

A list of Working Papers on the last pages

No. 214, 1988

AGENCY COSTS AND INNOVATION

by

Bengt Holmström

Paper prepared for the IUI Conference on
Markets for Innovation, Ownership and Control,
June 12–16, 1988, at Grand Hotel, Saltsjöbaden

September, 1989

Agency Costs and Innovation.

by Bengt Holmstrom*

Yale University

June 1988

June 1989, present version.

* This paper was written for a conference on "The Market for Innovation, Ownership and Control" held at Saltsjobaden, Stockholm, June 13-15, 1988 and organized by the Industrial Institute for Economic and Social Research, Stockholm. The research was partially funded by grants from NSF and the Sloan Foundation. I'm grateful to David Kreps for many stimulating discussions related to the subject.

AGENCY COSTS AND INNOVATION.

by

Bengt Holmstrom

Yale University

ABSTRACT

Stylized facts indicate that small firms are responsible for a disproportionate share of innovative research. There are many possible explanations for this fact. The paper seeks to understand this phenomena as the outcome of an optimal assignment of tasks across individuals and organizations. It is shown that incentive costs associated with a given task depend on the total portfolio of tasks that an individual or an organization undertakes. Mixing hard to measure activities (innovation) with easy to measure activities (routine) is particularly costly, since it will either lead to misallocation of attention across tasks or to misallocation of risk. Larger firms are at a comparative disadvantage in conducting highly innovative research, because of the costs associated with managing a heterogenous set of tasks. It is further argued that optimal organizational responses to coordination and control of routine tasks will lead to bureaucratization within the firm and to financial constraints imposed by capital markets, both of which are hostile to innovation.

Agency Costs and Innovation.

by Bengt Holmstrom.

June 1988 (revised June 1989).

1. Introduction.

For an increasing number of firms, innovation has become an urgent concern. With the lifetimes of products shortened due to an accelerating pace of technological change, the fight against obsolescence is raising new and unprecedented challenges everywhere. Business consultants, ever sensitive to the needs of the hands that feed them, have jumped into the act with a new gospel tailored to the management of innovation: Bureaucratic procedures are to be replaced by a culture that encourages action, allows freedom to experiment and exhibits substantial tolerance for errors; formal planning models, extensive information collection and centralized decision-making are all to be significantly curtailed.

Concerns about the innovativeness of U.S. firms have surfaced on the political front as well. The debate on a national industrial policy has touched on the need for a government agenda to encourage research and development. Such government intervention is based on a distrust for the ability of markets to steer capital to its best use, and understandably has met with skepticism. But whatever one's views, it is clear that the questions are important and topical. At stake is the pre-eminence of the U.S. as the industrial leader of the world.

My intention is not to join the debate on the need for increased innovation either on the national or the firm level, though in a larger historical perspective the attention seems appropriate; the West grew rich

primarily because it was willing to experiment (for a particularly compelling account, see Rosenberg and Birdzell, 1985). Rather I want to discuss the most efficient forms for organizing innovative activity in the private sector. If innovation is to be encouraged, which private institutions are best positioned to undertake it?

Stylized facts seem to indicate that small firms have been responsible for a disproportionate share of significant innovations in the past (see for instance Scherer, 1984, Chapter 11). I'm aware that the validity of this claim can be debated and that the conclusions are sensitive to what one counts as innovation. Also, a higher success rate does not prove a comparative advantage; small firms could be innovating too much.¹ Yet, the casual evidence suggests the hypothesis that large firms are at a comparative disadvantage in managing truly innovative research. My specific purpose is to study whether current theories of organization, particularly those based on transaction cost and incentive considerations, can lend support for such a claim.

I will discuss two types of reasons why large firms might (rationally) innovate less. One has to do with the internal organization of the firm and the other with the firm's relationship to the capital market. On the internal side the main theme is that the large corporation has emerged primarily to serve production and marketing goals and that in pursuing those objectives effectively it has to organize in a way that compromises innovation incentives. Providing incentives for both types of activities within one organization is more costly than providing them through separate organizations. Ultimately the reasons for this can be traced back to the loss

¹Small firms may be forced to innovate in order to distinguish themselves or be able to compete with large firms.

in reliable performance measures that attend integration. Weaker performance measures lead the firm to take other steps to insure proper behavior. In general terms they involve affecting the opportunity costs of the employees. More rigid rules and less discretion are primary consequences. Such bureaucratization is hostile to innovation both directly by restricting experimentation and indirectly by screening out innovative personalities.

On the external side I will argue that a concern for reputation in the capital market will lead a large firm to act more cautiously in taking risks. Past performance is an important signal for future potential and will determine the terms under which new capital will be made available. This has a tendency to make the firm myopic in its behavior. For a small firm which has less flexibility in choosing its activities, this problem is less severe. The extreme case is a start-up firm, singularly devoted to the development of a new product.

In presenting the details of these arguments I will draw unrestrainedly from the growing literature on transaction cost and incentive contracting; in particular, the recent efforts to identify the economic purpose of the firm.² I will spend a fair amount of time interpreting this literature before applying the insights to the innovation issues at hand. I hope some of the discussion will be of independent interest to institutional economics in general.

The paper is organized as follows. Section 2 identifies some of the incentive dilemmas. Section 3 makes preliminary observations through the lens of a simple principal agent model. Sections 4 through 6, the body of the paper, develop the details of the internal organization problems that

²My main sources include Williamson (1975, 1985), Grossman and Hart (1986), Milgrom and Roberts (1987), Milgrom (1988), Holmstrom and Tirole (1988) and Holmstrom and Milgrom (1989).

handicap innovation in large firms. Section 7 discusses the pressures from capital markets. A summary and conclusion are offered in section 8.

2. The Incentive dimensions of innovation.

In the simplest abstraction, innovation decisions are just investment decisions and as such part of the standard problem of how to allocate capital. One would expect modern finance theory to give good general advice on how to manage investments into research and development. But a quick look at finance textbooks reveals answers that are based on a very stylized conception of the problem and rather less illuminating than one would hope. Mainly, the student is told that the decision to invest should be based on a straightforward net present value calculation, in which the expected future return stream from the contemplated project is discounted using a cost of capital that reflects the appropriate social risk inherent in that stream. The most striking part of the advice is that projects should be evaluated this way without regard to individual portfolios or capital constraints.

Practitioners do not follow these simple rules. Firms do care about idiosyncratic risk and they are exceedingly conscious of capital constraints. In any given year there is a limited amount of money to be allocated among proposed projects and typically the demand for funding vastly exceeds the amount set aside for investment in that year. This imbalance is not resolved in a decentralized fashion by creating an internal cost of capital schedule. Rather, the dilemma is dealt with by various idiosyncratic rules that all turn on more or less intense centralized scrutiny of the proposals put forward. For financial indicators firms rely on internal rates of return (with hurdle rates substantially above the cost of capital implied by market

rates) as well as payback criteria more commonly than on net present values.

Likewise, capital markets do not allocate funds to firms in the stylized fashion prescribed by net present value theory. Rationing of capital is commonplace and the cost of funds is very much perceived to depend on the source from which they are obtained.

By now it is well recognized where the discrepancies between standard theory and practice lie. Reallocating shares in the stock market is a very different type of activity than supplying fresh funds to a company. Marginal transfers of shares have little or no impact on the operation of a firm, while new funds change the firm's production set. The source of funds, the amount of capital and the terms on which these funds are made available, all influence the operation of the firm and the behavior and prospects of its members. Potential incentive problems are easy to envision, which suggests that the variety of institutions mediating capital from markets to firms (as well as capital within firms) is best understood in terms of their effectiveness in monitoring and managing the incentives of the people involved in using that capital.³ As well, this is the appropriate perspective for discussing comparative advantage in innovation, since the process of innovation is mainly one of matching financial capital with the human capital behind the ideas.

The use of incentive contracts to reduce the costs of transacting under asymmetric information have been studied extensively. The standard framework considers a principal and typically one but sometimes many agents. Normally, the party with superior information is the agent, while the principal is the

³Those who think the problem is in management practice rather than the theory, should be reminded of the speed with which financial markets have adopted modern asset pricing theories. Understanding net present value analysis is trivial compared with learning option pricing and valuation of other derivative securities.

capital owner. In talking about investment incentives, the first question to ask is why there should be a problem. Why would the agent not want to act in the best interests of the principal? Suppose that the agent does nothing but give investment advice. Suppose the principal pays a flat fee for the advice. Then the agent would have no reason not to tell what he knows. This is not as uncommon a case as one might think. Evaluators of projects, accountants who are to judge the veracity of financial statements, for instance, are paid on a non-contingent basis. Contingent fees would do little but create distorting incentives to report honestly what's observed.

More than evaluation expertise is needed to make a straight salary contract undesirable. Three ways of introducing preference incongruities are commonly considered. The first one recognizes that investments require efforts by the agent that cannot be compensated directly, because of problems with observability. To motivate private expenditures, contingent fees based on what's observable, for instance the output of the project, will be necessary. Such incentive schemes introduce risk preferences for the agent, assuming that the agent is risk averse or does not have enough financial resources to buy out the principal.

A second possibility is that the agent owns part of the project, say the idea, and is shopping around for an equity partner. Since the agent knows the value of the project better than the potential partner, there is a problem in deciding on the right price. A contingent fee schedule is a means by which ex ante asymmetries in information can be reduced.

Finally, a third case recognizes that the agent may have a direct interest in the project, contingent fees notwithstanding. One plausible reason is that the agent's market value will depend on undertaking the project as well as on its outcome. Thus, investments commonly yield finan-

cial returns as well as human capital returns. Some kind of contract will be needed to align incentives more closely.

All three incentive dimensions seem relevant for managing innovation. Indeed, the agency costs associated with innovation are likely to be high, because innovation projects are: (a) risky -- there is a high probability of failure, but also prospects for extraordinary returns; (b) unpredictable -- many future contingencies are impossible to foresee; (c) long-term and multi-stage -- the project has an invention, a development and a completion stage, and can be terminated between those; (d) labor intensive -- all stages require substantial human effort; (e) idiosyncratic -- not easily comparable to other projects. It turns out that contracting under this set of circumstances is particularly demanding.

2. Some insights from agency theory.

As will become clearer in the next section, standard principal-agent models do not get to the core of the institutional choice question, because their results do not depend on the organizational location of the agency relationship. Yet, the models can shed light on a number of organizational issues. In this section I will work with a particularly simple principal-agent model to identify some trade-offs that I think are central to the innovation questions I have set to discuss. The model is described in full in Holmstrom and Milgrom (1987); my discussion relies in part on extensions contained in Holmstrom and Milgrom (1989).

Basic model. Consider a single innovation project, which yields an uncertain payoff x . As well, the payoff depends on what the agent does, but not so that it fully reveals the agent's role. Let the specific relationship

be:

$$(1) \quad x = e + \text{normal error term},$$

where e represents the agent's action (effort).

Assume that the principal and the agent have an identical understanding of the stochastic project returns. Both believe the normal error term has mean μ and variance σ^2 . Thus, the agent has no superior information about project returns before acting. This may be a reasonable assumption if we are at the initial stages of a research undertaking. Note that with this return specification, the agent is in effect choosing the mean of a normal distribution. If he acts e , the mean is $e + \mu$. The variance is not within the agent's control.

The principal is unable to observe directly what the agent does; that's what gives rise to the incentive problem. Let x be verifiable in the sense that enforceable contracts can be written on it. A contract specifies payments $s(x)$ to the agent when x occurs, leaving the principal with the residual $x - s(x)$.

The principal is risk neutral and the agent is risk averse with preferences described by the exponential utility function $-\exp\{-r(s(x)-c(e))\}$. Here r is the coefficient of absolute risk aversion and $c(e)$ the cost function. The natural interpretation of cost is in terms of opportunity loss. Working on the project limits the income that can be generated from other sources.

The problem is to choose s so that it encourages adequate effort, without overly burdening the agent with risk. Under suitable assumptions, the second-best contract (ie the best contract given the informational

restrictions of the model) takes the linear form: $s(x) = \alpha x + \beta$.⁴ The best choice of α is derived by maximizing the certain equivalent of joint surplus, which, given the linear rule, is $\mu + e - (1/2)r\alpha^2\sigma^2 - c(e)$, subject to the agent's first order condition $c'(e) = \alpha$. To give an explicit example of a solution, suppose $c(e) = ke^2/2$. Then, the best piece rate is:

$$(2) \quad \alpha = (1 + kr\sigma^2)^{-1},$$

which also is the agent's choice of effort. The salary component β will be set so that the agent is willing to participate in the project; it is a mere transfer of wealth between the principal and the agent and of little interest here.

Formula (2) precisely accords with one's intuition: the agent's share is higher, the lower is his aversion to risk (lower r), the lower is the risk (the variance) and the lower is the cost of action. The solution reflects a trade-off between risk sharing and incentives to supply effort. For optimal effort one should choose $\alpha = 1$ and for optimal risk sharing $\alpha = 0$. Only if risk is absent or the agent is risk neutral, can one avoid the agency costs associated with limited observability. The joint surplus (under the best contract) is in fact: $\mu + 1/2k^{-1}(1 + rk\sigma^2)^{-1}$, which is less than the first-best surplus: $\mu + 1/2k^{-1}$. Note that welfare varies positively with α .

Monitoring. This basic model can be enriched in various ways. One is to introduce a monitoring variable y . By a monitor I mean any signal of the agent's action other than the outcome. It could be the principal looking over the agent's activities or information obtained from observing agents in

⁴I'm simplifying the description of the model considerably. The actual model views the agent as choosing the drift rate of a stochastic process over time. I'm describing the reduced form.

related activities. If $y = e +$ a normal error term, as in (1), then the optimal schedule is linear in the two variables: $s(x,y) = \alpha x + \gamma y + \beta$. One finds that the higher is the risk of a project, the more intensive additional monitoring will be. Higher project risk forces a reduction in the coefficient α , that is in the agent's outcome share. This reduces effort. Consequently, the incentive coefficient γ on the monitoring variable will be raised, since the marginal cost of effort has gone down (convex cost function). But an increase in the monitoring coefficient will increase the risk stemming from errors in monitoring and it becomes valuable to invest more in monitoring to reduce that error.

An executive, trying to encourage more innovation, recently observed: "I try to give people a feeling that it's okay to fail, that it's important to fail". Indeed, incentives for innovation must provide for more tolerance, since innovation is intrinsically risky and progress more erratic than with standard investments. But the consequence is equally important to recognize: Direct, close-handed monitoring of the agent's activities must be introduced to compensate for the weaker output rewards. (However, as I will discuss later, monitoring information may often not be possible to use as effectively as envisioned here.)

Project choice. How do agency costs affect the choice of projects? And how do projects get assigned across agents? These are the most interesting questions in the context of innovation. Using my earlier example, a project is identified with the characteristics (μ, σ^2, k) . The best (single) project to choose maximizes the net welfare: $\mu + 1/2k^{-1}(1 + rk\sigma^2)^{-1}$. One notes that the best project is not determined by standard net present value rules. A

concern for risk is present even though the principal is risk neutral.⁵ The reason is that the worker is carrying (by design) some undiversified risk.

The implication is that, to a degree, advantages in technology (high μ) will be traded off against incentive considerations. A more uncertain technology is more costly from an incentive point of view and might be passed up in favor of more routine ones despite their lower returns. By the earlier monitoring logic, there would also be value in choosing projects that are correlated with each other. Overlapping or competing projects could make sense, since they would reduce the incentive costs even if duplication might otherwise be technologically wasteful. This would be an argument for carrying out many projects within the same firm. However, in the case of innovations, since their risks are rather unique, one would not expect significant gains from such integration.

Considering the cost function $c(e)$, it is interesting to observe that a low cost technology may be worse than a high one. Even if $c_1(e) < c_2(e)$ for every e , the latter cost function may be preferred. As an example, suppose $c_1(e) = e^2/2$ and the optimal incentive contract leads the agent to choose $e_2 = 1/2$ (this is optimal by appropriate choice of risk aversion and variance). Compare it to the following cost function: $c_2(e) = 1/8$ whenever $e < 1/2$ and infinity beyond that. Since $c_1(1/2) = 1/8$, the first cost function is uniformly lower than the second. Yet, the second cost function will be preferred, since with that one we can get the agent to choose $1/2$ without imposing any risk.

An important general point hides behind this trivial example. Sometimes

⁵Risk neutrality is a rather natural assumption here, if one interprets risk as being idiosyncratic. Formally, the assumption can be rationalized in a model where both systematic and idiosyncratic risk are present. It is easy to show that the agent's contract would filter out the systematic risk and effectively be based only on the idiosyncratic component.

it is more effective to provide incentives by changing the agent's opportunity cost than by offering financial rewards; and quite generally this is a valuable additional incentive instrument. How does one change the opportunity cost? By controlling the agent's options to spend time and effort on alternative activities. The situation above could be interpreted in these terms. The first cost function applies when the agent can divide his time between working for the principal and working on a private project (eg. watching TV at home). It corresponds to the opportunity cost of having less time to devote to the private project when more time is spent on the principal's. Under cost function c_2 the private option is removed. The agent's cost goes up uniformly, because the benefit from spending the balance of his time at home is no longer available. However, with the option removed, no further incentives for effort need be provided. Forcing the agent to come into the office is an example in point; indeed, fixed salaries are more prevalent in office jobs than in jobs performed at home.

This discussion shows that the cost of providing incentives for a given task importantly depends on what other activities the agent is allowed to engage in, be that in private or within the firm. Loosely speaking, the more flexibility the agent has, the costlier it is to induce him to work on a given project. As a consequence, it may be optimal to reduce the agent's flexibility by eliminating marginal tasks, that is, tasks which do not contribute enough in net receipts to off-set increased costs of providing incentives for more important tasks (for more on this, see Holmstrom and Milgrom, 1989).

A significant factor in determining how much flexibility to allow the agent, is how accurately one can measure the agent's performance in his major tasks. Let me illustrate this point by extending the earlier example a bit

further. Suppose the agent could allocate some effort (e') to an outside activity with non-stochastic return $f(e')$. There is no cost to effort, but the total amount of effort, $e + e'$, cannot exceed 1; (ie. $c(e + e') = 0$ for $e + e' \leq 1$ and infinity for $e + e' > 1$). Assume $f(e') = \lambda_1 e'$, for $e' \leq n < 1$ and $\lambda_1 n + \lambda_2(e' - n)$ for $e' > n$, where $\lambda_1 > 1 > \lambda_2 > 0$; in other words, f is piecewise linear with a kink at $e' = n$. Note that ideally, the agent should devote time n to the outside and the rest to the principal's project.

However, if the principal offers an incentive $\alpha < \lambda_2$ the agent will work only on the outside option. Thus, α must be at least λ_2 if the second option is around and it is desirable to have the agent work for the principal at all. The latter can be assured by choosing parameters so that $f(1) < 1$.

When there is little noise in output x (σ^2 small), the best alternative is to choose α just larger than λ_2 . The cost of imposing risk on the agent is $(r/2)(\lambda_2)^2 \sigma^2$, which is lower than the opportunity cost of forgoing the outside return: $(\lambda_1 - \lambda_2)n$. However, as σ^2 increases, it becomes increasingly costly to keep α above λ_2 . When $\sigma^2 > (\lambda_1 - \lambda_2)n((r/2)(\lambda_2)^2)^{-1}$, it is better to exclude the outside activity altogether, allowing the principal to set $\alpha = 0$.

The logic here can be generalized to several outside activities.⁶ As risk increases, more outside activities will be eliminated. The up-shot is that, the agent's flexibility will be more restricted the poorer the performance measures for the main tasks. In more familiar terms: responsibility and authority must be in a balanced relationship. This control principle will play a central role both in the argument for corporate bureaucratization and the capital market discussion that follow. In passing I note that there

⁶In Holmstrom and Milgrom (1989) a model with a continuum of outside projects is used.

is recent empirical evidence on the relevance of this prediction. Anderson and Schmittlein (1984) found that in the electronics industry, the use of in-house salespeople (rather than independent agents) correlates significantly and positively with the uncertainty of the environment in which they operate.

Project assignment. The second main principle I want to bring up is that of uniformity in tasks. Suppose there are several projects to be allocated between two identical agents. Assume only total cost of effort matters to the agents. The projects differ in their risk characteristics, but are similar in their expected returns. Let $x_i = f(e_i) + \text{normal error term}$, be the return of project i , where f is concave and error terms are independent. (The linearity result extends easily to this case.) Assume the principal only observes the aggregate output of each agent. Then the following holds: the best allocation of projects is such that the projects assigned to one agent are uniformly more risky than the projects assigned to the other.⁷

The proof is trivial. One can always switch around two projects between the agents without affecting output (since projects have identical return functions f). Thus, only risk considerations matter. Risk is minimized by assigning the agent with the lower incentive coefficient all projects with a variance above some cut-off level (determined endogenously by the cost function) and assigning all low risk projects to the agent with the higher incentive coefficient. (If it is optimal to give both agents the same incentive coefficient, then the assignment does not matter, but of course this is rarely the case.)

The intuition can be phrased as follows. If an agent is given both high variance projects and low variance projects (relative to the other agent),

⁷ Minahan (1988) has independently arrived at this same result.

the presence of high variance projects in the portfolio will force the incentive coefficient down, foregoing the opportunity to offer stronger incentives for the low variance projects. On the other hand, if projects are split according to risk, one of the two agents can be offered strong incentives, while only the other one will have to operate under weak incentives because of risk. (Since projects are independent and utility exponential, diversification issues do not arise.)⁸

Thus, the agency costs are lower if the projects offered to the agents are uniform rather than diverse. The insight here can be looked at another way: when agents choose between homogenous projects, incentive problems associated with the allocation of effort are eliminated - only overall effort remains an issue. The implication is that innovation activities may mix poorly with relatively routine activities in an organization. (The point remains even if the organization is hierarchical; at some level attention must be allocated between both kinds of projects.)

The idea that incentives for different activities need to be in balance is of course relevant more generally. Thus, where cooperation between agents is desirable, individual incentives must be reduced (see Lazear, 1989). For this reason it may again pay to separate tasks which require cooperation from tasks in which strong individual incentives are invaluable. Cooperation and competition do not coexist comfortably, at least within a narrow group.

Before closing this section, a qualification is in order. The model I have been using to guide my intuition is one in which effort incentives are paramount and no asymmetries in information are present, other than those

⁸With appropriate modifications, the result extends to cases in which individual project returns are observed, so that each project has its own incentive coefficient α_i . Also, one can allow individual return functions f_i and agents with different risk aversion and cost functions. See Holmstrom and Milgrom (1989).

regarding effort. The case in which the agent is better informed about project returns at the time of contracting may lead to somewhat different conclusions. (Information that arrives after a contract is signed and a project is selected does not change anything substantively). Adverse selection, as this case is referred to, also raises a new issue: How to elicit information from the agent so that relatively efficient decisions (eg project choices) can be made? I will return to discuss this later. Let me just make one observation for future reference: in an effort to restrict the information rents of the informed, allocational distortions will arise. For instance, an entrepreneur shopping for an equity partner, will receive too little funds and will end up bearing too much risk. More generally, when bargaining takes place under conditions of adverse selection, some of the surplus is dissipated, either because agreements will not be reached when they should or because there will be costly delays.

3. Incomplete contracts and institutional choice.

That incomplete contracting plays a central role in understanding institutional choice has been stressed by Williamson in his long line of research on organizations (Williamson, 1975, 1985). Were it the case that parties could sign comprehensive contracts, that is contracts which fully specify each side's responsibilities in all future contingencies, the organizational context of the contract would not matter, really.⁹ Grossman and Hart (1986) have sharpened the argument by introducing the important

⁹Traditional contract models are typically comprehensive. However, I note that the model in the previous section can be interpreted as a reduced form incomplete contract model if there are return streams that cannot be split between parties (eg $\alpha_i = 0$ or 1).

distinction between residual decision rights, implied by institutional choice, and specific rights conferred by explicit contract. In this terminology, the ownership of a firm is identified with the residual rights to the firm's assets: its tangible assets (machines, buildings, cash, etc.) as well as its intangible ones (patents, brand names, reputation, etc.). This is a significant extension of the traditional view of ownership as an entitlement to the firm's residual income stream.

Residual rights become important when one encounters situations not covered by specific contracts. If a disagreement arises on how to resolve the matter, parties will have to negotiate a solution (with the courts offering final arbitration if necessary). What each side will get will depend on its bargaining position, which in turn is a function of its residual rights. Thus residual rights and institutional choice play a central role in imputing returns to each side.¹⁰ Ownership of assets is an indirect way of choosing an incentive scheme for the transactors. The key point is that such variation in incentives could not be had by specific contract alone, since contracting is necessarily incomplete.

Relationship specific investments. The transactions cost literature has made asset specificity a key component in the analysis of integration. The central hypothesis is that integration -- the purchase of the decision rights to all relevant physical assets by one party -- will be an economically efficient arrangement in situations where asset specificity is high. A concentration of decision rights in one hand will provide the owner with all the rents and hence make initial ex ante investments worthwhile. By con-

¹⁰More often, of course, the knowledge of rights (and how they might affect potential bargaining) will control behavior without actual bargaining ever taking place. An implied governance structure comes to frame the relationship and affect the incentives of the transacting parties.

trast, non-integration prevents owners from fully appropriating the returns to their initial investments, causing inefficient decisions ex ante.

The simplest case is one in which there are two indispensable and entirely relationship specific assets. Under separate ownership, each side will capture only half of the investment returns (assuming an even split of rents), while under integrated ownership the full returns are appropriated by the single owner. (With a single asset, the corresponding comparison would be between joint ownership and individual ownership.)

There are several qualifications to this conclusion. Two are particularly relevant. First, it is essential that the parties cannot contract on the initial investment itself.¹¹ If they could, the distribution of rents would not affect efficiency. Second, it is implicitly assumed above that the costs of investment get transferred under integration; it is the single owner who pays for all investments. Were it so that the other party remains responsible for the costs of investing in the asset he manages, he will lose all incentives to invest under integration and this may become a much worse alternative.

Thus, whether ex ante investments will be more efficient under integration depends on what rights and responsibilities can be transferred. Certain assets are inalienable, most importantly those related to human capital. The right to decide whether to supply one's human services to an enterprise cannot be sold (because of involuntary servitude). Likewise, the decision on how much effort to expend remains a private matter, irrespective of organiza-

¹¹Actually, the requirement is stronger. Even if investment decisions cannot be directly observed, it may be the case that short-term contracts (based on imperfect information) can provide adequate investment incentives. On this, see Fudenberg et al (1988).

tional context (because effort cannot be observed).¹² Consequently, the incentives to invest effort into an enterprise can be expected to be diluted for an entrepreneur who sells his company but remains in his former managerial position. Benefits to integration rest on the assumption that ex ante investment decisions do not relate significantly to human efforts, but rather to decisions on expenditures of a monetary kind. I will return to this important point in the next section.

Coordination. In the example above, the distribution of decision rights matters, because it influences the division of ex post surplus. (Surplus is only a function of ex ante investments, not of the bargaining process, which costlessly is assumed to achieve an efficient outcome.)¹³ In general, of course, decision rights affect all aspects of bargaining, including importantly the costs of reaching an agreement.¹⁴ If there is surplus to be divided, and both sides can threaten to dissipate it by withdrawing the use of critical assets, haggling over the division of the surplus can be expected. Such bargaining can be productive, if it generates information about the best way to proceed. But if it is relatively clear what the right course of action is and haggling only occurs because parties try to enhance their own share at the expense of the other side's, then bargaining merely reduces surplus. Vesting decision rights with one party is a way of eliminating those costs.

¹²It is true that a change in ownership transfers the rights to design incentive mechanisms and that way gives indirect control over human capital decisions. This dimension has not been carefully studied.

¹³For more detailed models of this variety see Grossman and Hart (1986) and also Hart and Moore (1988). The latter presents a very accessible reduced form analysis, which covers cases with many assets and many interested parties.

¹⁴Bargaining costs are discussed at length in Milgrom and Roberts (1987).

Present modelling technology suggests that we introduce asymmetric information in order to capture such costs. So assume that one side has private information about the value of continuing the relationship in a contingency that lies outside an explicit contract. If both parties have ownership of an indispensable asset, we know from bargaining under asymmetric information that costly delays will result from haggling. The informed will try to extract rents from his private information and the uninformed will try to keep rents as low as possible. By contrast, if either party is given title to both assets, bargaining costs will go to zero. If the informed gets the assets, he will simply impose the best continuation given his information. And rather interestingly, the same is true if the uninformed is made the single owner. With nothing at stake, one can expect the informed to reveal his information quickly and in an unbiased fashion. Thus, for effective coordination of information, low-powered incentives may be necessary.

Note, however, that if one side has human or other non-transferrable assets that are instrumental to realizing the surplus, then that side should be given all assets. Simon's (1951) suggestion that employers should be given the right to decide on the allocation of workers to tasks because of their better information, implicitly assumes that the employers could otherwise hold up bargaining - or that bargaining among workers would be costly.

To sum up, the main arguments for integrating are two: incentives for investment in relationship specific assets and improved coordination of decision making. How relevant are these for innovation? It seems to me that they do not provide strong reasons for placing innovation activities in a large firm. Relationship specific investments are significant, but limited

to relatively small groups. They do not require large scale integration. Neither does there seem to be a great need to coordinate decisions among different research projects as the argument based on costly bargaining would demand. Indeed, it appears that large firms frequently make an effort to keep different innovation projects separate. By contrast, it is easy to understand that coordination gains can be significant in the production, marketing and development of established products. In this view, firms grow large with the increased size of product markets.¹⁵

4. Appropriation and measurement problems.

I stressed above that human capital is an asset that cannot be transferred and therefore incentives for effort may be significantly diluted by removing title to transferrable assets from those whose efforts are central to production. I believe innovation requires significant personal sacrifices particularly in acquiring information. The coordination benefits discussed above were conditional on fixed information. Coordination aids the elicitation of information, but makes the incentives to invest in information correspondingly worse. Appropriating the returns from such efforts is the major problem with integration.

Why can the firm not duplicate market incentives, for instance by giving the innovator the rights to patents that might come with the innovation? There are several difficulties. First, the employee is using the firm's assets, human and physical, in the process of innovating. If he would receive all the benefits, without having to bear the costs, a serious

¹⁵To explain horizontal integration and conglomerates appears much more difficult. The incomplete contract paradigm has not cast much light on that important puzzle.

misalignment would arise. He would undertake innovation with an eye on enhancing his human capital. Excess and wrong kinds of innovation would be likely: bad projects might be undertaken and cost intensive projects would be disproportionately favored. The firm might try to charge fees for using its assets, but the allocation of costs poses a dilemma. Even the best intentioned firm does not know capital costs, as accountants would be the first to tell.¹⁶ In addition, as an interested party, the firm has control of many levers to make accounting measures less reliable. The implication is that individual incentives based on innovation returns must be tempered to provide the agent with a properly balanced objective (as in the simple agency model of section 3). As Williamson (1985) has emphasized, low-powered incentives come to replace high-powered incentives (based on net receipts) upon integration.

Another appropriation problem concerns the rights to decide on the continuation of projects. Innovations occur in several stages. It is easy to imagine situations in which the project leader would like to continue the project when the firm would not. He has much human capital at stake, which the firm cannot appropriate in turn. He is also likely at this stage to have private information on the success probability. Renegotiation of continuation rights would have to take place under conditions of asymmetric information. As mentioned earlier, bargaining under asymmetric information is difficult and will not allow the informed to receive all the rents. This feeds back on the incentives to invest human capital in the project at the initial stage. It is true that some of the same problems could be encountered in market arrangements as well. A lender could refuse additional

¹⁶Since a firm does not put its individual assets (or divisions) regularly on sale, important market information on asset values is missing. See discussion below.

capital. But if it is deemed desirable ex ante, long-term financing can be arranged at the outset. Long-term financing within the firm might not be as easily arranged, partly because the innovator is not bearing research costs as he would as an independent entrepreneur.

Indirectly, the arguments above make reference to the problems of measuring marginal product. Difficulties in identifying relevant costs and benefits, so as to make the innovator bear his marginal share, are central. Of course, even as an individual entrepreneur, measurement problems are substantial. The entrepreneur does not know all the relevant figures either (for instance what future returns he will get). But the knowledge that the money will flow into his own pockets, that nothing will be taken away, still provides appropriate incentives. It is when financial accounts are integrated that the difficulties of measurement become consequential and severe.

To give an example of how integration destroys performance information, consider a scenario in which a smaller research oriented firm is bought up by a corporation.¹⁷ Assume both firms are publicly traded before the acquisition. After the acquisition, the small firm's stock would typically be withdrawn from the market. Evidently that would eliminate a critical information variable that could be used to evaluate managerial performance. The market would no longer monitor the purchased firm separately and the value of that portion of activities would be confused with contributions from the rest of the corporation.

The story is not complete, however, without asking why the corporation cannot continue to trade both stocks. In fact, it can as evidenced by the recent merger of EDS with General Motors. GM, after purchasing EDS, started

¹⁷This discussion is from Holmstrom and Tirole (1987).

to trade GM-E stock, a stock without voting rights, but with a value based on EDS performance. Presumably, this was done to maintain an outside monitor of EDS. Indeed, EDS management got hefty shares in the new stock, as a continuation of their extensive stock incentive plans in EDS when it still was an independent company.

The experiment has not had a very successful history, at least to date. The alleged difficulties relate to disagreements in transfer pricing. Apparently the two managements are trying to resolve disputes arising from an incomplete contract. Since GM holds the main control rights, EDS management does not have their pre-merger bargaining power to extract surplus when unspecified contingencies are encountered. Also, the price of GM-E shares become to some extent manipulable by GM. This does not by itself render the stock valueless. It can be protected by covenants, and in the GM-EDS case it was. However, covenants change the nature of the stock; a GM-EDS share is no longer a pure piece of the economic value of EDS (even with the transfer problems factored in). One would also expect a reduced market interest in information acquisition (indeed, trading in GM-E shares has been lackluster). For all these reasons, it is clear that today's price of GM-E stock does not reflect the contributions of its management as effectively as the pre-merger price of EDS would have done.

To compensate for the dilution in incentives that attend appropriation and measurement problems, the integrated firm could intensify monitoring, and I think it does.¹⁸ But as I will argue in the next section, there are

¹⁸In an influential article, Alchian and Demsetz (1972) have argued that firms emerge in response to problems of joint production. One needs a monitor to mete inputs. To give the monitor the right incentives to monitor, he should be made the owner (receive the residual returns). In the scenario I'm sketching, the argument goes reversely. When two firms are combined, a joint production problem is created (or made more serious), because cost and benefit streams will be confounded. Monitoring is a consequence of integra-

serious impediments that compromise the effectiveness of internal monitoring. The information may be accurate, but the firm cannot act on it as strongly, because of potential collusion between the monitor and the ones he monitors. The key point is that verifiability is an endogenous variable, which depends on the incentives of those who collect the information. What makes market information so powerful is not its accuracy relative to information within the firm (one would expect it to be less accurate, in fact), but rather that market monitors express their conviction by "putting their money where their mouth is".

5. The bureaucratization dilemma.

To say that increased size brings increased bureaucracy is a safe generalization. To note that bureaucracy is viewed as an organizational disease is equally accurate. The biggest patient is the government, whose bureaucratic manners are notorious, though some of the largest corporations appear almost as badly afflicted.

Undoubtedly, there are bureaucratic excesses in corporations. But the fact that bureaucracy is so universal and that it survives even in situations where choice of organizational form is free and subject to strong competitive forces should suggest some virtues. One is inclined to believe that if there were an easy way out of the dilemma, it would surely have been found by now.

Influence costs. In fact, Milgrom (1988) has argued that bureaucratic rules are a rational way of curbing detrimental influence activities in hierarchies. Concentrated authority will invite such activity. Subordinates will try hard to influence a superior's decisions to the extent such decision, not a reason for it.

ons impact on their welfare. Many decisions do. Plainly, some jobs are better than others and if the superior is in charge of allocating jobs, that's worth attention. Getting allocated more resources is also beneficial. It enhances one's value in the corporation as well as one's social status. Pleasing the superior may help in getting better wage raises. And so on. In general, the more discretion the superior has, the more intensive become the efforts to influence. As we all know, authority tends to engender remarkable attention, some of it less desirable.

Why should the boss pay attention to the influence attempts of his subordinates? An important reason is that the boss may be in charge of evaluating performance in order to make the right job-skill matches. He will have to observe signals of ability and cannot close his ears and eyes to efforts by the subordinate to look good. The result is that the subordinate may divert energies to prove his worth in ways that are less productive. The market is not immune to these afflictions either; career concerns give rise to distorted behavior in any situation where performance is being evaluated (see Holmstrom, 1982). But within the firm, the subordinate is being watched more closely and therefore receives more returns from signalling his value. Market authority, being more dispersed, offers less easily identified targets for influence activity. Thus, potential influence costs are higher in hierarchies.

Collusion. A less benign reason why the boss may care to be influenced is that bribes can be offered. In effect, the superior and the subordinate can collude (Tirole, 1986). Monetary transfers may be less common, because evidence of such transfers can leak out relatively easily. Also, people are not entirely unscrupulous. But transfers in kind tend to be viewed more innocently and are certainly prevalent. In exchange for personal services,

flattery and the like, more favorable decisions can be expected. Forgiveness will be a common response to mistakes. The strength of reciprocity can be assumed to grow with time and with the intensity of contact. Hierarchies are more fertile ground than markets in that regard.

Influence activities and collusion may severely compromise the integrity of subjective monitoring information. In its most extreme form collusion may render the monitoring information entirely useless. If the subordinate and superior form a team they can always extract the maximal bonus from the firm. Consequently, monitoring information will have to be ignored. But even with more scrupulous behavior, monitoring information is not apt to be as useful as objective evaluation measures. In order to reduce potential or actual collusion, the firm will want to place constraints on the monitor's scope of discretion. For instance, bonuses may be permitted only so often or to so many in a given time period. But, of course, this will reduce the degree to which monitoring information bears on actual performance. One is caught in an unpleasant trade-off between allowing some collusion or ignoring part of the relevant information. Either way the effectiveness of monitoring is reduced.¹⁹

Another response to collusion is to ask for documentation. On what basis were bonuses awarded? This tends to eliminate purely subjective impressions. It shifts the monitoring focus towards more verifiable but less performance related measures. Wage and promotion policies based on seniority

¹⁹It should be mentioned that collusion need not be bad. If two productive agents can monitor each other more effectively than the principal can, then it will often be desirable to allow them to collude, that is allow them to make cooperative agreements about how hard to work. The principal induces collusion by making the agents responsible for each others' outcomes (see Itoh, 1988 and Holmstrom and Milgrom, 1989). The general issue is what kinds of trades one should allow between agents. Since markets are incomplete, the principal will want to regulate the agents' trade.

and other objective factors can be understood in this light (Milgrom and Roberts, 1987). Direct constraints may be imposed on subordinate conduct (eg. time cards), to the point where checking for obvious errors and violations of the rules become the prime activity of the monitor. This is most common in government organizations, where comparison data and other performance measures are most lacking.

While explicit rules and policies are common, it is worth adding that the mere fear of being suspected of favoritism can lead a superior to discount performance in making wage and promotion decisions, even if he had the discretion to act freely on his information. Suspected favoritism would lower employee morale and invite attempts to purchase favors. Thus, one can expect organizations to be even more rule bound than the explicit evidence would suggest.

I want to stress that what gets this bureaucratic "misery" all started is the loss of performance measures higher up in the hierarchy. The integrity of subjective evaluations is a function of the monitor's incentives. An owner will not have to worry about bribes from an employee (this conclusion may change if there are many employees). If he accepts personal services in exchange for higher bonuses, this is merely an efficient trade. But when the monitor does not bear all the financial consequences of his actions, such trades will be excessive. The lesser his responsibility, the bigger the potential distortions. Stricter limitations on the use of information is implied.

In sum, the internal labor market in hierarchies will be pushed towards bureaucratic manners as a rational response to monitoring and influence

problems.²⁰ The basic principle at work here is the principle of inflexibility featured in section 3. The more difficult it is to reward agents, because performance information becomes unreliable or diffuse, the more heavily the agent's opportunity costs will have to be controlled. This is precisely what the bureaucratic constraints are meant to accomplish. That these tendencies will be hostile to innovation seems plausible. Let me elaborate on a few of the main points.

By definition, rules and rigidities inhibit or discourage activities that are exceptional. Extensive capital budgeting procedures are a particularly severe impediment. Funding requests have to pass many layers of approval in order to bring the decision to a level that carries sufficient responsibility. The chances that unfamiliar and innovative projects get accepted are diminished (the discussion in Sah and Stiglitz, 1986, is apposite). In part this bias will reflect attempts by superiors to protect their specialized investments in human capital. New products and production methods may be a potential threat to their position of leadership.

The move towards verifiable but less relevant performance measures is equally troublesome. Subjective monitoring would be particularly valuable for innovation, since success is so uncertain. Exceptional tolerance for failure is essential. But such performance cannot be checked by conformance to organizational rules or by evaluation reports that can be readily substantiated.

Monitoring limitations suggest that the firm seek out activities, which are more easily and objectively evaluated. Assignments will be chosen in a

²⁰A rather different reason why a firm may find it valuable to promote internal uniformity can be provided along the lines of Kreps (1984). Kreps argues that a uniform corporate culture is an important vehicle for nurturing the firm's reputation.

fashion that are conducive to more effective control. Authority and command systems work better in environments which are more predictable and can be directed with less investment in information. Routine tasks are the comparative advantage of a bureaucracy and its activities can be expected to reflect that.

Finally, in bureaucracies, promotions no longer serve the exclusive purpose of matching skills with tasks. The ability and responsibility of the firm to act as a human capital filter are being compromised as a consequence. Promotions based on measures weakly related to performance is one reason. A more important reason is that employees who have done well and deserve to be rewarded will have to be promoted rather than paid in cash. Promotion as a reward is less subject to misuse by a superior than cash rewards. By promoting the subordinate to a new, presumably more demanding task, the superior subjects himself to outside judgements. Favoritism will become more visible. But then one must make sure that promotion to a new task does not result in a serious misassignment. This forces the firm to operate with employees whose characteristics are of more general use. Employees have to conform to the general culture and objectives of the firm. But innovators are not necessarily good managers as Robert Sculley must have realized at the Apple company when he wanted Steve Jobs out. Uniformity in the treatment of employees translates into uniformity in their characteristics, which may well screen out innovative personalities.

7. Capital market effects.

One of the remarkable features of modern capital markets is that investors are willing to part with their money in such huge quantities with

so limited explicit assurances of getting anything back. Behind this magic is a sophisticated network of institutional arrangements that controls management behavior. Much of the control is indirect through reputation. When firms go to borrow more money, which most have to do with some frequency, a good credit record and a healthy condition are invaluable. For the manager of the publicly held company, the pressures to perform and look good are greatest, because of the continuous assessment that stock prices provide. In this section, I want to mention two reasons why reputation concerns may lead to conservatism and why the problem can be expected to be more severe for larger firms.

In large part, the market learns from a firm's past what to expect from it in the future. Extrapolation is rational since there are characteristics in the stochastic process of returns that have permanence. But the firm is not a passive player in this learning process. Management can make decisions that influence perceptions about the firm's potential. A simple scenario is the following. An investment decision has to be made today. The investment options vary in their return patterns. Some have returns far in the future and others in the nearer term. The market knows exactly the same as management: the investment options, the return distributions and so on. However, the market cannot observe the management's decision directly; it can only infer what the management will do.

In this situation there may be a tendency for management to act myopically. By choosing projects with faster paybacks, early returns are enhanced on average. This raises market expectations about management and firm potential. Of course, later returns will be lower on average, offsetting some of the early gains. Exactly how the trade-off works out in the management's mind requires a specification of its preferences. What one can show is that

if management is paid based on expected marginal product in each period, and management is less patient than capital markets (because of incomplete income smoothing), then there will be a bias towards the shorter term. Note well, that this happens even though the market in the end is not fooled by management's choice. The market expects the bias, but this only reinforces management's need to show short-term results.²¹

Because market expectations will be correct, this story is perfectly consistent with informationally efficient markets as Stein (1987) has observed. More strikingly, a management that tries to maximize the market value of the firm (that is current price) would be led to choose short-term projects. Thus, the common complaint about myopic and conservative American management may be well founded, but an inevitable consequence of our competitive system.

Since innovations tend to pay off in the distant rather than near future, this type of reputation story would suggest that innovations will not be undertaken sufficiently often. However, one has to explain why the problem is more severe in large firms. There are two reasons that fit the paradigm. One is that a disproportionate number of large firms are publicly traded. With the constant monitoring from markets large firms are forced to be more myopic. The second reason is that large firms have more flexibility and independence. The argument for myopia turns on the fact that the market cannot observe the actual actions of management. If management could validate what it is doing, the problem would disappear. That some evidence can be presented is clear and explains why behavior is not entirely myopic. But the more activities there are, the more opportunity there is for unob-

²¹For explicit models that feature myopic behavior, see Narayanan (1985), Campbell and Marino (1987), and Stein (1987).

served allocations that inflate early performance.

These problems of risk taking are of course also present within the firm. In an organization that uses promotion as the main vehicle for rewarding performance, an aversion to risk can be expected among those who see their chances for promotion to be good.

A related reputation story has been presented by Diamond (1987). In his model projects are financed by debt. Projects can be risky or safe and firms can be either good or bad risks. What he shows is that over time, as the firm establishes a reputation for being a good business, interest rates charged on its borrowing will come down. Because of the option feature of debt contracts firms will take risky projects when interest rates are high and safe projects when they are low. Putting the two together means that firms start off choosing risky projects but later revert to safe ones. The increased value of reputation makes it eventually not worth risking. Thus, established firms can be expected to guard their reputation by becoming conservative.²²

I believe both reputation stories touch on relevant dimensions of the problems of channeling capital from markets to firms. Certainly, they are not the only ones that bear on innovation. Unfortunately, our present understanding of investment processes is so limited that it is hard to present a very comprehensive picture. As the discussion in section 2 indicated, finance theory is at a loss in explaining the rich variety of institutions set up to intermediate capital flows. Nor is the role of firms

²²Another implication from a concern for market reputation is that uniformity in activities and products are conducive to reputation maintenance. A retailer who sells both high and low quality items will confuse the quality of observations. For the same reason, it seems plausible that pressures for uniformity in project choice may manifest itself when dealing with capital markets. Mixing innovations and routine projects may raise the cost of capital above the average of undertaking the two separately.

in this process clear. Why is centralized capital budgeting so universal a phenomenon? Is it so that capital budgeting reflects a comparative advantage in distributing funds or is it an inevitable consequence of the bureaucratization dilemma discussed before? These questions are very central in understanding how innovation is distributed across firms, but the answers are not yet available.

7. Conclusion.

According to the theses above integration is primarily motivated by coordination benefits and improved incentives for investment in non-human assets. Large scale production and marketing activities are the main beneficiaries. Innovation, being a small scale activity initially at least, will not gain much by this argument. On the cost side, integration suffers from weaker incentives to invest in human assets. These are the incentives most essential for successful innovation. Performance measures will be confounded and objective market assessments lost. Internal monitoring cannot compensate for this fully, because of potential collusion problems that attend weaker performance measures for the monitor. The firm will restrict freedom by bureaucratic rules in a rational effort to control incentives indirectly. More uniformity in activities and personnel will follow. Both are hostile to innovation.

Innovation activity requires exceptional tolerance for failure; lest tolerance translates into slack, monitoring has to be intensified. How can one improve the incentives for the monitor? By making him financially responsible for the consequences of his judgements. Venture capitalism is a solution of this kind. The venture capitalist is a specialist in evaluating

the quality of potential innovations. Also, he holds a substantial stake in the companies he oversees. The fact that the venture capitalist eventually withdraws, once the firm is up and running, suggests that his monitoring services are no longer as valuable. One reason is that the initial asymmetries in information have been reduced and the market can take over the role of monitoring; indeed, the venture capitalist often leaves when the firm goes public. Another reason may be that the firm is turning from innovation towards reaping the financial rewards from its successful discoveries, and the needs for monitoring are thereby reduced.

For an established corporation, turning up the rate of innovation will by this logic require decentralization. The innovative parts of the business have to be made more independent and financially more responsible. It appears that firms which are innovation oriented are also more decentralized. Not infrequently, departments or divisions of research and development are spun off.

The new organizational advice from business consultants to decentralize is supported by this analysis, but with important qualifications. The advice does not apply across the board, but only to firms that should intensify innovation because of obsolescence of products. Bureaucracy will continue to have a place as an efficient form of organizing large scale production.

REFERENCES

- Anderson E. and D. Schmittlein, 1984, Integration of the sales force: an empirical investigation, *Rand Journal of Economics*, 13-3, 385-395.
- Campbell, T. and A. Marino, 1988, On the incentives for managers to make myopic investment decisions, mimeographed, (University of Southern California).
- Diamond, D., 1987, Reputation acquisition in debt markets, mimeographed, (University of Chicago).
- Fudenberg, D., B. Holmstrom, and P. Milgrom, 1988, Short-term contracts and long-term agency relationships, forthcoming in *Journal of Economic Theory*.
- Grossman, S. and O. Hart, 1986, The costs and benefits of ownership: a theory of vertical and lateral integration, *Journal of Political Economy*, 94, 691-719.
- Hart O. and J. Moore, 1988, Property rights and the nature of the firm, mimeographed, (Massachusetts Institute of Technology).
- Holmstrom, B., 1982, Managerial incentives - a dynamic perspective, in *Essays in honor of Lars Wahlbeck*, Helsinki: (Swedish School of Economics).
- _____ and P. Milgrom, 1987, Aggregation and linearity in the provision of intertemporal incentives, *Econometrica*, 55-2, 303-28.
- _____ and P. Milgrom, 1989, Assigning tasks to agents, mimeographed, (Yale University).
- _____ and J. Tirole, 1989, The theory of the firm, manuscript to appear in *Handbook for Industrial Organization*, Eds. R. Schmalensee and R. Willig. (Elsevier Publishing Company).
- Itoh, H., 1988, Essays on the internal organization of the firm, unpublished Ph.D. thesis, (Graduate School of Business, Stanford University).
- Kreps, D., 1984, Corporate culture and economic theory, mimeographed, (Graduate School of Business, Stanford University).
- Lazear, E., 1989, Pay equality and industrial politics, *Journal of Political Economy*, 97-3, 561-580.
- Milgrom, P., 1988, Employment contracts, influence activities, and organization design, *Journal of Political Economy*, 96-1, 42-60.
- _____ and J. Roberts, 1987, Bargaining and influence costs and the organization of economic activity, RP 934, (Graduate School of Business, Stanford University).

- Minahan, J., 1988, Managerial incentive schemes and the divisional structure of firms, mimeographed, (University of Massachusetts at Boston).
- Narayanan, M., 1985, Managerial incentives for short-term results, *Journal of Finance*, XL-5, 1469-1484.
- Sah, R. and J. Stiglitz, 1986, The architecture of economic systems: hierarchies and polyarchies, *American Economic Review*, 76-4, 716-727.
- Scherer, F., 1984, *Innovation and growth*, 2nd Edition, (Cambridge: MIT Press).
- Simon, H., 1951, A formal theory of the employment relationship, *Econometrica*, 19, 293-305.
- Stein, J., 1987, Efficient stock prices, inefficient firms: a signal-jamming model of myopic corporate behavior, forthcoming in *Quarterly Journal of Economics*.
- Tirole, J., 1986, Hierarchies and bureaucracies, *Journal of Law, Economics and Organization*, 2, 181-214.
- Williamson, O., 1975, *Markets and hierarchies: analysis and antitrust implications*, (New York: The Free Press).
- _____ 1985, *The Economic institutions of capitalism*, (New York: The Free Press).