A SEQUENTIAL PRINCIPAL-AGENT APPROACH TO REGULATION

P. B. Linhart
R. Radner
F. W. Sinden

Bell Laboratories
Murray Hill, New Jersey 07974

1. Introduction

It is commonly believed that regulation should simulate competition in order to maximize social welfare. This is sometimes taken to imply that regulators should select prices so that the firm's economic profits are always zero. If, however, this rule were actually followed precisely, the firm would be deprived of all incentive to minimize cost. In particular, the firm would have no incentive to select an optimal mix of factor inputs, to obtain these inputs from the lowest-cost suppliers, or to invest in cost-reducing innovation.

In fact, the record of productivity improvement in regulated industries compares favorably on the whole with that of the entire economy (Houthakker, 1979, 1981). It would appear, then, that real regulation simulates real competition better than the above simple analysis suggests.

Both positive and normative theories of regulation have already received considerable attention in the economics literature. In this introductory section we indicate briefly why there is a need for further study of efficiency incentives under regulation, and sketch several properties that a feasible incentive mechanism should have.

Real competitive markets provide incentives for innovation through temporary monopoly profits, as described by Schumpeter (1942). More or less unintentionally, real regulation provides a similar incentive. Regulation is not exact and continuous, but takes the form of a discrete series of reviews, at which (in a simplified picture of the process) the prices are reset so that economic profit is zero. Between reviews, if the firm can reduce its costs, it can earn a positive economic profit. It has been pointed out (Baumol, 1968) that the imperfection of this process (loosely called "regulatory lag"), far from being a drawback, confers a positive benefit on society. Baumol recommends that it be institutionalized.

meet demand at the given prices, as long as it is feasible to do so, and that the shareholders and directors require the manager to pay out the cost of capital at a given rate, again as long as this is feasible. These two requirements can be met as long as the firm's cash reserve is positive.

In the context of a particular formal model, we show that, under this class of regulatory strategies, the management of the regulated firm will have an incentive to engage in productivity improvement. Furthermore, if the management's behavior is optimal from its own point of view, then the incentive phases will be long relative to the recovery phases, and the resulting long-run average rate of actual price decrease will be close to the regulator's target rate of price decrease, provided the management does not discount its own future benefits too strongly.

Thus, under suitable conditions, this class of regulatory strategies induces economically efficient behavior on the part of the manager, without placing a large informational burden on the regulators and their staff, and in particular without requiring the regulators to monitor the firm's rate of return except during the infrequent recovery phases.

It is, of course, difficult to say how efficient or close to the social optimum the incentives provided by regulatory lag are; one of the aims of the present paper is to provide a framework within which the efficiency of such regulatory mechanisms can be discussed more precisely.

When inflation raises the cost of the firm's inputs, regulatory lag as an incentive mechanism no longer works so well. This comes about as follows. In any industry, given its technological possibilities, there is a maximum achievable rate of cost-reduction; beyond this rate, the cost of developing and installing new cost-reducing technology exceeds the cost reduction. Suppose the rate of inflation exceeds this maximum achievable rate of productivity increase, as has arguably been the case for most American industry in recent years. Then if prices are fixed in nominal terms at regulatory reviews, as they customarily are, the firm will necessarily lose money between reviews no matter what cost-reducing measures it takes.

One would expect, therefore, that unusually rapid inflation would cause firms to apply for rate reviews with unusual frequency in order to reduce their losses. The data bear this out. Figure 1, for example, shows the remarkable correlation between the rate of inflation and the frequency of general rate orders for Bell Telephone companies (excluding Long Lines).

Frequent rate reviews have two social costs: the direct cost of the extra hearings and the indirect cost of reducing regulatory lag, and thus weakening incentives to improve productivity. In the limit as reviews become continuous, productivity incentives approach zero, and the rational firm invests nothing in productivity improvement.

In times of rapid inflation, then, conventional regulation fails in two ways: it leads to losses and it leads to weakened incentives. These problems have led both regulators and utilities to seek formulas ("inflation adjustment clauses") that correct for inflation and at the same time preserve productivity incentives.[1] This search has in turn led some investigators to examine more closely the game that underlies regulation.

In this paper we present a new approach to regulation, which is suggested by the theory of

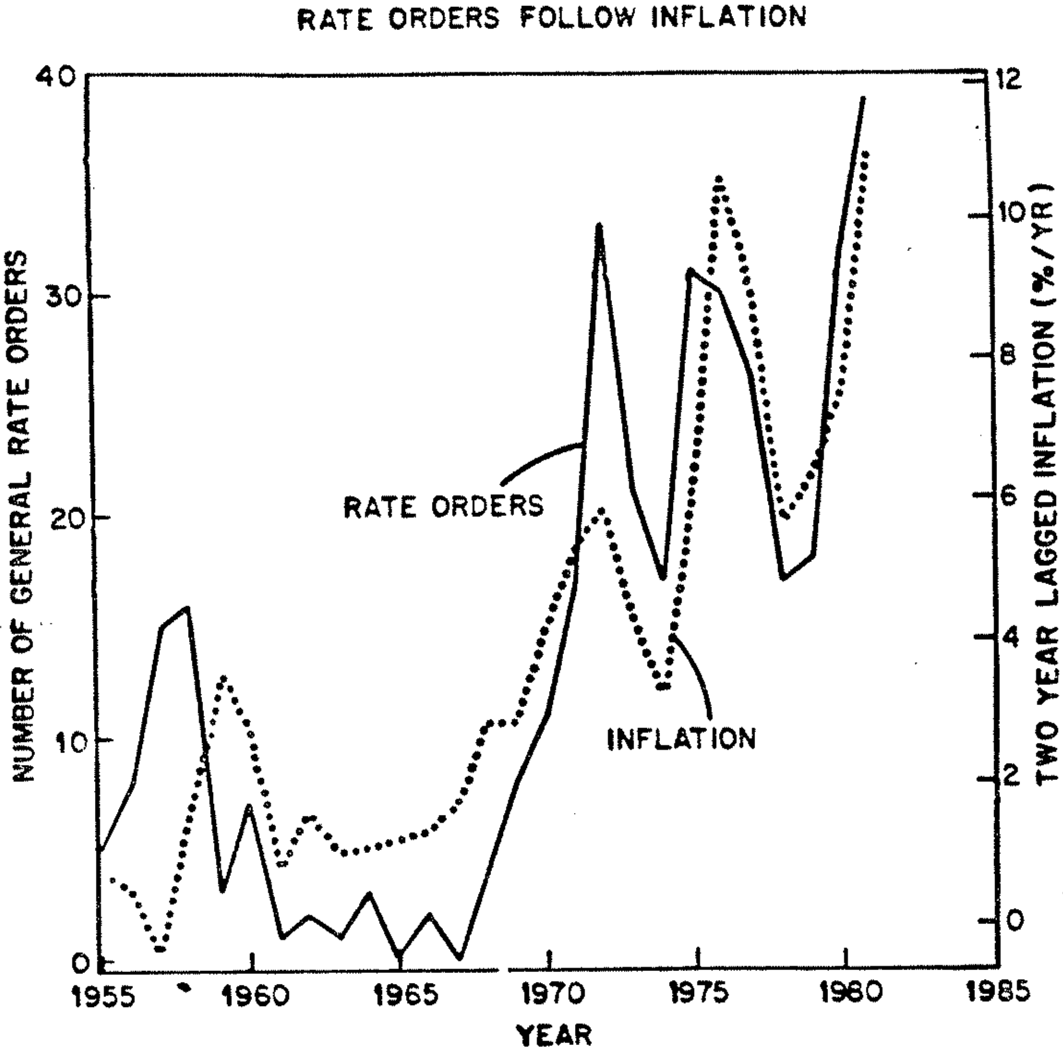

Figure 1

Data on the left hand vertical axis represent the annual number of general rate awards to Bell System companies; the source is AT&T Co. internal records. Data on the right depict percentage annual changes in the Consumer Price Index, lagged two years — the approximate average interval between rate cases being filed and awards granted. The 1981 plotted points include partial projections. The correlation coefficient of the two curves is 0.84.

sequential principal-agent games. Our model takes account of the following fundamental characteristics, among others, of the regulatory situation:

1. The regulator and the firm have different information. In particular, the firm has more information about the possibilities for productivity improvement than the regulators. In fact one of the firm's options is to invest in research in order to obtain more of this information. In principle, the regulator could also obtain more information at some cost, but matching the firm's information seldom appears to be part of the regulator's strategy.

2. The regulator and the firm's managers to some extent have different goals. The regulator may strive to provide incentives strong enough to overcome the difference, but in general we would not expect an equilibrium outcome to meet the regulator's goals entirely.

3. The service is deemed essential, so that its continued availability must be assured in spite of possible mismanagement and/or bankruptcy.

To be acceptable in the real world, a regulatory mechanism must not differ too radically from those that already exist. The strategies we discuss resemble conventional regulation in that periods of regulatory inaction alternate with periods of action that are intended to be corrective.

The essence of the regulator's problem is that he cannot directly observe the manager's actions, nor can he observe the exogenous random events that also affect productivity. He can, however, observe the consequences of those actions and events, namely the realized profits of the firm, and whether or not demand is met. He may also, with additional effort, be able to observe productivity changes, but we do not in our model rely on this possibility.

Suppose that the regulator provisionally fixes a sequence of prices that decline in real terms at a fixed "target rate" (which must be suitably chosen). Suppose further that the regulator requires the firm to meet demand at the given prices, as long as it is feasible to do so, and that the shareholders and directors require the manager to pay out the cost of capital at a given rate, again as long as this is feasible. These two requirements can be met as long as the firm's cash reserve is positive.

In the context of a particular formal model, we have shown that, under this class of regulatory strategies, the management of the regulated firm will have an incentive to engage in productivity improvement. Furthermore, if the management's behavior is optimal from its own point of view, then the incentive phases will be long relative to the recovery phases, and the resulting long-run average rate of actual price decrease will be close to the regulator's target rate of price decrease, provided the management does not discount its own future benefits too strongly.

Thus, under suitable conditions, this class of regulatory strategies induces economically efficient behavior on the part of the manager, without placing a large informational burden on the regulators and their staff, and in particular without requiring the regulators to monitor the firm's rate of return except during the infrequent recovery phases.

Although inflation has played a role in stimulating the present interest in incentives, inflation is not essential to the theoretical structure of the underlying game. We sidestep inflation problems by assuming that prices are always measured in real dollars. We do not address the nontrivial econometric problem of defining a suitable price index.

In Section 2 we discuss in general terms the relations between the so-called "principal-agent" problem and the problem of providing incentives for efficiency in a regulatory situation. In Section 3 we describe a particular formal model that embodies some of the considerations sketched in Section 2, and we give formal statements of the properties of the regulatory strategy.

Section 4 provides some remarks on possible generalizations and extensions of the model of Section 3.

2. Regulation and the Principal-Agent Problem

The problem of regulation can be looked at as a principal-agent problem -- more precisely as a sequential principal-agent game. In fact, several such games are ingredients of the true regulatory situation. Thus the principal-agent literature can be extended to give a reasonably realistic description of regulation and to derive results about the efficiency of regulatory strategies. We now review briefly some features of principal-agent games.

In a principal-agent situation, a principal hires an agent to act on his behalf, generally because the agent has better information about some enterprise of interest to the principal. The resulting outcome depends on a random state of the environment as well as on the agent's action. After observing the outcome, the principal makes a payment to the agent according to a pre-announced reward function, which depends directly only on the observed outcome. (This last restriction expresses the fact that the principal cannot directly observe the agent's action, nor can the principal observe the information on which the agent bases his action.)

If this action-reward situation occurs only once, we call it a short-run principal-agent relationship. The situation can be naturally modeled as a two-move game, in which the principal first announces a reward function to the agent, and then the agent chooses an action (or decision function if he has prior information about the environment).

An equilibrium of such a game is a reward function (by the principal) and a decision function (by the agent), each of which is optimal given the other. The equilibria of such a game are typically inefficient (unless the agent is neutral towards risk), in the sense that there will typically be another (but nonequilibrium) reward-decision pair that yields higher expected utilities to both players.

If such a game is repeated *infinitely* many times, the whole set of repetitions is said to constitute a "supergame". In such a supergame, more efficient equilibria exist, in the sense that both "players" are *on the average* better off than in the short-run relationship. This is essentially because in the supergame, the principal can design his reward strategy in such a way as to punish noncooperative behavior by the agent.

Of course, since the principal cannot observe the agent's action, but only the outcome of that action, he cannot be sure in case of a bad outcome whether the agent acted in bad faith or was the victim of bad luck (i.e., the random state of the environment was unfavorable). However, after many repetitions the principal can distinguish *statistically* between bad faith and bad luck and impose (or withhold) penalties accordingly.

Finally, if the agent's objective is to maximize the *discounted* sum (present worth) of future rewards, there is some loss in the efficiency of equilibria, since future penalties lose some of their force compared to present gains. But as the discount rate approaches zero, this loss in efficiency disappears.

Repeated principal-agent games, including those with discounting, are discussed in Radner (1981a, 1981b).

When we begin to look for principal-agent relationships in the regulatory scene, several immediately spring to mind. First of all, the regulatory body can be seen as an agent of society or the electorate. Regulators may have - doubtless do have - their own personal goals (see, e.g., Stigler (1971)). The problem for society, acting through the political process, is then to provide a system of rewards and penalties such that regulators will best further their personal interests by acting in the general interest.

Second, the firm can be seen as an agent of the regulators; this is the relationship on which we are concentrating in this paper. (If incentives for the regulators to act in the interest of society are correct, *and* if incentives for the firm to act in the interest of the regulators are correct, then the chain will hold and the firm will act in the interest of society.) The actions of the firm in which we are most interested (in this paper) consist in undertaking productivity-improving projects, typically R&D. Such projects cost money, involve the expenditure of effort, and may or may not result in lowering the firm's production costs. The regulators can observe the cost-reduction, if any, and the expenditure of dollars on R&D, but not the expenditure of effort. Rewards to the firm will take the form of increased allowed profit.

In reality, of course, the firm is not a monolith. A third principal-agent relation exists between the firm's owners and its managers. It is reasonable to assume that the owners want to maximize the firm's profit, but there is no reason to think that the managers internalize this objective; their direct interests may include salary and bonuses, longevity of employment, span of control, effort (to be avoided), and research interests. The owners' well-known problem is then to provide incentives for managers to act in the direction of profit maximization. In this connection, we may take the owners (shareholders) to be risk-neutral and the managers risk-averse.

Neither the regulators nor the owners can discriminate between bad luck and bad faith on the part of the managers, except statistically, so both regulatory rewards and penalties to the firm (in terms of profit) and the owners' rewards and penalties to the management (in terms of bonuses, promotions, etc.) will have to be based on statistical evidence.

Other principal-agent relationships can be discerned. For example, the board of directors may be interposed between the owners (shareholders) and the managers. Also, we have said nothing about those of the firm's employees who are not managers; a whole hierarchy of principal-agent relationships exists here. Moreover, these workers, if they are unionized, may employ union officials as their agents in negotiations with management.

For our present purposes, we find it expedient to collapse the electorate and the regulators into a single entity, and the owners, managers and workers into a single entity, and to examine a single principal-agent relationship.

3. Description of a Formal Model

We have claimed that a certain class of regulatory strategies (sketched in Section 1), which is consistent with the informational asymmetry between the manager and the regulator, can induce economically efficient behavior on the part of the manager. We now substantiate this claim by showing that it is so, and in what sense it is so, in the context of a particular model. This model, although much simplified for purposes of tractability, is plausibly representative of a wide class of regulatory models.

In the present treatment, we describe the model in some detail and state our results with some precision, but omit almost all mathematical notation and all proofs. The technical details may be found in Linhart, *et. al* (1982). We now specify:

(a) the technological and financial aspects of our simplified model,

(b) the distribution of information,

(c) the objectives of the regulator and the manager, and

(d) the strategy of the regulator, and the optimal responsive behavior of the manager.

Finally we state, without proof, our results, which make precise the sense in which the whole scheme is efficient, i.e., the sense in which the manager's self-interested behavior achieves the regulator's objective.

a. Financial and Technological Model

We assume that all the firm's inputs, including capital, can be adjusted in each period so that the firm can just meet demand in that period at minimum cost. Thus capital investment is reversible, and questions of excess capacity do not arise.

We assume that the real cost, per unit of output, to the firm of each of its factor inputs, including capital, is the same in each period. We take the unit cost of capital to be in a fixed ratio to the cost of all other inputs plus the depreciation of capital; this would be the case for example if the depreciation schedule were unchanging and if capital and all other inputs were optimally used in fixed proportions (technically, if the production function were homothetic).

We assume Hicks-neutral technical change, i.e., that changes in productivity affect all inputs proportionally, so that we may take as a measure of productivity either (i) the ratio of output to capital (where "output" is an index of the outputs of the firm and "capital" is an index of the capital of the firm, measured in constant dollars), or (ii) equivalently, total factor productivity.

We think of changes in productivity as brought about by R&D. We suppose, for simplicity, that the R&D budget is fixed, so that the manager's choices lie in how he manages the R&D budget and how he implements any resulting productivity improvement. Although in reality the conduct of R&D activities in any one period typically has consequences for productivity improvements in future periods, we suppose here that these consequences are entirely concentrated in the current period.

We recognize that increases in productivity are affected by the manager's actions, as just described, but are also affected by exogenous or random factors beyond the manager's control (e.g., the unpredictability of the outcomes of research projects). We assume the technology of productivity change to be stationary, in that given managerial actions combined with given exogenous conditions will achieve the same productivity change in any period.

Finally, we assume that the firm's net profits (i.e. revenues minus costs, including capital costs) are accumulated in a *cash reserve*, which is initially positive. This reserve earns some return, at a rate not necessarily equal to the cost of capital. Since net profits can be negative, the cash reserve can decrease and even become negative. Each time the cash reserve becomes negative, we say a *crisis* has occurred. We shall have more to say below on how the regulator deals with a crisis.

It turns out that this device, of hinging the regulator's action solely on the sign of the firm's cash reserve, is the key simplification that makes the model tractable.

b. Distribution of Information

We assume that the regulator cannot directly observe the manager's actions, nor can he observe the exogenous random events that also effect productivity changes. This limitation on the regulator's information constitutes the essence of his problem. (He may, with effort, be able to observe the resulting productivity changes, but *we do not in our model rely on this possibility.*) The regulator can, however, observe the consequences of these managerial actions and random events, namely the realized profits of the firm and whether or not demand is met.

The manager is of course aware of his own past actions, and of the past history of prices and realized productivity changes. Importantly, he is also aware of the regulator's policy, which the regulator has announced. Thus, if it is part of the regulator's policy to reduce (real) prices by a fixed percentage each period, the manager is aware of this projected schedule of prices (until the next crisis). He can base his action in each period on all this information. A managerial strategy is a rule (sequence of functions) according to which the manager determines his action in each period as a function of his information.

It remains only to add that we assume that both the manager and the regulator know the price-elasticities of demand for the firm's products.

c. Objectives

We assume that the regulator wants low prices for the firm's outputs. Because of increasing productivity, the firm can afford to lower its prices over time and still remain solvent. Thus, more precisely, the regulator's objective is to maximize the long-run average rate of actual price decrease.

With respect to the manager, we assume that his wage is fixed, and we make the perhaps rather Draconian assumption that at each crisis, when the firm first becomes insolvent, the manager is fired, and replaced. Thus one component of the manager's objective is to maximize his tenure, i.e., the period during which he receives his (fixed) wage.

On the other hand, we admit the possibility that the manager's preferences among the alternative actions he can take in one period would not, other things being equal, lead to maximum productivity increase (and hence maximum price decrease) in the long run. It is this possibility, of course, that motivates the regulator to devise a policy that will enhance the manager's incentive to increase productivity. These preferences form the other component of the manager's objective[2].

Both sorts of considerations — his desire to continue receiving his salary and his preference for certain managerial actions — will contribute to determining the manager's optimal strategy. For example, we would expect that when the cash reserve gets close to zero, the manager would act so as to increase productivity very rapidly, thus reducing costs more than the regulator is reducing prices. On the other hand, when the cash reserve is very large we would expect him to act in a way closer to his one-period preference.

d. Strategies

As we have said, the firm is required to adjust its output to meet demand exactly, and to pay all costs including the cost of capital. It is able to do these things as long as it is solvent, i.e., as long as its cash-reserve is positive. Through bad luck or bad management, the cash reserve can become negative. This event we call a crisis; when a crisis occurs, the manager is fired and replaced. The regulator must now provide some way for the firm to get back on its feet; this part of the regulator's strategy will be described below. Thus time is divided into alternating segments: *incentive phases* and *recovery phases*.

During each *incentive phase*, the regulator requires that prices decrease at a fixed percentage say — rate ϕ — in each period[3]. For the scheme to work, ϕ must be set at a level not greater than the greatest achievable level of productivity increase. The regulator will want to meet this condition, since he wants the firm to be solvent most of the time, and we assume that his information on the firm's past productivity behavior enables him not to choose ϕ too high.

We have considered two alternative, but similar, versions of a regulatory policy for the *recovery phase*. In the first of these policies, the regulator sets the price in each period so that the firm will achieve a fixed rate of return on capital, higher than the cost of capital. The resulting positive net profit will build up the cash reserve; when the cash reserve reaches its initial value, the recovery phase ends and a new incentive phase begins. In the second version, a fixed positive dollar net profit, rather than a fixed rate of return is set in each period. In this case the length of the recovery phase is determined[4]. In either version, the manager can affect prices by affecting productivity and hence costs. It can be shown that prices are not apt to fall as rapidly as the rate ϕ during the recovery phase; they may even increase. Thus if the system is to work at all well, the recovery phases must be infrequent, and short relative to the average length of an incentive phase; that this can indeed be the case is one of our results, which we now state.

Theorem 1.

Given that the regulator chooses a technologically feasible rate of price decrease, ϕ for the incentive phase,

Then as the manager's discount approaches zero (that is, as his preferences approach giving equal weight to present and future periods), the expected length of the incentive phase approaches infinity.

Since (at least for the second version of the regulator's policy during the recovery period) recovery periods are of fixed length, it follows form Theorem 1 that for a low managerial discount rate, almost all of the firm's time is on the average spent in the incentive phase[5].

Since, during the incentive phase, productivity is increasing enough for prices to decrease at the regulator's desired rate ϕ, it follows that the regulator's pricing objective is being achieved almost all the time. The next Theorem goes a little further in this direction.

Theorem 2.

Given the conditions of Theorem 1,

and that the maximum potential length of employment of each manager is large or infinite

and that the elasticity of demand is less than unity in absolute value,

Then as the manager's discount rate approaches zero, the long-run average rate of price decrease comes close to the regulator's incentive rate, ϕ, or some even greater value.

Note that the condition of less than unitary elasticity required for Theorem 2, is new. It may be possible to weaken this condition.

Theorem 2 states that the regulator's policy induces managerial behavior such that the regulator's pricing objective can approximately be met, not just during the long incentive phases but also in the long-run average over both incentive and recovery phases.

4. Remarks

In the model of Section 3, we have simplified the real regulatory situation in many ways. Some of these simplifications were made for ease of exposition, others for mathematical tractability, and still others to avoid problems of a more fundamental nature. In this concluding section we comment on several possible generalizations and extensions of the model.

Three extensions which seem not to be fundamental are the following:

a. We have supposed that the productivity consequences of R&D are all concentrated in the current period. Actually it is characteristic of R&D efforts that their benefits are realized after long and distributed lags, and that they are uncertain. Inclusion of this complication, although perhaps difficult, would not entail a basic change in the model.

b. We have included in the model the action the regulators take if the firm runs out of cash, but have not mentioned various possibilities for regulatory action if the firm prospers too greatly; i.e., if S_t and R_t become too large. Some action must be taken in this case, because:

   (i) enormous profits for the firm will be politically unpopular, even if the firm's prices are dropping rapidly;

   (ii) it would clearly be possible to distribute part of a large cash reserve to the customers either in the form of still further reduced prices or as a lump-sum refund.

It would be a natural extension to include this aspect of the regulators' strategy in the equilibrium behavior.

c. We have assumed that the firm's output just meets demand. The firm may provide several products, in which case we speak of output and price indices. It is easy to see how, by counting outputs of different quality as different outputs, quality can be included in the output index; thus when we assume demand is met we also assume quality is maintained.

In reality, of course, utilities may fail to meet demand, in a quantitative sense, and may fail to meet standards of quality. It would be reasonable to let such failures trigger a recovery phase, just as cash insolvency triggers it in our example. In fact, in one sense our model may be re-interpreted as already having this feature: namely, if there is in the model no reason other than insolvency for the firm to fail to meet demand or quality standards, then such failures are subsumed in the triggering mechanism we have already described.

Two much more fundamental extensions are the following:

d. We have assumed amounts of capital and other inputs can be optimally adjusted in each period. This ignores any irreversibility of capital investment by a regulated firm, as well as difficulties associated with laying off employees. Actually, the problem of how to best perform even "conventional" rate-of-return regulation when demand falls off and there is irreversible investment may not have been treated in the literature. To treat this requires a more general sequential principal-agent model.

e. We have mentioned, but not included in the model, the possibility of treating the distribution of information as endogenous. That is, we could include as a parameter of the regulators' strategy the amount of information they obtain, at a cost, about the firm's technological possibilities, and perhaps even about the manager's "private" utility.

In addition:

f. We have not attempted numerical estimates of the *magnitude* of the effects we are discussing. We have shown, in a simplified example, that there are equilibria of the sequential game (corresponding to productivity incentive clauses) that are more efficient than short-term equilibria (corresponding to conventional rate-of-return regulation), but how much more efficient are they?

   Such questions are difficult to answer, for two reasons. First, the technological possibilities open to the firm will (in the real world) usually be known to the managers only very incompletely, and will be even less well known to the regulators. Thus it is hard to estimate the maximum achievable growth in productivity, and hence the maximum achievable rate of price decrease. Second, it is hard to know how to estimate the manager's private utility (or disutility) for various actions, and hence how great a share of the cost-savings the managers must receive in order to be persuaded to act in a cost-minimizing way. (The managers may for example receive this share in the form of increased longevity in office, which requires a larger cash reserve, undistributed to customers.) Still, examination of the productivity growth of utilities as a function of the intervals between regulatory reviews might yield some empirical evidence on these points.

   A final consideration, which would take us into a rather different area of economic modeling is:

g. The sequential games we are considering have a multiplicity of equilibria. If one of these is better than all others for both the firm and society, it will be adopted. Otherwise, the choice of an equilibrium strategy pair will have to be established by bargaining between the regulators and the firm. Such bargaining mechanisms are not discussed in this paper.

REFERENCES

Houthakker, H. S. (1979). Growth and Inflation: Analysis by Industry. *Brookings Papers on Economic Activity 1*, 241-256.

Houthakker, H. S. (1981). Competition, Regulation and Efficiency. *In* Jules Backman (Ed.), "Regulation and Deregulation," pp. 35-48, Bobbs Merrill.

Linhart, P. B., and Sinden, F. W. (1982). Productivity Incentives Under Rate Regulation. *Bell Labs Economic Discussion Papers 236*.

Linhart, P. B., Radner, R., and Sinden, F. W. (1982). A Sequential Principal-Agent Approach to Regulation. Bell Labs Economic Discussion Paper (forthcoming).

Radner, R. (1981a). A Repeated Principal-Agent Game with Discounting. *Bell Labs Economic Discussion Papers 233*.

Radner, R. (1981b). Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship. *Econometrica 49*, 1127-1148.

Schumpeter, J. A. (1942). "Capitalism, Socialism, and Democracy." p. 84, New York: Harper and Row.

Stigler, G. T. (1971). The Theory of Economic Regulation. *Bell Journal of Economics 2*, 3-21.

REFERENCES

Houthakker, H. S. (1979). Growth and Inflation: Analysis by Industry. *Brookings Papers on Economic Activity 1*, 241-256.

Houthakker, H. S. (1981). Competition, Regulation and Efficiency. *In* Jules Backman (Ed.), "Regulation and Deregulation," pp. 35-48, Bobbs Merrill.

Linhart, P. B., and Sinden, F. W. (1982). Productivity Incentives Under Rate Regulation. *Bell Labs Economic Discussion Papers 236.*

Linhart, P. B., Radner, R., and Sinden, F. W. (1982). A Sequential Principal-Agent Approach to Regulation. Bell Labs Economic Discussion Paper (forthcoming).

Radner, R. (1981a). A Repeated Principal-Agent Game with Discounting. *Bell Labs Economic Discussion Papers 233.*

Radner, R. (1981b). Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship. *Econometrica 49,* 1127-1148.

Schumpeter, J. A. (1942). "Capitalism, Socialism, and Democracy." p. 84, New York: Harper and Row.

Stigler, G. T. (1971). The Theory of Economic Regulation. *Bell Journal of Economics 2,* 3-21.

FOOTNOTES

1. One such clause is in effect for Michigan Bell. The Michigan plan requires that a price index of the firm's output that drop in real terms by a fixed percentage every year (subject to a rate of return ceiling) to allow for expected productivity improvement. This allowance (initially 4%/year), along with the rest of the formula, is to be reviewed every three years.

2. Formally, if the manager receives a salary w per period and chooses action A_t in period t, then his utility for period t is

$$U_t = m_0(w) + m_1(A_t) \ ,$$

where m_0 and m_1 are two given functions that characterize the manager's preferences.

Let δ be the manager's discount rate, and T his tenure in his job (a random variable). Then we assume that he will choose his strategy to maximize his expected discounted utility for the duration of his employment, i.e.,

$$v = (1-\delta)E \sum_{t=1}^{T} \delta^{t-1} U_t \ ,$$

where the symbol E denotes mathematical expectation.

3. We mean either that the firm has a single output whose price decreases at this rate, or that the firm has multiple outputs and that an index of their prices decreases at this rate (in real terms, of course). We do not go into the complex econometric problems associated with the construction of such price indices, or with the conversion of prices from nominal to real terms.

4. It is of course necessary to specify the conditions under which a price exists which will achieve these results, so that the recovery phase is feasible.

5. The model that has been discussed so far is unrealistic in that it assumes that the potential length of the manager's employment is infinite. However, a result corresponding to

Theorem 1 can be derived in the case in which there is an exogenous upper bound on the length of the manager's employment. This result, while similar to Theorem 1, is somewhat technical to state, and we do not do so here.