Moral hazard and observability

Bengt Holmström
Swedish School of Economics and Business Administration

The role of imperfect information in a principal-agent relationship subject to moral hazard is considered. A necessary and sufficient condition for imperfect information to improve on contracts based on the payoff alone is derived, and a characterization of the optimal use of such information is given.

1. Introduction

■ It has long been recognized that a problem of moral hazard may arise when individuals engage in risk sharing under conditions such that their privately taken actions affect the probability distribution of the outcome.[1] This situation is common in insurance, labor contracting, and the delegation of decisionmaking responsibility, to give a few examples. In these instances Pareto-optimal risk sharing is generally precluded, because it will not induce proper incentives for taking correct actions. Instead, only a second-best solution, which trades off some of the risk-sharing benefits for provision of incentives, can be achieved.

The source of this moral hazard or incentive problem is an asymmetry of information among individuals that results because individual actions cannot be observed and hence contracted upon. A natural remedy to the problem is to invest resources into monitoring of actions and use this information in the contract. In simple situations complete monitoring may be possible, in which case a first-best solution (entailing optimal risk sharing) can be achieved by employing a forcing contract that penalizes dysfunctional behavior. Generally, however, full observation of actions is either impossible or prohibitively costly. In such situations interest centers around the use of imperfect estimators of actions in contracting. Casual observation indicates that imperfect information is extensively used in practice to alleviate moral hazard, for instance in the supervision of employees or in various forms of managerial accounting.

A natural question then arises: when can imperfect information about actions be used to improve on a contract which initially is based on the payoff alone? Secondly, how should such additional information be used optimally?

This paper is partly based on Chapter 4 of the author's unpublished dissertation, "On Incentives and Control in Organizations," submitted to Stanford University, December 1977. It was written while the author was visiting the Center for Operations Research and Econometrics, Université Catholique de Louvain, Belgium. An earlier version was presented at the European Meeting of the Econometric Society in Geneva, 1978. I am much indebted to Joel Demski, Frøystein Gjesdal, Charles Holloway, David Kreps, and Robert Wilson for many helpful discussions and to David Baron and Gerald Kramer for detailed comments on an earlier manuscript.

[1] See for instance Arrow (1970), Zeckhauser (1970), Pauly (1974), and Spence and Zeckhauser (1971).

A recent interesting paper by Harris and Raviv (1976) addresses these questions in the context of a principal-agent relationship in which the agent provides a productive input (e.g., effort) that cannot be observed by the principal directly.[2] Their results relate to a very specific kind of imperfect monitoring of the agent's action. They study monitors which provide information that is independent of the state of nature and allows the principal to detect any shirking by the agent with positive probability. Such monitors are of limited interest, however, since they are essentially equivalent to observing the agent's action directly, because a first-best solution can be approximated arbitrarily closely in this case.[3] Clearly, one cannot expect imperfect monitoring to possess such strong characteristics in general.

Employing a different problem formulation from Harris and Raviv's, we are able to simplify their analysis and generalize their results substantially. Both questions posed above are given complete answers (in our particular model). It is shown that any additional information about the agent's action, however imperfect, can be used to improve the welfare of both the principal and the agent. This result, which formalizes earlier references to the value of monitoring in agency relationships (Stiglitz, 1975; Williamson, 1975), serves to explain the extensive use of imperfect information in contracting. Furthermore, we characterize optimal contracts based on such imperfect information in a way which yields considerable insight into the complex structure of actual contracts.

The formulation we use is an extension of that introduced by Mirrlees (1974, 1976). We start by presenting a slightly modified version of Mirrlees' model (Section 2), along with some improved statements about the nature of optimal contracts when the payoff alone is observed. In Section 3 a detour is made to show how these results can be applied to prove the optimality of deductibles in accident insurance when moral hazard is present. Section 4 gives the characterization of the optimal use of imperfect information and Section 5 presents the result when imperfect information is valuable. Up to this point homogeneous beliefs are assumed, but in Section 6 this assumption is relaxed to the extent that we allow the agent to be more informed at the time he chooses his action. The analysis is brief, but indicates that qualitatively the same results obtain as for the case with homogeneous beliefs. Section 7 contains a summary and points out some directions for further research.

2. Optimal sharing rules when the payoff alone is observed

■ We study a principal-agent relationship, where the agent privately takes an action $a \in A \subseteq \mathcal{R}$, A being the set of all possible actions, and a together with a random state of nature θ, determines a monetary outcome or payoff $x = x(a, \theta)$. The problem is to determine how this payoff should be shared optimally between the principal and the agent. The principal's utility function is $G(w)$, defined over wealth alone, and the agent's utility function is $H(w, a)$, defined over wealth

[2] The main results of Harris and Raviv (1976) are reported in their 1978 paper. For earlier work on principal-agent models, see Wilson (1969), Ross (1973), and Mirrlees (1976).

[3] This fact, which is not observed by Harris and Raviv (1976), can be verified by using an argument similar to the one given by Mirrlees (1974, p. 249), or by Gjesdal (1976) (cf. example in footnote 7). Obviously, it implies that monitoring, which satisfies Harris and Raviv's conditions, is valuable. This is their partial answer to the first question raised above.

and action. The model is further restricted by assuming that $H(w,a) = U(w) - V(a)$, with $V' > 0$ and $x_a \geqslant 0$.[4] The interpretation is that a is a productive input with direct disutility for the agent and this creates an inherent difference in objectives between the principal and the agent. It is convenient to think of a as effort and this term will be used interchangeably with action. Since the problem of moral hazard can be avoided when the agent is risk-neutral (Harris and Raviv, 1976), we shall assume $U'' < 0$. The principal may or may not be risk-neutral, i.e., $G'' \leqslant 0$.

In this section, we consider the case where the principal observes only the outcome x. Thus, sharing rules have to be functions of x alone. Let $s(x)$ denote the share of x that goes to the agent and $r(x) = x - s(x)$ denote the share that goes to the principal. It is assumed that both parties agree on the probability distribution of θ and that the agent chooses $\underset{\sim}{a}$ before θ is known.[5] In this case (constrained) Pareto-optimal sharing rules $s(x)$ are generated by the program:

$$\max_{s(x),a} E\{G(x - s(x))\} \tag{1}$$

$$\text{subject to } E\{H(s(x),a)\} \geq \bar{H}, \tag{2}$$

$$a \in \underset{a' \epsilon A}{\text{argmax}}\, E\{H(s(x),a')\}, \tag{3}$$

where the notation "argmax" denotes the set of arguments that maximize the objective function that follows.[6]

Constraint (2) guarantees the agent a minimum expected utility (attained via a market or negotiation process). Constraint (3) reflects the restriction that the principal can observe x but not $\underset{\sim}{a}$. If he also could observe $\underset{\sim}{a}$, a forcing contract could be used to guarantee that the agent selects a proper action even when $s(x)$ is chosen to solve (1)–(2) ignoring (3). The latter we will refer to as the *first-best solution*, which entails optimal risk sharing. It differs in general from the solution of (1) subject to (2) *and* (3), which we call a *second-best solution*.

Two approaches can be used to solve the program above. The earlier one, used by Spence and Zeckhauser (1971), Ross (1973), and Harris and Raviv (1976), recognizes explicitly the dependence of x on $\underset{\sim}{a}$ and θ, so that the expectations in (1)–(3) are taken with respect to the distribution of θ. They proceed to characterize an optimal solution by replacing (3) with the first-order constraint $E\{H_1 \cdot s' \cdot x_a + H_2\} = 0$, and then apply the calculus of variations. To validate these steps one has to assume that an optimum exists and is differentiable. However, as an example by Mirrlees (1974) shows, there may commonly exist no optimal solution among the class of unbounded sharing rules, and for this reason $s(x)$ has to be restricted to a finite interval in general. As a result, the solution will become nondifferentiable and the above-mentioned approach can no longer be applied.[7]

---

[4] Subscripts denote partial derivatives with respect to corresponding variables.

[5] This assumption corresponds to model 1 in Harris and Raviv (1976), which is the model they use for studying imperfect information. We shall relax it in Section 6.

[6] As usual, E denotes the expectation operator. Since $E\{H(s(x),a)\}$ need not be concave in a, there may exist multiple solutions, hence the inclusion symbol.

[7] Even when an optimal solution exists among unbounded sharing rules, it may be nondifferentiable. This has been observed by Gjesdal (1976). To illustrate his ideas one can look at the follow-

A better approach to solving (1)–(3), which also gives a more intuitive characterization of an optimum, has been introduced by Mirrlees (1974, 1976). He suppresses θ and views x as a random variable with a distribution $F(x,a)$, parameterized by the agent's action. Given a distribution of θ, $F(x, a)$ is simply the distribution induced on x via the relationship $x = x(a,\theta)$.[8] It is easy to see that $x_a \geqslant 0$ implies $F_a(x, a) \leqslant 0$. It will be assumed that for every $\underset{\sim}{a}$, $F_a(x, a) < 0$ for some x-values, so that a change in $\underset{\sim}{a}$ has a nontrivial effect on the distribution of x. In particular, it will shift the distribution of x to the right in the sense of first-order stochastic dominance.

For the moment, assume F has a density function $f(x, a)$ with f_a and f_{aa} well defined for all (x, a).[9] Replacing (3) with a first-order constraint yields the program:

$$\max_{s(x)\in[c,d+x],a} \quad \int G(x - s(x))f(x,a)dx \tag{4}$$

$$\text{subject to} \quad \int [U(s(x)) - V(a)]f(x,a)dx \geq \bar{H}, \tag{5}$$

$$\int U(s(x))f_a(x,a)dx = V'(a). \tag{6}$$

Note that $s(x)$ is restricted to lie in the interval $[c, d + x]$ to avoid nonexistence of a solution.[10] This restriction is natural from a pragmatic point of view as well, since the agent's wealth puts a lower bound, and the principal's wealth (augmented with x) an upper bound on $s(x)$.

Let λ be the multiplier for (5) and μ the multiplier for (6). Pointwise optimization of the Lagrangian yields the following characterization of an optimal sharing rule:

$$\frac{G'(x - s(x))}{U'(s(x))} = \lambda + \mu \cdot \frac{f_a(x,a)}{f(x,a)}, \tag{7}$$

---

ing insightful example. Let $x(a,z) = a + z$ and $z \sim \text{Unif}(0,1)$, so that $x \sim \text{Unif}(a,a + 1)$. If $(a^*,s^*(x))$ is a first-best solution it is easy to see that a contract of the form $s(x) = s^*(x)$ when $x \geqslant a^*$, $s(x) = w$ otherwise, will make the agent choose $a = a^*$ for w sufficiently low. But in that case $x \geqslant a^*$ for all outcomes of $\tilde{z}$, and the first-best solution $s(x) = s^*(x)$ is effectively realized. In other words, a nondifferentiable sharing rule, which penalizes the agent for outcomes $x < a^*$, will give both the principal and the agent the same expected utility as a first-best solution. In this example no optimal differentiable sharing rule exists for (1)–(3).

Gjesdal's analysis shows that both Spence and Zeckhauser (1971, p. 383, footnote 5) and Harris and Raviv (1976, pp. 36–37) err in giving incorrect characterizations (based on the Euler equation) for examples similar to this.

We will avoid situations like these by essentially assuming that the support of the distribution of x will not change with a, as explained below. For a more detailed comparison of the state-space approach with Mirrlees' approach, see Holmström (1977).

[8] Thus, it is always possible to go from the state space approach to Mirrlees' approach, while the reverse is not always true.

[9] In Section 3 we shall allow discrete distributions as well. The crucial assumption is that f_a exists. Note that this assumption is not satisfied by the example in footnote 7.

[10] More precisely, existence of a solution to (1)–(3) can be proved for the class of functions:

$$S_K = \{s(x) \in [c,d + x] \mid V_b^{b'}(s) \leq K\cdot(b' - b)\},$$

where $V_b^{b'}(s)$ is the total variation of s in the interval $[b,b']$ (Kolmogorov and Fomin, 1970), under some technical assumptions about integrability and the behavioral assumption that the agent, for almost every x for which (7) has a solution $s(x) \in [c, d + x]$; otherwise $s(x) = c$ or $d + x$, depending on whether the right-hand side $\lesseqgtr$ left-hand side throughout the interval. Furthermore, μ is given as the solution to the adjoint equation,

$$\int G(x - s(x))f_a(x,a)dx + \mu \cdot \left\{ \int U(s(x))f_{aa}(x,a)dx - V''(a) \right\} = 0, \quad (8)$$

and $\underset{\sim}{a}$ is determined by (6).[11]

From Borch's (1962) work, we know that $s(x)$ will be Pareto optimal from a risk-sharing point of view only if the right-hand side in (7) is constant. Now, $f_a/f = k$, a constant, implies $0 = \int f_a = \int f \cdot k = k$, since $\int f = 1$ for all $\underset{\sim}{a}$. Hence, $f_a \equiv 0$ would follow, which contradicts the assumption that $F_a < 0$ for some x. Consequently, perfect risk sharing could only obtain if $\mu = 0$. But, in fact, one can prove the following:[12]

Proposition 1. Assume $V' > 0$ and $F_a \leq 0$ (with strict inequality for some x-values), then $\mu > 0$, or equivalently: The principal would like to see the agent increase his effort given the second-best sharing rule.

Proof: See Appendix.

Two immediate corollaries follow:

Corollary 1. Under the assumption of Proposition 1, one has the following relationship between the second-best solution $s(x)$ and the first-best solution $s_\lambda(x)$, for a given λ:

$$\begin{cases} s(x) \geq s_\lambda(x), & \text{on} \quad X_+ = \{x \,|\, f_a(x,a) \geqslant 0\}, \\ s(x) < s_\lambda(x), & \text{on} \quad X_- = \{x \,|\, f_a(x,a) < 0\}. \end{cases} \quad (9)$$

Proof: See Appendix.

Corollary 2. Under the assumption of Proposition 1, the second-best solution is strictly inferior to a first-best solution.[13]

Proof: See Appendix.

---

in case he is indifferent, chooses his action according to the principal's preferences. By taking K large enough, the characterization in (6)–(8) will be valid for this solution, and S_K will contain all functions of practical relevance.

[11] The characterization can be proved rigorously as in Holmström (1977) using proposition 9.6.1 in Luenberger (1969). Some technical assumptions which we do not spell out are needed. More important is the fact that one has to assume that the agent's optimal choice of action is unique for the optimal $s(x)$. This assumption seems very difficult to validate except in specific problems and regrettably we have to leave the question about its validity open.

Mirrlees (1974) was the first to give a characterization of an optimum in the form above (without bounds). Earlier Spence and Zeckhauser (1971) and Ross (1973) gave alternative characterizations based on the state space formulation.

[12] This proposition generalizes Mirrlees' (1976) conclusion that $\mu > 0$ when f_a/f is increasing in x.

[13] It is worthwhile stressing the difference between Corollary 2 and the example in footnote 7. A first-best solution can be achieved in that example because f_a does not exist at the endpoints of the uniform distribution. Whenever f_a exists, Corollary 2 indicates that a first-best solution cannot be achieved. Also, note that $V' > 0$ is essential. The role played by V in the characterization is obscured by the complexity of the relationships between (6)–(8), but generally one expects that the larger V'' is, the smaller is μ and the accompanying welfare loss. At an extreme, if $V' = 0$ for $a \leq \bar{a}$ and $V' = \infty$ for $a > \bar{a}$ then the first-best outcome can be achieved since it entails $a = \bar{a}$, which the agent will choose given an optimal risk-sharing rule.

The characterization in (7) has an intuitive interpretation in terms of deviating from optimal risk sharing to provide incentives for increased effort on the part of the agent. This is accomplished by taking $s(x) \geqslant s_\lambda(x)$ when the marginal return from effort is positive to the agent, and $s(x) < s_\lambda(x)$ when it is negative (see Corollary 1). The incentive effect of deviating from optimal risk sharing is stronger the larger is $|f_a|$, and it is more costly (in terms of lost risk-sharing benefits) the greater is f. Thus $|f_a|/f$ may be interpreted as a benefit-cost ratio for deviation from optimal risk sharing, and (7) states that such deviations should be made in proportion to this ratio, with individual risk aversion taken into account.

In contrast to perfect risk sharing, the second-best solution is crucially dependent on the distribution of x and its functional relation to $\underset{\sim}{a}$. This occurs because the outcome x can be used as a signal about the action which is not directly observed. We note that f_a/f is the derivative of the maximum likelihood function $\log f$, when $\underset{\sim}{a}$ is viewed as an unknown parameter. In this sense f_a/f measures how strongly one is inclined to infer from x that the agent did not take the assumed action, and (7) says that penalties or bonuses (as expressed by deviations from first-best risk sharing) should be paid in proportion to this measure.[14]

The deviation from perfect risk sharing implies that the agent is forced to carry excess responsibility for the outcome and this points to the implicit costs involved in contracting under imperfect information (Corollary 2). Consequently, there are positive gains to observing the agent's action, since in that case a first-best solution can be achieved by using a forcing contract. This provides the basis for discussing ways to realize part of these gains by using imperfect monitoring, which is the subject of Sections 4 and 5.

To illustrate the formula in (7) and the interpretations, consider the following example: $G(w) = w$, $U(w) = 2\sqrt{w}$, $V(a) = a^2$, $x \sim \exp(1/a)$. In this example, the agent could be a machine repairman, whose effort $\underset{\sim}{a}$ will determine the expected time before the machine will break down. The monetary return x is proportional to the length of time the machine will remain operative; (here the proportionality factor has been taken = 1).

From (7), the optimal share is:[15]

$$s(x) = \left[\lambda + \mu \cdot \frac{(x - a)}{a^2}\right]^2, \tag{10}$$

and some simple calculations yield $\mu = a^3$, and the equation $4a^3 + 2\lambda \cdot a = 1$ for $\underset{\sim}{a}$ (using (6) and (8)). As one would expect, μ is increasing in $\underset{\sim}{a}$, since it is more costly to induce higher effort. The first-best solution is $s_\lambda(x) = \lambda^2$, $a_\lambda = 1/2\lambda$.[16]

For a numerical solution let $\lambda = \frac{1}{2}$. Then $s(x) = \frac{1}{4}(x + \frac{1}{2})^2$, $a = \frac{1}{2}$, $s_\lambda(x) = \frac{1}{4}$, $a_\lambda = 1$, as pictured below. The welfare measure for the first-best solution is $\frac{3}{4}$ and for the second-best it is $\frac{9}{16}$. (See Figure 1.) In this example,

[14] Of course, this interpretation is not quite accurate, since given $s(x)$, the principal knows from (4) what action a rational agent will take according to the model. But I think this interpretation corresponds well with reasoning in practice.

[15] This is an exceptional example in that no bounds need to be imposed on the sharing rule, and an explicit solution can thus be obtained.

[16] In this example the question of uniqueness, referred to in footnote 11, is no problem. For any fixed a, (10) gives the appropriate solution to (4)–(6), which is a relaxation of (1)–(3) when a is fixed at its optimal value. But it is easily checked that $s(x)$ in (10) makes (6) strictly concave in a, and hence the relaxed problem also solves the original one.

FIGURE 1

THE SECOND-BEST SOLUTION FOR THE REPAIRMAN EXAMPLE WITHOUT MONITORING

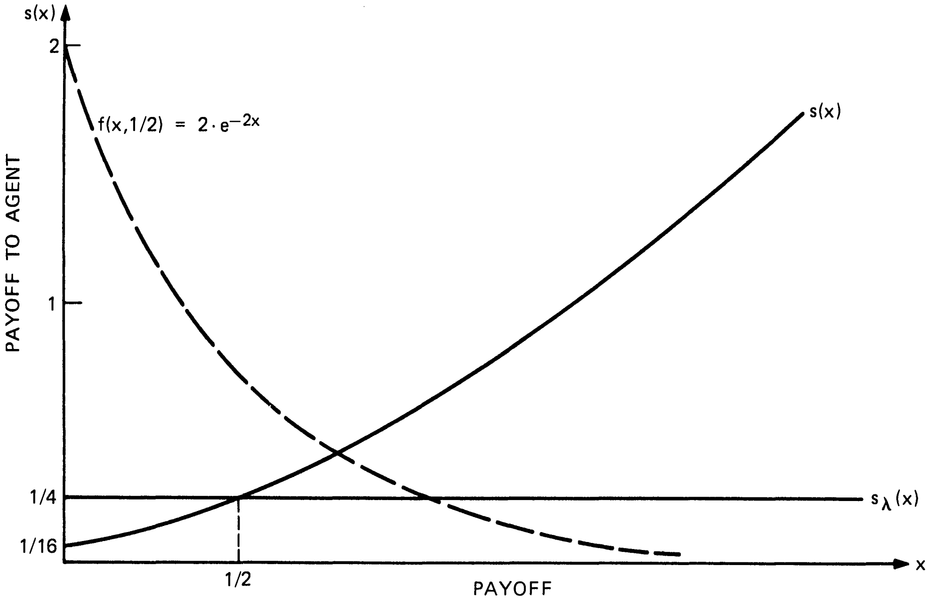

the penalties imposed on the agent for $x < \frac{1}{2}$, which is the mean of x, are relatively small owing to the high values of $f(x, \frac{1}{2})$ in this region (and, of course, owing to risk aversion), and bonuses for $x > \frac{1}{2}$ are correspondingly large. In view of risk-sharing benefits, the convexity of the second-best solution may be surprising, but this is in no way exceptional (cf. Mirrlees (1976); also in Wilson (1969), convex sharing rules may be optimal). Examples for which sharing rules are concave or linear or even two-peaked can be easily generated as well.

3. Deductibles in insurance

■ The characterization in (7) can be applied to the insurance setting to conclude that optimal accident insurance policies necessarily entail deductibles in the presence of moral hazard. To demonstrate this, the assumption that the distribution $F(x,a)$ possesses a density function will be relaxed. Since (7) is derived via point-wise optimization, a mixture of a continuous and discrete distribution can be used as well, provided the support of the discrete distribution is left unchanged by the action (cf. footnote 7). In that case simply interpret $f(x,a)$ in (7) as the probability mass rather than the value of the density function whenever x is a mass point (and correspondingly for f_a and f_{aa}, which both are assumed to exist as before).

Mixed distributions are characteristic in accident insurance. First, there is a probability that no accident occurs and this generates a mass point at $x = 0$; and conditional on an accident, there is a damage distribution over $x < 0$, which usually can be assumed continuous. If $\underset{\sim}{a}$ represents a precautionary action, it is natural to assume that this mixed distribution satisfies:

$$f_a(0,a) > 0, f_a(x,a) < 0. \tag{11}$$

This assumption says that the probability of an accident decreases with $\underset{\sim}{a}$ so that each outcome $x < 0$ is less likely. For instance, driving a car more carefully will presumably decrease the probability of both small and large accidents.

Because $\mu > 0$ and the left-hand side in (7) is continuous, (11) clearly implies that the optimal sharing rule $s(x)$ is discontinuous at $x = 0$. In fact, $s(0) > s_\lambda(0) > s(x)$ for all $x < 0$, since $G'(x - s(x))/(U'(s(x))$ is increasing in $s(x)$ and nonincreasing in x (here s_λ is the solution to (7) with $\mu = 0$). If $d = \min_{x<0} \{s(0) - s(x)\} > 0$, we can write:[17]

$$s(x) = \begin{cases} k, & \text{if} \quad x = 0, \\ k - d - t(x), & \text{if} \quad x < 0, \end{cases} \tag{12}$$

where k is the agent's wealth after paying the premium, d is the deductible which is paid when an accident occurs, and $t(x) \geq 0$ is the agent's additional share in the costs of an accident. One would expect $t(x)$ to be increasing in x. This is the case if, for instance, f_a/f is increasing in x (which holds for surprisingly many standard distributions; see Holmström (1977)).

In many situations it is approximately true that the agent's action will only affect the probability of an accident and not the size of losses, given that an accident occurs. In that case one can write $f(0,a) = 1 - p(a), f(x,a) = p(a)\cdot g(x)$, $x < 0$, where $p(a)$ is the probability of an accident ($p' < 0$) and $g(x)$ is a damage distribution independent of $\underset{\sim}{a}$. This implies $f_a(0,a)/f(0,a) = -p'(a)/(1 - p(a)) > 0$ and $f_a(x,a)/f(x,a) = p'(a)/p(a) < 0$, for $x < 0$. Hence f_a/f is independent of x for $x < 0$, which means that for $x < 0$ we have first-best risk sharing. In particular, if the insurance company is risk-neutral, only a deductible will be charged when an accident occurs.

To summarize the discussion we have:

Proposition 2. Given the assumptions in (11), optimal accident insurance policies entail a deductible. If the insured's action only affects the probability of an accident but not the size of damage and the insurance company is risk-neutral, a deductible alone is optimal.

This proposition lends additional support to the frequent use of deductibles in accident insurance. However, the reasoning is quite different from that behind the well-known proposition by Arrow (1970), which holds that pure deductibles are always optimal. Arrow does not consider moral hazard aspects, and in his case deductibles arise for instance if the firm uses loading to determine the premium (Mossin, 1968).

4. Optimal sharing rules based on additional information

■ One of the main conclusions from Section 2 is that the optimal solution under moral hazard is not first-best and, hence, that there would be gains to observing the agent's action (see Corollary 2 and the subsequent discussion). Since perfect observation of the agent's action is generally precluded, interest centers on the use of imperfect information for improvements of the contract. This issue can be studied using a straightforward extension of the model in Section 2.

---

[17] If x can be observed only at the option of the insured, (12) is not enforceable. In that case, the optimal contract is $\bar{s}(x) = \max(x, s(x))$, with $s(x)$ as in (12). This is still a contract with a deductible.

Let y be a signal (possibly vector-valued), which in addition to x, is observed by both parties and hence can be used in constructing the sharing rule. Let $F(x,y,a)$ be the joint distribution of x and y given $\underset{\sim}{a}$. As in Section 3, let $f(x,y,a)$ be either the value of the density function of the continuous part of F or the probability mass of the mass point (x,y), if such exists. As before, f_a and f_{aa} are assumed to exist. The following extension of (7) obtains for an optimal sharing rule $s(x,y)$:

$$\frac{G'(x - s(x,y))}{U'(s(x,y))} = \lambda + \mu \cdot \frac{f_a(x,y,a)}{f(x,y,a)}, \tag{13}$$

for almost every (x,y) such that (13) has a solution $s(x,y) \in [c, d + x]$; otherwise $s(x,y) = c$ or $d + x$ depending on whether the right-hand side $\gtrless$ the left-hand side throughout the interval. Here μ is the multiplier of the agent's first-order constraint and satisfies (8) and $\underset{\sim}{a}$ satisfies (6) (with obvious changes in notation).

Again $\mu > 0$ follows as in Proposition 1, and consequently the second-best solution $s(x,y)$ will be strictly worse than a first-best solution.[18] The interpretation of f_a/f in Section 2 can be repeated for (13). A new, important feature, however, is that $f_a(x,y,a)/f(x,y,a)$ may change with y. Thus, for the same value of x, but under different contingencies signalled by y, the agent should generally receive different remuneration. In particular, if for one value of y it is possible to infer less about $\underset{\sim}{a}$ via x, then the deviation from optimal risk sharing should be smaller, and *vice versa*. At an extreme, a realization of the signal y could be such that $f_a(x,y,a) \equiv 0$ for all x (which means that nothing about the action can be inferred from the payoff), and in this case the optimal risk-sharing rule should be employed. In sharecropping, for example, if a natural disaster destroys the crop, farm workers should not be held responsible for the outcome (beyond optimal risk sharing).

This is quite intuitive and corresponds well with observed practice. Equation (13) would predict that contracts are elaborate and contain a variety of provisions for unexpected events. Certainly, there is substantial empirical support for this conclusion. Contracts, at least between external parties, tend to be detailed, spelling out different responsibilities in different contingencies (e.g., strikes, accidents, natural disasters, etc.). Not doing so would be inefficient and add to the implicit costs of contracting. In the same way managers are not held responsible for events one can observe are outside their control, and implicitly at least, their performance is always judged against information about what should be achievable given, say, the current economic situation.[19]

To illustrate the point we can look at an extension of the example in Section 2. Suppose now that the machine can also break down because of a failure in a component over which the repairman has no control. Let this event have an exponential probability distribution with constant parameter $(1/k)$, and assume it is independent of the event that the components which the repairman controls will break down. The latter event still has the same probability distribution as before, namely exponential with parameter $(1/a)$.

---

[18] As in Section 2, this result depends crucially on the assumption that f_a exists for all (x,y,a); cf. footnote 13.

[19] Note, however, that internal labor contracts rarely contain explicit reference to monitoring information, and presumably this information is often unknown to the agent. Yet such information is and should be used. The reason the principal (i.e., the firm) will not default on such an implicit contract is its concern for reputation in the labor market.

If it is not possible to determine whether the failure occurred in a component outside the repairman's control, the optimal solution is to employ a sharing rule:

$$s(x) = \left(\lambda + \mu_1 \cdot \left[\frac{x}{a^2} - \frac{k}{a(a+k)}\right]\right)^2. \tag{14}$$

This follows from (7), since $x \sim \exp((a + k)/ak)$.

On the other hand, if one can determine which component failed, this information can be used to improve the contract. Set $y = 1$ if the failure was outside the repairman's control and $y = 0$ otherwise. Employing (13), one has:[20]

$$\begin{aligned} s(x,0) &= \left(\lambda + \mu_2 \cdot \left[\frac{x}{a^2} - \frac{1}{a}\right]\right)^2, \\ s(x,1) &= \left(\lambda + \mu_2 \cdot \frac{x}{a^2}\right)^2. \end{aligned} \tag{15}$$

Here $\mu_1 = a^3(1 + (a/k))^2$, $\mu_2 = a^3(1 + (a/k))$ from (6). Hence, $\mu_1 > \mu_2$, indicating that it is more costly to induce a particular action a when y cannot be observed.

The interesting comparison is between $s(x,0)$ and $s(x,1)$. One can see that $s(x,0)$ is a translation of $s(x,1)$ to the right as indicated in Figure 2. Confirming our intuition, the repairman receives higher pay if it is found that the failure was outside his control than if it is found that a component that he controls failed. The optimal solution when y is not observed will lie initially between $s(x,0)$ and $s(x,1)$ and eventually go above $s(x,1)$, since $\mu_1 > \mu_2$. Notice that as $k \to \infty$, $s(x) \to s(x,0)$, since it becomes all the less likely that the failure will be caused by anything outside the repairman's control.

5. Value of information

■ Before proceeding with a discussion of the value of imperfect information, the notion of a valuable signal needs to be made precise. A signal y is said to be *valuable* if both the principal and the agent can be made strictly better off with a contract of the form $s(x,y)$ than they are with a contract of the form $s(x)$.

Equation (13) suggests that y will be valuable if and only if it is *false* that

$$\frac{f_a(x,y,a)}{f(x,y,a)} = \bar{h}(x,a), \tag{16}$$

for almost every (x,y). The reason is that when (16) holds, a contract $s(x)$ will satisfy (13), whereas if (16) is false, it must necessarily take the form $s(x,y)$. We shall prove this proposition formally below as it is the main result of the paper and no proof of (13) was given. Before doing so, let us rewrite (16) in a way which allows a surprisingly simple interpretation of this necessary

[20] Simple calculations show that

$$f(x,0,a) = \frac{1}{a} \cdot \exp\left\{-\frac{a+k}{ak} \cdot x\right\}$$

and

$$f(x,1,a) = \frac{1}{k} \cdot \exp\left\{-\frac{a+k}{ak} \cdot x\right\}.$$

FIGURE 2

THE SECOND-BEST SOLUTION FOR THE REPAIRMAN EXAMPLE WITH MONITORING

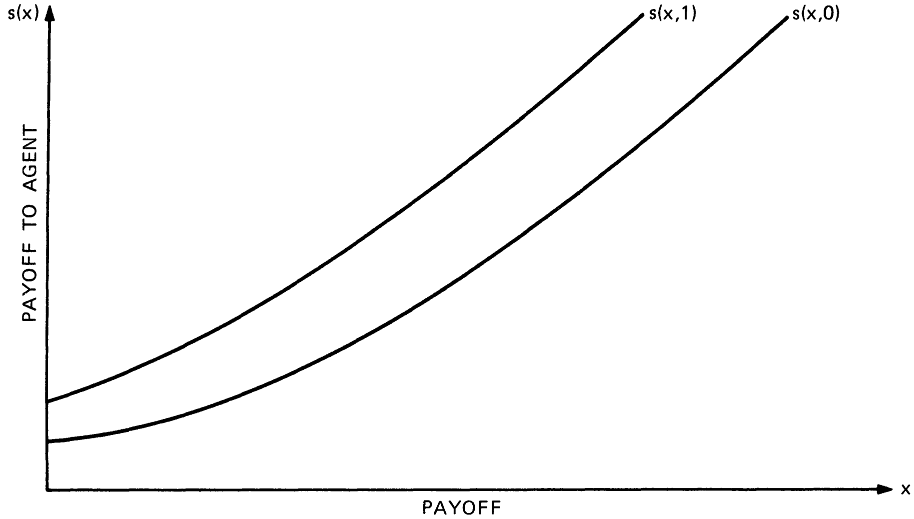

and sufficient condition. Suppose (16) holds for all $\underset{\sim}{a}$. Solving it as a differential equation in $\underset{\sim}{a}$ yields

$$f(x,y;\, a) = g(x,y)\cdot h(x,a), \qquad \text{for almost every } (x,y), \tag{17}$$

where h and g can be taken nonnegative. Conversely, (17) implies (16). Thus (16) and (17) are equivalent.[21]

Equation (17) has a natural interpretation, since it is precisely the condition for a sufficient statistic, if one views $\underset{\sim}{a}$ as a random parameter (de Groot, 1970). That is, when (17) holds, x is a sufficient statistic for the pair (x,y) with respect to $\underset{\sim}{a}$, which means that x carries all the relevant information about $\underset{\sim}{a}$, and y adds nothing to the power of inference. The signal y could only be used for risk-sharing purposes, but optimal risk sharing is independent of the distribution of the random variables when agents have homogeneous beliefs. Consequently, y should be valueless when (17) holds, which is what (13) says. On the other hand, when (17) is false, y contains some information about $\underset{\sim}{a}$ beyond that conveyed by x. In accordance with (13), y should then be used in the contract to improve welfare.

This discussion suggests the following:

Definition: A signal y is said to be *informative* about $\underset{\sim}{a}$ when (17) is false, and *noninformative* otherwise.

With this definition the main result can be stated as follows:

Proposition 3. Let $s(x)$ be an optimal sharing rule for which the agent's choice of action is unique and interior in A. Then there exists a sharing rule $s(x,y)$ which strictly Pareto dominates $s(x)$ if and only if (17) is false; or more concisely, *a signal is valuable if and only if it is informative*.

---

[21] This is not necessarily true if (16) only holds for a single value of a, because then we cannot integrate (16) to get (17). Such an exceptional case is of little interest, however, and in the subsequent analysis, we will only deal with distributions for which (16) is true for either all a or no a.

Proof: Suppose y is noninformative. Then, if $s(x,y)$ is an arbitrary sharing rule, a sharing rule $s(x)$ which is at least as good as $s(x,y)$ will be constructed, establishing the claim that y is of no value.

For every x, define $s(x)$ so that

$$\int U(s(x,y))g(x,y)dy = \int U(s(x))g(x,y)dy = U(s(x))\cdot \int g(x,y)dy. \quad (18)$$

Then using (17) and (18),

$$\int U(s(x,y))f(x,y;\, a)dxdy = \int U(s(x,y))h(x;\, a)g(x,y)dxdy = \int U(s(x))h(x;\, a)g(x,y)dxdy.$$

Consequently, $s(x)$ will result in the same action and welfare for the agent. By Jensen's inequality, (18) implies

$$\int s(x,y)g(x,y)dy \geqslant \int s(x)g(x,y)dy,$$

or

$$\int (x - s(x,y))g(x,y)dy \leqslant \int (x - s(x))g(x,y)dy.$$

This implies, using Jensen's inequality a second time, that:

$$\int G(x - s(x,y))g(x,y)dy \leqslant \int G(x - s(x))g(x,y)dy.$$

Since this is true for every x, and $h(x;\, a) \geqslant 0$, one obtains, by integrating,

$$\int G(x - s(x,y))f(x,y;\, a)dxdy \leqslant \int G(x - s(x))f(x,y;\, a)dxdy.$$

Since the agent takes the same act with $s(x)$ as with $s(x,y)$ by construction, this shows that the principal is at least as well off with $s(x)$ as with $s(x,y)$. The agent's utility is the same for both $s(x)$ and $s(x,y)$, and thus $s(x)$ is weakly Pareto superior to $s(x,y)$, which proves the first part of the proposition.

To prove the second part, let $s(x)$ be a second-best solution with the properties assumed in the proposition. Fix x for a moment. Since the agent's response is unique and interior in A, the principal's and the agent's marginal returns δE^P- and δE^A-conditional on x, from an additive variation $\delta s(x,y)$ in the sharing rule $s(x)$, are [see proposition 9.6.1 in Luenberger (1969)]:

$$\delta E^P = -G'(x - s(x))\int \delta s(x,y)f(x,y;a)dy + \mu\cdot U'(s(x))\int \delta s(x,y)f_a(x,y;a)dy,$$

$$\delta E^A = U'(s(x))\int \delta s(x,y)f(x,y;\, a)dy. \quad (19)$$

Here μ is the solution to (8) corresponding to $s(x)$.

Suppose y is informative. From (16) it follows that there exists a set Y in the range of y, with $\int_Y f(x,y; a)dy \equiv f(x,Y; a) \neq 0$, and correspondingly for the complement Y^c, such that:

$$\frac{f_a(x,Y; a)}{f(x,Y; a)} > \frac{f_a(x,Y^c; a)}{f(x,Y^c; a)}, \tag{20}$$

Choose a variation $\delta s(x,y)$ such that $\delta s(x,Y) > 0$ and

$$\delta s(x,Y)\cdot f(x,Y; a) + \delta s(x,Y^c)f(x,Y^c; a) = 0; \tag{21}$$

($\delta s(x,Y)$ is constant for all $y \in Y$ and correspondingly for $\delta s(x,Y^c)$). From (19) and (21) it follows that:

$$\delta E^P = \mu\cdot U'(s(x))[\delta s(x,Y)\cdot f_a(x,Y; a) + \delta s(x,Y^c)\cdot f_a(x,Y^c; a)],$$

and

$$\delta E^A = 0.$$

Substituting from (21), we have:

$$\delta E^P = \mu\cdot U'(s(x))\cdot\delta s(x,Y)\cdot f(x,Y; a)\left[\frac{f_a(x,Y; a)}{f(x,Y; a)} - \frac{f_a(x,Y^c; a)}{f(x,Y^c; a)}\right] > 0,$$

since $\mu > 0$ (Proposition 1), $\delta s(x,Y) > 0$ as chosen, and the expression in brackets is positive by (20). The procedure can be repeated for a set of x-values with positive mass, since y is informative, which guarantees that one can make the principal strictly better off and the agent no worse off, for a small enough variation. Finally, utilities are continuous, so part of the principal's gain can be transferred to the agent (e.g., use the same argument as above, taking $\delta E^p = 0$), and this proves the sufficiency part of the proposition. *Q.E.D.*

Remarks:

(1) The sufficiency argument can be appropriately modified to apply to the case where the agent's utility function H is nonseparable.
(2) If, for administrative reasons, one has restricted attention *a priori* to a limited class of contracts (e.g., linear price functions or instruction-like step-functions), then informativeness may not be sufficient for improvements within this class.
(3) From the proof of the proposition one can see that if f_a/f is continuous in (x,y), then there will exist a single region Y (independent of x) such that the indicator function on Y is a valuable signal whenever y is.[22] This implies that $s(x)$ can be improved upon by a dichotomous contract of the form $(s(x,Y),s(x,Y^c))$, which does not use all the information contained in y. Since dichotomous contracts are simpler to administer, this result suggests an explanation of their frequent use.
(4) It is clear that informativeness can be directly extended to cover cases where one already observes a signal y_1 in addition to x and is interested in the value of an additional signal y_2. The necessary and sufficient condition becomes $f(x,y_1,y_2,a) \neq h(x,y_1,a)\cdot g(x,y_1,y_2)$.

The conclusion that a noninformative signal will have no value may not be surprising (even if our terminology is chosen to make this statement appear more

[22] The indicator function on Y is a function which has the signal y as an argument, and equals 1 on Y and 0 otherwise.

obvious than it is). Basically, it tells us that pure randomization does not pay. The more important part of the proposition is the result that any informative signal, regardless of how noisy it is, will have positive value (if costlessly obtained and administered into the contract). As in Harris and Raviv (1978), one might conjecture that in some situations a sufficiently noisy, yet informative, signal would add too much randomness to the contract to be acceptable by risk-averse parties. But as the proof of Proposition 3 indicates, since both parties are on the margin risk-neutral towards randomness in y, given x, the new contract can be designed so that marginally it does not increase risk, but still improves incentives for action.[23] Alternatively, equation (13) indicates that one can improve risk sharing for each y separately while at the same time retaining incentives for action. This pointwise improvement results, of course, in an overall improvement.

It is of interest to look at a few special cases of informativeness. Suppose first that y is independent of x. This could be the case if the agent is directly monitored or supervised. Then we can write

$$f(x,y,a) = h(x,a)\cdot g(y,a).$$

From this it follows that

$$\frac{f_a(x,y,a)}{f(x,y,a)} = \frac{h_a(x,a)}{h(x,a)} + \frac{g_a(y,a)}{g(y,a)} .$$

Hence y is noninformative if and only if g_a/g is constant, which readily is seen to imply $g_a \equiv 0$ (since $\int g_a = 0$). Thus, whenever g depends at all on $\underset{\sim}{a}$, it is informative and consequently valuable. Even the most casual supervision of an agent can be used to the benefit of both parties.

Second, suppose y is informative. Then we can construct another information system as follows:

$$\hat{y} = \begin{cases} y, & \text{if} \quad x \leq \bar{x}, \\ 0, & \text{if} \quad x > \bar{x}. \end{cases}$$

This signal is a conditional information system, where resources are invested to find out y only if the outcome is sufficiently bad (below $\bar{x}$). It is readily seen that $\hat{y}$ is also informative and, depending on the costs of obtaining y, the net benefits of using $\hat{y}$ may exceed those of y.[24] Conditional information systems are widely used in practice, which indicates that their cost savings are often sufficient to cover the information loss they engender.

Finally, one can construct an informative signal $\hat{y}$ from y by simply deciding randomly whether or not to find out y.[25] Again, this would save costs and is quite effective, particularly if y is a very precise signal about $\underset{\sim}{a}$.[26]

The last two examples bring attention to the fact that Proposition 3 says nothing about *how* valuable y is, which would be important whenever costs for information acquisition and administration of more complex contracts are considered. An upper bound for the value is, of course, provided by the value one

---

[23] This line of argument was first used in Gjesdal (1976) for the case where x and y are independent. It has also been used by Shavell (1978), who independently of us proves the sufficiency part of Proposition 3, but without employing the same notion of informativeness.

[24] Demski and Feltham (1978) discuss conditional information systems.

[25] Feltham (1977) gives an example of this kind of information system.

[26] In the limit, if $y = a$ and high penalties are allowed, we are very much in the same situation as in the example in footnote 7. An arbitrary low probability of checking y will suffice to induce the agent to take the correct action.

gets from observing $\underset{\sim}{a}$ itself. As Mirrlees' (1974) example (p. 248) indicates, this value may occasionally be negligible.

Some indications of the value of the signal can be found by studying (13). Roughly speaking, the more variation a signal causes in f_a/f, the more valuable it will be. This seems difficult to formalize, and I believe that on a general level signals can only be compared by using Blackwell's notion of fineness (see Blackwell (1951) and also remark 4 above).

6. Asymmetric information

■ In many respects the model we have analyzed is very primitive. One unrealistic feature is the assumption that the agent chooses his action having the same information as the principal, that is, before anything about θ is revealed. Commonly this will not be the case. After the sharing rule is fixed, the agent will often learn something new about the difficulty of his task or the environment in which it is to be performed. The following extension of our model applies to such cases.[27]

Let z be a signal about θ which the agent observes prior to choosing $\underset{\sim}{a}$, so that his choice becomes a function $a(z)$. As before, we suppress θ and write $f(x,y,z,a)$ for the joint density function, where y is some additional information observed by both parties. The best sharing rule $s(x,y)$ can be determined by solving the program:

$$\max_{s(x,y),a(z)} \int G(x - s(x,y))f(x,y\,|\,z,a(z))p(z)dxdydz \tag{22}$$

$$\text{subject to} \quad \int U(s(x,y))f(x,y\,|\,z,a(z))p(z)dxdydz - \int V(a(z))p(z)dz \geq \bar{H}, \tag{23}$$

$$a(z) \in \underset{a' \in A}{\text{argmax}} \int U(s(x,y))f(x,y\,|\,z,a')dxdy - V(a'), \forall z. \tag{24}$$

Here $f(x,y\,|\,z,a)$ is the conditional density of x and y, given z and the action a, and $p(z)$ is the marginal density of z. Letting $\mu(z)p(z)$ be the multiplier function for (24) and λ the multiplier for (23), point-wise optimization gives the characterization:

$$\frac{G'(x - s(x,y))}{U'(s(x,y))} = \lambda + \frac{\int \mu(z)\cdot f_a(x,y\,|\,z,a(z))p(z)dz}{\int f(x,y\,|\,z,a(z))p(z)dz}. \tag{25}$$

This equation closely resembles equation (13). Again the second term on the right-hand side indicates deviations from a first-best solution, and qualitatively one can draw conclusions similar to those for the earlier model. The difference is that the deviation from first-best risk sharing is determined by a weighted average of the incentive effects in the various states z, with the weight $\mu(z)p(z)$ being dependent on the probability of z and the desirability (or cost) of forcing

[27] This corresponds to Model 2 in Harris and Raviv (1976).

the action $a(z)$. It is easy to show that $\mu(z) \equiv 0$ is impossible (since $\mu(z)$ is determined by an equation similar to (8)), and hence again we have a second-best solution. However, we may have $\mu(z) > 0$ for some z, and $\mu(z) < 0$ for others, since $s'(x) > 1$ is possible (cf. our repairman example) in some x-region.

The necessary part of Proposition 3, namely that a noninformative signal is valueless, extends readily to the asymmetric case. Here noninformativeness is defined by the condition:

$$f(x,y,z; a) = g(x,y) \cdot h(x,z; a), \qquad \text{for almost every } (x,y,z). \tag{26}$$

For the sufficiency part of the proposition, an additional but insignificant qualification is needed. When (26) is false, that is, when y is informative, f_a/f will depend on y as before. Yet, when integrating as in (25), it is conceivable that the right-hand side of (25) would become independent of y, making a function $s(x)$ optimal and y valueless. However, this is extremely unlikely and will not happen generically; any small change in the problem data would take us out of such a situation. Thus, we can safely say that for all that matters, Proposition 3 is also valid in the asymmetric case.

7. Concluding remarks

■ We have studied efficient contractual agreements in a principal-agent relationship under various assumptions about what can be observed, and hence contracted upon, by both parties. When the payoff alone is observable, optimal contracts will be second-best owing to a problem of moral hazard. By creating additional information systems (as in cost accounting, for instance), or by using other available information about the agent's action or the state of nature, contracts can generally be improved. A simple necessary and sufficient condition for such imperfect information to be of value was given as well as a characterization of optimal contracts which use such information.

Principal-agent relationships are prevalent in economic organizations. The analysis presented here improves our understanding of the functioning of this basic organizational form. In view of our result that essentially any imperfect information about actions or states of nature[28] can be used to improve contracts, we have an explanation of the observed complexity of real contracts (as evidenced for instance in insurance arrangements). Additional information is of value because it allows a more accurate judgment of the performance of the agent; or viewed differently, it provides the same incentives for effort with less loss of risk-sharing benefits.

Our analysis also provides a basis for studying the design of contracts and information systems in more specific contexts. An application of this kind has recently been given by Baron and De Bondt (1978) in the context of automatic fuel adjustment clauses. Other fields of applications have been discussed in Harris and Raviv (1978) and recently Demski (1977) has used the model for a theoretical study of financial reporting.

Of course, the analysis presented here leaves unanswered many interesting questions in contracting. One important aspect of the problem, which we have

[28] Note that our analysis shows that from a theoretical point of view there is no distinction to be made between a signal which provides information about actions and one which provides information about states of nature, since these pieces of information are inherently linked via the outcome function.

not considered, is that many contracts are based on long-term relationships. When the same situation repeats itself over time, the effects of uncertainty tend to be reduced and dysfunctional behavior is more accurately revealed, thus alleviating the problem of moral hazard. Such long-term effects could be analyzed in an extension of our model. Another extension would recognize that asymmetry of information as discussed in Section 6 may warrant a re-negotiation of the contract. One can view management by objectives and the New Soviet Incentive Scheme (Weitzman, 1976) as examples of this. In both cases, after observing the difficulty of his task, the agent can change the contract within certain limits to the benefit of both parties. A preliminary discussion of this kind of contracting is given in Holmström (1977), where it is seen as a special case of delegation of decisionmaking responsibility to an agent with superior information.

Appendix

■ **Proof of proposition 1.** Let $s(x)$ be a second-best sharing rule for $\lambda > 0$ and write $r(x) = x - s(x)$. If $\mu \leqslant 0$, contrary to our claim, then

$$\frac{G'(r(x))}{U'(x - r(x))} = \lambda + \mu \cdot \frac{f_a(x,a)}{f(x,a)} \leqslant \lambda = \frac{G'(r_\lambda(x))}{U'(x - r_\lambda(x))}, \tag{A1}$$

for $x \in X_+ = \{x \mid f_a(x,a) \geqslant 0\}$. Here $r_\lambda(x)$ is the first-best sharing rule (in terms of the principal's share), corresponding to λ; see Wilson (1968). Since $G'(r(x))/U'(x - r(x))$ is decreasing in $r(x)$ for fixed x, $r_\lambda(x)$ is an increasing function, and from (A1) it follows that $r(x) \geqslant r_\lambda(x)$ for $x \in X_+$.

Correspondingly, $r(x) \leqslant r_\lambda(x)$ on $X_- = \{x \mid f_a(x,a) < 0\}$. We have then,

$$\int G(r(x))f_a(x,a)dx \geqslant \int G(r_\lambda(x))f_a(x,a)dx > 0, \tag{A2}$$

where the last inequality follows, by first-order stochastic dominance, from the assumption $F_a(x,a) \leqslant 0$ (with strict inequality for some x), and the fact that $r_\lambda(x)$ is increasing.

The expression in braces in equation (8) is the second-order condition for the agent's maximization problem, and hence is <0. (It cannot be $=0$, since then (A2) and (8) would be inconsistent). Combining (8) and (A2), this implies $\mu > 0$, which contradicts our contrapositive assumption $\mu \leqslant 0$. We have arrived at a contradiction assuming $\mu \leqslant 0$ and conclude that $\mu > 0$. *Q.E.D.*

□ **Proof of corollary 1.** The proof follows from Proposition 1 and the fact that $G'(x - s(x))/U'(s(x))$ is increasing in $s(x)$ for fixed x. *Q.E.D.*

□ **Proof of corollary 2.** The solutions will differ on a set of nonzero measure, since $\mu > 0$ and f_a/f is nonconstant. *Q.E.D.*

References

Arrow, K. *Essays in the Theory of Risk-Bearing*. Amsterdam: North-Holland Publishing Company, 1970.

Baron, D. and De Bondt, R. "The Design of Automatic Price Adjustment Mechanisms." Onderzoeksrapport N° 7810, Departement voor Toegepaste Economische Wetenschappen, Katholieke Universiteit Leuven, Leuven, Belgium, 1978.

BLACKWELL, D. "Comparison of Experiments" in J. Neyman, ed., *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley: University of California Press, 1951.

BORCH, K. "Equilibrium in a Reinsurance Market." *Econometrica*, Vol. 30, No. 3 (1962), pp. 424–444.

DEGROOT, M. *Optimal Statistical Decisions*. New York: McGraw-Hill Book Company, 1970.

DEMSKI, J. "A Simple Case of Indeterminate Financial Reporting." Mimeo, Graduate School of Business, Stanford University, 1977.

——— AND FELTHAM, G. "Economic Incentives in Budgetary Control Systems." *The Accounting Review*, Vol. 53, No. 2 (1978), pp. 336–359.

FELTHAM, G. "Optimal Incentive Contracts: Penalties, Costly Information and Multiple Workers." Working paper, University of British Columbia, 1977.

GJESDAL, F. "Accounting in Agencies." Mimeo, Graduate School of Business, Stanford University, 1976.

HARRIS, M. AND RAVIV, A. "Optimal Incentive Contracts with Imperfect Information." Working Paper #70-75-76, Graduate School of Industrial Administration, Carnegie-Mellon University, April 1976 (revised December 1977).

——— AND ———. "Some Results on Incentive Contracts with Applications to Education and Employment, Health Insurance, and Law Enforcement." *The American Economic Review*, Vol. 68 (1978), pp. 20–30.

HOLMSTRÖM, B. "On Incentives and Control in Organizations." Unpublished Ph.D. dissertation, Graduate School of Business, Stanford University, 1977.

KOLMOGOROV, A.N. AND FOMIN, S.V. *Introductory Real Analysis*. New York: Dover Publications, 1970.

LUENBERGER, D. *Optimization by Vector Space Methods*. New York: John Wiley & Sons, 1969.

MIRRLEES, J. "The Optimal Structure of Incentives and Authority within an Organization." *The Bell Journal of Economics*, Vol. 7, No. 1 (Spring 1976), pp. 105–131.

———. "Notes on Welfare Economics, Information, and Uncertainty" in Balch, McFadden, and Wu, eds., *Essays on Economic Behavior under Uncertainty*, Amsterdam: North Holland Publishing Co., 1974.

MOSSIN, J. "Aspects of Rational Insurance Purchasing." *Journal of Political Economy*, Vol. 76 (1968), pp. 553–568.

PAULY, M. "Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection." *Quarterly Journal of Economics*, Vol. 68 (1974), pp. 44–62.

ROSS, S. "The Economic Theory of Agency: The Principal's Problem." *The American Economic Review*, Vol. 63 (1973), pp. 134–139.

SHAVELL, S. "Risk Sharing and Incentives in the Principal and Agent Relationship." *The Bell Journal of Economics*, Vol. 10, No. 1 (Spring, 1979), pp. 55–73.

SPENCE, M. AND ZECKHAUSER, R. "Insurance, Information, and Individual Action." *The American Economic Review*, Vol. 61 (1971), pp. 380–387.

STIGLITZ, J. "Incentives, Risk, and Information: Notes toward a Theory of Hierarchy." *The Bell Journal of Economics*, Vol. 6, No. 2 (Autumn 1975), pp. 552–579.

WEITZMAN, M. "The New Soviet Incentive Model." *The Bell Journal of Economics*, Vol. 7, No. 1 (Spring 1976), pp. 251–257.

WILLIAMSON, O.E. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: The Free Press, 1975.

WILSON, R. "The Theory of Syndicates." *Econometrica*, Vol. 36 (1968), pp. 119–132.

———. "The Structure of Incentives for Decentralization under Uncertainty." *La Decision*, No. 171 (1969).

ZECKHAUSER, R. "Medical Insurance: A Case Study of the Trade-off between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory*, Vol. 2 (1970), pp. 10–26.