



An Economic Approach to Influence Activities in Organizations

Author(s): Paul Milgrom and John Roberts

Source: *American Journal of Sociology*, Vol. 94, Supplement: Organizations and Institutions: Sociological and Economic Approaches to the Analysis of Social Structure (1988), pp. S154-S179

Published by: [The University of Chicago Press](#)

Stable URL: <http://www.jstor.org/stable/2780245>

Accessed: 08/08/2013 18:05

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *American Journal of Sociology*.

<http://www.jstor.org>

An Economic Approach to Influence Activities in Organizations¹

Paul Milgrom and John Roberts

Stanford University

Members of organizations spend considerable time, effort, and ingenuity attempting to influence decision makers. Such influence activities may bring benefits to the organization, but they also involve real costs. This essay offers an economic rationale for such influence activity as representing rational, self-interested behavior in the presence of informational asymmetries and an analysis of how the design of the organization's structure and policies should respond to the incentives for attempting influence. It is posited that information valuable for the organization's decision making is directly available only to members of the organization who have some personal stake in the decisions. These individuals may then have an incentive to try to manipulate the information they develop and provide in order to influence the resulting decisions to their benefits. This can be costly both in degrading the quality of decision making and in diverting the attention and effort of the organization's members from more productive activities. The organization has three different methods it can employ to discourage excessive influence activities and to encourage more directly productive uses of time and effort. It can limit access to decision makers and participation in decision making; it can alter its decision-making criteria to favor those performing well in productive activities; and it can provide direct financial incentives to encourage the desired allocation of effort. It is shown that an efficiently designed organization will use such financial incentives only as a last resort. Instead, it will always first alter its decision-participation policies and decision-making criteria.

This paper is concerned with influence activities in organizations. As has been well documented (see, e.g., Madison et al. 1980), members of organizations spend large amounts of time, effort, and ingenuity attempting

¹ This work was done while Milgrom was visiting the University of California, Berkeley. It was supported by the National Science Foundation and the John Simon Guggenheim Memorial Foundation. We thank Chris Winship and two anonymous referees for their comments and suggestions. Requests for reprints should be sent to Paul Milgrom, Department of Economics, Stanford University, Stanford, California 94305.

© 1988 by The University of Chicago. All rights reserved.
0002-9602/89/9407-0014\$1.50

S154 *AJS* Volume 94 Supplement S154–S179

to influence decision makers. Sometimes, this activity is aimed at realizing well-specified, immediate results: individuals campaign for pay increases, choice job assignments, and promotions, and groups push for larger budgets, acceptance of their policy proposals, adoption and continuation of their programs and projects, and rejection or abandonment of those that are harmful to their interests. In other cases, these efforts seem to be intended to create a general good impression with decision makers, which may lead them to favor the individual or group in as-yet-unforeseen decisions. Examples here would include both activities that are clearly productive (doing a better job than necessary on some visible task) and ones that are not of obvious direct value to the organization (the sort of activity described as “buttering up the boss”).

As we shall argue, such influence activities may bring benefits to the organization, but they also involve very real costs. Overall, their effect on the efficiency of the organization can be significantly negative. Thus, from the standpoint of standard economic analysis, it is not immediately obvious why they are countenanced when they detract from the organization’s performance. Our objective in this essay is to offer an economic rationale for such influence activity and an analysis of its effects on the design of the organization’s structure and policies. In particular, we offer an explanation of the occurrence of influence activities as the outcome of the interplay among rational, self-interested members of the organization. Our analysis of the organization’s response to influence activities is similarly based on optimizing behavior, with those who are charged with designing the governance of the organization and with selecting and implementing its policies accounting correctly for how their decisions will affect behavior.

This approach to the analysis of organizations is characteristic of the relatively new, but already very extensive, literature on the economics of organization. This literature addresses a wide array of phenomena, ranging from the choices of organizing as a corporation, a partnership, or a not-for-profit (Fama and Jensen 1983) and of what extent of vertical integration to employ (Williamson 1985; Grossman and Hart 1986), through the design of the number of levels of hierarchy and the spans of control (Rosen 1982; Bohn 1987; Geanakoplos and Milgrom 1985), to the adoption of “golden parachutes” for managers in takeovers (Baron 1983). However, the main focus of this work has been on issues of incentives and compensation. These questions have most often been studied in terms of formal models of the principal-agent relationship (see Arrow 1985; Hart and Holmstrom 1987; Holmstrom and Tirole 1987). In this literature, one or more “agents” act on behalf of a “principal”: workers on behalf of owners, lawyers on behalf of clients, suppliers on behalf of customers. The agents have an informational advantage: only the agent knows what

action he has taken in pursuit of his or the principal's goals, or only the agent has access to the specialized knowledge on which his action is based. The principal's problem is to design a compensation and control (monitoring) system that attracts and retains good agents and motivates them to behave appropriately (in the principal's interest). The asymmetry of information prevents easy determination of whether a particular observed action or outcome corresponds to desirable behavior and thus renders the problem nontrivial.

The basis of our explanation is an argument involving informational asymmetries. We take it as given that some of the information that is important for the organization to make good decisions is not directly available to those charged with making the decisions. Instead, it is lodged with or producible only by other individuals or groups that are not empowered to make the decisions but may have a direct interest in the resulting outcomes. Examples of such information are many. The directors of a firm may have the final say on whether a new plant will be built, but only the division whose products will be made in the plant can generate important parts of the relevant information on the likely profitability of the new facility. The tenure and appointments committee of a university may have the final decision on professorial appointments, but it must rely on the department proposing the appointment for the information needed to make the right decision. The manager of a production line may be charged with setting the speed of the line, but only the workers on the line know how onerous they find the work at different line speeds and how much they must reduce the care they take in doing their tasks (and, therefore, degrade quality) as the line's speed increases. Or an employer may have the power to set compensation, but she lacks direct information on the market opportunities available to her employees and can obtain such information only from them.

In such situations, the members of the organization may have an incentive to try to manipulate the information they develop and provide in order to influence decisions to their benefits. Such manipulation can take many forms, ranging from conscious lies concerning facts, through suppression of unfavorable information, to simply presenting the information in a way that accentuates the points supporting the interested party's preferred decision and then insisting on these points at every opportunity.

This influence activity can be costly to the organization in a number of ways. First, to the extent it is successful in biasing the decision maker's information, it may lead to decisions being taken that are inefficient from the organization's point of view.² Second, the time and effort spent on

² Note, however, that decision makers who properly account for the incentives of interested parties may be able to arrive at correct decisions despite attempts to in-

influence activities (and in dealing with them) are resources with valuable alternative uses. Yet, to the extent that influence activities are aimed at shifting the distribution of the net benefits of decisions among the members of the organization, these activities need bring no efficiency gain to the organization that offsets the costs involved. Of course, if the influence activities actually lower the quality of decisions, their net effect is even more negative.

Thus, one would expect the organization to attempt to discourage excessive influence activities. The most direct way to do this would be to ignore or even punish attempts at organizationally unproductive influence. This might be effective in some circumstances, especially when it is clear that the behavior in question is designed to advance personal objectives at the expense of organizational ones. However, in other circumstances this approach may be extremely expensive or impossible. In particular, in many situations the decision makers may strongly suspect that individuals or groups are attempting to manipulate their decisions by distorting the information they produce and provide. Yet, it may be impossible to be sure that this is going on (so that punishment is essentially infeasible) or to undo the distortion so as to “ignore” the attempt at influence.

There are three options the organization can adopt. First, it can limit influence activities simply by closing the communication channels, at least for some decisions. For example, the n th time a department chair shows up in the dean’s office to complain about some decision, the dean may feign having another appointment and effectively refuse to listen to the complaints. However, even when this approach works, it has the cost that valuable information is not made available to support decision making. Moreover, in many situations it may not be feasible to close the communication channels selectively; the chair may make an appointment to discuss another matter, then bring up the old complaints again. As long as the communication channels are open between decision makers and other members of the organization, there will inevitably be opportunities for influence activities.

Second, given that the channels are open, the organization’s decision makers can reduce the returns to influence activities by committing themselves in various ways that limit their discretion and constrain their ability to respond to the information supplied by others. For example, a strict seniority-based promotion scheme may be adopted, or provisions may be included in the corporate charter that limit the ability of those charged

fluence their decision making by manipulating the information with which they are provided. Even so, the interested parties may still have an incentive to try to influence the decision making. (See, e.g., Milgrom and Roberts 1982, 1986.)

with operational decisions to change the organization's fundamental structure and policies.

The third option is for the organization to adjust its compensation, promotion, investment, and other decision criteria from what would otherwise be optimal in an attempt to ensure that organizational and individual goals are congruent and that individuals can gain only by promoting the objectives of the organization. This, too, is costly. Thus, even if such measures could eliminate organizationally unproductive influence activities, the measures themselves would still be costly for the organization.

An important example of this last approach concerns compensating differentials in pay. In some situations, it may be possible to adjust monetary compensation to insulate employees fully against any non-monetary effects of the organization's decisions. They would then be indifferent among the various decisions the organization might take, and they would have no reason not to cooperate fully in promoting the organization's objectives. In fact, those familiar with the economic theory of labor markets may already be wondering why compensating differentials do not automatically arise and eliminate any incentives for influence activities.

If we were to make the textbook economic assumption that labor markets are perfectly competitive,³ with compensation adjusting constantly to balance supply and demand, and if we further assume that the costs of changing jobs are trivial, then we would deduce that compensating differentials should arise and that there ought to be no motive for influence activities. For example, competition for jobs that involve high prestige or attractive perquisites ("on-the-job consumption") would lower the monetary rewards associated with these positions. Correspondingly, dirty or dangerous work would command premium wages. In this case, no employee would prefer any job that he or she is capable of performing to the one he or she holds. Thus, there would be no reason to try to influence assignments, job design, or similar decisions. If, however, we move away from such extreme perfect-market assumptions and consider more realistic environments, then the theory of labor contracting provides some possible reasons for the incompleteness of compensating differentials. (See Rosen 1985; Hart and Holstrom 1987, for surveys of this topic.)

Once markets are assumed not to clear continuously but to be mediated by multiperiod, explicit or implicit contracts, then, even if we assume that there are no informational asymmetries among individuals that would complicate writing complete contracts, optimal contracting results

³ This assumption rules out any significant differences in information among individuals.

in full compensating differentials only when one or the other of two very special assumptions holds (Milgrom 1988). Under the first of these, preferences must be additively separable in income.⁴ This very strong condition implies that the compensation demanded by an employee for providing effort or accepting undesirable job characteristics does not depend on his level of income or wealth. Under the second, there must be no costs—in lost wages, search and moving expenses, personal disruption, or whatever—in changing jobs. Since neither assumption appears empirically very plausible, and since even these strong assumptions cannot assure that compensating differentials are paid when employees have private information, one must expect that compensating differentials will not be complete. Consequently, employees will care about such issues as the characteristics of the jobs to which they are assigned and the size of their departments' budgets.

Thus, compensating differentials ought not to be expected to arise automatically and to eliminate the incentives for trying to influence decisions. Nevertheless, the possibility of actively designing compensation and other policies so as to control influence activities still may be an option that is available to, and used by, the organization.

We develop our arguments more formally in the next section with a model of a specific decision problem. In this model, a firm must select one of two current employees to fill a "key" job. The employees can devote time and attention either to increasing output in their current assignments or to establishing their qualifications for the key job.⁵ Two key assumptions are that the allocation of effort between these two activities is not observable (nor can it be inferred from any observables)⁶ and that the organization has no independent source of information about the employees' abilities. The firm can observe differences between employees in their performances in the current assignments and can also observe differences in their qualifications for the key job. However, it cannot determine whether observed differences in qualifications reflect actual differences in the employees' expected productivities in the key job or are merely the result of one of them having devoted too much time to building his

⁴ That is, preferences between income, x , and other determinants of satisfaction, y , must be representable by a utility function of the form $x + v(y)$.

⁵ These qualifications need not be obviously related to productivity in the key job in any very direct way: an employee who adopts the bosses' dress and hobbies, laughs at their jokes, and is pleasant to their spouses is demonstrating a commonality of interests with the bosses, an ability to cooperate, an eagerness to please, or a taste for self-abasement. Any of these may be valuable characteristics.

⁶ In particular, if employees devote excessive time to establishing their credentials, they can selectively suppress information to prevent this fact from being evident. Naturally, the information that is suppressed will be chosen carefully to suit the employee's objectives.

credentials. Also, low levels of performance in the current jobs are not proof of employees' having slighted their responsibilities in order to campaign for the key job; poor performance may reflect nothing more than bad luck.

In this case, the employees' pursuit of their self-interests may lead them to choose their allocations of effort differently from what is optimal for the organization. For example, if the key job is very desirable and the promotion decision is made solely on the basis of which employee seems best qualified, then they will expend too much effort on establishing their credentials for the job. Then, performance in the current job will be degraded.⁷

In this context, we examine how the organization can optimally respond to maximize its expected returns, counting both the returns it receives from the employees' current tasks and those from the key job. The instruments we consider as being available to the organization are the wages to be paid to the employee assigned to the key job, those paid to the other employees, the circumstances in which the employees' information on their qualifications is gathered from them, the criteria used in promotion (including direct qualifications, current job performance, and other factors—such as seniority—that, in terms of the model, appear random), and the extent to which compensation in the current job is linked to performance.

In this model, we assume that individuals at subordinate levels in the organization are motivated solely by concern for their expected money incomes and that the decision makers in the organization seek to maximize the organization's net returns, which we also take to be expressed in monetary terms as profits. These represent extreme assumptions, even in an economic model. Not only are we excluding most of the concepts and variables that psychologists and sociologists have found useful in their analyses of organization behavior (such as social norms, status, or finding satisfaction in work); we even rule out many standard elements of economic treatments, including aversion to risk, aversion to effort, and possible lack of profit maximization by top-level managers. We do this not because we think our assumptions are descriptively accurate. Rather, our intent is to explore the extent to which we can obtain useful insights from a simple, easily understandable structure that may later be enriched to make it more realistic. In the same spirit, we assume that information can be costlessly and accurately transmitted, assimilated, and evaluated, so that all actors are capable of making arbitrarily complicated calculations and of carrying out extremely subtle reasoning.

⁷ This model is thus one in which employees' career concerns have negative effects on organizational performance. This general possibility is suggested by Holmstrom (1982); see also Holmstrom and Ricart i Costa (1986).

These assumptions militate against the emergence of influence activities. First, because of our assumptions on preferences, it is always possible to arrange compensation so that the members of the organization are indifferent about the assignment to the key job: earnings can be made independent of whether one is given the key job. With wage equalization in place, there would be no reason to attempt to influence the job-assignment decision. Moreover, there is no room in our assumptions for top decision makers to act in a fashion inconsonant with profit maximization, and thus no opportunity for influence activities designed to win favor by appealing to a superior's personal self-interest at the possible expense of organizational goals. Also, since employee satisfaction is of interest to the organization in this model only if it improves profits and, in any case, is a function only of compensation, various other possible avenues by which influence activities might enter are also blocked. For example, allowing employees a say in decisions because they directly enjoy having such a role or because it gives greater legitimacy to the resultant decisions has no place, given our assumptions.

Our model involves two elements of private information. In particular, the employees' allocations of effort between directly productive work and credential building are unobservable to the organization's decision makers, as is the "quality" of any outside job offers that the employees receive. The latter means that the key job, which is defined by the organization's incurring especially high turnover costs if the person occupying this position should quit, would command a premium wage, absent concern with influence activities. The role of this premium is to reduce the probability of turnover by making it less likely that outside offers will be successful.⁸

We will show that, even when influence activities are taken into account, it will typically still be optimal to pay more (on average) in the key job, though this creates an incentive problem that could be eliminated by wage equalization. Because the key job pays better, employees will want this job, and, under the promotion and incentive-pay policies that would otherwise be optimal, they will have an incentive to spend too much of their time campaigning for the key job. Because their misallocation of time and effort is not directly observable, the firm must use indirect methods to offset this incentive. In particular, the firm cannot simply require that a proper allocation be made and enforce this rule through disciplinary measures. Instead, it will usually reduce the pay differential

⁸ Note that the offers received are assumed to be independent of the wage currently being paid. Note too that if the quality of outside offers were observable by the firm, then it could prevent turnover simply by meeting outside offers received by employees it wanted to keep. However, this would encourage spending time trying to generate attractive offers, which is another influence cost.

between the key and non-key jobs, thereby reducing the relative attractiveness of the former and the incentives for using influence activities, and it will often tie compensation in the current position to performance through performance bonuses, which reward spending time on directly productive activities.

More striking is that it is *always* optimal for the firm to adjust its promotion criteria and information collection rules from what would otherwise be optimal. In particular, it sometimes will promote first on the basis of on-the-job performance, even though (in our model) this is unrelated to the employees' abilities in the key job;⁹ it will sometimes not solicit the employees' information at all, even when having it would allow better decisions; and sometimes it will ignore the information it has on qualifications and promote the less-qualified person as a reward for performance. Thus, the very design of the organization—in terms of decision-participation policies and promotion criteria—becomes a variable that is used to achieve performance in the face of informational problems.

In the next section, we develop our model and the results. In the following section, we examine some applications of the ideas in this paper. In particular, we offer some comments on stylized differences in the compensation, promotion, and decision-participation policies of U.S. and Japanese firms. We also sketch some other applications of the logic of influence costs that are related, but less directly connected, to the analysis in the next section. We close with a few concluding comments.

A FORMAL MODEL OF INFLUENCE ACTIVITIES AND ORGANIZATIONAL DESIGN

In this section, we offer a completely specified, formal model that captures the essence of the arguments sketched above. Like any such model, this one is necessarily “unrealistic” on important dimensions and thus “wrong.” This is the inherent nature of a formal model: it abstracts from reality, focusing on what seem to the modeler to be the crucial phenomena, so as to yield a structure that is simple enough to be easily analyzed and understood. In the process of constructing such a model, one is forced to delineate carefully the meaning of various terms, to specify the logical relationships among the concepts one has defined, and to recognize the assumptions that lie behind one's intuition. The discipline imposed by this process is the pride and strength of economic analysis; the concomitant narrowness or lack of realism is the cost.

⁹ More realistically, we would expect that past performance would carry at least some information regarding productivity in the key job. In this context, our result might be interpreted as indicating that controlling influence activities will lead to “excessive” weight being given to past performance in promotion decisions.

Consider an organization in which upper-level managers seek to maximize expected profits in all situations through whatever instruments are available. Given this objective, it is natural to refer to the organization as a firm.¹⁰ At regular intervals, certain “key jobs” become available in the firm, and the firm’s policy is to fill these from within. The distinguishing feature of a key job is that the firm will incur high costs if the occupant of the job quits. These costs might include the expenses of finding and training a replacement and the reduced output during the transition, as well as the wage paid to the replacement. They may vary among different key jobs. These key jobs also vary in the gains that can be realized from appointing a better-qualified person to the job: we refer to this characteristic as the job’s “importance.” We assume that both the importance of a particular key job and the replacement costs associated with its occupant’s quitting are a priori random variables whose values do not become known until the next vacancy occurs.

At any given time, there are a number of employees in the firm who are candidates for the next key job; for simplicity, we will assume this number is two. They are currently assigned to positions in which the resultant profits to the firm are uncertain, with greater effort or attention from the employee to his job increasing the probability that the realized profits will be high rather than low.¹¹ The employees are identically productive in their current jobs but differ in their productivities for the key job in a fashion that is not directly observable either by them or by their superiors. Each employee is equally likely a priori to be the more productive in the key job.

We quantify all this by assuming that if an individual with productivity q is assigned to a key job with importance I and replacement-cost parameter C , then the maximum resulting profits are $\pi(q, I)$, and it is from this quantity that an amount $R(C)$ is subtracted if the individual quits. We assume that $R(C)$ is strictly positive and nondecreasing for all C , so that higher levels of C mean higher turnover costs, and that the rate at which π increases with q is increasing in I , so that higher levels of I mean that the value of assigning a more productive worker to the key job is greater. The random variables I and C are scaled so that values of zero denote minimal importance and turnover costs and values of one denote maximal importance or costs. We assume that the density of I is bounded below and that $\partial\pi/\partial q = 0$ at $I = 0$. This means that there are some jobs

¹⁰ Bendor, Taylor, and van Gaalen (1987) analyze incentives for certain types of influence activities in public agencies, and Walder (1986) discusses the proliferation of influence attempts in industry in communist China.

¹¹ Both high and low profits are possible at any level of effort, so low profits are not proof that the employee neglected his responsibilities.

for which the gains from assigning a more productive employee to the job are negligible.

Once the assignment of an employee to the key job is made, the employee may receive outside employment offers. We assume that the total compensation associated with such an offer cannot be verifiably reported to the firm, even though the terms of the offer may be completely clear to both the employee and the outside employer. The idea here is that the current employer cannot be sure of how attractive the new job is to the employee compared with the current one. We assume that, before the fact, the employer and employees agree on the probability distribution, G , over the offers that each employee might receive.

In any case, we assume that employees must always receive at least some base wage w^0 , which can be thought of as the minimum pay needed to attract a worker to the firm. Then the distribution of outside offers can be taken to be such that there is probability one of receiving an offer of at least w^0 . We will also assume that there are large numbers of workers available at the wage w^0 , so that the firm need not pay more than this to hire a replacement worker, and that there are no special costs in replacing a worker who has not been assigned to a key job.

We have not assumed that employees prefer to occupy key jobs; instead, we want this to be an endogenous outcome of the firm's policies. To this end, we assume that the distribution of outside offers that a worker might receive does not depend on whether he is assigned to the key job. If this were not the case, the key job could be valued just because it gave access to better outside opportunities.

As indicated before, each employee is assumed to act to maximize his expected income. We will focus on just a single episode of selecting one employee for one key job and will treat both the firm and employees as ignoring any possible future promotion possibilities in their current decision making. In this case, the employee's expected income is the sum of three terms. The first is the expected income received in the current assignment; the second is the product of the probability of being assigned to a key job times the expected income if so assigned; and the third is the probability of not getting the key job times the income received in this case. (Note that these incomes may come from another job if the employee leaves the firm.) The variable that each such employee controls is his allocation of time or effort between productive activities that improve the chance of a high-profit outcome in the current position and influence activities designed to affect his chances of getting the key job. This allocation must be made before either the importance of the key job or the turnover costs associated with it are known. However, when making their allocation decisions, the employees do know the firm's compensation and promotion policies.

Note that employees will not have any reason to deviate from whatever allocation of time the firm might desire unless the key jobs are financially desirable and the firm permits influence activities to be effective. In particular, if the financial rewards (including outside opportunities) are the same in key and non-key jobs, or if the firm does not permit and reward influence activities, the employees will be willing to devote all their efforts to productive activities if the firm so directs.

As suggested earlier, the source of influence is the employees' abilities to provide useful information about their productivities in a key job. However, the employees are not freely endowed with this information; it must be created, and this may take time and effort. We refer to this information about an employee's prospective productivity as his "qualifications" or "credentials." The time and effort devoted to establishing qualifications and developing credentials then comes at the expense of attention to the employee's other responsibilities. Thus, credential building has the cost to the organization of reducing the expected profit contribution from the employee's current position.

The credentials of an employee are modeled as a random variable, \hat{q} , whose distribution depends on q , what his productivity would be in the key job, and s , the amount of time spent on establishing credentials.¹² The employee knows s (and learns \hat{q}) but does not observe q , while the firm neither knows nor directly observes any of these variables. However, the employee can report the realized value of \hat{q} in a verifiable fashion if the firm permits this. Both the employees and the decision makers in the firm know the distribution of \hat{q} depends on q and s . If \hat{q} is reported, then, given some hypothesis about the level of s chosen by the employee, the decision makers in the firm can make an inference about the actual value of q . We assume that increasing s probabilistically increases the estimate of q that will result from this inference. Thus, increasing s increases the chance that this employee will appear to be the better qualified. Of course, increasing s reduces the time and effort, t , available for the current job and thus probabilistically reduces performance in this position.¹³

Those charged with designing the organization and setting its policies so as to maximize expected total profits are constrained by the self-interested behavior of the employees, both in their time allocations and in their decisions about whether to accept outside offers. Of course, what constitutes self-interested behavior depends on the design of the organization and its policies. We model the policymakers as attempting to maximize expected profits, subject to these constraints, by choosing from a set

¹² We assume that both q and \hat{q} have distributions that contain no mass points.

¹³ Henceforth, "performance" will be reserved for use in referring to the current job, and "productivity" for use in connection with the key job.

of decision-participation, promotion, and wage and incentive-pay policies that are described by parameters α , β , γ , b , and w_N and a function $w_K(C)$.

The decision-participation policy is represented by a parameter β with the interpretation that the employees will be asked to report their qualifications for the key job only if the job is sufficiently important; that is, $I \geq \beta$. Otherwise, the employees' information is simply not solicited.

The promotion policy is assumed to be as follows. If the realized importance of the key job is sufficiently high ($I \geq \gamma$, where we assume that $\gamma \geq \beta$), the policy assigns the key job to the worker with the best credentials and pays attention to past performance or other factors only if two candidates are exactly equally well qualified. (However, this will occur with probability zero in our model, so that the second-level criteria are essentially irrelevant.) If the job is of "middling" importance ($\alpha \leq I < \gamma$), then primacy is given to good performance. In this case, credentials are considered if they are reported ($\beta \leq I$) and the employees' performance levels are equal. If $\beta > I$, so that qualifications are not reported, ties in performance are broken in some apparently random fashion. Finally, if the job turns out to be quite unimportant ($0 \leq I < \alpha$), then both performance and qualifications are ignored, and the assignment is made at random. (Note that we may assume that $\beta \geq \alpha$, so that information is not sought if it will not be used.)

Finally, the wage and incentive-pay scheme involves the wage w_N to be paid to the worker not assigned to the key job, the monetary bonus b to be paid if the employee's performance in the current job is high, and the pay $w_K(C)$ in the key job.¹⁴

The only part of the compensation scheme that we allow to depend on the realized value of C is the pay $w_K(C)$ in the key job.¹⁵ We further assume that the firm can commit to its policies and will follow them even after it learns the importance of the key job, the turnover costs associated with it, the employees' performances and qualifications, and whether the person in the key job receives an outside offer. The structure of the decision-participation and promotion policies is shown in figure 1.

This class of policies is fairly rich, but it certainly is not fully general. For example, more generally, the relative weights on performance, credentials, and other factors in the promotion policy could vary with each realized value of the parameter I giving the importance of the key job. Similarly, w_K could depend on the employee's credentials and on I , as

¹⁴ Note that we have said nothing about the pay level in the current job. This variable is not central to our analysis, and a variety of ways of including it lead to qualitatively similar results. Thus, in the interests of simplicity, we ignore this variable.

¹⁵ More generally, both pay in the base job and the bonus for performance might depend on both C and I if these were known at the time these payments were due. Of course, if they are not yet known, payments cannot be made contingent on them.

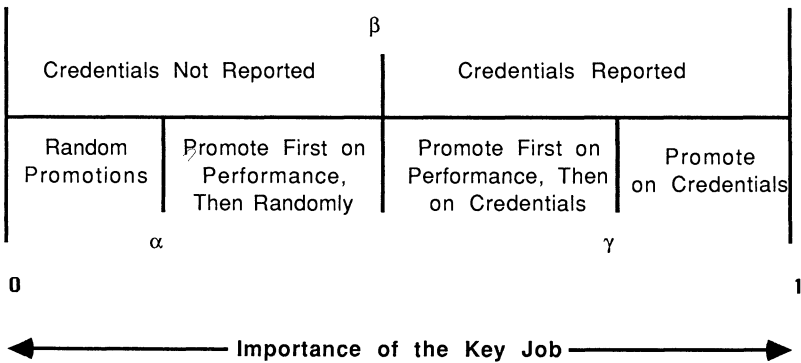


FIG. 1.—Structure of the decision-participation and promotion policies

well as on C . However, solving the full, general design problem is far beyond the scope of this paper. Moreover, the restricted set of policies we consider does seem to capture many of what seem to be the essential elements of the problem.

First, the decision-participation problem allows the firm to commit to ignoring the employee's information if it so chooses, even though this information is valuable and not otherwise available. The potential advantage of not allowing the employees to participate in the promotion decision by providing information is that they then will have no incentive to divert effort from productive activities. A decision-participation policy in which β is set at one means that the firm ignores any information that employees have or generate. If $0 < \beta < 1$, the firm allows employees to provide information on their productivity when the job is relatively important and the effect on profit of assigning the better candidate to the key job is correspondingly significant. A policy of $\beta = 0$ means that employees' information is always sought, even though this may set the stage for an incentive problem.

If α is set at a positive level, then noneconomic factors are allowed to determine some promotion decisions completely. Given the senior management's profit-maximizing behavior, we might expect that α will in fact be set at zero, since doing so while keeping β fixed increases the range over which performance is rewarded and so costlessly increases the incentives for devoting effort to productive activity.

If γ is set at zero, then the firm always promotes the candidate with the best credentials. Thus, all promotions are made on the basis of maximizing expected profits from the key job. However, setting $\gamma = 0$ also gives maximal play for influence activities because improving one's credentials will always increase the chances of being promoted. Setting γ strictly greater than α means that some promotions are made as a reward for

performance for incentive purposes, even though past performance (by assumption) has no relation to productivity in the key job.

If γ is set at zero, then promotion is based only on qualifications and, effectively, not at all on performance.¹⁶ Then, the only positive incentive to put any effort into productive activities comes from giving bonus payments for high performance in the current job.¹⁷ Even when γ is positive, bonuses provide additional incentives for increasing productive activity and, correspondingly, limiting influence activities.

Higher pay levels in the key job help hold key employees and reduce turnover costs, but increases in $w_K(C)$ relative to w_N increase the relative desirability of these jobs with concomitant effects on incentives and behavior in the current job. To the extent that these effects are costly, higher retention of key employees is thus doubly costly.

Before we consider the firm's optimal choice of policies, it will be useful to establish, as a baseline, what the optimal values of these parameters would be if the firm could monitor and control the employees' allocations of time. In this case, the firm and employees could design a mutually acceptable contract that would specify the employees' allocations of their time between s and t . This specification would be determined exclusively by the criterion of maximizing the resultant expected gross profit, with only the distribution of these returns being affected by such factors as the relative bargaining power of the firm and the workers¹⁸ (see, e.g., Milgrom and Roberts 1987). The contract would include penalties to be applied if the employees allocate their time differently from what the contract specifies or, equivalently, rewards to them if they abide by the contract provisions. These agreed-on penalties and rewards ensure that the contract will be followed, because the penalties or rewards are sufficient to make following the contract more attractive than deviating and being caught and the observability of s and t means that any cheating will be detected. We will refer to the contract and policies that would be followed if time and effort allocations could be observed as the "first-best."

It is clear that the first-best will involve the firm's setting $\alpha = \beta = \gamma = 0$; any information the employees generate would be communicated and used, with the key job assigned to the best-qualified employee. This is because there is no incentive problem in motivating the correct choice of s . Also, the firm will set $b = 0$, again because the observability of the

¹⁶ Recall that ties on qualifications have probability zero.

¹⁷ This assumes that the non-key job is not the more attractive one. We shall see that this is the case.

¹⁸ In our model, bargaining power is determined by the availability of replacement workers at the wage w^0 and by the employees' outside options.

effort allocation ensures that there is no incentive to slight current responsibilities, and so there is no need to provide incentives indirectly by rewarding high performance through bonus pay.

Finally, to determine the optimal wage payments, we should recall first that the wage in both the key and non-key jobs must be at least w^0 . In the non-key jobs, there is no reason to pay more than this because there are plenty of workers available at this wage and the costs of replacing an employee in a non-key job are trivial. Thus, the wage in the non-key jobs will be w^0 .

Determination of the wage in the key job will involve the firm's trading off the costs of paying a higher wage against the gains it receives from reducing the number of quits. As can be shown (Milgrom 1988), the optimal solution is for the firm to set a wage that is increasing in C , the costs it incurs if the employee quits. Further, if the costs of replacing a key worker are sufficiently large for all C or if the chances of getting an outside wage offer exceeding w^0 are sufficiently high, then the wage in the key job will exceed w^0 , and workers will see this job as desirable.

Finally, in this first-best case the firm would ask employees to set s at whatever level was optimal, given the costs in probabilistically reduced output that come with higher values of s and the corresponding benefits in better information and resultant job-assignment choices. If we assume that the effect of increasing effort on raising the probability of high performance is sufficiently large,¹⁹ then this allocation will involve s 's being strictly less than T , the total time available (so that t is strictly positive).

With the employees' allocations of time between productive and influence activities being unobservable, the firm's problem becomes one of maximizing its expected profit (as before), but subject to the constraint that employees must find it to be in their own self-interests to choose the levels of influence and productive activities that the firm desires. This constraint means that the firm will have to design its policies to provide incentives for employees to adopt the behavior it seeks.

A key point to recognize is that the first-best policy of $\alpha = \beta = \gamma = 0$, $w_N = w^0$, and $w_K(C) = W(C) > w^0$ is not "incentive compatible" if s is not observable; that is, employees will not find it in their interests to select the levels of s and t that would be specified in the first-best contract. With the key job commanding a premium wage, employees will covet it. With appointment to the job fully determined by qualifications, the employees have an incentive to spend all their time on building their credentials (i.e., $s = T$, $t = 0$). This is typically suboptimal for the firm because expected current output is depressed by too great an amount.

¹⁹ Formally, it is sufficient that $dp/dt \rightarrow \infty$, as $t \rightarrow 0$, where p is the probability of high output and t is the time devoted to current job responsibilities.

We now offer a partial characterization of the expected-profit-maximizing policies²⁰ when the choice of s is not observable and interpret the results. We assume that the incentive problem is to keep employees from spending too *much* time on influence activities, just as it would be if the firm tried to implement the first-best policies.²¹

First, it is easily shown that the expected earnings from being assigned to the key job are no greater than they would be under the first-best and that they weakly exceed the expected earnings associated with the non-key job.²² Thus, the key job is still at least as attractive as the non-key one because it pays at least as well. Also, because w_N still must be at least as large as w^0 , the relative attractiveness of the key job is no greater than that of the first-best. Moreover, the optimal policy never increases and sometimes reduces wage disparities (in expectation) relative to the first-best.

Reducing wage disparities reduces the relative desirability of the key job and the incentives to build credentials, and we expect there to be some closing of this gap. Thus, an apparent organizational concern with pay equity can be explained as being derived solely from an efficiency concern with controlling costly influence activities.

If the expected incomes in the two jobs are in fact equalized, then employees do not care about the job assignment and will have no reason to distort their effort allocations. In this case, the rest of the organization's policies would be set as under full observability: no performance bonuses would be paid, employees' information would always be solicited ($\beta = 0$), all promotions would be on the basis of qualifications ($\gamma = 0$), and the employees would be asked to set s and t at their first-best levels.

If, however, the losses suffered when an employee quits are sufficiently high in very important key jobs, (i.e., $R(C)$ is large for C close to one), then completely solving the incentive problem by equating expected earnings across jobs is too expensive: either too high a level of turnover is

²⁰ Within the class of policies delineated above.

²¹ It is theoretically possible that the cheapest ways to mute employees' incentives for excessive influence activities go too far, leading to too few attempts at influence and too little information being provided for decision making. Our analysis assumes that this possibility does not arise. For example, the analysis applies when the optimum calls for the employees to spend no time at influence activities but simply to report whatever information arises in the course of doing their jobs.

²² These expected earnings include what one receives by leaving the firm to take a better outside offer. Formally, the expected earnings associated with a wage of w are defined as the sum of two terms. The first term is the given wage, w , times the probability, $G(w)$, that the employee gets no outside offer exceeding w and so stays with the firm, receiving this wage. The second term is equal to the probability of receiving a better offer and thus quitting, $1 - G(w)$, times the expected value of the wage in the outside offer, given that it exceeds w .

experienced, or non-key employees are “grossly” overpaid. Thus, “compensating differentials” (equal wages) are not paid—even though they could be—and the key job will be strictly more attractive. In this case, at least some of the other policy instruments will deviate from their first-best levels to control the induced incentive problem.

As suggested earlier, however, setting α at its first-best level, zero, is in fact optimal: if α were positive, so that some promotions were made purely on a random basis, reducing it (promoting more often on the basis of performance) would improve performance incentives, which would then allow the firm to adjust one or more of its other control variables in an advantageous fashion. For example, it might adjust w_N and $w_K(C)$ back toward their first-best levels, thereby reducing labor or turnover costs.

The most striking result is that it can never be optimal to use only performance bonuses to overcome the incentive problem; it can never be optimal to set $b > 0$ and $\alpha = \beta = \gamma = 0$. We have already seen that, if earnings are equalized, then b is set at zero because there is no incentive problem. Suppose then that the compensation policy makes the key job desirable. Increasing b has a marginal cost to the firm that is bounded away from zero.²³ On the other hand, because the expected gain from making superior appointments to key jobs of minimal importance is zero, the marginal cost of providing incentives through promotions based on performance is zero when $\beta = \gamma = 0$. If the only motivation for allocating time to current job responsibilities, as desired by the firm, is the possibility of receiving a bonus, the same performance incentives can be achieved by increasing γ (and perhaps β) from zero, so that promotions to the high-paying job reward performance while lowering b . This shift saves the firm money while maintaining incentives. Thus, adjusting promotion (and possibly decision-participation) policy is more effective than using cash payments alone. Here, the structure of the organization is a central policy element.

Further, it is worth noting explicitly that, if bonuses are optimally set at zero, then wages will be set at their first-best levels, $w_N = w^0$ and $w_K(C) = W(C)$, with *all* the incentives being provided through promotion and decision-participation policies. The reason is that, if b is zero, then the objective function of the employees is of the form $p(s) \times \Delta$, where p is the probability of getting the key job as a function of the time spent on building credentials and Δ is the expected gain in income from getting the key job, which is independent of s . If this gain is strictly positive,²⁴ the

²³ If p is the equilibrium probability of a given employee producing high output, the expected cost of bonuses is $2pb$.

²⁴ As it will be if quits from the key job are very costly.

optimal level of s for the employee is independent of the magnitude of the gain. Thus, the firm will set $w_K(C)$ and w_N at their first-best levels because this minimizes the costs of wages and turnovers without affecting incentives.

In this case, with $b = 0$, $w_N = w^0$, and $w_K(C) = W(C)$, the firm will set β at a strictly positive level, so that useful information is sometimes not solicited and some promotions are made first on performance and then on other, apparently random factors, all of which have no connection to whether the promoted employee is well qualified for the key job.²⁵ Also, γ will exceed β , so that there is a further range over which performance is given precedence over qualifications in promotions.²⁶ If the expected pay-off from appointing the best-qualified candidate gets very large as the job importance increases, then γ will be less than one: for very important key jobs, qualifications become the sole criterion for promotion (as they are in the first-best). However, if this condition does not hold, then it is possible that γ is set at one, and the organization is always prepared to promote on criteria that have nothing to do with the expected productivity of the employee in the key job.

If, on the other hand, performance bonuses are paid, then it is still the case that β and γ will both be set at strictly positive levels, with $\beta < \gamma$. Also, the organization will strictly reduce the disparity between the earnings in key and non-key jobs relative to the first-best, with $w_N > w^0$ and $w_K(C) < W(C)$ for all C . Thus, the organization uses all the instruments available to it to provide incentives and reduce influence activities. The logic here is that paying bonuses has a marginal cost that is strictly positive, while changing the other controls slightly from their first-best levels is, at the margin, free. Thus, if bonuses are actually used, so too will all the other instruments.

A central question thus becomes, Under what conditions will perfor-

²⁵ The Peter Principle is thus rationalized if we think of a hierarchy of key jobs of increasing importance, with the corresponding γ values going to zero as the costs of promoting an unqualified person increase.

²⁶ The proof that γ exceeds β is quite subtle. The essential idea involves recognizing that a high level of q is evidence not only that q is large but also that s may have been set too high. Since the firm's problem involves controlling the incentives to overinvest in building credentials, it is costly (in terms of incentives) to promote on qualifications rather than on a random basis and even more so when credentials are used instead of performance. Setting β optimally involves trading off the gain from appointing the better-qualified person to the key job against the incentive costs of promoting on credentials rather than randomly when the employees' performance levels are equal. Setting γ involves the same sorts of benefits but also the higher costs of ignoring the evidence of cheating on s (as well as the evidence of not having cheated that is provided by having higher performance). These higher costs will be incurred only if the benefits justify them, which means that the gains from appointing the better-qualified candidate must be greater at γ than at β .

mance bonuses be used? We cannot yet give a full characterization of these conditions, but we can indicate two factors. First, if large amounts of information about q are generated in the normal course of events and, as a result, the estimate of q is unlikely to be much affected by increasing s , then bonuses will not be used. The idea is that the free information about q is likely to be very complete, and estimated productivity is consequently relatively insensitive to extra time spent developing credentials. Thus, there is little chance for any one employee that increasing s from zero will be critical in improving his qualifications enough to make them appear best. Consequently, there is little incentive problem, and what incentives there are to indulge in influence activities can be overcome easily by small increases in either bonuses or the β and γ parameters from their first-best level of zero. However, as noted earlier, marginal increases of β and γ from zero are free, while the marginal costs of increasing b are strictly positive. Thus, all the incentives will be provided through the design of the promotion policies.

The second case in which bonuses will not be paid arises when $\partial\pi/\partial q$ varies sufficiently little with I over a broad range of (low) values of I . In this case, the possible effect on profit of a degradation in the quality of the job assignment (by raising γ) is slight over a broad range. In particular, the marginal cost of raising γ over this range is less than that of raising b . Thus, incentives can be provided most cheaply by using promotion criteria again.

APPLICATIONS

In this section, we briefly discuss three applications of the ideas associated with the model in the previous section. The first of these, which concerns an interpretation of some of the stylized facts regarding differences between Japanese and U.S. management practices, relates quite directly to our formal model. The other two build on prominent features of the organizational response to influence activities identified in the solution in the previous section. Specifically, we interpret certain characteristics of Personnel Departments' behavior and organization in terms of the organization's committing itself to ignoring decision-relevant information, and we use the idea of compression of wage differentials as a response to influence activities to offer some predictions regarding the compensation policy of fast-growing firms.

Among the major stylized facts about the environment, policies, and organization of the large Japanese firms that are the focus of most of the attention paid to Japanese management are the following:²⁷

²⁷ See Abegglen and Stalk (1985).

- The opportunities for improving one's position by changing employers are relatively limited, at least for (both blue-collar and white-collar) permanent employees of the large firms that offer lifetime employment to those joining the firm at the start of their careers. A permanent employee of one large company is very unlikely to be able to move to another large company that offers comparable pay, job security, fringe benefits, and so forth, except as a temporary employee.²⁸
- The differential in earnings between the highest- and lowest-paid permanent employees of large Japanese firms is much smaller than the corresponding differential in the United States.
- Decision making in Japanese firms tends to be much more participative than in the United States, with all concerned being allowed to have their say before a decision is reached.

Our model suggests that these factors are crucially interrelated. In particular, suppose that in our model there were actually no chance of getting an outside offer at a wage exceeding w^0 , the amount that is available to anyone in the "secondary" outside labor market. In that case, the firm need not worry about turnover in the key job. Consequently, it can set $w_K(C)$ equal to w_N without inducing costly quits. As a result, there would be no incentive problem, and the firm can then permit the employees to provide information in all circumstances without fear of inducing influence activities and the associated costs.

Further, if we "compress" the distribution, G , of possible offers so that the probability of getting an offer above w^0 becomes small, then we would expect the firm to be increasingly less concerned with quits because these are less and less likely. Correspondingly, we would expect the differential between the pay in key and non-key jobs to narrow. As long as there is some direct reward to good performance in the current job, so that employees are not simply trying to maximize their probabilities of being promoted, this narrowing will reduce the incentive problem. As a result, the firm can give greater attention to employees' information in its decision making (reduce β and γ). Thus, the three stylized facts fit together: the weaker outside labor market in Japan relative to that in the United States explains the narrower wage differentials and the greater involvement in decision making. Incidentally, the greater use of employees' information ought to lead to better decisions, which would then be manifested in better firm performance in Japan than in the United States.

This analysis, if correct, suggests that a piecemeal approach of adopt-

²⁸ Companies foster this attachment by encouraging employees to live and socialize together and avoid contact with competitors.

ing various Japanese management practices in U.S. firms may be dangerously counterproductive. The differences between U.S. and Japanese firms' policies may reflect appropriately differing responses to the economic environments in the two countries. If so, then, for example, adopting Japanese participative decision making in the United States may lead to great increases in influence costs that are avoided in Japan because of the narrower wage differentials there.

This analysis does not directly address one of the other most frequently noted differences between Japanese and U.S. employment policies—the practice in Japan of promoting more slowly and much more on the basis of seniority than in the United States. However, this feature too can be in part interpreted in terms of influence costs and the relatively less developed outside labor market in Japan.

Unlike American firms, Japanese firms need not fear losing bright young employees to other firms. Thus, they need not promote quickly in response to outside pressure. In this context, promoting on seniority reduces the payoff to influence activity because promotion decisions are made more on standardized, nonmanipulable criteria. The long period between promotion decisions also tends to reduce the payoff to influence activities, at least when the decision is still far off, because the rewards to the activity are deferred. However, the opposite effect might obtain when the decision is near at hand because failure to be promoted at this round means that there will not soon be another chance. One student of Japanese management has indicated to us that it is his impression that the pattern of behavior suggested by this analysis in fact prevails: there is little influence activity practiced by people in the early stages of their careers, but, when the promotion decision approaches, the amount of politicking increases dramatically (Okuno-Fujiwara, personal communication). It would be interesting to see if this impression is borne out by careful observation.

A key feature of the solution developed in the last section is that the firm's decision makers commit themselves to ignoring valuable information and to making decisions that they know are "wrong" in order to limit influence activities. When I is less than β , the information that employees have about their qualifications for the key job is not collected, and, when I is between β and γ , this information is solicited but is ignored if the candidates' performance levels differ.

In general, one would expect it to be difficult to ignore decision-relevant information, especially once one has been provided with it; think of the difficulty of appointing one worker to a key job as a reward for past performance when you know both that this worker is grossly unqualified for the job and that there is another worker who is superbly qualified and

whose lower level of past performance was almost surely simply a matter of relatively bad luck.²⁹ Of course, if the commitment to ignoring information is not credible, then the policy described above will no longer provide the appropriate incentives because workers will expect that promotions will actually be made on the basis of qualifications, no matter what the official policy might say. Consequently, they will devote too much time to building their credentials.

One way to achieve this commitment is to take the promotion decision out of the hands of decision makers who have an interest in making the best possible assignment and turn it over to someone whose interests are in following the stated policy. This observation may provide a rationalization of some features of the behavior of Personnel Departments that have been frequently noted.

In every organization with which we have been associated, and in most of those of which we have heard, the Personnel Department is viewed by line managers and employees as unresponsive, rule-bound, and bureaucratic. It takes forever to get a decision from Personnel, and the decisions seem aimed more at maintaining the Personnel Department's precious rules, procedures, and job classification/earnings and experience/earnings curves than at attracting, rewarding, and retaining the best people for the organization. Moreover, protests fall on deaf ears: Personnel people are always in meetings when you try to reach them, and they do not return your calls.

Our analysis suggests that this may, in fact, be efficient. If the Personnel Department were routinely responsive to information about employees' special circumstances, talents, and opportunities and to managers' requests for raises for their people or for expedited decisions on hiring, promotions, or termination, then there would be much greater return to influence activities. Managers would spend much more of their time trying to get exceptions and special deals for their people, and the cost would be diversion of their attention from more productive activities. Moreover, if managers were able to get Personnel to respond to their requests, lower-level employees would have an incentive to spend their time trying to influence managers. However, under current arrangements, managers do not waste their time trying to influence Personnel Department decisions because it is usually pointless. Correspondingly, lower-level employees do not devote as much time to attempting to get managers to exert influence on their behalves, as they otherwise would, because the managers have

²⁹ Performance differences in our model are purely a reflection of random chance because the policies of the firm, by design, provided sufficient incentives to ensure that workers will not slight their responsibilities in their current jobs and because the employees do not differ in their productivities in the key job.

no means of influencing Personnel. Only if Personnel's decision is egregious is it worthwhile to attempt to reverse it, so only in the most important cases is influence attempted and the concomitant cost incurred.

Our final application relates to the narrowing of wage differentials as a means of controlling influence costs that featured in the solution in the last section. It leads to a prediction that high-growth firms should have narrower disparities in compensation among jobs than do more stable organizations and that they should also have higher overall levels of compensation.

When the environment in which an organization operates is changing rapidly, or when growth means that the organization itself is undergoing rapid change, the organization will have frequent and, perhaps, especially intense needs for access to information for decision making. In such circumstances, influence activities can be particularly costly because the frequent need to rely on information from interested parties in the firm means that the opportunities to practice influence activities are especially rich. Just as in our model, narrowing compensation differentials helps overcome these difficulties, because it means that employees are more nearly indifferent to the various options from which the firm must choose. With less personally at stake, employees are less likely to indulge in costly influence activities. Thus, to overcome the especially intense problems of influence activities to which high-growth firms would seem to be subject, we would expect them to narrow compensation disparities among jobs more than do other firms.³⁰

More broadly, any change that has the potential to shift power and rewards in the organization may lead to influence activities, with potential losers resisting the change by whatever means they have available, and potential winners pushing for the change. Note that the incentives for such behavior will depend to a large extent on the redistributive effects of the change and that these need bear no necessary relation to its efficiency effects. Again, reducing the differences in rewards among different jobs in the organization should reduce the incentives to attempt to promote or combat changes because of their redistributive effects. In a high-growth firm, this is especially easy: changes typically enlarge the pie, so narrow differentials can be obtained by giving more to the winners without taking from the losers in a particular change. This approach will, of course, also lead to relatively high levels of compensation.

In situations where growth does not permit everyone to gain from changes, then influence activities in the form of resistance to change

³⁰ More precisely, compensating differentials in such firms should be more nearly complete. However, to the extent that nonmonetary job characteristics do not vary significantly within a firm, this would show up as a narrowing of earnings disparities.

ought to be especially significant. This leads then to a further prediction that, in general, one ought to observe a narrowing of differences in wages and other compensation in firms as the pace of environmental change increases and, with it, the need to make organizational changes that carry distributional implications.

CONCLUDING COMMENTS

We have shown how influence activities—politicking—in organizations can be incorporated into an economic model and that there may be efficiency considerations that give rise to these activities and that follow from them. We have also shown how, in a particular context, the organization's policies and structure will be adapted to the possibility of influence activities. These adaptations, and the motivation behind them, seem (to us) to offer some new insights into aspects of organizational behavior.

Of course, our model has been very “bare-bones” and could be enriched in many ways. Some directions are suggested by organization theory and others by economics; both should be pursued, and the robustness of our results to such enrichments explored. Moreover, our analysis of even this model is incomplete because we have not derived the optimum among all possible decision-participation and promotion policies but only for a restricted set.

Despite the incompleteness of this work, we hope that it will prove useful in creating greater contact between sociologists and economists. If it has this effect, then the attempt at influence that this paper itself represents will have been successful.

REFERENCES

- Abegglen, James C., and George Stalk, Jr. 1985. *Kaisha: The Japanese Corporation*. New York: Norton.
- Arrow, Kenneth J. 1985. “The Economics of Agency.” Pp. 37–51 in *Principals and Agents: The Structure of Business*, edited by John Pratt and Richard Zeckhauser. Boston: Harvard Business School Press.
- Baron, David P. 1983. “Tender Offers and Management Resistance.” *Journal of Finance* 38 (2): 331–43.
- Bendor, Jonathan, Serge Taylor, and Roland van Gaalen. 1986. “Stacking the Deck: Bureaucratic Missions and the Search for Alternatives.” *American Political Science Review* 81 (3): 87–96.
- Bohn, Henning. 1987. “Monitoring Multiple Agents: The Role of Hierarchies.” *Journal of Economic Behavior and Organization* 8 (2): 279–306.
- Fama, Eugene F., and Michael C. Jensen. 1983. “Agency Problems and Residual Claims.” *Journal of Law and Economics* 26 (2): 327–49.
- Geanakoplos, John, and Paul Milgrom. 1985. “A Theory of Hierarchies Based on

- Limited Managerial Attention." Discussion paper, Yale University, Cowles Foundation.
- Grossman, Sanford, and Oliver Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy* 94 (4): 691–719.
- Hart, Oliver, and Bengt Holmstrom. 1987. "The Theory of Contracts." Pp. 71–155 in *Advances in Economic Theory—Fifth World Congress* (Economic Society Monograph no. 12), edited by Truman Bewley. Cambridge: Cambridge University Press.
- Holmstrom, Bengt. 1982. "Managerial Incentive Problems—A Dynamic Perspective." Pp. 209–30 in *Essays in Economics and Management in Honor of Lars Wahlbeck*, edited by B. Walross. Helsinki: Swedish School of Economics.
- Holmstrom, Bengt, and Joan Ricart i Costa. 1986. "Managerial Incentives and Capital Management." *Quarterly Journal of Economics* 101 (4): 835–60.
- Holmstrom, Bengt, and Jean Tirole. 1987. "The Theory of the Firm." In *The Handbook of Industrial Organization*, edited by Richard Schmalensee and Robert Willig. Amsterdam: North-Holland. (In press.)
- Madison, Dan L., Robert W. Allen, Lyman W. Porter, Patricia A. Renwick, and Bronston T. Mayes. 1980. "Organizational Politics: An Exploration of Managers' Perceptions." *Human Resources* 33 (2): 79–100.
- Milgrom, Paul. 1988. "Employment Contracts, Influence Activities and Organization Design." *Journal of Political Economy* 96 (1): 42–60.
- Milgrom, Paul, and John Roberts. 1982. "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis." *Econometrica* 50 (2): 443–59.
- . 1986. "Relying on the Information of Interested Parties." *Rand Journal of Economics* 17 (1): 18–32.
- . 1987. "Bargaining and Influence Costs and the Organization of Economic Activity." Research paper 934, Stanford University, Graduate School of Business.
- Rosen, Sherwin. 1982. "Authority, Control and the Distribution of Earnings." *Bell Journal of Economics* 13:311–23.
- . 1985. "Implicit Contracts—A Survey." *Journal of Economic Literature* 23 (3): 1144–75.
- Walder, Andrew. 1986. *Communist Neo-Traditionalism: Work and Authority in Chinese Industry*. Berkeley and Los Angeles: University of California Press.
- Williamson, Oliver E. 1985. *The Economic Institutions of Capitalism*. New York: Free Press.