

EFFICIENCY AND OBSERVABILITY WITH LONG-RUN AND SHORT-RUN PLAYERS*

Drew Fudenberg

and

David K. Levine**

May 1989

Revised: September 1991

*We are grateful to Eddie Dekel-Tabak, David M. Kreps and Eric Maskin. NSF Grants SES 87-088016 and SES 86-09697 provided financial support.

** Departments of Economics, MIT and UCLA, respectively.

Abstract

We present a general algorithm for computing the limit, as $\delta \rightarrow 1$, of the set of payoffs of perfect public equilibria of repeated games with long-run and short-run players, allowing for the possibility that the players' actions are not observable by their opponents. We illustrate the algorithm with two economic examples. In a simple partnership we show how to compute the equilibrium payoffs when the folk theorem fails. In an investment game, we show that two competing capitalists subject to moral hazard may both become worse off if their firms are merged and they split the profits from the merger. Finally, we show that with short-run players each long-run player's highest equilibrium payoff is generally greater when their realized actions are observed.

1. Introduction

The folk theorem for repeated games shows that any feasible, individually rational payoff can be supported as a perfect equilibrium if the players are sufficiently patient. Fudenberg, Levine and Maskin [1990] show that the same conclusion obtains in repeated games with imperfectly observed actions, provided that the information players receive allows them to statistically distinguish between the actions of any pair of opponents. However, the folk theorem does not obtain in repeated games where some players are "short run." For example, with a single firm facing a sequence of short-run consumers, no subgame-perfect equilibrium need yield a payoff as high as could be obtained if the firm could precommit to a mixed strategy. The set of equilibrium payoffs in the case of a single long-run player facing short-run opponents with perfectly observable pure actions was characterized by Fudenberg, Kreps and Maskin [1990].

This paper considers games with several long-run and short-run players, and imperfectly observed actions. (Perfect observability is a special case.) We give an exact characterization of the limit of the payoffs of "perfect public" equilibria, where strategies depend only on public information. This characterization does not require assumptions on the statistical quality of the players' information, and yields the theorems of Fudenberg, Levine and Maskin [1991], and Fudenberg, Kreps and Maskin [1990] as special cases, as well as an analog of the theorem of Matsushima [1989] in the finite-action case. Our methods build on those developed on these previous papers, and also on those of Abreu, Pearce, and Stachett [1990].

If there are no short-run players, and there is enough information to statistically identify the actions played, the limit set of equilibrium payoffs is independent of whether or not players' actions are observable: In either case, it is the full socially feasible individually rational set. This is no longer true once there are short-run players in the game. Even if there is enough information to statistically identify the players' actions, the long-run players' highest equilibrium payoff in the limit as their discount factors approach one can be strictly lower when actions are imperfectly observed than when they are not. Thus, the introduction of imperfect information into a game with short-run players has important economic consequences, even in the limit as the

discount factor of the long-run players approaches one.

We provide two economic applications of the techniques we develop. In addition, we give more general results for two classes of games: games with a product structure and moral hazard mixing games. In a game with a product structure there is an independent signal of each long-run player's action. In such game we show that the limit set of payoffs is the subset of socially feasible individually rational sets in which each long-run player's payoff is bounded above by a number that may easily be computed. The investment example discussed below is an example of a game with a product structure.

Moral hazard mixing games are defined as games with a product structure in which every possible outcome has positive probability under every action profile, and in which the incentive constraints due to moral hazard bind in a sense we make precise. We show that in such game long-run players cannot do as well even in the limit as $\delta \rightarrow 1$ as they could if their actions were observable.

Our first economic application is to the type of partnership game studied by Radner, Myerson and Maskin [1986]. They showed that the limit set of payoffs in this type of game can be strictly less than that obtained if players have time-average preferences, but they did not characterize the limit set of equilibrium payoffs. Our methods permit a simple characterization.

Our second, more extensive, application examines the economic consequences of competition between different long-lived borrowers. In this example there are infinitely-lived capitalists, short-lived investors, and economies of scale in intertemporal production. The short-run investors do not live long enough to make credible promises to repay loans, so the capitalists do all the production. However, the behavior of capitalists is imperfectly monitored by the investors. Each period the capitalist consumes some of the output and distributes the remainder to the investors as dividends. The investors observe the dividend payment but do not observe the realized output, so that the capitalist's action is subject to moral hazard whenever the distribution of outcomes is stochastic.

If there is a single capitalist and output is deterministic, so there is no moral hazard, then a

first-best payoff is attainable when the capitalist's discount factor is close to 1. When output is stochastic, though, the first best cannot be approximated for any discount factor, and the greater the moral hazard, the greater the loss. Our main conclusion is that the loss vanishes if there are enough capitalists. In particular, competition between capitalists can make them better off, in the sense that a merger between them can make them both worse off.

This is true despite the fact that neither the monitoring technology nor the utility possibilities are changed by adding a capitalist. Rather, adding more capitalists promotes efficiency because the investors can "punish" one capitalist by switching their lending to another capitalist, rather than withdrawing from the market. This is important when there is moral hazard, because investors cannot distinguish between an unlucky capitalist with low output, and a greedy one. Consequently, to induce the capitalist to pay dividends when output is high, the capitalist must be punished for failing to make dividend payments when output is low. This means that even an honest capitalist faces positive probability of punishment. With a single capitalist, punishment is socially inefficient, while with several capitalists the strategy of investing in another firm when dividends are not paid provides incentives without this social cost.

The effect of competition here is different than in the papers of Ausabel and Denecker [1989] and Gul [1986], who study sales of a durable good. In their models, a single monopolist may be unable to credibly keep prices high, while with two sellers there is an equilibrium where each seller keeps price high for fear of provoking a price war. Consequently, adding a seller reduces social efficiency by pricing willing buyers out of the market, but increases the greatest possible profit that either or both sellers can obtain. In our example, adding a capitalist (or subdividing a capitalist into two competing firms) increases social efficiency by increasing the probability of full employment and increases their maximum combined total profit, but has no effect on the maximum profit that can be attained by a single capitalist. From an economic point of view, our example is most closely related to those of Atkeson [1987] and Marimon [1988], who study issues of wealth accumulation and borrowing in a game with two long-run players and moral hazard.

2. The Model

In the stage game, each player $i = 1$ to n simultaneously chooses a (pure) action a_i from a finite set A_i with m_i elements. Each action profile $a \in A \equiv \times_{i=1}^n A_i$ induces a probability distribution over publicly observed outcomes y and privately observed outcomes $z = (z_1, \dots, z_n)$. For a given action profile a let $\pi_{yz}(a)$ denote the probability of (y, z) and let $\pi_y(a)$ denote the marginal probability for the public outcome y . The public outcomes lie in a finite set Y with m elements. Each player i 's realized payoff $r_i(z_i, y)$ depends only on the realized outcome: actions matter only in their influence on the distribution over outcomes. This model includes as a special case games where $y = a$, so the public information perfectly reveals the actions chosen. It also includes games where y conveys only imperfect information about the actions. For example, y can be the realized market price and the $z_i = a_i$ can be choices of output levels as in Green-Porter [1984], or y can be the realized quantity of a good and $z_i = a_i$ the care with which the good is manufactured. Other examples with a mix of public and private information are analyzed in detail below. The key assumption is that there is some outcome y that is publicly observed. This rules out games where the players receive only private signals, a case is considered in Lehrer [1988] and Fudenberg and Levine [1991a].

Player i 's expected payoff to an action profile a is

$$g_i(a) = \sum_{\substack{y \in Y \\ z \in Z}} \pi_{yz}(a) r_i(z_i, y),$$

which gives the normal form of the game.

We will also consider mixed actions α_i for each player i . For each profile $\alpha = (\alpha_1, \dots, \alpha_n)$ of mixed actions, we can compute the induced distribution over outcomes,

$$\pi_{yz}(\alpha) = \sum_{a \in A} \pi_{yz}(a) \alpha(a), \quad \pi_y(\alpha) = \sum_{a \in A} \pi_y(a) \alpha(a)$$

and the expected payoffs

$$g_i(\alpha) = \sum_{\substack{y \in Y \\ z \in Z}} \sum_{a \in A} \pi_{yz}(a) \alpha(a) r_i(z_i, y). \quad 3$$

We denote the profile where player i plays α_i and all other players follow profile α by (α_i, α_{-i}) ; $\Pi_i(\alpha_i, \alpha_{-i})$ and $g_i(\alpha_i, \alpha_{-i})$ are defined in the obvious way.

As in Fudenberg, Kreps and Maskin [1990], we label players so that $i \in LR = \{1, 2, \dots, L\}$, $L \leq n$, are long-run players, whose objective is to maximize the average discounted value of per-period payoffs using the common discount factor δ . If $\{g_i(t)\}$ is a sequence of payoffs for long-run player i , player i 's average discounted payoff is

$$(1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} g_i(t) \quad 4$$

where we normalize by $(1-\delta)$ to measure the repeated game payoffs in the same units as payoffs in the stage game.

The remaining players $j \in SR = \{L+1, \dots, n\}$ represent different types of short-run players, each representative of which plays only once. One example of a model with long and short-run players is Selten's [1977] chain-store game, where a single long-run incumbent faces a sequence of short-run opponents. A more extensive example is given below.

In the repeated game, in each period $t = 1, 2, \dots$, the stage game is played, and the corresponding public outcome is then revealed. The public history at the end of period t is $h(t) = (y(1), \dots, y(t))$. The private history for long-run player i at the end of period t is $h_i(t) = (a_i(1), z_i(1), \dots, a_i(t), z_i(t))$. We also let $h(0)$ and $h_i(0)$ denote the null public and private histories in which nothing has happened. A strategy for long-run player i is a sequence of maps mapping public and private histories $(h(t-1), h_i(t-1))$ to mixed actions. A strategy for the period- t players of type j , $j \in SR$, is a map from the public information $h(t-1)$ to mixed actions. Note that the short-run player j observes the public information at t , but not the private history corresponding to the play of his predecessors. However, for the class of equilibria we will consider, this assumption does not matter.

Each strategy profile generates probability distributions over histories in the obvious way,

and thus also generates a distribution over histories of the players' per-period payoffs.

Let A_i denote the space of mixed actions, and

$$B: A_1 \times \dots \times A_L \rightarrow A_{L+1} \times \dots \times A_n$$

be the correspondence that maps any mixed action profile $(\alpha_1, \dots, \alpha_L)$ for the long run players to the corresponding static equilibria for the short run players. That is, for each $\alpha \in \text{graph}(B)$, and each $j > L$, α_j maximizes $g_j(\alpha, \alpha_{-j})$.

Our focus is on a special class of the Nash equilibria that we call perfect public equilibria. A strategy for long-run player i is public if at each time t , it depends only on the public information $h(t-1)$ and not on the private information $h_i(t-1)$. A perfect public equilibrium is a profile of public strategies such that at every date t and for every history $h(t-1)$ the strategies are a Nash equilibrium from that date on. Note that in a public equilibrium the players' beliefs about each others' past play are irrelevant: No matter how player i plays, all nodes in the same information set for i at the beginning of period t lead to the same probability distribution over other player's current and future play. One consequence of this is that the assumption that short-run players do not observe the private information of earlier short-run players is irrelevant. The other is that when we define "perfection", it made sense to speak of "Nash equilibrium" as though each period began a new subgame.

It is easy to show that perfect public equilibria are sequential, and that the set of public equilibrium payoffs is stationary, that is, the set of perfect public equilibrium payoffs in any period t and for any public history $h(t-1)$ is independent of t and $h(t-1)$. However, when the players' actions are not perfectly observed, there may be sequential equilibrium payoffs that are not obtainable with public strategies. (With observed actions all equilibria are public.)

In any period t of a public equilibria, the players' period- t mixed actions $\alpha(h(t-1))$ are common knowledge. In particular, each short-run player j believes he will face the mixed action $\alpha_{-j}(h(t-1))$. Since the short-run players care only about their one-period payoffs, in a public equilibrium each period's mixed action must lie in the graph of B . (This need not be true of the

sequential equilibria that are not public, for subsets of the players can use their knowledge of their own past actions to correlate their play. See Fudenberg, Levine and Maskin [1990] for an example of the way this correlation can generate additional equilibrium payoffs.)

We define $E(\bar{\delta}) \subset \mathbb{R}^L$ to be the set of average present values for the long-run player that can arise in perfect public equilibria when the discount factor is $\bar{\delta}$. We define V to be the subset of \mathbb{R}^L generated as convex combinations of payoffs to long-run players $(g_1(\alpha), \dots, g_L(\alpha))$, where $\alpha \in \text{graph}(B)$. Note that this set is closed (since B is), convex, and of course $E(\bar{\delta}) \subseteq V$.

3. Enforceability on Halfspaces and the Set of Limit Equilibria

This section describes an algorithm that computes the limiting value of the set $E(\bar{\delta})$ as $\bar{\delta}$ tends to one. We consider various applications below.

Definition 3.1: Let $\bar{\delta}, W \subseteq \mathbb{R}^L$ and $v \in \mathbb{R}^L$ be given. Action α is enforceable with respect to W and $\bar{\delta}$ if there exists a map $w: Y \rightarrow W$ such that $\alpha \in \text{graph}(B)$ and for all $i \in \text{LR}$ and $a_i \in A_i$,

$$v_i = (1-\bar{\delta})g_i(a_i, \alpha_{-i}) + \bar{\delta} \sum_y \pi_y(a_i, \alpha_{-i}) w_i(y) \text{ for } a_i \text{ s.t. } \alpha_i(a_i) > 0$$

(3.1)

$$v_i \geq (1-\bar{\delta})g_i(a_i, \alpha_{-i}) + \bar{\delta} \sum_y \pi_y(a_i, \alpha_{-i}) w_i(y) \text{ for } a_i \text{ s.t. } \alpha_i(a_i) = 0.$$

We also say that w enforces (α, v) with respect to W and $\bar{\delta}$. If for a given α , a v exists satisfying (3.1), we say that α is enforceable with respect to W and $\bar{\delta}$. Note that each $w(y)$ specifies continuation payoffs only for the long-run players. If such $w(y)$'s exist, we say that v is generated by W and $\bar{\delta}$.

Equation 3.1 says that if all players $j \neq i$ play α_{-i} then (i) no short-run player i can increase his payoff by deviating from α_i , and (ii) if long-run player i 's discounted average payoff

starting tomorrow on when y occurs today is given by $w_i(y)$, then player i receives exactly average payoff v from any of the actions in the support of α_i , and no choice of action yields a higher average payoff.

The key to the algorithm for finding the limit of $E(\delta)$ is the study of the payoffs that can be generated using continuation payoffs that lie in half-spaces H of \mathbb{R}^L : these are sets of the form $H(\lambda, k) = \{v \mid \lambda \bullet v \leq k\}$ for $\lambda \in \mathbb{R}^L$ and $k \in \mathbb{R}$.

Definition 3.2: The maximal score attainable by action α in direction λ with discount factor δ , denoted $k^*(\alpha, \lambda, \delta)$, is the maximum of $\lambda \bullet v$ such that there exists $v \in \mathbb{R}^L$ and $k \in \mathbb{R}$ with (α, v) enforceable with respect to δ and $H(\lambda, k)$ and such that $k = \lambda \bullet v$. The halfspace associated with this maximum is $H^*(\alpha, \lambda, \delta)$; its boundary is the hyperplane $h^*(\alpha, \lambda, \delta)$ given by $\lambda \bullet v = k^*(\alpha, \lambda, \delta)$.

This definition asks us to fix a mixed action α and a direction λ , and then find the "highest" halfspace in direction λ such that a point on the boundary of the hyperplane can just be generated with action α and continuation payoffs in the halfspace. Since a halfspace is defined by a system of linear inequalities, the maximal score k^* may also be expressed as the solution to the linear programming problem:

$$\begin{aligned} & \max_{v, w} \lambda \bullet v \text{ subject to} \\ & \text{(a) } v_i = (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum \Pi_y(a_i, \alpha_{-i}) w_i(y) \text{ for } a_i \text{ s.t. } \alpha_i(a_i) > 0 \\ & \text{(3.2)} \\ & \text{(b) } v_i \geq (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum \Pi_y(a_i, \alpha_{-i}) w_i(y) \text{ for } a_i \text{ s.t. } \alpha_i(a_i) = 0 \end{aligned}$$

for $i = 1, \dots, L$

$$\text{(c) } \lambda \bullet v \geq \lambda \bullet w(y) \text{ for } y \in Y.$$

Lemma 3.1: (i) $k^*(\alpha, \lambda, \delta) = k^*(\alpha, \lambda)$ independent of δ .

(ii) $k^*(\alpha, \lambda) \leq \lambda \bullet g(\alpha)$;

(iii) $k^*(\alpha, \lambda) = \lambda \bullet g(\alpha)$ if and only if $g(\alpha)$ is enforceable with respect to the hyperplane orthogonal to λ at $g(\alpha)$.

Proof: (i) Suppose w' and v are feasible in (3.2) for $(\alpha, \lambda, \bar{\delta}')$. Define

$$w''(y) = [(\bar{\delta}'' - \bar{\delta}')/\bar{\delta}''(1 - \bar{\delta}')]v + [\bar{\delta}'(1 - \bar{\delta}'')/\bar{\delta}''(1 - \bar{\delta}')]w'(y).$$

It may then be checked that w'' and v are feasible in (3.2) for $(\alpha, \lambda, \bar{\delta}'')$. This implies that the set of feasible v in (3.2) is independent of $\bar{\delta}$, so the maximum is as well.

(ii) $k^*(\alpha, \lambda) > \lambda \bullet g(\alpha)$ would imply that $g(\alpha)$ is in the interior of the maximal halfspace.

Since each $v^* \in H^*$ is a strict convex combination of $g(\alpha)$ and points in H^* , this is impossible.

(iii) $k^*(\alpha, \lambda) = g(\alpha)$ requires that $(\alpha, g(\alpha))$ be enforceable with continuation payoffs on the hyperplane orthogonal to λ and passing through $g(\alpha)$. ■

We should point out that $k^*(\alpha, \lambda)$ is not necessarily upper semi-continuous in α . Although the constraints in the definition of enforceability (3.1) have the closed graph property, half-spaces are not compact, so the set of payoffs enforceable with continuations on a given half-space is not generally upper hemi-continuous in α .

Definition 3.3: The maximal score in direction λ , $k^*(\lambda)$, solves

$$k^*(\lambda) = \sup_{\alpha \in \text{graph}(B)} k^*(\alpha, \lambda). \quad \text{The maximal halfspace in direction } \lambda \text{ is } H^*(\lambda)$$

$$= H(\lambda, k^*(\lambda)).$$

The following lemma is sometimes used in characterizing the equilibrium set. Let e_i be the vector whose i^{th} component equals 1, and whose other components are all zero.

Lemma 3.2: The maxima score $k^*(-e_j)$ in the direction $-e_j$ that corresponds to minimizing player j 's payoff is no greater than

$$-\min_{\alpha \in \text{graph } B} \max_{\alpha_j \in A_j} g_j(\alpha_j, \alpha_j).$$

Proof: For $\lambda = -e_j$, constraint 3.2(c) simplifies to $w_j(y) \geq v_j$. Constraint 3.2(a) and (b) for player j

then imply that $v_j \geq g_j(a_j, \alpha_j)$ for all a_j , and here in particular that $v_j \geq \max_{a_j} g_j(a_j, \alpha_j)$. Hence for

$$\text{any } \alpha, k^*(-e_j, \alpha) \text{ is no greater than } -\min_{\alpha \in \text{graph } B} \max_{\alpha_j \in A_j} g_j(\alpha, \alpha_j).$$

Definition 3.4: $Q = \bigcap_{\lambda} H^*(\lambda)$.

Note that Q is convex, and that each point q on the boundary of Q is enforceable on all of the halfspaces that are tangent to Q at q .

Theorem 3.1: (i) For all δ , $E(\delta) \subseteq Q$.

(ii) If the dimension of $Q \subset \mathbb{R}^L$ is L , then $\lim_{\delta \rightarrow 1} E(\delta) = Q$.

Proof: See the Appendix.

4. A Partnership Game

Our first application of Theorem 3.1 will be to a game involving only long-run players. If a deviation by each player can be statistically distinguished from that of any other player, it is an implication of Theorem 3.1 that the folk theorem holds, and indeed this has been shown by Fudenberg, Levine and Maskin [1991]. This leads us to study a partnership example in which deviations

by different players cannot be distinguished, and one closely related to the example of Radner, Myerson and Maskin [1986] in which the folk theorem is known to fail with discounting.

However, while Radner, Myerson and Maskin showed that the folk theorem failed, they did not show what the actual set of limit equilibrium payoffs is, which our algorithm enables us to do.

The example we study is a two-player partnership game. Each identical partner must choose one of two effort levels $a_i = 0$ or 1 . Output per person is either 0 or 1 . The probability of the high output level is

$(a_1 + a_2)$; the additivity is important as we discuss below. Effort is unobserved, so each player's private signal z_i is identified with his action; per capita output is publicly observed. Effort has a disutility q , so payoffs are

$$r_i(a_i, y) = \begin{cases} y & \text{if } a_i = 0 \\ y - q & \text{if } a_i = 1. \end{cases}$$

The strategic form of this game is given in Figure 4.1.

		partner 2	
	effort	1	0
partner 1	1	[2p-q, 2p-q	p-q, p]
	0	[p, p-q	0, 0]

FIGURE 4.1

We assume $p \leq q$, so that the effort level 0 is the unique dominant strategy Nash equilibrium in the stage game. We also assume $q \leq 2p$, so that both partners providing effort maximizes the sum of the player's payoffs.

At the profile where both partners work, the probability that output is high is $2p$. If either partner deviates, the probability that output is high is p , so it is impossible to determine which

partner deviated even if the probability of y is observed. In Fudenberg, Levine and Maskin [1990], we showed that the ability to make such determinations is one of several conditions that are sufficient for the folk theorem to hold. In this example deviations cannot be distinguished in part because the game is symmetric and we are considering a profile in which both players behave in the same way. However, in some symmetric games it is possible to endogenously introduce an asymmetry by having the players mix in different ways.¹

This game has the special feature that the probability of high output is linear in the players' actions, so that changing a_i from 0 to 1 increases the probability that $y = 1$ by p , regardless of the play of player $-i$. Thus, whether player i prefers $a_i = 1$ or $a_i = 0$ depends only on the specification of the continuation payoffs $w_i(\bullet)$, and is independent of a_{-i} . In particular, player i prefers $a_i = 1$ if the reduction of $(1-\delta)(p-q)$ in his immediate payoff is outweighed by the change of $\delta p[w_i(1)-w_i(0)]$ in his expected continuation payoff. To simplify future notation, we define $\Delta = (q-p)/p$, and note that player i prefers $a_i = 1$ or $a_i = 0$ as $w_i(1)-w_i(0) >$ or $<$ $\Delta(1-\delta)/\delta$.

Moreover, and this is a key simplification, the linear structure implies that in determining the maximal scores $k^*(\lambda)$ in each direction λ , we need only consider the four pure action profiles, as shown by Lemma 4.1.

Say that continuation payoffs $w_i:Y \rightarrow R$ enforce profile α for player i if they satisfy the incentive constraint (3.1) for player i .

Lemma 4.1: In a game with only long-run players, suppose that for each player i , for all a_i , and all α_{-i} , whenever continuation payoffs $w_i(\bullet)$ enforce (a_i, α_{-i}) for player i they also enforce (a_i, α_i) for player i

¹Indeed if there are three outcomes and we fix the strategic form payoffs, Fudenberg, Levine and Maskin show that almost all perturbations of the matrix of outcome probabilities result in the folk theorem. Here, however, the additive structure of probabilities are a non-generic case in which mixing does not introduce an asymmetry. (It should be noted that perturbing the marginal probabilities of an additive structure will generally change the strategic form payoffs.)

for all α_i . Then to determine the maximal scores $k^*(\lambda) = \sup_{\alpha} v(\alpha, \lambda)$, it suffices to consider only

the pure action profiles.

Proof: If v and w solve the LP problem 3.2 for a given α and λ , then for each $a_i \in \text{support}(\alpha_i)$, (a_i, α_{-i}) is enforced by w . On the hypothesis of the lemma, this implies that each $a_i \in \text{support}(\alpha_{-i})$ is enforced by w for any α_{-i} , and so any profile $a \in \text{support}(\alpha)$ is enforced by w . Define $v(a) = (1-\delta)v(a) + \delta \sum_y \pi_y(a) w_i(y)$, and let a be an element of $\arg\max_{a \in \text{support}(\alpha)} v(a)$. Since $v(a) = v(a) \leq v(a)$, a gives at least as high a score, so we need only show that (v, w) is a solution of problem 3.2 for profile a . The first equality is satisfied from the definition of v , the incentive constraints are satisfied as we observed above, and the inequalities $v(a) \geq v(a)$ follow from $v(a) \geq v(a) \geq v(a)$. ■

To determine the set of limit payoffs, we will compute the maximal scores $k^*(\lambda)$, the maximal half-spaces $H^*(\lambda)$, and the set $Q = \bigcap_{\lambda} H^*(\lambda)$. Since each player's minmax value is zero, Lemma 3.2

implies that the maximal feasible score in directions $-e_1$ and $-e_2$ is zero, so Q contains at most the nonnegative feasible payoffs. Moreover, since the minmax profile $(0,0)$ is a static equilibrium, the score of 0 can be obtained on any half space, so the maximal half spaces in directions λ with both λ_1 and λ_2 negative contain the point $(0,0)$ and hence imply no further restrictions on Q .

For λ with $\lambda_2 > 0 > \lambda_1$, the maximum feasible value of v is $\lambda_2 p + \lambda_1(p-q)$; since this value can be obtained as a solution to problem 3.2 (by using the profile $(1,0)$ and continuation payoffs $w_1(0) = p-q$, $w_1(1) = p-q + \Delta(1-\delta)/\delta$, $w_2(0) = p$, and $w_2(1) = p - (\lambda_2/\lambda_1)\Delta(1-\delta)/\delta$) it is the maximum score $k^*(\lambda)$ in these directions. A symmetric argument shows that $k^*(\lambda) = \lambda_1 p + \lambda_2(p-q)$ in directions with $\lambda_1 > 0 > \lambda_2$.

The only remaining directions are those λ with both λ_1 and λ_2 positive. Using Lemma

4.1, to find the highest score in the directions it suffices to consider the pure action profiles; since (0,0) yields payoff of only 0, and (0,1) and (1,0) are symmetric, it will suffice to consider the profiles (1,1) and (1,0). Using our earlier observation that player i prefers $a_i = 1$ or $a_i = 0$ as $w_i(1) - w_i(0) > \text{or}$ $< \Delta(1-\delta)/\delta$, the program 3.2 for profile (1,1) simplifies to

$$\max_{v, w} \lambda^{\circ} v \quad \text{s.t.}$$

$$v_i = (1-\delta)[2p-q] + \delta[2pw_i(1) + (1-2p)w_i(0)] \\ \} \quad i = 1, 2 \\ w_i(1) - w_i(0) \geq \Delta(1-\delta)/\delta, \quad J$$

and

$$\lambda^{\circ} v \geq \lambda^{\circ} w(1).$$

$$\lambda^{\circ} v \geq \lambda^{\circ} w(0).$$

We claim that for any positive λ_1 and λ_2 , the solution to this program is $= (\lambda_1 + \lambda_2)[(2p-q) - \Delta(1-2p)]$.

We show first that this value is obtainable: Set

$$w_1(1) = w_2(1) = 2p-q-\Delta(1-2p), \text{ and set}$$

$w_1(0) = w_2(0) = 2p-q-\Delta(1-2p) - \Delta(1-\delta)/\delta$, so that $w_i(1) - w_i(0) = \Delta(1-\delta)/\delta$. Then $v_i = (1-\delta)[2p-q] + \delta[(2p-q)-\Delta(1-2p)-(1-2p)\Delta(1-\delta)/\delta] = 2p-q - \Delta(1-2p)$, so $= \lambda^{\circ} v = \lambda^{\circ} w(0)$, and all of the constraints are satisfied.

Next we argue that no value above is attainable. Assume to the contrary that the solution is some $k' > .$ For a fixed value of $w_1(1)$, the program is maximized by setting

$$w_2(1) = (k' - \lambda_1 w_1(1)) / \lambda_2,$$

$$w_1(0) = w_1(1) - \Delta(1-\delta)/\delta, \text{ and}$$

$$w_2(0) = w_2(1) - \Delta(1-\delta)/\delta = (k' - \lambda_1 w_1(1))/\lambda_2 - \Delta(1-\delta)/\delta.$$

Substituting for v_i yields

$$k' \leq (1-\delta)^{(\lambda_1+\lambda_2)}[2p-q] + \delta\lambda_1\{w_1(1) - (1-2p)\Delta(1-\delta)/\delta\} + \delta\lambda_2[(k' - \lambda_1 w_1(1))/\lambda_2 - (1-2p)\Delta(1-\delta)/\delta],$$

which, after collecting terms, implies

$$k' \leq (\lambda_1+\lambda_2)[2p-q-(1-2p)\Delta] = .$$

Now we determine the highest score that can be attained using profile $(1,0)$. We claim that this score is at $\lambda^\circ g(0) = \lambda_1(p-q) + \lambda_2 p$. From Lemma 3.1 (iii), this will follow if $(1,0)$ can be enforced with continuation payoffs w on the hyperplane $\lambda^\circ w = \lambda_1(p-q) + \lambda_2 p$ orthogonal to $\lambda^\circ g(0)$. One specification of the continuation payoffs that does this is

$$w_1(0) = p-q, \quad w_1(1) = p-q - \Delta(1-\delta)/\delta,$$

$w_2(0) = p, \quad w_2(1) = p - (\lambda_1/\lambda_2)\Delta(1-\delta)/\delta$. A similar argument shows that the highest score that can be obtained with profile $(0,1)$ is $\lambda_1 p + \lambda_2(p-q)$.

The last step is to compare the scores to each action profile, and determine the overall highest score in each direction. If $2p - q > 2[2p-q-(1-2p)\Delta]$, then for all λ , $\max\{\lambda^\circ(p,p-q), \lambda^\circ(q,q-p)\} > \lambda^\circ(2p-q-(1-2p)\Delta, 2p-q-(1-2p)\Delta)$, and none of the extremal equilibrium payoffs are generated by the profile $(1,1)$. This is because the cost of providing effort $((1-2p)\Delta)$ is so large that when incentive costs are considered it is more efficient to have only one player work. This is the case depicted in Figure 4.2.

Figure 4.3 depicts the opposite case $2p - q < 2[2p-q-(1-2p)\Delta]$. Here the score $(\lambda_1+\lambda_2)(2p-q-(1-2p)\Delta)$ generated by $(1,1)$ exceeds the scores to $(1,0)$ and $(0,1)$ for intermediate values of λ_1/λ_2 .

5. Games With a Product Structure

As noted, in the investment game there is a separate public signal for each long-run player

which is independent of the signal for other long-run players, and depends only on his own play and that of the short-run players. We describe this by saying that (information in) the game has a product structure. Such games have strong properties, which we apply to the investment game in the next section.

Let $a_{\text{SR}} = (a_{L+1}, \dots, a_n)$ be the vector of short run players play. Formally, we say that a game has a product structure if $y = (y_1, \dots, y_L)$ (with probability one), and $\Pi_y(a) = \prod_{y_1} \Pi_{y_1}(a_1, a_{\text{SR}}) \dots \prod_{y_L} \Pi_{y_L}(a_L, a_{\text{SR}})$ where Π_{y_i} is the marginal distribution on y_i . In other words, each long-run player's action influences only the distribution of his "own" outcome y_i , and the y_i are statistically independent.

We now specialize to games with a product structure. A special case of these games is the case in which actions are observable but mixed strategies are not. In these games Fudenberg and Maskin [1986] show that the folk theorem holds when there are no short-run players. With a single long-run player, Fudenberg, Kreps and Maskin [1990] show that the folk theorem fails, but are nevertheless able to characterize equilibrium payoffs. Using Theorem 3.1, we can extend this latter characterization to encompass many long-run and short-run players, and moral hazard as well.

In a game with a product structure, the marginal distribution of long-run player i 's signal does not depend on the actions of other long-run players, so it may be denoted by $\Pi_{y_i}(a_i, \alpha_{\text{SR}})$. For

$i = 1, \dots, L$ we then define the matrix $\Pi_i(\alpha_{\text{SR}})$ with rows corresponding to $a_i \in A_i$ and columns to $y_i \in Y_i$, and with the (a_i, y_i) component equal to $\Pi_{y_i}(a_i, \alpha_{\text{SR}})$. Similarly, we may define $g_i(\alpha_{\text{SR}})$ to be

the vector indexed by a_i , with components equal to $g_i(a_i, \alpha_{\text{SR}})$. We say that a game with a product structure satisfies the sufficient rank condition at α if every left null vector of $\Pi_i(\alpha_{\text{SR}})$ is orthogonal to the vector $g_i(\alpha_{\text{SR}})$. This is certainly true if the rank of $\Pi_i(\alpha_{\text{SR}})$ is m_i , the number of player i 's actions, so that it has full row rank. In this case the only left null vector is 0 , which is certainly orthogonal to $g_i(\alpha_{\text{SR}})$. In this case we say that the full rank condition is satisfied.

Lemma 5.1: In a game with a product structure, if α satisfies the sufficient rank condition, then

for all players i , payoffs v_i and discount factors $\delta > 0$ there exist $w_i: Y_i \rightarrow \mathbb{R}$ such that

$$(5.1) \quad v_i = (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in \Pi_y} (\alpha_{i, \text{SR}}) w_i(y_i) \text{ for all } a_i \in A_i.$$

Remark: The condition (4.3) asserts that the incentive constraints (3.1) may be satisfied with exact equality using continuation payoffs w_i that depend only on player i 's outcome, and not those of other long-run players.

Proof: Let e be the vector that consists entirely of ones. Writing out (3.3) in matrix form and rearranging terms, we see that we must find w_i (which may be viewed as a vector with components corresponding to rows of Π_i) to solve

$$\delta \Pi_i(\alpha_{\text{SR}}) \cdot w_i = v_i \cdot e - (1-\delta) g_i(\alpha_{-i}).$$

This is possible if and only if every null vector on the left is orthogonal to the vector on the right. Since $\Pi_i(\alpha_{\text{SR}})$ is a stochastic matrix, every left null vector is orthogonal to e , and by the assumption of sufficient rank, every such null vector is orthogonal to $g_i(\alpha_{-i})$. ■

If the sufficient rank (or full rank) condition is satisfied for every $\alpha \in \text{graph}(B)$, we say that the game satisfies the sufficient rank (or full rank) condition.² Notice that these conditions generalize the idea of a game with observable actions. Observable actions means that y_i is isomorphic to a_i , so that Π_i is the identity matrix.

Let e_i be the unit vector in the direction of the i^{th} coordinate axis, and recall that $k^*(e_i)$ corresponds to maximizing player i 's payoff and $k^*(-e_i)$ to minimizing it. Recall that V is the subset of \mathbb{R}^L generated as convex combinations of payoffs to long-run players $(g_1(\alpha), \dots, g_L(\alpha))$, where $\alpha \in \text{graph}(B)$, and that this set is closed. Our goal is to prove:

Theorem 5.1: In a game with a product structure satisfying the sufficient rank condition, Q is the intersection of V with the 2L constraints

²Fudenberg, Levine and Maskin [1990] distinguish the full rank condition as the "individual full rank" condition to distinguish it from a stronger condition called "pairwise full rank".

$$k^*(-e_i) \leq v_i \leq k^*(e_i).$$

If, following Fudenberg and Maskin [1986], we assume that Q has dimension L , it follows from Theorem 3.1 that Q is the limit of $E(\delta)$ as $\delta \rightarrow 1$.

Although this gives a characterization of limit payoffs as $\delta \rightarrow 1$, it is not interpretable as a folk theorem, as in general $k^*(e_i)$ is not the best payoff for player i in the graph(B).

In proving Theorem 5.1 there are two cases: directions λ that are parallel to a coordinate axis, and those that are not. We refer to the former as coordinate directions, the latter as non-coordinate directions. From the definitions of Q and k^* , Q must satisfy the constraint $k^*(-e_i) \leq v_i \leq k^*(e_i)$. To prove Theorem 6.1 it therefore suffices to show that for non-coordinate directions λ we can obtain scores $k^*(\lambda)$ on the boundary of V .

Lemma 5.2: In a game with a product structure satisfying the sufficient rank condition, if λ is a non-coordinate direction, then $k^*(\lambda) = \lambda \bullet v$, where $v \in V$ and $V \subseteq H(\lambda, k^*(\lambda))$.

Proof: By Lemma 3.1(iii), it suffices to show that for some $v \in V$ and k' such that $h(\lambda, k')$ is tangent to V at v , that there is an α with $(g_1(\alpha), \dots, g_L(\alpha)) = v$ that may be enforced on $h(\lambda, k')$. By the definition of V we may find $\alpha \in \text{graph}(B)$ and a k' such that $(g_1(\alpha), \dots, g_L(\alpha)) = v$ and $h(\lambda, k')$ is tangent to V at v . Consequently, we need only construct w satisfying $\lambda \bullet w(y) = k'$, and such that the incentive constraints (4.1) are satisfied. Since $y = (y_1, \dots, y_L)$, we initially restrict attention to continuation payoffs that depend only on each player's own outcome, that is, $w_i(y) = w_i(y_i)$, and such that the incentive constraints all hold with exact equality. They may be written as

$$(5.2) \quad v_i = (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y_i \in Y_i} \Pi_{y_i} (a_i, \alpha_{-i}) w_i(y_i).$$

By Lemma 5.1 there exists a solution $w_i(y_i)$.

Since λ is non-coordinate, $\lambda_i \neq 0$ for at least two players, say $i = 1, 2$. Consequently, $\lambda \bullet (y) = k'$ provided that

$$\lambda_1 \hat{w}_1(y) + \lambda_2 \hat{w}_1(y) = k' - \sum_{i=3}^L \lambda_i \hat{w}_i(y_i). \quad (13)$$

Clearly, setting $w_1(y) = w_1(y_1) + w(y_{-1})$, $w_2(y) = w_2(y_2) + w(y_1)$, and $w_i(y) = w_i(y_i)$ for $i \neq 1, 2$ preserves incentive compatibility for any functions w of y_{-1} and w of y_1 . If we choose $w(y_1) = -(\lambda_1/\lambda_2)w_1(y_1)$, and $w(y_{-1}) = (1/\lambda_1)(k' - \sum_{i=2}^L \lambda_i w_i(y_i))$, the resulting $w(y)$ completes the proof. ■

It is worth noting that when the sufficient rank condition is satisfied, the linear programming problem (3.2) for vectors λ corresponding to coordinate directions simplifies considerably. Since all components of λ are zero except the i th (say), v_i is the only component of v to appear in the objective. Moreover, the only constraints involving v_j , $j \neq i$, are those from (3.1). However, we already showed that if the full rank condition is satisfied, we may choose w_j to satisfy these constraints. The remaining constraints do not involve player j at all, except that $w_i(y)$ may depend on y_j . However, if $w_i(y)$ satisfy all the constraints, so does

$$w(y) = w_i(y_i) \equiv E[w_i(y) \mid y_i, \alpha_i].$$

Consequently, the problem (4.2) is equivalent to the problem

$$\begin{aligned} \max \quad & \lambda \bullet v \text{ subject to} \\ & v_i, w_i \text{ } i \text{ } i \\ & v_i = (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in \Pi_y} (a_i, \alpha_{SR}) w_i(y_i) \text{ for } a_i \text{ s.t. } \alpha_i(a_i) > 0 \\ (5.3) \quad & i \\ & v_i \geq (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in \Pi_y} (a_i, \alpha_{SR}) w_i(y_i) \text{ for } a_i \text{ s.t. } \alpha_i(a_i) = 0 \\ & i \\ & \lambda_i \bullet v_i \geq \lambda_i \bullet w(y) \text{ for } y \in Y. \end{aligned}$$

In other words, the score in the direction of maximizing or minimizing player i 's payoff is determined as if he were the only long-lived player in the game: the other long-lived players figure into this problem only through g_i .

6. Moral Hazard Mixing Games

Now we show how to determine the constraints $k^*(e_i)$ and $k^*(-e_i)$ in a broad class of games with a product structure. Following Fudenberg, Kreps and Maskin [1990], we define two bounds on the payoffs of the long-run player $i = 1, \dots, L$. First,

$$\underline{v}_i \equiv \min_{\alpha \in \text{graph}(B)} \max_{a_i \in A_i} g_i(a_i, \alpha_{-i}) \quad 14$$

is the minmax, incorporating the constraint that short-run players must play best responses to some play of the long-run players. Lemma 3.1 shows that every long-run player's payoff is at least \underline{v}_i in every perfect public equilibrium. The second bound is

$$\bar{v} \equiv \max_{\alpha \in \text{graph}(B)} \min_{\alpha_i \in \text{support}(\alpha_i)} g_i(a_i, \alpha_{-i}), \quad 15$$

which is the most player i can get if he cannot be trusted not to maximize within the support of a mixed strategy. Let $\underline{\alpha}_i$ be the strategies for i 's opponents that lead to \underline{v}_i , and let α^{*i} be strategies for all players that solve the problem defining \bar{v} .

Fudenberg, Kreps and Maskin show that for games with observed actions and a single long-run player the limit set of perfect equilibrium payoffs is the interval $[\underline{v}_i, \bar{v}]$. Fudenberg, Levine and Maskin [1991] consider games with unobserved actions where all players are long-run. In these games all strategy profiles are in $\text{graph}(B)$, so $\bar{v} = \max_{\alpha} g_i(\alpha)$, and any strategy profile α that gives player i payoff \bar{v}_i necessarily has α^i a static best response to α . This means that player i can be induced not to deviate from a^i with continuation payoffs that are orthogonal to player i 's coordinate axis, so that $k^*(e_i) = \max_{\alpha} g_i(\alpha)$. When some players are short run, α^i need not be a static best response to α , which is why $k^*(e_i)$ can be less than $\max_{\alpha} g_i(\alpha)$.

An important class of games with a product structure that do not have perfect observability are moral hazard mixing games. These are games for which $\prod_{y_i} (a_i, a_{-i}^{\text{SR}})$ is strictly positive for all

$$a_i \in A_i, y_i \in Y_i,$$

$i = 1, \dots, L$, and $a_{-i}^{\text{SR}} \in \times_{j=L+1}^n A_j$ so there is genuine moral hazard. In addition we require that if

$\alpha \in \text{graph}(B)$ and $g_i(\alpha) = v$ then α_i is not a best response to α_{-i} . This will imply that the

incentive constraint due to moral hazard binds on player i at his optimal equilibrium.

Theorem 6.1: (i) In games with a product structure $k^*(e_i) \leq v$ and

$$k^*(-e_i) \leq -v_i.$$

(ii) In games with observable actions $k^*(e_i) = v$ and

$$k^*(-e_i) = -v_i.$$

(iii) In moral hazard mixing games $k^*(e_i) < v$.

Proof: We calculate for each $\alpha \in \text{graph}(B)$, the scores $k^*(\alpha, e_i)$ and $k^*(-\alpha, -e_i)$. In the e_i case, $\lambda_i = +1$ and in the $-e_i$ case $\lambda_i = -1$. In either case, the score is determined by solving the linear programming problem (3.2).

Proof of (i): Suppose $k^*(e_i) > v$. Then, since k^* is independent of δ , for each $\delta \in [0, 1]$ there exists $\varepsilon > 0$ and $\alpha \in \text{graph}(B)$ such that

$k^*(\alpha, e_i) > v + [\delta/(1-\delta)]\varepsilon$ and $k^*(\alpha, e_i) = k^*(e_i) - \varepsilon$. It follows that for all a_i with $\alpha(a_i) > 0$

$$k^*(\alpha, e_i) = (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y_i \in Y_i} \prod_{y_i} (a_i, \alpha_{-i}^{\text{SR}})^{w_i(y_i)}$$

with $w_i(y_i) \leq k(e_i, \alpha) + \varepsilon$. Let a_i be such that $\alpha_i(a_i) > 0$ and $g_i(a_i, \alpha_{-i}) \leq g_i(a_i, \alpha_{-i}^*)$ for all $\alpha(a) > 0$. By the definition of

$v_i^* g_i(a_i, \alpha_{-i}^*) \leq v_i^*$. Consequently,

$$k^*(\alpha, e_i) \leq (1-\delta)v + \delta(k(e_i, \alpha) + \varepsilon)$$

or $k^*(\alpha, e_i) \leq v + [\delta/(1-\delta)]\varepsilon$. This contradicts the assumption that $k^*(\alpha, e_i) > v + [\delta/(1-\delta)]\varepsilon$, so $k^*(e_i)$

$\leq v$.

Next, suppose $k^*(-e_i) > -\underline{v}_i$. As above, there is an $\alpha \in \text{graph}(B)$ such that $k^*(\alpha, -e_i) > -\underline{v}_i + [\delta/(1-\delta)]\varepsilon$ and $k^*(\alpha, -e_i) = k^*(-e_i) - \varepsilon$. It follows that for all $a_i \in A_i$

$$-k^*(e_i, \alpha) \geq (1-\delta)g_i(a_i, \alpha_{-i}) + \sum_{y_i \in Y_i} \lambda_i^{17} g_i(a_i, \alpha_{-i}, y_i) w_i(y_i)$$

with $w_i(y_i) \geq k(e_i, \alpha) - \varepsilon$. Let a_i be such that $g_i(a_i, \alpha_{-i}) \geq g_i(a, \alpha_{-i})$ for all $a_i \in A_i$. By the definition of \underline{v}_i , $g_i(a_i, \alpha_{-i}) \geq \underline{v}_i$. Consequently

$$-k^*(e_i, \alpha) \geq (1-\delta) + \delta(-k(e_i, \alpha) - \varepsilon),$$

or $k^*(e_i, \alpha) \leq -\underline{v}_i + [\delta/(1-\delta)]\varepsilon$. Again, this is a contradiction.

Proof of (ii): Under perfect observability, the constraints (4.4) simplify to

$$\begin{aligned} v_i &= (1-\delta)g_i(a_i, \alpha_{-i}) + \delta w_i(a_i) \quad \alpha_i(a_i) > 0 \\ (6.1) \quad v_i &\geq (1-\delta)g_i(a_i, \alpha_{-i}) + \delta w_i(a_i) \quad \alpha_i(a_i) = 0 \end{aligned}$$

$$\lambda_i v_i \geq \lambda_i w_i(y_i).$$

For $\lambda_i = 1$, set $v_i = \min_{\alpha_i(a_i) > 0} g_i(a_i, \alpha_{-i})$, and

$w_i(a_i) = [v_i - (1-\delta)g_i(a_i, \alpha_i)]/\delta$ for a_i with $\alpha_i(a_i) > 0$;
 $w_i(a_i) = \min_{a \in A_i} g_i(a)$ for a_i with $\alpha_i(a_i) = 0$. Since these continuation payoffs satisfy constraints

(8.1), we conclude that

$$k^*(\alpha, e_i) \geq \min_{a_i \in \text{support}(\alpha_i)} 18g_i(a_i, \alpha_{-i}),$$

*

and $k^*(e_i) = \max_{\alpha} k^*(\alpha, e_i) \geq \max_{\alpha} \min_{\alpha_i \in \text{support}(\alpha_i)} g_i(a_i, \alpha_{-i}) = v$.¹⁹ Combining this

inequality with that of part (i) we have $k^*(e_i) = v$.

For $\lambda_i = -1$, set $v_i = \max_{a_i} g_i(a_i, m_i)$, where m_i is the minmax profile against player i , and set $w_i(a_i) = [v_i - (1-\delta)g_i(a_i, m_i)]/\delta$ for all a_i . Once again these continuation payoffs satisfy (6.1), and hold player i 's payoff to \underline{v}_i regardless of how he plays. Thus $k^*(-e_i) \leq \underline{v}_i$, and so combining with

part (i) yields $k^*(-e_i) = \underline{v}_i$.

Proof of (iii): We turn finally to moral hazard mixing games and $\lambda_i = 1$. Let $\Pi_i =$

$$\min_{a_i \in A_i, y_i \in Y_i, a_{SR}} \sum_{y_i \in Y_i} 20 \Pi_i (a_i, a_{SR})^{y_i} > 0 \text{ from the assumption that all outcomes have positive}$$

probability under all profiles.

Now for each α consider the problem (?). At the solution to this program, $v_i = k^*(\alpha, e_i) \geq w_i(y_i)$, so that

$$\sum_{y_i \in Y_i} 21 \Pi_i (a_i, \alpha_{SR})^{y_i} w_i(y_i) \leq k^*(\alpha, e_i) + \Pi_i \min_{y_i \in Y_i} 22 w_i(y_i).$$

Choose a_i so that $\alpha_i(a_i) > 0$. Then

$$k^*(\alpha, e_i) \leq (1-\delta)g_i(a_i, \alpha_{-i}) + \delta v_i + \delta \Pi_i \min_{y_i \in Y_i} w_i(y_i),$$

or

$$w_i(y_i) \geq \Pi_i [(1-\delta)/\delta] [k^*(\alpha, e_i) - \max g_i].$$

Since we know from above that $k^*(\alpha, e_i)$ satisfies

$\max g_i \geq \max v_i \geq \min g_i$, we may add this constraint to the LP problem and find also

$$\max g_i \geq w_i(y_i) \geq \Pi_i [(1-\gamma)/\delta] [\min g_i - \max g_i].$$

Consequently the relevant constraint set of continuation payoffs is bounded independent of α ,

and it follows that there exists α with

$k^*(e_i, \alpha) = k^*(e_i)$; that is, the supremum over α is attained.

Suppose then that $k^*(\alpha, e_i) = v$. For some a_i with $\alpha_i(a_i) > 0$, $g_i(a_i, \alpha_{-i}) \leq v$, by the definition of v . Since $\alpha_i(a_i) > 0$

$$v = (1-\delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y_i \in Y_i} \pi_{y_i(a_i, \alpha_{SR}, w_i(y_i))} \cdot 23$$

Since $g_i(a_i, \alpha_{-i}) \leq v$ and $w_i(y_i) \leq v$, it follows that $g_i(a_i, \alpha_{-i}) = v$, and since $\pi_{y_i}(a_i) > 0$ for all y_i ,

$w_i(y_i) = v$ for all $y_i \in Y_i$. But we then have

$$v^i * \geq (1-\delta)g_i(a_i, \alpha_{-i}) + \delta v$$

for all $a_i \in A_i$, with exact equality if and only if $\alpha_i(a_i) > 0$. In other words, α_i is a best response to α_{-i} yielding the payoff v , and contradicting the definition of a moral hazard mixing game. ■

Remark: One interpretation of part (iii) of the theorem is that player i 's moral hazard reduces his best equilibrium payoff relative to the case of observed actions unless the best observed action payoff v can be attained in the stage game with a profile where player i does not have an incentive to deviate. It is intuitive that moral hazard should not be costly in this case, as there is no need to "keep track" of player i 's action. The theorem shows that this is the only case where moral hazard has no additional cost.

We now give an example of a moral hazard mixing game.

	L	M	R
U	4,0	0,1	1,-100
D	2,2	1,1	0,3

FIGURE 6.1

The payoff matrix in Figure 6.1 gives the expected payoffs as a function of the actions. Player 1 is a long-run player facing a sequence of short-run player 2's. If player 1's actions are observable,

but not his mixed strategy, then by Theorem 5.1, his maximum equilibrium payoff is $v = 2$, which corresponds to player 1 playing D with probability $p \in [1/2, 100/101]$.

Now imagine that there are two outcomes y' and y'' , and that $\pi_{y'}(U) = \pi_{y''}(D) = 1 - \varepsilon$. It is clear that we may define payoffs $r_i(y^i, a^j)$ that give rise to the normal form in Figure 1, for example

$$r_1(y, L) = 4 - 6\varepsilon / (1 - 2\varepsilon)$$

$$r_1(y, R) = 2 - 6\varepsilon / (1 - 2\varepsilon).$$

This class of games obviously has a product structure and for $\varepsilon \neq 1/2$ satisfies the full-rank condition.

These games also satisfy the pairwise full rank condition described in Fudenberg, Levine and Maskin [1990]. According to that paper, if player 2 were a long-run player the limit set of equilibria would be the same as that with observable actions. However, with a short run player, the limit set is strictly smaller when $\varepsilon > 0$. This follows from the fact that the game is a moral hazard mixing game:

	U	D	
U	$1 - \varepsilon, \varepsilon$		
D			

is strictly positive, and the most player 1 can get from a pure strategy is $0 < v = 2$, while if player 1 is indifferent between up and down, he must clearly get no more than one.

Note that it can be shown that as $\varepsilon \rightarrow 0$, $k^*(\varepsilon)$ converges $v = 2$. This is quite generally

true: if we fix a payoff matrix and consider a sequence of corresponding information structures that converge to perfect observability, the maximum payoffs converge to v .

7. An Investment Game

We now consider an extended example of a product structure game. The long-run players are identical capitalists. The short-run players are investors who live two rounds. At the beginning of the stage game, all investors are young, and during the second round they are all old. In other words, in the repeated version of the game, the generations of investors do not overlap. Later we will discuss the consequences of overlapping generations of investors. It is convenient to assume that there are $\underline{x}v$ identical investors in each stage game, so that $n = L + \underline{x}v$.

Each investor is endowed with one indivisible unit of investment input when young. He has no endowment when old, and consumes only then. Capitalists have no endowment. Investment and the single consumption good are perfect substitutes in consumption for investors, while capitalists get no utility from the investment good. Consumption is completely perishable. All agents are risk neutral.

The production technology is available to all agents. It converts $x \geq \underline{x}$ units of investment input in the first round into consumption available during the second round. There are two possible levels of output, low output of 0 and high output equal to xf/p , where $f > 0$ and $0 \leq p \leq 1$ is the probability of high output. Notice that there are constant returns to scale (above the threshold level of \underline{x}), and that the expected amount of output per unit of input remains equal to f as p changes. We will assume that $f > 1$ so that the expected output is greater than the investment input. Since all agents are risk neutral, efficiency therefore requires that no investment be consumed directly.

We assume that if several different agents are producing, the outcome of their production processes are independent of each other. This assumption is important because it means that adding producers does not improve the ability to monitor production. If, for example, output were perfectly correlated among producers, it would be possible to design incentive schemes by

comparing the output of different producers, and punishing those that claimed to have less output.

On the other hand, independence means that while mean output is not changed by spreading inputs among more producers, the variance of output is reduced. Consequently, it is important that we have assumed that all agents are risk neutral: the reduced variance has no social value, and is irrelevant to the analysis.

An important aspect of this technology is that it involves an indivisibility: less than \underline{x} units of investment produce no output at all. If $\underline{x} = 1$, investors can carry out production using their own investment, and this will clearly be the unique equilibrium. We wish to focus on the case where production must be centralized due to economies of scale, so we always assume $\underline{x} > 1$. For convenience we have assumed that there are \underline{x}^v investors, so that at most v different producers may operate, and all investors may contribute equally to production.

It is evident that old investors will consume any of the consumption good they have. Consequently, only capitalists can purchase the investment input, for only they have an incentive to pay back investors. Indeed, since capitalists are the only agents who can be trusted to repay, they will be the only producers even if the investors have access to a superior technology.

When young, each investor must decide to whom (including himself) to give his indivisible investment input. Since investors cannot be trusted to repay loans, we simplify the analysis by assuming the feasible actions A_j for each investor $j = 1, 2, \dots, \underline{x}^v$ are the $L+1$ points corresponding to either consuming the investment input, or giving it to one of the L capitalists. In the second round, capitalist i must decide what to do with his output. (Old investors need not decide how to dispose of their output: they never have any.) We make several assumptions that substantially simplify the analysis, although it can be shown that they do not change our conclusions. First, we assume there is limited anonymity. A capitalist can tell if someone invested with him last period, but no further identification is possible. Let J_i be the realized number of investors providing investment to capitalist i . If output is zero, no decision is involved. If output is high the capitalist consumes a fraction $(1-\varphi)$ of it, with the remainder being equally divided as dividends among his investors. Because of limited anonymity, it is not possible to pay them different

amounts. In principle, the choice of φ can depend on the total number of investors J_i , although not on who they are. However, we assume that the capitalist must decide on the compensation scheme before observing the outcome of the investors' decisions.

Since the high level of output is Jif/p , for a given choice of φ , the capitalist gets $(1-\varphi)Jif/p$ units to consume, and his investors get each get $\varphi f/p$. To keep the game finite, we assume that φ lies in Φ a finite subset of the unit interval. We also assume that $0 \in \Phi$, so that the capitalist always has the option of paying no dividends. The action set for the capitalist is therefore simply $A_i = \Phi$.

Notice that we have assumed that the capitalist cannot pay either investors who do not invest with him, or other capitalists. We discuss this assumption further below.

All players observe to and by whom investment is made, and the amount of payment made (if any). However, the level of output and the amount consumed by each capitalist his his private information. While the natural definition of the public outcome y_i corresponding to capitalist i is his total dividend payment D_i , we will find it convenient to define $y_i = D_i/(Jif/p)$, so that y_i lies in the set Φ of capitalists' actions. In this way the distribution on outcomes induced by a given long-run player's action is independent of J_i , provided $J_i \geq \underline{x}$. Of course, if there is no output, or too few investors ($J_i < \underline{x}$), then $y_i = 0$ regardless of a_i , as the investors do not observe the amount the capitalist would have paid if output had been high.

For the short-run investor j , the observed outcome y_j is simply his chosen action $a_j \in A_j$.

In addition to the public signal $y = (y_1, \dots, y_n)$, a capitalist also receives the private signal $z_i \in \Phi$ corresponding to his chosen action. The probability distribution $\Pi_{yz}(a)$ places probability one on $z_i = a_i$ for $i \in LR$, and $y_j = a_j$ for $j \in SR$, and the signals y_i are independent of one another for $i \in LR$. The random variable y_i takes on the values a_i with probability p , and 0 with probability $1-p$. Note that the distribution of public signals corresponding to capitalist i is independent of the signals corresponding to other capitalists, and depends only on his own action and the actions of the short run players.

Remark: Because there are indivisibilities, lotteries may be welfare improving in this economy, as pointed out by Prescott and Townsend [1984]. Indeed, the type of indivisibility assumed here has been modelled using lotteries by Rogerson [1985], and in the work on real business cycles by Hansen [1985]. Lotteries may be represented in this game by introducing a public randomization device at the beginning of each period. In the case of primary interest, when the discount factor of the capitalists is close to one, our results do not depend on whether or not there are lotteries: as Sorin [1986] and Fudenberg and Maskin [1991] point out, deterministic alternation between different allocations can serve as well. However, public randomization does strengthen and simplify the analysis for discount factors that are not near one, and when we consider non-asymptotic results for the one capitalist case, we will assume that there is a public randomization device.

8. Analysis of the Investment Game

We will now show that the investment game satisfies the sufficient rank condition, so we apply Theorem 5.1. We can then solve the LP problem (5.3) to find the limit set Q . Since this has dimension L , we have an explicit computation of the limit equilibrium payoffs. In the case of a single capitalist we also explicitly calculate an equilibrium giving the capitalist the greatest possible rent, and determine the range of discount factors for which this is a solution. We also consider the economic interpretation of the results.

To see that the game satisfies the sufficient rank condition, fix a strategy profile α_{SR} for the investors, and define P_i to be the probability that $J_i \geq \underline{x}$, that is, capitalist i has enough investors to produce. Then the corresponding matrix $\Pi_i(\alpha_{SR})$ is

$$\begin{array}{c|c} 0 & 1, & 0, & 0, & \dots & \gamma \\ \hline \varphi & 1, & 0, & 0, & \dots & 1 \end{array}$$

$$\begin{array}{c|cccc}
 a_i & \varphi_i & 1-P_i + P_i(1-p) & P_i p & 0 & \dots \\
 \hline
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
 \hline
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
 \hline
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
 \hline
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots
 \end{array}$$

For $p \neq 0$, this matrix has full row rank.

Second, note that the unique equilibrium of the stage game is for the capitalists not to pay anything and for the investors not to make any loans. In this case the capitalists all get 0, and this is also their individually rational payoff level, since they can guarantee themselves 0 by never paying. It is clear therefore that $k^*(-e_i) = 0$.

To find the capitalist i 's maximum equilibrium payoff, $k^*(e_i)$, we write out the LP problem (5.3) that determines $k^*(\alpha, e_i)$. Corresponding to P_i , which is the probability of sufficient investors investing given the strategies of the short run players, define α_i to be the expectation conditional on α_{SR} of the random variable that is 0 if fewer than \underline{x} investors invest with capitalist i , and the number of investors who invest with him otherwise. The linear programming problem (5.3) may be written as

(8.1) $\max v_i$ subject to

(8.2) $v_i = (1-\delta)(1-\alpha_i)\text{if} + \delta[(1-P_i+P_i(1-p))w_i(0) + P_i p w_i(a_i)]$

if $\alpha_i(a_i) > 0$, then

$$v_i \geq (1-\delta)(1-\alpha_i)\text{if} + \delta [(1-P_i+P_i(1-p))w_i(0) + P_i p w_i(a_i)]$$

(8.3) $v_i \geq w_i(a_i)$.

The solution defines $k^*(\alpha, e_i)$.

Recall that to find $k^*(e_i)$, we will maximize over $\alpha \in \text{graph}(B)$. Examining (8.2) and (8.3), we see that the constraints are homogeneous in α_i , w_i and v_i . Consequently, if we begin with a feasible solution, increasing α_i by a factor of λ leads to a new feasible solution in which w_i and v_i are also scaled up by λ . Moreover, if $\alpha_i > 0$ is a best response for the investors, then so is $\lambda \alpha_i$.

$= \underline{x}v$. Consequently, either $k^*(e_i)$ is zero, or we may assume $i = \underline{x}v$ and $P_i = 1$. We proceed to solve (8.1) to (8.3) in the latter case, in which case (8.2) simplifies to

$$v_i = (1-\delta)(1-a_i)\underline{x}vf + \delta[(1-p)w_i(0) + pw_i(a_i)] \text{ if } \alpha_i(a_i) > 0. \quad (8.2')$$

$$v_i \geq (1-\delta)(1-a_i)\underline{x}vf + \delta[(1-p)w_i(0) + pw_i(a_i)].$$

Let \bar{a}_i be the largest element in the support of α_i . For any a_i in the support of α_i , (8.2') shows that

$$w_i(a_i) = w_i(i) - [(1-\delta)/\delta p] (i-a_i) \underline{x}vf.$$

Moreover, (8.2') requires that for any a_i

$$w_i(a_i) \leq w_i(i) = [(1-\delta)/\delta p] (i-a_i) \underline{x}vf.$$

From this, we see that if we solve the problem subject to the constraints corresponding to $a_i \neq i$, 0 in (8.2') and (8.3), then we may certainly also find $w_i(a_i)$ the remaining constraints in (8.2') and (8.3). Since (8.2') has v_i increasing in $w_i(i)$ and $w_i(0)$, (8.3) must hold with exact equality for i , as must the incentive constraint for $a_i = 0$. This gives us three linear equations to solve for v_i , $w_i(i)$, and $w_i(0)$, and we find

$$(8.4) \quad w_i(0) = [v_i - (1-\delta)\underline{x}vf]/\delta,$$

$$(8.5) \quad v_i = w_i(i) = \underline{x}vf[1-a_i/p].$$

We now wish to maximize $k^*(\alpha, e_i)$ over α that are incentive compatible for investors. The incentive constraint for the investors implies that the expected value of a_i must be at least $1/f$ if $P_i > 0$. (Otherwise, $v_i = 0$.) Since (8.4) is decreasing in a_i , the maximizing strategy for the capitalist strategy should assign probability one to

$$(8.6) \quad \underline{\varphi} \equiv \min \{\varphi \in \Phi \mid \varphi \geq 1/f\}.$$

One consequence of this is that the capitalist need not use a mixed strategy. The only potential gain from mixing is if $\underline{\varphi} > 1/f$, in which case the capitalist might try to mix in order to hold the investors to their reservation utilities. However, roughly speaking, mixing introduces a form of

moral hazard which keeps the capitalist from benefiting from the reduction in payments. Indeed, in the case $p=1$, Fudenberg, Kreps and Maskin [1990] show that the best the capitalist can do is to get the worst payoff in the support of his mixed strategy. We have just showed that this result holds true even with exogenous moral hazard.

Substituting (8.6) into (8.5), we conclude

Theorem 8.1: $k^*(e_i) = \underline{x}^v [(f-1) - (1/p-1) - (f/p)(\underline{\phi}-1/f)]$.

One consequence is that if $pf < 1$ then no single capitalist can get the first best payoff, $v^* = \underline{x}^v(f-1)$. There are two reasons for this, corresponding to the last two terms in the expression of $k^*(e_i)$. The first term, $\underline{x}^v(1/p-1)$ reflects the cost of providing the capitalist with incentives under moral hazard; if $p = 1$, this term is equal to 0. The second term, $\underline{x}^v(f/p)(\underline{\phi}-1/f)$, is less interesting. It reflects the fact that the capitalist cannot hold the investors to their reservation utility using a pure strategy, and as we observed above, the capitalist cannot improve his payoff by mixing. If we make the not unreasonable assumption that $1/f \in \Phi$, this term vanishes.

If $p = 1$ and $1/f \in \Phi$, then neither moral hazard nor discreteness poses a problem, and each capitalist can attain the first best. Conversely, if these conditions fail, his equilibrium payoff remains bounded away from the first best, as $\bar{\delta} \rightarrow 1$.

Even though each individual capitalist's highest equilibrium profit is no greater than if he were the only capitalist, the aggregate profit to all capitalists can be greater than if there is a single capitalist. This reflects the fact that investors can now "punish" each capitalist by investing with a different one, rather than by shutting down the whole economy. In fact, if there are enough capitalists, the limit set E will contain a point where the capitalists' profit is the first-best level $v^* = \underline{x}^v(f-1)$. In this case, we say that the first best is attainable. From Theorem 8.1, we may calculate

Corollary 8.2: The first best is attainable as $\delta \rightarrow 1$ with one capitalist if and only if $1/f \in \Phi$ and $p = 1$. With $m > 1$ capitalists, the first best is also attainable if $p \geq \underline{\varphi}$ and $m \geq (f-1)/f(1-\underline{\varphi}/p)$ which is necessary and sufficient.

This shows that the greater the number of capitalists, the broader the range of parameters p and f for which the first best is attainable.

In the one capitalist case, we may explicitly find an equilibrium yielding $k^*(e_i)$. Indeed, the fact that δ is near one is used only in the proof of Theorem 3.1, and there only to ensure that $w_i(a_i)$ is in fact an equilibrium payoff vector. However, in the one capitalist case, $w_i(a_i)$ can be realized as a public randomization between the best and worst equilibrium, provided it lies in between them. Since (8.4) to (8.5) show that $w_i(a_i)$ is always less than or equal to the best equilibrium payoff, it is sufficient that $w_i(a_i) \geq 0$. Only $w_i(0)$ is at issue, and it is calculated from (8.4) and (8.5). Consequently, we see that an equilibrium yielding the asymptotic rent is possible provided

$$(8.7) \quad \delta \geq a_i/p.$$

substituting in $\underline{\varphi}$, we find

Corollary 8.3: With one capitalist if

$$\delta \geq 1/pf + (1/p)(\underline{\varphi}-1/f)$$

then there is an equilibrium in which he gets $k^*(e_i)$.

9. Extensions to the Investment Game

We now want to consider various extensions of the investment model. Although the discussion is specific to this model, many of the same issues are relevant in other economics settings.

Welfare of the Short-Run Players: Previously, we argued that the first best is not

attainable. This simply means that the capitalists can not get as much rent as they could from precommitment. We now argue that the terminology is not misleading: in fact with moral hazard and a single capitalist, from the broader welfare point of view that includes the investors, the first best is not attainable.

The first best from the joint perspective of investors and capitalists requires full employment: that is, that all \underline{x}^v units of investment available to the economy be employed each period. If this is not the case there is said to be unemployment. We now argue that in any equilibrium it must occur infinitely often, contradicting full efficiency.

Suppose in fact that full employment always occurs in equilibrium. If v_i is an equilibrium present value when there is full employment, it must not be optimal to pay zero. If $w_i(0)$ is the continuation present value when no payment is made, this means that

$$(9.1) \quad v_i \geq (1-\delta)\underline{x}^v + \delta w_i(0).$$

This may be rewritten as

$$(9.2) \quad v_i \geq (1-\delta)\underline{x}^v + \delta(v_i - [v_i - w_i(0)]).$$

Rearranging terms then yields the inequality

$$(9.3) \quad v_i - w_i(0) \geq ((1-\delta)/\delta) [\underline{x}^v - v_i].$$

However we know that because investors are individually rational, in any equilibrium, $v_i \leq \underline{x}^v(f-1)$. It follows from (9.3) that

$$(9.4) \quad v_i - w_i(0) \geq ((1-\delta)/\delta) \underline{x}^v.$$

If $p < 1$, it then follows that every period in which there is full employment, there is probability $(1-p)$ that the continuation present value drops by at least $((1-\delta)/\delta)\underline{x}^v$. By assumption, the continuation equilibrium with present value $w_i(0)$ also has full employment. This can occur only a fixed number of times T (the number depending on the discount factor) before the constraint $w_i(0) \geq 0$ is violated. Since the probability that this happens is $(1-p)^T$, there is a positive probability that there is less than full employment. We conclude that with moral hazard, unemployment must occur infinitely often.

In this calculation, it is clear that as $\delta \rightarrow 1$ it is possible to construct equilibria in which there

is full employment for increasingly long periods of time. It is tempting to argue that this means "approximate efficiency" is attained as we approach the limit. However, as the period of full employment grows longer, the capitalist's discount factor grows closer to one at exactly the same rate, so if we use it to discount future unemployment we see that the present value of future unemployment remains unchanged.

We should note that it is not true that in all equilibria investors get their reservation utility. If Φ contains only two points, $0, \varphi \in \Phi$, and $\varphi > 1/f$, we know from our solution in the one capitalist case that if δ is close enough to one, there exists an equilibrium in which φ is played in the initial period. Moreover, as long as $\varphi \in \Phi$, regardless of how many points Φ has, this same equilibrium can be enforced for δ near enough one by punishing the capitalist with zero if he ever plays anything other than φ or 0 .

We should also point out that the inefficiency conclusion derives from the fact that the only methods of punishing a single capitalist is either to stop providing him with inputs, which is inefficient, or to transfer output to the investors, which is subject to moral hazard. If there were a method of carrying out transfer payments that was perfectly observed, and so not subject to moral hazard, it could be used as an alternative method of punishing capitalists who default on loans, and the inefficiency would not arise. (Although the upper bound on each individual capitalists' payoff would not change.)

A simple, if artificial, way to modify the example to make this point is to endow each capitalist with $m > 2\varphi/vf/p$ units of a pure consumption good ("honey") valued only by the capitalists and some third group of players ("drones") who have no role in the game other than accepting honey from the capitalists. (If the investors valued honey, then the capitalists could pay them in honey, and this would circumvent the moral hazard problem.)

Capitalists make their transfers (if any) to the drones at the start of each period before the investors make lending decisions, and these transfers are observed by the investors.

Consider a strategy profile of the following form: along the equilibrium path all of the investors lend to the capitalist, and the capitalist always chooses $a_i = \varphi$. Any period in which the

capitalist pays back less than $\frac{\varphi_x v f}{p}$, he make a transfer of honey of size R (to be determined) to the drones. If he fails to do so, we revert to the no trade equilibrium forever. We claim for δ sufficiently large, we can find an R so that these strategies are an equilibrium.

The investors' behavior is clearly optimal given the strategy of the capitalist. If we set $R = \frac{\varphi_x v f}{p} \delta$, which is less than the capitalist's endowment whenever $\delta \geq 1/2$, then the capitalist does not have any incentive to deviate from paying the investors. Can he gain by refusing to pay the honey to the drones? For $\delta \geq 1/2$ the fine paid the drones is less than $2\frac{\varphi_x v f}{p}$, while the loss of profit from punishment is a fixed amount forever. It follows that for δ near enough one, the capitalist should pay the drones.

These strategies therefore form an equilibrium for δ near enough to one. Moreover, there is full employment on the equilibrium path, so this equilibrium is efficient. The key here is that the transfers of honey are an efficient way of punishing the capitalist, while the threat of shutting down the economy to force the capitalist to accept his punishment does not have to be executed, and so has no negative efficiency implications.

Lending Between Capitalists: Bulow and Rogoff [1989] argue that in the case of sovereign debt, the ability of a borrower to make deposits in a bank at the competitive interest rate means that reputation alone is not enough to create a market for debt: there is always a point at which it would be optimal for the borrower to renege on the debt and deposit the proceeds in a bank. This might suggest that if we expand the model by allowing permitting capitalists to freely transfer consumption to one another, and there are two or more of them, the equilibrium set would collapse to the static equilibrium, but this is not the case. Any equilibrium of the model without transfers is also an equilibrium of the model with transfers, provided all capitalists follow the strategy of never repaying any loans. This contrasts with the Bulow-Rogoff assumption that banks must repay all depositors. In our model, this amounts to picking one equilibrium from the equilibrium set.

Since many features of our model are unrealistic, we do not contend that our conclusions are more plausible than theirs. However, since in our model lending is enforced by the threat of

exclusion from future borrowing their assertion that "under fairly general conditions, lending ... cannot be supported by a country's reputation for repayment," is too strong.

Overlapping Generations: So far, we have assumed that investors live in non-overlapping generations. Although not especially realistic, this has the virtue that it makes the environment a repeated game. Moreover, it is equivalent to an economy in which there are two-period lived overlapping generations of investors and the completely perishable consumption good may not be used as a substitute for investment in the production process. Since investors wish to consume only when old, and the consumption good is perishable, there is no purpose in transferring the consumption good to young investors. As a result, there is no possible trade between old and young investors, and the economy is in effect one with non-overlapping generations of investors.

If generations live more than two periods, however, additional possibilities arise. With non-overlapping generations, extending the life of investors does not lead to any great complication. However, in the overlapping generations case, capitalists are no longer necessary in order to arrange production. Suppose that investors live three periods, have endowment only when young and wish to consume only when old, but that they are also middle aged in between. This has the implication that they can engage in production when middle aged. Consider the following strategies: Each period, a middle-aged investor is picked at random from the pool of "eligible" middle-aged investors, and all young investors give him their investment input. The lucky designee then produces and consumes all the output himself. Young investors agree to participate in this scheme because if they do not they are not "eligible" when they become middle aged. Under our extreme assumption of risk neutrality, this scheme actually implements the first best.

Once risk aversion is introduced a lottery scheme of this type becomes less desirable. If the indivisibilities are large, the inefficiency is large as well: Imagine compensating employees at General Motors by a small chance at all of General Motors profits for a year. Consequently, while capitalists are not necessary for production in this more general setup, they will have a role to play if indivisibilities and risk aversion are large compared to the losses caused by the capitalists' moral hazard.

The Role of Reputation: The analysis so far has been in terms of a repeated game. The equilibrium involves "reputation" in the loose sense that capitalists repay loans so that they will be allowed to borrow in the future. However, reputation here is purely an equilibrium phenomenon, in the sense that repayment of a loan need not necessarily signal a future willingness to repay. Kreps and Wilson [1982] and Milgrom and Roberts [1982] have introduced a more explicit sort of reputation. If the investors a priori suspect that there is a possibility that the capitalists will choose to repay regardless of future consequences, then capitalists can maintain this belief by repaying. Fudenberg and Levine [1991b] have shown that typically with a single patient capitalist, this type of explicit reputation leads him to do as well as he could through precommitment, so long as the stage game has simultaneous moves.

However, when the stage game has sequential moves, we obtain a lower bound on the power of reputation effects that can be lower than the commitment payoff. In examples with sequential moves but no moral hazard, we can show that this lower bound is in fact tight by constructing equilibria that hold the long-run player to the lower bound, but whether the lower bound is tight in games with sequential moves and moral hazard is still an open question.

To see why, we construct an equilibrium in the investment game that holds the capitalist to payoff zero even when explicit reputation is considered, for the case where $p = 1$ so there is no moral hazard. The best equilibrium payoff with reputation effects here is $\underline{v}[f-1]$, and this can be obtained even without reputation with the investor strategies of never investing with a capitalist who defaults. Now suppose that there is another type of capitalist, the "strong" capitalist who always pays $1/f$. Consider the following strategy profile: the original capitalist does not repay the first loan he receives, and no investor ever invests. If some investor does invest and then regardless of what the capitalist does, switch to the equilibrium in which the investors always invests and the capitalist always repays. This means that there is no advantage in imitating the strong capitalist.

This construction relies on the fact that without moral hazard the commitment payoff is attainable as an equilibrium even without reputation, so that we can construct an equilibrium in

which an initial default has no cost. When there is moral hazard, an initial default may have a cost because it reveals that the capitalist is not "strong" therefore restricting him to the most he can get without reputation. Exactly what the best and worst equilibrium are in this model is as yet unknown.

REFERENCES

- Abreu, D., D. Pearce and E. Stacchetti [1986], "Optimal Cartel Monitoring with Imperfect Information," Journal of Economic Theory, 39, 251-69.
- Abreu, D., D. Pearce and E. Stacchetti (1990), "Toward a Theory of Discounted Repeated Games With Imperfect Monitoring," Econometrica, 58, 1041-64.
- Atkeson, A. [1987], "International lending with moral hazard and risk of repudiation," Federal Reserve Bank of Minneapolis WP #381.
- Ausabel, L. and R. Deneckere [1989], "Reputation in Bargaining and Durable Goods Monopoly," Econometrica, 57, 511-32.
- Bulow, J. and K. Rogoff [1989], "Sovereign Debt: Is to Forgive to Forget," American Economic Review, (March).
- Friedman, J. [1971], "A non-cooperative equilibrium for supergames," Review of Economic Studies, 38, 1-12.
- Fudenberg, D., D. Kreps and E. Maskin (1990), "Repeated Games With Long-Run and Short-Run Players," Review of Economic Studies, 57, 555-74.
- Fudenberg, D. and D. Levine [1983], "Subgame Perfect Equilibria of Finite and Infinite Horizon Games," Journal of Economic Theory 31, 251-258.
- Fudenberg, D. and D. Levine [1989], "Reputation and Equilibrium Selection in Games with a

Patient Player," Econometrica, 57, 759-778.

Fudenberg, D., and D. Levine (1991a), "An Approximate Folk Theorem in Repeated Games With Imperfect Private Information," Journal of Economic Theory, 54, 26-41.

Fudenberg, D. and D. Levine [1991b], "Maintaining a Reputation when Strategies are Imperfectly Observed," forthcoming, Review of Economic Studies.

Fudenberg, D., D. Levine, and E. Maskin (1991), "The Folk Theorem in Discounted Repeated Games With Imperfect Public Information," mimeo.

Fudenberg, D., and E. Maskin (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," Econometrica, 54, 532-54.

Fudenberg, D. and E. Maskin [1991], "On the Dispensability of Public Randomization," Journal of Economic Theory, 53, 428-38.

Green, E., and R. Porter [1984], "Non-Cooperative Collusion Under Imperfect Price Information," Econometrica, 52, 975-94.

Gul, F. [1986], "Two Essays on the Bargaining Foundations of Value Theory," Ph.D. dissertation, Princeton University.

Hansen, G. [1985], "Indivisible Labor and the Business Cycle," Journal of Monetary Economics, 16, 309-27.

Kreps, D. and R. Wilson [1982], "Reputation and Imperfect Information," Journal of Economic

Theory, 50, 253-79.

Lehrer, E. [1988], "Nash Equilibria of n-Player Repeated Games With Semi-Standard Information," mimeo, Northwestern.

Marimon, R. [1988], "Wealth Accumulation with Moral Hazard," mimeo.

Matsushima, H. [1989], "Efficiency in Repeated Games with Imperfect Monitoring," Journal of Economic Theory, 48, 428-442.

Milgrom P. and J. Roberts [1982], "Predation, Reputation and Entry Deterrence," Econometrica, 50, 443-60.

Prescott, E. and R. Townsend [1984], "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," Econometrica 52, 21-45.

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, 59-70.

Rogerson, R. [1985], "Indivisible Labor, Lotteries and Equilibrium," mimeo.

Selten, R. [1977], "The Chain-Store Paradox," Theory and Decision, 9, 127-59.

Sorin, S. [1986], "On Repeated Games with Complete Information," Mathematics of Operations Research, 4, 147-160.

APPENDIX: Characterizing Equilibria with Halfspaces

Abreu, Pearce and Stachetti [1990] introduced the notion of a self-generated set of payoffs for games with only long-run players, and showed that a sufficient condition for a set of payoffs to be in $E(\bar{\delta})$ is that it be self-generated. Fudenberg, Levine and Maskin [1990] used a local version of this result, and we prove Theorem 3.1 using a similar technique. The proof makes use of two definitions.

Definition A.1: If for a given v , W , and $\bar{\delta}$ an α exists so that (α, v) is enforceable with respect to W and $\bar{\delta}$, we say that v is generated by W . $P(\bar{\delta}, W)$ is the set of all points generated by W .

Definition A.2: A subset $W \subseteq \mathbb{R}^L$ is smooth if it is closed, has non-empty interior with respect to \mathbb{R}^L , and the boundary of W is C^2 submanifold of \mathbb{R}^L .

Theorem 3.1: (i) For all $\bar{\delta}$, $E(\bar{\delta}) \subseteq Q$.

(ii) If the dimensions of $Q \subset \mathbb{R}^L$ is L , then $\lim_{\bar{\delta} \rightarrow 1} E(\bar{\delta}) = Q$.

Proof of (i): We will show more strongly that $E^*(\bar{\delta})$, the convex hull of $E(\bar{\delta})$, is contained in Q . If not, we may find a halfspace $H(\lambda, k)$ and a point $v \in E(\bar{\delta})$ such that $\lambda \bullet v = k > k^*(\lambda)$ and $\lambda \bullet v \leq k$ for all

$v' \in E^*(\bar{\delta})$. Then v must be enforceable with continuation payoffs in

$E(\bar{\delta}) \subseteq H(\lambda, k)$, contradicting the definition of $k^*(\lambda)$.

Proof of (ii): Since any compact convex set with non-empty interior Q can be approximated arbitrarily closely by smooth convex sets $W \subset \text{interior}(Q)$, this will follow provided that for any such set W there exists a $\bar{\delta}$ such that $E(\bar{\delta}) \supseteq W$. We claim in fact that it is enough that for each $v \in W$ there exists a $\bar{\delta} < 1$ and an open set U containing v such that $U \subseteq P(\bar{\delta}, W)$. If so, we may find an open cover $\{U\}$ of W together with

$\delta_U < 1$, such that $U \subseteq P(\delta_U, W)$. Since W is compact, choose a finite subcover, and let $\bar{\delta}$ be the maximum of δ_U over this subcover. Let

$\delta' \geq \bar{\delta}$. Then $\delta' \geq \delta_U$ for each U in the subcover. Since W is convex, it may easily be shown that $U \subseteq P(\delta', W)$. Since the U 's cover W , this implies $W \subseteq P(\delta', W)$. Finally, since W is bounded, it is straightforward to use the principle of optimality to conclude that since each point in W can be enforced using continuation payoffs in W , that $W \subseteq E(\delta')$.

To conclude the proof, we must show that for all $v \in W$ there is a $\delta < 1$ and an open neighborhood U of v with $U \subseteq P(\delta, W)$. This is easy to do for $v \in \text{interior}(W)$. As in Fudenberg, Levine and Maskin [1990], let (α, u) be a static Nash equilibrium. Fix a U containing v and with closure in the interior of W . This implies each $u \in U$ can be expressed as $(1-\delta)g() + \delta u$, where $g() \in W$ for some $\delta < 1$. Moreover, the continuation payoffs $w(y) = (1-\delta)g() + \delta u$ clearly enforce (α, u) so $U \subseteq P(\delta, W)$.

Next we consider points v on the boundary of W . Fix such a v , and let λ be normal to W at v . Let $k = \lambda \bullet v$, and let $H = H(\lambda, k)$ be the unique halfspace in the direction λ that contains W and whose boundary h is tangent to W at v . Since $W \subset \text{interior}(Q)$, it follows that H is a proper subset of the maximal halfspace $H^*(\lambda)$. Let (α, v) be a strategy that generates a boundary point of $H^*(\lambda)$ using continuation payoffs in $H^*(\lambda)$. Since v is a boundary point of H , which is a proper subset of $H^*(\lambda)$, for some $\delta' < 1$ and $\varepsilon > 0$, (α, v) can be enforced with respect to $H(\lambda, k-\varepsilon)$.

It may be shown directly from the incentive constraints that for $\delta'' \geq \delta'$, we may find $w(y, \delta'')$ that enforce (α, v) and a $\delta'' > 0$ such that

$$w(y, \delta'') \in H(\lambda, k - [\delta'(1-\delta'')/\delta''(1-\delta')]\varepsilon),$$

and $|w(y, \delta'') - v| \leq (1-\delta'')$.

Consider, then, the ball $U(\delta'')$ around v of radius $2(10\delta'')$. Since W is smooth, for δ'' sufficiently close to one the ball is translated by a small constant independent of y generating incentive compatible payoffs in a neighborhood U of v . ■

³This observation originates with Matsushima [1988].