

338.70151  
Ar69e

# THE ECONOMICS OF AGENCY

by

KENNETH J. ARROW

OPEN SHELF

TECHNICAL REPORT NO. 451

October 1984

A REPORT OF THE  
CENTER FOR RESEARCH ON ORGANIZATIONAL EFFICIENCY  
STANFORD UNIVERSITY  
Contract ONR-N00014-79-C-0685, United States Office of Naval Research

THE ECONOMICS SERIES  
INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES  
FOURTH FLOOR, ENCINA HALL  
STANFORD UNIVERSITY  
STANFORD, CALIFORNIA



338.70151  
Ar69e

Commerce

THE ECONOMICS OF AGENCY\*

by

Kenneth J. Arrow\*\*

The agency relationship is a pervasive fact of economic life. Even in the limited sense in which the concept has traditionally been understood in ordinary and in legal discourse, the principal-agent relation is a phenomenon of significant scope and economic magnitude. But economic theory has recently recognized that analogous interactions are virtually universal in the economy, at least as one significant component of almost all transactions.

The common element is the presence of two individuals. One (the agent) must choose an action from a number of alternative possibilities. The action affects the welfare of both the agent and another person, the principal. The principal, at least in the simplest cases, has the additional function of prescribing payoff rules; that is, before the agent chooses the action, the principal determines a rule that specifies the fee to be paid to the agent as a function of the principal's observations of the results of the action. The problem acquires interest only when there is uncertainty at some point, and, in particular, when the information available to the two participants is unequal. The

---

\*To be published in J. W. Pratt and R. Zeckhauser (eds.) Agency: The Structure of Business, Cambridge, Mass.: Harvard Business School Press, 1985.

\*\*This research was supported by the Office of Naval Research Grant ONR-N00014-79-C-0685 at the Center for Research on Organizational Efficiency, Institute for Mathematical Studies in the Social Sciences, Stanford University. I wish to express my gratitude to John W. Pratt, John G. Riley, and Richard Zeckhauser, whose comments have materially improved the exposition of this paper.

economics literature has focused primarily (but not exclusively) on the case in which (1) the agent's action is not directly observable by the principal, and (2) the outcome is affected but not completely determined by the agent's action. (Were it not for the second condition, the principal could infer the agent's action by observing the outcome.) In technical language, the outcome is a random variable whose distribution depends on the action taken.

More generally, a single principal may have many agents. Each takes an action, and the output of the system is a random function of all the actions. The principal cannot observe the actions themselves but may make some observations, for example, of the output. Again the principal sets in advance a schedule stating the fees to be paid to the individual agents as a function of the observations made by the principal.

A similar but not identical principal-agent relation occurs when the agent makes an observation not shared with the principal and bases his/her action on that observation. In this case, the principal may be able to observe the action itself, but does not know whether it is the most appropriate one.

The principal-agent theory is in the standard economic tradition. Both principal and agent are assumed to be making their decisions optimally in view of their constraints; intended transactions are realized. As is usual in economic theory, the theory functions both normatively and descriptively. It offers insights used in the construction of contracts to guide and influence principal-agent relations in the real

world, at the same time it represents an attempt to explain observed phenomena in the empirical economic world, particularly exchange relations that are not explained by more standard economic theory.

1. The Two Types: Hidden Action and Hidden Information

Before specifying the model more completely, it is useful to distinguish a few examples of each of the two kinds of principal-agent problems and give a few examples of each. As will be seen, many situations that are not classified under that heading in ordinary discourse can be considered as such. I will call the two types of principal-agent problems hidden-action and hidden-information, respectively. In the literature, they are frequently referred to as moral hazard and adverse selection. These terms have been borrowed from the practice and theory of insurance and are really applicable only to special cases.

The most typical hidden action is the effort of the agent. Effort is a disutility to the agent, but it has a value to the principal in the sense that it increases the likelihood of a favorable outcome (technically, the distribution of the outcome to a higher effort stochastically dominates that to a lower effort, i.e., the probability of achieving an outcome that exceeds any given level is higher with higher effort). The physician-patient relation is a notorious case. Here, the physician is the agent who chooses actions affecting the welfare of the principal (the patient). The very basis of the relation is the superior knowledge of the physician. Hence, the patient cannot check to see if the actions of the physician are as diligent as they could be.

A second less obvious example is that of torts. One individual takes an action that results in damage to another - for example, an automobile collision. The care taken by the first driver cannot easily be observed, but the outcome is very visible indeed. Although it may seem an odd use of language, one has to consider the damager as the agent and the one damaged as the principal. Again, in pollution control, society may be regarded as the principal, and the polluter, whose actions cannot be fully monitored, as the agent.

An example of very special economic importance is the relation between stockholders and management. The stockholders are principals, who certainly cannot observe in detail whether the management, their agent, is making appropriate decisions. A formally similar relation, though in a different context, is that of sharecropping. Instead of paying straight wages, the landlord (the principal here) prefers a relation that supplies incentives for better production, since the tenant's diligence cannot be directly observed; on the other side, the tenant, too poor to bear excessive risks, wants to avoid a fixed rent, which would maximize incentives but would expose him or her to all the risks of weather and price.

Fire insurance dulls incentives for caution and even creates incentives for arson; this is the origin of the term moral hazard. Health insurance creates similar problems, for it creates an incentive to excessive medical care. If medical fees are paid by the insurer, the patient may elect more costly treatments than he or she would be willing to pay for individually. The employment relation in general is one in

which effort and ability acquired through training and self-improvement are hard to observe. In one view, firms exist as a means of measuring effort.

In the hidden-knowledge problems, the agent has made some observation that the principal has not made. The agent uses (and should use) this observation in making decisions; however, the principal cannot check whether the agent has used his or her information in the way that best serves the principal's interest. A case much studied from various points of view in the economic literature is that of a decentralized socialist economy. Because the knowledge of productivity cannot be centralized, the individual productive units have information about the possibilities of production not available to the central planning unit. The productive units may well have incentives not to reveal their full potentiality, because it will be easier to operate with less taxing requirements. The problem for the central planning unit (the principal) is how to tap the agents' information. A similar problem occurs in decentralization within a firm. This topic in the literature has acquired the name of incentive compatibility.

The problem of adverse selection was originally noted in insurance of several kinds. The population being insured is heterogeneous with respect to risk; in the case of life insurance, for example, some have a higher probability than others of dying young. In at least some cases, the insured have better knowledge of this probability than the insurance company, which is unable to differentiate. If the same premium is charged to everyone, then the high-risk individuals will purchase more

insurance and the low-risk ones less. This will lead to an inefficient allocation of risk-bearing (Rothschild and Stiglitz [1976]). Public utilities, such as telephones, also face heterogeneous populations; as in insurance, the utility provider cannot know to which class a given purchaser belongs. Nevertheless, as has been pointed out in recent literature, some differentiation can be made by offering alternative rate schedules and letting the customers choose which to follow. In these cases, the insurance company or the public utility is the principal, the customer, with more knowledge not available to the principal, is the agent (Spence [1977], Roberts [1979], Maskin and Riley [1983]).

## 2. Example: Public Utility Rate Setting

To illustrate the theoretical issues for the hidden-knowledge model, consider a monopolistic public utility facing two types of customers, designated H and L for high and low demanders respectively. Assume there are no income effects. Let  $U_t(x)$  be the money equivalent of amount  $x$  of the public utility for type  $t$  ( $t = H, L$ ), so that  $U_t(0) = 0$ , and characterize high and low demand by the condition that  $U'_H(x) > U'_L(x)$  for all  $x$ . It is assumed that the characteristics of the product preclude resale.

The public utility knows the proportion of high demanders but not the identity of these individuals. It offers a total payment schedule,  $T(x)$ , a function of the amount purchased. Assuming a constant marginal cost of production,  $c$ , the monopolist's markup,  $M$ , for  $x$  units is,

$$M(x) = T(x) - cx .$$

For convenience, let  $V_t(x) = U_t(x) - cx$ , the consumer's surplus over social cost. Since  $V'_H(x) > V'_L(x)$ , for all  $x$ , there is a difference in willingness to pay which the monopolist can exploit.

Since individuals are free to refrain from purchase, there will never be a negative consumer's surplus, no matter what pricing scheme the utility adopts. The monopolist can try to extract all consumer's surplus by all-or-none offers. Let  $\bar{x}_t$  maximize  $V_t(x)$ . If the monopolist can identify the type to which each consumer belongs, it will offer buyers of type  $t$   $\bar{x}_t$  units and charge a markup of  $\bar{M}_t = V_t(\bar{x}_t)$ .

If identification is not possible, however, this scheme breaks down. If the monopolist offers the consumer a choice between  $(\bar{x}_L, \bar{M}_L)$  and  $(\bar{x}_H, \bar{M}_H)$ , the high demanders will always choose the former. Since  $V'_H(x) > V'_L(x)$ , it follows that  $V_H(\bar{x}_L) > V_L(\bar{x}_L) = \bar{M}_L$ , so that type H individuals get a positive consumer's surplus by choosing the offer appropriate to type L individuals; if they accept  $(\bar{x}_H, \bar{M}_H)$ , their consumer's surplus is zero. To induce type H individuals to buy  $\bar{x}_H$ , the markup must be reduced so that they are no worse off than they would be choosing  $(\bar{x}_L, \bar{M}_L)$ . That is, the markup demanded must satisfy the condition,

$$(2.1) \quad V_H(\bar{x}_H) - M_H^0 = V_H(\bar{x}_L) - \bar{M}_L .$$

This can be accomplished without knowing individual consumer types by choosing  $M(x) = \bar{M}_L$  for  $x \leq \bar{x}_L$ ,  $= M_H^0$  for  $x > \bar{x}_L$ .

This allocation is Pareto efficient, since all consumers are paying marginal cost. The monopolist is extracting all surplus from the low demanders but not from the high demanders. To realize maximum profits, however, the monopolist must set prices in a manner that creates inefficiency. The amount to be bought by the low demanders will be reduced by a small amount. This will reduce the surplus to be extracted from them. On the other hand, the constraint imposed on extraction of surplus from the high demanders to prevent them from switching to the offer intended for the low demanders will become easier to satisfy. It turns out that the loss in profits due to the reduction in amount purchased by the low demanders is much smaller than the gain from higher markup obtainable from the higher demanders. In symbols, let the amount to be purchased by type L consumers be reduced from  $\bar{x}_L$  to  $\bar{x}_L - dx$ . This is enforced by locating the discontinuous increase in markup at that point. The markup must be reduced correspondingly; choose  $M_L^* = V_L(\bar{x}_L - dx)$ . Since  $V_L$  is maximized at  $\bar{x}_L$ , it must be that the difference,  $M_L^* - \bar{M}_L$  is of the second order in  $dx$ .

To induce the type H consumers to choose  $\bar{x}_H$  rather than  $(\bar{x}_L - dx, M_H^*)$ , the markup to them must be set so that

$$V_H(\bar{x}_H) - M_H^* = V_H(\bar{x}_L - dx) - M_L^* .$$

By comparison with (2.1), it is seen that,

$$M_H^* - M_H^0 = (M_L^* - \bar{M}_L) + [V_H(\bar{x}_L) - V_H(\bar{x}_L - dx)] .$$

The first term on the right is, as stated, of the second order in  $dx$ . But since  $V_H'(\bar{x}_L) > 0$ , the second term is positive and of the first order. Hence, for  $dx$  sufficiently small, the loss in markup from the type L consumers is of the second order, the gain in markup from the type H consumers is of the first order, and there is a net gain. This is true no matter what the proportions of the two types of consumers are, though of course the optimal policy of the monopolist depends on those proportions. The optimal monopoly policy can be enforced without identification of the types of consumers by letting  $M(x) = M_L^*$  for  $x < \bar{x}_L - dx$ ,  $= M_H^*$  for  $x > \bar{x}_L - dx$ .

Constraints such as (2.1) that ensure that the different types are induced to accept the allocations allotted to them are referred to as self-selection constraints. The example illustrates a very general principle in hidden-knowledge models; the optimal incentive schedule typically requires distortions (deviations from first-best Pareto-optimal) at all but one point.

Another instance of hidden knowledge in economic decisionmaking is the auction with private information (Vickrey [1961], Maskin and Riley [1984], Milgrom and Weber [1982]). Bidders for oil leases, for example may be permitted to engage in exploratory drilling and other geophysical studies. Each then has an observation unknown to the others and to the seller, which in the United States today is most usually the government. The problem is to design auction rules to achieve some objective. Much of the current literature is devoted to maximizing the seller's revenues, rather than social welfare in some broader sense.

A final illustration of hidden knowledge is the problem of optimal income taxation (Mirrlees [1971]). Any income tax distorts the choice between labor and leisure. This deficiency could in principle be overcome completely if the social price of leisure (i.e., the productivity or wage rate of the individual) were observable. In general, however, this information is available to the taxpayer but not to the government. Like the geophysical estimates of oil field size in the earlier example, individual wage rates are the private information of the agents and therefore hidden knowledge to the principal.

### 3. Multiple Principals

The hidden-knowledge principal-agent problem becomes more complicated when multiple principals compete for agents (Spence [1983], Rothschild and Stiglitz [1975], Riley [1975]). Suppose a large number of potential principals will enter the market to exploit any profitable alternative. This might be the case, for example, in an insurance market with a large number of competing insurance companies, each of which, because of risk pooling, is approximately risk neutral. As argued earlier, any given offer (so much coverage at such a premium) will be more attractive to those with higher loss probabilities; insurance companies will then have an incentive to sort risk classes by offering lower premiums per dollar of coverage to those willing to accept higher deductibles. However, unlike the monopoly utility, each insurance company must take into account the effect of other available alternatives on the type of individuals attracted to its own offerings. To use

Spence's terminology, it is not enough that low risk classes can "signal" their differences by accepting larger deductibles; such signals must also be competitively viable.

The issue of what kind of signaling survives competitive pressures turns out to be a delicate one. In general, there does not exist a Walrasian (or Nash) equilibrium with the property that no principal has an incentive to introduce new profitable alternatives. However, recent work by Wilson [1977] and Riley [1979] has argued that equilibrium can be sustained if principals rationally anticipate certain responses to their behavior.

#### 4. The Hidden-Action Model

Let me now turn to a simple formulation of the hidden-action model. The agent (for the moment, assume there is only one) chooses an action  $a$ . The result of this choice is an outcome,  $x$ , which is a random variable whose distribution depends on  $a$ . The principal has chosen beforehand a fee function,  $s(x)$ , to be paid to the agent. For the simplest case, assume that the outcome  $x$  is income, i.e., a transferable and measurable quantity. Then the net receipts of the principal will be  $x - s(x)$ . The principal and agent are both, in general, risk averters. Hence, each values whatever income he or she receives by a utility function with diminishing marginal utility. Let  $U$  be the utility function of the principal,  $V$  that of the agent. Further, let  $W(a)$  be the disutility the agent attaches to action  $a$ . It will be assumed separable from the utility of income, i.e., the marginal utility

of income is independent of the action taken (the amount of effort). Note that the action is taken before the realization of the uncertainty and is therefore not uncertain to the agent, though it is unknown to the principal.

Since, even for a given action, the outcome,  $x$ , is uncertain, both principal and agent are motivated to maximize the expected value of utility. Given the principal's choice of fee function,  $s(x)$ , the agent wishes to maximize the expected value of  $V[s(x)] - W(a)$ . In effect, therefore, the principal can predict the action taken for any given fee schedule. The choice of fee schedules is, however, restricted by competition for agents, who have alternative uses for their time. Hence, the principal must choose a fee schedule that offers the agent a utility at least equal to what he or she could achieve in other activities. The literature has usually referred to this condition as that of individual rationality, a term first used by J. von Neumann and O. Morgenstern, but this name is easily misinterpreted. The term, participation constraint, has come into use recently and seems more appropriate.

The principal-agent relation defined, as here, by a fee function is a significant departure from the arm's length fixed-price relation among economic agents usually postulated in economic theory. The principal does not buy the agent's services at a fixed price set by the competitive market nor does the principal simply buy output from the agent. The relation cannot even be described as a contingent contract, in which payments and services rendered are agreed-on functions of an exogenous random variable; rather the principal observes the outcome but

cannot analyze it into its two components, the agent's action and the exogenous uncertainty. Even though the underlying principles are impeccably neoclassical, in that both parties are acting in their own self-interest and are subject to the influence of the market, the variable to be determined is not a price but a complicated functional relationship.

The principal-agent problem combines two inextricable elements, risk-sharing and differential information. Even if there were no problem of differential information, there would be some sharing of the outcome if both parties are risk-averse. Indeed, if the agent were risk-neutral, the principal-agent problem would have a trivial solution: the agent would bear all the risks, and then the differential information would not matter. That is, the principal would retain a fixed amount and pay all the remainder to the agent, who therefore would have no dilution of incentives (Shavell [1979]). In the terminology used above, the fee function would equal the outcome less a fixed amount,  $s(x) = x - c$ , where the constant  $c$  is determined by the participation constraint. Thus a landlord renting land to a tenant farmer would simply charge a fixed rent independent of output, which in general depends on both the tenant's effort, unobservable to the landlord, and the vagaries of the weather. However, this solution is not optimal if the agent is risk-averse. Since all individuals are averse to sufficiently large risks, the simple solution of preserving incentives by assigning all risks to the agent fails as soon as the risks are large compared with the agent's wealth. The president of a large corporation can hardly be held responsible for its income fluctuations.

In the general case of a risk-averse agent, the fee will be a function of the outcome, in order to supply incentives, but the risk will be shared. If the ability of the agent to affect outcomes approaches either zero or infinity, then the efficiency level that could be achieved under full information to the principal can be approached with an optimally chosen fee function. More generally, there is a trade-off between incentives and efficiency of the system (considering both principal and agent) (Shavell [1979]).

For an application, consider the case of insurance with moral hazard. Some insurance will be written, but it will not be complete. In the terminology of the insurance industry, there will be coinsurance, that is, the insured will bear some of the losses against which the insurance is written. Coinsurance is customary in health insurance policies, where the insured has considerable control over the amount of health expenditures. Similarly, in a system of legal liability for torts (assuming no insurance), the required payment should increase with the amount of damages inflicted, to provide incentive for avoiding the inflicting of damages, but by an amount less than the increase in damages, so that the unavoidable risks are shared.

##### 5. Monitoring

More recent literature has stressed the possibility of monitoring. By this is meant that the principal has certain information in addition to the outcome. If this observation,  $y$ , conveys any information about the unobserved action,  $a$ , beyond that revealed by the outcome,  $x$

(technically, if  $x$  is not a sufficient statistic for the pair  $x, y$  with respect to  $a$ ), then one can always improve by making the fee depend upon  $y$  as well as  $x$ . In the case of torts, it is held in many cases that a damager is not liable if due care has been exercised. Therefore the plaintiff is required to show negligence on the part of the defendant, so that additional knowledge beyond the outcome is available. It turns out that if the liable party (the agent in this interpretation) is risk neutral, then a strict liability standard, which requires only knowledge of the outcome, is optimal (in the sense of economic efficiency). But otherwise an appropriate negligence standard is an improvement (Shavell [1979], Holström [1979]). Harris and Raviv [1978, 1979] have argued that the custom of paying lawyers (in most circumstances) on the basis of time spent, as well as by a contingent fee, is an example of monitoring. If this idea were applied to health insurance, it would suggest that an improvement could be achieved by making insurance payments to the provider of care depend on some measure of the amount of medical services provided, such as frequency of visits.

It has been shown that if the monitoring information is essentially an imperfect measure of the action taken, i.e.,  $y = a + u$ , where  $u$  is a random variable with mean zero, then an optimal fee policy is to pay a very low figure, independent of outcome, if the measured action is sufficiently low, and to pay according to a more complicated schedule otherwise.

6. Multiple Agents and Repeated Relations

New possibilities for incentives arise when there are many agents for a single principal or repeated relations between agent and principal. The many-agent case offers new opportunities for inference of hidden actions (or of hidden information) if the uncertainty of the relation between the action (or the agent's observation) is the same for all the agents. In that situation, the uncertainty can be estimated by comparing the performances of the different agents; thus individual actions can be approximately identified. One can meaningfully compare the performance of each agent with the average, for example, or use the ordinal ranking of the agents' outcomes as a basis for fees (Holmström [1982]).

A different and as yet only slightly explored problem can arise in the case of many agents with a single principal. Suppose the principal cannot observe the outcome of each individual agent's action but only the output of the group of agents as a whole. This is obviously an important case in production carried out jointly, with many complementary workers. Even when the relation between actions and collective outcomes is certain, there are difficulties. Holmström [1982] has considered the problem of a team, whose output depends on the unobservable actions of all members. Each team member has a disutility for his or her action. Assume for simplicity that utility is linear in the output. Then one can speak of a social optimum, that vector of actions which maximizes total output minus the sum of disutilities for actions. The question is whether the team can devise some incentive scheme that

will induce the members to perform the socially optimal actions. This will necessarily be a game, since the reward to each is a function of the output and therefore of the actions of all. When there is no uncertainty, an incentive scheme can be devised with the desired outcome in mind. Let  $a_i$  be the action to be chosen by individual  $i$ ,  $x(a_1, \dots, a_n)$  the production function that gives the output of the team as a function of the actions of all members, and  $W_i(a_i)$  the disutility of individual  $i$  as a function of his or her action. Then the socially optimal set of actions is that which maximizes  $x(a_1, \dots, a_n) - \sum W_i(a_i)$ . Call the actions so defined,  $a_1^*, \dots, a_n^*$ , and let  $x^* = x(a_1^*, \dots, a_n^*)$  be the output at this optimum. Choose any set of lump-sum rewards,  $b_1, \dots, b_n$ , which add up to  $x^*$ , subject to the condition that  $b_i > W_i(a_i^*)$  for each  $i$ . Then set up the following game: Each individual  $i$  chooses  $a_i$ . If the result of all these actions is to produce an output that is less than optimal, no one receives anything. If the total output,  $x(a_1, \dots, a_n)$  is greater than or equal to  $x^*$ , then individual  $i$  receives  $b_i$ . It is easy to see that a Nash equilibrium of the game is for each individual to choose the appropriate action,  $a_i^*$ ; that is, for each individual  $i$ , choosing  $a_i^*$  is optimal given the payoffs, providing each other individual,  $j$ , chooses  $a_j^*$ . But the proposed game is hardly satisfactory. It involves in effect collective punishment. More analytically, there are many Nash equilibria, of which  $(a_1^*, \dots, a_n^*)$  is only one. If some individuals shirk a little, it pays the others to work somewhat harder to achieve the same output. Hence, the scheme does not enforce the optimal outcome, though it permits it.

Repeated relations between a principal and an agent provide new opportunities for incentives. Experience rating in insurance illustrates the situation; the premium rate charged today depends on past outcomes. In effect, the information on which the fee function is based is greatly enriched. Radner [1981] has demonstrated the possibilities for achieving almost fully optimal outcomes in hidden-action situations. Suppose the principal wishes the agent to implement a certain level of action,  $a^*$ . In any one trial, the action is hidden, in that the outcome differs from the action by a random variable, i.e.,  $x_t = a_t + u_t$ , where the random variables,  $u_t$ , are identically and independently distributed, with mean zero. If the agent is in fact performing the desired action  $a^*$ , then the distribution of the  $x_t$ 's is known. Hence if enough are observed, the principal should be able to detect statistically whether the agent is performing actions below the desired level. Specifically, the principal can keep track of the cumulative sum of the outcomes. If it ever falls below a known function of time, then the principal can assume that the performance of the agent is below that desired. More exactly the principal imposes a very severe penalty if there is some time,  $T$ , such that,

$$\sum_{t=1}^T x_t < Ta^* - k \log \log T .$$

For properly chosen  $k$ , the probability of imposing a penalty when the agent is in fact carrying out the desired action can be made very low, while the probability of eventually imposing the penalty if the agent is shirking is one.

7. An Evaluation of Agency Theory

I have sketched some of the leading ideas in the rapidly burgeoning literature on the economic theory of the principal-agent relation. We may step back a bit from the pure theory and ask in a general way to what extent our understanding of economic processes has been enhanced. On the positive side, there is little question that a good many economic relations inexplicable in previously standard analysis can now be understood. Contractual relations are frequently a good deal more complicated than the simple models of exchange of commodities and services at fixed prices would suggest. Sharecropping, incentive compensation to executives and other employees, the role of dismissal as an incentive, coinsurance, and other aspects of insurance all find a place in this literature not found in standard economic analysis.

But it is perhaps more useful to consider the extent to which the principal-agent relation in actuality differs from that in the models developed to date. Most importantly, the theory tends to lead to very complex fee functions. It turns out to be difficult to establish even what would appear to be common-sense properties of monotonicity and the like. We do not find such complex relations in reality. Principal-agent theory gives a good reason for the existence of sharecrop contracts, but it is a very poor guide to their actual terms. Indeed, as John Stuart Mill pointed out long ago, the terms tend to be regulated by custom. They are remarkably uniform from farm to farm and from region to region. Principal-agent theory would suggest that the way the produce is divided between landlord and tenant would depend on the

probability distribution of weather and other exogenous uncertainties and on the relation between effort and output, both of which certainly vary from one region to another; the latter has varied over time as well. Similarly, the coinsurance provisions in health insurance policies are much simpler than could possibly be accounted for by principal-agent theory.

In some cases where principal-agent theory seems clearly applicable, real-world practice is very different from the model. In many respects, the physician-patient relation exemplifies the principal-agent relation almost perfectly. The principal (the patient) is certainly unable to monitor the efforts of the agent (the physician). The relation between effort and outcome is random, but presumably there is some connection. Yet in practice the physician's fee schedule is in no way related to outcome. Liability for malpractice can be seen as a modification of the fee schedule in the direction indicated by principal-agent theory; but it is not applicable to what might be termed run-of-the-mill shirking, and it requires very special kinds of evidence. In general, indeed, compensation of professionals shows only a few traces of the complex fee schedules implied by theory.

Even in situations where compensation systems seem closer in form to the theoretical, there are significant differences. Consider the incentive compensation schemes for corporate executives. They invariably have a large discretionary component. What is the purpose of this? Why should the incentive payment not be based entirely on observable magnitudes, profits, rates of return, and the like?

These difficulties can be explained within the terms of the principal-agent logic but in a way that points beyond the usual bounds of economic analysis. One basic problem is the cost of specifying complex relations. There is a large, though not easily defined, cost to a contract that specifies payments that depend on many variables. There is a cost to the very statement of the contract, a cost to understanding it and its implications, and a cost to verifying which terms apply in a given situation. Hence, there is a pressure for simple contracts, the more so since any of our models is actually much too simple to capture all aspects of a relation that would be thought relevant by those in it.

There are a variety of means of monitoring, and it is difficult to define exactly what they are. The world is full of performance evaluations based on some kind of direct observations. These evaluations may not always be objective, reproducible observations of the kind used in our theories (perhaps the only kind about which it is possible to construct a theory). Executives are judged by their superiors and students by professors on criteria that could not have been stated in advance. Outcomes and even supplementary objective measures simply do not exhaust the information available on which to base rewards.

A third limitation of the present models is the restricted reward or penalty system used. It is always stated in terms of monetary payments. Actually, the present literature has already begun to go beyond this limit by considering the possibility of dismissal. Still further extensions are needed to capture some aspects of reality. Clearly, there is a whole world of rewards and penalties in social rather than

monetary form. Professional responsibility is clearly enforced in good measure by systems of ethics, internalized during the education process and enforced in some measure by formal punishments and more broadly by reputations. Ultimately, of course, these social systems have economic consequences, but they are not the immediate ones of current principal-agent models.

All three of these limiting elements - cost of communication, variety and vagueness of monitoring, and socially mediated rewards - go beyond the usual boundaries of economic analysis. It may ultimately be one of the greatest accomplishments of the principal-agent literature to provide some structure for the much-sought goal of integrating these elements with the impressive structure of economic analysis.

References

- Gjesdal, F. [1982], "Information and Incentives: The Agency Information Problem." Review of Economic Studies 49: 373-390.
- Grossman, S., and O. D. Hart [1983], "An Analysis of the Principal-Agent Problem." Econometrica 51: 7-46.
- Harris, M. and A. Raviv [1978], "Some Results of Incentive Contracts with Applications to Education and Employment, Health Insurance and Law Enforcement." American Economic Review 68: 20-30.
- Harris, M. and A. Raviv [1979], "Optimal Incentive Contracts with Imperfect Information." Journal of Economic Theory 20: 231-259.
- Holmström, B. [1979], "Moral Hazard and Observability." Bell Journal of Economics 74-91.
- Holmström, B. [1982], "Moral Hazard in Teams." Bell Journal of Economics 13: 324-340.
- Marcus, A. J. [1982], "Risk Sharing and the Theory of the Firm." Bell Journal of Economics 13: 369-378.
- Maskin, E. S. and J. G. Riley [1983], "Monopoly with Incomplete Information." University of California in Los Angeles Discussion Paper.
- Maskin, E. S. and J. G. Riley [1984], "Optimal Auctions with Risk Averse Buyers." Econometrica, forthcoming.
- Milgrom, P. R. and R. J. Weber [1982], "A Theory of Auctions and Competitive Bidding." Econometrica 50: 1089-1122.
- Mirrlees, J. [1971], "An Exploration in the Theory of Optimum Income Taxation." Review of Economic Studies 38: 175-208.
- Radner, R. [1981], "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship." Econometrica 49: 1127-1148.
- Riley, J. G. [1975], "Competitive Signalling." Journal of Economic Theory 10: 174-186.
- Riley, J. G. [1979], "Informational Equilibrium." Econometrica 47: 331-360.
- Roberts, K. [1979], "Welfare Considerations of Nonlinear Pricing." Economic Journal 89: 66-83.

- Rothschild, M. and J. E. Stiglitz [1976], "Equilibrium in Competitive Insurance Markets." Quarterly Journal of Economics 90: 630-649.
- Shavell, S. [1979], "Risk Sharing and Incentives in the Principal and Agent Relationship." Bell Journal of Economics 10: 55-73.
- Spence, A. M. [1973], Market Signalling: Information Transfer in Hiring and Related Processes. Cambridge, Mass.: Harvard University Press.
- Spence, A. M. [1977], "Nonlinear Prices and Economic Welfare." Journal of Public Economics 8:1-18.
- Vickrey, W. [1961], "Counterspeculation, Auctions, and Competitive Sealed Tenders." Journal of Finance 16: 8-37.
- Wilson, C. A., [1977], "A Model of Insurance Markets with Incomplete Information." Journal of Economic Theory 16: 167-207.

207. "Optimal Depletion of an Uncertain Stock," by Richard Gilbert.  
208. "Some Minimum Chi-Square Estimators and Comparisons of Normal and Logistical Models in Qualitative Response Analysis," by Kunio Morimune.

160. "The Structure and Stability of Competitive Dynamical Systems," by David Cass and Karl Shell.

REPORTS IN THIS SERIES