A SURVEY OF AGENCY MODELS OF ORGANIZATIONS

by

DANIEL LEVINTHAL

TECHNICAL REPORT NO. 443
March 1984

A REPORT OF THE
CENTER FOR RESEARCH ON ORGANIZATIONAL EFFICIENCY
STANFORD UNIVERSITY

THE ECONOMICS SERIES

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES

FOURTH FLOOR, ENCINA HALL

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

A SURVEY OF AGENCY MODELS OF ORGANIZATIONS*

by

Daniel Levinthal**

Technical Report No. 443

March 1984

A REPORT OF THE
CENTER FOR RESEARCH ON ORGANIZATIONAL EFFICIENCY
STANFORD UNIVERSITY

THE ECONOMICS SERIES

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES
Fourth Floor, Encina Hall
Stanford University
Stanford, California
94305

A SURVEY OF AGENCY MODELS OF ORGANIZATIONS*

by

Daniel Levinthal**

1. Introduction

Spence [1975] indicates two approaches to the study of organizations within economics. One focuses on the boundaries of firms and markets (Arrow [1975]; Williamson [1975]) and the other examines the structure and decision-making processes of organizations directly. Agency models of organizations, and correspondingly this essay, are primarily directed at the latter of these two approaches. These models incorporate two basic features of organizations: goal conflict among members of the organization and incomplete information.

Incomplete information is classified as taking one of two forms. The first, referred to as adverse selection, reflects uncertainty regarding a characteristic of an individual, such as productivity (Guasch and Weiss [1980]) or propensity to leave the organization (Salop and Salop [1976]). The second category of uncertainty, referred to as moral hazard, reflects the unobservability of the decisions made by the members of the organization. Without these elements of incomplete information, goal conflict could be overcome by explicit contracting.

The inclusion of goal conflict distinguishes these models from the theory of teams (Marschak and Radner [1972]). Marschak and Radner

[1972] exclude goal conflict on the grounds of mathematical tractability and on the substantive grounds that sufficient side payments are made to the members of the organization such that cooperation is elicited in carrying out the task of the organization. The inducements/contributions theory of the employment relation (Barnard [1938]; Simon [1951]; March and Simon [1958]) is consistent with this view; in this framework, the contract between the employee and the organization does not specify the particular tasks that the employee is to perform but places bounds on what the employee may be compelled to do. Within the agreed bounds, the employee acts with respect to the authority in a manner consistent with a member of a team in the Marschak and Radner [1972] framework, although functioning quite differently in terms of his decision-making process (Simon [1955] and [1959]). Wilson's [1968] work on the theory of syndicates specifies the conditions under which a collection of individuals may be appropriately viewed as forming a team. He shows that similarity of goals is not sufficient but that the members' attitude towards risk must be compatible if the collectivity is to agree on the ranking of outcomes.

The inclusion of incomplete information distinguishes these models of organizations from earlier economic models of the firm that rely on the existence of authority relationships to justify the avoidance of incentive issues. Jensen and Meckling [1976] describe the standard neoclassical theory of the firm as "not a theory of the firm but actually a theory of markets in which firms are important actors. The firm is a 'black box' operated so as to meet the relevant marginal conditions with respect to inputs and outputs, thereby maximizing profits, or more

accurately, present value."[1] To the extent to which the theory of the firm succeeds in describing the allocation of resources by a price system, its gross assumptions will be justified (Friedman [1953]). However as Cyert and March [1963] point out, there are a number of important and interesting questions relating specifically to firm behavior that the theory cannot answer and was never developed to answer, especially with regard to the internal allocation of resources and the process of setting prices and outputs. They go on to say that "ultimately a new theory of firm decision-making behavior might be used as a basis for a theory of markets, but at least in the short run we should distinguish between a theory of microbehavior, on the one hand, and the microassumptions appropriate to a theory of aggregate economic behavior on the other."[2]

Agency models provide a framework for the analysis of intra-organizational issues that cannot be addressed by the neoclassical theory of the firm. The agency relationship is quite general. Ross [1973] states that: "The relationship of agency is one of the oldest and commonest codified modes of social interaction. We will say that an agency relationship has arisen between two (or more) parties when one, designated as the agent, acts for, on behalf of, or as representative for the other, designated the principal, in a particular domain of decision problems. Examples of agency are universal."[3] Along the same lines, Jensen and Meckling [1976] note: "The problem of inducing an 'agent' to behave as if he were maximizing the 'principal's' welfare is quite general. It exists in all organizations and in all cooperative efforts--at every level of management in firms."[4]

Despite (or perhaps as a result of) its generality, agency models fail to capture many important aspects of behavior within organizations. Arrow [1971] points out that, in situations in which moral hazard arises, there is potentially general advantage in moral behavior, i.e. behavior that is not motivated by narrow self-interest, and that such behavior occurs. The models discussed here assume that contracts, explicit and implicit, are exploited by the parties in their own interest so that promises and claims about unobservable behavior are not admissible--"A theory that overemphasizes self-interested behavior in this way deserves to fail in predicting various features of actual organizations; but it would be surprising if it were wholly irrelevant."5/

Models of the employment relation are grounded on particular assumptions of the observability of actions, uncertainty of employee and organizational characteristics, the technology of the task environment and the goals of both the individual and the organization. Visions of the employment relation as a simple authority relationship are sensible in a world in which actions are observable and the nature of the task is known. Much of the apparatus of personnel management that we observe can be explained by uncertainty about worker characteristics in combination with a complex technology.

Thompson [1967] notes that "instrumental or purposive action is rooted on one hand in desired outcomes and on the other hand in beliefs about cause/effect relationships, and we would also expect these to be the basic variables for the assessment of organizational action and

readiness for action."[6]/ Goals may range from crystallized to ambiguous and the understanding of cause/effect relations from complete to incomplete. Combining these two dimensions and considering only the extreme values, Thompson [1967] considers what assessment technique might be expected in each of the four cases. Where cause/effect understanding is believed complete and a standard of desirability is crystallized, it is possible to assess whether a given effect was produced with least cost, or, whether a given amount of resources were used in a way to achieve the greatest result.[7]/ Agency models of organizations are developed in a setting in which standards of desirability are crystallized and there is incomplete knowledge of cause/effect relationships. When standards of desirability are ambiguous, the assessor may refer to social reference groups (Merton [1957]) for evaluation purposes. Rational choice models of behavior have made only tentative steps in incorporating ambiguity of preferences (March [1978]), and agency models, in particular, have little to say with respect to organizational decision-making in the context of ambiguous preferences.

The existence of incomplete information and goal conflict provides the impetus for the study of many of the basic issues in organizational theory, such as motivation of workers (Maslow [1943], Vroom [1964], Locke [1968]); organizational design--in particular, decentralization and coordination (Galbraith [1973]) and centralization (Blau and Schoenher [1971]); assessing effectiveness--the setting of standards and the selection of indicators for assessing organizational effectiveness and decisions regarding the selection of the performance sample (Scott,

Dornbush, Busching and Laing [1967]); the strategic manipulation of evaluation mechanisms (March [1980]; Pfeffer [1982]); and, organizational conflict (March and Simon [1958]).

While motivated by the same concerns, agency models, at present, focus on the more micro-organizational issues, such as performance evaluation systems. As Eisenhardt [1983] points out, agency problems "concern organizational design; however, their solutions are not the usual structural ones such as matrix, decentralization, or divisionalization. Rather the solutions relate to control and the concrete design tools of control such as reward structures and information systems...the problem is how to obtain coordinated action, which is problematic to observe, from a diverse set of people with differing objectives."[8] The discussion of these issues in the behavioral literature has proceeded largely independently of the recent contributions of economists to these questions. The intent of this essay is to explicate the structure of the agency model of organizations and to discuss the results of this literature in the context of the above issues in organizational theory.

Section 2 develops the basic mathematical formulation of the agency relationship and discusses the application of the revelation principal to agency problems. The implications of the optimal compensation function for risk-sharing and the possibility of achieving the first-best outcome are discussed in Section 3. A salient feature of employment relationships is that there are "winners and losers" (Rosenbaum [1979]). Section 4 explores the role of dichotomous contracts and, more generally, tournament models of the labor market. Hierarachy is

perhaps the most pervasive feature of organizations; it is this feature that, in large part, distinguishes organizational from market relationships (Williamson [1975]). Section 5 analyzes formal models of hierarchical structures and explores issues of the span of control and the assignment of personnel within the organization. Personnel policy, in the sense of screening workers, is further explored in Section 6. The possible contractual relationships are constrained by the extent of the monitoring structure within the organization. Section 7 discusses the value of monitoring, the ranking of monitoring systems, the nature of optimal conditional investigation strategies and the value of communication within an agency relationship. Participative goal setting is a particularly important form of communication within an organization. The role of such processes in situations of divergent goals and incomplete information is explored in Section 8. Participatory goal setting is an example of delegation of decision-making; a more general analysis of delegation is carried out in Section 9. Finally, the possibility of applying agency models in less structured decision settings--in particular, uncertainty regarding the implementation of organization plans, ambiguity concerning the agents' preferences, and uncertainty regarding the efficacy of effort--is discussed in Section 10.

2. Structure of Agency Model

The basic agency model is directed at the issue of employee motivation and, more generally, moral hazard. Since the agent's action is not observable and his goals are not perfectly congruent with those of

the organization, the agent, in general, will not act as desired. If the technology of the task and the goals of the organization are well defined, the organization's problem is the design of an appropriate incentive structure. Some outcome measure of the agent's action is observable, but this outcome indicator is likely to depend on the vagaries of nature, as well as the agent's choice of action. This joint determination of outcomes confounds inferences concerning the action taken by the agent.

The problem of incentive design is solved by first determining the optimal action on the part of the agent for a given sharing rule, that states the agent's payoff as a function of the observable outcome. For any sharing rule the principal may offer, the agent chooses the action that maximizes his own utililty. Given the principal's beliefs concerning what action the agent will select in response to a sharing rule, the principal chooses the sharing rule to maximize his profits. More formally, the principal-agent relationship is modeled as a non-cooperative game, using the solution concept of a perfect equilibrium (Selten [1975]). The analysis is a partial equilibrium one in that there is an exogenously specified minimum level of utility that the contract must, in expectation, provide the agent. This utility represents the agent's opportunity cost of working on behalf of the principal.

Uncertainty is modeled by representing the observable outcome as a random variable whose distribution depends on the action chosen by the agent; $f(x|a)$ represents the probability distribution of output, x, conditional on the agent's action, a. Increased effort shifts the

outcome distribution to the right in the sense of first-order stochastic dominance.[9]

The principal's utility function is denoted by $G(\cdot)$; the agent's utility is assumed to be separable into the utility for money, $U(\cdot)$, and disutility for effort, $V(\cdot)$. The agent's compensation is a function of the observed outcome, x and is represented by $s(x)$. The principal is constrained by the labor market to offer the agent compensation that, in expectation, will equal or exceed the utility he may obtain through alternative employment; $\bar{U}$ represents this opportunity utility level. Having developed the notation, the mathematical representation of the basic problem can be stated as follows:[10]

$$\underset{s(x)}{\text{Max}} \quad \int G[x - s(x)]f(x|a)dx$$

subject to

$$\int U[s(x)]f(x|a)dx - V(a) \geq \bar{U}$$

$$\int U[s(x)]f_a(x|a)dx - V'(a) = 0 \quad .$$

The optimal sharing rule satisfies the following equality:[11]

$$(2.1) \qquad \frac{G'[x - s(x)]}{U'[s(x)]} = \lambda + \mu \, \frac{f_a(x|a)}{f(x|a)} \quad ,$$

where λ is the Lagrange multiplier on the agent's outside utility constraint and μ is the Lagrange multiplier on the agent's first-order condition on his choice of effort. Equation (2.1) characterizes the optimal sharing rule, given that the agent's action is unobservable.

This solution is referred to as the second-best; distinguishing it from the complete information outcome.

Revelation Principle

The basic agency model outline above addresses the problem of moral hazard. Recently, Myerson [1979] and others (Dasgupta, Hammond and Maskin [1979]; Harris and Townsend [1981]) have developed a powerful device, the revelation principal, for analyzing issues of adverse selection. The revelation principle states that the only incentive scheme that a principal need consider is one in which the agents' message to the principal consists of a statement of their type, those characteristics that distinguish them and their particular setting in the organization. Of course, a statement of who one is and what one's place in the world is is not trivial; the message may be extraordinarily complex and it may not be realistic to suppose that it can be communicated in all settings. Whether or not the message may be communicated, the revelation principal may still be used as a methodological device to characterize the outcomes achieved from the actual mechanisms employed.

Myerson [1979] proves that the set of revelation equilibria (i.e., Bayesian-Nash equilibria in which it is incentive-compatible to reveal one's true type) is the same as the set of response plan equilibria, where a response plan is an arbitrary strategy. If a choice mechanism is Bayesian incentive-compatible, then no agent can expect positive gains from being the only one to lie about his true type when all other agents plan to tell the truth: "universal honesty is an equilibrium for the players if and only if the choice mechanism is Bayesian incentive-

compatible."[12] In this structure, it is assumed that each player does
not know the messages sent by the other players and, therefore, compares
his expected utility conditional only on his own type. Let F repre-
sent the feasible set of expected allocations and F* be the set of
Bayesian incentive-compatible expected allocations. By definition, it
is not possible to obtain revelation equilibrium outside the set F, but
it might be possible that a choice mechanism could have other equilibria
(i.e., one that involved some 'dishonesty' in the revelation of type)
that would generate expected allocations with greater utility than those
in F*. Consider a general response plan for player i, σ_i, that maps
each player's type, t_i, onto a probability distribution over his set of
possible responses. A collection of response plans forms an equilibrium
if no player would ever expect to gain from unilaterally changing his
plans. Myerson [1979] defines this set of equilibrilum feasible expec-
ted allocations to be F**.

Myerson [1979] proves that equilibrium feasibility is not more
general than incentive feasibility. For any response-plan equilibrium
of any choice mechanism, there is an equivalent incentive-compatible
mechanism giving players of all types the same expected payoffs.
The principal can implement any expected allocation in F** as an
equilibrium of a revelation game employing the policy $\sigma^{**}(t_i)$, where
$\sigma^{**}(t_i)$ is the agent's optimal strategy in the response plan equili-
brium and t_i is the agent's reported type; the agent has no incentive
to lie to the principal if $\sigma^{**}(t_i)$ is an equilibrium to the original
game. Therefore, there is no loss in generality in restricting the

principal to select an incentive-compatible mechanism with the response sets consisting of the possible types of agents -- F* is equivalent to F**. Myerson [1979] refers to this as the revelation principle, since it asserts that there is no loss of generality in assuming that the principal should structure his incentive system so that all agents will be willing to reveal all of their information to him honestly.

Myerson [1979] points out that the concept of Pareto-efficiency must be re-interpreted in a world of incomplete information: "It is unreasonable to base normative standards on comparisons with plans that are known to be unimplementable."[13] The Pareto-efficiency criterion makes sense with respect to F* but not with respect to F.[14] Incentive-compatibility may be a severe restriction in the sense that F* may be a much smaller set than F, with much of the Pareto frontier of F lying outside the incentive-compatible set F* (Hurwicz [1972]). Myerson [1979] refers to mechanisms that generate allocations on the frontier of F* as incentive-efficient. A mechanism is interim incentive-efficient if and only if it is a Bayesian incentive-compatible choice mechanism and is not strictly dominated by any other Bayesian incentive-compatible mechanism.[15]

Myerson [1980] formulates the general principal-agent problem in which agents have both private information and private decisions, unobservable to the principal. The set of all possible types for agent i is denoted T_i, where each $t_i \in T_i$ represents a complete description of all the private information i might have such as his environment, abilities or preferences. The set of possible decisions of agent i is

represented by D_i where each $d_i \in D_i$ represents the agent's decision, such as level of effort. The principal's decision domain is denoted D_o. A decision alternative, $d_o \in D_o$, may represent an allocation of resources among the agents or a compensation scheme; in addition, any actions by an agent that the principal can observe and thereby control directly are appropriately viewed as part of the principal's decision domain. Myerson [1980] assumes that the principal can control all communication among the agents and that he can send messages and recommendations regarding d_i to the agents. Given this structure, the principal's problem is to coordinate his decision with those of the agents in order to maximize his expected utility.

Myerson [1980] shows that the principal can restrict himself to incentive-compatible direct coordination mechanisms in which agents report their information to the principal, who then recommends to them decisions forming a correlated equilibrium (Aumann [1974]).16/ A coordination mechanism is considered direct (Dasgupta, Hammond and Maskin [1979]) if and only if the set of messages, M_i, equals the set of decisions, D_i, and the set of responses, R_i, equals the set of types, T_i. In a direct coordination mechanism, each agent is asked to report his type to the principal and in return the principal will send the agent a suggested decision. A direct mechanism is incentive compatible if and only if the honest-obedient participation strategies form an equilibrium, where honest means that $m_i = t_i$ and obedient means that $\sigma_i(d_i, t_i) = d_i$, where the agent's optimal response, $\sigma_i(\cdot)$, is a function of the principal's suggested decision and his true type.17/ Given any

equilibrium in any coordination mechanism, there exists an incentive-compatible direct coordination mechanism, in which the agents are honest and obedient and the principal gets the same expected utility as in the given equilibrium of the given mechanism--the optimal incentive-compatible direct coordination mechanism is also optimal in the class of all coordination mechanisms.

If the principal himself has private information, then it may be possible for the agents to make inferences about the principal's type from the principal's choice of mechanism. With this new information, the agents may find advantage in lying and disobedience. More formally, the coordination mechanism may not be incentive-compatible even though it satisfies the existing incentive-compatibility condition, for the case of an uninformed principal. In a subsequent paper, Myerson [1981] takes up this more subtle question of mechanism design by an informed principal. The basic dilemma that an informed principal faces is that to conceal his information the choice of mechanism must not depend on his information, but his preferences over alternative mechanism will depend on his information. Myerson [1981] develops a theory of inscrutable mechanisms to resolve this choice problem; a mechanism must be incentive-compatible after all other individuals have inferred whatever information might be implicit in the establishment of the mechanism itself.

3.   Risk-Sharing and the First-Best

A division of analytic treatments of agency relationships among substantive domains, of necessity, must be somewhat arbitrary due to the interconnectedness of many of the issues.   In particular, essentially all agency models address, either implicitly or explicitly, the issue of risk-sharing.   Risk-sharing under conditions of moral hazard is the primary focus and motivating question of the early work on agency relationships (Harris and Raviv [1978]; Ross [1973]; Spence and Zeckhauser [1971]; Stiglitz [1974]).   The relationships between equity holders and management, landowner and tenant farmers, and insurers and insured have provided the motivation for much of these analyses.

The issue of risk-sharing can be examined in its purest form when the agent incurs no direct disutility from his choice of action (Ross [1973]; Wilson [1968]).   The family of first-best sharing rules can be characterized by assuming that the agent and the principal cooperate to choose a fee schedule that maximizes a weighted sum of utilities:

$$\underset{s(x)}{\text{Max}} \quad E\{G[x - s(x)] + \lambda U[s(x)]\} \quad ,$$

where λ is a relative weighting factor.   If $G(\cdot)$ and $U(\cdot)$ are monotone and concave, the solution is obtained by maximizing the function internal to the expectation, resulting in the following equality:

(3.1) $G'[x - s(x)] = \lambda U'[s(x)]$.

Borch [1962] refers to this as the Pareto-efficiency condition and it defines the fee schedule as a function of the payoff and the welfare

weight, λ. The Pareto surface can be mapped out by varying the weight λ; this is equivalent to varying $\bar{U}$, the utility constraint on the agent.

An alternative approach to characterizing the first-best fee schedule, referred to as the similarity condition, was proposed by Wilson [1969]. This solves for the fee schedule by setting

$$(3.2) \qquad G[x - s(x)] = aU[s(X)] + b \quad ,$$

for constants a and b, where a > 0. If the fee schedule satisfies the similarity condition, then the agent and principal have identical attitudes towards risky payoffs, since the utility of one is an affine transformation of the utility of the other; therefore, the agent would choose the act that the principal desires. Ross [1974] shows that the class of utility functions such that both the Pareto-efficiency and similarity conditions are satisfied are the $<G(x),U(x)>$ pairs with linear risk tolerance; in particular, $G(x)$ and $U(x)$ are such that the following condition holds:

$$- \frac{G'\left(x - S(x)\right)}{G''(x)} = cx + d \quad \text{and} \quad - \frac{U'(x)}{U''(x)} = cx + e \quad ,$$

where c, d, and e are constants.

Ross [1973] analyzes what set of assumptions regarding the set of utility functions or outcome function lead to a first-best outcome. The solution of the principal's problem (equation (2.1)) coincides with the first-best if and only if μ is zero, or, if $\mu \neq 0$ and $f_a(x|a)/f(x|a)$ is a function of a alone. In order to achieve a

first-best for all outcome functions, the motivational constraint must not be binding, implying that μ is zero; the agent must choose the principal's most desired act as a result of solving his first-order condition. For any fee schedule, $s(x)$, the principal wants the act to be chosen so as to solve the following first-order condition:

$$\int G(x - s(x)) f_a(x|a) dx = 0 \quad .$$

If this is to be equivalent to the motivational constraint on the agent, then the similarity condition (equation (3.2)) must be satisfied.

Wilson [1979] examines a similar structure in which the agent's action consists of sharing private information. Sharing information is an important illustration of an action that can be reasonably viewed as not involving disutility of action on the part of the agent, but having important implications for the agency relationship. In this case, the information regards the appropriate investment or size of the enterprise. The output shares can be designed to depend on the information revealed so as to provide incentives for truthful revelation of information. In particular, if efficient risk-sharing can be achieved by linear sharing rules, then linearity of sharing rules results in truthful revelation of information.[18]/

When the agent derives disutility from the action itself (for instance if the action represents effort), then the prospects for first-best risk-sharing become more remote. Holmström [1979] shows that if the agent's disutility of action is increasing in the level of action, $V'(a) > 0$, and the cumulative distribution of outcomes, $F(x|a)$, is

non-increasing in the level of action and strictly decreasing for some outcome value, then $\mu > 0$, where μ is the Lagrange multiplier on the agent's first-order condition on his effort choice. A comparison of equation (2.1), that characterizes the second-best sharing rule, with equation (3.1), a characterization of the first-best sharing rule, reveals that if $\mu > 0$ the sharing rule does not achieve a first-best allocation of risk. The agent, in the second-best contract, must bear some portion of the risk associated with the uncertainty concerning the state of nature, even if the principal is risk neutral (Shavell [1979]). The reward scheme provides an incentive for the agent to work hard, but it comes at the expense of the benefits of risk sharing. Agents are, to some extent, held accountable for events over which they have no control. This basic feature of the optimal incentive contract violates established notions of equity; the notion of ex-ante fairness replaces the principle of equal pay for equal work.[19]

The incentive effect of deviating from first-best risk-sharing can be seen to be stronger the larger is $|f_a(x|a)|$, and it is more costly, in terms of lost risk-sharing benefits, the greater is $f(x|a)$. $|f_a(x|a)|/f(x|a)$ may be interpreted as a benefit-cost ratio for deviation from optimal risk-sharing; equation (2.1) states that such deviations should be made in proportion to this ratio (Holmström [1979]). $f_a(x|a)/f(x|a)$ is the derivative of the log of the likelihood function, $f(x|a)$ and is used in maximum likelihood procedures for estimating an unknown parameter a of a distribution given a sample observation x.[20] The principal imposes risk on the agent by

compensating the agent more when the outcome is "good" and less when the outcome is "bad." The sign of $f_a(x|a)/f(x|a)$ determines whether the outcome is "good" or "bad" news (Milgrom [1981]).

The observed outcome is used for two purposes: risk sharing and as a signal of the agent's action. Grossman and Hart [1983] point out that on the one hand the agent's output contributes positively to the principal's utility, and, therefore, the principal desires a high output. On the other hand, the output is a signal to the principal about the agent's level of effort. This informational role may be in conflict with the consumption role. For example, there may be a moderate output level which is achieved when the agent takes low effort levels that rarely occurs at other effort levels. If the agent is penalized whenever this moderate output occurs, he is discouraged from taking these low effort actions. However, there may be lower output levels which have some chance of occurring regardless of the agent's action. To encourage the agent to take high effort levels, it is then optimal to pay the agent more in low output states than in moderate output states, even though the principal prefers moderate output levels to low putput levels. If the outcome distribution satisfies the monotone likelihood ratio property, then $f_a(x|a)/f(x|a)$ is a strictly increasing function of x; for these distributions, the higher the outcome is the more likely it is that the agent chose a high level of effort.[21] If the outcome distribution satisfies the monotone likelihood ratio property, then $\mu > 0$ implies that the sharing rule given by equation (2.1) is an increasing fuction of x.

Harris and Raviv [1979] describe the contract between the principal and the agent as the trade of three "goods": "risky return," "nonrisky return," and the "agent's action." They show that if the agent's action can be traded (i.e., observed), an allocation can be achieved which is Pareto-superior to any allocation made when the action cannot be traded. In these circumstances, the sharing rule consists of a forcing contract in which the agent is compensated if a specified action is taken and receives no compensation if an alternative action is taken. Even if the agent's action cannot be "traded," under certain conditions it is still possible to design an incentive scheme that results in a Pareto-efficient outcome.

Mirrlees [1974] shows that if the support of the probability distribution of outcomes changes with the agent's action and payments (or punishments) to the agent are unbounded, then the first-best solution can be achieved; Holmström [1982c] proves a similar result in a multi-agent setting.[22/] Lewis [1980] extends the results for a single agent setting by showing that the use of a lump-sum penalty of any size, that is based on a measure of performance which varies continuously with the agent's effort, can improve a contract by reducing incentive for the agent to shirk. It is critical to this line of work (Mirrlees [1974]; Gjesdal [1978]; Lewis [1980]; Holmström [1982c]) that the support of the probability distribution of outcomes change with the agent's action. If the derived distributon of output changes but the support of the distribution is fixed, Holmström [1979] finds that the optimal incentive contract may be a continuous function of output. The importance of the

movement in the support of the outcome distribution is that it establishes the possibility of a probability one inference being made about the agent's choice of action. This possibility of a probability one inference, in conjunction with unbounded payment schemes, allows the principal to design a contract that forces the agent to take the first-best action.

The possible inefficiencies that result from the unobservability of the agent's action can be alleviated in a trivial manner if the agent is risk neutral. Risk neutrality of the agent implies that no social loss results from the agent absorbing all the risk associated with the uncertainty regarding the future states of nature. In this setting, the optimal contract takes the form of the principal receiving a fixed payment and the agent receiving the residual outcome (Shavel [1979]). Essentially, this amounts to the agent purchasing the firm from the principal--agency problems are resolved by ending the agency relationship.

Holmström [1982c] explores how the separation of ownership from production can be used to alleviate the inefficiencies that result from the agency relationship in a multi-agent setting, where the privately taken actions of agents jointly determine output. If agents form a partnership in which they are to share the joint outcome among themselves, then there exists no sharing rule based on the joint outcome alone which induces the choice of the Parteto-efficient action; sharing of the joint output falls into the general class of free rider problems. The separation of ownership and production permits the use of a

nonbudget balancing Groves scheme (Groves [1973]), and, if the joint outcome is a deterministic function of the agents' action, the first-best solution can be achieved.[23] In the case of uncertainty, Holmström [1982c] shows that if increasingly sharp inferences about chosen actions are possible as the joint outcome value declines, then the first-best solution can be approximated arbitarily closely by using group incentives in which all agents are penalized for low output.

Shavell [1979] provides asymptotic results regarding the attainment of the first-best solution as the "efficiency" of the agent's effort either tends towards zero or infinity. Let the index of the efficiency of the agent's effort be denoted γ and the level of effort by a; the probability density of outcomes is $f(x|\gamma a)$.[24] If $\gamma = 0$, $f(\cdot)$ is not affected by a and the problem is merely one of risk sharing, as a result, a first-best solution is achievable. When $\gamma > 0$, $f(\cdot)$ is affected by a and a first-best solution is not achievable; however, as $\gamma \to \infty$, the difference between the achievable solution and the first-best solution tends to zero. When γ is large, little effort is required to markedly change the probability density of outcomes, bringing it nearer to the first-best density; only a slight deviation of the sharing rule from the efficient risk-sharing schedule is needed to overcome the problem of unobservable action. Grossman and Hart [1983] provide a stronger version of this result, showing that the loss between the first and the second-best effort, as a function of γ, divided by γ, approaches zero as $\gamma \to \infty$. Since the marginal product of the agent's labor, that is, the increase in expected outcome resulting

from additional effort, is proportional to $1/\gamma$, the proposition can be interpreted as saying that the welfare loss is of a smaller order of magnitude than the reciprocal of the agent's marginal product of labor.

The repetition of an agency relationship over time will, in general, improve the efficiency of the relationship. Principals and agents often agree to wage contracts that establish dependence between periods so that the current wage depends not only on the current performance but on past performance as well. Lambert [1981] and Rogerson [1982] consider a case of complete but imperfect information; the principal is aware of the agent's abillity but does not observe the actions taken by the agent.[25/]

Lambert [1981] models a finitely, independently repeated agency relationship as a T period game. If neither the principal nor the agent is able to precommit to a multi-period contract, then the T period game is equivalent to T separate one period games. If precommitments are possible, then the optimal contract bases compensation in one period, in part, on performance in earlier periods. Precommitment may not be possible on the part of the agent, but, as long as the principal is able to precommit to a multi-period contract, it is optimal to supplement current period incentives with future period incentives, though not as extensively as in the case in which the agent is able to precommit to remaining with the organizaiton. Rogerson [1982] assumes that both the principal and the agent can commit to a long-term contract and, as Lambert [1981] does, shows that in a Pareto-efficient contract whenever an outcome has any effect on the current wage it must also have an

effect on future period wages.[26/] The spreading of incentives over subsequent periods reduces the variance of the wage payment while still maintaining the agent's incentives. The explanation of the role of memory in repeated contracts with imperfect information lies in the dual role of the wage contract: providing incentives for the agent to choose the desired action and allocating risk away from the agent to the less risk averse principal. These goals are necessarily in conflict since incentives to the agent can only be provided by differentially rewarding outcomes.

Providing the agent with incentives exposes the agent to risk and reduces his utility relative to that which he would obtain from a constant wage. This utililty loss to incentive maintenance decreases (increases) with the expected value of the wage, depending on whether the agent's inverse marginal utility is concave (convex). Since the Pareto-efficient contract always spreads incentive payments into the future, in some sense more incentive maintenance is occurring in later periods of the relationship. If the utility cost decreases with the expected wage, expected wages are optimally chosen to be larger in later periods, where more incentives are being provided; Rogerson [1982] shows that the expected wage payment will rise (fall) over time if the agent's inverse marginal utility is concave (convex). For the class of HARA utility functions, expected wages increase (decrease) over time if the agent's risk tolerance increases at a rate greater than (less than) 1 (Rogerson [1982]).[27/]

The work on repeated agency relationships with imperfect information also provides asymptotic results regarding the achievement of the first-best outcome (Radner [1981b]; Rubinstein [1979]; Rubinstein and Yaari [1980]). Radner [1981b] states that: "The theory of repeated games explores in a formal way...the conventional wisdom...that when members of an organization have long lasting relationships they can encourage and maintain cooperative behavior (without the device of binding commitments) by signalling intentions to punish defectors from informal agreements."[28] As the relationship is extended through time, increasingly sharp inferences can be made about the agent's action if the structure of the relationship is stable.

Radner [1981b] examines the conditions under which members of an organization will cooperate; in particular, how, with no discounting of future income, the duration of the relationship can encourage cooperative behavior. If attention is restricted to perfect Nash equilibrium, then the game must be repeated an infinite number of times in order for cooperation to be induced (Selten [1975]). Radner [1981b], using the concept of an epsilon equilibrium (Radner [1979]), shows that, for any Pareto-efficient cooperative outcome of the one-period game that dominates a one-period Nash equilibrium, there exists a time horizon sufficiently long such that the non-cooperative epsilon equilibrium of the repeated game yields the principal and the agent an average per period expected utility that is arbitrarily close to the expected utility in the one-period cooperative solution.[29] If cooperative behavior is to be sustained as an epsilon elquilibria of the repeated game, the

principal must have some statistical method of detecting non-cooperative behavior on the part of the agent that is sufficiently rapid so as to deter non-cooperation and, at the same time, avoid repeatedly falsely punishing the agent for non-cooperatiave behavior. Radner and others (Rubinstein and Yaari [1980]; Rubinstein [1979]) use the law of the iterated logarithm in constructing detection schemes that punish non-cooperation with sufficiently high probability and punish cooperation with sufficiently low probability that the strategy will have the desired limiting properties.[30/] Radner [1981a] extends these results by showing that a first-best is 'nearly achievable' if that agent's discount rate is close to 1.

4.   Winners and Losers

The notion of winners and losers may be interpreted as discontinuities in reward (Harris and Raviv [1979]; Lewis [1980]) or as rewards based on an ordinal ranking of outcomes (Lazear and Rosen [1981]; Mookherjee [1981]; Green and Stokey [1982]; Nalebuff and Stiglitz [1983]). A particular form of discontinuity is a dichotomous contract.[31/] If the results of monitoring reveal the action to be acceptable, the agent is paid according to a predetermined schedule; otherwise, the agent receives a less preferred fixed payment, such as being fired. Harris and Raviv [1979] show that such a dichotomous contract is optimal given a monitoring technology of the form $\alpha = a + \delta$, where δ is a random variable with support $[-\alpha_0, \alpha_1]$ and α_0, α_1 are finite, positive values. For a particular case, a separable utility function

that is a power function of money and effort, Harris and Raviv [1979] show that as the monitoring technology becomes less precise the worker's wage contingent on the outcome being acceptable must increase, the terms on which the worker is judged become more lenient, the worker's effort decreases, the probability of the outcome being unacceptable increases, and the gains to monitoring decrease.

A common form of dichotomous contracting is to fire workers for unsatisfactory performance. Stiglitz, in conjunction with Shapiro [1981] and Weiss [1983], has examined the impact of terminating the contractual relationship in the context of general equilibrium models of the economy. Stiglitz and Weiss [1983] note that: "Firms often respond to low output of a worker not by lowering that worker's wage but by firing the worker."[32] In the Stiglitz and Weiss [1983] framework, the threat of termination is used to encourage higher effort in the initial period if the second-period utility constraint is binding for the workers who have low performance in the initial period.

In most organizations, favorable rewards take the form of a promotion. In the typical hierarchy, employees at one echelon are competing for a fixed, smaller number of positions at the next higher echelon (Green and Stokey [1982]). The goal of these employees is not just to do well, but to do better than their peers. Independently, but roughly contemporaneously, the metaphor of career patterns within an organization as a tournament emerged in both the behavioral (Rosenbaum [1979]) and economic literature (Lazear and Rosen [1981]). Rosenbaum [1979] tests an ahistoric (path independent) model of mobility versus a

historical, or tournament model. In Rosenbaum's tournament mobility model, careers are conceptualized as a sequence of competitions, each of which has important implications for an individual's mobility chances in all subsequent selections—winners have the opportunity to compete for high levels while losers are permitted to compete only for low levels, or are denied the opportunity to compete any further at all. His empirical analysis supports the tournament model, finding that mobility in the earliest period of one's career has a significant relationship with many important parameters of one's later career, such as career ceilings, career floors and the probability of promotion in each successive period. In the economic literature, the notion of a tournament refers to the evaluation of workers based on their performance **relative** to their co-workers.

In these formulations of the agency relationship, the agents' strategies are modeled as being in a Nash equilibrium with respect to each other. The principal then selects the sharing rule so as to elicit the most desirable Nash equilibrium in the game among the agents. Mookherjee [1981] points out that there are situations in which, at the principal's optimizing choice of incentive scheme, the actions that are required to be sustained as a Nash equilibrium among the agents, are Pareto-dominated by other Nash equilibria. In such cases, the Nash equilibrium selected by the principal may be viewed as not being implementable. Mookherjee [1981] analyses the multi-agent game when the principal is restricted to implementing Nash equilibria that are unique, or, are Pareto-efficient. Under this restriction, the optimal contract

for an individual agent may depend on the output of other agents, even when the random variables affecting the outcomes are independent. This dependence results from the need to obtain Pareto-efficient risk sharing among the agents.

In a situation with independently producing agents who face correlated uncertainties, moral hazard can be mitigated by evaluating each agent against average performance (Holmström [1982c]; Lazear and Rosen [1981] and Mookherjee [1981]). If agents share a common uncertain factor of production, then, as the number of agents grows large, the comparison with average performance will remove the commonly shared uncertainty from the relationship, leaving the uncertainty that pertains to the agent's idiosyncratic risks (Holmström [1982c]). For either additive or multiplicative disturbance terms, optimal contracts can be designed in which each agent's reward depends only on his own output level and on the mean output level for the group of agents; the mean output level is a sufficient statistic for all of the information about the common disturbance term (Holmström [1982c]).

Lazear and Rosen [1981] show that when risk neutral agents are subject to the same uncertainties, compensation schemes that reward agents on the basis of their ordinal ranking in terms of output will induce the same allocation of resources as an incentive scheme based on individual actions; an optimal two-person tournament is equivalent to offering the first-best incentive contract to each agent independently.[33/] Mookherjee [1981] generalizes this work and shows that perfect correlation of the disturbance terms is alone sufficient to achieve

the first-best outcome. Even if the function relating the principal's payoff for the agents' output exhibits decreasing returns, it may benefit the principal to have multiple agents--the benefits of competition among the agents may outweigh the cost of the redundancy of their effort (Mookherjee [1981]). Mookherjee [1981] also establishes that the principal's payoff is continuous in the value of the correlation between the disturbance terms; the results regarding the value of tournaments are not knife edge properties of the correlation of the disturbance terms. Lazear and Rosen [1981] also compare linear piece rates and tournaments for the case of risk-averse agents and a normally distributed noise term. They show that if the variance of the noise term is sufficiently large, the optimal tournament yields higher expected utilities.

Green and Stokey [1983] compare the efficiency of independent contracts and tournaments. Under the former, each agent's reward depends only on his own output level, while under the latter the reward is assumed to depend only on his rank order. They consider a situation in which a risk-neutral principal employs a group of identical risk-averse agents. As in the Lazear and Rosen [1982] model, each agent's output is assumed to depend stochastically on his own effort and a common additive disturbance term. Agents are assumed to observe private signals, correlated with the common disturbance term before they choose their effort levels. The realized output of each agent then is a stochastic function of his effort and the value of the common disturbance term; the principal only observes the output levels of the agents. For

any finite number of agents, in the absence of a common disturbance term, using the optimal tournament is dominated by using optimal independent contracts (Green and Stokey [1982]). In the absence of a common disturbance term, the output levels of the other agents convey no information about the effort of an agent and using a tournament in this case only introduces extraneous noise into the payoff function that the agent faces. Conversely, given any group of at least two agents, if the distribution of the common error term is sufficiently diffuse, then the optimal tournament dominates using optimal independent contracts (Green and Stokey [1982]). In this situation, using tournaments eliminates a major source of noise, while adding a relatively minor one. In addition, given any fixed distribution for the common disturbance term, for a sufficiently large number of agents, using the optimal tournament dominates using optimal independent contracts. In fact, if the number of agents is sufficiently large, a principal, who cannot observe the value of the common disturbance terms using an optimal tournament, can do as well as a principal who can observe the value of the disturbance term and uses independent contracts. For a large group of agents, the rank order is an extremely accurate signal about an agent's output level net of the common additive disturbance term.

In addition to providing an effective incentive in a static sense, Nalebuff and Stiglitz [1983] point out that "compensation schemes based on relative performance have the further advantage of automatically adjusting incentives to changes in the economic environment."[34] They refer to this feature as "built-in flexibility." When a task is easier,

the individual's rewards for performing the task should be reduced. If pay is based on relative performance, although all individuals perform better when they exert the same level of effort, their compensation is automatically adjusted.

5.   Hierarchy

Alchian and Demsetz [1972] explain the payment of the residual output to a manager, basing their arguments on the free rider problem in situations of joint production. The role of the manager (principal), in the Alchian and Demsetz [1972] framework, is to monitor the activities of the other agents and meter out payments in accordance with the contractual agreement; the payment of the residual to the manager provides the incentive for the manager's supervision of the other agents. They claim that this is the only model appropriate to firms, and that pure authority relationships do not occur.

In the Alchian and Demsetz [1972] vision of the employment relation, workers voluntarily agree to be supervised and work harder than a non-hierarchical incentive system would induce them to work. The workers submit to being compelled to work harder than direct incentives provide for, because it results in a higher level of expected utility; firms that impose some degree of compulsion are able to pay higher wages. Although each worker may resent this element of compulsion and feel that it is unnecessary on his own part, he prefers to work for firms that impose this compulsion, recognizing that without it some of his colleagues would not exert as much effort. Along similar lines,

Stiglitz [1975a] suggests an additional rationale for supervision may be that individuals with information about their abilities are willing to pay to have these abilities recognized so that they can capture their ability rents.

In deciding on whether to accept a contract with this element of compulsion, individuals who dislike authority relationships may require considerable additional compensation to induce them to accept this type of employment relationship. One of the disadvantages of this type of relationship is that it may not allow for individual variability or variability over time (Stiglitz [1975a]). The viability of hierarchical structures and monitoring reflect the fact that for some individuals the disadvantages are outweighed by the greater wages and more certain incomes that firms who use these contracts can pay. The commonly expressed attributes of self-proprietors reflects this trade-off--those who choose to work for themselves are more tolerant of risk, place a premium on flexibility and have a distaste for supervision.

If monitoring is to resolve the possible inefficiencies in the agency relationship, a sufficiently rich set of observable measures must be available. Holmström [1982c], in the case of joint production under certainty, investigates the conditions for a monitoring scheme to be sufficiently rich. It suffices that the principal be able to detect when an agent is the only one who is deviating from the first-best action. If the joint outcome is monotone with respect to the agents' action, then as many measures are needed as there are agents.

It is not obvious that the asymmetric solution commonly used and assumed optimal by Alchian and Demsetz [1972] of one party specializing

in monitoring is in fact optimal, when the means of production are owned in common. Clearly, implicit in this solution is some notion of economies of scale in monitoring. Mirrlees [1976] suggests investigating under what circumstances it would pay a group of workers to have one of their number undertake all performance observation, and when it would pay instead to have a symmetric solution in which each worker devotes some of his time to monitoring one of the others. Stiglitz [1975a] views the supervisor as obtaining information that otherwise would be lacking; whether there should be a supervisor will depend on the value of this information relative to the costs of the supervisor. Another function of the supervisor is his role as a decision-maker. Stiglitz [1975a] develops an argument for centralized decision-making based on non-convexities in the value of information (Radner and Stiglitz [1975]).35/ The value of a given amount of information increases with how frequently it can be used. Even if the initial cost of acquisition of information of the employee and the supervisor were the same, in a hierarchical structure the supervisor makes use of the information more frequently and is therefore of greater value to him. This advantage of specialization may be offset by the fact that workers may find it easier to acquire certain categories of information; a natural joint product of the performance of their tasks may be the acquisition of certain information (Stiglitz [1975a]).

Mirrlees [1976] develops a model of a hierarchy that provides an illustration of the limits to the desired size of an organization. In Mirrlees's [1976] model, it is assumed that aggregate output is

correctly measured at each level of the organization and that accurate accountants set the wage bill off against output in every department. If

      (i)   each supervisor has the same number of immediate subordinates,

    (ii)   errors of observation are proportional to the mean size of the observed variable,

  (iii)   the payment rule is proportional to output at each level,

   (iv)   all individuals are identical,

profit per worker diminishes as the number of workers increases. The income of the proprietor is approximately:

$$Z \sim (AN + B + N^{1/2}\varepsilon)z \quad ,$$

where z is the value of output per worker and N is the number of workers, A and B are constants, and ε is a random variable that is independent of N. Despite the decreasing returns, if the cost of other inputs, such as capital, is sufficiently low per worker (less than zA), the entrepreneur would prefer a large organization. Another implication of the model is that if members of the organization have constant or decreasing relative risk aversion, then payment schedules should be less concave, less like instructions and more like profit-sharing at higher levels of the organizaiton, since the relative riskiness of income is smaller.

Rosen [1982] models the hierarchical structure of firms and examines the assignment of a diverse pool of personnel to hierarchical positions within and among firms, and the distribution of firm size and managerial rewards. To understand these issues it is necessary to explain the production function of the firm, deriving it form the underlying technological and control relationships. Three production functions of the firm are distinguished: management, supervision, and production. The top layer of the organization represents management, whose function involves discrete and indivisible choices--such as which goods to produce, in what varieties and volume, and how to produce them. These commands are issued to the second level, that supervises the bottom level, the instrument of actual production.[36/]

A central characteristic of all hierarchical organizations is the delegation of authority within and between levels. In this simplified structure, these aspects of the firm are described by the number of levels in the hierarchy and the breadth of control within each level. Decisions at each level spread over a wider range as one moves up the hierarchy, consistent with the common notion of increasing "responsibility" at higher levels. Indivisibilities inherent in the command system imply increasing returns to scale to decision-making as one moves from the bottom of the organization towards the top. The loss of control that arises when decisions are delegated through a larger organization, due to opportunistic behavior, transactions costs and limitations on information transfer, provides a countervailing force to the increasing returns to decision-making. The form of the actual organization strikes

a balance between these two forces that depends on the technology of production and talents of workers, supervisors, and managers.

Rosen [1982] regards the output of each level as an intermediate product that is processed or enhanced by the activities of the next highest level. The productivity of worker i in the production activity is denoted q_i; q_i is a measure of skill, or labor in efficiency units, varying from worker to worker.[37] The skill of a supervisor is denoted s_i and is also measured in efficiency units. The amount of supervisory time allocated to worker q_i is t_i. The product attributable to supervisor of type r, supervising q_i is:

$$x = g(r)f(rt_i, q_i) \quad ,$$

where $f(\cdot)$ has the standard properties of neoclassical production functions of positive and diminishing marginal products and a positive cross-partial derivative (complementarity).

The supervisory skill parameter r is both time augmenting and total factor productivity improving. A supervisor with greater skill has more effective time or a greater efficiency of time to allocate. Total factor productivity $g(r)$ represents the effects of skills that are independent of time and which apply equally to all workers who are controlled by that particular person; $g(r)$ reflects the extent of "management" as opposed to supervision.[38] The amount of this public good, $g(r)$, is bounded by the necessity of devoting time to actual supervision and the limit to a supervisor's time.

The third level of the hierarchy controls the output of multiple supervisors, analogous to the locus of control exercised by a supervisor over a group of workers. The output of a manager of talent s who manages supervisory output X_j (the output of a supervisor of talent r_j) is:

$$y_j = G(s)F(sv_j, X_j) \quad ,$$

where v_j is the amount of time manager s allocates to control of X_j. The functions $G(\cdot)$ and $F(\cdot)$ have the same general properties and interpretations as $g(\cdot)$ and $f(\cdot)$ -- $G(\cdot)$ representing the discrete managerial elements of decision-making that are independent of to whom the orders are given and $F(\cdot)$ representing the limits to the span of control through the necessity for monitoring and supervision.

Increasing returns to managerial decision-making are not sufficient to explain a convex salary structure by rank; there may be many people willing to work at high-level positions at a low price. In order to explore the earnings profile, supply factors must also be considered. The matching of authority with talent through the labor market, in conjunction with the technology of production, determines the final production outcomes and the internal structure of decision-making of the organization. Each worker is described by a vector of latent skills, (q,r,s), indicating the amount of skill potentially available to each level of activity. The actual skill that is used is determined by the rank to which the person is assigned, with all other skills remaining latent and unutilized. The available or potential skills in the overall

labor market are given by the distribution $h(q,r,s)$. A special case of interest arises if latent skills follow a homogeneous, one-factor structure such that:

$$q = a_q + b_q \epsilon$$

$$r = a_r + b_r \epsilon$$

$$s = a_s + b_s \epsilon \quad ,$$

where a_i and b_i are positive constants and ϵ has the natural interpretation of general ability. The problem can now be stated as finding an assignment from the distribution of latent skills to ranks in a set of hierarchical firm structures that maximizes the value of total output in the economy. This requires specifying a set of firms and their internal structure, with respect to both the breadth and depth of the hierarchy, and a corresponding allocation of the distribution of latent skills, indicating those skills that are actually utilized.

Each worker is a unique factor of production, whose productivity depends on both the number of other workers in the firm and the precise skills embodied in each of them. In general, it matters who works with whom even though there are no direct interactions in production. Supervisory effort is allocated to equate its marginal value over all workers; diminishing returns and complementarity imply that more supervisory effort is allocated to the more able workers. In the case of constant returns to scale, supervisory time is allocated in strict proportion to worker skill: $t_i = kq_i$. The ratio of supervisory talent to total labor

resources controlled is denoted by β; $1/\beta$ can be thought of as the span of control. The supervisory span of control depends upon the technology (i.e., the production function), supervisory skill, and the price of a unit of labor. Rosen [1982] shows that the span of control rises as the price of labor falls; greater supervisory talent commands greater resources; and the span of control increases in r, except when $g'(r) = 0$, in which case it is a constant. Because of the assumption of perfect quality-quantity substitution, there is no systematic correlation between r and the size of the firm as measured by the number of workers; however, there is systematic correlation between r and total efficiency units of labor controlled. There is also a systematic correlation between skill r and firm size as measured by X -- more capable supervisors produce larger outputs.

The distributions of earnings within groups have greater weight in their right-hand tails than the corresponding distributions of measured abilities. This results from the fact that more able persons control greater resources. The ability distribution is partitioned such that all persons with general ability greater than some critical value ε^* are assigned to supervisory positions and those with less ability become production workers. The partition itself, ε^*, is determined by the relative prices of labor and the production technology. The relation between income or productivity and ability is flatter for workers, implying that there is less scope for talent to stand out. Anyone can do about the same amount of work in the production activity, whereas talent stands out and has a large marginal effect in the "higher" skill

positions. The overall distribution of earnings in the entire population of persons must be more skewed to the right than the underlying distribution of talent, ε.

When $g(r)$ is increasing in r, this creates additional impetus for expecting greater relative weight in the right-hand tail of the observed distribution of earnings among second-level personnel compared with the distribution of those at the bottom. In this case, the span of control increases with r and firm size increases more than proportionately with r. Therefore, the distribution of firm output is right skewed relative to the true distribution of supervisory talent and even more skewed relative to the distribution of innate ability.

In a three level hierarchy, if $G(s)$ is increasing in s, then both output and income distributions among the firms are more skewed than the distribution of true managerial talent. There is an indivisibility of management and scale economy; this explains why top managers of large firms are so highly paid and why their incomes are closely associated with firm size. Higher ability managers control more than proportionately larger firms and have more than proportionately larger marginal products.[39]/ Correspondingly, there is an enormous increase in salary as productivity increases. Though less able people could manage the firm, the productivity of their services would not yield the same total output; output would fall by more than the opportunity cost of their services in some other position, and they optimally go to a lower position in a smaller firm.

6. <u>Personnel Policy</u>

Rosen's [1982] work focuses on the design of hierarchical structures and assumes away informational issues regarding the worker's ability or choice of action. The rationale in Rosen's [1982] model for a hierarchical pay structure within organizations is based on the demand and supply of workers of particular abilities. Alternatively, hierarchical pay structures within organizations may be justified by the incentives they provide; consideration of the incentives created by the pay structure should be an important element in explaining it. To explore these matters, Mirrlees [1976] suggests a rather simple and extreme model in which workers are completely aware of their own abilities and choose how hard to work on that basis, while employers cannot distinguish among job applicants and, therefore, set pay schedules which are applied to all comers. The behavior of workers, which they determine themselves in light of the pay structure, is assumed to determine the firm's output.

Mirrlees [1976] shows that the firm must have some monopsony power in the labor market if wages are to differ from the worker's true marginal product. A particularly interesting example of monopsony power is that of a stagnant or declining firm, in which case there would be expected to be some inelasticity in response to wage reductions. Mirrlees [1976] points out that a firm with an established labor force may be able to increase its profits by changing to a payment schedule which will no longer attract new workers. This would be possible if the firm has acquired information about its workers that allows them to be

placed in the organization appropriately and, thereby, makes them more valuable to their current firm than to other firms. It is not in the interests of the current employer to tell a prospective employer what grades or positions the worker has achieved and, therefore, the employees cannot expect the same wage in another firm as those who occupy similar positions to their own.

Thus, at least in the case of employees with some seniority, the employer may have some monopsony advantage. Exercising this power will reduce recruitment of workers who have sufficient ability to aspire to higher paid positions, but the policy may still be appropriate for a declining firm. The policy results in a less steep wage payment schedule than would otherwise be instituted; the question remains of how flat a payment schedule it could profit a firm to adopt, given the disincentive effects on effort. If incentive considerations were ignored, the optimal mark-up of marginal product over the worker's wage would be proportional to their level of ability. The incentive considerations result in a larger mark-up than monopsony theory suggests for the less skilled and a lower mark-up for the more skilled; yet, the **absolute** difference between marginal product and wage is higher for the more skilled.[40] This result depends on assuming that higher ability is associated with lower elasticity of supply. Mirrlees [1976] argues that this should be the case because of the investment in reputation and knowledge of the firm, which a worker of higher ability normally makes.

Individuals are often screened by means of their choice from a menu of options offered by employers; such a device is referred to as a

self-selection mechanism and causes an individual to reveal information about himself by his choice behavior. There may be benefits to both the organization and the individual to screening workers at the initial stage of the employment relation. These benefits may result from improved resource allocation to on-the-job training, or better job placement. Salop and Salop [1976] derive a self-selection equilibrium for a model in which individuals differ in their propensity to leave the organization. The firm discourages high turnover individuals from applying and encourages low turnover workers to apply for employment by predictably increasing an employee's wage with his tenure at the firm. This has the effect of the applicant essentially guaranteeing his longevity with the firm, since he pays the consequences in terms of foregone higher earnings if he should quit prematurely. A similar self-selection structure may be used to validate abililty (Guasch and Weiss [1980]). These self-selection mechanisms are based on performance tests: "Although the employer is using information from self-selection, self-selection only works because of the performance tests."[41]/ If there were no possibility of failures, everyone would make the choices associated with the more able workers.

Holmström [1982b] explicitly models the process of the principal's learning about the agent's ability over time. In this framework, the agent chooses his effort level in an attempt to influence the principal's perception of his ability. Holmström [1982b] shows that the stationary level of effort is never greater than the first-best level and approaches the first-best level as the agent's ability is more

subject to stochastic influences over time, the observations on output become more accurate, and the discount rate declines. The first two properties make the updating of the ability estimate each period more informative and, thereby, make labor investments in influencing perceptions of ability more valuable.

Holmström [1982b] shows that the sequence of optimal labor supply converges monotonically to the stationary state. If the initial precision of information about ability is less than the stationary precision level, then the convergence of effort is from above. Conversely, if the initial precision is greater than the stationary level, the convergence is from below. One would normally expect the precision of information about ability to increase over time; as a result, young workers invest more in labor supply because the returns from building a reputation are highest when the market information is most diffuse.

In the same paper, Holmström [1982b] examines the incentives for risk taking in the context of the firm's personnel policy. Concern is frequently expressed over the risk attitudes of management--particularly common is the complaint that managers are overly risk averse. Holmström [1982b] conjectures that a major source of this incongruity in risk preferences between managers and equity holders stems from the manager's career concerns. An important element of managerial talent relates to projecting investment returns and choosing attractive prospects. If talent is not fully known, investment decisions become tests that provide information about talent; investments are risky from the manager's perspective due to their impacts on perceptions of his talent, even if his income is not explicitly tied to profits.

Holmström [1982b] considers a manager who is in charge of choosing investment projects for a risk-neutral firm. Investments can either fail or succeed; let y_- be the payoff if a project fails and y_+ if it succeeds. Talent is associated with the likelihood that investments are successful; the likelihood that a project succeeds is ℓ_T if the manager is talented and ℓ_N if he is not. In this set-up, investment projects can be characterized by a vector $I = (y_+, y_-, \ell_T, \ell_N)$ and the pool of potential projects is a collection of such I's. The manager's expertise lies in observing this pool, while the principal does not. From the pool of potential projects, the manager is assumed to choose at most one and proposes it as an investment. Such a proposal involves presenting the information I in a verifiable way to his superiors who will make the final decision. Holmström [1982b] addresses the issue of whether the proposed project is the best available alternative from the firm's perspective and ignores the potential incentive problems associated with misrepresenting information about a proposed project.

If the manager is risk-neutral, then he is indifferent between all projects, and presumably proposes the project that the firm most prefers. The results are considerably different for a risk-averse manager. The expected return from undertaking an investment is no higher than abstaining from investments altogether. Since investing carries risk, it is then clear that the manager has no incentive to invest. The manager would claim that no worthwhile investment opportunity was present in the pool of potential investments and, under the informational assumptions made in the model, such a claim cannot be invalidated.

The risk facing the manager is quite different from the risk that is of concern to the firm. A key variable for the manager is the likelihood of success, ℓ_T. The manager dislikes investments, which will reveal accurately whether he is a talented manager or not, since these investments make his income most risky, and prefers investments which leave him protected by providing exogenous reasons for the failure of the investment. The firm, however, has no interst in ℓ_T, given the expected probability of success; it is primarily concerned about the actual payoffs, (y_-, y_+), of the project, which are irrelevant from the manager's perspective. This incentive problem seems to call for having the manager own some share of the firm in order to bring greater convergence of preferences; however, Holmström [1982b] points out that giving risk averse managers a share of the firm may not be a desirable policy, since it carries downside as well as upside risk. A stock option is suggested as a more valuabale incentive, since it removes the downside risk to the manager.

Holmström [1982b] goes on to show that if the manager cannot communicate the investment risks in a verifiable manner, the incentive problems become more severe. Beliefs about the manager's ability are updated under the assumption that the expected return from the investment is positive. If a risk neutral manager is rewarded according to expected marginal product, computed based on the rule that investment takes place if its expectation is non-negative, the manager would not conform to this rule. The manager would take less risk because of the negative talent evaluation that follows upon failure. The firm's

updating of beliefs about talent conditional on the general knowledge that the investment was positive in expectation, places the manager in an unfavorable position if the expected gain is close to zero. Without verification of investment risks, the only equilibrium outcome is for there to be no investments, even with a risk neutral manager. This incentive problem regarding the choice of investment project is mitigated if the manager has more information about his ability than the firm does. In this setting, an undervalued manager would be willing to take risks in order to prove himself and risk taking would become a signal of talent (Spence [1974]). Holmström's [1982b] work may help to explain why investment procedures in firms are so detailed and centralized; it may have as much to do with securing a proper evaluation of managerial talent as it has to do with controlling what projects get selected.

Marcus [1982] confirms Holmström's [1982a] conjecture that providing managers with some ownership in the firm will not resolve the problem of excessive risk aversion on the part of the managers. The dependence on the outcome (profits) of the firm can create a portfolio diversification problem for the manager who cannot sell short shares in the firm. Marcus [1982] shows that compensation schemes in which the no-sale constraint is binding lead managers to underinvest in risky assets and expend excessive resources on activities that reduce the variability in profits. Marcus [1982] shows that, under the optimal compensation scheme, the manager's compensation includes equity holdings in the firm that exceed the level in the manager's optimal portfolio: the portfolio

diversification constraint is always binding.[42] The existence of firm-specific human capital reinforces the effects of profit contingent compensation. If skills are more valuabale to the firm for which the manager currently works than in other firms, the manager acquires an interest in the success of the firm, since bankruptcy would eliminate the rents he can capture from these firm specific skills.

As Holmström [1982b] does, Marcus [1982] suggests that stock options might usefully be part of the manager's compensation.[43] Call options increase in value with the variance of the value of the firm (Merton [1973]), managers who own options have an incentive to increase the variance of firm value. This incentive can counteract the variance reducing tendency of the manager. As Marcus [1982] points out, options are not a panacea. Because option values depend on total (as opposed to systematic) risk, agents who hold options may take high variance projects that equity holders would reject on the basis of present value and systematic risk criteria. In general, any incentive scheme for which the compensation pattern across states of nature differs significantly from the pattern of stock returns is subject to distortion of investment incentives.

7. Information Systems

Baiman [1982] identifies three uses of managerial information systems: improve the manager's ex-ante decision, motivate subordinates and facilitate the sharing of risk. The latter two properties are sufficiently interrelated that they are lumped together under the label as

performance evaluation. Information may be gathered after the agent has chosen his action and the outcome has been generated in order to evaluate the agent's action choice and determine his compensation. Here, the information is used for motivational as well as risk-sharing purposes: Demski and Feltham [1976] have termed this use decision-influencing. Alternatively, information may be supplied to the agent before his action choice in order to improve his choice; this is referred to as information's decision facilitating role by Demski and Feltham [1976]. In the first case, information is gathered about the agent and publicly reported to both the agent and principal, while in the second case information about the state of nature is acquired and possibly supplied only to the agent. Preference rankings over information systems used strictly for performance evaluation purposes, need not be the same as rankings for the purposes of belief revision in decision-making (Gjesdal [1982]).

Even when there are gains to perfect information about the agent's action, it may not be possible to achieve gains by means of imperfect monitoring. The introduction of imperfect information of the agent's action produces two opposing effects on the welfare of the parties to the contract: First, the uncertainty introduced by the information will tend to reduce the welfare of the agent and a risk averse principal; second, the introduction of monitoring may motivate the agent to take an action which, neglecting the first effect, would make both parties better off.

Investigation along this line leads to one of the important early results of agency theory—that a signal of effort, that is a process measure of the employee's action, is of value if the observable outcome of the employee's action is not a sufficient statistic of the signal (Holmström [1979]).[44] That this is a necessary condition for a signal to have value is fairly intuitive. That the condition is also sufficient is not at all obvious and implies that an agent, regardless of his degree of risk aversion or the level of noise in the signal, will benefit if compensation is made contingent on a process measure of his input as long as the outcome measure is not a sufficient statistic of that process measure. To the extent that rewards may be made contingent on input measures, the agent is less subject to the unsystematic risk of the state of nature. Along related lines, Harris and Raviv [1979] show that if the state of nature is observable or the agent is risk neutral, then there are no gains to monitoring. Under these circumstances, the first-best outcome can be achieved **without** monitoring and, therefore, there is no informational value in monitoring. This has the important methodological implication that it is not possible to simplify analyses of monitoring by assuming away either exogenous uncertainty or risk aversion on the part of the agent.

In addition to showing that an information system has value relative to no information, it is important to compare the relative value of different information systems. Gjesdal [1982] compares mutually exclusive information systems and characterizes a ranking of information systems that does not depend on the particular agency problem. This

ranking may be viewed as a generalization of Blackwell's ranking of experiments (Blackwell [1951]).[45/] It is not, however, a trivial extension of Blackwell's theorem, since the incentive problem is a non-cooperative game in which the actions are chosen by players in the game rather than by nature. Choosing an information system for the agency relationship is not formally equivalent to choosing information systems in a Bayesian decision problem. As an example of the difference in the two situations, consider that for incentive purposes an information system that detects shirking may be replaced by a system that detects this shirking with positive probability.[46/] Indeed, Gjesdal [1982] shows that more information according to the Blackwell ranking is sufficient but not necessary for an information system to be as valuable as another system in the context of the agency problem. There are even cases in which randomization, less information in the Blackwell sense, is efficient.

The sufficiency part of the theorem regarding the value of monitoring generalizes to arbitrary current information systems and multidimensional actions (Gjesdal [1982]). However, when one information system is replaced by another system, marginal informativeness is no longer a necessary condition for marginal value, since an additional information system may be valuable as information about the outcome as well as about the agent's action. The ability of the agent and principal to share risk depends on the precision with which they can measure the principal's ex-post marginal utility; however, the transfer payment is a function only of the expected marginal utility, conditional on the

given signal from the information system. Incremental information provides information relevant to risk-sharing if and only if it is false that the expected marginal utility of the principal is conditionally independent of the given information.

Gjesdal [1982] demonstrates that an information system is marginally valuable whenever it provides information about the agent's action that is not already available. The increased risk imposed on the agency by the revised contract does not matter (at the margin). The reason for this is that the incentive effect of more information is a first-order effect (influences the first derivative of the objective function), whereas the risk effect is of second order (influences only the second derivative). However, the incentive effect of randomization is a second order effect as well. It follows that randomization is valuable only if the incentive effect is large enough relative to the negative effect of more risk.[47/] Risk has a positive incentive effect if the agent's risk aversion decreases when **a** changes in a direction which is preferred by the principal.

Gjesdal's [1982] main result on randomization as an incentive mechanism may be summed up as follows: if the incentive constraint on the agent is binding and the agent's optimal action is unique and interior, then the transfer payment should be randomized when signal y_i obtains if

$$\frac{1}{R(y_i)} \frac{\partial R_A(y_i)}{\partial a} < -\frac{1}{\mu} \frac{E\{U'(x - S(y_i))|y_i\}}{V'(S(y_i))} \quad ,$$

where R_A is the agent's absolute risk aversion, and the agency's risk aversion

$$R(y_i) \equiv \frac{\dfrac{\partial^2 EU(x - S(y_i))}{\partial S^2(y_i)}}{\dfrac{\partial EU(x - S(y_i))}{\partial S(y_i)} - \dfrac{V''(a)}{V'(a)}} \quad ,$$

where $S(y_i)$ is the optimal transfer given the signal y_i.

Harris and Raviv [1978], Holmström [1979] and Shavell [1979] focus on the demand for unconditional evaluation mechanisms. Efficiency gains may be possible if the decision to produce additional information is made conditional on preliminary performance data; conditional investigation of variances has the advantage of allowing the choice of whether to produce costly information to depend upon preliminary observations of performance. Baiman and Demski [1980] characterize optimal investigation strategies for a class of principal-agent problems; they investigate a situation in which the initial information consists of the outcome itself and a further signal of action is obtainable at some cost. While variance investigation has a long history in the quality control literature, it is important to note that the process being investigated in the quality control context is not responding to the control policy. In quality control problems, the optimal strategy takes the form of two-tailed investigation policy of investigating if the outcome is sufficiently far from the standard. Baiman and Demski [1980] consider whether an analogous policy is desirable in an agency setting, under the assumptions that the monitoring technology is conditionally independent

of the production technology, the principal is risk neutral and the agent's utility function is a member of the HARA class of utility functions.[48/] Given these restrictions, Baiman and Demski [1980] show that if the agent is relatively risk averse then a lower-tail investigation strategy is optimal and that if the agent has relatively little risk aversion then an upper-tail investigation strategy is optimal.

Given an arbitrary level of effort and outcome realization, the investigation is effectively a lottery from the agent's perspective (Baiman and Demski [1980]). Since the monitoring technology is imperfect, monitoring the agent will result in an uncertain signal of action and, accordingly, uncertain compensation to the agent. The threat of being investigated will make those outcomes triggering the investigation relatively less desirable for a very risk averse agent and, thereby, provide motivation for the agent to act so as to reduce the probability of those outcomes being realized. Since the firm's outcome is positively related to the agent's effort, the principal can best exploit the conditional investigation policy by lower-tail monitoring, that is, investigating only when the observed outcome is less than some cutoff value.

For a less risk averse agent, the threat of inducing an outcome which will trigger an investigation becomes less compelling, since the disutility from facing any given lottery is reduced. To keep the threat of being investigated effective, it has to be increased either by making the lottery more unfavorable or by increasing the probability of inducing the lottery. Since compensation functions are bounded from below

and investigation is costly, it may be more efficient to utilize the agent's risk-tolerance rather than to try to overcome it; it may be optimal to make the investigation a favorable lottery so that the agent prefers to be investigated. Thus, monitoring is most efficiently used as punishment for the more risk averse agents, but as a reward for the more risk tolerant agents.

The two-sided investigation strategy common to the quality control literature is not optimal when the "process" is a strategic actor rather than an exogenous process. If the principal were to choose a two-sided investigation strategy, there would exist low outcome/low effort realizations which a risk averse agent would prefer to high outcome/high effort realizations; such a situation clearly could never be efficient. Analogous comments can be made regarding the relatively risk tolerant agent.[49/]

Baiman [1981] points out that the role and value of a pre-decision information system is more complex than that of a post-decision information system. Expanding a post-decision information system to report additional information will always result in a weak Pareto improvement, since the principal and agent can agree to a sharing rule that ignores the additional information. However, expanding a pre-decision information system to report additional information need not result in even a weak Pareto improvement.[50/] This is true whether the additional pre-decision information is privately or publicly reported. With private pre-decision information, the principal no longer even knows what action is optimal for the agent, from either the agent's or principal's own

perspective. Baiman [1981] notes that: "While a better private pre-decision information system may allow the agent to be better informed and capable of making better...decisions, it may also reduce his motivation by reducing the risk...improving the agent's pre-decision information system may exacerbate the moral hazard problem."[51] Christensen [1979] constructs an example in which making the agent better informed results in the principal being made worse off. Even for the case where communication of the pre-decision information is allowed, Christensen [1979] provides an example in which the principal is made worse off if the agent has access to private pre-decision information. Baiman and Evans [1981] develop sufficient conditions under which the acquisition of private pre-decision information by the agent results in a weak Pareto improvement.

The agency framework has also been used to examine issues of information transfer (Crawford and Sobel [1982], Green and Stokey [1981a], [1981b]). In these papers, the decision maker (principal) and the transmitter of the information (agent) are treated symmetrically in that the principal's choice of decision rule and the agent's choice of revelation strategy is a Nash equilibrium.[52] Crawford and Sobel [1982] assume that the private information of the informed agent consists of observing the realization of a random variable, and that the random variable as well as the decision taken by the receiver of the message enters the utility function of both the transmitter of the information and the decision-maker. They further assume that the utility maximizing value of the decision variable is increasing in the realized value of

the random variable. Crawford and Sobel [1982] show that the agent's equilibrium strategy calls for him to partition the support of the probability distribution of the variable that represents his private information and, in effect, introduce noise into his signal by reporting only in which element of the partition his observation lies. This represents a compromise on the part of the agent between including enough information in the signal so that the principal responds to it and holding back sufficient information so that the principal's response is as favorable as possible. Using a similar framework Green and Stokey [1981a, 1981b] study the effects of improved information on the principal's and agent's utility. In general, it is not possible to give a Pareto ordering of pre-decision information systems.

Green and Laffont [1982] study the impact of limited communication possibilities on the design and performance of incentive compatible mechanisms. The complexity of information places some limits on the possibility of its full communication and utilization--costs of transmission, storage and information processing are among the factors that could cause a principal to limit the potential for information flows. Green and Laffont [1982] model these constraints by restricting the dimensionality of the message space in a resource allocation mechanism. Messages must condense the private information into some summary statistics, and it is these summary statistics that determine the decision outcome. The problem of incentive compatible design of an implementable plan under these restrictions encompasses both the choice of these summary statistics and the way in which they will be used. Even with

full communication, it is only under special conditions that the difference in objectives between principal and agent can be overcome and a fully efficient outcome achieved. When further constraints are imposed on the problem, it is natural to expect a greater departure from optimality. Green and Laffont [1982] discuss the loss due to the incentive compatibility constraints, in addition to that due to the restricted communication capability.

Green and Laffont [1982] assume that the agent's private information consists of the realization of two, normally distributed, variables. They solve the principal's problem with a one-dimensional and a two-dimensional message space and both with and without the incentive compatibility constraints on the agent's action. With a two-dimensional message space the presence of the incentive-compatibility constraints does not decrease the value of the principal's problem; however, with communication restricted to only a one-dimensional message, the incentive constraints become binding. In general, the principal is hurt by both the limited ability to receive information from the agent and the agent's ability to distort his private information.

The set of plans that are implementable depends only on the agent's utility function; while, the best plan within the set of implementable plans depends, in addition, on the distribution of the private information and on the principal's objective function. For agents with quadratic utility functions, only summary statistics of the form $a(\theta_1, \theta_2) = \lambda\theta_1 + \mu\theta_2$ can be induced as response rules by the agent, where θ_1 and θ_2 represent the agent's private information.

The correlation between the two bits of information, σ_{12}, can be interpreted as a parametric representation of the amount of information available to the agent relative to the amount of information that can be transmitted through the communication channel. When $\sigma_{12} = 1$, the problem is effectively one-dimensional and there is no loss due to the restriction on information transfer. If $\sigma_{12} = 0$, the principal would select the summary statistic based on just one of the two pieces of information. Since this is incentive compatible, there is no loss due to the incentive effects. For intermediate values of σ_{12}, the principal's solution of the problem with a one-dimensional channel of communication requires a choice of statistic and decision rule that are in conflict with the incentive compatibility constraint. Therefore, the loss due to the incentive constraints will be non-monotonic in the amount of information the agent has relative to what he can transmit.

8. Goal Setting

The goal setting process serves two purposes: one is to make plans for the organization's future operations and the other is to create incentives for members of the organization. One of the factors which makes the goal setting process difficult is that the available information is dispersed among the members of the organization. The principal has to rely on the agents as sources of information for purposes of predicting or decision-making, since they are closer to the production process and can be expected to better know the production prospects. The behavioral literature on participation considers

participation as conceptually divisible into act and result. The mere act of participating in the goal setting process might change the motivation of the agents, because the participants get to know each other and feel as members of a group (Vroom [1964]). The analytic models of participative goal setting do not consider this effect of participation; only the resulting outcome and its subsequent effect on motivation and planning is considered.

The idea behind participatory goal setting is that the principal tries to obtain some of the information to which the agents have private access. The agent is viewed as participating in setting his own standards if his compensation function is a non-constant function of the signal he reports to the principal. Similarly, the agent is participating in the decision process if the decision strategy adopted by the principal is a non-constant function of the agent's reported signal. If either the agent is participating in setting his own standards or in the decision process, the budgeting process is said to be participative (Christensen [1979]).

In the prior section, we examined when a Pareto improvement resulted from the agent acquiring private pre-decision information. In the context of participative goal setting, we consider when a strict Pareto improvement can be achieved by allowing for communicaion between the agent and the principal, subsequent to the agent receiving private information and prior to the firm's output being revealed. The distinction of a strict versus weak Pareto improvement is important; since the principal and agent can agree on a payment schedule that ignores the

agent's message, it can never be Pareto inferior to allow for communication.53/ If the agent's private pre-decision information is perfect, then communicaiton has no value; observing the firm's output in this case allows the principal to infer the agent's private pre-decision information. However, if the agent's information is imperfect, a necessary and sufficient condition for communicaton to be strictly valuable is for the truthful revelation of the agent's private pre-decision information to be strictly valuable (Baiman and Evans [1982]).54/

Since the direct statements of workers concerning the difficulty of the job which has been assigned to them lack credibility, the principal seeks to infer information about the difficulty of the task from the revealed behavior of the workers. In the absence of information, jobs of differing difficulties receive the same wage; a distortion results since the marginal rates of substitution between consumption (income) and effort will differ on different tasks. Although the information reduces this distortion, making inferences about the job difficulty on previous performance has a cost: it reduces the incentive to work by in effect lowering the wage rate. Stiglitz [1975a] shows that under a simple linear revision scheme, some revision of the wage rate will, if the base period is sufficiently long, be desirable.

Furthermore, Stiglitz [1975a] indicates that workers all can be made better off if they could collude not to work so hard; taking the payment schedule as given, the workers observe that by colluding they can increase their expected utility. Workers who are working hard are exerting a negative externality on the other workers. The firm,

observing that the supply of effort does not correspond to the competitive equilibrium among the workers, responds by changing the parameters of the payment schedule. Once the firm has adjusted the payments schedule to take account of collusion, agents will observe that if they fail to collude their expected utility will fall below their reservation utility level.

More generally, there are two basic incentive problems associated with standard schemes that reward quota overfulfillment or penalize quota underfulfillment. The first difficulty is, what Weitzman [1976] calls a **static** problem. It is in the interest of the manager to convince his superiors that production is likely to be small, thereby entitling him to a lower target and a bonus that is more likely to be attained. The second difficulty is the **dynamic** incentive problem, arising out of the tendency of planners to use current performance as a partial basis for setting future targets. The agent may be tempted to hold back output in the hope of inducing a lower quota the next period. A target that is set too low will not ordinarily lure agents to overfulfill by a conspicuous margin because next period's plan target may start off with this period's performance as a point of departure--the "ratchet effect."

The static incentive problem biases the agent's statement of production possibilities downward to induce a lower quota, whereas the dynamic incentive problem biases the agent's performance toward lower fulfillment levels, given the quota. Both incentive problems would be eliminated by doing away with a target and instead setting rewards in

some fixed relation to output. This is a satisfactory solution in a situation where planning is not necessary because coordination of activities is not needed; the basic rationale for planning is the need for coordination. Makng the bonus size depend on the plan target should mitigate the tendency of agents to underrepresent their potential in seeking low assignments.

Fan [1975], Bonin [1976], Weitzman [1976] and Thomson [1979] consider the situation in which the production possibilities of the organizaiton are independent of the agent's actions and analyze schemes that result in truthful revelation of the production possibilities. One can think of output next period as a random variable, whose true distribution is better known by the agent; the principal's aim is to elicit this distribution, or some statistic of it. After production has taken place, the principal observes the level of output and, therefore, the reward schemes may depend on the actual output as well as the agent's report. Fan [1975] proposes a scheme linear in two variables, a target level chosen by the agent and the final output; the scheme is such that it is optimal for the agent to choose as the target the level of output that can be produced, in his estimation, with 50 percent chance. In a subsequent contribution by Bonin [1976], a more general scheme is developed for which it is the manager's best strategy to announce the level of output that can be produced with probability r, where r is any value between 0 and 1 determined by the principal. In Weitzman's [1976] model a tentative target and a tentative bonus are chosen by the principal. The agent then proposes a new target, and a new bonus is

computed. Finally, production takes place and the bonus is adjusted again and the agent is compensated. The bonus scheme studied by Weitzman [1976] is linear in the three variables--tentative target, revised target and final output--and it is shown, as in Bonin [1976], that the manager's best response is to announce the level of output that can be produced with probability r, where r is a simple function of the parameters that determine the bonus. Thomson [1979] shows that revelation schemes that are bounded exist, so that limits can be placed on the size of necessary monetary transfers, and also establishes the existence of individually rational schemes, in which all transfers are from the principal to the agent.

The crucial assumptions underlying these models are:

(1)   the level of output is purely a random variable that cannot be influenced by the actions of either the principal or the agents;

(2)   the principal has no ability to observe (even probabilistically) the capacity of an agent; and

(3)   the agents are all expected-reward maximizers.

Thomson [1979] maintains these three basic assumptions and provides a characterization of all reward schemes that will generate a self-imposed quota that can be achieved with a given probability. Using this characterization, Conn [1979] is able to show that if resource allocation decisions depend on the agents' report, then it is impossible for an elicitation scheme to be an optimal incentive structure; when the

information generated by the production agents for the purpose of setting quotas is also used to allocate scarce resources, the Bonin-Weitzman procedures lose their incentive compatibility. Elicitation schemes induce producers not to divulge their expected output, but to reveal the output level that is achievable with a pre-specified probability. Even when behaving truthfully, the agents do not convey appropriate information; the report yields only median-like information that may not bear any particular relationship to expected output, the statistic needed to optimally allocate resources. Elicitation incentive compatibility can be retained in the presence of resource allocation if and only if resources are allocated randomly, indicating a complete incompatibility between elicitation procedures and rational resource allocation (Conn [1979]).

In addition, these elicitation schemes assume away the multi-period gaming problem associated with the planner's use of past performance as a basis for setting future standards. If the agent discounts future earnings, this dynamic incentive problem is mitigated to the extent that the lag between present performance and future target or coefficient setting is lengthened. Weitzman [1980] and Holmström [1982a] analyze the ratchet effect assuming that the principal can commit himself to an intertemporal scheme.[55/] The principal announces the current scheme as well as the revision procedure; the agent then solves the dynamic programming problem given the principal's plan. Holmström [1982a] shows that appropriate use of revision rules outperforms no revision, in and of itself, revision is good rather than bad.

Revision of goals is simply a way of utilizing information about production potential for mutually beneficial adjustments in targets. Holmström [1982a] criticizes the prior literature for focusing on conditions under which reported targets will correspond to true (or best estimates) of production potential. When the principal receives a signal from such a reported information system, he of course anticipates the subordinates' possible slack-creating behavior and takes that into account when and if the information is used; the existence of false information (presumably mitigating the organization's productive capabilities or exaggerating difficulties faced by the agent) need not imply that the organization loses any profit due to the subordinate's 'sandbagging' policy, if it is anticipated by the principal. Any scheme which leads the agent to give different reports depending on the content of his private information will provide valuable signals about production potential. Indeed, schemes with reporting of the same dimension are equally good from a purely informational point of view and should be compared in terms of their implications on final production choices (Holmström [1982a]).

If the agent's welfare does not depend on plant output directly, then he could be given a fixed wage and there would be no reason to expect the agent to falsify reports (Holmström [1982a]). Thus, the incentive issue is only meaningful in a context in which there are costs to production or some private benefits to certain outcomes; it is assumed that it is the cost of effort that creates the incentive problem. The appropriate incentives will depend on the production potential

and production costs which constitute the agent's private information. Holmström [1982a] formalizes these concerns with an explicit objective function that incorporates the actual output, information about what the output will be for coordination purposes, and the agent's costs. A simple formalization of these concerns leads to the following objective function:

$$W(x) = G(x) - H(x - t) - V(x) \quad ,$$

where $G(x)$ is the value of output gross of information benefits and production costs, $H(x - t)$ is the cost of diverging from expected output t and $V(x)$ is the firm's cost of production.

Holmström [1982a] examines whether there is any value in letting the firm decide on the target level of production in a one period setting, with $H(\cdot) \equiv 0$. It is assumed that the agent can determine x with certainty and that the principal does not know $V(x)$. Given this structure, Holmström [1982a] proves that there is an interval, (t_1, t_2), and a $\beta > 0$ (where β is the proportion of the difference between the agent's reported target and the initial target set by the principal added to the agent's bonus), such that letting the agent choose its target from this interval according to a Bonin-Weitzman incentive scheme will make both the agent and the principal better off than under the optimal, non-participatory goal setting scheme. The restriction to a certain range of outputs is necessary; in general, complete flexibility of goal setting need not dominate a fixed target.

The ratchet principle (i.e., the tendency of planners to revise targets based on previous performance) has been negatively viewed in the "Soviet incentive" literature (Weitzman [1976]; Snowberger [1977]; Ekern [1979]; Weitzman [1980]) that has focused only on its discouraging effect on high performance; yet, a compensation scheme that allows revision of the target may be preferable to one that does not, since reporting a single target is a very narrow channel of communication. The target cannot convey all the relevant information--some of the remaining information can be signaled via performance and used as a basis for revision of the reward scheme. Revision is a form of delegation and, if properly administered can dominate a fixed target (Holmström [1982a]).

The remaining question is whether delegation is preferable to revision. Delegation allows an immediate move to the jointly preferred reward structure without the cost of signaling and should, therefore, be preferable to revision, where the choice of production is mixed up with signaling. Revision and delegation are similar except that with delegation the target can be changed freely, whereas with revision the first-period cannot be changed at all and the second-period target changes only with output. Additionally, if coordination of output is important, then delegation would have a further advantage in that revision schemes do not communicate anything in the first period and their signals for the second period are presumably more noisy than direct delegation.

9.   Delegation

As discussed in the previous section, if agents possess private information, the principal may wish to let them participate in the decision-making process. When the principal's decision is dependent on the information transmitted by the agent, the decision itself is, in some sense, delegated. Thus, the discussion of information transmission (Crawford and Sobel [1982]; Green and Stokey [1981a, 1981b]) and of participative goal setting (Christensen [1979]; Conn [1979]; Holmström [1982a]; Thomson [1979]; Weitzman [1976, 1980]) can be interpreted as examining the issue of delegation. In this section, an explicit process of delegation is considered in which the agent chooses the decision variable within the bounds set by the principal. This joint decision-making process is typically a complete iterative process of proposals and counter proposals, leading up to the ultimate choice; yet, analytically, the process can be modeled as a mapping from a sufficiently complex message space to a decision (Holmström [1982a]). The model structure Holmström [1982a] proposes is quite general and may cover a wide range of situations--for example, the decision could be capital allocations within the organization and the decision mechanism a formal investment budgeting routine or the decision could be a set of reward schemes with the agent's messages representing suggested targets.

Despite its generality, the structure has important limitations. The decision is ultimately made by the principal and there is no uncertainty regarding the implementation of the decision. It is also assumed that the principal chooses the decision function prior to receiving the

signal regarding the state of nature; if this were not the case, the principal could not choose the decision rule arbitrarily (Myerson [1981]).[56/] Finally, the cost of information coordination is ignored. This cost may be sufficiently high that it is preferable for the organization to operate with the agents making decisions independently of one another; such a restriction corresponds to a decision function in which each agent's decision is a function solely of his own signal. Unless the agent has some private information, delegation of decision-making responsibility will yield no gains to the principal. At the other extreme, if the preferences of the principal and agents coincide, then full delegation is optimal (Holmström [1982a]). Apart from these extreme cases the use of the agent will, in general, entail some, but not full, delegation of choice.

The question then is when and to what extent it is worthwhile to delegate. One simple, seemingly sufficient condition is that if d^* would be the principal's choice if he acted alone and there exists a decision, d, such that when this decision is preferred to d^* by the agent, it is also preferred by the principal (given the information that the agent's choice of d would reveal to the principal), then letting the agent choose between d and d^* is preferred to having the principal act alone (Holmström [1982a]). Unfortunately, Holmström [1982a] shows that this test is **not** sufficient--there are cases where optimal delegation may lead the agent to make a choice which the principal immediately afterwards regrets. The choice is still optimal in an ex-ante sense; if the principal decided to enforce another decision after

hearing the agent's choice, then that would change the agent's behavior and the information revealed by his choice. Holmström [1982a] shows that the best the principal can do is to stick to his proposed decision policy, even if he occasionally regrets it.

In delegation, the agent is given freedom to make decisions subject to constraints imposed by the principal. A central issue is the relationship between the agent's freedom of choice, the information he possesses and the divergence of preferences. Holmström [1980] addresses this in the decision context of one-dimensional quantity controls. The delegation process proceeds as follows: the principal chooses a control set $C \subset D$ and the agent observes some signal, the agent then determines the final decision by choosing $d \in C$.[57/] The principal is concerned with controlling the production of a single good of which an amount d can be produced at a cost $C(d,z)$, for a benefit of $B(d)$. The principal knows $B(d)$, while the agent is aware of the value of z.

Holmström [1982a] follows Weitzman [1974] in approximating $B(d)$ and $C(d,z)$ by quadratic functions:

$$B(d) = b + B'(d - d^*) + \frac{1}{2} B''(d - d^*)^2$$

$$C(d,z) = c(z) + (C' - h(z))(d - d^*) + \frac{1}{2} C''(d - d^*)^2$$

where d^* is the decision an uninformed principal would make. Weitzman [1974] considers two modes of control: setting centrally a quantity d, or setting a price p and delegating the choice of d to the agent. He shows that the best price to set in the price mode is $p^* = B' - C'$

and that the price mode is better than the quantity mode when $B'' - C'' > 0$, independent of the distribution of z. Holmström [1980] considers a more elaborate form of delegation that mixes a policy of prices and quantities. The principal can set a price and in addition require that d $[d_L, d_U]$. If z has a symmetric distribution, then it is clear that the optimal price to set will again be $p^* = B' - C'$. Supressing the price decision, what remains is a one dimensional control problem, with D and preference functions:

$$U^P(d,z) = B(d) - C(d,z)$$

$$U^A(d,z) = p^*d - C(d,z) \quad .$$

With a normally distributed signal, y, the agent will be given a finite degree of freedom $[d_L^*, d_U^*]$ that will include the best centralized act, d^*. The production unit is given more freedom with an increase in information and the principal is made better off by the increase in information. In addition, the agent is given more freedom when there is either a decrease in the curvature of the benefit function or an increase in the curvature of the cost function.

The result regarding improved information is not general. Improved information may reduce the agent's freedom and make both parties worse off (see Christensen [1979] and Green and Stokey [1981b] for examples of this in other decision contexts). The result regarding the inclusion of d^* in the constraint set and the value of delegation is considerably more general and depends on the agent's and principal's preferences being, in some sense, coherent (Holmström [1982a]). Coherence is

essentially the requirement that if for one signal y the agent prefers d_1 to d_2 and the principal d_2 to d_1, then there cannot be another signal for which the reverse is true; the principal's and the agent's preferences move in the same direction as y changes. Without such conformity, delegation will frequently be valueless.

10.   Agency Models and Less-Structured Organizational Environments

The prior sections have examined complex, but relatively well structured incentive problems within oranizations. The question remains as to the extent of the domain of problems over which the agency structure is applicable. This section illustrates the applicability of agency models in settings of ambiguity about preferences and uncertain technology. We have already discussed the application of agency models in situations of problematic implementation of "technically" feasible plans in the context of Mookherjee's [1981] work on tournament models of performance evaluation. Mookherjee [1981] points out that there are situations in which, at the principal's optimizing choice of incentive scheme, the actions that are required to be sustained as a Nash equilibrium among the agents are Pareto-dominated by other Nash equilibria. In such cases, the Nash equilibrium selected by the principal may be viewed as not being implementable. Mookherjee [1981] analyses the multi-agent game when the principal is restricted to implementing Nash equilibrium that are unique, or, are Pareto-efficient.

Hurwicz and Shapiro [1977, 1978] examine a situation in which the principal does not know the agent's disutility of effort and

productivity parameters. If the principal had a prior probability distribution on the space of utility functions, he would have a well-defined Bayesian decision problem. However in their work, Hurwicz and Shapiro [1977, 1978] assume that only an a priori admissible class of utility functions is specified: the principal knows that the worker's utility function belongs to this class, but he has no information or beliefs about which element of this class he faces. For instance, the principal may know that the worker's utility is linear in reward and quadratic in effort, without having any information as to the coefficient values. In such a situation of incomplete information, the principal cannot, in general, persue simple gain maximization. Hurwicz and Shapiro [1977, 1978] assume that the principal will maximize the minimum of an efficiency measure related to the gain.[58] They show in Hurwicz and Shapiro [1977] that among a broad class of "smooth" reward functions, a 50-50 split is best for the principal when there are constant returns to effort and the agent's utility is linear in reward and quadratic in effort.

In existing models of the agency relationship, the technology of the task is assumed to be known by both the organization and the individual worker. Technology, in this context, is taken to mean the relationship between the worker's action and the realized outcomes. This relationship may be probabilistic, but both the worker and the organization know the distribution of outcomes for a given action. In fact, workers are often highly uncertain of the impact that their actions may have on the outcomes that are ultimately observed. If an individual

feels relatively ineffective in a particular task environment, then offering additional financial rewards for higher outcomes is a costly and inefficient incentive. In the extreme case, if the worker feels that outcomes are independent of his choice of action, then offers of additional rewards, no matter how extreme, would have no effect on the worker's choice. Changing the agent's perception of the relationship between action and outcome, in these instances, would be a less costly means of eliciting greater effort on the agent's part. The only credible way for the organizaton to change the agent's beliefs is for the organizaton itself, in some sense, to take seriously the outcomes of the agent's action, such as in the promotion and firing procedure.

The agency framework is capable of addressing a variety of organizational issues in a structured and systematic manner. Furthermore, the development of the revelation principal eases the analysis of adverse selection issues in the context of the basic moral hazard problem on which agency theorists have focused. Agency models have become central to recent work in managerial accounting (Baiman [1982]) and constitute the response of the mainstream of microeconomics to the gaps left by the neoclassical theory of the firm.

Footnotes

1/   Jensen and Meckling [1976], p. 306-307.

2/   Cyert and March [1963], pp. 15-16.

3/   Ross [1973], p. 134.

4/   Jensen and Meckling [1976], p. 309.

5/   Mirrlees [1976], p. 107.

6/   Thompson [1967], p. 84.

7/   In terms of agency models, this is a setting in which the first-best outcome is achievable by means of explicit contracting.

8/   Eisenhardt [1983], p. 1.

9/   A distribution F_1 is said to dominate F_2 in the sense of first-order stochastic dominance if for every function, $U(x)$, increasing in x the following inequality holds: $\int U(x)dF_1(x) > U(x)d_2(x)$. This inequality holds if and only if $F_1(x) \leqslant F_2(x)$, with strict inequality for at least one value of x. Intuitively, this can be thought of as stating that any decision-maker whose utility is increasing in x prefers the lottery F_1 to F_2 (Milgrom [1981]).

19/   In carrying out the analysis, it is assumed that the agent's choice of effort can be represented by setting the first-order condition on effort equal to zero. In order to represent the agent's choice of effort with the first-order condition, it is necessary that each reward structure elicit a unique action on the part of the agent (Mirrlees [1975]; Holmström [1979]). If there is not a one-to-one correspondence between reward structure and action, then the application of calculus and first-order conditions is not justified. In fact, it is not sufficient that the agent's choice of effort is unique and interior—it is necessary also to ensure that there are no actions that satisfy the agent's first-order condition on effort but are not global maximizers. Given a sharing rule, $s(x)$, the agent's optimal action may be unique and interior, but the principal is unable to distinguish among those actions that satisfy the first-order condition which one is optimal for the agent. This problem is discussed in Mirrlees [1975] and Grossman and Hart [1983].

11/   An optimal sharing rule may not exist. Holmström [1979] provides an example of $f(x|a)$ distributed normally with a mean of a and

variance of 1. For this example, $\left(\mathbf{f}_a(x|a)\right)/\mathbf{f}(x|a) = x - a$. If
$\mu \neq 0$, $\lambda + \mu\left(\mathbf{f}_a(x|a)\right)/\mathbf{f}(x|a)$ can be negative as x approaches
either plus or minus infinity. Since the left hand side of equa-
tion (2.1) must be positive, it is not possible to have $\mu \neq 0$.
However, if $\mu = 0$, the sharing rule will not give the proper
action incentive when $V'(a) > 0$. In order to avoid problems of
existence of an optimal sharing rule, bounds are placed on the
allowable payments to the agent and the set of possible sharing
rules is assumed to be compact. Holmström [1979] provides exam-
ples of compact sets of sharing rules such as
$S = \{c \leqslant s(x) \leqslant d \,|\, s$ is nondecreasing in $x\}$ and
$S = \{c \leqslant s(x) \leqslant d \,|\, s$ is continuous in $x\}$. If the sharing rule is
restricted to be a member of either of these two families, then a
solution to the optimization problem exists. The bounds on pos-
sible payments may be justified on the basis of institutional
features such as limited liabililty and bankruptcy. Within this
class of sharing rules, Holmström [1979] shows that when the share
is strictly between the upper and lower bound it will satisfy
equation (2.1).

12/ Myerson [1979] pp. 63-64.

13/ Myerson [1979] p. 67.

14/ The Pareto frontier of F are considered first-best outcomes,
while the Pareto frontier of F^* are regarded as the second-best
outcomes. When the incentive-compatibility constraint is binding
the frontier of second-best outcomes are strictly inferior to
those for the first-best outcomes.

15/ Interim efficiency is the appropriate standard since the agents
know their own type but have yet to reveal their information via
their actions (see Holmström and Myerson [1981]).

16/ Strategies of the players are correlated if they are statistically
dependent; the choice of action is dependent on the same mutually
observable random variable.

17/ This restriction is not equivalent to breaking the incentive
constraint into two parts: one a consideration of the agent's
message and the other a constraint on the agent's choice of
action. The joint incentive compatibility constraint implies that
the two considered separately would hold but the converse is not
true. There may be some lie that the agent would want to use, but
only in conjunction with some planned disobedience.

18/ Truthful revelation is not a dominant strategy, but is a Nash
equilibrium: if each other shareholder adopts the strategy of
truthful revelation then the best strategy for any one shareholder
is truthful revelation.

19/ In empirical work on evaluation systems, Dornbush and Scott [1974] find that a mismatch between responsiblity for performance and controllability of that performance leads to dissatisfaction-- expressed by turnover, non-compliance, and verbal expressions of dissatisfaction.

20/ This interpretation is not quite accurate; given $s(x)$, the principal knows what action a rational agent will take, but the intuition provided is useful.

21/ A density $f(x|a)$ has the monotone likelihood ratio property if for every $x_1 > x_2$ and $\bar{a} > a$, the following relation holds:

$f(x_1|\bar{a})f(x_2|a) - f(x_1|a)f(x_2|\bar{a}) > 0$. Among the families of densities with this property are the normal, exponential, Poisson and uniform.

22/ The support of the distribution are those outcomes that occur with positive probability. If the support of the outcome distribution shifts with the agent's choice of action, then there are outcomes that may serve as a perfect signal of the agent's action.

23/ In the Groves mechanism, each agent receives a payment from the principal consisting of the sum of all other agents' payoff, assuming they carry out the optimal action, minus a constant. Given this compensation function, it is in each agent's interest to choose the optimal joint action. Groves [1973] examines dominant strategy equilibrium; each agent's choice is optimal independent of the other's action. D'Aspremont and Gerard-Varet [1979] examine a similar structure in which they explicitly introduce agents' beliefs about each other's type and solve for the Bayesian-Nash equilibrium. They show that if the beliefs of the agents are the same for all agents, there exists an efficient mechanism that is both Bayesian incentive-compatible **and** budget balancing; they also establish this result for a slightly weaker condition on the compatibility of the agents' beliefs.

24/ The definition of efficiency is motivated as follows. Suppose that $f(\cdot)$ is determined by the quantity of some good purchased by the agent and that a is expenditure on the good. Then if p is the price of the good, the quantity purchased is a/p, so the density is $f(x|a/p)$ and $\gamma = 1/p$.

25/ When information is imperfect but complete, the agent's type is known with certainty and only his actions cannot be accurately monitored, there is no reason to fire the agent (Rogerson [1982]). The history of outcomes contains no new information about the agent nor about the agent's choice of action, since the principal can determine what action would be optimal for the agent to select. Creating incentives for the agent to act correctly by the

threat of firing results in a loss of the net surplus which the remainder of the relationship would have generated. Making the agent pay a penalty instead of firing him can generate the same incentives for the agent and result in the principal capturing the surplus associated with continuing the relationship. Therefore, if contracts are enforceable, contracts involving a firing rule are a result of incomplete rather than imperfect information.

26/ Rogerson [1982] points out that if the agent could not commit to remaining with the organization, then contracts with falling wages would not be enforceable.

27/ HARA utility functions are the class of utility functions that exhibit linear risk tolerance; there are constants a and b, such that $(-V'(x))/V''(x) = a + bx$.

28/ Radner [1981b], p. 1127.

29/ For any positive number epsilon, an epsilon equilibrium is a pair of strategies such that each player's strategy is within epsilon in average per period expected utility of being a best response to the other player's strategy.

30/ Radner [1981b] specifies a sequence of numbers, b_t, that indicate the tolerated discrepancy from the expected outcome; this sequence declines over time, since the average outcome over a longer time horizon is less subject to stochastic influences. In order to inhibit non-cooperation, this sequence must converge to zero in the limit, but the more quickly the sequence approaches zero, the higher the probability of mistakenly punishing cooperative behavior. Radner [1981c] uses sequences of the form: $b_t = k((\log \log t)/t)$, where k is at least $\sqrt{2}$ times the standard deviation of the one-period output.

31/ In Section 3 on risk sharing and Pareto-efficiency, there is some discussion of the use of lump-sum bonuses and punishments in the context of the work of Mirrlees [1975], Lewis [1980] and Holmström [1982c]. In particular, the importance of the movement in the support of the probability distribution of outcomes for the desirability of such contracts is made clear.

32/ Stiglitz and Weiss, [1983], p. 1.

33/ The use of a tournament contract is of interest, but the result itself is not important. The moral hazard problem can be avoided simply by shifting all the risk onto the risk neutral agents.

34/ Nalebuff and Stiglitz [1983], p. 278.

35/ Radner and Stiglitz [1975] show that, assuming that information gathering is a costly activity and that the derivative of the likelihood function is bounded at zero, then it does not pay to gather just a little information; therefore, either no monitoring should be done or it should be carried out in some finite amount.

36/ In Rosen's [1982] model all communication flows are from the top down, with no interactions going the other way around.

37/ Labor supply decisions are ignored, with all personnel assumed to work a fixed number of hours.

38/ It is possible to specify two separate skills in this process, one for the $g(\cdot)$ element and another for the monitoring, time augmenting role. Rosen [1982] adopts a single skill in his analysis for simplicity.

39/ This general result holds when $g(r)$ is increasing in r so long as its elasticity is less than that of $G(s)$.

40/ Mirrlees [1976] points out that this wage scheme amounts to applying a progressive tax to marginal products before paying wages, and thus does some of an egalitarian government's work for it.

41/ Stiglitz [1975b], p. 294.

42/ The optimality of the constraint on portfolio diversification being binding is precisely analogous to the deviation from perfect risk-sharing in Holmström's analysis (Marcus [1982]). Although the portfolio problem induces certain suboptimal decisions (from the equity holder's viewpont), the extra effort elicited by the profit contingent compensation balances the losses due to risk misallocation, as it does in the basic agency model (Holmström [1979]).

43/ In showing that under the optimal contract the portfolio diversification constraint is binding, Marcus [1982] restricted attention to contracts that consisted of salary and stock compensation.

44/ This result holds as a necessary condition and generically (i.e., any perturbation of the problem data takes you out of a situation in which it does not hold) as a sufficient condition if the agent has private, pre-decision information (Holmström [1979]). Harris and Raviv [1979] show that if the agent's private information consists of the true value of the state of nature then the same informativeness condition holds.

45/ The Blackwell theorem states that an information system Y is at least as valuable as Y', given Y, if and only if there are nonnegative numbers $\beta_{y'y}$ such that:

$\pi(y'|z) = \sum_{y \in Y} \beta_{y'y} \pi(y|z) \; \forall \; z \in Z, y' \in Y'$ and that

$\sum_{y' \in Y'} \beta_{y'y} = 1$. The situation is analogous if the agent's action is taken to be the unknown state of nature and the incentive scheme regarded as the decision function.

46/ If infinite penalties are not available, this probability will have a strictly lower bound. An interesting example of less information being as valuable as more, is the case of conditional investigation (Demski and Feltham [1979], Baiman and Demski [1980], Townsend [1979]). Rather than buying the information system Z, the information system Y -- buy Z with probability p, use Y^0 otherwise -- is chosen.

47/ These results depend on the assumption that a^* is unique. If the agent is indifferent between several actions, a marginal randomization may make him jump to an action where risk is less harmful.

48/ The monitoring technology is conditionally independent of the production technology (i.e., the outcome function) if $h(x,y|a) = f(x|a)g(y|a)b(a)$ for some $b(a)$.

49/ The use of a risk-neutral principal appears to strongly influence the results. Risk-sharing considerations may result in a non-convex investigation region.

50/ Ex-ante, one or more individuals may be made worse off by increasing the information.

51/ Baiman [1982], p. 82.

52/ The timing of decision-making is critical. If, in the basic agency structure, the compensation choice and effort decision were made simultaneously, the equilibrium outcome would consist of no effort and zero compensation to the agent. Given any level of effort on the part of the agent, the best response is to give zero compensation. The agent's best response to that would obviously be zero effort (the same equilibrium would result if the analysis were started from the agent's perspective). It is the ability of the principal to commit to a compensation structure that allows for positive levels of effort in equilibrium. In these papers, the principal is not making a zero-sum monetary allocation (as is in the case of the standard agency analysis given the agent's level of effort) but is choosing some action from which the agent may benefit. It is this fact that results in the Nash equilibrium to be interesting here, while it is degenerate in the standard agency framework.

53/   As an example, Demski, et al. [1983] develop a model in which the principal prefers that the agent choose the accounting system despite the fact that the system is to be used to evaluate the agent's performance. Delegating choice of the accounting system to the agent expands the range of possible contracts without eliminating any contractual opportunities previously available and, therefore, cannot make the principal worse off.

54/   In order for the agent to work for the principal the contract must offer him expected utility at least as great as he could get in the labor market. Since the agent will not receive any private information unless he joins the organization, his expected utility is based on his beliefs prior to receiving the private information; the model assumes that the agent cannot leave the organization after learning the private information and before making his choice of action. Baiman [1981] notes that it would be interesting to see if the results would remain valid if the organizaiton had to offer a contract that not only induced the agent to join the organization but to stay even after receiving his private pre-decision information.

55/   The principal may wish to **revise** the revision procedure after observing the first period outcome. Thus, commitment often may not be credible; the principal is free to design incentive schemes, but also to change them over time at his discretion. Freixas, et al. [1982] show that when the principal cannot commit to an incentive scheme, under incomplete information about the productivity of the agent, the optimal bonus is smaller than the shadow price of output.

56/   See the discussion of the revelation principal with an informed principal in Section 2.

57/   This planning problem is taken from Weitzman [1974].

58/   It is not useful to maximize the minimum of the possible gains since this is always zero in the problem they are considering. Instead they make use of the concept of regret rather than payoff. The regret is usually computed by finding the best payoff assuming complete information and subtracting the actual payoff from it. Hurwicz and Shapiro [1978] use division instead of subtraction because of the multiplicative relationships between quantities in their model. In this form, regret can be interpreted as a measure of the efficiency of a policy.

References

Alchian, Armen and Harold Demsetz [1972], "Production, Information Costs and Economic Organization," American Economic Review, Vol. 62, pp. 777-795.

Arrow, Kenneth J. [1971], Essays in the Theory of Risk-Bearing, Chicago: Markham.

Arrow, Kenneth J. [1975], "Vertical Integration and Communication," The Bell Journal of Economics, Vol. 6, pp. 173-183.

Aumann, Robert [1974], "Subjectivity and Correlation in Randomized Strategies," Journal of Mathematical Economics, Vol. 1, pp. 67-96.

Baiman, Stanley [1982], "Agency Research in Managerial Accounting: A Survey," unpublished, University of Pittsburgh.

Baiman, Stanley and Joel Demski [1980], "Economically Optimal Performance Evaluation and Control Systems," Journal of Accounting Research, Vol. 18, pp. 184-220.

Barnard, Chester [1938], The Functions of the Executive, Cambridge, Massachusetts: Harvard University Press.

Blackwell, F. [1951], "Comparisons of Experiments." In J. Neyman, Proceedings of the Second Berkeley Symposium in Mathematical Statistics and Probability, University of California Press.

Blau, Peter and Richard Schoenherr [1971], The Structure of Organizations, New York: Basic Books.

Bonin, John P. [1976], "On the Design of Managerial Incentive Structures in a Decentralized Planning Environment," American Economic Review, Vol. 66, pp. 682-688.

Borch, Karl [1962], "Equilibrium in a Reinsurance Market," Econometrica, Vol. 30, pp. 424-444.

Christensen, John [1979], Communication and Coordination in Agencies, Ph.D. Thesis, Stanford University.

Conn, David [1979], "A Comparison of Alternative Incentive Structures for Centrally Planned Economic Systems," Journal of Comparative Economics, Vol. 3, pp. 261-276.

Crawford, Vincent and Joel Sobel [1982], "Strategic Information Transmission," Econometrica, Vol. 50, pp. 1431-1451.

Cyert, Richard M. and James G. March [1963], A Behavioral Theory of the Firm, Englewood Cliffs, N. J.: Prentice-Hall.

D'Aspremont, Claude and L. Gerard-Varet [1979], "Incentives and Incomplete Information," Journal of Public Economics, Vol. 11, pp. 25-45.

Dasgupta, Partha, Peter Hammond and Eric Maskin [1979], "The Implementation of Social Choice Rules," Review of Economic Studies, Vol. 46, pp. 185-216.

Demski, Joel and George Feltham [1976], Cost Determination: A Conceptual Approach," Ames: Iowa State University Press.

Demski, Joel and George Feltham [1978], "Economic Incentives in Budgetary Control Systems," The Accounting Review, Vol. 563, pp. 336-359.

Demski, Joel, James Patell and Mark Wolfson [1983], "Decentralized Choice of Monitoring Systems," Research Paper 674, Graduate School of Business, Stanford University.

Eisenhardt, Kathleen [1983], "Theoretical and Empirical Comparison of Organizational and Economic Approaches to Control," unpublished, Stanford University.

Ekern, Steinar [1979], "The New Soviet Incentive Model: Comment," The Bell Journal of Economics, Vol. 10, pp. 720-725.

Fan, L. S. [1975], "On the Reward System," American Economic Review, Vol. 65, pp. 226-229.

Freixas, Xavier, Roger Guesnerie and Jean Tirole [1982], "Planning Under Incomplete Information and the Ratchet Effect," unpublished.

Friedman, Milton [1953], Essays in Positive Economics, University of Chicago Press.

Galbraith, Jay [1973], Designing Complex Organizations, Reading, Mass.: Addison-Wesley.

Gjesdal, Froystein [1978], Stewardship Accounting: Controlling Informational Externalities, Ph.D. Thesis, Stanford University.

Gjesdal, Froystein [1982], "Information and Incentives: The Agency Information Problem," Review of Economic Studies, Vol. 49, pp. 373-390.

Green, Jerry and Jean-Jacques Laffont [1982], "Limited Communication and Incentive-Compatibility," Working Paper 919, Harvard Institute of Economic Research.

Green, Jerry and Nancy Stokey [1981a], "A Two-Person Game of Information Transmission," Working Paper 751, Harvard Institute of Economic Research.

Green, Jerry and Nancy Stokey [1981b], "The Value of Information in the Delegation Problem," Working Paper 776, Harvard Institute of Economic Research.

Green, Jerry and Nancy Stokey [1982], "A Comparison of Tournaments and Contracts," Working Paper 840, National Bureau of Economic Research.

Grossman, Sanford and Oliver Hart [1983], "An Analysis of the Principal-Agent Problem," Econometrica, Vol. 51, pp. 7-45.

Groves, Theodore [1973], "Incentives in Teams," Econometrica, Vol. 41, pp. 617-31.

Guasch, J. Luis and Andrew Weiss [1980], "Wages as Sorting Mechanisms in Competitive Markets with Asymmetric Information: A Theory of Testing," Review of Economic Studies, Vol. 47, pp. 653-664.

Harris, Milton and Arthur Raviv [1978], "Some Results on Incentive Contracts with Applications to Education and Employment, Health Insurance, and Law Enforcement," American Economic Review, Vol. 68, pp. 20-30.

Harris, Milton and Arthur Raviv [1979], "Optimal Incentive Contracts with Imperfect Information," Journal of Economic Theory, Vol. 20, pp. 231-255.

Harris, Milton and Robert Townsend [1981], "Resource Allocation under Asymmetric Information," Econometrica, Vol. 49, pp. 33-64.

Holmström, Bengt [1977], On Incentives and Control in Organizations, Ph.D. Thesis, Stanford University.

Holmström, Bengt [1979], "Moral Hazard and Observability," The Bell Journal of Economics, Vol. 10, pp. 74-91.

Holmström, Bengt [1980], "On the Theory of Delegation," Working Paper 438, Northwestern University.

Holmström, Bengt [1982a], "Design of Incentive Schemes and the New Soviet Incentive Model," European Economic Review, Vol. 17, pp. 127-148.

Holmström, Bengt [1982b], "Managerial Incentive Problems--A Dynamic Perspective." In Essays in Economics and Management in Honor of Lars Wahlbeck, Helsinki: Swedish School of Economics.

Holmström, Bengt [1982c], "Moral Hazard in Teams," _Bell Journal of Economics_, Volume 13, pp. 324-340.

Holmström, Bengt and Roger Myerson [1981], "Efficient and Durable Decision Rules with Incomplete Information," Working Paper 495, Northwestern University.

Hurwicz, Leonid [1972], "On Informationally Decentralized Systems." In Roy Radner and B. McGuire, _Decision and Organization_, Amsterdam: North-Holland Press.

Hurwicz, Leonid and Leonard Shapiro [1977], "Incentive Structures Maximizing Residual Gain under Incomplete Information," Research Paper 77-83, University of Minnesota, Center for Economic Research.

Hurwicz, Leonid and Leonard Shapiro [1978], "Incentive Structures and Maximizing Residual Gain under Incomplete Information," _Bell Journal of Economics_, Vol. 9, pp. 180-191.

Jensen, Michael and William Meckling [1976], "Theory of the Firm: Managerial Behavior, Agency Costs, and Ownership Structure," _Journal of Financial Economics_, Vol. 3, pp. 305-360.

Lambert, Richard [1981], _Managerial Incentives in Multiperiod Agency Relationships_, Ph.d. Thesis, Stanford University.

Lazear, Edward and Sherwin Rosen [1981], "Rank-Order Tournaments as Optimum Labor Contracts," _Journal of Political Economy_, Vol. 89, pp. 1045-1064.

Lewis, Tracy R. [1980], "Bonus and Penalties in Incentive Contracting," _Bell Journal of Economics_, Vol. 11, pp. 292-301.

Lippman, Steven A. and John J. McCall [1976], "The Economics of Job Search: A Survey," _Economic Inquiry_, Vol. 14, pp. 155-189, 347-368.

Locke, Edwin [1968], "Toward a Theory of Task Motivation and Incentive," _Organizational Behavior and Human Performance_, Vol. 3, pp. 157-189.

March, James G. [1978], "Bounded Rationality, Ambiguity, and the Engineering of Choice," _Bell Journal of Economics_, Vol. 9, pp. 578-608.

March, James G. and Herbert A. Simon [1958], _Organizations_, New York: Wiley.

Marcus, Alan [1982], "Risk-Sharing and the Theory of the Firm," _Bell Journal of Economics_, Vol. 13, pp. 369-378.

Marschak, Jacob and Roy Radner [1972], The Economic Theory of Teams, New Haven and London.

Maslow, Abraham [1943], "A Theory of Human Motivation," Psychological Review, Vol. 50, pp. 370-396.

Merton, Robert C. [1973], "Theory of Rational Option Pricing," Bell Journal of Economics and Management Science, Vol. 4, pp. 141-183.

Merton, Robert K. [1957], "The Role-Set: Problems in Sociological Theory," British Journal of Sociology, Vol. 8, pp. 106-120.

Milgrom, Paul R. [1981], "Good News and Bad News; Representation Theorems and Applications," The Bell Journal of Economics, Vol. 12, pp. 380-391.

Mirrlees, James [1974], "Notes on Welfare Economics, Information and Uncertainty." In M. Balsch, D. McFadden and S. Wu, Essays on Economic Behavior under Uncertainty, Amsterdam: North-Holland.

Mirrlees, James [1975]. "On Moral Hazard and the Theory of Unobservable Behavior," Oxford University.

Mirrlees, James [1976], "The Optimal Structure of Incentives and Authority within an Organization," The Bell Journal of Economics, Vol. 7, pp. 105-131.

Mookherjee, Dilip [1981], "Optimal Incentive Contracts in Multiagent Situations," unpublished, London School of Economics.

Myerson, Roger [1979], "Incentive Compatibility and the Bargaining Problem," Econometrica, Vol. 47, pp. 61-73.

Myerson, Roger [1980], "Optimal Coordination Mechanisms in Principal-Agent Problems," Working Paper 429, Northwestern University.

Myerson, Roger [1981], "Mechanism Design by an Informed Principal," Working Paper 481, Northwestern University.

Nalebuff, Barry and Joseph Siglitz [1983], "Information, Competition, and Markets," American Economic Review, Vol. 73, pp. 278-283.

Pfeffer, Jeffrey [1982], "Management as Symbolic Action: The Creation and Maintenance of Organizational Paradigms." In Larry Cummings and Barry Staw, Research in Organizational Behavior, Vol. 3, JAI Press.

Pratt, John W. [1964], "Risk Aversion in the Small and in the Large," Econometrica, Vol. 32, pp. 122-136.

Radner, Roy [1979], "Can Bounded Rationality Resolve the Prisoner's Dilemma?," Working Paper, Bell Laboratories Discussion Paper.

Radner, Roy [1981a], "A Repeated Principal-Agent Game with Discounting," Working Paper, Bell Laboratories Discussion Paper.

Radner, Roy [1981b], "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," Econometrica, Vol. 49, pp. 1127-1148.

Radner, Roy [1981c], "Optimal Equilibria in a Class of Repeated Games with Imperfect Monitoring," Research Paper, Bell Laboratories.

Radner, Roy and Joseph Stiglitz [1976], " A Nonconcavity in the Value of Information," unpublished, Stanford University.

Rogerson, William [1982], "The Structure of Wage Contracts in Repeated Agency Models," Research Paper, Institute of Mathematical Studies in the Social Sciences, Stanford University.

Rosen, Sherwin [1982], "Authority, Control, and the Distribution of Earnings," Bell Journal of Economics, Vol. 13, pp. 311-323.

Rosenbaum, James [1978], "Tournament Mobility: Career Patterns in a Corporation," Administrative Science Quarterly, Vol. 24, pp. 220-241.

Ross, Stephen [1973], "The Economic Theory of Agency: The Principal's Problem," American Economic Review, Vol. 63, pp. 134-139.

Ross, Stephen [1974], "The Economic Theory of Agency and the Principal of Similarity." In M. Balch et al., Essays on Economic Behavior Under Uncertainty, North-Holland.

Rubinstein, Ariel [1981], "An Optimal Conviction Policy for Offenses that May Have Been Committed by Accident." In Brahms, Schotter and Schwodiauer, Applied Game Theory, Wien. Physica Verlag.

Rubinstein, Ariel and M. Yaari [1980], "Repeated Insurance Contracts and Moral Hazard," Working Paper, Hebrew University.

Salop, Joanne and Steven Salop [1976], "Self-Selection and Turnover in the Labor Market," Quarterly Journal of Economics, Vol. 90, pp. 619-627.

Scott, Richard, Sanford Dornbusch, Bruce Busching, and James Laing [1967], "Organizational Evaluation and Authority," Administrative Science Quarterly, Vol. 12, pp. 93-117.

Selten, Reinhart [1975], "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory, Vol. 4, pp. 25-55.

Shapiro, Carl and Joseph Stiglitz [1982], "Equilibrium Unemployment as a Worker Discipline Device," Working Paper, Princeton University.

Shavell, Steven [1979], "Risk Sharing and Incentives in the Principal and Agent Relationship," The Bell Journal of Economics, Vol. 10, pp. 55-73.

Simon, Herbert [1951], "A Formal Theory of the Employment Relationship," Econometrica, Vol. 19, pp. 293-305.

Simon, Herbert [1955], "A Behavioral Model of Rational Choice," Quarterly Journal of Economics, Vol. 69, pp. 99-118.

Simon, Herbert [1959], "Theories of Decision-Making in Economic and Behavioral Science," American Economic Review, Vol. 49, pp. 253-283.

Snowberger, Vinson [1977], "The New Soviet Incentive Model: Comment," Bell Journal of Economics, Vol. 8, pp. 591-600.

Spence, A. Michael [1974], Market Signaling, Harvard University Press.

Spence, A. Michael [1975], "The Economics of Internal Organization: An Introduction," Bell Journal of Economics, Vol. 6, pp. 163-172.

Spence, A. Michael and Richard Zeckhauser [1971], "Insurance, Information and Individual Action," The American Economic Review, Vol. 61, pp. 380-387.

Stiglitz, Joseph [1974], "Risk Sharing and Incentives in Sharecropping," Review of Economic Studies, Vol. 61, pp. 219-256.

Stiglitz, Joseph [1975a], "Incentives, Risk and Information: Notes toward a Theory of Hierarchy," The Bell Journal of Economics, Vol. 6, pp. 552-579.

Stiglitz, Joseph [1975b], "The Theory of 'Screening,' Education, and the Distribution of Income," American Economic Review, Vol. 65, pp. 283-300.

Stiglitz, Joseph and Andrew Weiss [1983], "Incentive Effects of Terminations: Applications to the Credit and Labor Markets," Working Paper, Bell Laboratories.

Thompson, James [1969], Organizations in Action, McGraw-Hill: New York.

Thompson, William [1979], "Eliciting Production Possibilities from a Well-Informed Manager," Journal of Economic Theory, Vol. 20, pp. 360-380.

Townsend, Robert [1979], "Optimal Incentive Contracts and Competitive Markets with Costly State Verification," Journal of Economic Theory, Vol. 21, pp. 265-293.

Vroom, Victor [1964], Work and Motivation, New York: John Wiley.

Weitzman, Martin [1974], "Prices vs. Quantities," Review of Economic Studies, Vol. 41, pp. 477-491.

Weitzman, Martin [1976], "The New Soviet Incentive Model," The Bell Journal of Economics, Vol. 7, pp. 251-257.

Weitzman, Martin [1980], "The 'Ratchet Principle' and Performance Incentives," The Bell Journal of Economics, Vol. 11, pp. 302-308.

Williamson, Oliver [1975], Markets and Hierarchies: Analysis and Antitrust Implications, New York: The Free Press.

Wilson, Robert [1968], "On the Theory of Syndicates," Econometrica, Vol. 36, pp. 119-132.

Wilson, Robert [1979], "Incentive Compatible Risk Sharing," Research Paper, Institute of Mathematical Studies in the Social Sciences, Stanford University.

List of Reports