example, $\pi_2>\pi_1$ and the reliability level, θ, of the prediction observed by the portfolio manager is such that $\pi_2>\theta>\pi_1$, then the portfolio manager will always choose a=0 and his expected utility will be

$$u(m_E)\pi_1+u(m_C)\pi_2.$$

Note that when $\pi_2>\theta>\pi_1$, the probability that the prediction will be correct is lower than the probability of a low return, r_2. Thus, the portfolio manager is more likely to earn the high fee m_c if ignores his relatively unreliable information and simply chooses the portfolio that is appropriate when the return is low. Note also that when the portfolio manager ignores the information he may choose a portfolio different from the one preferred by the investor. If, for example, the inequalities $\pi_2>\theta>\pi_1$ and $r_1\pi_1+r_2\pi_2>0$ hold simultaneously as they may, and if the investor knows that the information will be ignored when the portfolio is chosen, he will prefer the portfolio with all wealth invested in the risky asset. The portfolio manager will, indeed, ignore the information when he chooses the portfolio, but he will prefer to invest all wealth in the riskless asset.

Let us now suppose that $\pi_2>\pi_1$ and ask how the portfolio manager's choice of θ will be made when his compensation is determined by the revised security analyst contract. We will use $\phi(\theta)$ to denote the portfolio manager's expected utility. Because the analyst uses the information correctly when $\pi_2<\theta$ and ignores the information by always choosing a=0 when $\pi_2>\theta$, his expected utility will be

$$\phi(\theta)\equiv\begin{cases} u(m_E)\pi_1+u(m_C)\pi_2-C(\theta), & \text{if } \theta<\pi_2 \\ \\ u(m_C)\theta+u(m_E)(1-\theta)-C(\theta), & \text{if } \theta>\pi_2. \end{cases}$$

Note that if $\hat{\theta}>\pi_2$, then $\hat{\theta}$ maximizes ϕ, so the portfolio manager chooses the optimal level of reliability and makes optimal use of the information provided by the prediction. In this case. the potential problem mentioned above is avoided; the portfolio manager is not discouraged from making efficient use of the optimal amount of information and the revised security analyst contract is the optimal contract for the portfolio manager.

When $\hat{\theta}<\pi_2$, the revised security analyst contract does create an incentive for the analyst to ignore the optimal amount of information. In this case, the problem of choosing an optimal contract for the portfolio manager requires a more complicated analysis that involves dropping the conditions that

(12) $m_{11}=m_{22}=m_C$

and

(13) $m_{12}=m_{21}=m_E$

and then imposing the resulting analogues of the inequalities (20)>(21) and (24)>(25) as incentive compatibility conditions when $\pi_2>\theta$. We will not carry out this analysis in the present paper.

IV. ADVERSE SELECTION AND OPTIMAL CONTRACTS FOR SECURITY ANALYSTS

For the purpose of investigating the problem of adverse selection in the investor-analyst relationship, the parameter θ is interpreted as a measure of the analyst's predictive ability. The analyst's ability is assumed to be unknown to the investor. For simplicity we assume that only two ability levels are possible. If the analyst is a reliable forecaster, then $\theta=\theta_1>\theta^*$. If his forecasts are unreliable, then $\theta=\theta_2<\theta^*$. Thus, if the investor knew that θ were equal to θ_1, he would use the investment rule $a(x_1)=\bar{a}$ and $a(x_2)=0$. If, on the other hand, the investor knew that θ were equal to θ_2, and if the unconditional expectation $E\tilde{r}$ were, for example, negative, then the investor would ignore the prediction and use the investment rule $a(x_1)=a(x_2)=0$.

It is assumed that, although the investor does not know whether θ equals θ_1 or θ_2, he does know the probability of each possible θ. The probability of θ_j is denoted by μ_j.

The security analyst's utility function depends only on wealth. There is no cost of acquiring information. We let m represent the analyst's monetary income and denote his utility function by u(m) where $u'>0$ and $u''<0$. The analyst is, therefore, risk averse.

In designing a compensation scheme for the analyst, the investor is constrained by the fact that the analyst has other options that can be exercised if he finds the investor's offer unattractive. The alternative option is assumed to be more attractive for the analyst whose information is reliable. Formally, the analyst whose reliability level is θ_j can achieve a utility level $\bar{u}_j$ if he rejects the contract offered by the investor, and $\bar{u}_1 > \bar{u}_2$.

If the analyst's ability, θ, were observable to the investor and if the investor observed that θ equaled θ_1, he would offer the analyst a noncontingent fee m_1 which satisfied the constraint

(27) $u(m_1) = \bar{u}_1$.

This offer would be just attractive enough to induce the reliable analyst to foresake his alternative option and provide information to the investor. If the analyst were observed to lack ability, the investor would not offer him anything for his information. An optimal allocation of resources to the provision of information results when the analyst is compensated in this way: the investor acquires information only when its value exceeds the opportunity cost of its acquisition. In this connection, note that although there is no cost of acquiring information, $\bar{u}_j$ does measure the opportunity cost incurred when the analyst provides information of reliability θ_j to the investor.

If the investor were unable to observe the analyst's ability, he could still induce the analyst to provide information when θ equaled θ_1 by offering the noncontingent fee m_1.
The problem is that, because $\bar{u}_1 > \bar{u}_2$, this offer would be even more attractive to the analyst when his prediction was useless than when it was reliable since (27) would imply
(28) $u(m_1) > \bar{u}_2$.
Thus, if the investor were to offer to pay the analyst a fee of m_1 for information, the analyst would always claim to have a reliable prediction. With probability μ_1 this claim would be correct and with probability μ_2, it would be false. The expected reliability of the prediction would therefore be
$\bar{\theta} = \theta_1\mu_1 + \theta_2\mu_2$.
The expected reliabilty level $\bar{\theta}$ would be above θ^* if μ_1 were large but it would be below θ^* if μ_1 were near zero.

If the investor had rational expectations, he would use the information provided by the prediction of an analyst who accepted the fee m_1 by following the investment strategy appropiate when θ equaled $\bar{\theta}$. That is, he would use the strategy $a(x_1)=\bar{a}$ and $a(x_2)=0$, if $\bar{\theta}>\theta^*$ and he would use the strategy $a(x_1)=a(x_2)=0$, if $\bar{\theta}<\theta^*$ (assuming that $E\tilde{r}$ is negative). The expected return to this investment strategy is $R(\bar{\theta})$. Thus, the expected return to an investor who offers the analyst a single contract with a guaranteed fee of m_1 is also $R(\bar{\theta})$.

The investor can, however, do better than $R(\bar{\theta})$ by offering a contingent fee schedule for a reliable prediction and simultaneously offering to pay the analyst a noncontingent fee for admitting that his information is useless when it, in fact, is. The contingent fee schedule would make the analyst's compensation depend on the prediction and on the return actually paid by the risky asset.

Formally, a contingent fee schedule will be represented by a vector $M^1=(m_{11},m_{21},m_{12},m_{22})$ where m_{jk} is the analyst's fee when $\tilde{x}=x_j$ and $\tilde{r}=r_k$. If such a fee schedule is to induce the analyst to supply information when he is capable of providing a reliable prediction, it must satisfy the constraint

(29) $U(M^1,\theta_1)=\bar{u}_1$,

where

(30) $U(M^1,\theta)=[u(m_{11})\theta+u(m_{21})(1-\theta)]\pi_1$
$$+[u(m_{12})(1-\theta)+u(m_{22})\theta]\pi_2.$$

This constraint is the analog of (27) for the case of a contingent compensation schedule. It guarantees that the reliable analyst will not reject the contract to pursue his alternative option.

The noncontingent fee paid to the analyst when he admits to having access to unreliable information will be denoted by M^2. This fee will have to satisfy the constraint

(31) $u(M^2)=U(M^1,\theta_2)$

if it is to be an effective inducement to truthful revelation by the analyst when his information is unreliable. The noncontingent fee M^2 will also have to satisfy the constraint

(32) $u(M^2) \leq U(M^1, \theta_1)$

if it is not to be too attractive to the analyst when his prediction is reliable. Conditions (31) and (32) combine to imply a restriction that, in fact, involves only the contingent fee schedule M^1. Specifically, (31) and (32) imply that M^1 must be such that

(33) $[\bar{u}_C - \bar{u}_E][\theta_1 - \theta_2] > 0$

where, as in the moral hazard case,

(34) $\bar{u}_C = u(m_{11})\pi_1 + u(m_{22})\pi_2$

and

(35) $\bar{u}_E = u(m_{21})\pi_1 + u(m_{12})\pi_2$.

Recall that $\bar{u}_C$ is the conditional expected utility of the analyst when he knows that his prediction will be correct. Remember also, that $\bar{u}_E$ is the conditional expected utility of the analyst when he knows that his prediction will be wrong.

Since θ_1 exceeds θ_2, condition (33) can hold only if M^1 guarantees the analyst a higher conditional expected utility when he is right than when he is wrong. Thus, in the presence of (31),

(36) $\bar{u}_C > \bar{u}_E$

is a necessary condition for (32). In fact, because θ_1

exceeds θ_2, and because $U(M^1,\theta)$ can be expressed as

(37) $U(M^1,\theta)=[\bar{u}_C-\bar{u}_E]\theta+\bar{u}_E$

condition (36) is also sufficient for (32) in the presence of (31).

Any pair of fee schedules $\langle M^1,M^2\rangle$ that satisfies the constraints (29), (31) and (32) can be used to induce the analyst to reveal the quality of his prediction to the investor. Such a pair of schedules will be referred to as separating contracts. When faced with the choice between the noncontingent fee M^2 and the contingent fee schedule M^1, the analyst will opt for M^1 only if his information is useful to the investor. Thus, the investor will use the information of the analyst only if the analyst chooses the fee schedule M^1. The expected return to the investor's portfolio choice strategy in that case will be $R(\theta_1)$. The expected cost of the fee schedule M_1 chosen in that case will be

(38) $E[M^1|\theta_1]=[\theta_1 m_{11}+(1-\theta_1)m_{21}]\pi_1$

$$+[\theta_1 m_{22}+(1-\theta_1)m_{12}]\pi_2.$$

Thus, when the analyst is observed to choose the fee schedule M^1, the investor's expected profits will be

(39) $R(\theta_1)-E[M^1|\theta_1]$.

Since the analyst will opt for the fee schedule M^1 only when $\theta=\theta_1$, and since the probability of this event is μ_1, μ_1 will also be the probability with which the investor earns the profits (39).

When the analyst is observed to choose the noncontingent fee M^2, the investor knows that the analyst's prediction is unreliable and the return to the portfolio strategy he uses in that case is $R(\theta_2)=R(1/2)$. The investor's profits in this case are

$$R(1/2)-M^2.$$

These profits are obtained with probability μ_2.

The *ex ante* expected profits of the investor who offers a separating contract $\langle M^1,M^2\rangle$ are, therefore,

$$\{\mu_1 R(\theta_1)+\mu_2 R(1/2)\}$$
$$-\{\mu_1 E[M^1|\theta_1]+\mu_2 M^2\}.$$

As noted above, the investor can also induce the analyst to supply information when his information is reliable by offering the noncontingent fee m_1. This fee will, however, attract the analyst whose information is useless as well as the analyst with useful information. The offer of the noncontingent fee m_1 will be referred to as a *pooling contract*. The expected profits of this strategy are

(40) $R(\bar{\theta})-m_1$.

The investor will choose to offer the noncontingent fee m_1 to both types of analysts if the expected profits (40) of this strategy exceed the ex ante expected profits of the best separating contract $\langle M^1,M^2\rangle$.

Since all separating contracts result in the same expected returns to the portfolio strategy, these contracts are distinguished only by their expected costs. The optimal separating contract is that which minimizes the expected fee

(41) $\mu_1 E[M^1|\theta_1]+\mu_2 M^2$.

subject to the constraints (29), (31) and (32). In searching for the expected cost minimizing contract we can simply assume that condition (36) holds and ignore the constraint (32). Once the expected cost minimizing separating contract has been found we can compare the expected profits of this contract with the expected profits of the pooling contract. Since

(42) $\{\mu_1 R(\theta_1)+\mu_2 R(1/2)\}>R(\bar{\theta})$,

the only possible advantage the pooling contract can have over the best separating contract is that m_1 might be less than the expected fee

(43) $\mu_1 E[M^1|\theta_1]+\mu_2 M^2$.

It is possible to show, however, that the minimum expected fee associated with the best separating contract is smaller than m_1.

As a first step in the derivation of an optimal separating contract, note that the constraints (29) and (31) can be rewritten as

(29') $[\bar{u}_C-\bar{u}_E]\theta_1+\bar{u}_E=\bar{u}_1$

and

(31') $[\bar{u}_C-\bar{u}_E]\theta_2+\bar{u}_E=u(M^2)$.

Since the contingent fee schedule enters the constraints and condition (36) only through the conditional expected utilities $\bar{u}_C$ and $\bar{u}_E$, the argument used for the moral hazard case can be applied here to demonstrate that, in in an optimal contract,

$$m_{11}=m_{22}$$

and

$$m_{12}=m_{21}.$$

If we denote the common value taken by m_{11} and m_{22} by m_C, and if we use m_E to denote the value at which m_{12} and m_{21} are equated, then the problem of finding an optimal separating contract is reduced to one of choosing M^2 and $M^1=(m_C,m_E)$ so as to minimize the expected fee

(41′) $\mu_1[m_C\theta_1+m_E(1-\theta_1)]+\mu_2M^2$

subject to the constraints

(29′′) $[u(m_C)-u(m_E)]\theta_1+u(m_E)=\bar{u}_1$

and

(31′′) $[u(m_C)-u(m_E)]\theta_2+u(m_E)=u(M^2)$.

In the reduced form just described, the mathematical problem of finding an optimal pair of contracts $\langle M^1,M^2\rangle$ is formally equivalent to that solved in Stiglitz's [1977] analysis of the monopoly provision of insurance with adverse selection. Although Stiglitz's analysis can be reinterpreted to provide a solution to this problem, the intuitive nature of the solution is more clearly revealed by a direct analysis.

The analysis is structured to yield a simple comparison of the costs and benefits of alternative separating contracts. The same cost-benefit approach is used to compare the best separating contract with the pooling contract. In fact, a useful way of introducing the discussion of costs and benefits is to compare the expected costs of a typical seaprating contract with the cost of the pooling contract. The pooling contract pays analysts of both types a fixed fee m_1. We can interpret the introduction of the contingent contract as a way of reducing, to a level below m_1, the fixed fee M^2 paid to unreliable analysts. (The observation that M^2 is lower than m_1, follows from the fact that $\theta_1 > \theta_2$ and from (27), (29'') and (31'').) The cost of this reduction in M^2 is an increase, to a level above m_1, of the expected contingent fee paid to reliable analysts. (The result that the expected contingent fee exceeds m_1 follows from (27) and (29'') when $m_C > m_E$.) The question of whether the expected cost of a pair of separating contracts exceeds the cost of the pooling contract can, therefore, be viewed as a question of whether the reduction in the unreliable analyst's fixed fee achieved with the separating contracts is worth the corresponding increase in the reliable analyst's expected contingent fee.

Each pair of separating contracts can also be identified with the benefit it achieves in the form of a reduction in M^2, the fixed fee paid to unreliable analysts, and the cost it imposes in the form of an increase in the expected contingent fee paid to reliable analysts. Different pairs of separating contracts can then be compared by comparing these costs and benefits.

Formally, for each level of the fixed fee M^2, let $M^1(M^2) = \langle m_C(M^2), m_E(M^2)\rangle$ be the contingent contract that solves the constraints (29'') and (31''). The expected value of the contingent fee schedule $M^1(M^2)$,

$$\bar{m}(M^2)=m_C(M^2)\theta_1+m_E(M^2)(1-\theta_1),$$

represents the cost of reducing the fixed fee to the level M^2. The marginal cost of a further reduction in M^2 is $-\bar{m}'(M^2)$. A formal expression for $\bar{m}'(M^2)$ can be obtained by implicitly differentiating (29'') and (31'') to obtain expressions for $m_C'(M^2)$ and $m_E'(M^2)$. In fact, the implicit differentiation is most easily carried out if we first rearrange the constraints and observe that $m_C(M^2)$ satisfies

(44) $$u(m_C(M^2)) =[(1-\theta_2)\bar{u}_1-(1-\theta_1)u(M^2)]/(\theta_1-\theta_2)$$

and that $m_E(M^2)$ satisfies

(45) $$u(m_E(M^2)) =[\theta_2\bar{u}_1-\theta_1 u(M^2)]/(\theta_2-\theta_1).$$

Implicit differentiation of (44) now implies that

(46) $$m_C'(M^2) =[(\theta_1-1)u'(M^2)]/[(\theta_1-\theta_2)u'(m_C)] < 0,$$

and implicit differentiation of (45) implies that

(47) $$m_E'(M^2) =[\theta_1 u'(M^2)]/[(\theta_1-\theta_2)u'(m_E)] > 0.$$

Expressions (46) and (47) imply that additional reductions in M^2 can only be achieved by increasing the risk of error imposed on reliable analysts by the contingent contract.

Specifically, (46) and (47) assert that decreases in M^2 are accompanied by increases in $m_C(M^2)$, the fee paid for correct predictions, and decreases in $m_E(M^2)$, the fee paid in the event of error. Since (29'') must continue to hold after the increase in the riskiness of the contingent contract, i.e., since the reliable analyst must be compensated for the increase in risk so that he is still indifferent between the contingent contract and his alternative option, the expected fee $\bar{m}(M^2)$ will have to increase when M^2 decreases. Using (46) and (47), the expression for the increase in $\bar{m}(M^2)$ associated with a decrease in M^2 is the negative of

(48) $\bar{m}'(M^2)$

$$= \frac{[u'(m_E)^{-1}-u'(m_C)^{-1}]u'(M^2)\theta_1(1-\theta_1)}{(\theta_1-\theta_2)} \gtrless 0.$$

The optimal pair of separating contracts is the one associated with the fixed fee M^2 that minimizes the expected fee

(41'') $\bar{M}(M^2) \equiv \mu_1\bar{m}(M^2)+\mu_2M^2.$

The necessary first-order condition satisfied at the optimal M^2 is

(49) $\bar{M}'(M^2) = \mu_1\bar{m}'(M^2) + \mu_2 = 0.$

This condition simply asserts that, at the optimal M^2, the benefit of a further decrease in M^2 is just equal to the marginal cost of acheiving that benefit, i.e., the increase in the expected contingent fee that must accompany the decrease in M^2.

While the intuitive meaning of (49) is clear, care must be taken in interpreting solutions of (49) as cost minimizing M^2 levels. The problem, which will become clear in the course of the subsequent discussion, is that $\bar{M}(M^2)$ is a nonconvex function of M^2.

Having described the necessary first order conditon satisfied at the best pair of separating contracts, we can now prove that this contract is better than the pooling contract. In the proof, we can ignore the fact that separating contracts permit the investor to achieve an expected return

$\{\mu_1 R(\theta_1)+\mu_2 R(1/2)\}$

that exceeds $R(\bar{\theta})$, the expected return that can be achieved without separation. In other words, we can simply show that the expected cost,

(41'') $\bar{M}(M^2) = \mu_1 \bar{m}(M^2)+\mu_2 M^2$,

of the best separating contract is lower than the cost, m_1, of the pooling contract.

For the purpose of comparing the cost of the pooling contract with that of any separating contract, it is useful to note that the pooling contract can be regarded as a pair of separating contracts, one of which is a degenerate contingent contract. Specifically, when the fixed fee contract pays $M^2 = m_1$, the contingent contract that satisfies (29'') and (31'') is $\langle m_C(m_1), m_E(m_1)\rangle = \langle m_1, m_1\rangle$. This observation permits us to compute the cost of the pooling contract by substituting $M^2 = m_1$ in (41''). The result is, as it must be, $\bar{M}(m_1) = m_1$.

We now use (48) to observe that, at $M^2 = m_1$,

$\bar{m}'(M^2) = 0$ and

$$\bar{M}'(M^2) = \mu_2 > 0,$$

Since this derivative is positive, reducing M^2 below m_1 will reduce $\bar{M}(M^2)$ below m_1. In particular, the cost minimizing M^2 level is below $M^2 = m_1$. This, of course, means that the best separating contract is less costly than the pooling contract as asserted above.

The fact that $\bar{m}'(M^2)$ is zero at $M^2 = m_1$ and positive when $M^2 < m_1$, implies that $\bar{m}(M^2)$ and, hence, $\bar{M}(M^2)$ are nonconvex functions of M^2. This is the nonconvexity mentioned earlier.

V. SUMMARY

The security analyst considered in this paper is a supplier of information about security returns. As we have seen, when the reliability of the information supplied by the analyst is determined by the amount of effort he invests in its acquisition and when the analyst's client is unable to monitor the analyst's efforts, a moral hazard problem arises. As we have also observed, if it is the analyst's ability that determines the reliability of the information he provides and if the client is less well informed about the analyst's ability than the analyst, then an adverse selection problem arises.

In a moral hazard setting, the investor would like to pay the analyst a fee that depends only on the level of effort he exerts. Such a contract would induce the analyst to exert the optimal amount of effort and it would not transfer any risk to the risk averse analyst from the risk neutral investor. The unobservability of effort makes such contracts unenforceable. The investor is limited to using contracts that tie the analyst's compensation to predicted and realized returns. By imposing some of the risk of prediction errors on the analyst such contracts do provide the analyst with an incentive to exert effort. Because these contracts impose risk on the analyst, they also raise the expected fee the analyst has to be paid for any given level of effort he expends. When the optimal enforceable contract is used, the marginal benefit of additional effort is just equal to the marginal cost of additional effort as measured by the increase in the analyst's expected fee. When there is moral hazard, the increase in the expected fee covers not only the cost of additional effort it also covers the cost of the added risk that must be borne by the analyst if he is to have the incentive to exert more effort.

With moral hazard, the optimal contract has the feature that the analyst's fee depends only on whether he is correct or not. His fee does not depend on what he correctly predicts nor on what kind of mistake he makes. Fee schedules that pay different fees for different kinds of errors or for different correct predictions, simply increase the risk borne by the analyst without improving his incentive to exert effort.

When there is adverse selection, the optimal fee schedule has the same structure as that just described for the case of moral hazard. All types of errors are penalized with the same low fee and all correct predictions are rewarded with the same high fee. This contract is used for a different purpose in cases of adverse selection. In such cases, the problem is not one of providing the analyst with an incentive to exert effort. The problem is rather one of inducing the analyst to reveal his ability through his choice of a contract. Thus, the investor finds it advantageous to offer two kinds of contracts. In addition to the contingent contract that penalizes errors and rewards correct predictions in the way just described, the investor also offers a fixed fee contract. When the fees paid by these two contracts are properly structured, the analyst will reveal his ability by his contract choice. The able analyst for whom the risk of a prediction error is low will choose to take the risks imposed by the contingent contract. The analyst who lacks ability will reveal his inability by choosing the fixed fee contract. This self selection process works because the amount of risk imposed on the analyst by the contingent contract depends on the analyst's probability of error. For able analysts this probability is low, for analysts without ability the error probability is high.

The knowledge of the analyst's ability, acquired by observing the contract choice, is used by the investor to determine how the return prediction should influence the portfolio decision. The predictions of able analysts are used, those of analysts who lack ability are ignored.

If, in the adverse selection setting, the investor knew the analyst's ability, he would be able to avoid paying for the useless information provided by an analyst who lacked ability, and he would also be able to pay the able analyst a fixed fee. This fixed fee contract would impose no risk on the risk averse analyst. The fixed fee would, therefore, be lower than the expected fee paid by the contingent contract used to compensate able analysts when the investor is forced to learn the analyst's ability by observing his contract choice. Adverse selection, therefore, imposes two costs on the investor. He pays for unreliable information even though he doesn't use it. He also pays a risk premium to able analysts who bear risk because their fee is contingent on the accuracy of their predictions.

The results obtained, in the case of moral hazard, for a security analyst can be reinterpreted to provide an optimal fee schedule for a portfolio manager whose relationship to his client is also marked by the presence of moral hazard. As modeled here, the distinction between a portfolio manager and a security analyst is that the prediction of a portfolio manager is not observed by the manager's client. This can be interpreted to mean that the portfolio manager has information about the return on the risky asset that cannot be made available to the client in the same way that a security analyst's prediction can. The client can use the portfolio manager's information only by delegating the portfolio choice to the manager.

Under certain conditions the optimal security analyst contract can be rewritten so that the results obtained with the portfolio manager duplicate those obtained with the security analyst. The restructured contract makes the portfolio manager's compensation contingent on the portfolio he chooses and the return actually paid by the risky security. Specifically, when the portfolio the manager chooses is inappropriate for the risky security return actually earned, he is punished as though he were a security analyst who had made an incorrect prediction. Similarly, if the portfolio chosen by the manager is appropriate to the risky security return actually earned, the manager is rewarded as though he had made a correct prediction.

The problem with this kind of contract is that it might give the analyst an incentive to ignore information he should use. This occurs when either the probability of a high return is very likely or the probability of a low return is very likely. If, for example, the low return on the risky security is very likely, the portfolio manager is better off simply ignoring his information and always investing nothing in the risky security. When the high and low returns are approximately equally likely, it is a better bet to use information than to ignore it when choosing a portfolio.

APPENDIX

PROPOSITION: <u>The contract</u> $\langle m_C(\theta), m_E(\theta)\rangle$ <u>defined implicitly as the solution to the constraints</u>

(5'') $u(m_C)\theta+u(m_E)(1-\theta)-C(\theta)=\bar{u}$,

<u>and</u>

(10''') $[u(m_C)-u(m_E)]-C'(\theta)=0$.

<u>is such that</u>

(18) $m_E'(\theta)=-\theta C''(\theta)/u'(m_E(\theta))$

<u>and</u>

(19) $m_C'(\theta)=(1-\theta)C''(\theta)/u'(m_C(\theta))$.

PROOF: Note that the constraints (5'') and (10''') can be combined to obtain

(A1) $u(m_E)+C'(\theta)\theta-C(\theta)-\bar{u}=0$,

a single equation that determines $m_E(\theta)$ and that can be implicitly differentiated to obtain (18). The same two constraints can be similarly rearranged to yield

(A3) $u(m_C)-C'(\theta)[1-\theta]-C(\theta)-\bar{u}=0$,

a single equation that determines $m_C(\theta)$ and that can be implicitly differentiated to obtain (19).■

FOOTNOTES

1. Actually, (3) is concave only on the domain $[\theta^*,\infty)$. We will assume, however, that the optimal effort level θ_1 is larger than one-half. This effectively means that θ_1 maximizes (3) on the interval $[\theta^*,\infty)$ and that (3) evaluated at $\theta = \theta_1$ exceeds (3) evaluated at $\theta = 1/2$.

2. Care must be taken in interpreting solutions of (16) as solutions to the problem of maximizing (15). The difficulty is one that usually arises in the analysis of moral hazard problems: the maximand in (15) is not, in general, concave.

3. Differentiation of the expression (17′) for $\bar{m}'(\theta)$ reveals that, in general, $\bar{m}''(\theta)$ may not be positive. Thus, $\bar{m}(\theta)$ may not be convex and, as noted earlier, (15) may not be concave.

REFERENCES

Admati, A. and P. Pfleiderer (1985a), "A Monopolistic Market for Information" Research Paper # 797R, Graduate School of Business, Stanford University, August.

____________________________ (1985b), "Viable Allocations of Information in Financial Markets" Research Paper # 847, Graduate School of Business, Stanford University, November.

Allen, F. (1985a), "Contracts to Sell Information" Working Paper 12-85, Rodney L. White Center, University of Pennsylvania.

_________ (1985b) "Information-Contracting in Financial Markets" forthcoming in S. Bhattacharya and G. Constantinides (editors) Classics in Modern Financial Theory.

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

Blackwell, D. (1951), "The Comparison of Experiments", Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, Berkeley, 93-102.

_____________ (1953), "Equivalent Comparisons of Experiments", Annals of Mathematical Statistics, June, 265-272.

Grossman, S. and O. Hart (1983), "An Analysis of the Principal-Agent Problem", Econometrica, 51, 7-46.

Grossman, S. and J. Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets", American Economic Review, 70, 393-408.

Holmstrom, B. (1979), "Moral Hazard and Observability" The Bell Journalof Economics, Spring, 74-91.
_____________ (1982), "Moral Hazard in Teams" Bell Journal of Economics, 13, Autumn, 324-340.

Kihlstrom, R.E. (1984), "A Simple Example of the Radner-Stiglitz Nonconcavity in the Value of Information" Bayesian Models in Economic Theory, M. Boyer and R. Kihlstrom (editors), North-Holland, Amsterdam.

Lambert, R. (1986), "Executive Effort and Selection of Risky Prospects" RAND Journal Of Economics,17, Spring, 77-88.

Marshall, J. (1976), "Moral Hazard", American Economic Review, 66, December, 880-890.

Merton R. (1981), "On Market Tinming and Investment Performance. I. An Equilibrium Theory of Value for Market Forecasts" Journal of Business, 363-406.

Mirrlees, J. (1974), "Notes on Welfare Economics, Information and Uncertainty" in Balch, McFadden and Wu (editors), Essays on Economic Behavior under Uncertainty, North Holland, Amsterdam.

_____________ (1976), "The Optimal Structure of Incentives and Authority within an Organization" Bell Journal of Economics, 7, Spring.

Myerson, R. (1979), Incentive Compatibility and the Bargaining Problem" Econometrica, 47, 61-73.

Radner, R. and J. Stiglitz (1984), " A Nonconcavity in the Value of Information", Bayesian Models in Economic Theory, M. Boyer and R. Kihlstrom (editors), North-Holland, Amsterdam.

Rogerson, W. (1984), "The Value of Private Information in R&D Contracting" unpublished paper, Stanford University.

Shavell, S. (1979), "Risk Sharing and Incentives in the Principal and Agent Relationship" The Bell Journal of Economics, Spring, 55-73.

Stiglitz, J. (1977) "Monopoly, Non-linear Pricing and Imperfect Information: the Insurance Market" Review of Economic Studies, 44, October, 407-430.

FIGURE 1

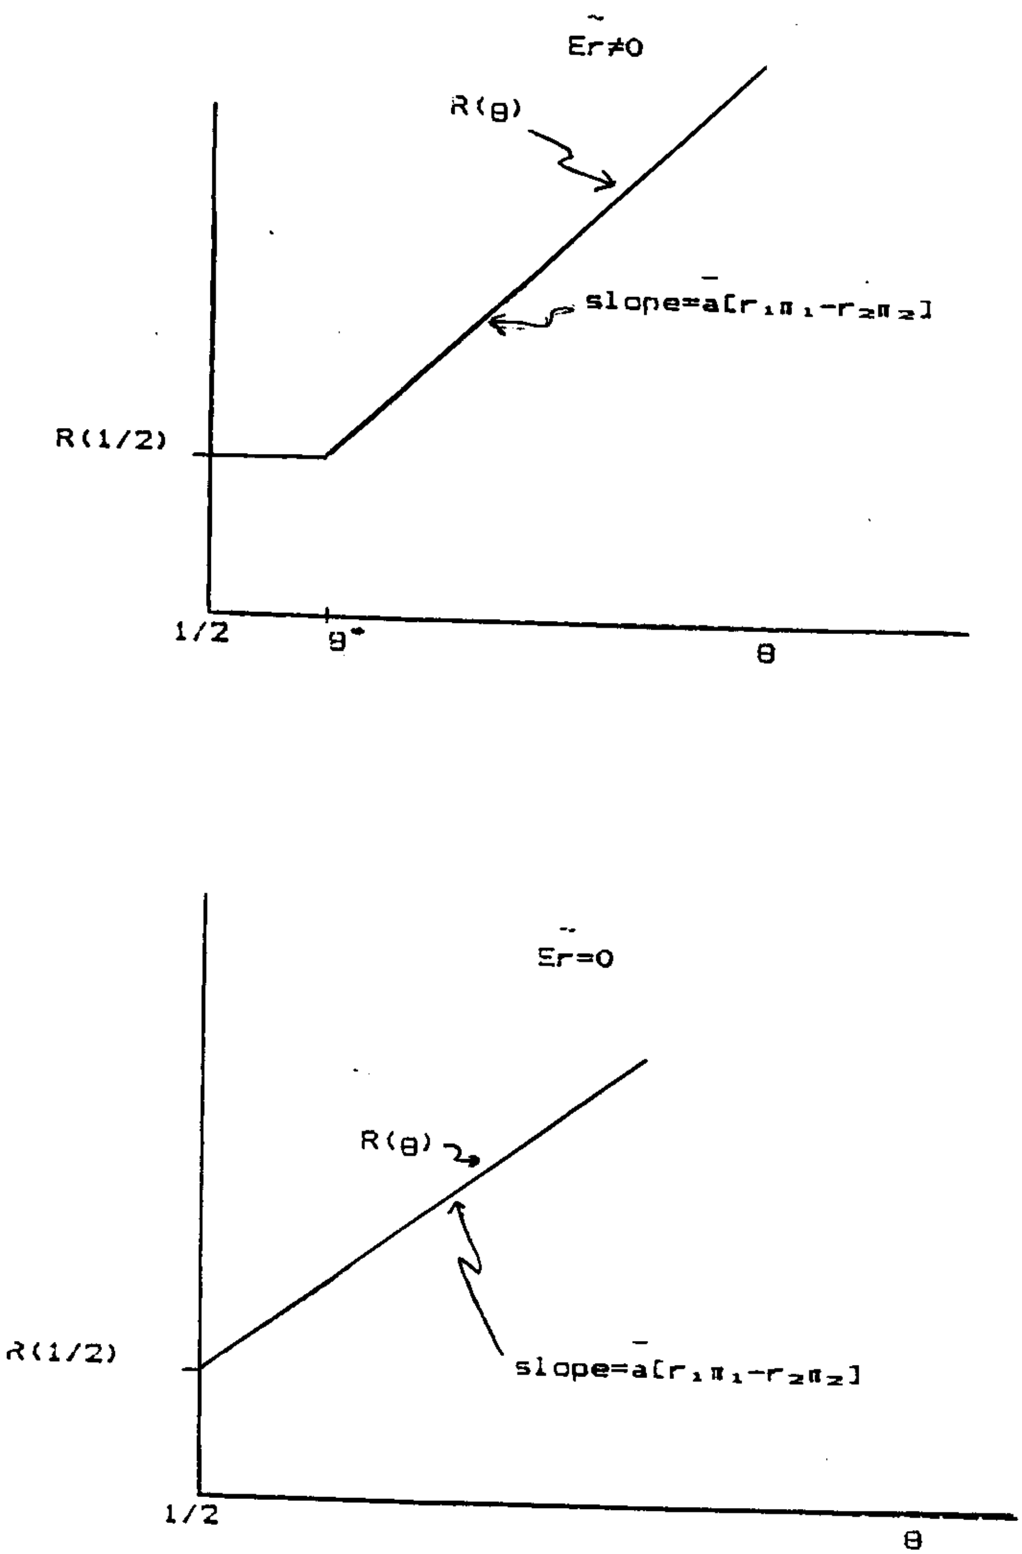

$E\tilde{r}\neq0$
$R(\theta)$
slope=$\bar{a}[r_1\pi_1-r_2\pi_2]$
R(1/2)
1/2
θ^*
θ
$E\tilde{r}=0$
$R(\theta)$
slope=$\bar{a}[r_1\pi_1-r_2\pi_2]$
R(1/2)
1/2
θ

FIGURE 2

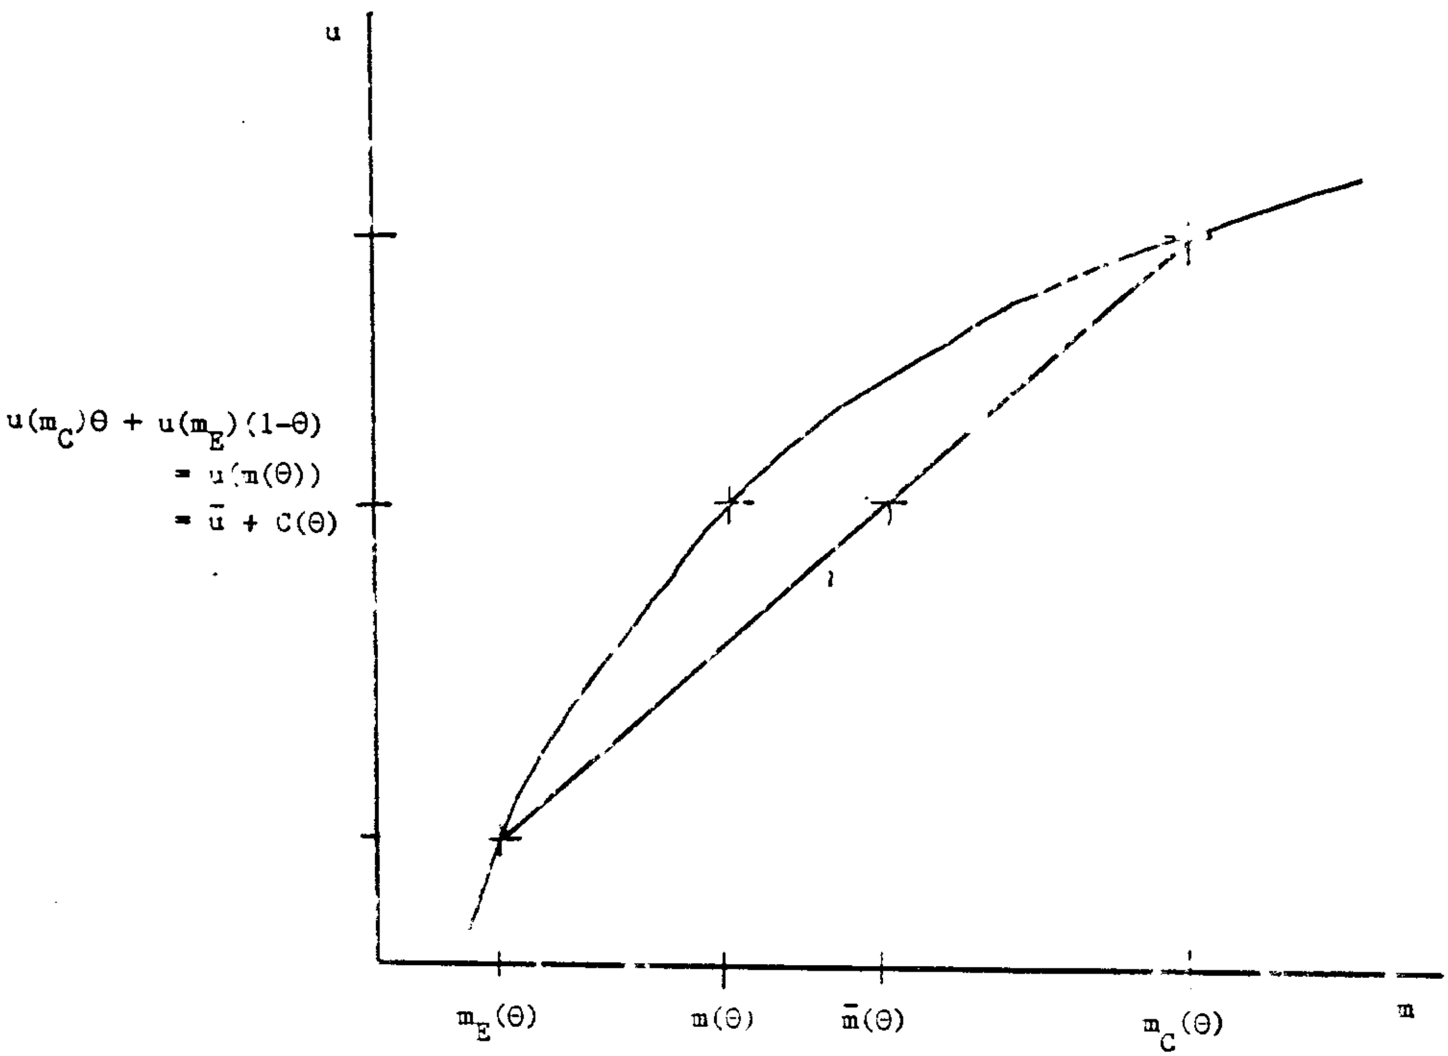

FIGURE 3

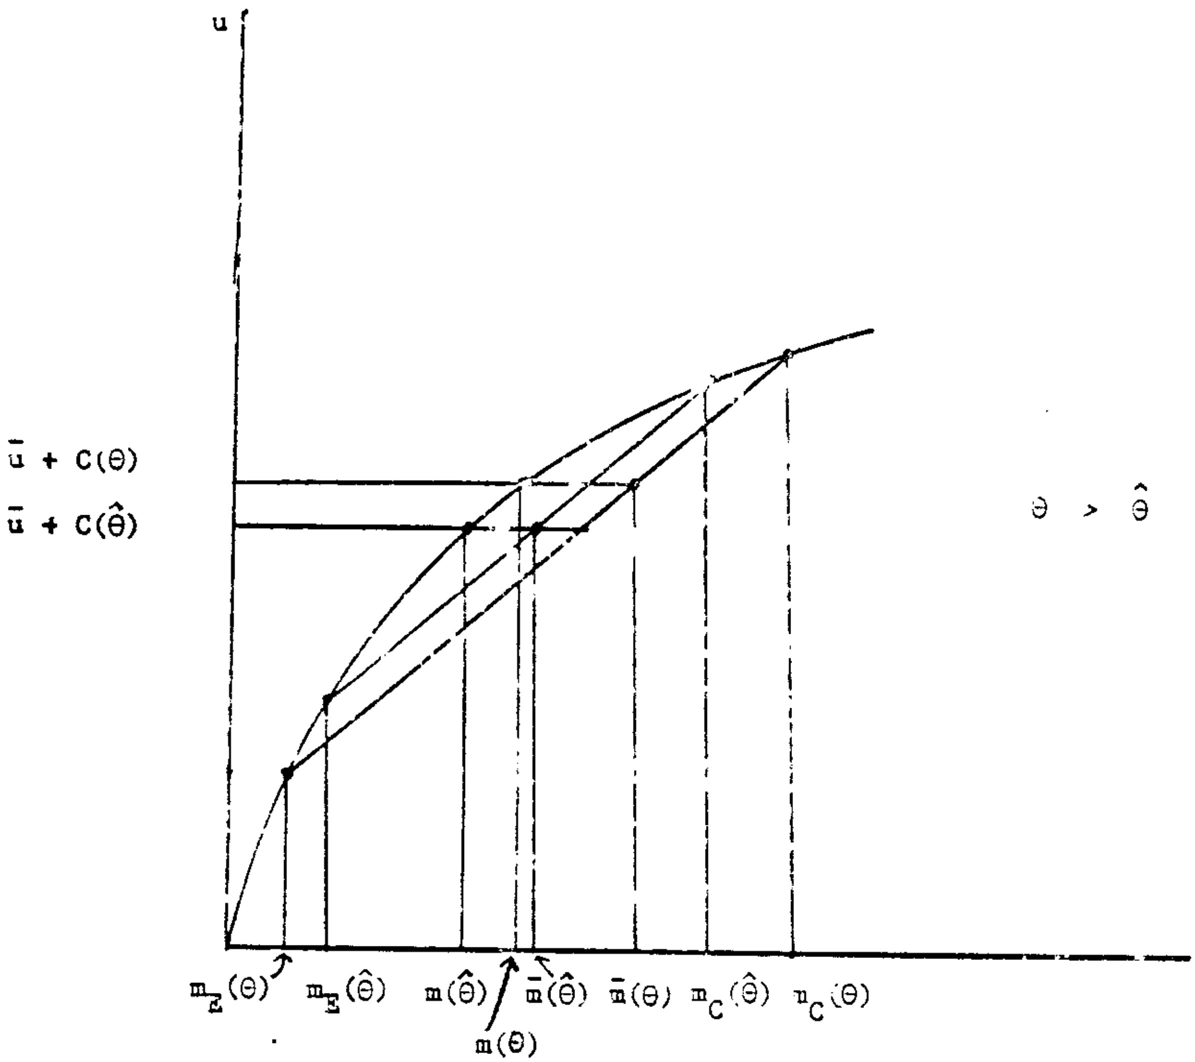

FIGURE 3

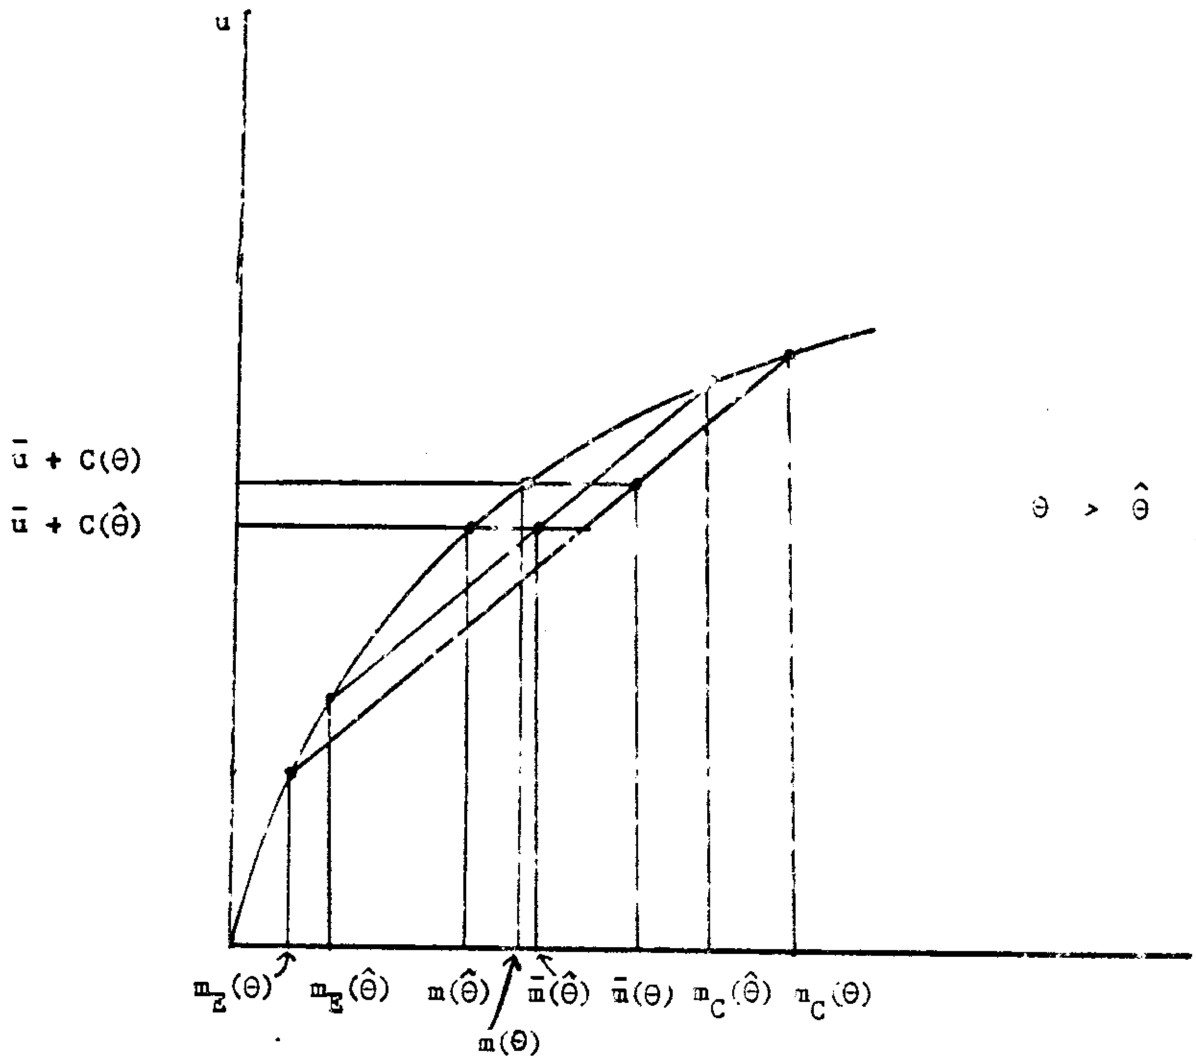

u
$\bar{u}$ + C(Θ)
$\bar{u}$ + C($\hat{\Theta}$)
Θ > $\hat{\Theta}$
m_E(Θ)
m_E($\hat{\Theta}$)
m($\hat{\Theta}$)
$\bar{m}$($\hat{\Theta}$)
$\bar{m}$(Θ)
m_C($\hat{\Theta}$)
m_C(Θ)
m(Θ)