I E S E

University of Navarra

CONTROL AND INCENTIVES IN ORGANIZATIONAL DESIGN

Joan E. Ricart*
Josep M. Rosanas**

RESEARCH PAPER Nº 285
January, 1995

* Professor of General Management, IESE
** Professor of Control, IESE

Research Division
IESE
University of Navarra
Av. Pearson, 21
08034 Barcelona - Spain

CONTROL AND INCENTIVES IN ORGANIZATIONAL DESIGN

I. Introduction

Agency theory deals with the problem of motivating one person to act on behalf of another. The typical agency theory model assumes that (i) the argument of the principal's utility function is «profit», and (ii) the arguments of the agent's utility function are his compensation and «effort».

These assumptions are very useful in some contexts. For instance, when interpreting the principal as a purchaser of specialized services supplied by the agent; or when interpreting the principal as being the «owner» of a business firm, and the agent as being the «manager». In other situations, however, they are much less realistic.

Specifically, in an organizational context, there are a lot of principal-agent relationships at different echelons of the «chain of command». The Chief Executive Officer is the principal to certain agent(s), but these agents are themselves principals to other agents one echelon below, who, in turn, are principals to someone else, and so on. In these situations, the «principals» at the different levels are no more interested (directly) in the «profit» figure than their agents, with the possible exception of the Chief Executive Officer herself. Indirectly, of course, they may be interested in «profit» precisely because of the incentive effects of the compensation payment made to them. But, in general, any «principal» anywhere in the organizational ladder is going to be interested in a number of variables, which may range from the compensation payment she receives, to «pride in a job well done» or «loyalty to the organization». Agents, of course, are interested in similar variables, but also (very explicitly) in the compensations they receive, besides wanting perhaps to avoid «effort».

This paper is an attempt to setup a more general formulation of the principal-agent model that can accommodate these organizational situations, which are rather common in the world of business management, for the purpose of deriving useful propositions both about organizational design and about the process of influencing the decisions made by agents on behalf of principals.

We will proceed as follows. In the next section we will describe the process of delegation of authority, and structure it along the lines of agency theory. Section III will examine the information requirements for decision-making, and their implications in the organizational context. In Section IV we establish the formal model of the problems we want to study; and section V will explore the implications and consequences of such a model.

II. Utilities, evaluation, incentives

We focus, then, on the agency relationship in the organizational context, at the higher levels of management, where the basic financial variables are directly related to performance. The principal may be, for instance, the General Manager of Manufacturing, and the agent, a Plant Manager; or the principal may be a Divisional Manager, and the agent the Divisional Marketing Manager.

As stated above, in these contexts the principal may be no more interested in profit «per se» (or in any other measure of economic success) than the agent, because neither is entitled to receive any part of it, except perhaps through the incentive systems. The principal, however, may be interested in profit as a guide for action, because she likes «well-done» jobs, and thinks that a measure of achievement should include such a variable.

In general, principals at higher levels of the hierarchy may be interested in different variables. Surely, it is realistic to assume that they are interested in their compensation, possibly based on their performance, which (of course) includes the performance of their subordinates. Also, «loyalty to the organization», for instance, which would, in turn, imply adherence and commitment to the organization's various goals (e.g., «profitability», «personnel development», «product leadership», «good labor relations», «organizational climate», «public responsibilities» (1), and so on), which would thus enter as arguments of the principal's utility function. Finally, «pride in well-done jobs», or some beliefs about how jobs should be done, may lead a principal to have utility for some kinds of activities or even specific actions that the agent may undertake to achieve the organizational purpose.

The agent's preferences can be sketched similarly. Again, the agent is interested in his compensation (which is one of the essential variables to be analyzed in our inquiry), and in values similar to those of the principal, although possibly with different intensity, if only because one would expect the agent's utility function to be more sensitive than that of the principal to the last kind of variable (i.e., those that have to do with the actions and activities undertaken by the agent).

Some of the variables entering the utility functions of principals and agents may be very difficult to measure objectively. Non-quantitative business goals like some of those mentioned above (e.g., «product leadership», «organizational climate», or «public responsibilities») cannot be measured with a high degree of precision. Therefore, such measurements: (i) are likely to be noisy (i.e., random factors will influence the measure), and (ii) can be influenced by the agent by means of actions that do not contribute in any way to actually achieving organizational objectives and goals.

In the usual agency theory models, the possible conflicts between principal and agent arise because of two reasons:

1) the payment to the agent decreases the residual profit of the principal, and therefore his utility, while obviously increasing the utility of the agent, and

2) the agent's utility function is assumed to be a decreasing function of «effort». The profit function, on the other hand, is the only argument of the principal's utility function, and it increases with «effort».

In our analysis, in contrast, and since we allow for several dimensions, the preferences of the principal and of the agent in some of these dimensions may be perfectly

aligned. For instance, if the agent likes doing a job «well done», his utility will increase with the quality of the product, which may coincide with the interest or the preferences of the principal. Obviously, this is not going to happen with all variables, and in some others there may be a straight conflict, as we will see.

III. Information

Agency costs also originate in asymmetry of information. Usually, it is assumed that the agent is better informed than the principal, because of his closer contact with the problem. This is generally a realistic assumption, but it is also true that the principal has superior information in some other aspects. In fact, the two kinds of information are complementary, as we will show next.

Two kinds of information are needed for decision-making. One is knowledge of the specific circumstances surrounding the decision, i.e., the initial conditions that the decision-maker has to face and that can be changed in order to obtain the desired result. But a different kind of information is also needed: the general statements about how the real world works, and its reaction to the initial conditions introduced, which makes it possible to obtain the result. «A priori», this knowledge is crucial to predict the state of affairs that will obtain if some initial conditions are introduced. Popper (1957) has called this kind of prediction «technological», those that «form a basis of engineering. They are, so to speak, constructive, intimating the steps open to us if we want to achieve certain results».

When there is a well-established body of empirical science applying to a real-world problem, the technological prediction should logically follow by deduction from a combination of the initial conditions and the appropriate empirical laws. Normative reasoning, i.e., the derivation of technological predictions, will therefore follow this schema (Christenson, 1983):

To Find:	(1)	$L_1, L_2, \ldots L_n$,	(Empirical Laws)
	(2)	$C_1, C_2, \ldots C_n$,	(Controllable Initial Conditions)
Given:		S	(Desired Final State)

Of course, for that purpose, knowledge of the initial conditions that actually prevail has to be (close to) perfect as well.

Under these conditions, decision-making is limited to choosing between desired states of the world and introducing the necessary initial conditions so that the desired state is obtained.

In most business problems, though, such a body of well-established empirical science does not exist, or exists only partially. The «general» knowledge needed to make the decision must therefore originate in the «expertise» of the decision-maker. This expertise can thus be considered «pre-scientific» knowledge, or knowledge that has not passed certain authentication procedures.

A «perfect» knowledge of the initial conditions is also rare. To be sure, knowledge of some of the variables included will, in general, be easy to obtain and transmit, and most of these are probably measurable. But other variables are more difficult to evaluate: some aspects are highly qualitative and depend very much on the personal point of view.

For instance, in a production setting, the number of machines available, the number of hours that they can work in practice, and the machine-time needed for each unit of product are typically variables that are easy to obtain, to measure and to transmit to other people, if only because they can be summarized in numbers. On the other hand, circumstances that are also relevant for possible production decisions, such as the relative ability of different individuals to solve problems, the organizational climate in the factory, the specific arrangements that can be made to accommodate a special order, and so on, can only be assessed in a highly subjective way. To some extent, one could think of these as unauthenticated knowledge (because it is subjective) of the initial conditions.

Hence, we have in summary four kinds of information in decision-making:

1) Scientific knowledge about the way the world works. This is, generally speaking, a type of knowledge that is very difficult and costly to generate, but that can be authenticated and transmitted at a reasonable cost.

2) Unauthenticated knowledge about the way the real world works («expertise»). It is usually costly to obtain, and almost as costly to transmit.

3) «Objective» knowledge about the way the real world is, which can typically be summarized in figures and statistics.

4) «Subjective» knowledge about the way the real world is, which is not very costly to obtain, but has a very high cost of transmission, partly because it cannot be authenticated.

Hayek (1945) was a pioneer in recognizing that part of the knowledge needed for decision-making at the level of the individual firm is detailed information about «the particular circumstances of time and place», which is obtained in the daily practice of business (p. 521). An example would be the knowledge that there is an under-used machine that can be put to work, or that a certain person has particular skills that can be put to better use, or that there is a surplus stock which can be drawn upon if supply is interrupted. This «sort of knowledge …/… is knowledge of the kind which by its nature cannot enter into statistics and therefore cannot be conveyed to any central authority in statistical form» (p. 524).

Hayek's purpose was precisely to show that, because of this last point, the economic problem of society (which is defined by him as one of rapid adaptation to changes in specific circumstances) cannot be solved by any central authority with decision-making power. Instead, it must be solved by some form of decentralization, with a minimum communication of information.

Jensen and Meckling (1992) distinguish between knowledge that is costly to transmit («specific knowledge») and knowledge that is relatively inexpensive to transmit («general knowledge»). «Specific knowledge», according to these authors' classification, includes «idiosyncratic knowledge», which is (at least in part) the knowledge of «time and place» that Hayek had in mind; however, «specific knowledge» also includes scientific knowledge.

Jensen and Meckling use this concept for the purpose of studying organizational design and the allocation of decision rights within the organization. What they do is, in some sense, a corollary of Hayek's conclusion: concentrating all relevant information in one place

(presumably at the top) is not only impossible for the economy as a whole, but within organizations as well. Hence, from this point of view, decision-making should be placed as close as possible to where the information is. Of course, an aspect that works against this is the fact that the objectives of the different decision-makers within the organization may be inconsistent. Therefore, there are trade-offs between the two criteria that should be weighted. We will analyze this point in greater detail below.

Our four-way distinction is intended to build on the concepts of Hayek and Jensen and Meckling by explicitly including the possibility of having differences of information that are not related to the specific circumstances of time and place, but to the different decision models that different individuals in the organization may have. Specifically, we are interested in the possibility that the principal, while having less «time and place» information than the agent regarding the specific problem to be solved, is nevertheless better informed than the agent in other respects; for example, he may be better equipped to solve the kind of problem the agent is dealing with because of his «experience». As we have shown, part of this knowledge is surely scientific, but a great deal of it is not, in the sense that it has not passed any authentication procedure. Therefore, not only is it difficult to transmit, but the credibility to the transmitter may be rather low.

For instance, a General Manager of Manufacturing may have under her several Plant Managers, who, according to our analysis, will be better informed that she is about the specific circumstances of time and place in each plant. However, it may be the case (indeed, it very often is) that the General Manager is more experienced than the Plant Managers, perhaps because she has been a Plant Manager herself for a number of years before being promoted to her current job. To be «more experienced» is, of course, to have certain information that the Plant Managers do not, but this information is not about specific circumstances of time and place, but about «better» decision-models. Thus, it is not inconsistent to say that both the principal and the agent may have information that the other does not.

It is particularly important for our purposes to distinguish between «authenticated» and «unauthenticated» information (2). Both may be costly to transmit, but authenticated information is, by definition, credible, whereas unauthenticated information might not be, depending on the mutual respect and consideration of the two participants. This, we will argue, introduces agency costs, even if there is a good alignment between the objectives of the principal and those of the agent.

Hayek's analysis contains an additional consideration that is important in this context, although it is sometimes overlooked. It is that «every individual has some advantage over all others in that he possesses unique information of which beneficial use can be made, but of which use can be made only if the decisions depending on it are left to him or are made with his active cooperation» (p. 521). We will argue that, where it is not optimal to leave the decision to the agent, it is of the greatest importance to obtain his «active cooperation» if the information of «time and place» is to be fully utilized.

IV. The model

In this section we will present a tentative model that represents many of the concerns raised in the previous sections. The model itself is an elaboration of the classical approach to agency relationships (Holmstrom, 1979), with the necessary additions. In this way we try to connect with the prevalent tradition in organizational economics.

The agent's action choice will be represented by two variables, which we shall call «a» and «b». These two variables belong to the sets A and B, respectively. The first component, «a», is the traditional agent's decision. However, it may not necessarily be interpreted as an effort variable. It may have an effort content, or it may simply be a managerial decision, such as an investment decision (see Holmstrom and Ricart i Costa (1986) for an example of such a decision). Whatever its meaning, «a» is the relevant decision in the organization, the one that justifies the principal-agent relationship.

There is, of course, an implicit assumption in the above comments. Given our concern for organizational design, one would have to justify the existence of the specific organizational arrangement. It is assumed that the agent is needed either because he has to exert some effort or take care of an activity (e.g., looking for alternative decisions), or else because he possesses or can generate a particular knowledge that is very difficult to transmit in a credible way. Therefore, an organizational arrangement where the principal tries to elicit the information from the agent and takes the decision himself is ruled out by our supposition (3).

The second component of the action, «b», is of a different nature. The company is assumed to have a measurement system (to be described later on), which generates the observable variables that are relevant for the relationship. Let them be represented by the n-tuple $m = (m_1, m_2, ..., m_n)$. However, the agent may in some way manipulate or influence the value of the observations obtained by the system. That is, m is a function of b and other variables to be defined later on.

Similarly, the principal also makes two choices. One corresponds to the classical compensation system for the agent. The agent's compensation will be contingent on all observable variables and is represented by s(m). It is assumed to be monetary compensation. The second choice, represented by p, tries to model the principal's ability to affect the agent's perceptions about how the world works. We will be more precise on this later on.

The state of the world is also represented by two variables. We do this to simplify the exposition, since any of our variables could perfectly well be multi-dimensional. It will be denoted by $ø = (ø_1, ø_2)$, with an ex-ante probability distribution given by $G(ø_1, ø_2)$. The supposition behind this duality is that the first variable represents *initial conditions* about which the agent would have imperfect information, while the second component represents the *empirical laws* used by the principal.

The results of the agent's action will be denoted by $x = (x_1, x_2, ..., x_n)$. The value of x is given by the traditional representation: $x = X(a, ø)$.

However, x is multi-dimensional. One of the dimensions may represent the monetary result, but the other variables may be aspects as soft as «product leadership», «organizational climate», «a job well done», etc.; all consequences of the action. Our protagonists in this story care about almost everything.

Now we have to build the information structure into this model. We have two types of information: public and private. Public information is the information generated by the formal measurement system in the company. Obviously, this measurement system is imperfect and may be very noisy. It is intended to represent the accounting and reporting system in any organization. We denote it by:

$(m_1, m_2, ..., m_n) = I(x_1, x_2, ..., x_n, b, §)$.

This means that the observable variables depend on the real values of x, the manipulative action of the agent, b, and some noise, §, that follows a probability distribution F(.).

But both parties also have some private information. We assume that the agent observes a noisy signal y_1 of $ø_1$, while the principal observes a noisy signal y_2 of $ø_2$. Given that each party has some private information about the state of the world, the corresponding posterior distributions, after observing their corresponding signals, will obviously be different. Therefore, both parties will use different expectations in their decision-making, and furthermore, in a repeated model, they will be acting with incomplete information because the most they can do is try to infer what information the other party has. They will be playing a complex game of incomplete information.

Now that all the elements in play in this model have been defined, it is time to set up the corresponding decision problems for each party. Let us start with the agent. The agent's choice was represented by the two variables (a, b). His utility will obviously depend on these variables, plus all public information, m, the compensation system, s(m), and his subjective perception of the relevant variables x. We mentioned before that the principal has some means of influencing such perceptions. Therefore, we will denote this process by defining the agent's perceptions about the world by the n-variable function:

$$\pi = (\pi_1, \pi_2, ..., \pi_n) = \prod(x,p) = \prod(x_1, x_2, ..., x_n,p).$$

The agent will choose a and b to «optimize» his expected utility, which is given by:

$$E|_{y1} [U(a, b, m, \pi, s(m)].$$

The traditional model assumes that the agent will maximize this expression by choosing a and b in the corresponding sets A and B. However, we left this process unspecified. For many managerial decisions, it is not the choice that it is relevant, it is the search for alternative courses of action that will improve the current situation. Realistically, A and B are unknown sets that the agent «discovers» gradually over time, thus improving his decision-making.

Similarly, the principal’s choices, p and s(m), are made in such a way as to optimize:

$$E|_{y2} [V(m, x, s(m)]$$

where, in fact, one can interpret the result n-tuple x as the principal's perceptions of these variables. Again the optimization problem is left undefined, for the same reasons. We insist, however, on the fact that the model is one of incomplete information.

This is obviously a double optimization problem, as usual. Traditionally, it is assumed that the principal chooses the compensation scheme s(m) first, and then, knowing this choice, the agent chooses his action. Therefore, the principal must anticipate the agent’s choice for each possible compensation scheme. Thus, Grossman and Hart (1983) break this problem down into two parts. First they calculate the minimum cost for the agent of «implementing» any action, and then they calculate which «implemented action» maximizes the principal's utility.

Our reasoning is essentially the same but in an incomplete information world. The principal makes his choice about the compensation scheme. However, in doing so he is

unable to anticipate any of the agent's actions because of the incomplete information and the complex nature of the perceptions involved. Furthermore, the activities designed to affect the agent's perception of x will be undertaken once x has been realized and, again, full anticipation is very difficult.

Similarly, the agent chooses a and b knowing s(m), but recognizing that the principal has private information and is able to influence his choices once x and y are realized. Furthermore, he has to search for possible actions and learn about their effect. All together the problem is totally undefined.

Our inability to solve the model does not mean that it cannot be useful. We did not set up the model to prove any specific theorem, but to explain the nature of agency relationships in organizational design. Let us therefore try to draw some conclusions from our formulation of the relationship.

V. Implications

V.I. Bounded rationality and Learning

The first set of reflections has to do with the different components of the model. Agency relationships inside organizations are characterized by the fact that they involve a multiplicity of elements. One could try to aggregate all these aspects in a single monetary unit. However, the very nature of organizations makes this impossible. Organizations are characterized by the lack of alienability (Jensen and Meckling, 1992). Consequently, the distribution of decision rights in the organization does not run parallel to the measurement and reward systems implied by property and alienability rights in an economy. The exclusive use of monetary measures in organizational design misses a fundamental difference between organizations and markets, and this difference should be reflected in the model.

The second point about the components of the model has to do with the presence of incomplete information. Any agency model is based on the presence of uncertainty. However, perfectly rational agents have no problem dealing with uncertainty. They use subjective probability and update their prior distribution, based on available information, by using Bayes' Rule. Nevertheless, in a world of incomplete information, matters are much more complex. We have to make inferences and define an infinite set of beliefs. Still, perfect rationality will allow us to «solve» the problem by using Bayesian perfect equilibrium (Harsanyi, 1968). However, a minimum of realism will obviously raise the question of bounded rationality. We could assume a human ability to deal with uncertainty (even knowing that we make many mistakes doing so), but it is very difficult to assume an ability to deal efficiently with incomplete information.

Furthermore, matters are even more complex. The nature of managerial action is such that it is not credible to assume that the manager is capable of knowing all the alternative actions possible and the consequences of these actions in a complex set of variables they care for. Again, bounded rationality drives us to much simpler «satisfying» models of trial and error, as well as learning from alternative actions and consequences.

Our model does not describe the learning process. It shows, however, that no static model would ever be able to represent the key elements of organizational design. We need «approximate learning models», and learning is a complex issue, too (Ghemawat and Ricart i

Costa, 1993). By «approximate» we mean that, given the complexity of the model, neither the agent nor the principal is able to «solve his problem». What they could do is set up a learning process in which they act, measure, check, learn and correct. In this way they try to «approximate» a good solution. However, if they learn on the basis of a given set of constraints or beliefs, they may only reach static efficiency. To reach dynamic efficiency they need to optimize by reconsidering all assumptions and beliefs. They should build a «learning organization» (Senge, 1990).

V.II. Knowledge and organizational design

The second conclusion to be drawn from our model concerns the nature of knowledge in organizations. In fact, one of the consequences of the lack of alienability in organizations is that one has to allocate decision rights where the relevant knowledge is located. We have already discussed the fact that some types of knowledge are easier to transfer and communicate than others. Furthermore, the knowledge may be more o less relevant for the decision, depending on the sort of decision it is.

In our model we assumed that the agent has some relevant information about the nature of the initial conditions. This is the knowledge of «time and place», in Hayek's (1945) terms. The sort of organizational design where the agent makes the decision is good if the agent's information is very important and difficult to transfer. This means either that the principal's information about empirical laws is less relevant for the decision or that it can be transferred to the agent, either directly, via the measurement system, because it is easy to code and transmit or because it can be embedded in the influence activities that the principal carries out; or indirectly, via the compensation system.

There are, obviously, alternative organizational designs. For instance, the principal may try to elicit the agent's information and keep the decision for herself. Or both may transmit the information to a third party who then decides. The decision can even be «automated» using the transmitted information. Each of these solutions to the problem of how to assign this decision would entail a different set of information and control problems (Jensen and Meckling, 1992).

The problem of how to allocate decisions is a consequence of the fact that some knowledge is difficult to transfer. We can try to allocate decisions wherever the appropriate knowledge resides, but this is not always possible. The main reason is that we have a trade-off to consider. This is the control or agency problem. Any allocation of decision rights raises an agency problem due to inconsistent objectives, lack of alienability, and the principal’s knowledge of empirical laws. The principal would like the agent to make the decision following a given model, but he cannot impose his will; he can only try to influence the decision. Therefore, for every relevant decision, the design of the organization has to strike a balance between the information and the control problem, aiming for a compromise that is good enough.

As not all information can be transmitted, one has to weigh the cost of making worse decisions due to lack of information (of one kind or another) against the cost of providing the right incentives in a decentralized organization. If the General Manager makes the decision, we shall say that the organization is completely centralized. Alternatively, if the decision is delegated to the unit that holds the relevant specific knowledge, we shall say that it is completely decentralized.

Conceptually speaking, there may be an optimal point of equilibrium between these two extremes that would indicate the right degree of decentralization. This point is the one that minimizes the total cost for the organization, as the marginal cost due to inconsistent objectives or inexperience is equal to the marginal profit due to better information. Besides, the degree of decentralization depends on factors such as the size of the organization, the stability of the environment, the organization's orientation towards innovation, strategy, government regulation, the control technology available, information technology, etc. Additionally, it should be remembered that the way the rules of the game are defined influences not only decisions but also the incentives to produce and transfer information. However, in spite of all these limitations, our analysis allows us to refine our intuition by comparing organizational alternatives.

Such a compromise between the information and control problem is affected by the existing technology, both information and control technology. That is, organization «technology» in general. For instance, if the principal can incorporate his empirical laws in an expert system, the cost of delegating is drastically reduced. Alternatively, if one were able to define a measurement system that transmitted credible «time and place» information on initial conditions, the principal would be able to keep centralized control of the relevant decisions in the organization.

In fact, part of the time and place information might be transmissible through informal contacts between principal and agent.

For each allocation of decision rights, we have a different optimization problem. That is to say that for each organizational design we have a different trade-off between the cost associated with inconsistency in objectives (agency cost) and the cost associated with lack of information. Our model presupposes that due to such trade-offs we want the agent to decide. This is obviously not always the best alternative. It will depend on the type of knowledge required for the decision, and also on the control and incentive systems that one is able to construct for each organizational setup. That is what we study in this model.

V.III. Conflicts of interest in organizations

As we stated in Section II, in the usual formulation of agency problems, possible conflicts between principal and agent, which result in an imperfect alignment between their objectives, arise essentially for two reasons: (i) conflicting economic interests: the more the agent is paid, the higher his utility, and the lower the utility of the principal; and (ii) «profit» to the principal is supposed to increase with «effort», whereas the agent's utility decreases with the same variable.

Not surprisingly, then, the results are mainly about the efficiency of the effort, and the efficiency of risk-sharing. Our formulation can be viewed as a generalization of the previous one, whose results it includes as particular cases where the utility functions have certain specific forms. We can thus reinterpret those results in the light of our model.

(i) Agents with utility for monetary rewards only

To start with, let us assume that the agent's utility function has as its only argument the monetary reward (i.e., there is no disutility for effort, or for anything else). This is

U(s(m)). The utility function of the principal, however, includes other variables (the n-tuple x_1, x_2,…x_n). Then,

(1) Assume that formal incentives are put on the measure of one variable only (say, m_1), perhaps because it is the only one that has a relatively objective measure (i.e., it is the one that has the lowest noise and is least open to influence by the agent's action (b). The agent will then attempt to maximize his expected utility in terms of the expected reward (according to his probability distribution of the variable). The other variables will be regarded by the agent as irrelevant, and he will make no attempt to achieve even a minimum value for them.

At the extreme, we may have $m_1 = x_1$, unaffected by the noise or the manipulation variable b. As $s(m_1)$, the agent cares only about m_1 and will disregard any other consideration in making his choice. Quite obviously, the solution will not be efficient, even if the principal has increasing utility for this variable.

(2) Again, assume now that only one variable, m1, is included in the agent's incentive payment, but this time that variable is difficult to measure, which essentially means: (a) that there is sizable noise in the measurement of the variable (and, therefore, that an important part of the agent's compensation is random), and (b) that the agent can influence the measurement of the variable. Formally, both § and b affect m_1.

Of course, the existence of noise may give the agent a very high risk premium because the rewards have a random component unrelated to his achievements, but the size of this premium will depend on the specific problem. The consequences of (b) are easy to forecast: very likely, the agent will devote time, effort and material resources to obtaining good performance measures, regardless of the underlying reality. Instead of using the variable «a» to get a good value of x_1, and thereby of m_1, he will spend time on «b» to modify m_1, even when x_1 is unsatisfactory.

Note that an agent can do this in at least two qualitatively different ways. First, he can take actions that increase the value of the measure, without a proportionate increase in the variable itself. We can find plenty of examples of this behavior in everyday business practice. For instance, suppose the head of a production department is evaluated on the basis of standard costs and variance analysis, with the (explicit or implicit) incentive to minimize costs. This objective can be achieved by trying to make the department as efficient as possible, but it can also be achieved by allowing the quality of the goods produced to deteriorate, by using only experienced workers and delaying delivery, or by accepting only large orders so as to reduce setup costs.

In less structured settings, things can only be worse. Suppose, for instance, that a sales job has to be evaluated and an incentive put on the measure of performance. Then, if the measure is, say, the number of sales calls, the salesman can always artificially increase the number of unprofitable calls; or, if evaluated on sales, increase the number of unprofitable sales (i.e., sales with no margin); or, if evaluated on margin, increase the sales to doubtful accounts; and so on.

The former Soviet Union has been a fertile source of examples of this kind of behavior. For instance, when Moscow taxi drivers were put on a bonus system based on mileage, the Moscow suburbs were full of empty taxis driving extra miles to increase their bonuses. Also, when bonuses were based on tonnage of production, chandeliers and armchairs became heavier and heavier until they started pulling ceilings down (Horngren, 1972, p. 156).

Hofstede (1978) has described this kind of situation as «pseudo-control»:

«Pseudocontrol is a state of affairs in which a system is under control on paper (the figures look right), but not in reality. There are several ways to achieve pseudo-control. For example, by correcting the standards (rather than the process) whenever an important variance occurs, by choosing one's measures of performance so as to reflect the desired result (there are many ways to bend the figures without actually cheating), or *by adjusting one element in the process at the expense of another that doesn't show up in the figures* (4) (reducing cost at the expense of quality).»

Second, of course, the agent can try to cheat by plain manipulation of the measure. Of course, there may be penalties, but their enforcement is always less than perfect.

(3) Assume now that the incentive formula includes more than one variable, all of them with no measurement problems. These variables enter as arguments into the utility function of the principal; and, typically, there will be trade-offs between them, i.e., one may have to choose between a higher value of one variable and a lower value of another, or vice versa. Then, unless the agent's reward function is directly the principal's utility function, the trade-offs to the agent will not, in general, be the same. If the agent attempts to maximize the incentive payment, the result will, in general, be inefficient.

(4) Finally, assume that the incentive formula includes more than one variable and that at least one of them is hard to measure. One would expect to find all the problems mentioned above. First, unless the compensation formula reflects the principal's utility function, there may be undesirable trade-offs between variables. Second, even if the compensation formula does reflect the principal's untility function, but some of the variables are hard to measure, one may find a displacement of resources from the productive activities to activities that improve the measures of these variables.

Notice that for an agent it may sometimes be easier to achieve an inflated level of one of the hard-to-measure variables than a higher value of a measurable variable. Therefore, actual resources may even be taken from productive activities directed towards achieving these measurable objectives, and consumed in improving inflated measures of performance. Thus, the influence of the possible perverse effects is not limited to the variable itself, but can extend to the whole domain.

Agents with utility for a job well done

We assume now that the agent, besides having utility for monetary rewards, has utility for well-done jobs as well. The principal also has positive utility for a well-done job, since he owns the product. This is a factor of alignment between the utility functions of agent and principal: both functions increase with the degree of perfection of a given job.

This by itself, however, does not guarantee that the result will be optimal, essentially for the same reasons as in (3) above. Technical perfection may increase the value to the principal, but her utility is defined in several variables, of which this is only one. It might well be, for instance, that from an economic point of view, an alternative result that is «not so good» would be better that the «perfect» one. An incentive payment may help to align the objectives of principal and agent, but this alignment will generally be less than perfect, for reasons stated in the previous paragraph.

Agents with disutility for effort

This is the case most often studied in agency theory, with well-known results.

Agents: the general case

The utility function of the principal is defined in n variables. The utility function of the agent is defined in the same variables, in principle, although it may be constant in some of them (i.e., the agent may be indifferent to some of the variables). Of the other variables, in some (e.g., «pride») the interests of the principal and of the agent are «aligned» (i.e., the utility functions are both either increasing or decreasing), while in the rest they are opposed (e.g., «effort»).

The solution to this problem is, of course, complex. One can, however, indicate some opportunities and dangers. On the one hand, given the existence of «opposed» variables, a compensation scheme based on results seems to be logical, as one of the possible ways to motivate the agent to take actions that he would not otherwise take and that are in the best interest of the principal. On the other hand, and given what has been discussed in the previous sections, the incentive payments may have dangerous side-effects. For instance, they may encourage the agent to focus on some of the more measurable variables at the expense of the others. Also, the agent may engage in costly activities aimed at increasing the measures of perormance (and the payments) without any corresponding increase in the values of the underlying variables. These are some of the points raised by Kohn (1993) in criticizing incentive systems from a practical point of view. On the other hand, his critics have argued basically that it could not possibly be better to pay the same amount no matter what the actions and the results (G. Bennett Stewart III et al., 1993).

Milgrom and Roberts (1990, 1992) use the term «influence activities» to refer to the agent's attempts to influence the measurement system. «Influence costs», then, limit the growth of organizations: one would expect these costs to grow faster than the size of the organization, so that beyond a certain size it would already be better to have two (or more) separate organizations instead of just one.

Designing incentive systems is a very complex task and cannot be done without taking other considerations into account. Our analysis shows that formal systems of management control (i.e., measurement and compensation systems) are generally incapable of solving the problem on their own. There is a need for additional actions on the part of the principal to ensure a closer alignment between her objectives and those of the agent.

The main conclusion to be drawn from our model has to do with the limitations of formal systems. It throws light on the open debate on the use of incentive systems in organizations (see, for instance, Kohn, 1988, 1993a, 1993b, and Lawler III, 1990). To illustrate, let us imagine for a moment that the principal does not have any way to influence the agent's perceptions on x. She can only influence the agent’s choice by means of the compensation system s(m). Obviously, this is contingent only on observable, measurable, verifiable information provided by the measurement system. But then we have two problems:

First, the principal has only one variable, the compensation, with which to influence a multi-dimensional tuple x by influencing the action «a». Any form of compensation will induce behavior that, given incomplete information, is highly unlikely to affect all relevant variables in the required way. This line of reasoning is similar to the multitask setup in

Holmstrom and Milgrom (1991). Furthermore, the agent has some capacity to affect the measurement system, via «b», and will have every incentive to act in such a way as to neutralize the incentive effect of the compensation. In the end, the principal may discover that the effort is directed not towards better decision-making, as intended, but towards modifying the measurement system.

The problem we are trying to show is how difficult it is to be precise about the behavior that the incentives are supposed to encourage and how to avoid undesirable behavior. If professors were paid according to their students' evaluation of their performance, they might be tempted to try to influence those evaluations instead of getting on with the job of teaching. Also, if they were responsible for handing out and collecting in the evaluations, they might try to remove the unfavorable evaluations before they entered the system. This sort of thing goes on every day in organizations. It is not that incentives are not efficient. They are too efficient. But we cannot prevent all unintended actions and consequences.

Under what conditions will s(m) work more efficiently as an incentive device? If it is possible to define precisely how the agent should act and prevent him from manipulating the measurement systems, then the incentive systems will work properly. This is why some companies driven by cost leadership are able to use their incentive system efficiently. Two well known examples are Lincoln Electric Co. (5) and Nucor Corporation (6) (Ghemawat, 1993). While their incentive approach is very different, they have one thing in common. Both are cost-driven companies where strategy is very clear; what the agents have to do in order for the company to succeed is well established and, consequently, incentives work properly (7).

Good incentive systems are not only to be found in cost leadership cases. Nordstrom Company (8) is a Department Store company that successfully uses high powered incentives for its salespeople by paying them on the basis of sales per hour. This compensation policy is one component of a system in which sales staff are empowered to make their own decisions when serving customers. Thanks to this empowerment and other policies, the message regarding Nordstrom's strategy is clearly conveyed to the salesperson: «the salesperson should do anything for the customer». Such simplicity allows high powered incentives to do their job properly. All the same, the company should be very careful to ensure that the system is not manipulated and to check that undesirable behavior is not taking place. This is a typical situation with empowered workers: rewards based on group or individual performance work very well, but companies should be sure to set up a safety net of control to avoid undesirable behavior.

Generally, however, managerial decisions are difficult to define with precision; they have many unintended consequences, and managers have the means and the ability to influence measurement systems. Therefore, formal systems are not enough to «control» the agent (see also Ricart et. al, 1991). This justifies our including in the model the principal's ability to influence the agent's perceptions.

These influence activities can be of very different kinds. We think they could be classified in three components:

(1) *Information transmission.* One way the principal can influence the agent is by giving him enough information about the company's strategy, the way his actions affect organizational performance, the current state of the organization, his perception of the world, etc. This communication allows the agent to learn, to understand the effect of his actions and to share in the organization's objectives and policies.

This is coherent with the prevailing recommendation in management that firms should try to be participative in strategy formulation, that they should clearly communicate the strategy to the organization and try to develop the capacity for independent thought among their executives. One needs leaders to transmit «vision».

Furthermore, the importance of communication is also made clear in the literature on «empowerment» (Lawler III, 1992, or Galbraith, Lawler III & Associates, 1993). To get high involvement in organizations the typical recommendations are: (1) Give information; (2) Give training; (3) Empower decision; and (4) Reward good performance. Note that elements (1) and (2) mostly mean transmit enough information and knowledge so that the agent is able, based on the principal's scientific and empirical laws, to understand the alternative actions and the effect they can have on performance variables.

(2) *Socialization culture and leadership.* A second way of influencing perceptions has to do with developing strong cultural values in the organization. Developing missions, credos, a corporate philosophy, a «way of doing things here», a particular culture, etc. Culture, as a general concept, is becoming increasingly important as an integrative device in decentralized organizations.

Note a current trend in organizational design that uses a very high degree of decentralization so as to make better use of idiosyncratic particular knowledge that is generated by teams working in close contact with the client. In order to integrate the largely autonomous, self-managed teams, firms use: (1) extensive training; (2) information diffusion via a sophisticated, usually hard-to-manipulate measurement system; (3) a clear statement of mission and strategy; and (4) a strong culture.

Banc One Corporation (9) is a good example. Banc One is one of the largest US banks. It has grown mostly through the acquisition of small regional banks. After an acquisition, Banc One gives considerable autonomy to each affiliate bank in carrying out its commercial activities, but always within a set of constraints: (1) A clear strategic focus that makes all affiliate banks very similar; (2) A very strong, centralized back office that provides the small banks with the services and products of a very large bank and a shared information system; (3) A very strong cohesive culture based on personal values; and (4) A very strong network of relationships derived from their socialization process, which allows them to get a lot of internal benchmarking and sharing of experience.

(3) *Utility manipulation.* Any influence activity obviously has its associated excesses. The principal could try to manipulate the agent by affecting his utility through the use of «psychological pressures», «fears», or other counterproductive systems. The boundary between a strong positive culture and a manipulative one is very difficult to define, but it is a very important distinction that goes straight to the heart of human nature and personal freedom, as well as the limits to the interference of organizations in the private life of individuals. Also, this is not unrelated to the debate about managing diversity in organizations, where some firms are trying to develop competitive advantage based on having a very diverse workforce, while others argue in favor of a very uniform workforce.

Whatever the nature of the influence, managers need to rely on some of these elements because formal measurement and incentive systems are not enough to induce the «right» behavior in agents. However, once all this comes into play, we are introducing so much noise into the system that our compensation system will have to be very simple and clear; otherwise it will be very inefficient.

We find ourselves in a «no way out» loop. Except in very simple and clear situations, incentive and measurement systems (formal systems) are not enough; their efficiency as motivations to choose the «right action» is very limited. Firms have to rely on other mechanisms such as strategy, culture, information, training, and leadership. However, once all this has been introduced, the amount of noise in the chain reaction from compensation to action is such that incentives will tend to be simple and general, even team or organization-based. (This reasoning is similar to that used in the Milgrom Holmstrom 1985 paper.) This divides the world into two sets: stable, simple cases where static efficiency is good enough and formal systems work properly, and cases where dynamic efficiency is prevalent and complexity calls for simple incentive schemes and complex organizational arrangements (Ghemawat and Ricart, 1993).

VI. Conclusions

What we have attempted to do in this article is to formulate a general model of agency theory that could be useful in principal-agent situations in an organizational context, i.e., when the agent is the principal to someone else, and the principal is likewise the agent to another person. We argued that in such a context the arguments of the utility function of both the principal and the agent are multidimensional, and are not necessarily related to profit or any single measure of economic performance. Besides, we introduced the issue of measurement and perception problems on the part of both the principal and the agent.

Our conclusions can only be tentative at this time. We examined the implications of such a model in terms of bounded rationality and learning, organizational design, and conflicts of interest and incentives under different sets of assumptions. Finally, we suggested different ways to reduce agency costs in these circumstances. □

---

(1) Most of these expressions have been taken from the measurement project that General Electric Company undertook in the ’50s.

(2) It is common in the agency literature to qualify information as «verifiable» or «unverifiable». We have not used this distinction here because verifiable means that something can be done to ascertain whether the information is true or not, whereas what interests us here is whether, in fact, anything has been done to authenticate the information and so make it credible for all participants.

(3) The model in Holmstrom and Ricart i Costa (1986) is an example of an organizational design in which the agent’s role is to propose investments while the principal may veto them.

(4) Our italics.

(5) Harvard Business School case 9-376-028, Lincoln Electric Co.

(6) Harvard Business School case 9-793-039, Nucor at Crossroads.

(7) As expected, things are more complex because high power incentive systems are usually just one element among a number of complex human resource policies.

(8) Harvard Business School case 9-191-002, Nordstrom: Dissension in the Ranks? (A).

(9) Harvard Business School case 9-390-208, Banc One Corporation 1989 (Abridged).

References

G. Bennett Stewart III et al., «Rethinking rewards», Harvard Business Review, November/December 1993

Christenson, Charles, «The Methodology of Positive Accounting», The Accounting Review, vol LVIII, no 1, January 1983.

Galbraith, Jay R., Edward E. Lawler III & Associates, Organizing for the Future: The New Logic for Managing Complex Organizations, Jossey-Bass Publishers, San Francisco, 1993.

Ghemawat, Pankaj, «Commitment to a Process Innovation: Nucor, USX, and Thin-Slab Casting». Journal of Economics and Management Strategy, Spring 1993.

Ghemawat, Pankaj, and Joan E. Ricart, «The Organizational Tension Between Static and Dynamic Efficiency», Strategic Management Journal, Vol.14, 1993.

Grossman, S. and O. Hart, «An Analysis of the Principal-Agent Problem», Econometrica, Vol 51, Jan. 1983.

Harsanyi, J.C., «Games with Incomplete Information Played by "Bayesian Players"» Management Sciences, 14, 1967-68, pp. 159-189, 320-334, and 486-502.

Hayek, F.A. von, «The Use of Knowledge in Society», American Economic Review, Vol.35, No. 4, Sept. 1945.

Hofstede, Geert, «The Poverty of Management Control Philosophy», Academy of Management Review, July 1978, pp 450-461.

Holmstrom, Bengt, «Moral Hazard and Observability», The Bell Journal of Economics, 1979, pp. 74-91.

Holmstrom, Bengt, and Paul Milgrom, «Aggregation and Linearity in the Provision of Intertemporal Incentives», Econometrica, 1985.

Holmstrom, Bengt, and Paul Milgrom, «Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design», Journal of Law, Economics, & Organization, vol 7, Special Issue, 1991, pp. 24-52.

Holmstrom, Bengt, and Joan E. Ricart i Costa, «Managerial Incentives and Capital Management», Quarterly Journal of Economics, 1986, pp. 835-860.

Horngren, Charles T. (1972), Cost Accounting: a Managerial Emphasis, Prentice-Hall.

Jensen, Michael C., and William H. Meckling, «Specific and General Knowledge and Organization Structure», in Lars Werin and Hans Wijkander (eds.), Contract Economics, Blackwell, 1992.

Kohn, Alfie, Incentives Can Be Bad for Business, INC., 1988.

Kohn, Alfie, «Why Incentive Plans Cannot Work», Harvard Business Review, Sep.-Oct. 1993.

Kohn, Alfie, Punished by Rewards, Hoguton Mifflin Company, 1993.

Lawler III, Edward E., Strategic Pay: Aligning Organizational Strategies and Pay Systems, Jossey-Bass Pub., San Francisco, 1990.

Lawler III, Edward E., The Ultimate Advantage: Creating the High Involvement Organization, Jossey-Bass Pub., San Francisco, 1992.

Milgrom, P., and J. Roberts, Bargaining Costs, Influence Costs, and the Organization of Economic Activity, in Perspectives on Positive Political Economy, A. Alt and K. Shepsle, eds., Cambridge University Press, 1990.

Milgrom, P., and J. Roberts, Economics, Organization and Management, Prentice-Hall, 1992.

Popper, Karl, The poverty of historicism, Routledge and Kegan Paul, 1957.

Ricart, Joan E., Jordi Gual, Guillem López, Josep Mª Rosanas, and Josep Valor, Incentivos y Control en la Empresa Pública Española, Ariel Economía, Barcelona, 1991.

Senge, Peter M., The Fifth Discipline: The Art and Practice of the Learning Organization, Doubleday/Currency, New York, 1990.

IESE

DOCUMENTOS DE INVESTIGACION – RESEARCH PAPERS

No.	TlTULO	AUTOR
D/ 267	Sector español de la alimentación y bebidas: Empresas familiares y no familiares (II) Mayo, 1994, 43 Págs.	Gallo M.A. Estapé M.J.
D/ 267 BIS	The family business in the Spanish food and beverage industry (II) April 1994, 43 Pages	Gallo M.A. Estapé M.J.
D/ 268	Alliance networks in European banking. June 1994, 26 Pages	García Pont C.
D/ 269	Las relaciones laborales en Alemania. Junio 1994, 61 Págs.	Gómez S. Pons Mª
D/ 270	Las relaciones laborales en Francia Junio 1994, 66 Págs.	Gómez S. Pons Mª
D/ 271	La estructura interna de la empresa en la visión de Coase. Julio 1994, 19 Págs.	Argandoña A.
D/ 272	Teaching business ethics using the case method. July 1994, 17 Pages	Argandoña A.
D/ 273	Weak form market efficiency: A comparison between the Spanish and the U.S. stock markets. July 1994, 39 Pages	Martínez Abascal E. Pregel G.
D/ 274	Investing in Russia. July 1994, 36 Pages	Taylor G.
D/ 275	Aspectos éticos en la consultoría de búsqueda de directivos. Septiembre 1994, 16 Págs.	Melé D. Roig B.

IESE

DOCUMENTOS DE INVESTIGACION – RESEARCH PAPERS

No.	TlTULO	AUTOR
D/ 276	Factores determinantes del ahorro. Septiembre 1994, 88 Págs.	Argandoña A.
D/ 277	Spain and the European social charter. September 1994, 33 Pages	Argandoña A.
D/ 278	Technology management in Spain. October 1994, 14 Pages	Riverola J. Muñoz-Seca B.
D/ 279	Códigos de conducta empresarial: ¿Pueden contribuir al desarrollo ético de los empleados? Octubre 1994, 14 Págs.	Melé D.
D/ 280	Las relaciones laborales en Italia Noviembre 1994, 61 Págs.	Gómez S. Pons M.
D/ 281	Training the next generation of owners and managers: A possible key role for directors in family businesses. November 1994, 17 Pages	Tomaselli S.
D/ 282	Los servicios: El binomio privatización- desregulación. El caso de la educación. Diciembre 1994, 15 Págs.	Argandoña A.
D/ 283	Las relaciones laborales en el Reino Unido. Diciembre 1994, 75 Págs.	Gómez S. Pons Mª
D/ 284	Isomorphic pressures on identity: The case of learning partnerships with business schools. December 1994, 21 Pages	Enrione A. Knief C. Mazza C.
D/ 285	Control and incentives in organizational design. January 1995, 22 Pages	Ricart J.E. Rosanas J.M.