working paper
department
of economics

EQUILIBRIUM PAYOFFS WITH LONG-RUN
AND SHORT-RUN PLAYERS AND
IMPERFECT PUBLIC INFORMATION

Drew Fudenberg
and
David K. Levine

No. 524   June 1989

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

EQUILIBRIUM PAYOFFS WITH LONG-RUN
AND SHORT-RUN PLAYERS AND
IMPERFECT PUBLIC INFORMATION

Drew Fudenberg
and
David K. Levine

No. 524   June 1989

EQUILIBRIUM PAYOFFS WITH LONG-RUN AND SHORT-RUN PLAYERS

AND IMPERFECT PUBLIC INFORMATION[*]

Drew Fudenberg

and

David K. Levine[**]

May 1989

Abstract

We present a general algorithm for computing the limit, as $\delta \to 1$, 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 use the algorithm to obtain an exact characterization of the limit set when the players' realized actions, but not their choices of mixed strategies, are observable. We show each that long-run player's highest equilibrium payoff is generally greater in this case than when their actions are unobserved.

1. Introduction

Fudenberg-Kreps-Maskin [1989] (here referred to as FKM) show that the folk theorem extends in the natural way to repeated games with long run and short run players if the players' choices of mixed strategies are observed at the end of each period. They further show that the folk theorem need not obtain when the only information revealed at the end of each period is the realized choice of action as opposed to the intended mixed strategy. In the latter case, with a single long-run player, they characterize the equilibrium payoffs in the limit as the discount factor goes to one. Their proof explicitly constructs equilibria in "review strategies" whose form is simple.

This paper considers the case of multiple long run players. Rather than explicitly constructing equilibrium strategies, we extend the techniques of Fudenberg-Levine-Maskin [1989] (FLM) to study games with several long-run players facing a sequence of short-run opponents, allowing the possibility that players' strategies need not be observable. As in FLM, at the end of each period all players observe a public signal that is imperfectly correlated with players strategies. We present a general algorithm for computing the limit of the set of payoffs of perfect public equilibria, and use it to provide an exact characterization of the limit set in the case where the players' realized actions, but not their choices of mixed strategies, are observed by their opponents.

When all players are long-run, the limit set of payoffs is the same regardless of whether or not the players' actions are observable, provided a "pairwise full rank" condition is satisfied. We show that this is not the case in games where some of the players are short-run: the introduction of moral hazard may lower each long-run player's highest equilibrium payoff.

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 $a \in A = \times_{i=1}^{n} A_i$ induces a probability distribution $\pi_{yz}(a)$ over publicly observed outcomes y and privately observed outcomes $z = (z_1, \ldots, z_n)$. We let $\pi_y(a)$ denote the marginal for 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 on his own private information and on the realized outcome: his opponents' strategies matter only in their influence on the distribution over outcomes. This model includes as a special case games where the public information y 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. The key to the folk theorem is that there is an outcome y is publicly observed, which rules out games where the players receive only private signals. The latter case is considered in Lehrer [1988] and Fudenberg and Levine [1989].

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 $\alpha = (\alpha_1, \ldots, \alpha_n)$ of mixed actions, we can compute the induced distribution over outcomes,

$$\pi_{yz}(\alpha) - \sum_{a\in A} \pi_{yz}(a)\alpha(a), \; \pi_y(\alpha) - \sum_{a\in A} \pi_y(a)\alpha(a)$$

and the expected payoffs

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

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

As in FKM, we label players so that $i \in LR - \{1,2,\ldots,\ell\}$, $\ell \leq n$, are long-run players, whose objective is to maximize the normalized 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 present value is

$$(1-\delta) \sum_{t-0}^{\infty} \delta^t g_i(t)$$

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 $j \in SR - \{\ell+1,\ldots,n\}$ 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 (see FKM for other examples).

In the repeated game, in each period $t - 0,1,\ldots,$ the stage game is played, and the corresponding public outcome is then revealed. The <u>public history</u> at the end of period t is $h(t) - (y(0),\ldots,y(t))$. The <u>private history for long-run player i</u> at the end of period t is $h_i(t) -$

$(a_i(0), z_i(0), \ldots, a_i(t), z_i(t))$. A <u>strategy for long-run player i</u> is a
sequence of maps mapping public and private histories $(h(t-1), h_i(t-1))$ to
mixed actions. A <u>strategy for the period-t players of type j</u>, $j \in SR$, 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, if
the public information reveals the past choices of action, as in FKM then
these private histories are extraneous.

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. We denote the repeated game
with discount factor δ by $G(\delta)$.

Let A_i denote the space of mixed actions, and

$$B: A_1 \times \ldots \times A_\ell \to A_{\ell+1} \times \ldots \times A_n$$

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

Our focus is on a special class of the Nash equilibria that we call
<u>perfect public equilibria</u>. 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 i's payoffs. This is why we can 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 actuibs $\alpha(h(t-1))$ 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 FLM for an example of the way this correlation can generate additional equilibrium payoffs.)

3. Dynamic Programming and Local Generation

Let $E(\delta) \subset \mathbb{R}^{\ell}$ be the set of normalized values that can arise in perfect public equilibria when the discount factor if δ. Abreu-Pearce-Stachetti [1987] introduced the notion of a self-generated set of payoffs for games with long-run players only, and showed that a sufficient condition for a set of payoffs to be in $E(\delta)$ is that it be self-generated. FLM showed that a sufficient condition for a set of payoffs W to be in $E(\delta)$

for all δ sufficiently close to 1 is that W be locally generated. This section observes that those definitions and results carry over immediately to games with some short-run players.

<u>Definition 3.1</u>: Let δ, $W \subseteq \mathbb{R}^\ell$ and $v \in \mathbb{R}^\ell$ be given. Action α is <u>enforceable with respect to v, W and δ</u> if there exist vectors $w(y) \in \mathbb{R}^\ell$, $y \in Y$, such that $\alpha \in$ graph(B) and for all $i \in LR$ and $a_i \in A_i$,

$$
\begin{aligned}
(3.1) \quad & v_i = (1-\delta)g_i(a_i,\alpha_{-i}) + \delta \sum \pi_y(a_i,\alpha_{-i})w_i(y) \quad \text{for } a_i \text{ s.t. } \alpha_i(a_i) > 0 \\
& v_i \geq (1-\delta)g_i(a_i,\alpha_{-i}) + \delta \sum \pi_y(a_i,\alpha_{-i})w_i(y) \quad \text{for } a_i \text{ s.t. } \alpha_i(a_i) = 0
\end{aligned}
$$

and

$$(3.2) \qquad w(y) \in W.$$

We also say that the $w(y)$'s <u>enforce</u> (α,v) with respect to W and δ. If for a given α, a v exists satisfying (3.1), we say that α is <u>enforceable with respect to W and δ</u>. Note that each $w(y)$ specifies continuation payoffs only for the long-run players.

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 present value starting tomorrow on when y occurs today is given by $w_i(y)$, then player i receives exactly present value v from any of the actions in the support of α_i, and no choice of action yields a higher present value. We can exploit the linearity of the incentive constraints to show:

<u>Lemma 3.2 (Enforcement With Small Variation)</u>: Let C be a convex subset of $\mathbb{R}^\ell$. Suppose (α,v) is enforced by $W(y) \in C$ and δ. Let $\bar{w} = \Sigma_{y \in Y} \pi_y(\alpha)w(y)$. Then there is a constant κ such that for all $\delta' \geq \delta$

there is a $v' \in \mathbb{R}^\ell$ and there are continuation payoffs $w'(y) \in C$ that enforce (α, v'). Moreover,

(i) $\|w'(y) - \bar{w}\| \leq \kappa(1-\delta')$

(ii) $\Sigma_{y \in Y} \pi_y(\alpha) w'(y) = \bar{w}$

(iii) $w'(y) + (\delta')^{-1} (v'-v) \in W$.

(iv) $w'(y) = \bar{w} + \dfrac{\delta(1-\delta')}{\delta'(1-\delta)} (w(y)-\bar{w})$.

Proof: This is just a matter of checking that $w'(y)$ defined by (iv) are in C and satisfy (3.1), (3.2) and the other properties; see for example FLM.   ∎

Definition 3.3: If for a given v, W, and δ an α exists so that (α, v) is enforceable with respect to W and δ, we say that v is generated by W. $P(\delta, W)$ is the set of all points generated by W. A subset W of $\mathbb{R}^\ell$ is locally generated if for each $v \in W$ there exists a $\delta < 1$ and an open set U containing v such that $U \subseteq P(\delta, W)$.

Lemma 3.4: If $W \subseteq \mathbb{R}^\ell$ is compact, convex and locally generated, then there exists a $\delta < 1$ such that for all $\delta' \geq \delta$, $W \subseteq E(\delta')$.

Proof: By local generation 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 δ be the maximum of δ_U over this subcover. Let $\delta' \geq \delta$. Then $\delta' \geq \delta_U$ for each U in the subcover. Since $U \subseteq P(\delta_U, W)$, by Lemma 3.2(iii) for each $v \in U$ we may find $w(y) \in W$ that enforce v with respect to the discount factor δ'. So $U \subseteq P(\delta', W)$, and since the U's cover W, this implies $W \subseteq P(\delta', W)$. Finally, since W is compact, 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 $W \subseteq E(\delta)$.   ∎

4. Enforceability on Halfspaces and the Set of Limit Equilibria

This section describes an algorithm that computes the limiting value of the set $E(\delta)$ as δ tends to one. Section 5 uses this algorithm to characterize the limit set of equilibria.

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

__Definition 4.1:__ The maximal score attainable by action α in direction λ with discount factor δ, denoted $k^*(\alpha,\lambda,\delta)$ is the maximum of $\lambda\cdot v$ such that there exists $v \in \mathbb{R}^\ell$ and $k \in \mathbb{R}$ with (α,v) enforceable with respect to δ and $H(\lambda,k)$ and such that $k = \lambda \cdot v$. The halfspace associated with this maximum is $H^*(\alpha,\lambda,\delta)$; its boundary is the hyperplane $h^*(\alpha,\lambda,\delta)$ given by $\lambda \circ v = k^*(\alpha,\lambda,\delta)$.

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.

__Lemma 4.1:__ (i) $k^*(\alpha,\lambda,\delta) = k^*(\alpha,\lambda)$ independent of δ.

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

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

__Proof:__ (i) if $k^*(\alpha,\lambda,\delta') = k$, then there is a v with $\lambda \cdot v = k$ and continuation payoffs $w'(y) \in H(\lambda,k)$ that enforce (α,v) with respect to $(\delta',H(\lambda,k))$. Now consider the payoffs

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

Since $\delta"$ is less than one, the coefficient of v in the above equations is less than one as well. Then since v is on the boundary of $H(\lambda,k)$ and each $w'(y)$ is in $H(\lambda,k)$, $w'(y)$ is in $H(\lambda,k)$ as well. Moreover, it is easy to check from (3.1) and (3.2) that the $w"(y)$ enforce (α,v) with respect to $(\delta",H(\lambda,k))$.

(ii) $k^*(\alpha,\lambda) > \lambda \cdot g(\alpha)$ 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) $k^*(\alpha,\lambda) - g(\alpha)$ clearly 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 necessary 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.

<u>Definition 4.2</u>: The maximal score in direction λ, $k^*(\lambda)$, solves $k^*(\lambda) - \sup_{\alpha\in\text{graph}(B)} k^*(\alpha,\lambda)$. The maximal halfspace in direction λ is $H^*(\lambda) - H(\lambda,k^*(\lambda))$.

<u>Definition 4.3</u>: $Q - \cap H^*(\lambda)$.

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.

<u>Theorem 4.1</u>: (i) For all δ, $E(\delta) \subseteq Q$.

(ii) If the dimensions of $Q \subset \mathbb{R}^\ell$ is ℓ, then $\lim_{\delta\to 1} E(\delta) - Q$.

Proof of (i): We will show more strongly that $E^*(\delta)$, the convex hull of

$E(\delta)$, is contained in Q. If not, we may find a halfspace $H(\lambda,k)$ and a

point $v \in E(\delta)$ such that $\lambda \cdot v = k > k*(\lambda)$ and $\lambda \cdot v \leq k$ for all

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

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

Some preliminaries are needed before proving (ii).

Definition 4.4: A subset $W \subseteq \mathbb{R}^\ell$ is smooth if it is closed, has non-empty

interior with respect to $\mathbb{R}^\ell$, and the boundary of W is C^2 submanifold

of $\mathbb{R}^\ell$.

Since any compact convex set with non-empty interior Q can be

approximated arbitrarily closely by smooth convex sets $W \subset$ interior(Q),

part (ii) of the theorem will follow once we show that each such W in the

interior of Q is locally generated.

Lemma 4.4: If $W \subseteq$ interior(Q) is a smooth convex set, then Q is locally

generated.

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

$v \in$ interior(W). As in FLM, let $\hat{\alpha}$ 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 $(1-\delta)g(\hat{\alpha}) + \delta\bar{w}$, where $\bar{w} \in W$ for some

$\delta < 1$. Moreover, the continuation payoffs $w(y) - \bar{w}$ clearly enforce $(\hat{\alpha},u)$

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 $k = \lambda \cdot v$, 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 ⊂ interior(Q), it follows

that H is a proper subset of the maximal halfspace $H^*(\lambda)$. Let α 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 δ' < 1 and ε > 0, (α,v) can actually be

enforced with respect to H(λ,k-ε).

By enforcement with small variation (Lemma 3.2(iv)), for δ" ≥ δ', we

may find w(y,δ") that enforce (α,v) and a $\bar{\kappa} > 0$ such that

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

and $|w(y,\delta")-v| \le \bar{\kappa}(1-\delta")$.

Consider, then, the ball U(δ") around v of radius $2\bar{\kappa}(1-\delta")$. Since

W is smooth,[1] for δ" sufficiently close to one there exists a $\bar{\kappa} > 0$ the

difference between H and W in U(δ) is at most $\bar{\kappa}(1-\delta")^2$. It follows

that there exists a δ < 1 such that v can be enforced by continuation

payoffs w(y) ∈ interior W. Since w(y) are interior they may be trans-

lated by a small constant independent of y generating incentive compatible

payoffs in a neighborhood U of v.   ∎

5. The Characterization of E(δ) for Games with a Product Structure

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 short run players and

a single long-run player FKM show that the folk theorem fails, but are never-

theless able to characterize equilibrium payoffs. Using Theorem 4.1, we can

---

[1]This observation originates with Matsushima [1988].

extend this latter characterization to encompass many long-run and short-run players, and moral hazard as well.

Formally, a game has a <u>product structure</u> if $y = (y_1, \ldots, y_\ell, y_{\ell+1})$ with probability one, and $\pi_y(a) = \pi_{y_1}(a_1) \ldots \pi_{y_\ell}(a_\ell) \pi_{y_{\ell+1}}(a_{\ell+1})$ 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 hiw "own" outcome y_i, and the y_i are statistically independent. In addition, we require a <u>full rank condition</u>: for $i = 1, \ldots, L$ the matrix $\pi_i = \{\pi_{y_i}(a_i)\}$ with rows $a_i \in A_i$ and columns $y_i \in Y_i$ should have rank m_i equal to the number of player i's actions. It is easy to see that under this condition any $\alpha \in \text{graph}(B)$ can be enforced by some (not necessarily feasible) specification of the continuation payoffs.[2] Notice that this is a generalization of a game with <u>observable actions</u>. 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. We also define V to be the subset of $\mathbb{R}^\ell$ generated as convex combinations of payoffs to long-run players $(g_1(\alpha), \ldots, g_\ell(\alpha))$, where $\alpha \in \text{graph}(B)$. Note that this set is closed, since B is, and convex. Our goal is to prove:

<u>Theorem 5.1</u>: In a game with a product structure satisfying the full rank condition, Q is the intersection of V with the 2ℓ constraints $k^*(-e_i) \leq v_i \leq k^*(e_i)$.

---

[2] FLM call this the "individual full rank" condition to distinguish it from a stronger condition called "pairwise full rank".

If, following Fudenberg and Maskin [1986], we assume that Q has dimension ℓ, it follows from Theorem 4.1 that Q is the limit of $E(\delta)$ as $\delta \to 1$.

Although this gives a characterization of limit payoffs as $\delta \to 1$, 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 $k^*(-e_i) \le v_i \le k^*(e_i)$. To prove Theorem 5.1 in turn suffices to show that for non-coordinate directions λ we can obtain scores $k^*(\lambda)$ on the boundary of v.

<u>Lemma 5.2</u>: In a game with a product structure, if λ is a non-coordinate direction, then $k^*(\lambda) - \lambda \cdot v$, where $v \in V$ and $V \subseteq H(\lambda, k^*(\lambda))$.

<u>Proof</u>: By Lemma 4.1(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 $(g_1(\alpha), \ldots, g_\ell(\alpha)) - v$ that may be enforced on $h(\lambda, k')$. By the definition of V we may find $\alpha \in$ graph(B) and a k' such that $(g_1(\alpha), \ldots, g_\ell(\alpha)) - v$ and $h(\lambda, k')$ is tangent to V at v. Consequently, we need only construct $w(y)$ satisfying $\lambda \cdot w(y) - k'$, and such that the incentive constraints (3.1) and (3.2) are satisfied. Since $y - (y_1, \ldots, y_\ell, y_{\ell+1})$, we initially restrict attention to $w_i(y) - w_i(y_i)$, and such that the incentive constraints all hold with exact equality. They may be written as

(5.1) $\qquad v_i - (1-\delta)g_i(a_i, \alpha_{-i}) + \delta\Sigma_{y_i \in y_i} \pi_{y_i}(a_i)w_i(y_i)$.

Since π_i has rank m_i, there exist 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, $\lambda \cdot \hat{w}(y) = k'$ provided that

$$\lambda_1 \hat{w}_1(y) + \lambda_2 \hat{w}_2(y) = k' - \sum_{i=3}^{\ell} \lambda_i \hat{w}_i(y_i).$$

Clearly, setting $\hat{w}_1(y) = w_1(y_1) + w_1'(y_{-1})$, $\hat{w}_2(y) = w_2(y_2) + w_2'(y_1)$, and $\hat{w}_i(y) = w_i(y_i)$ for $i \neq 1,2$ preserves incentive compatibility for any functions w_1' of y_{-1} and w_2' of y_1. If we choose $w_2'(y_1) = -(\lambda_1/\lambda_2)w_1(y_1)$, and $w_1'(y_{-1}) = (1/\lambda_1)(k' - \sum_{i=2}^{\ell} \lambda_i w_i(y_i))$, the resulting $\hat{w}(y)$ completes the proof. ∎

Now we show how to determine the constraints $k^*(e_i)$ and $k^*(-e_i)$ in some classes of games of economic interest. Following FKM, we define two bounds on the payoffs of the long-run player $i = 1,\ldots,\ell$. First,

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

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

$$v_i^* \equiv \max_{\alpha \in \text{graph}(B)} \min_{a_i \in \text{support}(\alpha_i)} g_i(a_i, \alpha_{-i}),$$

which is the most player i can get if he cannot be trusted not to maximize within the support of a mixed strategy. Let $\underline{\alpha}_{-i}^i$ 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 v_i^*.

FKM show that for games with observed actions and a single long-run player the limit set of equilibrium payoffs is the interval $[\underline{v}_1, v_1^*]$. FLM

consider games with unobserved actions where all players are long-run. In these games all strategy profiles are in graph(B), so $v_i^* - \max g_i(\alpha)$, and any strategy profile α^{*i} that gives player i payoff v_i^* necessarily has α_i^{*i} a static best response to α_{-i}^{*i}. This means that player i can be induced not to deviate from α^{*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^{*i} need not be a static best response to α_{-i}^{*i}, which is why $k^*(e_i)$ can be less than $\max_\alpha g_i(\alpha)$.

An important class of games that do not have perfect observability are <u>moral hazard mixing games</u>. These are games for which $\pi_{y_i}(a_i)$ is strictly positive for all $a_i \in A_i$, $y_i \in Y_i$ and $i - 1,\ldots,\ell$, so there is genuine moral hazard. In addition we require that if $\alpha \in$ graph(B) and $g_i(\alpha) - v_i^*$ 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.

<u>Theorem 5.3</u>: (i) In games with a product structure $k^*(e_i) \leq v_i^*$ and $k^*(-e_i) \leq -\underline{v}_i$.

(ii) In games with observable actions $k^*(e_i) - v_i^*$ and $k^*(-e_i) - - \underline{v}_i$.

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

<u>Proof</u>: We calculate for each $\alpha \in$ graph(B), the scores $k^*(\alpha, e_i)$ and $k^*(-\alpha, -e_i)$. In the e_i case, let $\gamma - +1$ and in the $-e_i$ case let $\gamma - -1$. Then $k^*(\alpha, \gamma e_i)$ is the solution to the linear programming problem

$$\max \gamma v_i$$

subject to:

$$v_i = (1-\delta)g_i(a_i,\alpha_{-i}) + \delta\Sigma_{y_i\in y_i} \pi_{y_i}(a_i)w_i(y_i) \quad \alpha(a_i) > 0$$

(5.2)
$$v_i \geq (1-\delta)g_i(a_i,\alpha_{-i}) + \delta\Sigma_{y_i\in y_i} \pi_{y_i}(a_i)w_i(y_i) \quad \alpha(a_i) = 0$$

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

To see that this program determines $k^*(\alpha,e_i)$, note that the incentive constraints for players other than i can always be satisfied for some continuation payoffs $w(y)$ by the full rank condition, and that the condition that $w(y_i) \in H(\gamma e_i, v_i)$ is simply $\gamma v_i \geq \gamma w_i(y_i)$.

<u>Proof of (i)</u>: Suppose $k^*(e_i) > v_i^*$. Then, since k^* is independent of δ, for each $\delta \in [0,1]$ there exists $\epsilon > 0$ and $\alpha \in \text{graph}(B)$ such that $k^*(\alpha,e_i) > v_i^* + [\delta/(1-\delta)]\epsilon$ and $k^*(\alpha,e_i) = k^*(e_i) - \epsilon$. It follows that for all a_i with $\alpha(a_i) > 0$

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

with $w_i(y_i) \leq k(e_i,\alpha) + \epsilon$. 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 $\alpha(a_i') > 0$. By the definition of v_i^*, $g_i(a_i,\alpha_{-i}) \leq v_i^*$. Consequently,

$$k^*(\alpha,e_i) \leq (1-\delta)v_i^* + \delta(k(e_i,\alpha)+\epsilon)$$

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

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

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

with $w_i(y_i) \geq k(e_i,\alpha) - \epsilon$. Let a_i be such that $g_i(a_i,\alpha_{-i}) \geq g_i(a_i',\alpha_{-i})$

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

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

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

Proof of (ii): Under perfect observability, the constraints (5.2) simplify to

$$v_i = (1-\delta)g_i(a_i, \alpha_{-i}) + \delta w_i(a_i) \quad \alpha(a_i) > 0$$

(5.3) $$v_i \geq (1-\delta)g_i(a_i, \alpha_{-i}) + \delta w_i(a_i) \quad \alpha(a_i) = 0$$

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

For $\gamma = 1$, 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 (5.3), we conclude that

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

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

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

For $\gamma = -1$, set $v_i = \max_{a_i} g_i(a_i, m_{-i}^i)$, where m^i is the minmax profile against player i, and set $w_i(a_i) = [v_i - (1-\delta)g_i(a_i, m_{-i}^i)]/\delta$ for all a_i. Once again these continuation payoffs satisfy (5.3), 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 $\gamma = 1$. Let $\underline{\pi}_i = \min_{a_i \in A_i, y_i \in y_i} \pi_{y_i}(a_i) > 0$ from the assumption that all outcomes have positive probability under all profiles.

Now for each α consider the problem of finding max v_i subject to the constraints (5.2). At the solution to this program, $v_i = k^*(\alpha,e_i) \geq w_i(y_i)$, so that

$$\Sigma_{y_i \in y_i} \pi_{y_i}(a_i)w_i(y_i) \leq k^*(\alpha,e_i) + \underline{\pi}_i \min_{y_i \in y_i} w_i(y_i).$$

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

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

or

$$w_i(y_i) \geq \underline{\pi}_i^{-1}[(1-\delta)/\delta][k^*(\alpha,e_i)-\max g_i].$$

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

$$\max g_i \geq w_i(y_i) \geq \underline{\pi}_i^{-1}[(1-\gamma)/\delta][\min g_i-\max g_i].$$

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 sup over α is attained.

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

$$v_i^* - (1-\delta)g_i(a_i,\alpha_{-i}) + \delta \sum_{y_i \in Y_i} \pi_{y_i}(a_i)w_i(y_i).$$

Since $g_i(a_i,\alpha_{-i}) \leq v_i^*$ and $w_i(y_i) \leq v_i^*$, it follows that $g_i(a_i,\alpha_{-i}) - v_i^*$, and since $\pi_{y_i}(a_i) > 0$ for all y_i, $w_i(y_i) - v_i^*$ for all $y_i \in Y_i$. But we then have

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

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_i^*, and contradicting the definition of a moral hazard mixing game. ▌

Remark: An alternative 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_i^* 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 1

The payoff matrix in Figure 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 $v_i^* = 2$, which corresponds to player 1 playing D with probability $p \in [1/2, 100/101]$.

Now imagine that there are two outcomes y_1' and y_2'', and that $\pi_{y'}(U) = \pi_{y''}(D) = 1 - \epsilon$. It is clear that we may define payoffs $r_i(y_1, a_2)$ that give rise to the normal form in Figure 1, for example

$$r_1(y_1', L) = 4 - 6\epsilon/(1-2\epsilon)$$

$$r_1(y_1'', L) = 2 - 6\epsilon/(1-2\epsilon).$$

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

These games also satisfy the pairwise full rank condition described in FLM. According to FLM, 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 $\epsilon > 0$. This follows from the fact that the game is a moral hazard mixing game:

$$\begin{bmatrix} \pi_{y'}(U) & \pi_{y''}(U) \\ \pi_{y'}(D) & \pi_{y''}(D) \end{bmatrix} = \begin{bmatrix} 1-\epsilon, & \epsilon \\ \epsilon & 1-\epsilon \end{bmatrix},$$

is strictly positive, and the most player 1 can get from a pure strategy is $0 < v_1^* = 2$, 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 $\epsilon \to 0$, $k^*(e_i)$ converges $v_1^* = 2$. 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_i^*.

6. Conclusion

We have investigated the limit of the set $E(\delta)$ of public equilibria. When this limit is smaller than the limit set of feasible payoffs, it raises the possibility that a strictly larger limit set could be obtained by considering a larger class of equilibria. In particular, some action by a short-run player might not be a best response to any independent randomization by his opponents, but be a best response to a correlated distribution over the opponents' strategies. When the players' actions are perfectly observed, this possibility is irrelevant, as all equilibria are public equilibria. However, when actions are imperfectly observed, two

players can condition on the public information and their own past play in such a way that the distribution over their actions given only the public information is a correlated distribution. Thus, there can be equilibria where the short-run players choose responses that are not in the range of B. The characterization set of _all_ the equilibrium payoffs in these games remains an open problem.

REFERENCES

Abreu, D., D. Pearce and E. Stacchetti (1987), "Toward a Theory of Discounted Repeated Games With Imperfect Monitoring," Harvard.

Fudenberg, D., D. Kreps and E. Maskin (1989), "Repeated Games With Long-Run and Short-Run Players," this journal.

Fudenberg, D., and D. Levine (1989), "An Approximate Folk Theorem in Repeated Games With Privately Observable Outcomes," mimeo.

Fudenberg, D., D. Levine, and E. Maskin (1989), "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.

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

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

Matsushima, H. [1988], "Efficiency in Repeated Games with Imperfect Monitoring," forthcoming, Journal of Economic Theory.

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

Date Due