A list of Working Papers on the last pages

No. 228, 1989

**DYNAMIC GAMES IN ORGANIZATION
THEORY**

by
Roy Radner

Paper prepared for **The Second International Workshop on Dynamic
Sciences**, June 5–16, 1989, at IUI, Stockholm

November, 1989
Revised February 1991

Dynamic Games in Organization Theory

Roy Radner

AT&T Bell Laboratories
Murray Hill, New Jersey
and
New York University

1. Introduction*

In any but the smallest human organizations, no one person has all of the information relevant to the organization's activities, nor can he directly control all of those activities. This is so even in organizations that are described as highly "centralized." It follows that individual members of the organization – I shall call them *agents* – have some freedom to choose their own actions. If, in addition, there is some divergence among the agents' goals or objectives, then one can expect some inefficiencies to arise in the organization's operations. The analysis of these inefficiencies, and the possible remedies by means of organizational design, is the subject of this paper.

If the behavior of the agents is "rational" in the sense typically used by economists and decision theorists, then the appropriate formal model would appear to be the theory of games, especially games of incomplete information, as developed in the past two decades.¹ This is the methodology that I shall use here, although some aspects of "bounded rationality" will be touched on during the course of my exposition. Furthermore, the relationships among members of an economic organization are typically long-lived, calling for an analysis of *dynamic games*.

Two special paradigmatic models have arisen in the game-theoretic analysis of organization. In the first, which I have elsewhere called a *partnership*, the agents act together to produce a joint outcome (output, profit). This outcome can be observed by the agents, but they cannot directly observe each others' actions, nor do they completely share each others' information. In the most

* The views expressed here are those of the author, and not necessarily those of AT&T Bell Laboratories or New York University. This is a revision of lecture notes prepared for the Second International Workshop on Dynamic Sciences, IUI, Stockholm, June 5-16, 1989. In preparing the present version I benefited from comments by Joseph A. Doucet, Richard H. Day, and Gunnar Eliasson. However, to satisfactorily answer Professor Eliasson's probing questions I shall have to do additional research.

1. See Harsanyi (1967,1968) and Myerson (1985). Complete references are gathered in the bibliography at the end of the paper, together with additional bibliographical notes.

general – and realistic – case of this model, the outcome is also influenced by random variables that are only partially observed, if at all. The incompleteness of observation leads to what the statisticians call a “confounding” of the sources of variation of the outcomes, making it difficult to assign responsibility to the individual agents for the occurrence of unsatisfactory outcomes. It is this confounding that leads to organizational inefficiency, if the goals of the agents are not identical (and not identical with the goal of the organization).

The second special model is suggested by the hierarchical structure of many organizations. In this model, there is a particular agent, called the “principal,” who performs no immediately useful actions himself, but “supervises” the activities of the other agents, rewarding them according to their individual outcomes (and other information), and retaining the residual (output or profit) for himself. This is the so-called “principal-agent” model, the word “agent” here denoting a member of the organization who is not the principal.

In fact, most organizations combine aspects of both the “partnership” and the “principal-agent” models. A hierarchy can be thought of as a cascade of principal-agent relationships, each supervisor acting as a principal in relation to the persons he supervises, and as an agent in relation to his own supervisor. On the other hand, in most cases the valued outcomes of organizational activity depend on the joint actions of several agents, as in the partnership model, so that the attribution of specific outcome variables to specific individuals (as required by the principal-agent model) may not be strictly justified. Unfortunately, I am not aware of significant progress on more comprehensive models of organization that combine these two submodels in a systematic way. This is one of the main challenges that organization theorists face today.

In the present exposition, I shall start with the partnership model. In fact, I shall start with the special case in which there is complete information and no uncertainty (Section 2). As in most of the models I shall discuss, the behavior of the agents is assumed to be a (Nash) equilibrium of the corresponding game, i.e., a combination of actions (or strategies), one for each agent, such that no agent can increase his own utility by *unilaterally* changing his own action. Even in this special case, it is typically true that in the static or one-period game the equilibria are *inefficient*. Efficiency here is defined in the Pareto sense; a combination of actions is *efficient* if there is no other combination that yields each agent at least as much utility, and yields at least one agent strictly more. On the other hand, if the partnership situation is repeated, leading to a dynamic game, then there will typically be many equilibria. Many of these dynamic equilibria may be inefficient; for example, the repetition of the one-period equilibrium will be a dynamic equilibrium. On the other hand, if the agents do not discount the future too heavily (are not too

impatient), then there will typically be equilibria of the infinitely repeated game that *are* efficient.

In Section 4 I introduce uncertainty and incomplete information into the partnership model.² Equilibria of the one-period game are again typically inefficient, but in contrast with the certainty case, in the repeated game one cannot guarantee the existence of efficient equilibria when the agents' discount rates are sufficiently small. Indeed, equilibrium outcomes may be *uniformly* bounded away from efficiency as the agents' discount rates approach zero.

I should point out here that the game played by the agents is not well-defined unless one specifies a particular rule for sharing the outcome among the agents (partners) as a function of the observed outcome. The specification of the the sharing rule is thus one of the design problems for the organization – or the organizer. From this perspective, the “uniform inefficiency” result alluded to above is quite strong, since it holds uniformly in the choice of sharing rule as well as in the agents' discount rate.

A special case of interest is the one in which the agents are neutral towards risk (Section 4.3). Here, with uncertainty, if the agents are suitably “different” then it is possible to design sharing rules such that an equilibrium of the corresponding game is efficient. These efficiency-inducing sharing rules must, however, be tailored to the agents' particular utility functions, which limits the practical implications of this result.

In Section 4.4 I explore the case in which the agents can change their actions rapidly. This is done by embedding the problem in a continuous-time model. For the risk-neutral case one can provide explicit calculations of the efficiency-inducing sharing rules, and show that the corresponding outcomes converge (in a particular sense) as the time between actions converges to zero. On the other hand, if the sharing rule divides the outcome among the partners in fixed proportions (which is a natural method), then there is an interval of time between actions sufficiently small so that for all smaller intervals the players cannot attain an efficiency higher than that of the corresponding (inefficient) static equilibrium.

In Section 5 I turn to the principal-agent model. The exposition here parallels that of the partnership model, starting with the static case and moving to the repeated-game formulation. The

2. Because of limitations of space and time, I shall discuss moral hazard but not adverse selection or strategic misrepresentation of information. For an explanation of this distinction, see Section 4.

latter, however, provides a contrast to the partnership. Here, as the players' discount rate approaches zero, one can find equilibria of the repeated game that approach efficiency in the limit. Such approximately efficient equilibria can be characterized in some detail, depending on the specific model, and have interesting behavioral interpretations. Again, embedding the problem in a continuous-time model allows one to obtain sharper characterizations of the equilibrium strategies.

In fact, these equilibria of the principal-agent game lead to optimization problems for the agent that might be called problems of "survival." This prompts me to devote a special section (7) to the study of such problems, which also have an independent economic interest outside of the field of organization theory.

Finally, it should be recognized that in realistic settings the organizational decision problems are not strictly repeated. Typically, there are one or more state variables that evolve in response to both organizational activities and exogenous random variables; for example, this is characteristic of situations involving investment. Although a comprehensive theory is not yet available, I illustrate this phenomenon in Sections 3 and 6, as well as in the section on economic survival. In Section 3, I discuss a partnership model – with certainty – of the joint exploitation of a productive asset as exemplified by "fishing wars." In Section 6, I sketch a principal-agent model of the regulation of a public utility, in which the principal is the regulator, and the agent is the firm's manager who is engaging in risky research and development with the goal of reducing costs. In both cases the methods used for the repeated-game case can be extended and adapted to construct efficient – or approximately efficient – equilibria.

In this exposition, I shall not attempt any great level of generality. Instead, I shall illustrate the key ideas with a series of elementary mathematical examples, only sketching the directions in which further analysis has been successful. The interested reader may consult the corresponding references for treatments of greater depth and generality.

2. Simple Partnerships with Certainty

2.1 Introduction

In a simple partnership game, with certainty, the output of the partners is jointly determined as a function of the individual inputs of the partners. This output is divided among the partners according to some fixed rule. The utility to each partner in any one period is the difference between his compensation (i.e., his share of the output) and the cost (or disutility) of his input.

If the situation is repeated, then the resulting game is called a *supergame*. In the supergame, a strategy of a partner is a sequence of functions, one for each period, that determines his input in each period as a function of the history of all inputs and outputs in all previous periods. His utility for the supergame is the sum of his one-period utilities, typically discounted at some fixed (exogenous) rate. (In a variation on this definition, one may prohibit the partners from ever observing the inputs of the other players. This variation will be considered in Section 4 below.)

In the one-period game, a strategy for a partner is simply an input – a single nonnegative number. The description of the game is completed by specifying the rule according to which the output is shared among the partners. For example, the sharing rule might specify that the output is to be shared equally among the partners. A combination (vector) of inputs is an *equilibrium* (or noncooperative Nash equilibrium) if no individual partner can increase his utility by unilaterally changing his input. A combination of inputs is *efficient* (Pareto optimal) if no other combination of inputs yields each partner at least as much utility, and yields at least one partner strictly greater utility.

With natural assumptions about the output function and the individual cost functions, it is intuitively plausible that an equilibrium cannot be efficient. For example, suppose that the partners share the output equally. At an equilibrium, a small increase in one partner's input will result in an increase in his compensation that is approximately matched by the corresponding increase in his individual cost. On the other hand, the small increase in his input will also increase the compensation of every other partner, without corresponding increases in their own costs. Thus, starting from an equilibrium, a small increase in each partner's input will make all the partners better off.

In economic jargon, each partner's input produces a positive "externality" for the other partners, which he does not take into account in his own (equilibrium) behavior. Another way of putting it is that each partner tries (up to a point) to be a *free rider* on the inputs of the other partners. The result in equilibrium is that each partner's input is smaller than it should be for efficiency.

Example 2.1. Suppose there are two partners, and denote partner i 's input by a_i . The corresponding output is

$$y = R(a_1 + a_2) , \quad (2.1)$$

and i 's share of the output is $S_i(y)$, where for every y ,

$$S_1(y) + S_2(y) = y . \quad (2.2)$$

Here a_1 and a_2 must be nonnegative, and R is a positive constant. Denote i 's individual cost by $Q_i(a_i)$, then i 's utility is

$$u_i = S_i(y) - Q_i(a_i) . \quad (2.3)$$

In particular suppose that

$$S_i(y) = \frac{y}{2} , \quad (2.4)$$

$$Q_i(a_i) = q a_i^2 ,$$

in which case

$$u_i = \frac{R(a_1 + a_2)}{2} - q a_i^2 . \quad (2.5)$$

A one-period equilibrium is characterized by the first-order conditions:

$$\frac{R}{2} - 2 q a_i = 0 , \quad i = 1, 2 ,$$

so that the equilibrium inputs and utilities are

$$a_i^* = \frac{R}{4q} , \quad (2.6)$$

$$u_i^* = \frac{3R^2}{16q} .$$

To characterize the efficient input combinations, first note that if the utility pair (u_1, u_2) is feasible, then so is any pair (u'_1, u'_2) with

$$u'_1 = u_1 + s ,$$

$$u'_2 = u_2 - s .$$

Hence (a_1, a_2) is efficient if and only if it maximizes

$$\begin{aligned}
 u_1 + u_2 &= S_1(y) - Q_1(a_1) \\
 &\quad + S_2(y) - Q_2(a_2) \\
 &= y - Q_1(a_1) - Q_2(a_2) \\
 &= R(a_1 + a_2) - q a_1^2 - q a_2^2 .
 \end{aligned}$$

This uniquely determines the efficient inputs $(\hat{a}_1, \hat{a}_2)$,

$$\hat{a}_i = \frac{R}{2q}, \quad i = 1, 2, \quad (2.7)$$

as well as the *sum* of the utilities,

$$\hat{u}_1 + \hat{u}_2 = \frac{R^2}{2q}. \quad (2.8)$$

The various efficient utility pairs $(\hat{u}_1, \hat{u}_2)$ are now determined by varying the sharing rule. In particular, the family of sharing rules

$$S_1(y) = \frac{y}{2} + s, \quad (2.9)$$

$$S_2(y) = \frac{y}{2} - s,$$

yields all the efficient utility pairs by varying the parameter s , provided the partners use the inputs $\hat{a}_1$ and $\hat{a}_2$.

Comparing (2.6) with (2.7) and (2.8), we see that the efficient inputs are twice as large as the equilibrium inputs, and that the efficient sum of utilities is 4/3 times the equilibrium sum of utilities. Of course not *every* efficient utility pair is Pareto-superior to the equilibrium; see Fig. 1, where the efficiency frontier is the line with slope -1 .

2.2 The Repeated Game

Suppose now that the situation of the one-period game is repeated infinitely often. As described above, a partner's strategy determines his input in each period as a function of the previous history of inputs and outputs. A partner's utility for the supgame is the sum of his discounted one-period utilities.

Extending the notions of equilibrium and efficiency to the supgame in the obvious way, one can show that the supgame typically has many equilibria, some of which may be efficient.

For example, consider a particular equilibrium of the one-period game, and define the *stubborn strategy* of a partner to be the one in which he plays his one-period-equilibrium input in every period, no matter what the previous history is. It is easy to see that the strategy combination in which each partner plays his stubborn strategy is an equilibrium of the supergame; I shall call this the *stubborn equilibrium*. (To each equilibrium of the one-period game, there will correspond a stubborn equilibrium of the supergame.) If the one-period equilibrium is inefficient, then so is the stubborn equilibrium.

To construct efficient equilibria, I shall now consider so-called *trigger strategies*. Having singled out, as before, some (inefficient) one-period equilibrium, let us also single out some efficient input combination that makes every partner better off; call this the “target input” combination, and call the corresponding output the “target output.” A *trigger strategy* for a partner is defined as follows: the partner uses his own target input until the first period, if ever, in which some partner does not use his corresponding target input; thereafter he uses his stubborn strategy.

In order for the combination of trigger strategies to form an equilibrium, it must be that case that, for each partner, the one-period gain he gets from deviating (optimally) from the target input combination is less than the subsequent loss due to everyone switching to their respective stubborn strategies. However, as the partner’s discount rate approaches zero, the ratio of his one-period gain to his subsequent loss also approaches zero. Thus, for sufficiently low discount rates the trigger-strategy combination will be an equilibrium of the supergame.

Example 2.2. Extending Example 2.1, let a_{it} and u_{it} denote i ’s input and utility, respectively, in period t ($t = 1, 2, \dots$, and ∞); then i ’s supergame utility is

$$U_i = \sum_{t=1}^{\infty} (1-\delta) \delta^{t-1} u_{it} , \quad (2.10)$$

where δ , the *discount factor*, is between 0 and 1. (The *discount rate* is defined as $(1-\delta)/\delta$.) Note that both partners have the same discount factor.

Let $\hat{a}_i$ and a_i^* be defined as in Example 2.1 (efficient and equilibrium inputs). Let the *target utilities* be (cf. (2.8))

$$\hat{u}_i = \frac{R^2}{4q} , \quad i = 1, 2 ,$$

so that the total efficient utility is divided equally. Define i ’s *trigger strategy* by:

- 1) $\hat{a}_{i1} = \hat{a}_1$;
- 2) if for some j and t , $a_{jt} \neq \hat{a}_j$, then $a_{is} = a_i^*$ for all $s > t$.

Consider partner 1's optimal deviation from the trigger-strategy pair; without loss of generality, we can take this to occur in period 1. If 1 deviates in period 1, then 2 will use a_2^* from period 2 on, and so 1 should use a_1^* from period 2 on. Therefore 1's optimal deviation in period 1 is his optimal one-period input given that 2 uses $\hat{a}_2$; it is easy to verify that this is a_1^* , and that 1's corresponding one-period utility is

$$u'_{11} = \frac{5R^2}{16q} > \hat{u}_1 .$$

(The fact that a_1^* is optimal against $\hat{a}_2$ is special to this example.) Hence if 1 deviates optimally in period 1 his supgame utility will be

$$\begin{aligned} & (1-\delta)u'_{11} + \sum_{t=2}^{\infty} (1-\delta)\delta^{t-1}u_1^* \\ &= (1-\delta)u'_{11} + \delta u_1^* . \end{aligned} \tag{2.11}$$

On the other hand, if he stays with his target input, $\hat{a}_1$, then his supgame utility will be

$$\sum_{t=1}^{\infty} (1-\delta)\delta^{t-1} \hat{u}_1 = \hat{u}_1 . \tag{2.12}$$

Since $\hat{u}_1 > u_1^*$, (2.12) will exceed (2.11) when δ is sufficiently close to 1. Hence for δ sufficiently close to 1, it will not be optimal for 1 to deviate from the trigger-strategy pair, and so the latter is an equilibrium of the supgame.

Note that the above argument is quite general, since it used only the fact that $\hat{u}_1 > u_1^*$. Since the one-period equilibrium is in general inefficient, it will be possible to find an efficient pair of one-period utilities $(\hat{u}_1, \hat{u}_2)$ that is Pareto-superior to the equilibrium (u_i^*) , and any such pair can be sustained in a supgame equilibrium for δ sufficiently close to 1.

In fact, using other methods, it can be shown that the set of equilibria of the supgame is quite large. Define the *utility outcome* of a game to be the vector of the players' utilities (one for each player). A utility outcome is *feasible* if it is yielded by some combination of strategies, and it is *individually rational* if it gives each player at least as much utility as he could "guarantee" for himself. One can show that, under fairly general conditions, as the partners' discount factor approaches 1, the set of equilibrium utility outcomes of the supgame approaches the set of

feasible and individually rational utility outcomes of the one-period game. (This result is true for a large class of repeated games; see Fudenberg and Maskin (1986) and the references cited there. There is a corresponding result for the limit case in which each player's supgame utility is the long-run average of his one-period utilities; this is sometimes called the "Folk Theorem" for repeated games.)

3. A Partnership Game with Investment: 'The Dynamic Inefficiency of Capitalism'

As noted in the Introduction, for most "realistic" models of long-term relationships the one-period game is not strictly repeated; rather there are one or more state variables that evolve through time as a function of the players' actions and possibly exogenous factors. In this section I shall illustrate this phenomenon with a model of the joint exploitation of a productive and producible asset.³

The phrase "tragedy of the common" evokes an image of an overgrazed pasture used in common by many husbandmen. By extension, it refers to a situation in which a productive asset is exploited jointly by several economic agents whose "noncooperative" behavior results in an overexploitation of the asset, i.e., an exploitation that is not efficient. Other than grazing, examples of this situation included fishing, forestry, and hunting. A novel example, and the one that first attracted the attention of J. Benhabib and myself, was studied by Lancaster (1973), who viewed the assets of a modern capitalist firm as being jointly exploited by the firm's owners and its unionized workers. For various reasons, the owners and the workers cannot or do not bind themselves to long-term cooperative behavior, leading to what Lancaster called "the dynamic inefficiency of capitalism."

Following the direction suggested by the work of Lancaster, Levhari and Mirman (1980), and others, Benhabib and I analyzed a fairly general model of the joint exploitation of a productive asset as a dynamic, noncooperative game. Here the state variable is the stock of the productive asset, which changes through time as a result of its own productivity and the actions of the players. Our goal was to understand the variety of equilibria of this game, and in particular to understand the conditions under which there are equilibria that are efficient.

3. This section is based on Benhabib and Radner (1988).

In the continuous-time model that we study, the (positive) stock at date t , $Y(t)$, evolves according to the differential equation,

$$Y'(t) = \eta[Y(t)] - c_1(t) - c_2(t) ,$$

where (for the case of two players), $c_1(t)$ and $c_2(t)$ are the rates of consumption of the asset by players 1 and 2, respectively. The “production function,” η , is assumed to be concave, and to take the value zero at both zero and some positive stock level. The strategy of each player determines his consumption rate at each time as a function of the previous history of the process, possibly with some delay. We assume that each player’s utility for the game is equal to his total discounted consumption over the duration of the game. The game ends when the stock becomes zero, if ever. The linearity of a player’s utility in his consumption is the main special assumption of the model. We also assume that each player’s rate of consumption is nonnegative and bounded.

At an efficient equilibrium the weighted sum of the players’ total utility is maximized. Since the instantaneous utilities of the players are linear in consumption, this is equivalent to maximizing the discounted sum of the total consumption of the players. We show that the efficient consumption policy of the two players is to consume nothing until a certain level of the stock is reached. After that the total consumption of the players is equal to the output of the stock, so that the stock level is stationary. We call a consumption policy of this type a “frugal” policy. By contrast, if a player follows a “prodigal” consumption policy he always consumes at the upper bound of his consumption rate.

The equilibria of this dynamic game that correspond to the repeated static equilibria are those in which each player uses a strategy in which his action at any date is independent of the current stock of the asset; we might call these “extreme equilibria.” In these equilibria, the players run down the stock of the asset as fast as possible (“prodigal” consumption). By analogy with the terminology of repeated-game theory, we define a *trigger strategy equilibrium* to be a Nash equilibrium in which the players threaten to revert to an extreme equilibrium whenever a player is caught deviating from the target efficient path. The effectiveness of such threats depends, of course, on the “detection technology,” i.e., on how much extra utility the deviating player can gain before his deviation is detected by the other players. In the model we study, efficient trigger-strategy equilibria may exist from some starting states but not others. More precisely, there is a stock level, say y' , such that a trigger-strategy equilibrium exists from starting stocks greater than or equal to y' , but not from those strictly less than y' . (This statement is meant to include the cases in which y' is zero or infinite.)

Under some circumstances, there may exist a new kind of equilibrium, which we call a *switching equilibrium*. We show that, in our model, whenever y' is positive (and finite), there is an open interval I with upper endpoint y' such that, from any starting stock in I there is an equilibrium of the dynamic game with the following structure: the players follow an inefficient but growing path until the stock reaches the level y' , and then follow a trigger strategy (efficient) after that.

An important feature of our analysis is an explicit modelling of delayed information. In our treatment of trigger-strategy and switching equilibria we assume that each player can observe the state of the system (the stock of the asset) with a fixed delay, i.e., at time t each player can observe the *history* of the state variable up through time $(t-\tau)$, where the delay τ is a fixed, positive parameter of the model. The larger the delay, the more a player can benefit from a “defection” from a prescribed path before his defection is detected and the other player can respond. In previous discrete-time models, this delay has been implicitly equated to the length of the period between decision times. The use of a continuous-time model makes it convenient for us to vary the delay, τ , as an independent parameter, and we consider this to be an important contribution of our analysis.⁴

In fact, we show that, roughly speaking, for any fixed discount rate, (1) efficiency can be sustained by trigger-strategy equilibria from any positive initial stock, provided that the delay is sufficiently small, but (2) efficiency cannot be so sustained from any positive initial stock, provided that the delay is sufficiently large. A corresponding result holds for a fixed positive delay, as one varies the discount rate.

4. Simple Partnerships with Moral Hazard

4.1 Introduction

In the present section I introduce three new features into the model of a simple partnership that was described in Section 2:

4. For an analysis of information processing as a source of delay, see Radner (1989).

1. The joint output is influenced by exogenous random factors (the environment), as well as by the partners' inputs.
2. The partners cannot observe the random environmental factors.
3. The partners cannot observe each others' inputs.

A consequence of these new features is that, by observing the output alone, the partners cannot infer with certainty the cause of any departure from some "target" output. This situation is sometimes described as one of *moral hazard*.

In a more realistic model, features 2 and 3 above would be relaxed to allow for imperfect observation of the environment and the actions of other partners. In particular, if different partners had different information about the environment, then phenomena such as adverse selection, self-selection, misrepresentation, etc., might arise. For simplicity, however, I shall restrict my attention to the case of moral hazard.

If one assumes that the objective of each partner is to maximize his own expected utility, then the introduction of the above features does not essentially alter the analysis of the one-period game. On the other hand, the nature of the repeated game is changed in a fundamental way, as we shall see below.

Example 4.1. Modify Example 2.1 so that the outcome is a random variable, say Y , whose probability distribution depends on a_1 and a_2 . In particular, suppose that Y can take on only two possible values, y_1 and y_0 , with

$$\text{Prob}(Y = y_1) = \min \{ \alpha(a_1 + a_2), 1 \} , \quad (4.1)$$

and a_1 and a_2 are nonnegative. (Think of a_i as i 's "effort.") Without loss of generality, we may take $y_1 = 1, y_0 = 0$, and $\alpha = 1$.

Let s_{iy} denote partner i 's compensation if the outcome is y ,

$$\begin{aligned} s_{11} + s_{21} &= 1 , \\ s_{10} + s_{20} &= 0 . \end{aligned} \quad (4.2)$$

Partner i 's utility is assumed to be linear in compensation and quadratic in effort:

$$U_i = s_{iy} - qa_i^2 .$$

Hence his expected utility is, by (4.1)

$$\begin{aligned} u_i &= s_{i1}(a_1 + a_2) + s_{i0}(1 - a_1 - a_2) - q a_i^2 \\ &= (s_{i1} - s_{i0})(a_1 + a_2) - q a_i^2 - s_{i0} , \end{aligned} \quad (4.3)$$

(provided that $a_1 + a_2 \leq 1$). For example, if $s_{i1} = \frac{1}{2}$ and $s_{i0} = 0$, then

$$u_i = \frac{1}{2} (a_1 + a_2) - q a_i^2 . \quad (4.4)$$

Notice the formal similarity between (4.4) above and (2.5) in Example 2.1, the latter with $R = 1$. It follows that the analysis of efficiency and equilibrium in this example is the same as in Example 2.1.

With the introduction of uncertainty, one should take account of the attitudes of the players towards risk. In Example 4.1 the partners are represented as neutral towards risk, but of course this is not the general case. The special implications of the assumption of risk-neutrality will be explored in Sections 4.3 and 4.4.

4.2 Optimal Sharing Rules with Risk-Neutrality

In Section 2.1 I argued heuristically that, in the certainty case, an equilibrium cannot be efficient. That argument was based on the assumption that the partners shared the output equally. In fact, one can show that under quite general conditions (with certainty) *there is no sharing rule for which a corresponding equilibrium is efficient.*

With the introduction of uncertainty, the situation is changed. Following Williams and Radner (1988), in this section I shall sketch an argument to show that, if the number of possible outputs is at least 3, and if the partners are neutral towards risk, then – generically in the data of the game – there exists a sharing rule for which the corresponding equilibrium of the one-period game is efficient.

Basically, what is required is that the partners be sufficiently different in the effects that their actions have on the distribution of output. On the other hand, if the partnership is symmetric with respect to permutations of the partners, then an efficiency-inducing sharing rule will typically not exist. (Generically, the partnership will not be symmetric.)

I begin by describing the Williams-Radner model, which is a generalization of Example 4.1. There are $m > 1$ partners. The i^{th} partner chooses his input a_i from some closed and bounded subinterval A_i of the real line. His choice is his own private information. Let $a \equiv (a_1, \dots, a_m)$

denote an input profile, and let a_{-i} denote the $(m-1)$ -tuple $(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_m)$.

Once the partners have chosen their inputs, one of several levels of output results. This output is publicly observable. Let $\mathbf{Y}$ denote the range of output levels of the partnership. Except where otherwise noted, the reader should assume that Q is some finite set of real numbers with $n \geq 2$ elements,

$$\mathbf{Y} = \{y_1 < y_2 < \dots < y_n\}.$$

The partners' inputs determine a probability distribution over $\mathbf{Y}$. For the input profile a , let $F(\cdot, a)$ denote the cumulative distribution that is determined by a , and let $f(\cdot, a)$ denote the corresponding density function. These functions are common knowledge, and each is a C^1 function of the inputs. For simplicity, let $F_i(y, a) \equiv \frac{\partial F}{\partial a_i}(y, a)$ and $f_i(y, a) \equiv \frac{\partial f}{\partial a_i}(y, a)$.

The i^{th} partner's utility u_i consists of whatever share $s_i(y)$ he receives of the observed output y , minus the disutility $Q_i(a_i)$ of his contribution of the input a_i :

$$u_i \equiv s_i(y) - Q_i(a_i).$$

By assumption, $Q_i(\cdot)$ is a C^1 function of the i^{th} partner's contribution. Let $q_i(\cdot) \equiv \partial Q_i / \partial a_i(\cdot)$. We assume that $q_i(\cdot)$ is strictly positive.

Since utility is transferable, an input profile $\hat{a}$ is Pareto optimal if and only if it maximizes the difference between the expected total output and the total disutility of the input contributions:

$$\hat{a} \in \operatorname{argmax}_a \left\{ E(y|a) - \sum_{i=1}^m Q_i(a_i) \right\}. \quad (4.5)$$

We assume that there exists a solution to this maximization problem in the interior of $\prod_{i=1}^m A_i$.

Efficiency therefore requires each partner to make a positive input contribution.

Our concern is the existence of a sharing rule $s_1(\cdot), \dots, s_m(\cdot)$, that satisfies the budget constraint

$$\sum_{i=1}^m s_i(y) = y, \quad \text{all } y \in \mathbf{Y}, \quad (4.6)$$

and that also makes the efficient profile $\hat{a}$ into a Nash equilibrium,

$$\hat{a}_i \in \operatorname{argmax}_{a_i} \left\{ E(s_i(y) \mid (a_i - \hat{a}_{-i})) - Q_i(a_i) \right\} \quad \text{for } 1 \leq i \leq m.$$

The problem of devising a sharing rule with these properties is the *partnership problem*. The following first order conditions are necessary: if $\hat{a}$ is efficient, then for each partner the marginal expected *total output* must equal the marginal disutility of his contribution at $\hat{a}$,

$$\sum_{j=1}^n y_j f_i(y_j, \hat{a}) = q_i(\hat{a}_i) \quad \text{for all } 1 \leq i \leq m. \quad (4.7)$$

On the other hand, if $\hat{a}$ is a Nash equilibrium, then each partner's marginal expected *compensation* must equal the marginal disutility of his contribution at $\hat{a}$,

$$\sum_{j=1}^n s_i(y_j) f_i(y_j, \hat{a}) = q_i(\hat{a}_i) \quad \text{for all } 1 \leq i \leq m. \quad (4.8)$$

The problem of devising a sharing rule that satisfies the first order condition (4.8) and the budget constraint (4.6) is the *first order problem*.

Our approach is to solve the first order problem and then to determine whether or not this solution also solves the partnership problem. The main result is that the first order problem is solvable for a generic choice of $F(\cdot)$ and $Q_1(\cdot), \dots, Q_m(\cdot)$ when $F(\cdot, a)$ defines a probability distribution over at least three output levels. Efficiency plays a relatively minor role in the proof; we actually prove the stronger result that in a generic problem with at least three output levels, the budget constraint and the first order conditions for a Nash equilibrium are solvable at a *generic input profile*. (In other words, for a generic input profile, there will be a sharing rule for which the input profile is an equilibrium.)

To get an idea of the proof, suppose that there are n possible output levels y_j . The “unknowns” in equations (4.6) and (4.8) are the mn numbers $s_i(y_j)$. Thus we can regard (4.8) and (4.6) as a system of $m+n$ linear equations in the mn variables $(s_i(y_j))_{1 \leq i \leq m, 1 \leq j \leq n}$ whose coefficients are determined by the efficient profile $\hat{a}$. The left-hand sides of the first m equations (from (4.8)) are the marginal expected payments to the partners, and the last n equations (from (4.6)) form the budget constraint. Efficiency is used in the analysis of the $n \geq 3$ case only to determine the values of these coefficients; the argument could be carried out using coefficients determined by a generic input profile. The main task of the proof is to show that, generically in the data of the problem, these equations have full rank, provided $n \geq 3$. On the other hand, it is not difficult to show that, if $n=2$, or if the equations are symmetric in the partners, then no solution is possible.

It is easy to construct linear examples in which every sharing rule that satisfies both the budget constraint and the first order conditions for a Nash equilibrium also solves the partnership problem. More generally, Williams and Radner derive conditions on the output function under which at least one solution to the first order problem also solves the partnership problem. For simplicity, they derive these conditions only when there are two partners and three levels of output. The conditions are awkward, and at present they have no economic interpretation. "Reasonable" examples exist, however, that satisfy them (see Section 4.4).

We also prove a paradoxical result that concerns the nature of the solutions to the first order conditions when the output function satisfies stochastic dominance with respect to each partner's input. (For each partner, given the inputs of the other partners, stochastic dominance holds if the observation of a higher level of output allows one to infer, in a probabilistic sense, that the selected partner contributed a greater level of input; e.g., see Whitt (1980).) When stochastic dominance holds, one might expect that a partner's payment should increase with the output; as Alchian and Demsetz (1972, p. 778) suggested in their analysis of the internal structure of firms, a partner may have an incentive to "sabotage" the organization if his reward and the output are inversely related. In fact, the opposite is true: when stochastic dominance holds, for any sharing rule that satisfies the first order conditions, some (at least two) of the partners' payments must be nonincreasing over some subsets of the range of outputs levels. Thus moral hazard can be overcome in some problems in which stochastic dominance holds, but only if some partners do not always benefit when the joint output increases.

This paradox may explain why these results seem surprising, and why they have been overlooked in the literature on partnership. One can show that the Nash equilibria of partnerships are typically inefficient when the budget is balanced, provided that (i) each partner's payment increases with the output; (ii) for each state of a random environment, the output is an increasing function of each partner's input. This second assumption implies that the output function satisfies stochastic dominance.

Each of the above results can be extended to the case where the set of output levels is a subinterval of the real line. Thus it is possible to solve the partnership problem in our model, not because of any special assumption about the range of output levels, but because the joint output is uncertain.

This section focused on the case in which the partners are risk neutral. It can be shown that, in a generic problem with risk aversion, the partnership problem is unsolvable.

4.3 The Repeated Game

I have alluded above to the fact that the introduction of moral hazard into the partnership situation changes the repeated game in a fundamental way. The trigger strategies described in Section 2.2 can no longer be completely effective (even with discount factors close to 1), because departures from the target output can be caused by random variations in the environment as well as by deviations of the partners' inputs from their target values.

However, the partners are not completely powerless to monitor each others' inputs, since they will have *statistical* evidence from the sequence of observed outputs. For example, suppose that, over time, the random environmental factors are independent and identically distributed. It follows that, if each partner uses his target input in each period, then the sequence of outputs will be independent and identically distributed as well; call this the *target distribution* of the outputs. Each partner could now use a statistical procedure to test whether the other partners are adhering to their target inputs, in a manner analogous to the statistical quality control of a production process. A "failure" of the test would trigger a reversion by all of the partners to their respective stubborn strategies.

Notice, however, that a procedure that had any chance of detecting deviations from target inputs would also produce a "false alarm" from time to time. In other words, even if the partners always used their target inputs, there would be a positive relative frequency of test failures, so that the reversion to stubborn strategies would eventually be triggered, with probability one. The inefficiency caused by this could be mitigated, but not entirely eliminated, by making the reversions to stubborn strategies last only a finite length of time (so-called "relenting" strategies).⁵ In this case, there would be an infinite sequence of "phases" of two types, one in which the partners used their target inputs, and one in which they used their stubborn strategies.

Of course, this does not settle the question whether, as the players' discount factor approaches 1, efficiency can be approached by equilibria of the supergame. The conditions for this to be true are somewhat complicated, and I shall not attempt to describe them precisely here. Roughly speaking, what is required is that, for every pair of partners, the probability distributions of

5. Note that relenting strategies could also have been used in the case of certainty, but would not have produced any further increase in efficiency.

outcomes corresponding to different pairs of deviations from the target inputs should be linearly independent.⁶

4.4 A Continuous-Time Model

In this section I summarize the results of a study by Radner and Rustichini (1988) of a partnership model with uncertainty, in a framework that allows the analysis of the effect of varying the reaction time of the partners, including the limiting case of instantaneous adjustment. The one-period outputs are normally distributed random variables, with means and variances depending on the inputs of the partners. The sequence of outputs is a stochastic process of Wiener type, which can be thought of as the discretization of a diffusion-type process. As the reaction time tends to zero this process tends to the solution of a stochastic differential equation. The sample paths are (almost surely) continuous. It may be objected that in real-world partnerships the reaction time and the flow of information are always, for practical purposes, different from zero; so “real” partners can not adjust instantaneously all the time. In fact, the same objection may be raised against any model of a dynamic game in continuous time. But since a universal lower bound on the reaction time would certainly be artificial, the question arises whether the properties of the set of equilibria and of the strategies approach some limit when the reaction time becomes arbitrarily small. A related nontrivial question is the existence of equilibria. In other words, the analysis of a continuous-time model may be considered as a way of testing the robustness of the results for a discrete time (finite reaction) model; the study of the limit situation should clarify which properties of a discrete time model depend critically on the fixed delay in the reaction of the players when that delay is “small.”

The first question we analyze is the characterization of *efficient sharing rules*, i.e., sharing rules that have the efficient outcome as a (Nash) equilibrium. This question was first examined in Williams and Radner (1987), where the generic existence of efficient sharing rules was demonstrated in a class of partnership models (see Section 4.3). Like Williams and Radner, we assume that the partners are risk-neutral. For our model, we provide a complete characterization of the partnerships for which the design of efficient sharing rules is possible, and a characterization of such rules. This characterization has a particularly simple formulation in the

6. See Fudenberg, Levine, and Maskin (1989). I should mention here that their analysis makes use of a more general class of supgame strategies than those described in the preceding paragraph.

case where the outcome is a random variable with a normal distribution. Simply stated, the condition requires that at least two of the partners are different enough, in the sense that the variance and the mean of the outcome vary differently as the efforts of these two partners vary. It is interesting to note that the preferences of the players (in our case, the cost or disutility of the input) plays no role. We then present a general procedure to design such sharing rules. From this very construction, it will be apparent that the set of possible sharing rules is very large.

We then examine the problem of the existence of a limit for the optimization problem of each partner. The existence of such a limit is important from the point of view of the robustness of the equilibrium, as the repeated game becomes (in the limit) a continuous-time game. In fact the existence of such a limit is a necessary condition for the concept of equilibrium to be well defined. We prove that it is always possible to construct sharing rules that are both efficient and stable (with respect to this limit process). Indeed, very simple sharing rules can be formed, even with quadratic functions.

Lastly, we discuss the performance of fixed-proportion sharing rules, as the reaction time tends to zero. We examine the case of two identical partners, with the sharing rule given by equal splitting of the outcome, and examine upper bounds on the efficiency of symmetric equilibria. The main result is that, when the reaction time becomes shorter than a *fixed positive quantity* the only equilibrium of the repeated game is the equilibrium of the one-period game. A similar question has already been examined, with similar conclusions, in the paper of Abreu et al. (1987), for the not necessarily symmetric case. However, in their model the outcome is a stochastic process of Poisson type (rather than of Wiener type, as in our case), and the action space of the partners consists of two points. The model of the present paper and the one in Abreu et al. thus cover together all processes in continuous time for which: (1) increments over nonoverlapping time intervals are independent, (2) sample paths have at most discontinuities of the first kind, and (3) for any fixed time t , the sample paths are continuous in probability at t . (Such processes are called Lévy processes; see Ito (1985) for an analysis of such processes and a proof of the fact that Lévy processes are compositions of constant processes, Wiener processes and Poisson processes.)

5. Principal-Agent Games

5.1 Introduction

I turn now to the principal-agent model, which is suggested by the hierarchical or supervisory relationships that are common in organizations. From a formal point of view, we may consider

the principal-agent model as a special case of the partnership model, in which one of the partners, called the “principal,” effectively takes no action (formally, the outcome is independent of his action); the other partners are called the “agents.” In fact, most of the literature deals with the case of only one “agent,” which is the case I shall discuss here.⁷ Also, most analyses assume that it is the principal who chooses the compensation function (sharing rule), subject to some constraints; this choice becomes the strategy of the principal. This is the approach I shall follow in the present section.

To summarize, we shall be considering the following situation. The “enterprise” comprises the principal and the agent. The output of the enterprise depends on the agent’s action and on a stochastic environment, but the principal cannot fully monitor the agent’s information and action, nor can he fully monitor the environment. The principal can monitor the outcome, however, and in the simplest form of the principal-agent model – the one we shall study here – this is the only thing he can monitor. Thus in this simplest case the principal can make the agent’s compensation depend at most on the outcome. More generally, the compensation can depend on anything else that the principal can observe, e.g., some incomplete information about the agent’s information, action, or environment.

Table 1 lists some principal-agent relationships that can be more or less accurately represented by the general principal-agent model. The insurer-insured relationship is the one that gave rise to the term “moral hazard.” The action of the insured (agent) is the care he takes to prevent an accident (say to property), and the outcome is the occurrence or nonoccurrence of the accident. The compensation that the principal (insurer) pays to the agent is negative (the premium) if the accident does not occur, and is typically positive (the claim minus the premium) if the accident does occur. If the preventive care is costly to the agent, then the fact that he has insurance may lead him to lower his level of care, and this is the phenomenon called *moral hazard*. In this relationship, the insured party is the *agent*, since he is the actor whose actions (care) are unobserved, and the insurer is the *principal*, who compensates the insured according to the outcome.

7. This is clearly a limitation from the point of view of organization theory. See Radner (1987) and Groves (1973) for a more general discussion.

Although much of the literature on the principal-agent model refers to market or regulatory relationships, my concern here will be primarily with principal-agent relationships within organizations, such as those listed above the dotted line in Table 1.

In this section I shall use a simple example of a one-period principal-agent model to illustrate how moral hazard can lead to inefficiency. Suppose that the (stochastic) outcome of the enterprise is either "success" or "failure," and that the probability of success depends on the agent's action. In the case of success, the principal earns one unit of money (say one million dollars), but in the case of failure, he earns nothing. The principal will compensate the agent according to the outcome, giving him a compensation of w_1 for a successful outcome and a compensation of w_0 for a failure. (In principle, a compensation may be negative, although institutional constraints might rule that out.) The principal's utility is assumed to equal the difference between the outcome and the compensation he pays the agent. (Thus the principal is neutral towards risk.) The agent's utility is assumed to depend both on his action and on his compensation. (He may be neutral towards risk or averse to it.)

Table 1. Examples of Principal-Agent Relationships

board of directors	chief executive officer
manager	subordinate
foreman	worker
<hr/>	
client	lawyer
customer	supplier
regulator	public utility
insurer	insured

I shall represent this situation as a two-move game. The principal moves first, announcing a pair of compensations, (w_0, w_1) , to which he is committed. The agent moves second, choosing his action. The outcome is then observed by both players, and the agent is compensated accordingly. In this game the principal's *strategy* is the same as his move, namely the compensation-pair; but the agent's strategy is a *decision-rule* that determines his action corresponding to each alternative compensation-pair that the principal could announce.

An *equilibrium*⁸ of the game is a pair of strategies, one for the principal and one for the agent,

8. The game-theorist will recognize that I have added the condition of subgame-perfection.

such that

1. Given the announced compensation-pair, the agent chooses his action so as to maximize his own expected utility.
2. Given the optimizing behavior of the agent described in item 1, the principal chooses a compensation-pair that maximizes his own expected utility.

In the formulation of a principal-agent model one typically adds one or both of the following constraints on the compensation-pair that the principal may announce:

1. The compensation-pair must enable the agent to attain (ex ante) an "acceptable" expected utility.
2. The individual compensations are bounded below by some exogenously given bound.

The first constraint can be interpreted as requiring that the principal must offer the agent an expected utility at least as large as what the agent could obtain in other employment. The second constraint recognizes that the agent's wealth is finite, and so the agent cannot pay the principal arbitrarily large amounts of money (negative compensations).

A strategy-pair is defined to be *efficient* if no other strategy-pair yields one of the players more expected utility and yields the other no less. The main proposition of this section is, with one interesting exception, that *under "realistic" conditions, an equilibrium is not efficient*. Precise mathematical statements of the model and the proposition are given at the end of this section. I shall try to make the proposition plausible here with an informal argument.

Suppose that the agent is averse to risk. First, I shall argue that in an efficient strategy-pair the agent's compensation must be independent of the outcome, that is, w_0 must equal w_1 . Suppose, to the contrary, that the two compensations were different ($w_0 \neq w_1$), and let $\bar{w}$ be the expected compensation corresponding to the agent's action. Since the agent is averse to risk, he would be better off if he used the same action but received a compensation equal to $\bar{w}$ regardless of the outcome. The principal, on the other hand, would be no worse off in this new situation, since he is neutral towards risk. Indeed, if one wanted to make both players strictly better off, the principal could pay the agent a constant compensation that is slightly less than $\bar{w}$.

On the other hand, a strategy-pair in which the agent's compensation does not depend on the outcome typically cannot be an equilibrium, unless by a coincidence the action that the agent most prefers in itself is also part of an efficient strategy-pair. For example, if increasing the

probability of success requires more “effort” by the agent, and the agent prefers less effort to more, then if the compensation is independent of the outcome the agent will have no incentive to exert any effort at all! Thus in an equilibrium the agent typically must get a larger compensation for success than for failure. The incentive requirements for equilibrium, therefore, are incompatible with the conditions for efficiency.

An exception to the proposition occurs if the agent is neutral towards risk and is sufficiently wealthy. In this case, an efficient equilibrium is obtained if the principal sells the agent a “franchise” to the enterprise, that is, the agent pays the principal a fixed fee, and then keeps the entire outcome. (It is easy to see that this is equivalent to making the compensation for failure negative, and to making the compensation for success one unit higher than the compensation for failure.)

Are there any remedies for the inefficiency of equilibrium in the principal-agent relationship? One possible remedy is for the principal to expend resources to monitor the agent’s action (and, more generally, his information and environment). Whether this will improve net efficiency will depend, of course, on the cost of monitoring. The prevalence of de facto decentralization in large organizations suggests that accurate monitoring of agents’ actions is too costly to be efficient, or even practicable.

Another remedy for inefficiency of equilibrium may be available if the principal-agent relationship is a long-term one. This topic is discussed in the next subsection.

Example 5.1. I start with a formal model of the example of the principal-agent game discussed in this section. The notation is chosen, as far as possible, to indicate how this example is related to the model of Section 4. The action of the agent is a nonnegative real number, a , and the resulting output is

$$Y = G(a, X) ,$$

where X is a random variable. In this example, X is distributed uniformly on the unit interval, and

$$G(a, X) = \begin{cases} 1, & \text{if } a \geq X , \\ 0, & \text{if } a < X . \end{cases}$$

We may interpret Y as success or failure, X as the difficulty of the agent’s task, and a as his effort. From the specification of G ,

$$\text{Prob}(Y = 1) = \min(a, 1) .$$

The agent has no information about X (other than its distribution) when he chooses his action.

The agent's compensation depends on the output Y according to

$$S(Y) = \begin{cases} w_0, & \text{if } Y = 0, \\ w_1, & \text{if } Y = 1. \end{cases}$$

The agent's resulting utility is

$$U_1 = P[S(Y)] - Q(a),$$

where P and Q are differentiable and strictly increasing functions, P is strictly concave, and Q is strictly convex. Hence we may assume $a \leq 1$. Notice that I have assumed that the agent is averse to risk. Without loss of generality I make the convention that

$$P(0) = Q(0) = 0.$$

The principal receives what is left of the output after compensating the agent. Assume that his utility is equal to what he receives, that is,

$$U_0 = Y - S(Y).$$

(Thus the principal is neutral toward risk.)

In this game, the principal moves first, choosing a compensation function S , and then the agent moves, choosing an action a after learning what S is. Conditional on the agent's action, the resulting expected utility to the principal is

$$v_0 = a(1 - w_1) - (1 - a)w_0, \quad (5.1)$$

and to the agent is

$$v_1 = aP(w_1) + (1 - a)P(w_0) - Q(a). \quad (5.2)$$

The principal's strategy is the compensation function S , and the agent's strategy is a mapping α from compensation functions to actions:

$$a = \alpha(S).$$

An equilibrium of the game is a pair of strategies, (S^*, α^*) , such that:

- (1) S^* maximizes v_0 given α^* .
- (2) $\alpha^*(S^*)$ maximizes v_1 given S^* .

In addition, I shall require that in an equilibrium, for *every* S (not just S^*), $\alpha^*(S)$ maximizes the agent's expected utility given S . (Thus I require that equilibria be *perfect*; in this two-move game such equilibria are also called Stackelberg.) I shall write $a^* = \alpha^*(S^*)$.

Typically, it is realistic to impose two constraints on the compensations. The first constraint is that the principal may not impose arbitrarily large penalties on the agent; in other words, the compensations are constrained so that the agent's disutility is bounded from below. The second constraint expresses the condition that the agent is free to refuse to enter into the relationship (i.e., to play the game). For this, w_0 and w_1 must be such as to enable the agent to achieve some minimum *expected* utility. For the purposes of this chapter, it is sufficient to impose a constraint of the first type; the addition of the second constraint would slightly complicate the exposition; however, it would not change the results in any essential way. To express the first constraint, we can assume that the compensations are bounded from below (and that the function P is finite everywhere); without loss of generality I assume that they are nonnegative:

$$(w_0, w_1) \geq 0 .$$

Space limitations do not permit a complete analysis of this game. We can verify easily from (5.2) that if $w_0 = w_1 = w$, then the agent will have no incentive to work, that is, $\alpha^*(w, w) = 0$. In addition, we see from (5.2) that if

$$Q'(0) \geq P(1) ,$$

then $\alpha^*(w_0, w_1) = 0$ for all w_0 and w_1 between 0 and 1; in this case the only equilibrium has $S^* = (0, 0)$ and $a^* = 0$. On the other hand, if

$$Q'(0) < P(1) , \tag{5.3}$$

then the equilibrium is characterized by

$$0 = w_0^* < w_1^* < 1 , \tag{5.4}$$

$$a^* > 0 ;$$

also, $\alpha^*(0, w_1)$ is strictly increasing in w_1 whenever $\alpha^*(0, w_1)$ is strictly between 0 and 1. This is the case I shall discuss from now on.

A pair $(\hat{S}, \hat{a})$ is *efficient* (Pareto optimal) if no other pair (S, a) yields each player as much expected utility and at least one player strictly more. From the concavity of the function P , it follows that, for the same level of effort, the agent prefers the compensation function $(\bar{w}, \bar{w})$ to

the compensation function (w_0, w_1) , where

$$\bar{w} = aw_1 + (1-a)w_0 ,$$

whereas the principal is indifferent between the two (recall that the agent is risk-averse and the principal is risk-neutral). Hence, if $[(\hat{w}_0, \hat{w}_1), a]$ is efficient, then $\hat{w}_0 = \hat{w}_1$. Together with (5.4) this shows that an equilibrium is not efficient.

There are, of course, many efficient pairs $[(\hat{w}, \hat{w}), a]$; we can show that for $0 < \hat{a} < 1$, they are characterized by the condition $P'(\hat{w}) = Q'(\hat{a})$.

5.2 Repeated Games

In this subsection, I examine some ways that the two players can exploit a *long-term* principal-agent relationship to escape, at least partially, from the inefficiency of short-term equilibria. The long-term relationship will be modeled as a situation in which the one-period situation is repeated over and over again. These repetitions give the principal an opportunity to observe the results of the agent's actions over a number of periods, and to use some statistical test to infer whether or not the agent was choosing the appropriate action. The repetitions also provide the principal with opportunities to "punish" the agent for apparent departures from the appropriate action. Finally, the fact that the agent's compensation in any one period can be made to depend on the outcomes in a number of previous periods (e.g., on the average over a number of periods) provides the principal with an indirect means of *insuring* the agent, at least partially, against random fluctuations in the outcomes that are not due to fluctuations in the agent's actions. Thus, the repetitions provide an opportunity to reduce the agent's risk without reducing his incentive to perform well.

The same beneficial results could be obtained, of course, if the agent had some means of self-insurance, for example, through access to a capital market or because his wealth was substantial. However, in many interesting cases (such as the owner-manager relationship), the random fluctuations in outcome are too large compared to the agent's wealth or borrowing power to make such self-insurance practical. With such cases in mind, I shall confine my attention to *nonnegative* compensation functions.

The decision rule that the principal uses to adjust the agent's compensation in any one period in the light of previous observations constitutes the principal's (*many-period*) strategy. Likewise, the agent will have a (*many-period*) strategy for adjusting his actions in the light of the past history of the process. In principle, the players' strategy spaces are very large and contain very

complex strategies. For this reason, I shall devote most of my attention to equilibria that are sustained with relatively simple strategies.

It may be helpful to have a stylized example in mind; I shall call this the ‘‘owner-manager’’ story. In the story, the owner is the principal and the manager is the agent. The owner of an enterprise wants to put it in the hands of a manager. In each of a number of successive periods (month, quarter, year) the profit of the enterprise will depend both on the actions of the manager and on the environment in which the enterprise is operating. The owner cannot directly monitor the manager’s actions, nor can the owner costlessly observe all of the relevant aspects of the environment. The owner and the manager will have to agree on how the manager is to be compensated, and the owner wants to pick a compensation mechanism that will motivate the manager to provide a good return on the owner’s investment, net of the payments to the manager.

I shall consider two kinds of long-term relationship. In the first, the principal ‘‘punishes’’ the agent by replacing him with another agent. I shall call this the *replacement model*. In this model, there may be an infinite sequence of agents, either because an agent has a maximum potential tenure, or because the players use strategies that imply that, with probability one, each agent will eventually be replaced. In the second type of long-term relationship, which I shall call the *nonreplacement model*, a single agent is associated with the principal forever. The players ‘‘punish’’ each other by changing their actions in response to the publicly available information, just as the partners do in the equilibria of Section 4.2. I shall discuss the replacement model first.

It is perhaps intuitively plausible that it makes a great difference whether or not the principal can commit himself in advance to a particular compensation strategy.⁹ I shall first discuss the case in which, in the context of the replacement model, the principal can so commit himself. One can show that, with simple strategies, the principal can induce the agent to behave in a way that yields both players discounted expected utilities that are close to one-period efficiency, provided that the players’ discount factors are close to 1, and the agent’s potential tenure is long. An important step in the analysis is the derivation of a lower bound on the expected tenure of the agent, as a function of the agent’s discount factor, his maximum potential tenure, and minimal information about his one-period utility function.

9. Since the agent’s actions cannot be observed by anyone else, there is no credible way in which the agent can commit himself in advance to a particular strategy.

Here is an informal description of one class of such simple strategies for the principal, which I call “bankruptcy strategies:” In this description, as in the remainder of the subsection, I shall use the language of the owner-manager story. The owner pays the manager a fixed compensation (wage) w per period until the end of the first period T in which the total of the gross returns in periods 1 through T fall below $T(r + w)$ by an amount at least s (where w , r , and s are parameters of the owner’s strategy). At the end of such a period T , the manager is replaced and the owner engages another one under the same regime. This can be interpreted as requiring the manager to produce a “paper” gross return of $(r + w)$ each period (a net return of r), and also allowing any surplus to be added to a (paper) “cash reserve” and requiring any deficit to be subtracted. The manager starts with a positive “cash reserve” equal to s and is replaced by a new manager the first time the cash reserve falls to zero.

Since the cash reserve is only an accounting fiction, the bankruptcy strategy is really only a scoring formula for evaluating the manager’s long-term performance, together with a criterion (based on the manager’s score) for ending his tenure.

One can show that, if the players’ discount rate is sufficiently close to 1, and if the manager’s potential tenure is sufficiently long, then the parameters of the bankruptcy strategy can be chosen so that the manager’s correspondingly optimal strategy yields a stochastic process of outcomes for which the pair of expected discounted utilities is close to one-period efficiency. (See Radner, 1986b, for the analysis and the statement of appropriate assumptions.)

I turn now to a brief discussion of another class of simple strategies for the owner, which will also play a role in the noncommitment, nonreplacement case. I call these *review strategies*; in these strategies, the owner periodically reviews the manager’s performance, and replaces the manager if his cumulative performance since the last review is unsatisfactory in a sense to be defined.

A review strategy for the owner has three parameters ℓ , r , and s , where:

1. ℓ is the number of periods covered by each review, and
2. the manager is replaced immediately after any review for which the total return during the ℓ periods preceding the review does not exceed $\ell r - s$.

Thus the first review occurs at the end of period ℓ , and the manager is replaced if

$$S_\ell = R_1 + \cdots + R_\ell \leq \ell r - s ;$$

otherwise the manager continues in office until the end of period 2ℓ , at which time he is replaced if

$$S_{2\ell} - S_{\ell} = R_{\ell+1} + \dots + R_{2\ell} \leq \ell r - s ,$$

and so on. The manager's total tenure, say T , will be some random multiple of ℓ , that is,

$$T = N\ell .$$

I shall call the periods from $[(n-1)\ell+1]$ to $n\ell$ the n^{th} *review phase*; thus ℓ is the length of each review phase and N is the number of review phases in the manager's tenure.

Assume that during his tenure the manager receives from the owner a fixed payment per period. Consider a period t that is in the n^{th} review phase, and take the "state of the system" at the end of period t to be the pair

$$(t - n\ell, S_t - S_{(n-1)\ell}) ;$$

then with this state space the manager faces a standard finite-state dynamic programming problem. We may therefore, without loss of generality, suppose that the manager uses a strategy that is "stationary" in the sense that:

1. in each period, action depends only on the state of the system at the end of the previous period;
2. in periods $1, \ell+1, 2\ell+1$, and so on, action is the same and independent of the history of the process.

In other words, the beginning of each review phase is a point of renewal of the process.

One can prove (see Radner, 1986b) results for review strategies that are similar to those for bankruptcy strategies.

Up to this point I have assumed that the owner (principal) could precommit to a particular strategy, even though the manager (agent) could not. In fact, such precommitments are the exception rather than the rule in owner-manager relations, although precommitment, in the form of contracts, can be found in other principal-agent relationships (e.g., customer-supplier and client-broker).

For the strategies that have been considered in previous sections, there are many situations in which the owner might be tempted to change strategy in mid-course. For example, in the case of the bankruptcy strategy, if the manager has accumulated an unusually large cash reserve, he can

be expected to “coast” for many periods while the cash reserve falls to a lower level (but one that is still “safe” from the manager’s point of view). Similarly, if the manager is near the end of the maximum potential tenure, and has a relatively safe cash reserve, he will have an incentive to coast. In both of these situations the owner would be tempted to replace the manager immediately with a new one. Analogous situations arise under the review strategies. The manager would be expected to move away from the actions that produce the highest returns if the reserve were sufficiently high or sufficiently low. In both cases the probability of passing review would be very little affected by the manager’s choice of actions during the remainder of the review period, and so the manager would have an incentive to choose actions that gave him higher one-period utility.

On the other side of the balance, there may be costs to the owner of replacing a manager, costs that have not been taken into account in the previous discussion. First, the owner may find it more difficult to find replacements for the manager’s position if it is known that the owner has departed in mid-course from a previously announced strategy, or in other words has “renege” on a promise or understanding. Second, there may be replacement costs that are incurred whether or not the replacement conforms to the announced strategy, due to a breaking-in period for the new manager, replacements of subordinates, interruptions of established routines, and so on. These costs would give the owner an incentive to avoid replacement as a deterrent even in the announced strategy and to find some other means of inducing good behavior by the manager.

These considerations lead one to consider a model in which the manager is *never* replaced, but the consequence of poor performance is a temporary reversion to a “noncooperative” or “adversarial” phase in which the manager receives a less satisfactory compensation than under the normal “cooperative” phase. To the extent that the noncooperative phases are also less favorable for the owner, the owner will be deterred from ending the cooperative phases prematurely. One can show that the review strategies described above can be transformed into *self-enforcing* agreements by prescribing the noncooperative phases to be equilibria of the one-period game that are inferior to the cooperative phases (in expected value) for both the owner and the manager. Furthermore, one can do this in such a way as to sustain supergame equilibria that are approximately efficient, as in the case of the replacement model. (For the details of this construction under various conditions, see Radner, 1981, 1985, 1986b.)

In the nonreplacement model, the owner and the manager are bound to each other “forever.” A more general model would incorporate explicitly the costs to the owner and manager of the owner replacing the manager and of the manager quitting. The analysis would show how the

structure of self-enforcing agreements (supergame equilibria) would depend on those costs and on the other parameters of the model. I am not aware of any such general formal analysis.

5.3 A Continuous-Time Model

In more specialized models of repeated principal-agent games one can obtain sharper characterizations of supergame equilibria. This is the case in a particular continuous-time version of the “replacement model” of the previous subsection. In this model (see Dutta and Radner, 1987):

1. The cumulative gross return is a controlled diffusion process, as in the continuous-time version of the repeated partnership game of Section 4.4.
2. At each instant of time, the manager’s action (control) can take on one of finitely many values. To each action is associated a drift and variability of the diffusion process.¹⁰ The manager’s strategy must satisfy certain measurability conditions as a function of the previous histories.
3. The owner uses a continuous-time analogue of a bankruptcy strategy, corresponding in an obvious way to the discrete-time bankruptcy strategies described in the previous subsection. During his tenure, the manager receives a wage rate that is constant. This wage rate (per unit time) is a parameter of the owner’s bankruptcy strategy. As before, the other parameters are the initial stock and the (constant) target rate of return.
4. The manager’s supergame utility is the expected integral of his discounted instantaneous utility over his tenure (the latter is a random variable.) His instantaneous utility is a function of his wage rate and his current action. Of the several actions, some have a higher instantaneous utility for him, but also have a smaller drift. It is this feature that creates the conflict between him and the owner.
5. The owner’s expected utility is the expected discounted integral of his net return, i.e., his “instantaneous” gross return minus the wage rate.¹¹

10. The continuous-time diffusion process corresponds, roughly speaking, to the limit case of the model of Section 4.4 in which the time between successive actions is “infinitesimal.” However, in the present model, only the manager controls the drift and the variability of the diffusion.

11. Strictly speaking, since the time-derivative of a diffusion process is almost-everywhere nonexistent, the “instantaneous” gross return is not well-defined. Nevertheless, the theory of stochastic integration provides a basis for defining the cumulated gross return, with or without discounting.

6. The two players have the same discount rate.

The first result of this analysis is that the manager's optimal policy is of the "switch-point" type: for example, if there are two actions then there is a critical stock, called the *switch-point*, such that the manager uses the higher-drift action when current stock is below the switch-point, and the lower-drift action (preferred by him) when the current stock is above the switch-point. The switch-point can be calculated explicitly as the solution of a transcendental equation. In certain "extreme" cases, the switch-point may be zero or infinite.

As the manager's discount rate approaches zero, his (optimal) switch-point increases without bound, but in the limit there is no optimal policy (i.e., under the expected long-run average objective). One can also characterize the dependence of the switch-point on the other parameters of the bankruptcy policy.

With this information, one can characterize in some detail the bankruptcy policy that is optimal for the owner, given the manager's reservation utility, i.e., the lowest expected utility that will induce him to take the job.

Finally, one can calculate how fast the players' respective (equilibrium) expected utilities approach efficiency as their discount rate approaches zero.

6. A Regulated Firm with Investment in Research and Development

A study by P. B. Linhart, F. W. Sinden, and myself provides an application of the ideas of principal-agent theory to the development of regulatory policy. The policy in question is the so-called "price-cap" method of regulation, which has recently been adopted in the U.K. and the U.S., replacing the rate-of-return method in the regulation of telecommunications.¹² From a formal point of view, this study extends the analysis of "bankruptcy strategies" (Section 5.2) to a non-stationary principal-agent model.

We consider the problem faced by the regulator of a monopoly firm who wants to provide incentives for the firm to effect cost reductions – and hence price reductions – through

12. This study is reported in more detail in Linhart *et al.* (1987). The latter paper is itself based upon research done at AT&T Bell Laboratories several years earlier as part of the process of developing the price-cap method; see, for example, Linhart and Radner (1983).

technological change and other means. For reasons explained elsewhere, we seek an alternative to conventional rate-of-return regulation. We model the manager of the firm as facing constraints imposed by the shareholders and other providers of capital, by the customers, and by the regulator. The regulator's ultimate objective is a secular real decrease in the firm's prices. However, the manager's private utility may not be maximized by activities that are maximally cost-reducing. Moreover, the regulator cannot directly observe all of the manager's actions, the outcomes of which are also influenced by random exogenous events. Hence a problem of moral hazard arises.

We propose a regulatory policy in which the regulator directly requires the firm to lower its real prices at (or faster than) some prescribed target annual rate. We suppose that the manager is replaced when he can no longer simultaneously repay the cost of capital, lower the prices at the rate prescribed by the regulator, and satisfy the market demand at those prices. Whenever a manager is replaced, the regulator reverts to conventional rate-of-return regulation for a period sufficient to enable the firm to build up a new cash reserve.

The resulting situation leads naturally to a model of a sequential principal-agent relationship, in which the regulator is the principal and the manager is the agent. This is not a repeated game, however, because both the firm's prices and its productivity are changing through time, endogenously and stochastically. Using new techniques for the analysis of this nonstationary process, we (1) derive a lower bound on the expected length of tenure of a manager, and (2) show that if the manager does not discount future utility very much, then the realized long-run rate of price decrease will be correspondingly close to the target rate.

The model takes account of the following fundamental characteristics, among others, of the regulatory situation:

1. The regulator and the firm's manager have different information. In particular, the manager has more information about the possibilities for productivity improvement than regulators. In fact one of the manager'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 manager's information seldom appears to be part of the regulator's strategy. In the present model the regulator does not even try to elicit information about the firm's costs, hence *misrepresentation* is not a problem.
2. The regulator and the firm's manager 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 assumed in spite of possible mismanagement and/or bankruptcy.
4. 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). If this sequence of prices is beyond the firm's control, then it has, essentially, the desirable incentive property of a lump-sum payment. 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. However, 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. Thus time is divided into alternating segments: *incentive phases* and *recovery phases*.

In the context of a particular formal model of a single-product firm, 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 approximately 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.

Several features of our approach should be emphasized. First, as mentioned above, we model the firm's manager as the active decision-maker in the firm, optimizing his own utility subject to constraints imposed by shareholders, customers, and regulators.

Second, we portray the regulators as seeking a mechanism that is easy to administer and that gives "satisfactory" results. In this case, "satisfactory" means achieving a target rate of price reduction, perhaps only approximately. Thus the regulator does not seek an "optimal" mechanism in any precise sense.

Third, we propose a regulatory mechanism that does away with explicit rate-of-return regulation. We are interested in alternatives to rate-of-return regulation because (1) we are concerned about the weakness of its incentive properties, and (2) its informational requirements are heavy. Rate-of-return regulation is also difficult to administer if some of the firm's activities are regulated and others are not, as in the case of telecommunications today; see Linhart and Radner (1983).

Fourth, from a technical point of view, our model requires an analysis that goes substantially beyond currently available results for repeated principal-agent games. The reason for this is that both the firm's productivity and its prices are changing from period to period, and these changes are both endogenous and stochastic. Thus our model leads to a sequential – but not repeated – principal-agent relationship, with endogenous state variables, namely the current prices and productivity.

7. Survival

7.1 Introduction

When responding to the "bankruptcy policies" described in Sections 5.2, 5.3, and 6, the agent is faced with a problem of survival. Similar problems arise in the context of economic development and in the situation of an indebted investor (see Majumdar and Radner, 1990). From a formal point of view, these problems have in common the features that (1) the agent controls (imperfectly) a stochastic process of gross returns, (2) the agent must pay or consume at a fixed rate out of his gross returns, (3) the resulting net return (which may be positive or negative) is cumulated in a stock of capital or cash reserve, and (4) the agent starts with an exogenously

determined stock, and “fails” or is “ruined” if the stock falls to or below some prescribed level, say zero.

Closed-form characterizations of optimal policies for such problems of survival under uncertainty are difficult to obtain without making fairly detailed assumptions. One example was sketched in Section 5.3. Another case is described in the present section,¹³ also using a continuous-time diffusion model, which allows for a fairly general model of a stochastic linear technology of “investment,” but pays for this generality by assuming that the agent does not discount the future, i.e., his objective is to maximize the expected long-run average of his utility. I shall say that the agent *survives* (forever) if he is never ruined, i.e., if the stock never falls to zero. If the agent has available a strategy that enables him to survive with positive probability, then the above objective reduces to the objective of maximizing that probability, which is the case on which I shall concentrate.

We consider the situation in which, at each instant of time, the agent can reallocate his total capital stock among a set of alternative investment opportunities. Each investment opportunity is a diffusion process, characterized by the drift and variability of its rate of return (this will be made more precise below). In the case in which the probability of survival is positive we are able to characterize the optimal policy of the agent in a relatively simple fashion. A striking feature of this optimal policy is that when the agent’s capital is below a critical level he uses investments that are “inefficient” in the mean-variance sense, namely, there are other investments that have the same mean but a lower variance. Another interpretation is that the agent exhibits “risk-loving” or “risk-averse” behavior according as his capital is below or above some critical level.

7.2 A Diffusion Process with Positive Probability of Survival

We start with a discrete-time model and then describe the corresponding diffusion model that arises in the limit as the time between successive returns and payouts approaches zero, i.e., as the investment returns and the payouts (consumption) occur in continuous time.

13. The material of this section is adapted from Majumdar and Radner (1989). In order to give a precise description of the optimal policies, this section is somewhat more mathematical than the previous ones. It also provides the reader with a heuristic derivation of the continuous-time controlled diffusion model as a limit of discrete-time models.

In the discrete-time model, let h denote the length of time between dates, starting at date 0. Let t be an integral multiple of h , say nh , and let $Y(t)$ denote the agent's capital at date t ; then the agent's capital at date $(t+h)$ is

$$Y(t+h) = e^{R_{n+1}} [Y(t) - ch] , \quad (7.1)$$

where c is a positive constant, and $\{R_n\}$ is a sequence of independent and identically distributed Gaussian random variables, each with mean mh and variance vh . The parameter m is the *drift* and v is the *variability*.

The sequence $\{Y(nh)\}$ is a Markov chain, and we are interested in the probability that $Y(nh)$ never becomes zero or negative, given $Y(0) = y > 0$. We shall be able to provide an explicit expression for this probability in the limiting case in which h tends to zero, and the process $\{Y(nh)\}$ tends to a *diffusion* process (see, e.g. Karlin and Taylor, 1981). To this end it will be more convenient to deal with the process

$$Z(t) \equiv \ln Y(t) . \quad (7.2)$$

More precisely, we define T to be the first t such that $Y(t) \leq 0$, and

$$Z(t) = \begin{cases} \ln Y(t) , & t < T , \\ -\infty , & t \geq T . \end{cases} \quad (7.2')$$

If the agent survives, then $T = +\infty$ and $Z(t)$ is always finite.

Since the process is Markovian, it suffices to consider the conditional distribution of $Z(h)$ given $Z(0) = z \equiv \ln y$. (Note that, since $y > 0$, z is well-defined.)

For a given y , the smaller h the smaller is the probability that $Y(h) \leq 0$, i.e., that $Z(h) = -\infty$. In the following heuristic argument we shall suppose that h is so small relative to y that the probability that $Y(h) \leq 0$ is "negligible."

From (7.1) and (7.2),

$$\begin{aligned} Z(h) &= R + \ln(y - ch) \\ &= R + \ln(e^z - ch) \\ &= R + z + \ln(1 - hce^{-z}) . \end{aligned}$$

Hence

$$Z(h) - z = R + \ln(1-hC) \quad (7.3)$$

$$C \equiv ce^{-z} ; \quad (7.4)$$

note that C depends on z .

It will be useful to have expressions for the conditional mean and variance of $[Z(h)-z]$, given $Z(0) = z$. (In what follows, all expectations are to be understood as conditional on $Z(0) = z$.) Expanding $\ln(1-hC)$ in powers of h , we get from (7.3),

$$Z(h) - z = R - hC - \frac{h^2 C^2}{2} + O(h^3) .$$

Hence

$$E[Z(h)-z] = (m-C)h + O(h^2) , \quad (7.6)$$

$$\text{Var} [Z(h)-z] = vh , \quad (7.7)$$

$$E[Z(h)-z]^2 = vh + O(h^2) , \quad (7.8)$$

$$E[Z(h)-z]^3 = (E[Z(h)-z])^3 + 3E[Z(h)-z] \text{Var} [Z(h)-z] = O(h^2) . \quad (7.9)$$

Define, for $y \geq 0$,

$$U(y) \equiv \text{Prob} (\text{Survival} \mid Y(0) = y) ,$$

$$V(z) \equiv U(e^z) ,$$

where $U(0) = V(-\infty) = 0$. Because the process is Markovian,

$$V(z) = EV[Z(h)] , \quad (7.10)$$

where, as before, the expectation is conditional on $Z(0) (\geq 0)$. Now assume $Z(0) = z > 0$. Supposing V is sufficiently smooth, we expand $V[Z(h)]$ in a Taylor's series:

$$V[Z(h)] = V(z) + V'(z)[Z(h)-z] + \left[\frac{1}{2} \right] V''(z)[Z(h)-z]^2 + O([Z(h)-z]^3) . \quad (7.11)$$

Taking the expectation of both sides of (7.11), we have from (7.6)-(7.9),

$$EV[Z(h)] = V(z) + V'(z)(m-C)h + \left[\frac{1}{2} \right] V''(z)vh + O(h^2) , \quad (7.12)$$

and so from (7.10),

$$\begin{aligned}
 V(z) &= V(z) + V'(z)(m-C)h + \left[\frac{1}{2} \right] V''(z)vh + O(h^2) , \\
 0 &= V'(z)(m-C) + \left[\frac{1}{2} \right] V''(z)v + \left[\frac{1}{2} \right] O(h^2) .
 \end{aligned} \tag{7.13}$$

Letting h tend to zero in (7.13) we get the differential equation

$$(m - ce^{-z})V'(z) + \left[\frac{v}{2} \right] V''(z) = 0 . \tag{7.14}$$

The remainder of this subsection is devoted to the case in which the agent has no choice of investment, i.e., m and v are fixed in time, and strictly positive. (The reader may skip to Section 7.3 without loss of continuity.) The general solution for V' of the differential equation (7.14) is:

$$V'(z) = He^{-(az + be^{-z})} , \tag{7.15}$$

where

$$a \equiv \frac{2m}{v} , \quad b \equiv \frac{2c}{v} , \tag{7.16}$$

and H is an arbitrary positive constant whose particular value will be determined by the boundary conditions in the problem.

Before solving for V , we list (without proof) some properties of V' . First, one can show that

$$V'(z) > 0 . \tag{7.17}$$

Second, $V'(z)$ is decreasing if and only if

$$z > \ln \left[\frac{c}{m} \right] , \tag{7.18}$$

or

$$y > \frac{c}{m} . \tag{7.19}$$

Third,

$$\lim_{z \rightarrow -\infty} V'(z) = \lim_{z \rightarrow +\infty} V'(z) = 0 . \tag{7.20}$$

In summary, $V'(z)$ is strictly positive, and increases monotonically from 0 as z increases from $-\infty$ to $\ln(c/m)$, and then decreases monotonically towards zero as z increases beyond $\ln(c/m)$. It

follows that

$$V''(z) \geq 0 \quad \text{a} \quad z \leq \ln \left(\frac{c}{m} \right). \quad (7.21)$$

We turn now to the study of the function U . Recall that

$$U(y) = V(\ln y).$$

Clearly $U(0) = 0$, and by an independent argument, one can show that

$$\lim_{y \rightarrow \infty} U(y) = 1.$$

This is the boundary condition that will determine H . It is now straightforward to show that

$$U(y) = \frac{1}{\Gamma(a)} \int_{b/y}^{\infty} x^{a-1} e^{-x} dx, \quad (7.22)$$

where

$$\int_0^{\infty} x^{a-1} e^{-x} dx \equiv \Gamma(a).$$

Equation (7.22) can also be written as

$$U(y) = 1 - \left[\frac{1}{\Gamma(a)} \right] \int_0^{b/y} x^{a-1} e^{-x} dx. \quad (7.23)$$

The integral in the right-hand side of (7.23) is of course the incomplete gamma function. Thus (7.23) – or (7.22) – completely solves the problem of determining the probability of survival, starting from an initial capital y .

7.3 The Optimal Investment Policy when the Probability of Survival is Positive

We shall now derive the agent's optimal investment policy when at each instant of time the agent can choose from a set of alternative investments, each of which is characterized by a pair (m, v) as in Section 7.2. Recall that for each (m, v) , the rate of return for the corresponding investment is a Brownian motion with drift m and variability parameter v . The agent's objective is to maximize the probability of survival.

For mathematical simplicity, we shall analyze in detail the case in which the set of alternative pairs (m, v) is a compact convex set with smooth boundary. This will lead to an optimal policy

that is continuous in the capital. From the analysis of this case the nature of the optimal policy for more general cases will be clear, although a rigorous treatment contains some difficulties because the optimal policy may be discontinuous.

Let A denote the set of alternative pairs (m, v) available to the agent at any instant of time, and assume that A is compact and strictly convex, with a smooth boundary. Define

$$v' \equiv \min \{ v \mid (m, v) \in A \} ,$$

$$v'' \equiv \max \{ v \mid (m, v) \in A \} ,$$

and assume that

$$0 < v' < v'' .$$

For any v in the closed interval $[v', v'']$ define

$$f(v) = \max \{ m \mid (m, v) \in A \} .$$

(If we plot v on the horizontal axis, and m on the vertical axis, then the graph of f is the upper boundary of the set A ; see Figure 2.) It follows that f is strictly concave and differentiable on (v', v'') , and that

$$\lim_{v \rightarrow v'} f'(v) = +\infty ,$$

$$\lim_{v \rightarrow v''} f'(v) = -\infty . \quad (7.24)$$

Let v^* denote the (unique) point in $[v', v'']$ at which f attains its maximum, i.e.,

$$f(v^*) = \max \{ f(v) \mid v' \leq v \leq v'' \} ;$$

it follows from (7.24) that

$$v' < v^* < v'' .$$

Our last assumption is that there is *some* investment in A that is favorable:

$$f(v^*) > 0 .$$

To simplify the exposition, we make the stronger assumption that *all* of the investments in the graph of f are favorable:

$$f(v) > 0 \quad \text{on} \quad [v', v''] .$$

We now describe the optimal investment policy. As in Section 7.2, as long as the capital $Y(t)$ is strictly positive, let $Z(t) \equiv \ln Y(t)$. If the current state is $Z(t) = z$, then the optimal current investment is

$$\begin{aligned} v &= \beta(z) , \\ m &= f[\beta(z)] , \end{aligned}$$

where $\beta(z)$ is the unique solution v of the equation

$$f(v) - vf'(v) = ce^{-z} . \quad (7.25)$$

We shall show that β is strictly decreasing and that

$$\begin{aligned} \lim_{z \rightarrow -\infty} \beta(z) &= v'' , \\ \lim_{z \rightarrow +\infty} \beta(z) &= v_0 , \end{aligned}$$

where $[f(v_0), v_0]$ maximizes the ratio (m/v) in the set A (see Figure 2).

These properties of the optimal policy have an interesting interpretation. Recall that f reaches a maximum at v^* , which is between v_0 and v'' . The part of the graph of f between v^* and v'' is “inefficient” in the usual treatment of mean-variance portfolio analysis, since from any of those points one can reduce the variance without decreasing the mean. Nevertheless, when the agent’s capital is sufficiently low, his optimal choice of (m, v) will be in this “inefficient” part of the boundary of A . Rather than call such choices “inefficient,” we shall say that for the corresponding values of z the agent exhibits “risk-loving” behavior.

The critical value of z , call it z^* , below which the agent exhibits “risk-loving” behavior is easily calculated. Observe that v^* and z^* must satisfy (7.25); also, since f attains its maximum at v^* , $f'(v^*) = 0$. Hence

$$f(v^*) = ce^{-z^*} .$$

Let

$$m^* \equiv f(v^*) , \quad y^* = e^{z^*} ;$$

then

$$y^* = \frac{c}{m^*} .$$

Compare this with (7.21).

For a proof of these results, the reader is referred to Majumdar and Radner (1989). In brief, at every instant of time the function V must satisfy the differential equation (7.14); in addition, the optimal choices of the instantaneous drift and variability must satisfy a continuous-time analogue of the Optimality Equation for dynamic programming. One can also extend the analysis in a straightforward way to the case in which the set of investment opportunities is not convex, and possibly contains points with negative drift and/or zero variability.

The reader is also referred to the above paper for a discussion of the case in which, for any policy of the agent, the probability of survival is zero. In this case, the agent may be supposed to want to maximize the expected time to failure. Again, for a fixed drift and variability, one can obtain an explicit expression for the objective function. However, in the general case in which the agent has a choice among alternative investment opportunities, no explicit solution appears to be available at the present time.

8. Concluding Remarks

The models of dynamic games sketched here hardly constitute a theory of economic organizations. For such a theory, the partnership and principal-agent models will have to be merged and subsumed in a richer and more comprehensive model capable of depicting the interactions among many persons. This is particularly true if one's ambition is to develop a theory of the internal workings of today's large industrial enterprises, many of which have tens – or even hundreds – of thousands of employees. In addition, it will probably be necessary to take account of the costs of communication and information processing, as well as other aspects of “bounded rationality.”¹⁴ Nevertheless, the insights revealed by these relatively primitive analyses of moral hazard and free riding in long-term relationships will probably prove to be durable.

9. References

1. Abreu, D., P. Milgrom, and D. Pearce (1988), “Information and Timing in Repeated Partnerships,” Harvard University (unpublished).

14. See Radner (1989) for further reflections on these problems.

2. Alchian, A. A., and H. Demsetz (1972), "Production, Information Costs, and Economic Organization," *Amer. Econ. Rev.*, 62, 777-795.
3. Benhabib, J., and R. Radner (1988), "Joint Exploitation of a Productive Asset," New York University (unpublished).
4. Dutta, P. K., and R. Radner (1987), "Principal-Agent Games in Continuous Time," AT&T Bell Laboratories, Murray Hill, N.J. (unpublished).
5. Fudenberg, D., and E. Maskin (1986), "Folk Theorems for Repeated Games with Discounting and Incomplete Information," *Econometrica*, 54, 533-54.
6. Fudenberg, D., D. K. Levine, and E. Maskin (1989), "The Folk Theorem with Imperfect Public Information," Massachusetts Institute of Technology (unpublished).
7. Fudenberg, D., and J. Tirole (1990), "Moral Hazard and Renegotiation in Agency Contracts," *Econometrica*, 58, 1279-1319.
8. Grossman, S. J., and O. D. Hart (1983), "An Analysis of the Principal-Agent Problem," *Econometrica*, 51, 7-45.
9. Groves, T. (1973), "Incentives in Teams," *Econometrica*, 41, 617-631.
10. Harsanyi, J. (1967-68), "Games with Incomplete Information Played by Bayesian Players," Parts I, II, and III, *Management Science*, 14, pp. 159-182, 320-334, 486-502.
11. Ito, K. (1985), *Lectures on Stochastic Processes*, Springer-Verlag, Berlin.
12. Karlin, S., and H. M. Taylor (1981). *A Second Course in Stochastic Processes*, Academic Press, New York.
13. Lancaster, K. (1973), "The Dynamic Inefficiency of Capitalism," *J. of Political Economy*, 81, 1098-1109.
14. Levhari, D., and L. J. Mirman (1980), "The Great Fish-War: An Example Using the Cournot-Nash Solution," *Bell J. of Economics*, 11, 322-334.
15. Linhart, P. B., and R. Radner (1983), "Deregulation of Long-Distance Telecommunications," in V. Mosco, ed., *Policy Research in Telecommunications*, ABLEX Publishing Co., Norwood, N.J.

16. Linhart, P. B., F. W. Sinden, and R. Radner (1987), "A Sequential Mechanism for Direct Price Regulation," AT&T Bell Laboratories, Murray Hill, N.J. (unpublished).
17. Majumdar, M. K., and R. Radner (1990), "Survival under Production Uncertainty," to appear in W. Brock and M. K. Majumdar, eds., *Essays in Honor of David Gale*, Macmillan, London.
18. Majumdar, M. K., and R. Radner (1989), "Notes on a Diffusion Model of Economic Survival under Uncertainty," AT&T Bell Laboratories, Murray Hill, N.J. (unpublished).
19. Mookherjee, D., and S. Reichelstein (1990), "Dominant-Strategy Implementation of Bayesian Incentive-Compatible Allocation Rules," Univ. of Calif., Berkeley (unpublished).
20. Myerson, R. B. (1983), "Mechanism Design by an Informed Principal," *Econometrica*, 51, 1767-1798.
21. Myerson, R. B. (1985), "Bayesian Equilibrium and Incentive Compatibility: An Introduction," in L. Hurwicz, D. Schmeidler, and H. Sonnenschein, eds., *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, Cambridge U. Press, Cambridge, pp. 229-259.
22. Radner, R. (1981), "Monitoring Cooperative Relationships in a Repeated Principal-Agent Relationship," *Econometrica*, 49, 1127-1148.
23. Radner, R. (1985), "Repeated Principal-Agent Games with Discounting," *Econometrica*, 53, 1173-1197.
24. Radner, R. (1986a), "Repeated Partnership Games with Imperfect Monitoring and No Discounting," *Review of Economic Studies*, 53, 43-57.
25. Radner, R. (1986b), "Repeated Moral Hazard with Low Discount Rates," Ch. 2 in W. P. Heller, R. M. Starr, and D. Starrett, eds., *Uncertainty, Information, and Communication, Essays in Honor of Kenneth J. Arrow*, Vol. 3, Cambridge University Press, Cambridge, pp. 25-64.
26. Radner, R. (1987), "Decentralization and Incentives", Ch. 1 of T. Groves, S. Reiter, and R. Radner, eds., *Information, Incentives, and Economic Mechanisms*, Basil Blackwell, Oxford, pp. 3-47.

27. Radner, R. (1989), "Hierarchy: The Economics of Managing," Marshall Lectures, Cambridge University, October, 1989.
28. Radner, R., R. B. Myerson, and E. Maskin (1986), "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria," *Review of Economic Studies*, 53, 59-69.
29. Rustichini, A., and R. Radner (1988), "The Design and Performance of Sharing Rules for a Partnership in Continuous Time," AT&T Bell Laboratories, Murray Hill, N.J. (unpublished).
30. Williams, S. R., and R. Radner (1987), "Efficiency in Partnerships when the Joint Output is Uncertain," Northwestern University (unpublished).
31. Whitt, W. (1980), "Uniform Conditional Stochastic Order," *J. Appl. Prob.*, 17, 112-123.

Fig. 1. One-Period Partnership

Figure 2