

Moral Hazard

Dutta, Prajit K. and Roy Radner

Working Paper No. 356  
August 1993

University of  
Rochester

## **Moral Hazard**

Prajit K. Dutta and Roy Radner

Rochester Center for Economic Research  
Working Paper No. 356



**Moral Hazard\***

**Prajit K. Dutta**

Department of Economics  
Columbia University  
New York, NY 10027  
and  
Department of Economics  
University of Wisconsin  
Madison, WI 53706

and

**Roy Radner**

Mathematical Sciences Research Center  
AT&T Bell Laboratories  
Murray Hill, NJ 07974  
and  
Department of Economics  
New York University  
New York, NY 10003

July 1993

\*Forthcoming as a chapter in the "Handbook of Game Theory - volume 3", edited by Robert Aumann and Sergiu Hart and published by North-Holland, Amsterdam. We thank the editors and two anonymous referees for helpful comments. The views expressed here are those of the authors and do not necessarily reflect the viewpoint of AT&T Bell Laboratories.



## 1. Introduction

The owner of an enterprise wants to put it in the hands of a manager. The profits of the enterprise will depend both on the actions of the manager as well as the environment within which he operates. The owner cannot directly monitor the agent's action nor can he costlessly observe all relevant aspects of the environment. This situation may also last a number of successive periods. 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. This is the well-known "**principal-agent**" problem with **moral hazard**. Some other principal-agent relationships in economic life are: client-lawyer, customer-supplier, insurer-insured and regulator-public utility.<sup>1</sup>

The principal-agent relationship embodies a special form of moral hazard, which one might call "one-sided", but moral hazard can also be "many-sided". The paradigmatic model of many-sided moral hazard is the **partnership**, in which there are many agents but no principal. The output of the partnership depends jointly on the actions of the partners and on the stochastic environment; each partner observes only the output (and his own action) but not the actions of the other partners nor the environment. This engenders a free-rider problem. As in the case of principal-agent relationships, a partnership, too, may last many periods.<sup>2</sup>

---

<sup>1</sup>The insurer-insured relationship is the one that gave rise to the term "moral hazard" and the first formal economic analysis of moral hazard was probably given by Arrow (1963, 1965).

<sup>2</sup>More complex informational models can be formulated for both the principal-agent as well as the partnership framework; models in which some agents obtain (incomplete) information about the environment or the actions of others. We do not directly discuss these generalizations although many of the results that follow can be extended to these more complex settings (see also the further discussion in Section 6). Note too that we do not treat here an important class of principal-agent models, the "adverse selection" models. The distinction between moral hazard and adverse selection models is that in the former framework, the

In Section 2 we present the principal-agent model formally. Section 3 discusses some salient features of optimal principal-agent contracts when the relationship lasts a single period. The first main point to make here is that in a large class of cases an equilibrium in the one-period game is Pareto-inefficient. This is the well-known problem involved in providing a risk-averse agent insurance while simultaneously giving him the incentives to take, from the principal's perspective, appropriate actions. We also discuss, in this section, other properties of static contracts such as monotonicity of the agent's compensation in observed profits.

In Section 4 we turn to repeated moral hazard models. Section 4.1 discusses some known properties of intertemporal contracts; the main points here are that an optimal contract will, typically, reward the agent on the basis of past performance as well as current profits. Furthermore, although a long-term contract allows better resolution of the incentives-insurance trade-off, in general, some of the inefficiency of static contracts will persist even when the principal-agent relationship is long-lived. However, if the principal and agent are very patient, then almost all inefficiency can, in fact, be resolved by long-term contracts - and, on occasion, simple long-term contracts. These results are discussed in Section 4.2.

Many-sided moral hazard is studied in Section 5. The static partnership model is discussed in Section 5.1. The main focus here is on the possible resolution of the free-rider problem when the partners are risk-neutral. We also discuss some properties of optimal sharing rules, such as monotonicity, and the effect of risk-aversion on partners' incentives. Again, in general, static partnership

---

principal is assumed to know all relevant characteristics of the agent (i.e., to know his "type") but not to know what action the agent chooses whereas in the latter model the principal is assumed not to know some relevant characteristic of the agent although he is able to observe the agent's actions. (See Section 6 for a further discussion).

contracts are unable to generate efficiency. This motivates a discussion of repeated partnership models. Such a model is a special case of a repeated game with imperfect monitoring; indeed results for repeated partnerships can be derived more readily from studying this more general class of games. Hence in Section 5.2 we present known results on the characterization of equilibria in repeated games with imperfect monitoring.

It should be noted that the principal–agent framework is in the spirit of mechanism design; the principal chooses a compensation scheme, i.e., chooses a game form in order to motivate the manager to take appropriate actions and thereby the principal maximizes his own equilibrium payoff. The static partnership model is similarly motivated; the partners' sharing rule is endogenous to the model. In contrast, one can take the compensation scheme or sharing rule as exogenously given, i.e., one can take the game form is given, and focus on the equilibria generated by this game form. In the second approach, therefore, a moral hazard or partnership model becomes a special case of a game with imperfect monitoring. This is the approach used in Section 5.2.

Section 6 brings together additional bibliographical notes and discusses some extensions of the models studied in this paper.

## 2. The Principal–Agent Model

### 2.1 The Static Model

A static (or stage–game) principal–agent model is defined by the quintuple  $(A, \varphi, G, U, W)$ .  $A$  is the set of actions that the agent can choose from. An action choice by the agent determines a distribution,  $\varphi(a)$ , over output (or profit)  $G$ ;  $G \in \mathbf{G}$ . The agent's action is unobservable to the principal whereas the output is observable. The agent is paid by the principal on the basis of that which is observable; hence, the compensation depends only on the output and is denoted



$I(G) = I$ .  $U$  will denote the utility function of the agent and its arguments are the action undertaken and the realized compensation;  $U(a, I)$ . Finally, the principal's payoff depends on his net return  $G - I$  and is denoted  $W(G - I)$ . (Note that  $G$  and  $I$  are real-valued).

The **maintained assumptions** will be:

- (A1) There are only a finite number of possible outputs;  $G_1, G_2, \dots, G_n$ .  
 (A2) The set of actions  $A$  is a compact subset of some Euclidean space.  
 (A3) The agent's utility function  $U$  is strictly increasing in  $I$  and the principal's payoff function  $W$  is also strictly increasing.

A compensation scheme for the agent will be denoted  $I_1, \dots, I_n$ . Furthermore, with some abuse of notation, we will write  $\varphi_j(a)$  for the probability that the realized output is  $G_j$ ,  $j = 1, \dots, n$ , when the action taken is  $a$ .

The time-structure is that of a two-move game. The principal moves first and announces the compensation function  $I$ . Then the agent chooses his action, after learning  $I$ . The expected utilities for principal and agent are, respectively,  $\sum_j \varphi_j(a)W(G_j - I_j)$  and  $\sum_j \varphi_j(a)U(a, I_j)$ . The **principal-agent problem** is to find a solution to the following optimization exercise:

$$\text{Max}_{I_1, \dots, I_n, \hat{a}} \sum_j \varphi_j(\hat{a})W(G_j - I_j) \quad (2.1)$$

$$\text{s.t.} \quad \sum_j \varphi_j(\hat{a})U(\hat{a}, I_j) \geq \sum_j \varphi_j(a)U(a, I_j), \quad \forall a \in A \quad (2.2)$$

$$\sum_j \varphi_j(\hat{a})U(\hat{a}, I_j) \geq \bar{U} \quad (2.3)$$

The constraint (2.2) is referred to as the **incentive-constraint**; the agent will only take those actions that are in his best interest. Constraint (2.3) is called the **individual-rationality constraint**; the agent will accept an arrangement only if his expected utility from such an arrangement is at least as large as his outside option

$\bar{U}$ . The objective function, maximizing the principal's expected payoff, is, in part, a matter of convention. One interpretation of (2.1) is that there are many agents and only one principal, who consequently gets all the surplus, over and above the outside options of principal and agent, generated by the relationship.<sup>3</sup>

If there is a  $\bar{U}$  such that  $(a^*, I^*)$  is a solution to the principal-agent problem, then  $(a^*, I^*)$  will be called a **second-best** solution. This terminology distinguishes  $(a^*, I^*)$  from a Pareto-optimal (or **first-best**) action-incentives pair that maximizes (2.1) subject only to the individual-rationality constraint (2.3).

## 2.2 The Dynamic Model

In a repeated principal-agent model, in each period  $t = 0, 1, \dots, T$ , the stage-game is played and the output observed by both principal and agent; denote the output realization and the compensation paid,  $G(t)$  and  $I(t)$  respectively. The relationship lasts for  $T$  ( $\leq \infty$ ) periods, where  $T$  may be endogenously determined. The public history at date  $t$ , that both principal and agent know, is  $h(t) \equiv (G(0), I(0), \dots, G(t-1), I(t-1))$ , whereas the private history of the agent is  $h_a(t) \equiv (a(0), G(0), I(0), \dots, a(t-1), G(t-1), I(t-1))$ . A strategy for the principal is a sequence of maps  $\sigma_p(t)$ , where  $\sigma_p(t)$  assigns to each public history,  $h(t)$ , a compensation function  $I(t)$ . A strategy for the agent is a sequence of maps  $\sigma_a(t)$ , where  $\sigma_a(t)$  assigns to each pair, a private history  $h_a(t)$  and the principal's compensation function  $I(t)$ , an action  $a(t)$ . A strategy choice by the principal and agent

---

<sup>3</sup>An alternative specification would be to maximize the agent's expected payoffs instead; in this case, the constraint (2.3) would be replaced by a constraint that guarantees the principal his outside option. Note furthermore the assumption, implicit in (2.1)–(2.3), that in the event of indifference the agent chooses the action which maximizes the principal's returns. This assumption is needed to ensure that the optimization problem has a solution. A common, albeit informal, justification for this assumption is that, for every  $\epsilon > 0$ , there is a compensation scheme similar to the one under consideration in which the agent has a strict preference and which yields the principal a net profit within  $\epsilon$  of the solution to (2.1)–(2.3).

induces, in the usual way, a distribution over the set of histories  $(h(t), h_\alpha(t))$ ; the pair of strategy choices therefore generate expected payoffs for principal and agent in period  $t$ ; denote these  $W(t; \sigma_p, \sigma_\alpha)$  and  $U(t; \sigma_p, \sigma_\alpha)$ . Lifetime payoffs are evaluated under discount factors  $\delta_p$  and  $\delta_\alpha$ , for principal and agent respectively, and equal  $(1-\delta_p) \sum_{t=0}^T \delta_p^t W(t; \sigma_p, \sigma_\alpha)$  and  $(1-\delta_\alpha) \sum_{t=0}^T \delta_\alpha^t U(t; \sigma_p, \sigma_\alpha)$ . The **dynamic principal-agent problem** is:<sup>4</sup>

$$\text{Max}_{\sigma_p, \sigma_\alpha} \quad (1-\delta_p) \sum_{t=0}^T \delta_p^t W(t; \sigma_p, \hat{\sigma}_\alpha) \quad (2.4)$$

$$\text{s.t.} \quad (1-\delta_\alpha) \sum_{t=0}^T \delta_\alpha^t U(t; \sigma_p, \hat{\sigma}_\alpha) \geq (1-\delta_\alpha) \sum_{t=0}^T \delta_\alpha^t U(t; \sigma_p, \sigma_\alpha), \quad \forall \sigma_\alpha \quad (2.5)$$

$$(1-\delta_\alpha) \sum_{t=0}^T \delta_\alpha^t U(t; \sigma_p, \sigma_\alpha) \geq \bar{U} \quad (2.6)$$

The incentive-constraint is (2.5) whereas the individual-rationality constraint is (2.6). Implicit in the formulation of the dynamic principal-agent problem is the idea that principal and agent are bound to the arrangement for the contract length  $T$ . Such a commitment is not necessary if we require that a) the continuations of  $\sigma_\alpha$  must satisfy (2.5) and (2.6) after **all** private histories  $h_\alpha(t)$  and principal's compensation choice  $I(t)$ , and b) that the continuations of  $\sigma_p$  must solve the optimization problem (2.4) after all public histories  $h(t)$ .

### 3. Analyses of the Static Principal-Agent Model

It is customary to assume that an agent, such as the manager of a firm, is

---

<sup>4</sup>In the specification that follows, we add the principal's (as well as the agent's) payoffs over the contract horizon  $0, \dots, T$  only. If  $T$  is less than the working lifetime of principal and agent, then the correct specification would be to add payoffs over the (longer) working lifetime in each case. Implicit in (2.5)–(2.6) is the normalization that the agent's aggregate payoffs, after the current contract expires, are zero. The principal's payoffs have to include his profits from the employment of subsequent agents. It is straightforward to extend (2.4) to do that and in Section 3.2 we will, in fact, do so formally.

more risk-averse than the principal, such as the shareholder(s). From a first-best perspective, this suggests an arrangement between principal and agent in which the former bears much of the risk, and indeed, if the principal is risk-neutral, bears all of the risk. However, since the agent's actions are unobservable, the provision of such insurance may remove incentives for the agent to take onerous, but profitable, actions that the principal prefers. The central issue consequently, in the design of optimal contracts under moral hazard, is how best to simultaneously resolve (possible) conflicts between insurance and incentive considerations.

To best understand the nature of the conflict imagine, first, that the agent is in fact risk-neutral. In this case first-best actions (and payoffs) can be attained as second-best outcomes, and in a very simple way. An effective arrangement is the following: the agent pays the principal a fixed fee, independent of the gross return, but gets to keep the entire gross return for himself. (The fixed fee can be interpreted as a "franchise fee.") This arrangement internalizes, for the agent, the incentives problem and leads to a choice of first-best actions. Since the agent is risk-neutral, bearing all of the risk imposes no additional burden on him.<sup>5</sup>

On the other hand, imagine that the agent is strictly risk-averse whereas the principal is risk-neutral. Without informational asymmetry, the first-best arrangement would require the principal to bear all of the risk (and pay the agent a constant compensation). However, such a compensation scheme only induces the agent to pick his most preferred action. If this is not a first-best action, then we can conclude that the second-best payoff for the principal is necessarily less than his first-best payoff. These ideas are formalized as:

**Proposition 3.1:**      i) Suppose that  $U(a, \cdot)$  exhibits risk-neutrality, for every  $a \in A$

---

<sup>5</sup>The above argument is valid regardless of whether the principal is risk-neutral or risk-averse.

(and the principal is either risk-averse or risk-neutral). Let  $(a_F, I_F)$  be any first-best pair of action and incentive scheme. Then, there is a second-best contract  $(a^*, I^*)$  such that the expected payoffs of both principal and agent are identical under  $(a_F, I_F)$  and  $(a^*, I^*)$ .

ii) Suppose that  $U(a, \cdot)$  exhibits strict risk-aversion for every  $a \in A$ , and furthermore that the principal is risk-neutral. Suppose at every first-best action,  $a_F$ , a)  $\varphi_j(a_F) > 0$ ,  $j=1, \dots, n$ , and b) for every  $I'$  there is  $a' \in A$  such that  $U(a', I') > U(a_F, I')$ . Then, the principal's expected payoffs in any solution to the principal-agent problem is strictly less than his expected first-best payoff.

Proof: i) Let  $(a_F, I_F)$  be a first-best pair of action and incentive scheme and let the average retained earnings for the principal be denoted  $\bar{G} - \bar{I} \equiv \sum_j \varphi_j(a_F)(G_j - I_{jF})$ . Consider the incentive scheme  $I^*$  in which the agent pays a fixed fee  $\bar{G} - \bar{I}$  to the principal, regardless of output. Since the agent is risk-neutral, his utility function is of the form,  $U(a, I) = H(a) + K(a)I$ . Simple substitution then establishes the fact that  $U(a_F, I^*) = U(a_F, I_F)$ . Hence, the new compensation scheme is individually rational for the agent. Moreover, since the principal is not risk-loving, his payoff under this scheme is at least as large as his payoff in the first-best solution;  $W(\bar{G} - \bar{I}) \geq \sum_j \varphi_j(a_F)W(G_j - I_{jF})$ . The scheme is also incentive compatible for the action  $a_F$ . For suppose, to the contrary, that there is an action  $a'$  such that  $H(a') + K(a')[\sum_j \varphi_j(a')G_j - (\bar{G} - \bar{I})] > H(a_F) + K(a_F)\bar{I}$ . Then there is evidently a fixed fee  $\bar{G} - \bar{I} + \epsilon$ , for some  $\epsilon > 0$ , that if paid by the agent to the principal, constitutes an individually rational compensation scheme. Further, the principal now makes strictly more than his first-best payoff; and that is a contradiction.

$(a_F, I^*)$  is a pair that satisfies constraints (2.2) and (2.3) and yields the principal at least as large a payoff as the first-best. Since, by definition, the

second-best payoff cannot be any more than the first-best payoff, in particular the two payoffs are equal and equal to that under  $(a_F, I^*)$ ,  $W(\bar{G} - \bar{I})$ .<sup>6</sup>

ii) Let  $(a^*, I^*)$  be a solution to the principal-agent problem. If this is also a solution to the first-best problem, then, given the hypothesis,  $\varphi_j(a^*) > 0$ ,  $j=1, \dots, n$ , and principal and agent attitudes to risk, it must be the case that  $I_j^* = I_{j'}^* \equiv I^*$ , for all  $j, j'$ . But then, by hypothesis,  $a^*$  is not an incentive-compatible action.  $\square$ <sup>7</sup>

Proposition 3.1ii) strongly suggests that whenever the agent is risk-averse, there will be some efficiency loss in that the principal will provide incomplete insurance, in order to maintain incentives. The results we now turn to provide some characterization of the exact trade-off between incentives and insurance in a second-best contract. The reader will see, however, that not too much can be said, in general, about the optimal contract. Part of the reason will be the fact that although, from an incentive standpoint, the principal would like to reward evidence of "good behavior" by the agent, such evidence is linked to observable outputs in a rather complex fashion.

Grossman and Hart (1983) introduced a device for analyzing principal-agent problems that we now discuss. Their approach is especially useful when the agent's preferences satisfy a separability property;  $U(a, I) \equiv H(a) + V(I)$ .<sup>8</sup> Suppose

---

<sup>6</sup>A corollary to the above arguments is clearly that if the principal is risk-averse, while the agent is risk-neutral, then the unique first (and second) - best arrangement is for the agent to completely insure the principal.

<sup>7</sup>Whether or not there is always a solution to the principal-agent problem is an issue that has been discussed in the literature. Mirrlees (1974) gave an example in which the first-best payoff can be approximated arbitrarily closely but cannot actually be attained. Sufficient conditions for a solution to the principal-agent problem to exist are given, for example, by Grossman and Hart (1983).

<sup>8</sup>Grossman and Hart admit a somewhat more general specification;  $U(a, I) = H(a) + K(a)V(I)$  where  $K(a) > 0$ . That specification is equivalent to the requirement that the agent's preferences over income lotteries be independent of his action. See Grossman and Hart (1983) for further details.

also that the principal is risk-neutral and the agent is risk-averse.<sup>9</sup> Now consider any action  $a \in A$  and let  $C(a)$  denote the minimum expected cost at which the principal can induce the agent to take this action, i.e.

$$C(a) \equiv \text{Min}_{v_1, \dots, v_n} \sum_j \varphi_j(a) V^{-1}(v_j) \quad (3.1)$$

$$\text{s.t.} \quad H(a) + \sum_j \varphi_j(a) v_j \geq H(a') + \sum_j \varphi_j(a') v_j \quad \forall a' \quad (3.2)$$

$$H(a) + \sum_j \varphi_j(a) v_j \geq \bar{U} \quad (3.3)$$

where  $v_j \equiv V(I_j)$ . (3.2) and (3.3) are simply the (rewritten) incentive and individual rationality constraints and the point to note is that the incentive constraints are linear in the variables  $v_1, \dots, v_n$ . Furthermore, if  $V$  is concave, then the objective function is convex and hence we have a convex programming problem.<sup>10</sup> The full principal-agent problem then is to find an action that maximizes the net benefits to the principal,  $\sum_j \varphi_j(a) G_j - C(a)$ .

Although the (full) principal-agent problem is typically not convex, analysis of the cost-minimization problem alone can yield some useful necessary conditions for an optimal contract. For example, suppose that the set of actions is, in fact, finite. Then the Kuhn-Tucker conditions yield:<sup>11</sup>

---

<sup>9</sup>Since Proposition 3.1 has shown that a risk-neutral agent can be straightforwardly induced to take first-best actions, for the rest of this section we will focus on the hypothesis that the agent is, in fact, risk-averse.

<sup>10</sup>The earlier literature on principal-agent models replaced the set of incentive constraints (2.8) by the single constraint that, when the compensation scheme  $I_1, \dots, I_n$  is used, the agent satisfies his first-order conditions at the action  $a$ . That this procedure is, in general, invalid was first pointed out by Mirrlees (1975). One advantage of the Grossman and Hart approach is, of course, that it avoids this "first-order approach".

<sup>11</sup>The expression that follows makes sense, of course, only when  $\varphi_j(a) > 0$  and  $V$  is differentiable.

$$[V'(I_j)]^{-1} = \lambda + \sum_{a' \neq a} \mu(a') \left(1 - \frac{\varphi_j(a')}{\varphi_j(a)}\right) \quad (3.4)$$

where  $\lambda$ ,  $\mu(a')$ , are (non-negative) Lagrange multipliers associated with, respectively, the individual rationality and incentive constraints (one for each  $a' \neq a$ ). The interpretation of (3.4) is as follows: the agent is paid a base wage,  $\lambda$ , which is adjusted if the  $j$ -th output is observed. In particular, if the incentive constraint for action  $a'$  is binding,  $\mu(a') > 0$ , then the adjustment is positive if and only if the  $j$ -th output is more likely under the desired action  $a$ , than under  $a'$ .

One further question of interest is whether there are conditions under which the optimal contract is **monotonically increasing** in that it rewards higher outputs with larger compensations; if we adopt the convention that outputs are ordered so that  $G_j \leq G_{j+1}$ , the question is, (when) is  $I_j \leq I_{j+1}$ ? This question makes sense when "higher" inputs do, in fact, make higher outputs more likely. So suppose that  $A \subset \mathbb{R}$  (for example, the agent's actions are effort levels) and, to begin with, that  $a' > a$  implies that the distribution function corresponding to  $a'$  first-order stochastically dominates that corresponding to  $a$ .

Now although the first-order stochastic monotonicity condition does imply that higher outputs are more likely when the agent works harder, we cannot necessarily infer, from seeing a higher output, that greater effort was in fact expended. The agent's reward is conditioned on precisely this inference and since the inference may be non-monotone so might the compensation.<sup>12</sup> Milgrom (1983)

---

<sup>12</sup>For example, suppose that there are two actions,  $a_1 > a_2$  and three outputs,  $G_j$ ,  $j=1, \dots, 3$ ,  $G_j \leq G_{j+1}$ . Suppose also that the probability of  $G_1$  is positive under both actions but the probability of  $G_2$  is zero when action  $a_1$  is employed (but positive under  $a_2$ ). It is obvious that if action  $a_1$  is to be implemented, then the compensation, if  $G_2$  is observed, must be the lowest possible. Here the posterior probability of  $a_1$ , given the higher output  $G_2$ , is smaller than the corresponding probability when the lowest output  $G_1$  is observed.



introduced into the principal-agent literature the following stronger condition under which higher output does, in fact, signal greater effort by the agent:

**Monotone Likelihood Ratio Condition (MLRC):** If  $a' > a$ , then the likelihood ratio  $\frac{\varphi_j(a')}{\varphi_j(a)}$  is increasing in  $j$ .

Under MLRC, the optimal compensation scheme will indeed be monotonically increasing **provided** the principal does in fact want the agent to exert the greatest effort. This can be easily seen from (3.4); the right-hand side is increasing in the output level. Since  $V^{-1}$  is convex, this implies that  $v_j$ , and hence  $I_j$ , is increasing in  $j$ . If, however, the principal does not want the agent to exert the greatest effort, rewarding higher output provides the wrong incentives and hence, even with MLRC, the optimal compensation need not be monotone.<sup>13</sup> Mirrlees (1975) introduced the following condition that, together with MLRC, implies monotonicity: (let  $F(a)$  denote the distribution function corresponding to  $a$ )

**Concavity of the Distribution Function (CDF):** For all  $a, a'$  and  $\theta \in (0,1)$ ,  $F(\theta a + (1-\theta)a')$  first-order stochastically dominates  $\theta F(a) + (1-\theta)F(a')$ .

It can be shown by standard arguments that, under CDF, the agent's expected payoffs are a concave function of his actions (for a fixed monotone compensation scheme). In turn this implies that whenever an action  $\hat{a}$  yields the agent higher payoffs than any  $a < \hat{a}$ , then, in fact, it yields higher payoffs than all other actions (including  $a > \hat{a}$ ). Formally, these ideas lead to:

**Proposition 3.2 (Grossman and Hart (1983)):** Assume that  $V$  is strictly concave and differentiable and that MLRC and CDF hold. Then a second-best incentive scheme  $(I_1, \dots, I_n)$  satisfies  $I_1 \leq I_2 \leq \dots \leq I_n$ .

<sup>13</sup>Note that if  $\mu(a') > 0$ , for some  $a' > a$ , then (3.4) shows that on account of this incentive constraint the right-hand side decreases with  $j$ .

Proposition 3.2 shows that the sufficient conditions on the distribution functions for, what may be considered, an elementary property of the incentive scheme, monotonicity, are fairly stringent. Not surprisingly, more detailed properties, such as convexity, are even more difficult to establish. The arguments leading to the proposition have, we hope, given the reader an appreciation of why this should be the case, namely the subtleties involved in inverting observed outcomes into informational inferences.<sup>14</sup>

One other conclusion emerges from the principal-agent literature: optimal contracts will be, in general, quite delicately conditioned on the parameters of the problem. This can be appreciated even from an inspection of the first-order condition (3.4). This is also a corollary of the work of Holmstrom (1979) and Shavell (1979). These authors asked the question: if the principal has available informational signals other than output, (when) will the optimal compensation scheme be conditioned on such signals? They showed that whenever output is **not** a sufficient statistic for these additional signals, i.e. whenever these signals do yield additional information about the agent's action, they should be contracted upon. Since a principal, typically, has many sources of information in addition to output, such as evidence from monitoring the agent or the performance of agents who manage related activities, these results suggest that such information should be used; in turn, this points towards quite complex optimal incentive schemes.

However, in reality contracts tend to be much simpler than those suggested by the above results. To explain this simplicity is clearly the biggest challenge for the theory in this area. Various authors have suggested that the simplicity of observable schemes can be attributed to some combination of: a) the costs of

---

<sup>14</sup>Grossman and Hart (1983) do establish certain other results on monotonicity and convexity of the optimal compensation scheme. They also show that the results can be tightened quite sharply when the agent has available to him only two actions.

writing and verifying complex schemes, b) the fact that the principal needs to design a scheme that will work well in a variety of circumstances and under the care of many different agents and c) the long-term nature of many observable incentive schemes. Of these explanations it is only c) that has been explored at any length. Those results will be presented in the next section within our discussion of dynamic principal-agent models.

#### 4. Analyses of the Dynamic Principal-Agent Model

In this section we turn to a discussion of repeated moral hazard. There are at least two reasons to examine the nature of long-term arrangements between principal and agent. The first is that many principal-agent relationships, such as that between a firm and its manager or that between insurer and insured or that between client and lawyer/doctor are, in fact, long-term. Indeed, observed contracts often exploit the potential of a long-term relationship; in many cases the contractual relationship continues only if the two parties have fulfilled prespecified obligations and met predesignated standards. It is clearly a matter of interest then to investigate how such observed long-term contractual arrangements resolve the trade-off between insurance and incentives that bedevils static contracts.

A second reason to analyze repeated moral hazard is that there are theoretical reasons to believe that repetition does, in fact, introduce a rich set of incentives that are absent in the static model. Repetition introduces the possibility of offering the agent intertemporal insurance, which is desirable given his aversion to risk, without (completely) destroying his incentives to act faithfully on the principal's behalf. The exact mechanisms through which insurance and incentives can be simultaneously addressed will become clear as we discuss the available results. In Section 4.1 we discuss characterizations of the second-best contract at fixed discount factors. Subsequently, in Section 4.2 we discuss the asymptotic case

where the discount factors of principal and agent tend to one.

#### 4.1 Second-Best Contracts

Lambert (1983) and Rogerson (1985a) have established necessary conditions for a second-best contract. We report here the result of Rogerson; the result is a condition that bears a family resemblance to the well-known Ramsey-Euler condition from optimal growth theory. It says that the principal will smooth the agent's utilities across time-periods in such a fashion as to equate his own marginal utility in the current period to his expected marginal utility in the next. We also present the proof of this result since it illustrates the richer incentives engendered by repeating the principal-agent relationship.<sup>15</sup>

Recall the notation for repeated moral hazard models from Section 2.2. A public (respectively, private) history of observable outputs and compensations (respectively, outputs, compensations and actions) up to but not including period  $t$  is denoted  $h(t)$  (respectively,  $h_\alpha(t)$ ). Denote the output that is realized in period  $G_j$ . Let the period  $t$  compensation paid by the principal, after the public history  $h(t)$  and then the observation of  $G_j$  be denoted  $I_j$ . After observing the private history  $h_\alpha(t)$  and the output/compensation realized in period  $t$ ,  $G_j/I_j$ ,  $j=1,\dots,n$ , the agent takes an action in period  $t+1$ ; denote this action  $a_j$ . Denote the output that is realized in period  $t+1$  (as a consequence of the agent's action  $a_j$ )  $G_k$ ,  $k=1,\dots,n$ . Finally denote the compensation paid to the agent in period  $t+1$  when this output is observed  $I_{jk}$ ,  $j=1,\dots,n$ ,  $k=1,\dots,n$ .

**Proposition 4.1 (Rogerson (1985a)):** Suppose that the principal is risk-neutral and the agent's utility function is separable in action and income. Let  $(\sigma_p, \sigma_\alpha)$  be a second-best contract. After every history  $(h(t), h_\alpha(t))$ , the actions

---

<sup>15</sup>This proof of the result is due to James Mirrlees.

taken by the agent and the compensation paid by the principal must be such that

$$[V'(I_j)]^{-1} = \frac{\delta_p}{\delta_\alpha} \sum_{k=1}^n \varphi_k(a_j) [V'(I_{jk})]^{-1} \quad j=1, \dots, n \quad (4.1)$$

Proof: Pick any history pair  $(h(t), h_\alpha(t))$  in the play of  $(\sigma_p, \sigma_\alpha)$ . As before let  $v_j \equiv V(I_j)$ . Construct a new incentive scheme  $\sigma_p^*$  that differs from  $\sigma_p$  only after  $(h(t), h_\alpha(t))$  and then too in the following special way:  $v_j = v_j^*$ ,  $v_{jk} = v_{jk}^*$  for all  $k$  and  $j \neq j$ , but  $v_j^* = v_j - y$ ,  $v_{jk}^* = v_{jk} + \frac{y}{\delta_\alpha}$  where  $y$  lies in any small interval around zero. In words, in the contract  $\sigma_p^*$ , after the history  $(h(t), G_j)$ , the principal offers a utility "smoothing" of  $y$  between periods  $t$  and  $t+1$ .

It is straightforward to check, given the additive separability of the agent's preferences that the new scheme continues to have a best response of  $\sigma_\alpha$ , the agent's utility is unchanged (and therefore, the scheme is individually rational). Since  $(\sigma_p, \sigma_\alpha)$  is a solution to the principal-agent problem,  $\sigma_p$  is, in fact, the least costly scheme for the principal that implements  $\sigma_\alpha$  (a la Grossman and Hart (1983)). In particular,  $y=0$  must solve the principal's cost minimization exercise along this history. The first-order condition for that to be the case is easily verified to be (4.1).<sup>16</sup>  $\square$

Since the principal can be equivalently imagined to be providing the agent monetary compensation,  $I_j$ , or the utility associated with such compensation  $v_j$ ,  $V^{-1}(v)$  can be thought to be the principal's "utility function". Equation (4.1), and the proof of the proposition, then says that the principal will maintain

<sup>16</sup>In the above argument it was necessary for the construction of the incentive scheme  $\sigma_p^*$  that the principal be able to offer a compensation strictly lower than  $\min(I_j, I_{jk})$ ,  $j=1, \dots, n$ ,  $k=1, \dots, n$ . This, in turn, is possible to do whenever there is unbounded liability which we have allowed. If we restrict the compensations to be at least as large as some lower bound  $\underline{I}$ , then the argument would require the additional condition that  $\min(I_j, I_{jk}) > \underline{I}$ .

intertemporal incentives and provide insurance so as to equate his (expected) marginal utilities across periods.

An immediate corollary of (4.1) is that second-best compensation schemes will be, in general, history-dependent; the compensation paid in the current period will depend not just on the observed current output, but also on past observations of output. To see this note that if  $I_{jk}$ , the compensation in period  $t+1$ , were independent of period  $t$  output,  $I_{jk} = I_{jk}$  for  $j \neq k$ , then the right-hand side of (4.1) is itself independent of  $j$  and hence so must the left-hand side be independent of  $j$ . If  $V$  is strictly concave this can be true only if  $I_j = I_j$  for  $j \neq k$ . But we know that a fixed compensation provides an agent with perverse incentives, from the principal's viewpoint.<sup>17</sup> History-dependence in the second-best contract is also quite intuitive; by conditioning future payoffs on current output, and varying these payoffs appropriately in the observed output, the principal adds a dimension of incentives that are absent in static contracts (which only allow for variations across current payments).

An unfortunate implication of history-dependence is that the optimal contract will be very complex, conditioning as it ought to on various elements of past outputs. Such complexities, as we argued above, fly in the face of reality. An important question then is whether there are environments in which optimal contracts are, in fact, simple in demonstrable ways. Holmstrom and Milgrom (1987) have shown that if the preferences of both principal and agent are multiplicatively separable across time, and if each period's utility is representable by a CARA function, then the optimal contract is particularly simple; the agent performs the same task throughout and his compensation is only based on current

---

<sup>17</sup>The result, that second-best contracts will be history-dependent, was also obtained by Lambert (1983).

output.<sup>18</sup> Since such simplification is to be greatly desired, an avenue to pursue would be to examine the robustness of their result within a larger class of "reasonable preferences."

## 4.2 Simple Contracts

The second-best contracts studied in the previous sub-section had two shortcomings: not all of the inefficiency due to moral hazard is resolved even with long-term contracts and furthermore, the best resolution of inefficiency required delicate conditioning on observable variables. In this sub-section we report some results that remedy these shortcomings. The price that has to be paid is that the results require both principal and agent to be very patient.

The general intuition that explains why efficiency gains are possible in a repeated moral hazard setting is similar to that which underlies the possibility of efficiency gains in any repeated game with imperfect monitoring. Since this is the subject of Section 5.2, we restrict ourselves, for now, to a brief remark. The lifetime payoffs of the agent (see (2.4)) can be decomposed into an immediate compensation and a "promised" future reward. The agent's incentives are affected by variations in each of these components and when the agent is very patient, variations in future payoffs are (relatively) the more important determinant of the agent's incentives. A long-term perspective allows principal and agent to focus on these dynamic incentives.

A more specific intuition arises from the fact that the repetition of the relationship gives the principal an opportunity to observe the results of the agent's

---

<sup>18</sup>Fudenberg, Holmstrom and Milgrom (1990) have shown the result to also hold with additively separable preferences and CARA utility, under some additional restrictions. In this context also see Fellingham, Newman and Suh (1986) who derive first-order conditions like (4.1) for alternative specifications of separability in preferences. They then show that utility functions obeying CARA and/or risk-neutrality satisfy these first-order conditions.

actions over a number of periods and obtain a more precise inference about the likelihood that the agent used an appropriate action. The repetition also allows the principal opportunity to "punish" the agent for perceived departures from the appropriate action. Finally, the fact that the agent's actions in any one period can be made to depend on the outcomes in a number of previous periods provides the principal with an indirect means to insure the agent against random fluctuations in the output that are not due to fluctuations in the agent's actions.

We now turn to a class of simple incentive schemes called **bankruptcy schemes**. These were introduced and analyzed by Radner (1986b); subsequently Dutta and Radner (1992) established some further properties of these schemes.

For the sake of concreteness, in describing a bankruptcy scheme, we will refer to the principal (respectively, the agent) as the owner (respectively, the manager). Suppose the owner pays the manager a fixed compensation per period, say  $w$ , as long as the manager's performance is "satisfactory" in a way that we define shortly; thereafter, the manager is fired and the owner hires a new manager. Satisfactory performance is defined as maintaining a positive "cash reserve", where the cash reserve is determined recursively as follows:

$$\begin{aligned} Y_0 &= y \\ Y_t &= Y_{t-1} + G_t - r, \quad t > 0 \end{aligned} \tag{4.2}$$

The numbers  $y$ ,  $r$  and  $w$  are parameters of the owner's strategy and are assumed to be positive.

The interpretation of a bankruptcy scheme is the following: the manager is given an initial cash reserve equal to  $y$ . In each period the manager must pay the owner a fixed "return", equal to  $r$ . Any excess of the actual return over  $r$  is added to the cash reserve, and any deficit is subtracted from it. The manager is declared "bankrupt" the first period, if any, in which the cash reserve becomes zero



or negative, and the manager is immediately fired. Note that the cash reserve can also be thought of as an accounting fiction, or "score"; the results do not change materially under this interpretation.

It is clear that bankruptcy schemes have some of the stylized features of observable contracts that employ the threat of dismissal as an incentive device and use a simple statistic of past performance to determine when an agent is dismissed. Many managerial compensation packages have a similar structure; evaluations may be based on an industry-average of profits. Insurance contracts in which full indemnity coverage is provided only if the number of past claims is no larger than a prespecified number is a second example.

The principal's strategy is very simple; it involves a choice of the triple  $(y, w, r)$ . The principal is assumed to be able to commit to a bankruptcy scheme. A strategy of the agent, say  $\sigma_\alpha$ , specifies the action to be chosen after every history  $h_\alpha(t)$  – and the agent makes each period's choice from a compact set  $A$ . Suppose that both principal and agent are infinitely-lived and suppose also that their discount factors are the same, i.e.  $\delta_p = \delta_\alpha \equiv \delta$ . Let  $T(\sigma_\alpha)$  denote the time-period at which the agent goes bankrupt; note that  $T(\sigma_\alpha)$  is a random variable whose distribution is determined by the agent's strategy  $\sigma_\alpha$  as well as the level of the initial cash reserve  $y$  and the required average rate of return  $r$ . Furthermore,  $T(\sigma_\alpha)$  may take the value infinity.

The manager's payoffs from a bankruptcy contract are denoted  $U(\sigma_\alpha; y, w, r)$ :

$$U(\sigma_\alpha; y, w, r) \equiv (1-\delta) \sum_{t=0}^{T(\sigma_\alpha)} \delta^t U(a(t), w).^{19}$$

In order to derive the owner's payoffs we shall suppose that each successive manager uses the same strategy. This assumption is justified if successive managers are identical in their characteristics;

---

<sup>19</sup>Implicit in this specification is a normalization which sets the agent's post contract utility level at zero.

the assumption then follows from the principle of optimality.<sup>20</sup> Denote the owner's payoffs  $\mathbf{W}(y, w, r; \sigma_\alpha)$ . Then

$$\mathbf{W}(y, w, r; \sigma_\alpha) = (1-\delta)E \sum_{t=0}^{T(\sigma_\alpha)} \delta^t [r - w - (1-\delta)y] + E \delta^{T(\sigma_\alpha)} [(1-\delta)y + \mathbf{W}(y, w, r; \sigma_\alpha)] \quad (4.3)$$

Collecting terms in (4.3) we get

$$\mathbf{W}(y, w, r; \sigma_\alpha) = r - w - \frac{(1-\delta)y}{1 - E \delta^{T(\sigma_\alpha)}} \quad (4.4)$$

The form of the principal's payoffs, (4.4), is very intuitive. Regardless of which generation of agent is currently employed, the principal always gets per period returns of  $r - w$ . Every time an agent goes bankrupt, however, the principal incurs the cost of setting up a new agent with an initial cash reserve of  $y$ . These expenses, evaluated according to the discount factor  $\delta$ , equal  $\frac{(1-\delta)y}{1 - E \delta^{T(\sigma_\alpha)}}$ ; note that as  $\delta \rightarrow 1$ , this cost converges to  $\frac{y}{E T(\sigma_\alpha)}$ , i.e. the cash cost divided by the frequency with which, on average, this expenditure is incurred.<sup>21</sup>

The dynamic principal-agent problem, (2.2)–(2.4) can then be restated as:

$$\text{Max}_{(y, w, r)} \mathbf{W}(y, w, r; \hat{\sigma}_\alpha) \quad (4.5)$$

$$\text{s.t.} \quad \mathbf{U}(\hat{\sigma}_\alpha; y, w, r) \geq \mathbf{U}(\sigma_\alpha; y, w, r) \quad \forall \sigma_\alpha \quad (4.6)$$

$$\mathbf{U}(\hat{\sigma}_\alpha; y, w, r) \geq \bar{\mathbf{U}} \quad (4.7)$$

Suppose that the first-best solution to the static model is the pair  $(a_F, w_F)$ , where the agent takes the action  $a_F$  and receives the (constant) compensation  $w_F$ . Let the principal's gross expected payoff be denoted  $r_F$ ;  $r_F = \sum_j \varphi_j(a_F) G_j$ . Since

<sup>20</sup>The results presented here can be extended, with some effort, to the case where successive managers have different characteristics.

<sup>21</sup>In our treatment of the cash cost we have made the implicit assumption that the principal can borrow at the same rate as that at which he discounts the future.

the dynamic model is simply a repetition of the static model, this is also the dynamic first-best solution.

We now show that there is a bankruptcy contract in which the payoffs of both principal and agent are arbitrarily close to their first-best payoffs, provided the discount factor is close to 1. In this contract,  $w = w_F$ ,  $r = r_F - \epsilon/2$ , for a small  $\epsilon > 0$ , and the initial cash reserve is chosen to be "large".

**Proposition 4.2 (Radner (1986b)):** For every  $\epsilon > 0$ , there is  $\delta(\epsilon) < 1$  such that for all  $\delta \geq \delta(\epsilon)$ , there is a bankruptcy contract  $(y(\delta), w_F, r_F - \epsilon/2)$  with the property that whenever the agent chooses his strategy optimally, the expected payoffs for both principal and agent are within  $\epsilon$  of the first-best payoffs. It follows that the corresponding second-best contract has this property as well (even if it is not a bankruptcy contract).<sup>22</sup>

Proof: Faced with a bankruptcy contract of the form  $(y(\delta), w_F, r_F - \epsilon/2)$ , one strategy that the agent can employ is to always pick the first-best action  $a_F$ . Therefore,

$$U(\hat{\sigma}_a; y(\delta), w, r) \geq U(a_F, w_F)(1 - E\delta^{T(F)}) \quad (4.8)$$

where  $T(F)$  is the time at which the agent goes bankrupt when his strategy is to use the fixed action  $a_F$ . As  $\delta \uparrow 1$ ,  $(1 - E\delta^{T(F)})$  converges to  $\text{Prob.}(T(F) = \infty)$ . Since the expected output, when the agent employs the action  $a_F$ , is  $r_F$  and the amount that the agent is required to pay out is only  $r_F - \epsilon/2$ , the random walk that the agent controls is a process with positive drift. Consequently,  $\text{Prob.}(T(F) = \infty) > 0$ , and indeed can be made as close to 1 as desired by taking the initial cash reserve,  $y(\delta)$ , sufficiently large (see, e.g. Spitzer (1976, pp. 217–218)). From

---

<sup>22</sup>Using a continuous time formulation for the principal-agent model, Dutta and Radner (1992) are able to, in fact, give an upper bound on the extent of efficiency loss from the optimal bankruptcy contract, i.e. they are able to give a rate of convergence to efficiency as  $\delta \rightarrow 1$ .

(4.8) it is then clear that the agent's payoffs are close to the first-best payoffs whenever  $\delta$  is large.

The principal's payoff will now be shown to be close to his first-best payoffs as well. From the derivation of the principal's payoffs, (4.4), it is evident that a sufficient condition for this to be the case is that the (appropriately discounted) expected cash outlay per period,  $\frac{(1-\delta)y}{1-E\delta T(\sigma_a)}$ , be close to zero at high discount factors (or that the representative agent's tenure,  $ET(\sigma_a)$ , be close to infinity). We demonstrate that whenever the agent plays a best response such is, in fact, a consequence. Write  $U^0$  for  $\max_a U(a, w_F)$ . Since the agent's post-bankruptcy utility level has been normalized to 0, the natural assumption is  $U^0 > 0$ . Note that

$$U^0_{(1-E\delta T(\hat{\sigma}(\delta)))} \geq U(\hat{\sigma}(\delta); y, w, r) \quad (4.9)$$

where  $T(\hat{\sigma}(\delta))$  is the time of bankruptcy under the optimal strategy  $\hat{\sigma}(\delta)$ . Hence,

$$U^0_{(1-E\delta T(\hat{\sigma}(\delta)))} \geq U(a_F, w_F)_{(1-E\delta T(F))}$$

from which it follows that

$$(1-E\delta T(\hat{\sigma}(\delta))) \geq c(1-E\delta T(F)) \quad (4.10)$$

where  $c \equiv U(a_F, w_F)/U^0$ . Substituting (4.10) into the principal's payoffs, (4.4), we get the desired conclusion; the principal's payoffs are close to his first-best payoffs,  $r_F - w_F$ , provided his discount factor is close to 1. The proposition is proved.<sup>23</sup>  $\square$

In a (constant wage) bankruptcy scheme the principal can extend increasing levels of insurance to the agent, as  $\delta \rightarrow 1$ , by specifying larger and larger levels of the initial cash reserve  $y(\delta)$ . The reason that this gives a patient agent the

---

<sup>23</sup>We have not shown that an optimal bankruptcy contract, i.e. a solution to (4.5)–(4.7), exists. A standard argument can be developed to do so (for details, see Dutta and Radner (1992)).

incentive to take actions close to the first-best is suggested by some results in permanent income theory. Yaari (1976) has shown that, for some specifications of bankruptcy, a patient risk-averse agent whose income fluctuates, but who has opportunity to save, would find it optimal to consume his expected income every period; i.e. would want and be able to smooth consumption completely. A bankruptcy scheme can be interpreted as forced consumption smoothing with the principal acting as a bank; an almost patient agent would like to (almost) follow such a strategy anyway.<sup>2425</sup>

The first study of simple contracts to sustain asymptotic efficiency, and indeed the first analyses of repeated moral hazard, were Radner (1981) and Rubinstein (1979). Radner (1981) showed that for sufficiently long but finite principal-agent games, with no discounting, one can sustain approximate efficiency by means of approximate equilibria. Rubinstein showed in an example how to sustain exact efficiency in an infinitely repeated situation with no discounting. For the case of discounting, Radner (1985) showed that approximate efficiency can be sustained, even without precommitment by the principal, by use of **review strategies**. Review strategies are a richer version of bankruptcy schemes. In these schemes the principal holds the agent to a similar performance standard, maintaining an acceptable average rate of return, but a) reviews the agent periodically (instead of every period) and b) in the event of the agent failing to meet the standard, "punishes" him for a length of time (instead of severing the

---

<sup>24</sup>There are some delicate issues that we gloss over; the Yaari model is a pure consumption model (whereas our agent works as well) and bankruptcy as defined in the finite-horizon Yaari model has no immediate analog in the infinite-horizon framework adopted here.

<sup>25</sup>That allowing the agent to save opens up self-insurance possibilities in a repeated moral hazard model, has been argued recently by a number of authors such as Allen (1985), Malcomson and Spinnewyn (1988) and Fudenberg, Holmstrom and Milgrom (1990). In particular, the last paper shows that, even if the agent runs a franchise, and is exposed to all short-term risk, he can guarantee himself an average utility level close to the first-best.

relationship forever). After the punishment phase, the arrangement reverts to the normal review phase. The insurance-incentive trade-off in these schemes is similar to those under bankruptcy schemes.

Radner (1986b) introduced the concept of bankruptcy strategies for the principal, which were described above, and showed that they yield efficient payoffs in the limit as the principal's and agent's discount factors go to 1. Dutta and Radner (1992), in a continuous time formulation, provide a characterization of the optimal contract, within the class of bankruptcy contracts, and establish a lower bound on the rate at which principal-agent values must go to efficiency as  $\delta \rightarrow 1$ .

Up to this point we have assumed that the principal can precommit himself to a particular strategy.<sup>26</sup> Although precommitments can be found in some principal-agent relationships (e.g. customer-supplier, client-broker), it may not be a satisfactory description of many other such relationships. This issue is particularly problematic in repeated moral hazard contracts since at some point the principal has an incentive to renege on his commitment (if his continuation strategy does not solve (4.5)-(4.7) at that history).<sup>27</sup> For example, in a bankruptcy contract the principal has an incentive to renege after the agent, by virtue of either hard work or luck, has built up a large cash reserve (and consequently will "coast" temporarily).

A bankruptcy (or review) strategy can be modified to ensure that the principal has no incentives to renege, i.e. that an equilibrium is perfect. One way to do so would be to modify the second feature of review strategies which was

---

<sup>26</sup>The agent need not however commit to his strategy. Although the incentive constraint, (4.6) may suggest that the agent is committed to his period 0 strategy choice, standard dynamic optimality arguments show that in fact all continuations of his strategy are themselves optimal.

<sup>27</sup>Note that a (limited) precommitment by the principal is also present in static moral hazard contracts; the principal has to abide by his announcement of the incentive scheme after the output consequences of the agent's action are revealed.

described above; an agent is never dismissed but principal and agent temporarily initiate a punishment phase whenever **either** the agent does not perform satisfactorily or the principal reneges on his scheme.<sup>28</sup> We shall not present those results here since in Section 5.2 we discuss the general issue of (perfect) folk theorems in games with imperfect monitoring.

## 5. Games of Imperfect Monitoring

In the principal-agent model of Sections 2-4 there is only one agent whose actions directly affect gross return; the moral hazard is due, therefore, to the unobservability of this agent's actions. In this section we analyze moral hazard issues that arise when there are multiple agents who affect gross returns and whose individual actions are hidden from each other.

The paradigmatic model of many-sided moral hazard is the **partnership** model, in which there is no principal but there are several agents – or partners – who jointly own the productive asset. Any given formula for sharing the return – or output – determines a game; the partners are typically presumed to choose their actions in a self-interested way. The equilibria of the static partnership model are discussed in Section 5.1.

It is not very difficult to see from the above discussion that many-sided moral hazard is an example of the more general class of **games with imperfect monitoring**; indeed, for some of the results that follow it is more instructive to take this general perspective. So Section 5.2 will, in fact, introduce the general framework for such games, present results from **repeated games** with imperfect monitoring and discuss their implication for the repeated partnership model.

---

<sup>28</sup>The standard which the agent is held to, in such a strategy, has to be chosen somewhat carefully. For details on the construction, see Radner (1985).

### 5.1 Partnership Model

A **static** (or stage-game) partnership model is defined by a quintuple  $(A_i, \varphi, G, S_i, U_i; i=1, \dots, m)$ ;  $i$  is the index for a partner, there are  $m$  such partners and each partner picks actions – e.g. inputs – from a set  $A_i$ . Let an action profile be denoted  $a$ ;  $a \equiv (a_1, \dots, a_m)$ . The partners' action profile determines a distribution  $\varphi$  over the set of possible outputs  $G$ . The  $m$ -tuple  $S = (S_1, \dots, S_m)$  is the sharing-rule; partner  $i$ 's share of the total output is  $S_i(G)$ . A sharing-rule must satisfy the **balanced budget** requirement:

$$\sum_i S_i(G) = G, \text{ for all } G \quad (5.1)$$

The appropriate versions of the assumptions (A1)–(A3) will continue to hold. In other words, we will assume that the range of  $G$  is finite, with  $\varphi_j(a)$  denoting the probability that the realized output is  $G_j$ ,  $j=1, \dots, n$ . Furthermore, each  $A_i$  is compact; indeed in the literature, the partners' action sets  $A_i$  have been assumed either to be finite or a (compact) continuum. For the results of this subsection, we will assume the latter; in particular,  $A_i$  will be an interval of the real line.<sup>29</sup> Given a sharing rule  $S$  and the actions of all the partners, partner  $i$ 's expected utility is  $EU_i(S_i(G), a_i)$ . Note that this expected utility depends on the actions chosen by the others only through the effect of such actions on the distribution of the output. Finally, in addition to assuming that  $U_i$  is strictly increasing in the  $i$ -th partner's share, we will also assume that  $\varphi_j$  and  $U_i(S_i(G), \cdot)$  are differentiable functions.

An  $m$ -tuple of inputs is a **Nash equilibrium** of the partnership game if no partner can increase his own payoff by unilaterally changing his input. An

---

<sup>29</sup>All of the results discussed in this sub-section continue to hold when  $A_i$  is a set with a finite number of elements.



**efficient**  $m$ -tuple of inputs is one for which no other feasible input tuple yields each partner at least as much expected utility and yields one partner strictly more.

Note that since each partner only observes the total output, their individual compensations can, at most, depend on this realized output. This is precisely what creates a potential **free-rider** problem; each partner's input generates a positive externality for the other partners and, especially since inputs are unobservable, each partner may therefore have an incentive to provide too little of his input. Two questions can then be posed: i) the normative one: is there a sharing rule under which the free-rider problem can be resolved in that the equilibria of the corresponding game are efficient? (This is the subject of Section 5.1). ii) The positive question: how much of inefficiency is caused by the free-rider problem if the sharing rule is fixed ex-ante? (This is the subject of Section 5.2).

We begin the discussion of partnership models with the case of risk-neutral partners. Recall that with one-sided moral hazard, there are efficient incentive-compatible arrangements between principal and agent when the latter is risk-neutral even in the static game. The first question of interest is whether this result can be generalized, i.e., are there sharing rules for risk-neutral partners such that the Nash equilibria of the resulting partnership game are efficient?

Let  $U(S_i(G), a_i) \equiv S_i(G) - Q_i(a_i)$ . As is well-known, risk neutrality in this case implies that utility is transferable. Hence, the efficiency problem can be written as:

$$\text{Max}_a \sum_j G_j \varphi_j(a) - \sum_i Q_i(a_i) \quad (5.2)$$

Suppose that  $\hat{a}$  is an interior solution to the efficiency problem (5.2). The question we are studying can be precisely stated as: is there a sharing rule  $S$  that satisfies the budget-balancing condition (5.1) and has the property that:

$$\hat{a}_i \in \operatorname{argmax} \sum_j S_i(G_j) \varphi_j(a_i, \hat{a}_{-i}) - Q_i(a_i), \quad \forall i \quad (5.3)$$

An early paper by Holmstrom (1982) suggested an inevitable conflict between budget balance and efficiency; (5.3) can be satisfied only if (5.1) is sacrificed by allowing a surplus in some states of the world ( $\sum_i S_i(G_j) < G_j$ , for some  $j$ ). If it is indeed the case that a residual claimant – or principal – is always required to ensure efficiency in a partnership model, then we could correctly conclude that there is an advantage to an organization with separation between owner and management. As it turns out, the environments in which efficiency and budget balance are incompatible are limited although they contain the (important) symmetric case where each partner's effect on output is identical.

Suppose the distribution of output is affected by the partners' inputs only through some **aggregate variable**; i.e. there are (differentiable) functions  $\theta: A \rightarrow \mathbb{R}$  and  $\xi_j: \mathbb{R} \rightarrow [0,1]$ , such that for all  $j$

$$\varphi_j(a) = \xi_j(\theta(a)) \quad (5.4)^{30}$$

**Proposition 5.1:** Suppose the aggregate effect condition (5.4) holds and suppose further that  $\delta\theta/\delta a_i \neq 0$ , for all  $a, i$ . Then there does not exist any sharing rules that both balances the budget and under which an efficient input profile  $\hat{a}$  is a Nash equilibrium of the partnership game.

**Proof:** Suppose to the contrary that there is such a sharing rule  $S$ . The first-order conditions for efficiency yield (from (5.2)):

---

<sup>30</sup>The condition (5.4) can be shown to be equivalent to the following condition on the derivatives of the likelihood functions: for any pair of partners  $(i,i)$  and any pair of outputs  $(j,j)$  and for all action tuples  $a$ ,  $\varphi_{ji}(a)/\varphi_{ji}(a) = \varphi_{ji}(a)/\varphi_{ji}(a)$ , where  $\varphi_{ji}(a) \equiv \delta\varphi_j/\delta a_i$ .

$$\sum_j G_j \cdot \xi'_j(\theta(\hat{a})) = Q'_i(\hat{a}_i)/\theta'_i(\hat{a}_i) \quad \forall i \quad (5.5)$$

where the notation is:  $\xi'_j \equiv \delta \xi_j / \delta \theta(a)$ ,  $Q'_i \equiv \delta Q_i / \delta a_i$  and  $\theta'_i \equiv \delta \theta / \delta a_i$ . Since  $\hat{a}$  is a Nash equilibrium under sharing rule S, the first-order condition for best response yields (from (5.3)):

$$\sum_j S_i(G_j) \cdot \xi'_j(\theta(\hat{a})) = Q'_i(\hat{a}_i)/\theta'_i(\hat{a}_i) \quad \forall i \quad (5.6)$$

Note that the right-hand sides of (5.5) and (5.6) are identical. Summing the left-hand side of (5.5) over the index  $i$  yields  $m \cdot \sum_j G_j \xi'_j(\theta(\hat{a}))$ . A similar summation over the left-hand side of (5.6) yields, after invoking the budget balance condition (5.1),  $\sum_j G_j \xi'_j(\theta(\hat{a}))$ . Since  $m > 1$ , the two sums do not agree and we have a contradiction.  $\square$

**Remark:** When the output is deterministic, i.e.  $G = \theta(a)$ , the aggregate condition is automatically satisfied.<sup>31</sup> Indeed, this was precisely the case studied by Holmstrom (1982). Similarly, if there are only two outcomes the condition is automatically satisfied as well; this can be seen, for example, from the equivalent representation for this condition (see footnote 3), that the derivatives of the likelihood ratios are equal across partners. Finally, if the agents are symmetric in their effect on output, the aggregate condition is immediate as well.<sup>32</sup>

Williams and Radner (1989) were the first to point out that in order to resolve the organization problem, even with risk-neutral partners, it is necessary that there be some **asymmetry** in the effect that each partner's input has on the

---

<sup>31</sup>In this case,  $\xi$  cannot, evidently, be a differentiable function. The appropriate modification of the proof of Proposition 5.1 is, however, immediate.

<sup>32</sup>The aggregate condition on output distribution is also at the heart of the Radner, Myerson and Maskin (1986) example of a repeated game with discounting in which efficiency cannot be sustained as an equilibrium outcome. This was pointed out by Matsushima (1989).

distribution of the aggregate output. The intuition is quite straightforward: if, at the margin, the effect of partner  $i$ 's input is more important in the determination of output  $j$  than the determination of output  $k$  (and the converse is true for partner  $\hat{i}$ 's input vis-a-vis outputs  $k$  and  $j$ ), then the share assigned to  $i$  for output  $j$  has relatively greater impact on his decisions than the share assigned for output  $k$ ; a parallel argument works for partner  $\hat{i}$  and outputs  $k$  and  $j$ . Consequently, shares can be assigned in order to give the partners' appropriate incentives. Indeed, Williams and Radner (1989) show that generically, in the space of distributions, there are sharing rules that do resolve the free-rider problem and balance the budget.

Since the aggregate condition (5.2) is no longer being assumed, the first-order condition (5.6), for the efficient input profile  $\hat{a}$  to be a Nash equilibrium, can be written as:

$$\sum_j S_i(G_j) \cdot \varphi_{ji}(\hat{a}) = Q'_i(\hat{a}_i) \quad \forall i \quad (5.7)$$

If we wish to design the sharing rule  $S$  so that  $\hat{a}$  satisfies the first-order conditions for an equilibrium, then the  $mn$  unknowns,  $[S_i(G_j)]$ , must satisfy the  $(m+n)$  equations implied by equations (5.7) and the budget balance condition (5.1). The basic lemma of Williams and Radner (1989), reported below, is that, generically in the data of the model, such a solution can be found if  $n > 2$ , and that in particular this can be done if  $\hat{a}$  is an efficient vector of inputs. Of course, to complete the argument it remains to show that there are reasonable conditions under which a solution to the "first-order" problem is actually an equilibrium.

**Theorem 5.2 (Williams and Radner (1989)):**     i) When  $n > 2$ , there exists a solution to the first-order conditions, (5.7), and the budget balance conditions, (5.1), for each pair of distribution and utility functions  $(\varphi_j, Q_i, j=1, \dots, m, i=1, \dots, n)$  in

some open dense subset (in the Whitney  $C^1$  topology) of the set of feasible problems.

ii) Suppose that  $m=2$  and  $n=3$ . Assume further that  $\varphi_j$  is first-order stochastically increasing in its arguments, for  $j=1, \dots, 3$ . Then there exists a one-parameter solution to the problem of finding a sharing rule whose Nash equilibrium is efficient if the following two additional hypotheses are satisfied at the efficient input profile  $\hat{a}$ :

a)  $\varphi_{11}(\hat{a})\varphi_{22}(\hat{a}) - \varphi_{21}(\hat{a})\varphi_{12}(\hat{a}) > 0$

b)  $\varphi_{21}(\cdot, \hat{a}_2)/\varphi_{11}(\cdot, \hat{a}_2)$  is an increasing function whereas  $\varphi_{22}(\hat{a}_1, \cdot)/\varphi_{21}(\hat{a}_1, \cdot)$  is a decreasing function.

Other conditions for solutions to this problem in static partnership models have been presented by Legros and Matsushima (1991), Legros (1989) and Matsushima (1989b). All of these conditions resonate with the idea that "symmetry is detrimental to efficiency" in partnership models.<sup>33</sup>

A question of some interest is what properties will efficiency-inducing sharing rules possess. In particular, as in the case of principal-agent models, we can ask: will the sharing rules be **monotonically increasing** in output, i.e. will a higher output increase the share of **all** partners? Of course, such a question makes sense only when higher inputs do, in fact, make higher outputs more likely – i.e.,  $\varphi_j(\cdot)$  is first-order stochastically increasing – since a higher output may then be taken as a signal of higher inputs and appropriately rewarded. It is easy to show, however, that such **monotonicity is incompatible with efficiency**. The intuition is straightforward: if all partners benefit from a higher output then the social benefit

---

<sup>33</sup>Interestingly, the results from the static model, with risk-neutral partners, will turn out to be very helpful in the subsequent analysis of the repeated partnership model with general (possibly risk-averse) utility functions. This is because intertemporal expected utility will be seen to have a decomposition very similar to that between the monetary transfer and the input-contingent expected utility in the static risk-neutral case; this point will be clearer after our discussion in Section 5.2.

to any one partner increasing his input is greater than that partner's private benefit. However, in an equilibrium that is efficient, the private and social benefits have to be equal. This idea is formalized as:

**Proposition 5.3:** Suppose that  $\varphi_j(a_1, \dots, a_m)$  is strictly first-order stochastically increasing in its arguments. Let  $S$  be a sharing rule for which the first-best profile of inputs  $\hat{a}$  is a Nash equilibrium. Then there is some partner, say  $i$ , whose share,  $S_i$ , does not always increase with output.

**Proof:** Suppose, to the contrary, that the sharing rules,  $S_i$ , are increasing for all partners. The social marginal benefit to increasing partner  $i$ 's input is:

$$\sum_i [\sum_j S_i(G_j) \cdot \varphi_{ji}(\hat{a})] - Q'_i(\hat{a}_i) \quad (5.8)$$

Since  $\hat{a}_i$  is a best response,  $\sum_j S_i(G_j) \cdot \varphi_{ji}(\hat{a}) = Q'_i(\hat{a}_i)$ . Substituting this into (5.8) yields  $\sum_{i \neq j} \sum_j S_i(G_j) \cdot \varphi_{ji}(\hat{a})$ ; the assumption on first-order stochastic dominance implies that  $\sum_j S_i(G_j) \cdot \varphi_{ji}(\hat{a}) > 0$ , for all  $i \neq j$ . Hence, social utility would be increased by expanding partner  $i$ 's input.  $\square^{34}$

**Remark:** One immediate corollary of the proposition is that the proportional sharing rule,  $S_i(G_j) = G_j/m$ , does not solve the organization problem by inducing efficient provision of inputs.

If partners' utility functions exhibit risk-aversion there will not be, in general, a sharing rule that sustains the efficient outcome as an equilibrium. This is because efficiency arrangements in this case requires efficient risk-sharing as well as efficient provision of inputs. To see this note that an efficient solution to the

---

<sup>34</sup>It is obvious, from the proof, that the proposition is true as long as  $\varphi_j$  is first-order stochastically increasing in the input level of at least one partner.

partnership problem is given by any solution to the following:  $\text{Max}_{a, S_i} \sum_i \lambda_i [\sum_j U_i(S_i(G_j), a_i) \cdot \varphi_j(a)]$ , where  $\lambda_i > 0$ ,  $i=1, \dots, m$ . (There may also be efficient expected utility vectors corresponding to some choices of  $(\lambda_1, \dots, \lambda_m)$  with some coordinates  $\lambda_i$  equal to zero.) A solution to the above maximization problem will typically involve not just the action profile, as with risk-neutrality, but also the sharing rules. Moreover, the rules that share risk efficiently may not be incentive-compatible. An alternative way of seeing this is to note that if the sharing rules are specified in the efficient solution then there is no further degree of freedom left to specify these rules such that they also satisfy the Nash equilibrium first-order conditions (5.7).

There are several questions that remain open in the static partnership model. The first involves the characterization of solutions to the second-best problem; what are the properties of the most efficient Nash equilibrium in a partnership game (when the efficient solution is unattainable)? In particular, it may be fruitful to employ the Grossman and Hart (1983) cost-minimization approach to the partnership problem to investigate properties such as monotonicity and convexity for the second-best solution.<sup>35</sup>

A related question to ask is whether input profiles (and sharing rules) (arbitrarily) close to the efficient vector can, in fact, be implemented in an incentive-compatible fashion, even if exact efficiency is unattainable. Recent developments in implementation theory have introduced weaker (and yet compelling) notions that involve implementability of profiles close to that desired; see Abreu and Matsushima (1992). This approach may even be able to dispense

---

<sup>35</sup>Mookherjee (1984) has studied the related problem in which there is a principal, in addition to the partners. He characterizes the second-best outcome from the principal's perspective.

with the requirement that partners be asymmetric.<sup>36</sup>

We now turn to the general class of games with imperfect monitoring.

## 5.2 Repeated Games with Imperfect Monitoring

A static (or stage) game with imperfect monitoring is defined by a triple  $(A_i, \varphi_j, U_i; i=1, \dots, m, j=1, \dots, n)$ ;  $i$  is the index for a player and each player picks actions  $a_i$  from a finite set  $A_i$ . This choice is not observed by any player other than  $i$ . An action profile  $a \equiv (a_1, \dots, a_m)$  induces a probability distribution on a public outcome  $G_j$ ,  $j=1, \dots, n$ ; the probability that the outcome is  $G_j$  when the action profile  $a$  is chosen is denoted  $\varphi_j(a)$ . Each player's realized payoff depends on the public outcome and his own action but not on the actions of the other players; the payoff is denoted  $U_i(G, a_i)$ . We will allow players to pick mixed actions as well; denote a generic mixed action by  $\alpha_i$ . For each profile of mixed actions  $\alpha = (\alpha_1, \dots, \alpha_m)$ , the conditional probability of public outcomes and the player's expected payoffs are computed in the obvious fashion. Abusing notation, we write  $\varphi_j(\alpha)$  to be the probability of the outcome  $G_j$  under the mixed action profile  $\alpha$ . It will be useful to denote player  $i$ 's expected payoffs as  $\Gamma_i(\alpha)$ .

It is clear that a partnership model, with a fixed sharing rule, is an example of a game of imperfect monitoring. So also is the principal-agent model of Sections 2–4. Imagine that player 2 is the principal and his action is the choice of a compensation scheme for the agent. Since the agent actually moves after the principal – whereas in the above game, moves are simultaneous –  $a_1$  now must be interpreted as a contingent effort rule that specifies the agent's effort for every compensation rule that the principal could choose. The public outcome is then the

---

<sup>36</sup>Legros (1989) shows that even in the deterministic partnership model (with risk-neutrality),  $\epsilon$ -efficiency can be attained if partners are allowed to randomize in their choice of inputs.



realized output level plus the principal's compensation scheme.<sup>37</sup>

In a repeated game with imperfect monitoring, in each period  $t = 0, 1, \dots$ , the stage game is played and the associated public outcome revealed. The public history at date  $t$  is  $h(t) = (G(0), G(1), \dots, G(t-1))$  whereas the private history of player  $i$  is  $h_i(t) = (a_i(0), G(0), a_i(1), G(1), \dots, a_i(t-1), G(t-1))$ . A strategy for player  $i$  is a sequence of maps  $\sigma_i(t)$ , where  $\sigma_i(t)$  assigns to each pair of public and private histories  $(h(t), h_i(t))$  a mixed action  $\alpha_i(t)$ . A strategy profile induces, in the usual way, a distribution over the set of histories  $(h(t), h_i(t))$  and hence an expected payoff for player  $i$  in the  $t$ -th period; denote this  $\Gamma_i(t)$ . Lifetime payoffs are evaluated under a (common) discount factor  $\delta$  ( $< 1$ ) and equal  $(1-\delta) \sum_{t=0}^{\infty} \delta^t \Gamma_i(t)$ .

Player  $i$  seeks to maximize his lifetime payoffs. We restrict attention to a subclass of Nash equilibria that have been called **perfect public equilibria** in the literature; a strategy profile  $(\sigma_1, \dots, \sigma_m)$  is a perfect public equilibrium if, a) for all time-periods  $t$  and all players  $i$ , the continuation of  $\sigma_i$  after history  $(h(t), h_i(t))$  only depends on the public history  $h(t)$  and b) the profile of continuation strategies constitute a Nash equilibrium after every history.<sup>38</sup> Suppose that the stage game has a Nash equilibrium; an infinite repetition of this stage-game equilibrium is an example of a perfect public equilibrium. Let  $V$  denote the set of payoff vectors corresponding to all perfect public equilibria in the repeated game.

---

<sup>37</sup>Fudenberg, Levine and Maskin (1989), who suggest the above interpretation of the principal-agent model, show that several other models can also be encompassed in the current framework. In particular, the oligopoly models of Porter (1983) and Green and Porter (1984) are easily accommodated.

<sup>38</sup>Note that a player is not restricted to choosing a strategy in which he can only condition on the public history. If every other player but  $i$  chooses such a strategy, elementary dynamic programming arguments can be used to show that player  $i$  cannot, in his best response problem, do any better by choosing a strategy that conditions on his private history as well. A second point to note is that were we to restrict attention to pure strategies only, then without any loss of generality we could in fact restrict players to choosing strategies which only condition on public histories (for this and related arguments see Abreu, Pearce and Stachetti (1990)).

### 5.2.1 A Characterization of the Set of Equilibrium Payoffs

In this subsection we will provide an informal discussion of the Abreu, Pearce and Stachetti (1986, 1990) recursive characterization of the equilibrium payoff set  $V$ . The heart of their analysis is to demonstrate that the set of equilibrium payoffs in repeated games has a Bellman-equation like representation similar to the one exhibited by the value function in dynamic programming.

Suppose, to begin with, that we have a perfect public equilibrium profile  $\sigma^*$ . Such an equilibrium can be decomposed into a) an action profile in period zero, say  $\alpha^*(0)$ , and b) an expected continuation payoff (or "promised future payoff") profile,  $v^j(1)$ ,  $j=1,..n$ , that is contingent on the public outcome  $G_j$  realized in period zero. Since  $\sigma^*$  is an equilibrium it follows that: i)  $\alpha^*(0)$  must satisfy the incentive constraint that no player can unilaterally improve his payoffs given the twin expectations of other players' actions in that period,  $\alpha_{-i}^*(0)$ , and the continuation payoffs,  $v^j(1)$ ; ii) the continuation payoffs must themselves be drawn from the set of equilibrium payoffs  $V$ . Moreover, an identical argument is true for every equilibrium strategy profile and after all histories, i.e. an equilibrium in the repeated game is a sequence of incentive-compatible "static" equilibria.

Now consider an arbitrary set of continuation payoffs  $W \subset \mathbb{R}^m$ ; these need not be equilibrium payoffs. Define an action profile  $\hat{\alpha}$  to be enforceable, with respect to  $W$ , if there are payoff profiles  $w^j \in W$ ,  $j=1,..n$  with the property that  $\hat{\alpha}$  is a Nash equilibrium of the "static" game with payoffs  $(1-\delta)\Gamma_i(\alpha) + \delta Ew_i(\alpha)$ . Let  $B(W)$  be the set of Nash equilibrium payoffs to these "static" games (with all possible continuation payoffs being drawn from  $W$ ). If a bounded set  $W$  has the property that  $W \subset B(W)$ , (and Abreu, Pearce and Stachetti call such a set **self-generating**) then it can be shown that all payoffs in  $B(W)$  are actually

repeated-game equilibrium payoffs, i.e.  $B(W) \subset V$ .<sup>39</sup> In other words a sequence of static Nash equilibria, all of whose payoffs are self-referential in the manner described above, is a perfect equilibrium of the repeated game.

More formally let us define, for any set  $W \subset \mathbb{R}^m$ :

$$B(W) \equiv \{w \in \mathbb{R}^m: \exists w^j \in \mathbb{R}^m, j=1, \dots, n, \text{ and } \exists \alpha \text{ s.t.}$$

$$w_i = (1-\delta)\Gamma_i(\alpha) + \delta \sum_j w_i^j \varphi_j(\alpha) \quad (5.9)$$

$$(1-\delta)\Gamma_i(\alpha) + \delta \sum_j w_i^j \varphi_j(\alpha) \geq (1-\delta)\Gamma_i(a_i, \alpha_{-i}) + \delta \sum_j w_i^j \varphi_j(a_i, \alpha_{-i}) \quad \forall i, a_i \} \quad (5.10)$$

**Theorem 5.4 (Abreu, Pearce and Stachetti (1990)):** i) (Sufficiency) Every bounded self-generating set is a subset of the equilibrium payoffs set; if  $W$  is bounded and  $W \subset B(W)$  then  $B(W) \subset V$ .

ii) (Necessity) The equilibrium payoffs set  $V$  is the largest self-generating set among the class of bounded self-generating sets;  $V = B(V)$ .

The recursive approach has two useful consequences. The sufficiency characterization, part i), says that if a subset of feasible payoffs can be shown to be self-generating, then all of its elements are equilibrium payoffs; this (constructive) approach can be used to provide upper bound on the difference between second-best and efficient payoffs. The constructive approach is used in proving the folk theorem that we discuss shortly.

---

<sup>39</sup>The argument is as follows: by definition, if  $\hat{w}$  is in  $B(W)$ , then it can be decomposed into an action profile  $\hat{\alpha}(0)$  and "continuation" payoffs  $\hat{w}(1)$  where  $\hat{\alpha}(0)$  is a Nash equilibrium in the "static" game with payoffs  $(1-\delta)\Gamma_i(\alpha) + \delta E w_i(\alpha)$ . Since  $\hat{w}(1) \in W \subset B(W)$ , it can also be similarly decomposed. In other words there is a strategy  $\hat{\sigma}$ , which can be deduced from these arguments, with the property that no one-shot deviation against it is profitable for any player. The unimprovability principle of discounted dynamic programming then implies that there are, in fact, no profitable deviations against  $\hat{\sigma}$ .

A second (inductive) approach can be employed to determine properties of the second-best equilibrium. In this approach one conjectures that the equilibrium payoff set has the properties one seeks to establish and then demonstrates that such properties are, in fact, maintained under the recursion. Abreu, Pearce and Stachetti (1986, 1990) have utilized this approach to provide conditions on the primitives under which the equilibrium payoff set is compact, convex and monotonically increasing in the discount factor.

### 5.2.2 A Folk Theorem with Imperfect Monitoring

We turn now to the second-best problem: how large is the inefficiency caused by free-riding in a game of imperfect monitoring? The repeated perspective allows the design of a richer set of incentives. This is immediate from the incentive constraint, (5.10), above; by way of different specifications of the promised future payoff,  $w_i^j$ , there is considerable room to fine-tune current incentives. However, for the future to have significant bearing on the players' decisions today it must be the case that there is sufficient weight attached to the future; indeed, the folk theorem, which answers the second-best question, is, in fact, an asymptotic result that obtains for players with  $\delta$  close to 1.

There is a formal similarity between the payoffs to risk-neutral players in the static partnership model and the intertemporal decomposition of payoffs in the repeated game, (5.9). In the static model, the player's (risk-neutral) payoffs are  $\sum_j S_i(G_j)\varphi_j(\alpha) - Q_i(\alpha_i)$ . The intertemporal payoff is  $(1-\delta)\Gamma_i(\alpha) + \delta\sum_j w_i^j \varphi_j(\alpha)$ , where  $w_i^j$  can be interpreted as player  $i$ 's "share" of future payoffs resulting from an output  $G_j$ .<sup>40</sup> The difference in the analyses is that the future payoffs  $w_i^j$ , have to

---

<sup>40</sup>Indeed it is easy to check that in all of the analyses of the static risk-neutral case, the own action-contingent utility,  $Q_i(\alpha_i)$ , could be replaced with the exact analogue in (5.9),  $\Gamma_i(\alpha) \equiv E[U_i(G, a_i) | \alpha]$ , without changing any of the results. Hence, the correspondence, between the intertemporal (possibly risk-averse) payoffs and the

be self-generated whereas the static shares  $S_i(G_j)$ , have to satisfy budget-balance,  $\sum_i S_i(G_j) = G_j$ .

It turns out, however, that the analogy between the two models goes a little further still. This is because the folk-theorem proof techniques of Fudenberg, Levine and Maskin (1989) and Matsushima (1989) critically employ a construction in which actions are enforced by continuation payoffs  $w_i^j$  that are restricted to lie on a hyperplane, i.e., are such that  $\sum_i w_i^j = 0$ . This is, of course, a restatement of the budget balance condition, after relabelling variables. This explains why one hypothesis of the folk theorem below is an asymmetry condition like the ones that were used in solving the static risk-neutral incentives problem:

**Pairwise Full Rank** The stage game satisfies the pairwise full rank condition if, for all pairs of players  $(i, i \neq j)$  there exists a profile  $\alpha$  such that the matrix

$$\left\{ \varphi_j(a_i, \alpha_{-i}), \varphi_i(a_i, \alpha_{-i}) \right\}$$

with rows corresponding to the elements of  $A_i \times A_i$  and columns corresponding to the outcomes  $G_j$ ,  $j = 1, \dots, n$ , has rank  $|A_i| + |A_i| - 1$ .<sup>41</sup>

**Theorem 5.5 (Fudenberg, Levine and Maskin (1989)):** Suppose the stage game satisfies the pairwise full rank condition and, additionally, the following two hypotheses:

- i) for all players  $i$  and all pure action profiles  $\hat{a}$ , the  $|A_i|$  vectors,  $\{\varphi_j(a_i, \hat{a}_{-i}), j=1, \dots, n\}$ , are linearly independent
- ii) the set of individually rational payoff vectors, say  $F^*$  has dimension equal to

---

static risk-neutral payoffs, is exact.

<sup>41</sup>This is really a full rank condition since the row vectors must always admit at least one linear dependency. Also, a necessary condition for the pairwise full rank condition to hold is clearly that the number of outcomes  $n \geq |A_i| + |A_i| - 1$ .

the number of players.<sup>42</sup>

Then, for every closed set  $W$  in the relative interior of  $F^*$  there is a  $\underline{\delta} < 1$  such that for all  $\delta > \underline{\delta}$ , every payoff in  $W$  is a perfect public equilibrium payoff.

**Remarks:** 1. An obvious implication of the above result is that, as  $\delta \rightarrow 1$ , any corresponding sequence of second-best equilibrium payoffs is asymptotically efficient.

2. In the absence of the pairwise full rank condition, asymptotic efficiency may fail to obtain; Radner, Myerson and Maskin (1986) contains the appropriate example. In this example, the aggregate condition of Section 5.2 (condition (5.4)) holds and hence, for reasons identical to the static risk-neutral case, the efficient solution cannot be sustained as equilibrium behavior.

3. Condition i) is required to ensure that player  $i$ , when called upon to play action  $\hat{a}_i$  cannot play a more profitable action  $\tilde{a}_i$  whose outcome consequences are identical to those of  $\hat{a}_i$ . Condition ii) is required to ensure that there are feasible asymmetric lifetime payoffs; when a deviation by player  $i$  (respectively  $i$ ) is inferred, there exists a continuation strategy  $\sigma^i$  whose payoffs are smaller than the payoffs to some other continuation strategy  $\sigma^i$  (and vice-versa for player  $i$ ).<sup>43</sup>

4. Since the principal-agent model of Sections 2-4 is a special case of the game with imperfect monitoring, Theorem 5.5 also yields a folk theorem, and asymptotic efficiency, in that model. However, the informational requirements of this result are evidently more stringent than those employed by Radner (1985) to prove asymptotic efficiency in the principal-agent model. Restricted to the

---

<sup>42</sup>Note that mixed strategies are admissible and hence an individually rational payoff vector is one whose components dominate the mixed strategy minimax for each player.

<sup>43</sup>A similar condition is also required in repeated games with perfect monitoring; see Fudenberg and Maskin (1986). Recent work (Abreu, Dutta and Smith (1992)) has shown that, under perfect monitoring, the full-dimensionality assumption can be replaced by the weaker requirement that players' preferences not be representable by an identical ordering over mixed action profiles. Whether full-dimensionality can be similarly weakened in the imperfect monitoring case remains an open question.

principal-agent model, Theorem 5.5 can, however, be proved under weaker hypotheses, and Radner's result can be generalized; see Fudenberg, Levine and Maskin (1989) for details.

To summarize, under certain conditions, repetition of a partnership allows the design of intertemporal incentives such that the free-rider problem can be asymptotically resolved by patient partners. Of course, the curse of the Folk Theorem, Theorem 5.5, is that it proves that a lot of other, less attractive, arrangements can also be dynamically sustained.

In many partnerships there are variables other than just the partners' actions that determine the output, as for example the usage of (commonly owned) capital stock; this is an additional source of information. One question of interest, for both static and especially repeated partnership models, is whether, and how much, such additional information alleviates the free-rider problem. A second open question is whether bounds can be derived for the rate of convergence to efficiency as the discount factor goes to one (in a spirit similar to the Dutta and Radner (1992) exercise for the principal-agent model).

## 6. Additional Bibliographical Notes

**Notes on Section 3:** In the field of economics, the first formal treatment of the principal-agent relationship and the phenomenon of moral hazard was probably given by Arrow (1963, 1965), although a paper by Simon (1951) was an early forerunner of the principal-agent literature.<sup>44</sup> Early work on one-sided moral

---

<sup>44</sup>The recognition of incentive problems is of much older vintage however. In an often quoted passage from the "Wealth of Nations" Adam Smith says, "The directors of such companies, being the managers rather of other peoples' money than their own, it cannot well be expected, that they should watch over it with the same anxious vigilance with which the partners in a private copartnery frequently watch over their own... Negligence and profusion therefore must always prevail in the management of the affairs of such a company." (p. 700 *ibid*).

hazard was done by Wilson (1969), Spence and Zeckhauser (1971) and Ross (1973). James Mirrlees contributed early, and sophisticated, analyses of the problem; much of his work is unpublished (but see Mirrlees (1974, 1976)). Holmstrom (1979) and Shavell (1979) investigated conditions under which it is beneficial for the principal to monitor the agent, or use any other sources of information about the agent's performance in writing the optimal contract. In addition to these papers other characterizations of the second-best contract, all of which employ the first-order approach, include Harris and Raviv (1979) and Stiglitz (1983). (For analyses that provide sufficient conditions under which the first-order approach is valid, see Rogerson (1985b) and Jewitt (1988)). The paper by Grossman and Hart (1983) provides a particularly thorough and systematic treatment of the one-period model.

The static principal-agent model has been widely applied in economics. As we have indicated earlier, the phenomenon, and indeed the term itself, came from the study of insurance markets. One other early application of the theory has been to agrarian markets in developing economies; a principal question here is to understand the prevalence, and uniformity, of sharecropping contracts. An influential paper is Stiglitz (1974) that has subsequently been extended in several directions by, for example, Braverman and Stiglitz (1982); for a recent survey of this literature, see Singh (1989). Other applications of the theory include managerial incentives to a) invest capital in productive activities, rather than perquisites (Grossman and Hart (1982)), b) to invest in human capital, (Holmstrom and Ricart i Costa (1986)) and c) to obtain information about and invest in risky assets, (Lambert (1986)).

Three topics in static moral hazard that we have not touched upon are: a) incentive issues when moral hazard is confounded with informational asymmetries due to adverse selection; see, for example, Foster and Wan (1984) who investigate involuntary unemployment due to such a mix of asymmetries, and defense



contracting issues as studied by Baron and Besanko (1987) and McAfee and McMillan (1986); b) the general equilibrium consequences of informational asymmetries, a topic that has been studied in different contexts by Joseph Stiglitz and his coauthors; see, for example, Arnott and Stiglitz (1986) and Greenwald and Stiglitz (1986) and c) the implications of contract renegotiation. The last topic asks the question, (when) will principal and agent wish to write a new contract to replace the current one and has been recently addressed by a number of authors including Fudenberg and Tirole (1990). The survey of principal-agent models by Hart and Holmstrom (1986) is a good overview of the literature; furthermore, it expands on certain other themes that we have not been able to address.

**Notes on Section 4:** Lambert (1983) derives a characterization of dynamic second-best contracts that also yields the history-dependence implied by Rogerson's result; he takes, however, a first-order approach to the problem. Spear and Srivastava (1988) employ the methods of Abreu, Pearce and Stachetti (1990) to derive further characterizations of the optimal compensation scheme, such as monotonicity in output.

The second strand in the literature on dynamic principal agent contracts has explored the implications of simple contracts that condition on history in a parsimonious fashion. Relevant papers here are Radner (1981, 1985, 1986b), Rubinstein (1979), Rubinstein and Yaari (1983), and Dutta and Radner (1992). These papers have been reviewed in some detail in Section 4.2.

Recently a number of papers have investigated the consequences of allowing the agent to borrow and lend and thereby provide himself with self-insurance. Indeed if the agent is able to transact at the same interest rates as the principal, an assumption that is plausible if capital markets are perfect (but only then), there exist simple output contingent schemes (that look a lot like franchises) which

approximate efficiency. Papers in this area include Allen (1985), Malcomson and Spinnewyn (1988), and Fudenberg, Holmstrom and Milgrom (1990).

In the study of labor contracts a number of authors have investigated some simple history-dependent incentive schemes under which an employee cannot be compensated on the basis of observed output but rather has to be paid a fixed wage; the employee may, however, be fired in the event that shirking is detected.<sup>45</sup> Shapiro and Stiglitz (1984) show that involuntary unemployment is necessary, and will emerge, in the operation of such incentive schemes. An application of these ideas to explain the existence of dual rural labor markets in agrarian economies can be found in Easwaran and Kotwal (1985).

**Notes on Section 5:** Static partnership models were first studied formally by Holmstrom (1982) – see also the less formal discussion in Alchian and Demsetz (1972). For characterizations of conditions under which the first-best is sustainable as an equilibrium by risk-neutral partners, see, in addition to the Williams and Radner (1989) paper that we have discussed, Legros (1989), Matsushima (1989b) and Legros and Matsushima (1991). For a discussion of the case of risk-averse partners see Rasmussen (1987).

Mookherjee (1984) generalized the Grossman and Hart (1983) approach to single-sided moral hazard problems to characterize the second-best contract when there is a principal and several agents (or partners). (His framework covers both the case of a partnership, where production is joint, as well as the case of independent production). He derived an optimality condition that is the analog of condition (3.4) above and used this to investigate conditions under which a) an

---

<sup>45</sup>These contracts therefore bear a family resemblance to the bankruptcy schemes we have discussed in this paper; the one difference is that in a bankruptcy scheme observed output is utilized in deciding whether or not to terminate an agent's contract whereas in the papers referred to here it is usually assumed that shirking can be directly observed (and penalized).

agent's compensation should be independent of other agents' output, and b) agents' compensations should be based solely on their "rank" (an ordinal measure of relative output). The attainability of first-best outcomes through rank order tournaments has also received extensive treatment in the context of labor contracts; see Lazear and Rosen (1981) for the first treatment and subsequent analyses by Green and Stokey (1983) and Nalebuff and Stiglitz (1983).

Radner (1986a) was the first paper to study repeated partnerships; in his model partners do not discount the future but rather employ the long-run average criterion to evaluate lifetime utility. This paper showed that the efficient expected utility vectors **can** be sustained as a perfect equilibrium of a repeated partnership game (even under risk-aversion) for a "large" class of partnership models.<sup>46</sup> Subsequent work, which has incorporated discounting of future payoffs, has included Radner, Myerson and Maskin (1986), Abreu, Milgrom and Pearce (1991) and Radner and Rustichini (1989). Radner, Myerson and Maskin (1986) gave an example in which equilibrium payoffs for the repeated game with discounting are uniformly bounded away from one-period efficiency for all discount factors strictly less than one.<sup>47</sup> A model formally similar to a repeated partnership is that of a

---

<sup>46</sup>The exact condition that needs to be satisfied is as follows: fix a sharing rule  $S$  and suppose  $\hat{a}$  is the (efficient) input profile that is to be sustained. Then there exist positive constants  $K_i$  such that  $[EU_i(S_i(G), \hat{a}) - EU_i(S_i(G), \hat{a}_{-i}, a_i)] + K_i[E(G|\hat{a}_{-i}, a_i) - E(G|\hat{a})] \leq 0$ , for all  $a_i$ , where the expectations are taken under the input profiles  $\hat{a}$  and  $(\hat{a}_{-i}, a_i)$ .

<sup>47</sup>As discussed above there is a formal similarity between the two models, static partnerships in which the sharing rule is endogenous and a discounted repeated partnership with a fixed sharing rule (but endogenous future compensation). In particular in both of these cases an ability to treat partners asymmetrically is essential to the sustainability of efficient outcomes. The Radner, Myerson and Maskin example restricts itself to the aggregate condition (5.4), much as the Holmstrom analysis did in the static partnership context (and furthermore only considers a symmetric sharing rule). Hence in both cases efficiency cannot be sustained. The undiscounted case is different in that the asymmetric treatment of

oligopoly with unobserved quantity choices by each firm; this model was studied by Porter (1983), Green and Porter (1984) and Abreu, Pearce and Stachetti (1986).

Abreu, Pearce and Stachetti (1986, 1990) contain the analyses, reported in in Section 5.2, that characterize the equilibrium payoff set in repeated games with discounting and imperfect monitoring. Fudenberg, Levine, and Maskin (1989), and Matsushima (1989) have employed the sufficiency part of the characterization to prove that efficiency is sustainable, even with many-sided moral hazard, for a large class of repeated games with imperfect monitoring. These results are important in that they make clear the conditions needed to give agents appropriate dynamic incentives in order to sustain efficiency.

As noted in Section 1, models of adverse selection and misrepresentation have not been discussed in the present article. For this topic the reader is referred to Melumad and Reichelstien (1989) and the references cited there. For a recent survey of the related topic of incentive compatibility and the revelation of preferences for public goods, see Groves and Ledyard (1987). An elementary exposition of the problem of misrepresentation and the Groves-Vickrey-Clark mechanism can be found in Radner (1987), and further work on incentive compatibility appears in the volume edited by Groves, Radner and Reiter (1987).

---

partners is inessential to the sustainability of efficient outcomes.

## References

- Abreu, D., P. K. Dutta and L. Smith, 1992, The Folk Theorem for Discounted Repeated Games: A NEU Condition, mimeo, Princeton University.
- Abreu, D. and H. Matsushima, 1992, Virtual Implementation in Iteratively Undominated Strategies: Complete Information, *Econometrica*, 60, 993–1008.
- Abreu, P. Milgrom and D. Pearce, 1991, Information and Timing in Repeated Partnerships, *Econometrica*, 59.
- Abreu, D., D. Pearce and E. Stachetti, 1986, Optimal Cartel monitoring with Imperfect Information, *Journal of Economic Theory*, 39, 251–269.
- Abreu, D., D. Pearce and E. Stachetti, 1990, Towards a theory of Discounted Repeated Games with Imperfect Monitoring, *Econometrica* 58, 1041–1063.
- Alchian, A. and H. Demsetz, 1972, Production, Information Costs and Economic Organizations, *American Economic Review*, 62, 777–795.
- Allen, F., 1985, Repeated Principal Agent Relationships with Lending and Borrowing, *Economic Letters*, 17, 27–31.
- Arnott, R. and J. Stiglitz, 1986, Labor Turnover, Wage Structures and Moral Hazard: The Inefficiency of Competitive Markets, *Journal of Labor Economics*, 3, 434–462.
- Arrow, K. J., 1963, Uncertainty and the Welfare Economics of Medical Care, *American Economic Review*, 53, 941–973.
- Arrow, K. J., 1965, Aspects of the Theory of Risk-Bearing, Yrjo Johansonin Saatio, Helsinki, Lecture 3.
- Baron, D. and D. Besanko, 1987, Monitoring, Moral Hazard, Asymmetric Information and Risk-Sharing in Procurement Contracting, *Rand Journal of Economics*, 18, 509–532.
- Braverman, A. and J. Stiglitz, Sharecropping and the Interlinking of Agrarian Markets, *American Economic Review*, 72, 695–715.
- Dutta, P. K. and R. Radner, 1992, Optimal Principal Agent Contracts for a Class of Incentive Schemes: A Characterization and the Rate of Approach to Efficiency, *Economic Theory*, forthcoming.
- Easwaran, M. and A. Kotwal, 1985, A Theory of Two-Tier Labor Markets in Agrarian Economies, *American Economic Review*, 75, 162–177.
- Fellingham, J., D. Newman and Y. Suh, 1985, Contracts Without Memory in Multiperiod Agency Models, *Journal of Economic Theory*, 37, 340–355.
- Foster, J. and H. Wan, 1984, Involuntary Unemployment as a Principal-Agent Equilibrium, *American Economic Review*, 74, 476–484.
- Fudenberg, D., B. Holmstrom and P. Milgrom, 1990, Short-Term Contracts and Long-Term Agency Relationships, *Journal of Economic Theory*, 51, 1–31.
- Fudenberg, D., D. Levine and E. Maskin, 1989, The Folk Theorem with Imperfect Public Information, mimeo, MIT.
- Fudenberg, D. and E. Maskin, 1986, The Folk Theorem in Repeated Games with Discounting or with Incomplete Information, *Econometrica*, 52, 975–994.
- Fudenberg, D. and J. Tirole, 1990, Moral Hazard and Renegotiation in Agency Contracts, *Econometrica*, 58, 1279–1301.
- Green, E. and R. Porter, 1984, Noncooperative Collusion Under Imperfect Price Information, *Econometrica*, 52, 87–100.
- Green, J. and N. Stokey, 1983, A Comparison of Tournaments and Contracts, *Journal of Political Economy*, 91, 349–364.
- Greenwald, B. and J. Stiglitz, 1986, Externalities in Economies with Imperfect Information and Incomplete Markets, *Quarterly Journal of Economics*, 101, 229–264.
- Grossman, S. J. and O. D. Hart, 1982, Corporate Financial Structure and

- Managerial Incentives, in J. McCall ed. "The Economics of Information and Uncertainty", University of Chicago Press.
- Grossman, S. J. and O. D. Hart, 1983, An Analysis of the Principal-Agent Problem, *Econometrica*, 51, 7-45.
- Groves, T. and J. Ledyard, 1987, Incentive Compatibility since 1972, in T. Groves, R. Radner and S. Reiter eds. "Information, Incentives and Economic Mechanisms", University of Minnesota Press.
- Groves, T., R. Radner and S. Reiter "Information, Incentives and Economic Mechanisms", University of Minnesota Press.
- Harris, M. and A. Raviv, 1979, Optimal Incentive Contracts with Imperfect Information, *Journal of Economic Theory*, 20, 231-259.
- Hart, O. D. and B. Holmstrom, 1986, The Theory of Contracts, in T. Bewley ed. "Advances in Economic Theory", Cambridge University Press.
- Holmstrom, B., 1979, Moral Hazard and Observability, *Bell Journal of Economics*, 10, 74-91.
- Holmstrom, B., 1982, Moral Hazard in Teams, *Bell Journal of Economics*, 13, 324-340.
- Holmstrom, B. and P. Milgrom, 1987, Aggregation and Linearity in the Provision of Intertemporal Incentives, *Econometrica*, 55, 303-329.
- Holmstrom, B. and J. Ricart i Costa, 1986, Managerial Incentives and Capital Management, *Quarterly Journal of Economics*, 101, 835-860.
- Jewitt, I., 1988, Justifying the First-Order Approach to Principal-Agent Problems, *Econometrica*, 56, 1177-1190.
- Lambert, R., 1983, Long-Term Contracts and Moral Hazard, *Bell Journal of Economics*, 14, 441-452.
- Lambert, R., 1986, Executive Effort and the Selection of Risky Projects, *Rand Journal of Economics*, 16, 77-88.
- Lazear, E. and S. Rosen, 1981, Rank Order Tournaments as Optimal Labor Contracts, *Journal of Political Economy*, 89, 841-864.
- Legros, P., 1989, Efficiency and Stability in Partnerships, Ph.D Dissertation, California Institute of Technology.
- Legros, P. and H. Matsushima, 1991, Efficiency in Partnerships, *Journal of Economic Theory*, 55, 296-322.
- Malcomson, J. and F. Spinnewyn, 1988, The Multiperiod Principal-Agent Problem, *Review of Economic Studies*, 55, 391-408.
- Matsushima, H., 1989a, Efficiency in Repeated Games with Imperfect Monitoring, *Journal of Economic Theory*, 48, 428-442.
- Matsushima, H., 1989b, Necessary and Sufficient Condition for the Existence of Penalty rules with Budget Balancing, mimeo, Institute of Socio-Economic Planning, University of Tsukuba.
- McAfee, R. and J. Mcmillan, 1986, Bidding for Contracts: A Principal-Agent Analysis, *Rand Journal of Economics*, 17, 326-338.
- Melumad, N. and S. Reichelstein, 1989, Value of Communication in Agencies, *Journal of Economic Theory*, 47, 334-368.
- Milgrom, P., 1981, Good News and Bad News: Representation Theorems and Applications, *Bell Journal of Economics*, 12, 380-391.
- Mirrlees, J., 1974, Note on Welfare Economics, Information and Uncertainty, in M. Balch, D. McFadden and S. Wu eds. "Essays on Economic Behavior Under Uncertainty, North-Holland.
- Mirrlees, J., 1975, The Theory of Moral Hazard and Unobservable Behavior, mimeo, Nuffield College, Oxford University.
- Mirrlees, J., 1976, The Optimal Structure of Incentives and Authority Within an Organization, *Bell Journal of Economics*, 7, 5-31.
- Mookherjee, D., 1984, Optimal Incentive Schemes with Many Agents, *Review of*

- Stiglitz, J., 1983, Risk, Incentives and the Pure Theory of Moral Hazard, **The Geneva Papers on Risk and Insurance**, 8, 4-33.
- Williams, S. and R. Radner, 1989, Efficiency in Partnerships when the Joint Output is Uncertain, mimeo, Northwestern University.
- Wilson, R., 1969, The Theory of Syndicates, **Econometrica**, 36, 119-132.
- Yaari, M., 1976, A Law of Large Numbers in the Theory of Consumption Choice Under Uncertainty, **Journal of Economic Theory**, 12, 202-217.

- Economic Studies**, 51, 433–446.
- Nalebuff, B. and J. Stiglitz, 1983, Prizes and Incentives: Towards a General Theory of Compensation and Competition, **Bell Journal of Economics**, 13, 21–43.
- Porter, R., 1983, Optimal Cartel Trigger Price Strategies, **Journal of Economic Theory**, 29, 313–338.
- Radner, R., 1981, Monitoring Cooperative Agreements in a Repeated Principal–Agent Relationship, **Econometrica**, 49, 1127–1148.
- Radner, R., 1985, Repeated Principal–Agent Games with Discounting, **Econometrica**, 53, 1173–1198.
- Radner, R., 1986a, Repeated Partnership Games with Imperfect Monitoring and No Discounting, **Review of Economic Studies**, 53, 43–57.
- Radner, R., 1986b, Repeated Moral Hazard with Low Discount Rates, pp 25–64 in W. P. Heller, R. Starr and D. Starett eds., "Uncertainty, Information and Communication: Essays in Honor of Kenneth Arrow", Vol. 3, Cambridge University Press.
- Radner, R., 1987, Decentralization and Incentives, in T. Groves, R. Radner and S. Reiter eds. "Information, Incentives and Economic Mechanisms", University of Minnesota Press.
- Radner, R., R. 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.
- Radner, R. and A. Rustichini, 1989, The Design and Sharing of Performance Rules for a Partnership in Continuous Time, mimeo, AT&T Bell Laboratories, Murray Hill, NJ.
- Rasmussen, E., 1987, Moral Hazard in Risk–Averse Teams, **Rand Journal of Economics**, 18, 428–435.
- Rogerson, W., 1985a, Repeated Moral Hazard, **Econometrica**, 53, 69–76.
- Rogerson, W., 1985b, The First–Order Approach to Principal–Agent Problems, **Econometrica**, 53, 1357–1368.
- Ross, S., 1973, The Economic Theory of Agency: The Principal's Problem, **American Economic Review**, 63, 134–139.
- Rubinstein, A., 1979, Offenses That May Have Been Committed by Accident – an Optimal Policy of Retribution, in S. Brams, A. Schotter and G. Schrodiauer eds., "Applied Game Theory", Physica–Verlag, Wurzburg.
- Rubinstein, A. and M Yaari, 1983, Repeated Insurance Contracts and Moral Hazard, **Journal of Economic Theory**, 30, 74–97.
- Shapiro, C. and J. Stiglitz, 1984, Equilibrium Unemployment as a Worker Discipline Device, **American Economic Review**, 74, 433–444.
- Shavell, S., 1979, Risk Sharing and Incentives in the Principal and Agent Relationship, **Bell Journal of Economics**, 10, 55–73.
- Simon, H., 1953, A Formal Theory of the Employment Relationship, **Econometrica**, 19, 293–305.
- Singh, N., 1989, Theories of Sharecropping, in P. Bardhan ed. "The Economic Theory of Agrarian Institutions", Oxford University Press.
- Smith, A., 1937, "An Inquiry into the Nature and Causes of the Wealth of Nations", ed. by E. Cannan, Modern Library, NY.
- Spear, S. and S. Srivastava, 1988, On Repeated Moral Hazard with Discounting, **Review of Economic Studies**, 55.
- Spence, A. M. and R. Zeckhauser, 1971, Insurance, Information and Individual Action, **American Economic Review**, 61, 380–387.
- Spitzer, F., 1976, "Principles of Random Walk, (2nd.ed.) Springer–Verlag, New York.
- Stiglitz, J., 1974, Incentives and Risk–Sharing in Sharecropping, **Review of Economic Studies**, 41, 219–255.