


MIT LIBRARIES



3 9080 00567414 5





Digitized by the Internet Archive
in 2011 with funding from
Boston Library Consortium Member Libraries

<http://www.archive.org/details/approximatefolkt00fude>

DEWEY

HB31
.M415
no. 525

**working paper
department
of economics**

INST TECH
AUG 23 1989
LIBRARIES

AN APPROXIMATE FOLK THEOREM WITH
IMPERFECT PRIVATE INFORMATION


Drew Fudenberg
and
David K. Levine

No. 525

March 1989

**massachusetts
institute of
technology**

**50 memorial drive
cambridge, mass. 02139**



AN APPROXIMATE FOLK THEOREM WITH
IMPERFECT PRIVATE INFORMATION

Drew Fudenberg
and
David K. Levine

No. 525

March 1989

MIT LIBRARIES
AUG 23 1989
RECEIVED

AN APPROXIMATE FOLK THEOREM WITH
IMPERFECT PRIVATE INFORMATION^{*}

by

Drew Fudenberg^{**}

and

David K. Levine^{***}

March 1989

* We are grateful to Andreu Mas-Colell. NSF Grants 87-08616 and 86-09697 and a grant from the UCLA Academic Senate provided financial support.

** Department of Economics, Massachusetts Institute of Technology

*** Department of Economics, University of California, Los Angeles; part of this research was conducted while at the University of Minnesota.

Abstract

We give a partial folk theorem for approximate equilibria of a class of discounted repeated games where each player receives a private signal of the play of his opponents. Our condition is that the game be "informationally connected," meaning that each player i has a deviation that can be statistically detected by player j regardless of the action of any third player k . Under the same condition, we obtain a partial folk theorem for the exact equilibria of the game with time-average payoffs.

JEL Classification numbers = 022, 026

Keywords: Repeated games, folk theorem, private information,
approximate equilibrium

1. Introduction

We give a partial folk theorem for the approximate equilibria of a class of repeated games with imperfect information satisfying an informational linkage condition. Every payoff vector which exceeds a mutual threat point and is generated by one-shot mixed strategies from which no player can profitably deviate without having some effect on other players, is approximately an approximate sequential equilibrium payoff if the discount factor is close enough to one.

The class of repeated games we consider has complete but imperfect information. Each period, player i chooses an action a_i , then observes his own payoff and also a signal z_i of the play of his opponents. This model includes as a special case games of imperfect public information as defined by Fudenberg-Levine-Maskin [1989]. In these games, there is a publicly observed variable \tilde{y} , and each player i observes $z_i = (\tilde{y}, a_i)$. The public information case includes the literature on repeated oligopoly (Green-Porter [1983], Abreu-Pearce Stacchetti [1986]), repeated principal-agent games (Radner [1981, 1985], Rubinstein-Yaari [1983]), repeated partnerships (Radner [1986], Radner-Myerson-Maskin [1986]), as well as the classic case of observable actions (Auman-Shapley, Friedman [1971], Rubinstein [1981], Fudenberg-Maskin [1986]). It also includes the examples studied in the literature of repeated games with "semi-standard information", as discussed in Sorin [1988].

While examples of repeated games that have been formally studied have public information, players have private information in some situations of economic interest. Indeed, the central feature in Stigler's [1961] model of secret price-cutting is that each firm's sales depend on the prices of its rivals, but firms only observe their own demand. The goal of this paper is

to show that a partial folk theorem extends to these cases, provided we relax the notion of equilibrium slightly.

The key element in any folk theorem is that deviators must be punished. This requires, with three or more players, that non-deviators coordinate their punishments, and may even require a player to cooperate in his own punishment. With public information, the necessary coordination can be accomplished by conditioning play on the commonly observed outcome. When players receive different private signals they may disagree on the need for punishment. This raises the possibility that some players may believe that punishment is required while others do not realize this. The most straightforward way to prevent such confusion from dominating play is to periodically "restart" the strategies at commonly known times, ending confusion, and permitting players to recoordinate their play. This, however, makes it difficult to punish deviators near the point at which play restarts and forces us toward approximate equilibrium.

The use of approximate equilibrium leads to an important simplification, because it allows the use of review strategies of the type introduced by Radner [1981]. Under these strategies, each player calculates a statistic indicating whether a deviation has occurred. If the player's statistic crosses a threshold level at a commonly known time, he communicates this fact during a communication stage of the type introduced by Lehrer [1986]. Our informational linkage condition ensures that all players correctly interpret this communication, so that the communications stage allows coordination of punishments. The importance of approximate equilibrium is that Radner-type review strategies do not punish small deviations, i.e., those that do not cause players statistics to cross the threshold level. Consequently, there will typically be small gains to deviating. A second

consequence of approximate equilibrium is that sequentiality loses its force, and the review strategies are perfect whenever they are Nash. This is because punishments end when the strategies are restarted so the cost to punishers is small. Thus the punishment are "credible" in the sense that carrying them out is an approximate equilibrium.

Our results are closely connected to those of Lehrer [1986, 1988a,b,c,d, 1989] who considers time-average Nash equilibrium of various classes of games with private information and imperfect monitoring but with non-stochastic outcomes. In addition to the deterministic nature of outcomes, Lehrer's work has a different emphasis than ours, focusing on completely characterizing equilibrium payoffs under alternative specifications of what happens when a time average fails to exist. Our work provides a partial characterization of payoffs for a special but important class of games.

A secondary goal of our paper is to clarify the connection between Lehrer's work and other work on repeated time-average games, and work on discounted games such as ours. In particular we exposit the connection between approximate discounted equilibrium and exact time-average equilibrium, which is to some extent already known to those who have studied repeated games. We emphasize the fact that the time-average equilibrium payoff set includes limits of approximate as well as exact discounted equilibria. Although we show that the converse is false in general, the type of construction used by Lehrer (and us) yields both an approximate discounted and a time-average theorem. For economists, who are typically hostile to the notion of time-average equilibrium, we hope to make clear that results on repeated games with time-average payoffs are relevant, provided the logic of approximate equilibrium is accepted.

If we consider the set of equilibrium payoffs as potential contracts in a mechanism design problem, there is an interpretation of approximate equilibrium that deserves emphasis. The set of equilibrium payoffs in our theorem in some cases strictly exceeds that for exact equilibria. This means that exact equilibrium may be substantially too pessimistic. The efficiency frontier for contracts may be substantially improved if people can be persuaded to forego the pursuit of very small private benefits. For example, even if ethical standards have only a small impact on behavior, they may significantly improve contracting possibilities.

Section 2 of the paper lays out the repeated games model. Section 3 establishes a connection between approximate sequential discounted equilibria with discount factor near one, and time-average equilibrium. In particular, both types of equilibria can be constructed from finite-horizon approximate Nash equilibria. Section 4 describes our informational linkage condition, and proves a partial folk theorem.

2. The Model

In the stage game, each player i , $i = 1$ to N , simultaneously chooses an action a_i from a finite set A_i with m_i elements. Each player observes an outcome $z_i \in Z_i$, a finite set with M_i elements. We let $z = (z_1, \dots, z_n)$ and $Z = \times_{i=1}^n Z_i$. Each action profile $a \in A = \times_{i=1}^n A_i$ induces a probability distribution $\pi_z(a)$ over outcomes z . Each player i 's realized payoff $r_i(z_i)$ depends on his own observed outcome only; the opponents' actions matter only in their influence on the distribution over outcomes.

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

$$g_i(a) = \sum_{z \in Z} \pi_z(a) r_i(z_i).$$

We also define $\bar{d} = \max_{1, a_i, a'_i, a_{-i}} g_i(a_i, a_{-i}) - g_i(a'_i, a_{-i})$ to be the maximum one-period gain any player could obtain by playing one of his actions a_i instead of another a'_i .

We will also wish to consider various types of correlated strategy profiles. A correlated action profile α is simply a probability distribution over A . We write α_{-i} for the induced probability distribution over $A_{-i} = \times_{j \neq i} A_j$. A belief by player i is such a probability distribution, and the vector (a_i, α_{-i}) represents the correlated action in which player i plays a_i and the other players correlate according to α_{-i} . We also write α_i for the marginal induced over A_i by α . If the play of players is independent, we refer to α as a mixed action profile, in which case α is characterized by and may be identified with $(\alpha_1, \dots, \alpha_n)$, the vector of marginals.

For any correlated action profile α , we can calculate the induced distribution over outcomes

$$\pi_z(\alpha) = \sum_{a \in A} \pi_z(a) \alpha(a).$$

We can also calculate the marginal distribution over player i 's outcomes

$$\pi_i(z_i, \alpha) = \sum_{(\bar{z} \in Z | \bar{z}_i = z_i)} \pi_z(\alpha).$$

Finally, we may calculate the expected payoff to player i

$$g_i(\alpha) = \sum_{z_i \in Z_i} \pi_i(z_i, \alpha) r_i(z_i) = \sum_{z \in Z} \pi_z(\alpha) r_i(z_i).$$

In the repeated game, in each period $t = 1, 2, \dots$, the stage game is played, and the corresponding outcome is then revealed. The history for player i at time t is $h_i(t) = (a_i(1), z_i(1), \dots, a_i(t), z_i(t))$. We also

let $h_i(0)$ denote the null history existing before play begins. A strategy for player i is a sequence of maps $\sigma_i(t)$ mapping his private history $h_i(t-1)$ to probability distributions over A_i .

A system of beliefs b_i for player i specifies for each time t and private history $h_i(t)$, a probability distribution $b_i(h_i(t))$ over private histories of other players of length t . A profile of beliefs b for all players is consistent with the strategy profile σ , if for every finite time horizon the truncation of b is consistent with the truncation of σ in the sense of Kreps-Wilson [1982]. This requires that there exists for each truncation T a sequence of truncated strategy profiles $\sigma_n^T \rightarrow \sigma^T$, putting strictly positive probability on every action, such that the unique sequence of truncated belief profiles derived from Bayes law, $b_n^T \rightarrow b^T$.

Given a system of beliefs b_i and a history $h_i(t)$ a distribution is induced over the history of all players play. Given this distribution and the strategy profile σ , a corresponding distribution is induced over play at all times τ . In turn this gives an expected payoff in period τ to player i , which we denote by $G_i(\tau, h_i(t), b_i, \sigma)$. The corresponding normalized present value to player i at discount factor $0 \leq \delta < 1$ is

$$W_i(h_i(t), b_i, \sigma, \delta) = (1-\delta) \sum_{\tau=1}^{\infty} \delta^{\tau-1} G_i(\tau, h_i(t), b_i, \sigma).$$

While if $\delta = 1$

$$W_i^T(h_i(t), b_i, \sigma) = (1/T) \sum_{\tau=1}^T G_i(\tau, h_i(t), b_i, \sigma),$$

and

$$W_i(h_i(t), b_i, \sigma, 1) = \limsup_{T \rightarrow \infty} W_i^T(h_i(t), b_i, \sigma).$$

Notice that in the case of the initial null history $h_i(0)$, the present

value of time average expected payoff is the same for all beliefs consistent with σ . To emphasize this, we write $W_i(\sigma, \delta)$ and $W_i^T(\sigma)$ in place of $W_i(h_i(0), b_i, \sigma, \delta)$ and $W_i^T(h_i(0), b_i, \sigma)$.

For $\delta < 1$, a strategy profile σ is an ϵ -Nash equilibrium if for each player i and strategy σ'_i

$$(2.1) \quad W_i((\sigma'_i, \sigma_{-i}), \delta) \leq W_i(\sigma, \delta) + \epsilon$$

There is also a corresponding notion of a truncated ϵ -Nash equilibrium, where we require

$$W_i^T(\sigma'_i, \sigma_{-i}) \leq W_i^T(\sigma) + \epsilon.$$

An ϵ -sequential equilibrium is defined in a similar way, except that we require there exist beliefs b consistent with σ such that

$$(2.2) \quad W_i(h_i(t), b_i, (\sigma'_i, \sigma_{-i}), \delta) \leq W_i(h_i(t), b_i, \sigma, \delta) + \delta^{t-1} \epsilon$$

for all players i , and all histories $h_i(t)$ (not merely the null initial history). Notice that W_i is measured as a present value at time 1, rather than time t . Consequently, an ϵ -sequential equilibrium is defined so that the gain to deviating at time t is of order ϵ , in time- t units. Notice that this definition is stronger than the usual version, found in Fudenberg and Levine [1983], for example: ordinarily the gain to deviating at time t is measured in time-0 units, not time- t units.

For $\delta = 1$, we will impose a regularity condition on the equilibrium. A uniform Nash equilibrium is a profile σ such that

$$(2.3) \quad W_i^T(\sigma) \text{ converges, and}$$

$$(2.4) \quad \text{for all } \rho > 0, \exists \tau, \text{ such that } T \geq \tau \text{ implies}$$

$$W_i^T(\sigma'_i, \sigma_{-i}) \leq W_i^T(\sigma) + \rho.$$

A uniform sequential equilibrium is a profile σ such that

$$(2.5) \quad W_i^T(h_i(t), b_i, \sigma) \text{ converges for all histories } h_i(t) \text{ and consistent beliefs } b, \text{ and}$$

$$(2.6) \quad \text{there exist consistent beliefs } b \text{ such that for all } \rho > 0, \exists \tau \text{ such that } T \geq \tau \text{ implies for all histories } h_i(t) \text{ and strategies } \sigma_i,$$

$$W_i^T(h_i(t), b_i, (\sigma'_i, \sigma_{-i})) \leq W_i^T(h_i(t), b_i, \sigma) + \rho.$$

When σ satisfies condition (2.3) we call it regular; if it satisfies (2.5) we say it is sequentially regular.

Our definitions of time average equilibrium require that the appropriate regularity condition be satisfied. In other words, we have required that the time-average payoff exist if no player deviates, and that there be a uniform bound on the gain to deviating in any sufficiently long subgame. Both conditions are responses to our unease in comparing lim infs or lim sups of payoffs when the limit does not exist. We impose the conditions because we interpret time averaging as the idealized version of a game with long finite horizon or very little discounting. Below we show that non-uniform equilibria cannot always be interpreted in this way.

3. Discounting and Time Averaging

Our goal is to characterize the set of equilibrium payoffs when players are very patient. In the next section we characterize truncated ϵ -Nash equilibrium payoffs with small ϵ as containing a certain set V^* . Since this notion of equilibrium is not the most economically meaningful, we first show that this is in fact a strong characterization of payoffs.

Specifically we show

Theorem 3.1: Suppose there exists a sequence of times T^n , non-negative numbers $\epsilon^n \rightarrow 0$ and strategy profiles σ^n such that σ^n is a T^n -truncated time-average ϵ^n -Nash equilibrium and $W_i^{T^n}(\sigma^n) \rightarrow v_i$. Then

- (A) There exists a sequence of discount factors $\delta^n \rightarrow 1$, non-negative numbers $\epsilon^n \rightarrow 0$ and strategy profiles σ^n such that σ^n is an ϵ^n -sequential equilibrium for δ^n and $W_i(\sigma^n, \delta^n) \rightarrow v_i$.
- (B) There exists a uniform sequential equilibrium strategy profile σ such that $W_i(\sigma, 1) = v_i$.

Remark 1: It is well known that the hypotheses of the theorem imply that σ is a time-average Nash equilibrium (see for example, Sorin [1988]).

Remark 2: The hypothesis is weak in the sense that only approximate Nash equilibria are required. However, as long as we consider approximate equilibrium, or infinite time-average equilibrium, (A) and (B) show that sequentiality has little force. Roughly, the reason is that after a deviation, no further deviations are anticipated, so the cost of punishment must be only paid once. This has negligible cost if players are very patient. It is easily seen in the following example, which is a special case of the main theorem in the next section. Consider a game in which players perfectly observe each others past play, and suppose there is a pure action profile a for which $g(a)$ strictly exceeds the minmax for all players. Clearly there is a discount factor $1 > \underline{\delta} \geq 0$ and number of periods K , such that if $\delta \geq \underline{\delta}$ the loss to each player of being minmaxed by his opponents for K periods exceeds any possible one-shot gain of deviating from a . Consider then a strategy profile consisting of a review phase in which all players play a , and n different punishment phases in the i^{th} of which, player i is minmaxed for K periods. Initially, the

game is in the review phase, and the review phase continues as long as no player deviates, or more than one player deviates simultaneously. Whenever a single player deviates, a punishment phase against that player immediately begins. The punishment phase continues K periods, regardless of whether any further deviation occurs, then the game restarts with a new review phase. Since no player can profitably deviate during a review phase for $\delta \geq \underline{\delta}$, these strategies form a Nash equilibrium with payoffs $g(a)$. It is equally clear that they are not generally sequential (subgame perfect), because players may profitably deviate during punishment phases.

In proving the folk theorem, Fudenberg-Maskin (1986a) construct strategies that induce the players to carry out punishments by providing "rewards" for doing so. But such rewards are not provided by our review strategies. Indeed, the Fudenberg-Maskin construction is only possible if the game satisfies a "full-dimensionality" condition that we have not imposed. Thus the conclusion of (A) cannot in general be strengthened to $\epsilon^n = 0$.

The review strategies do, however, form an $\epsilon(\delta)$ -sequential equilibrium with $\epsilon(\delta) \rightarrow 0$ as $\delta \rightarrow 1$. (Note that with observable actions, sequential equilibrium and subgame-perfect equilibrium are equivalent.) To show this, we need only calculate the greatest gain to deviating during a punishment phase. Recall that \bar{d} is the greatest one-shot gain to any player of deviating from any profile. Consequently, the gain is at most $(1-\delta)K\bar{d}$, which clearly goes to zero as $\delta \rightarrow 1$. The point is that the punishment lasts only K periods, and players never expect to engage in punishment again. Consequently, with extreme patience, the cost of the punishment is negligible. In particular, with time-averaging this is an exact equilibrium.

Proof of (A): Consider the strategy $\hat{\sigma}^n$ of playing σ^n from 1 to T^n , then starting over and playing σ^n from $T^n + 1$ to $2T^n$, and so forth.

We call period 1 to T^n "round 1", $T^n + 1$ to $2T^n$ "round 2" and so on. Fix $\delta < 1$ and any beliefs consistent with $\hat{\sigma}^n$. Let a history $(h_1(t), \dots, h_n(t))$ be given. Let us calculate an upper bound on any player's gain to deviating in time- t normalized present value. Fix a time t and a history $(h_1(t), \dots, h_n(t))$, and choose k so that $kT^n < t \leq (k+1)T^n$, so that t is in round k . Since the maximum of per-period gain to deviating is \bar{d} , we know that regardless of the history, no player by deviating in round k , can gain more than $(1-\delta)T^n \bar{d}$. Further, play from time $(k+1)T^n + 1$ on is independent of what happened during the previous rounds. Since round $k+1$ is a truncated time-average ϵ^n -Nash equilibrium, at most $(1-\delta)T^n \epsilon^n$ can be gained by deviating in this round. During round $k+2$ at most $\delta(1-\delta)T^n \epsilon^n$ can be gained, and so forth. Consequently, no player can gain more than $(1-\delta)T^n (\bar{d} + \epsilon^n / (1-\delta))$. As $\delta \rightarrow 1$, this quantity approaches ϵ^n by l'Hôpital's rule. Clearly, then, we may choose δ^n close enough to 1 that $\hat{\sigma}^n$ is a $2\epsilon^n$ -sequential equilibrium. In addition, because of the repeated structure of $\hat{\sigma}^n$, it is clear that $W_i(\hat{\sigma}^n, \delta^n) \rightarrow W_i^{T^n}(\sigma^n)$ as $\delta^n \rightarrow 1$ so $W_i(\hat{\sigma}^n, \delta^n) \rightarrow v_i$. ■

Proof of (B): Consider the strategy σ of playing σ^1 from 1 to T^1 , σ^2 from $T^1 + 1$ to $T^1 + T^2$, and so forth. Fix any beliefs consistent with σ . Let a history $(h_1(t), \dots, h_n(t))$ be given and let t be in round k . By deviating in the current round, player i can gain at most \bar{d} per period as play in the future is independent of play in this round. The per-period gain at $k+1$ is at most ϵ^{k+1} , in round $k+2$, ϵ^{k+2} . Since $\epsilon^k \rightarrow 0$, this implies that the time average gain is zero. Since $W_i^{T^k}(\sigma^k) \rightarrow v_i$, σ is a uniform sequential equilibrium, and this yields the desired conclusion. ■

Each of the two conclusions of Theorem 3.1 has strengths and weaknesses. The limit of discounting is more appealing than time-averaging, which suggests the interpretation (A), but exact equilibrium is more appealing than approximate equilibrium, which suggests the interpretation (B). In this context, it is worth emphasizing reiterating that the conclusion of (A) cannot be strengthened to $\epsilon^n = 0$. One interpretation of this fact is that the set of time-average equilibria includes not only the limits of exact discounted equilibria, but the limits of approximate discounted equilibria as well.

If we weaken definition (2.6) to allow the gain to deviating, ρ , to depend on T and t , we have a time-average sequential equilibrium, rather than a uniform one. That these equilibria include all limits of discounted sequential equilibria as shown by

Proposition 3.2: If σ is sequentially regular, and is an $\epsilon(\delta)$ -sequential equilibrium for discount factor δ with $\epsilon(\delta) \rightarrow 0$, then σ is a time-average sequential equilibrium.

Proof: This effectively follows from the fact that the limit points, as $\delta \rightarrow 1$, of the normalized present value of a sequence of payoffs are contained in the set of limit points of the finite time averages. This implies that if σ is sequentially regular, then the discounted value of payoffs along the continuation equilibrium approaches the value with time averaging. Deviations from the continuation payoff when evaluated with the discounting criterion yield values whose limits are no greater than the lim sup of the time average. ■

It is worth noting that the converse of this proposition is not true. An example of a time-average equilibrium that is not even an approximate equilibrium with discounting can be found in the following "guru" game:

Player 1 is the guru. He chooses one of three actions 0, 1, or 2, each period. Player 2 is a disciple. He also chooses one of two actions. Each period the guru receives zero, while the disciple receives an amount equal to the action chosen by the guru. In other words, the guru is completely indifferent, while the disciple cares only about what the guru does. Notice that sequentiality reduces to subgame perfection in this example, and that subgame perfection has no force. Any strategy by the guru is optimal in any subgame, while any Nash equilibrium need only be adjusted for the disciple so that he chooses an optimum in subgames not actually reached with positive probability in equilibrium. (There is a trivial complication in that the guru has strategies such that no optimum for the disciple exists.)

Consider the following equilibrium with time averaging. As long as the disciple never played 2 in the past, the guru plays 1. Let t denote the first period in which the disciple plays 2. Then the guru plays 2 in period $(t+1)$ to $2t$, plays 0, in periods $2t+1, \dots, 3t+1$, and in period $3t+2$ the guru reverts back to 1 forever. Regardless of the history the disciple always plays 1. With time averaging the disciple is completely indifferent between all his deviations, since they all yield a time average of 1. Similarly the guru is clearly indifferent between all his actions. Consequently this is a time-average equilibrium. (In fact, it is also an equilibrium when the players' time preference is represented by the overtaking criterion.) On the other hand, with discount factor $\delta \geq 1/2$, we can always find a time t such that δ^t is between $1/2$ and $1/4$, and by deviating at such a time player 2 always gains at least $1/32$. Consequently this configuration is only an ϵ -equilibrium for $\epsilon \geq 1/32$, and

in particular, ϵ does not converge to zero as δ goes to 1.¹

We would argue that time-averaging is of interest only insofar as it captures some sort of limit of discounting. Consequently, we argue that this equilibrium of the guru game does not make good economic sense. We do not know whether Proposition 3.2 or its converse holds for uniform time-average equilibrium. If so, it is a strong argument in favor of restricting attention to uniform equilibria.

4. A Folk Theorem

We now restrict attention to a limited, but important class of games, which we call informationally connected games. All two-player games are informationally connected. If $N > 2$, we say that player i is directly connected to player $j \neq i$ despite player $k \neq i, j$ if there exists a mixed action profile α and a mixed action $\hat{\alpha}_i$ for player i such that $\pi_i(\cdot, \hat{\alpha}_i, \alpha'_k, \alpha_{-i-k}) \neq \pi_j(\cdot, \alpha)$ regardless of the play α'_k of player k . In other words, at α , player i has a deviation that can be potentially detected by player j regardless of how player k plays. We say that player i is connected to player j if for every player $k \neq i, j$, there exists a sequence of players i_1, \dots, i_n with $i_1 = i, i_n = j$ and $i_p \neq k$ for any p , and such that player i_p is directly connected to player i_{p+1} despite player k . In other words, a message can always be passed from i to j , regardless of which single other player might try to interfere. A game is informationally connected if every player is connected to every

¹Note that the counterexample is to the strategy profile σ not being an approximate equilibrium with discounting. From the folk theorem, we know that there are other strategies that are exact equilibria and yield approximately the same payoff. This is not true in games with a Markov structure and absorbing states, as shown by the example of Sorin [1986]. In fact, his example has equilibrium payoff with time averaging that are not ever approximately equilibrium payoffs with discounting.

other player. For an example of such a game, suppose that each player determines an output level of zero or one. "Average" price is a strictly decreasing function of total output, and each player observes an "own price" that is an independent random function of average price, with the property that the distribution corresponding to a strictly higher average price strictly stochastically dominates that corresponding to the lower price. In this case, all players can choose zero output levels, and if any player produces a unit of output, this changes the distribution of prices for all other players, with a deviation by a second player merely enhancing the signal by lowering the distribution of prices still further.

A mutual threat point v is a payoff vector for which there exists a mutual punishment action α such that $g_i(\alpha'_i, \alpha_{-i}) \leq v_i$ for all players i and mixed actions α'_i . With three or more players vector v^* of minmax values need not be a mutual threat point, as there may not be a single action profile that simultaneously holds all of the players to their minmax values. When such a profile exists, the game is said to satisfy the "mutual minmax property" (Fudenberg-Maskin [1986a]). One example where this property obtains is a repeated quantity-setting oligopoly with capacity constraints where the profile "all players produce to capacity" serves to minmax all of the players.

A payoff vector that weakly Pareto-dominates a mutual threat point is called mutually punishable; the closure of the convex hull of such payoffs is the mutually punishable set.

We say that a payoff vector v is (independently) enforceable if there exists a mixed action profile α with $g(\alpha) = v$, and such that if for some player i and mixed action α'_i , $g_i(\alpha'_i, \alpha_{-i}) > v_i$, then for some other player $j \neq i$, $\pi_j(\cdot, \alpha'_i, \alpha_{-i}) \neq \pi_j(\cdot, \alpha)$. In other words, any improving

deviation for player i can potentially be detected by some other player j . The enforceable set \tilde{V} is the closure of the convex hull of the enforceable payoffs. Notice that in addition to the static Nash equilibrium, which is clearly enforceable, every extremal Pareto efficient payoff is. This is because any extremal payoff is generated by a pure action profile, and any efficient pure action profile is enforceable: if a action profile is not enforceable, one player can strictly improve himself without anyone else knowing or caring. Note also that if the unconditional play in any Nash equilibrium at time t is a mixed (rather than correlated) action profile, it is clear that the profile must be enforceable (except in the infinite time average case, for a negligible fraction of periods). However, it is possible that unenforceable payoffs can be achieved through correlation.

Finally, we define the enforceable mutually punishable set $V^* = \underline{V} \cap \tilde{V}$, which is closed, convex, and contains at least the convex hull of static Nash equilibrium payoffs. We can now prove:

Theorem 4.1: (Folk Theorem) In an informationally connected game, if $v \in V^*$, there exists a sequence of times T^n , of non-negative numbers $\epsilon^n \rightarrow 0$ and strategy profiles σ^n such that σ^n is a T^n -truncated time-average ϵ^n -Nash equilibrium and $W_i^{T^n}(\sigma^n) \rightarrow v_i$.

Remark: This is a partial folk theorem in that, as remarked earlier, the mutually enforceable payoffs may be a strictly smaller set than the individually rational socially feasible ones. Moreover, in games with imperfect observation of the opponents' actions and three or more players, even the individually rational payoffs may not be a lower bound on the set of payoffs that can arise as equilibria (see Fudenberg-Levine-Maskin for an example). Finally, we do not know whether the set of equilibria can include payoffs

not in V^* .

The proof constructs strategies that have three stages. In review stages, players play to obtain the target payoff v . At the end of review stages players "test" to see if their own payoff was sufficiently close to the target level. Then follows a "communication stage" where players use their actions to "communicate" whether or not the review was passed; the assumption of information connectedness is used to ensure that such communication is possible. Finally, if players learn that the test was failed, they revert to a "punishment state" for the remainder of the game.

The idea of using strategies with reviews and punishments was introduced by Radner [1981, 1986] who did not need the communication phase because he studied models with publicly observed outcomes. Lehrer [1986] introduces the idea of a communications phase to coordinate play. Our proof is in some ways simpler than Radner's as we establish only the existence of ϵ -equilibria of the finite-horizon games and the appeal to Proposition 3.1. Lehrer's [1986] proof is more complex than ours because he obtains a larger set of payoffs.

The key to proving Theorem 4.1 is a lemma proven in the Appendix:

Lemma 4.2: If $\bar{\alpha}^h$ is enforceable, then for $j \neq i$ there exists m_j -vectors of weights λ_j such that for all α_i , $\sum_{j \neq i} \lambda_j \pi_j(\cdot, \alpha_i, \bar{\alpha}_{-i}^h) \geq g_i(\alpha_i, \bar{\alpha}_{-i}^h) + \eta$ and $\sum_{j \neq i} \lambda_j \pi_j(\cdot, \bar{\alpha}^h) = g_i(\bar{\alpha}^h) + \eta$.

This shows how we can reduce the problem of detecting profitable deviations by player i to a one-dimensional linear test. We need only deter α_i 's that lead to the information vectors $\bar{\pi}_j(\cdot, \alpha_i, \bar{\alpha}_{-i}^h)$ that lie above the half-space whose existence is asserted in the Lemma. This makes the problem very similar to one with publicly observed signals.

Proof of Theorem 4.1: Fix $v \in V^*$, and suppose $\epsilon > 0$ is given. We show how to find T and σ such that σ is a T -truncated time-average 100ϵ -Nash equilibrium, and $|W_1^T(\sigma) - v_1| \leq 100\epsilon$. This clearly suffices. The proof proceeds in several steps. First we construct a class of strategy profiles s that depends on a vector of constants L determined by the game and v , and constants ℓ that are free parameters. The constants L will index the relative lengths of various phases of play, and the ℓ 's will determine both the absolute length of the phases and the number of times the phases are repeated. The length of the game T is implicitly determined by these constants. We then show how to choose the constants ℓ to make s a 100ϵ -equilibrium yielding payoffs within 100ϵ of v .

Step 1 (Payoffs): Since $v \in V^*$, $v \in \tilde{V}$ and $v \in \underline{V}$. This means we can find finite L' -dimensional vectors of non-negative coefficients $\mu^h, \underline{\mu}^h$ summing to 1, and of payoffs $\bar{v}^h, \underline{v}^h$, with $\sum_{h=1}^{L'} \mu^h \bar{v}^h = v$, $\sum_{h=1}^{L'} \underline{\mu}^h \underline{v}^h = v$ such that \bar{v}^h is enforceable and \underline{v}^h is mutually punishable. Corresponding to these are mixed action profiles $\bar{\alpha}^h$ and $\underline{\alpha}^h$, where $\bar{\alpha}^h$ yields payoff \bar{v}^h and $\underline{\alpha}^h$ is a mutual punishment that enforces \underline{v}^h . Moreover, it is clear that we can find L^R and non-negative L' -dimensional vectors of integers $\bar{L}^h, \underline{L}^h$ that sum to L^R and such that $|\sum_{h=1}^{L'} (\bar{L}^h/L^R) \bar{v}^h - v| \leq \epsilon$ and $|\sum_{h=1}^{L'} (\underline{L}^h/L^R) \underline{v}^h - v| \leq \epsilon$.

Step 2 (Temporal Structure): We first describe the temporal structure of the game, which consists of ℓ' repetitions of a review stage followed by a communications stage. Set $L^C = N(N-2)(N-1)!$. A review stage lasts $\ell^R L^R$ periods, and a communications stage lasts $\ell^C L^C$ periods. Each stage is further subdivided into phases; a review stage into L^R phases indexed by $h = 1, \dots, L^R$, and a communications stage into L^C phases, indexed as

described below. The temporal structure is outlined in Figure 3.1; the game lasts $T = \ell'(\ell^R_L + \ell^C_L)$ periods.

The length of a phase of a review stage depend on whether it is regarded as a reward or punishment stage; if it is a reward stage, the h^{th} phase lasts ℓ^{R-h}_L periods; if it is a punishment stage it lasts ℓ^{R-h}_L periods. Each phase of a communications stage lasts ℓ^C periods.

Each communications phase is assigned an index (i,j,k) corresponding to a triple of players. (Each index generally occurs more than once.) The first $(N-2)(N-1)!$ phases have indices $(i,j,1)$; the next $(N-2)(N-1)!$ have indices $(i,j,2)$, and so forth. Fixing k , the third index, the (i,j) indices are determined by taking the set of all players but k , and calculating every permutation of the set, giving rise to $(N-1)!$ blocks of $(N-2)$ periods. Fix a permutation, i_1, i_2, \dots, i_{N-1} . Then the first index in the block is (i_1, i_2, k) ; the second (i_2, i_3, k) and so forth up to (i_{n-2}, i_{n-1}, k) . The point of this rather complex structure is that because the game is informationally connected, each player has an opportunity to send a signal to all other players, without being blocked by any other single player.

Using the fact that the game is informationally connected, we associate distinct triples (i,j,k) with certain mixed strategies. If i is connected to j despite k , we say that (i,j,k) is an active link, and let α^{ijk} be the action allowing communication and $\hat{\alpha}_i^{ijk}$ the deviation for player i that enables him to communicate with j . If i is not connected to j despite k we say that (i,j,k) is an inactive link, and arbitrarily define α^{ijk} to be $\bar{\alpha}^{-1}$, and $\hat{\alpha}_i^{ijk}$ to be $\bar{\alpha}_i^{-1}$.

Step 3 (Strategies): We now describe player i 's strategy. He may be in one of two states, a reward state or a punishment state. He begins the game

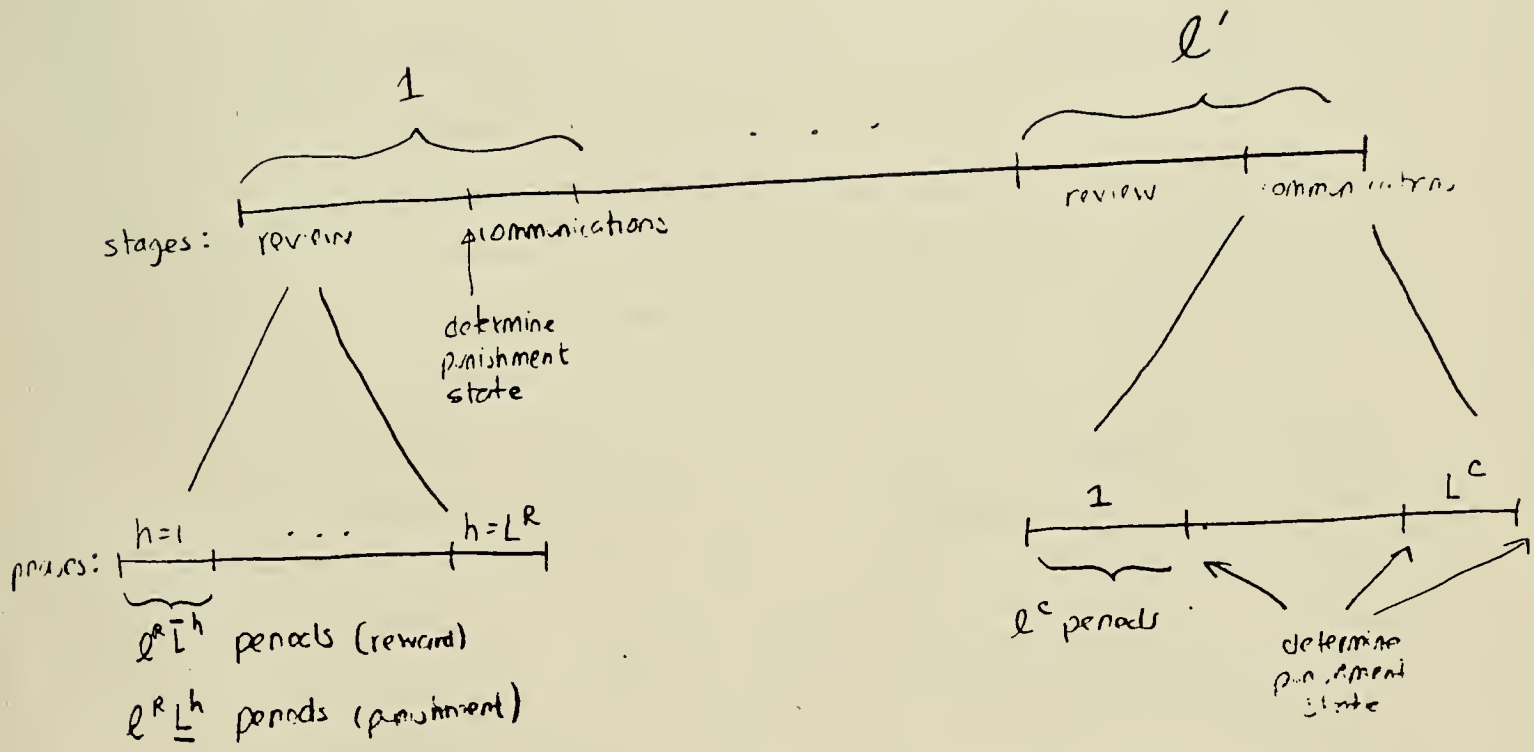


Figure 3.1: Temporal Structure

in the reward state. The punishment state is absorbing, so once reached, player i remains there forever. We must describe how a player plays in each state, and when he moves from the reward to punishment state.

In the reward state and review stage, player i regards the stage as divided into L' phases, with phase h lasting $\ell_{L'}^{R,h}$ periods. During the h^{th} phase he plays α_i^{-h} . In the reward state and communications stage, in the phase indexed by (i',j,k) he plays $\alpha_i^{i'jk}$.

In the punishment state and review stage, player i regards the stage as subdivided into L' phases, with phase h lasting $\ell_{L'}^{R,h}$ periods. During the h^{th} phase he plays α_i^h . In the punishment state and communications stage, in the phase indexed by (i',j,k) he plays $\alpha_i^{i'jk}$ if $i \neq i'$, and he plays $\hat{\alpha}_i^{ijk}$ if $i = i'$.

Player i can change states only at the end of a review stage, or at the end of a communications phase (j,i,k) that is an active link. The transition is determined by a parameter ϵ' .

At the end of each reward phase of the review stage player i calculates $\bar{\pi}_i^{-h}(z_i)$ to be the fraction of the $\ell_{L'}^{R,h}$ periods in which z_i occurred. If at the end of the review stage $|\bar{\pi}_i^{-h}(\cdot) - \pi_i(\cdot, \alpha_i^{-h})| > \epsilon'$ for any h , player i switches to the punishment state.

At the end of the p^{th} communications phase corresponding to a triple (j,i,k) that is an active link, player i calculates $\bar{\pi}_i^P(z_i)$ to be the fraction of the ℓ^C periods in which z_i occurred. If $|\bar{\pi}_i^P(\cdot) - \pi_i(\cdot, \alpha^{jik})| > \epsilon'$, player i switches to the punishment state.

In all other cases player i 's state does not change; in particular the punishment state is absorbing.

We must show how to choose ℓ' , ℓ^R , ℓ^C and ϵ' to give expected payoffs within 100ϵ of v , and no greater than an 100ϵ gain to any

player from deviating.

Step 4 (Payoffs): First we consider the equilibrium payoffs. Let $\bar{\pi}'$ denote the probability that no player ever enters the punishment state and let \bar{v}' be the expected payoff vector conditional on this. Let $\bar{v} = \sum_{h=1}^{L'} (\bar{L}^h/L^R) \bar{v}^h$. The payoff \bar{v}' differs from \bar{v} solely due to communications phases, yielding

$$(4.1) \quad |\bar{v}' - \bar{v}| \leq \frac{\ell^C_L \ell^C \bar{d}}{\ell^R_L \ell^R + \ell^C_L \ell^C} = \frac{\bar{d}}{1 + (\ell^R/\ell^C)(L^R/L^C)}$$

where recall that \bar{d} is the greatest difference between any payoffs. Since $|\bar{v} - v| \leq \epsilon$ by construction, we conclude

$$(4.2) \quad |W_i^T(\sigma) - v_i| \leq \epsilon + \bar{d} \left[\frac{1}{1 + (\ell^R/\ell^C)(L^R/L^C)} + (1 - \bar{\pi}') \right].$$

This shows we must take ℓ^R/ℓ^C very large and $\bar{\pi}'$ close to one. We will show that the weak law of large numbers implies that $\bar{\pi}'$ close to one if ℓ^R and ℓ^C are both sufficiently large (relative to ϵ').

Step 5 (Gain to Deviating): Next, we consider how much a player might gain by deviating. During a communications stage (or in any period for that matter) the greatest per period gain is \bar{d} . However, only $\ell^C_L \ell^C / (\ell^R_L \ell^R + \ell^C_L \ell^C)$ of the periods in the game lie in communications phases so we have the greatest per period gain over the whole game due to deviations during communications is

$$(4.3) \quad [\text{communications}] \frac{\bar{d}}{1 + (\ell^R/\ell^C)(L^R/L^C)}.$$

Next we consider review stages. Fixing a deviator, there are three possibilities: all other players are in the punishment state; all are in the review state, or they disagree on the state. Of course a player

contemplating deviating may not be certain which of these is true, but since he can only benefit from this extra information, we may suppose he knows which case he faces.

If all other players are in the punishment state they will remain there regardless of player i 's play. Regardless of how he plays during a punishment phase, player i gets at most \underline{v}_i^h . Since $|\Sigma(L^h/L^R)\underline{v}_i^h - v_i| \leq \epsilon$, the largest gain over $W_i^T(\sigma)$ that player i can obtain by deviating through the punishment state is

$$(4.4) \quad (\text{punishment}) \quad \epsilon + v_i - W_i^T(\sigma) \leq 2\epsilon + \bar{d} \left[\frac{1}{1+(\ell^R/\ell^C)(L^R/L^C)} + (1-\bar{\pi}') \right]$$

where the inequality follows from (4.2).

Next, suppose that other players disagree about the state (so in particular, $N > 2$). Then player i can possibly get \bar{d} per period. Let $\hat{\pi}$ be a lower bound on the probability that all opponents agree on a punishment state at the end of the subsequent communications stage (and therefore agree for the rest of the game). Then the deviating player can at best hope to get \bar{d} for one stage with probability $\hat{\pi}$, and \bar{d} in all periods with probability $(1-\hat{\pi})$. Since there are ℓ' review phases, \bar{d} per period for one phase is actually worth \bar{d}/ℓ' yielding

$$(4.5) \quad [\text{confusion}] \quad \frac{\hat{\pi}\bar{d}}{\ell'} + (1-\hat{\pi})\bar{d}.$$

This leaves review stages where all opponents are in the reward state. Let $1-\pi'$ be an upper bound on the probability that player i both gains more than 2ϵ and no opponent enters the punishment state at the end of the stage. Then player i can at best gain 2ϵ in all periods, and can gain

more than 2ϵ and remain with all opponents in the reward state next stage with probability no more than $1-\pi'$. Since at best he can gain \bar{d} , player i gains at most

$$(4.6) \quad [\text{reward}] \quad 2\epsilon + \frac{\pi'\bar{d}}{\ell'} + (1-\pi')\bar{d}.$$

in reward stages.

Adding these bounds, we find an upper bound on the gain to deviating

$$(4.7) \quad [\text{total}] \quad \frac{2\bar{d}}{1+(\ell^R/\ell^C)L^R/L^C} + 4\epsilon + \frac{\hat{\pi}\bar{d}}{\ell'} + (1-\hat{\pi})\bar{d} + \frac{\pi'\bar{d}}{\ell'} + \bar{d}(1-\pi') + \bar{d}(1-\bar{\pi}').$$

Our goal is to find $\ell', \ell^R, \ell^C, \epsilon'$ to simultaneously make (4.2) and (4.6) smaller than 100ϵ .

Step 6: Fix ℓ' so that $\bar{d}/\ell' < \epsilon$. Let $\ell^C(\ell^R)$ be the largest integer smaller than $\gamma\ell^R$ where

$$\frac{\bar{d}}{1+(1/\gamma)(L^R/L^C)} = \epsilon.$$

We may then simplify (4.2) to

$$(4.2') \quad |w_i^T(\sigma) - v_i| \leq 2\epsilon + (1-\bar{\pi}')\bar{d},$$

and (4.7) to

$$(4.7') \quad 8\epsilon + (1-\hat{\pi})\bar{d} + (1-\pi')\bar{d} + (1-\bar{\pi}')\bar{d}.$$

Consequently, it suffices to show that we can choose ℓ^R and ϵ' so that when $\ell^C = \ell^C(\ell^R)$, $(1-\bar{\pi}')$, $(1-\bar{\pi})$, $(1-\pi')$ are all less than or equal to ϵ/\bar{d} .

By Lemma A.1 in the Appendix, for the h^{th} reward phase there exists ϵ^h and ℓ^h such that for all $\epsilon' \leq \epsilon^h$ and $\ell^R \geq \ell^h$ the probability that either i gains less than or equal 2ϵ or at least one opponent switches

to punishment at the end of the stage is at least $(1-\epsilon/\bar{d})^{1/L^R}$.

Consequently, for $\epsilon' \leq \min \epsilon^h$, $\ell^R \geq \max \ell^h$, we have $\pi' \geq 1 - \epsilon/\bar{d}$.

By Lemma A.2 in the Appendix, for the (j,k,i) that are active communications links there exist ϵ^{jki} and ℓ^{jki} such that for all $\epsilon' \leq \epsilon^{jki}$ and $\ell^C \geq \ell^{jki}$ if player j plays \hat{a}_j^{jki} and players $k' \neq i$ play $\alpha_{k'}^{jki}$ the probability of k switching to punishment at the end of the phase is at least $(1-\epsilon/\bar{d})^{1/L^C}$. Consequently, for $\epsilon' \leq \min \epsilon^{ijk}$, $\ell^C \geq \max \ell^{ijk}$, we have $\hat{\pi} \geq 1 - \epsilon/\bar{d}$. Note that $\ell^C \geq \max \ell^{ijk}$, provided $\ell^C = \ell^C(\ell^R)$ and $\ell^R \geq (1+\gamma)\ell^{ijk}/\gamma$.

Fix then $\epsilon' = \min\{\epsilon^h, \epsilon^{ijk}\}$, and $\ell^* = \max\{\ell^h, (1+\gamma)\ell^{ijk}/\gamma\}$. By the weak law of large numbers, there exists an ℓ^{**} such that for $\ell' \geq \ell^{**}$, $\bar{\pi}' \geq 1 - \epsilon/\bar{d}$. Choosing $\ell' = \max\{\ell^*, \ell^{**}\}$ then completes the proof. ■

Remark: How frequently does punishment occur? As $T \rightarrow \infty$, the proof shows that the probability of punishment goes to zero. Examining the proof of Theorem 3.1, we see that infinite equilibria are constructed by splicing together a sequence of truncated equilibria. For approximate discounted equilibria, the construction repeats the same equilibrium over and over, so that the probability of punishment in each round is constant. By the zero-one law, this means punishment occurs infinitely often.

In the case of time average equilibria, it is irrelevant whether punishment occurs infinitely often or not. The truncated equilibria that are spliced together have decreasing probability of punishment, say $\pi_1, \pi_2, \dots \rightarrow 0$. By the zero-one law, if $\sum_{t=1}^{\infty} \pi_t = \infty$, punishment occurs infinitely often, while if $\sum_{t=1}^{\infty} \pi_t < \infty$, punishment almost surely stops in finite time. However, we can choose a subsequence $\{\pi'_t\}$ such that

$\sum_{t=1}^{\infty} \pi'_t < \infty$. The corresponding truncated equilibria when spliced together form a time average equilibrium in which punishment almost surely ceases. On the other hand, we may form a sequence in which the t^{th} truncated equilibrium is repeated $1/\pi_t$ times. Splicing this sequence together clearly yields a time-average equilibrium with $\sum_{t=1}^{\infty} \pi'_t = \infty$, and so punishment occurs infinite often.

This discussion should be contrasted with the results of Radner [1981] in which punishment is infinitely often, and Fudenberg, Kreps and Maskin [1986a], in which it almost surely ceases.

APPENDIX

Proof of Lemma 4.2: We may regard $\sum_{j \neq i} M_j + 1$ dimensional Euclidean space as having components π_j , $j \neq i$ and g_i . Define F to be the subset of this space such that there exists a mixed strategy μ_i for player i with $\pi_j = \pi_j(\cdot, \alpha_i, \bar{\alpha}_i^h)$ and $g_i \leq g_i(\alpha_i, \bar{\alpha}_i^h)$. Let $\bar{\pi}$ be the $\sum_{j \neq i} M_j$ vector with components $\pi_j(\cdot, \bar{\alpha}^h)$, and let $\bar{g}_i = g_i(\bar{\alpha}^h)$. Then $(\bar{\pi}, \bar{g}_i) \in F$ and for $\Delta > 0$, enforceability implies $(\bar{\pi}, \bar{g}_i + \Delta) \notin F$. Moreover, F is a convex polyhedral set and its extreme points correspond to pure strategies for player i .

Since F is convex polyhedral, we may characterize $(\pi, g_i) \in F$ by the linear inequalities $A(\pi - \bar{\pi}) + b(g_i - \bar{g}_i) \geq c$, where A is a matrix, b and c are vectors. Consider minimizing $\lambda(\pi - \bar{\pi}) - (g_i - \bar{g}_i)$ subject to this constraint. Suppose this problem has a solution equal to zero. Then $\lambda \pi \geq g_i + \eta$ where $\eta = \lambda \bar{\pi} - \bar{g}_i$, and $\lambda \bar{\pi} = \bar{g}_i + \eta$, the desired conclusion.

Since $(\bar{\pi}, \bar{g}_i) \in F$, a feasible solution to the primal exists, so a minimum of zero exists, if and only if the dual has a feasible plan yielding zero. The dual is to maximize ωc subject to

$$(\omega A, \omega b) = (\lambda, -1)$$

$$\omega \geq 0.$$

In other words, if we can find $\omega \geq 0$, $\omega c = 0$, $\omega b = -1$, then $\lambda = \omega A$ is the desired solution.

Recall that $(\bar{\pi}, \bar{g}_i) \in F$, so $0 \geq c$. Moreover, if $A(\pi - \bar{\pi}) + b(g_i - \bar{g}_i) \geq c$ and $g'_i \geq g_i$, $A(\pi - \bar{\pi}) + b(g'_i - \bar{g}_i) \geq c$ implying that $0 \geq b$. Finally, for $\Delta > 0$, $b\Delta \not\geq c$. We conclude for some component p , $c_p = 0$, $b_p < 0$. Choose then $\omega_q = 0$ $q \neq p$ and $\omega_p = 1$. This is the desired solution. \blacksquare

Lemma A.1: Suppose players $j \neq i$ are in the reward state and follow their equilibrium strategies in review phase h . For every $\beta > 0$ there exists an ϵ^h and ℓ^h such that for $\epsilon' \leq \epsilon^h$ and $\ell^R \geq \ell^h$,

$$\text{Prob}(\exists j \neq i \mid \bar{\pi}_j^h(\cdot) - \pi_j(\cdot, \alpha^{-h}) \mid \leq \epsilon' \text{ and } (1/\bar{L}^h \ell^R) \sum_{t=1}^{\bar{L}^h \ell^R} r_i(t) > g_i(\bar{\alpha}^h) + 2\epsilon) \leq \beta$$

regardless of the strategy used by player i . (Recall that $r_i(t) = r_i(z_i(t))$ is player i 's realized payoff in period t .)

Proof: From Lemma 4.2 there exist for $j \neq i$ m_j -vectors of weights λ_j and a scalar η such that $\sum_{j \neq i} \lambda_j \pi_j(\cdot, \alpha_i, \bar{\alpha}_{-i}^h) \geq g_i(\alpha_i, \bar{\alpha}_{-i}^h) + \eta$ and $\sum_{j \neq i} \lambda_j \pi_j(\cdot, \bar{\alpha}^h) = g_i(\bar{\alpha}^h) + \eta$. Consider then the random variable $x(t) = \sum_{j \neq i} \lambda_j(z_j(t))$, where $\lambda_j(z_j(t))$ is the component of $z_j(t)$ corresponding to z_j . The idea of this construction is to use a single random variable $x(t)$ to summarize the information received by all other players. Let

$$\sum_{j \neq i} \lambda_j \bar{\pi}_j^h(\cdot) = (1/\bar{L}^h \ell^R) \sum_{t=1}^{\bar{L}^h \ell^R} x(t) = \bar{x}.$$

It is clear then that there exists ϵ^h such that if $\bar{x} > g_i(\bar{\alpha}^h) + \eta + \epsilon$ then $|\bar{\pi}_j^h - \pi_j(\cdot, \bar{\alpha}^h)| > \epsilon^h$ for some j . That is, if the sample mean \bar{x} is far from its theoretical distribution under $\bar{\alpha}_j^n$ then some player $j \neq i$ will observe that his empirical information $\bar{\pi}_j^h$ is far from what it would be if player i had played $\bar{\alpha}^n$. Consequently, it suffices to show with

$$\bar{\rho}_i = \sum_{t=1}^{\bar{L}^h \ell^R} r_i(t) \text{ that}$$

$$\text{Prob}(\bar{x} \leq g_i(\bar{\alpha}^h) + \eta + \epsilon \text{ and } \bar{\rho}_i > g_i(\bar{\alpha}^h) + 2\epsilon) \leq 1 - \beta.$$

Fix a strategy for player i and consider

$$\tilde{x}(t) \equiv x(t) - E[x(t) | x(t-1), \dots, r(t-1), \dots].$$

These are uncorrelated random variables with zero mean and are bounded independent of the particular strategy of player i . Similarly considerations apply to $\tilde{r}(t) = r(t) - E[r(t) | x(t-1), \dots, r(t-1), \dots]$. Consequently the weak law of large numbers shows that as $\ell^R \rightarrow \infty$, $\bar{\tilde{x}}, \bar{\tilde{r}} \rightarrow 0$ in probability uniformly over strategies for player i . Let $\alpha_i(t)$ be player i 's mixed action at t conditional on $x(t-1), \dots, r(t-1), \dots$. Then

$$\begin{aligned} E[x(t) | x(t-1), \dots, r(t-1), \dots] &= \sum_{j \neq i} \lambda_j \pi_j(\cdot, \alpha_i(t), \bar{\alpha}_{-i}^h) \geq g_i(\alpha_i(t), \bar{\alpha}_{-i}^h) + \eta \\ &= E[r(t) | x(t-1), \dots, r(t-1), \dots] + \eta. \end{aligned}$$

Consequently $\text{Prob}(\bar{x} - \bar{r} - \eta < -\epsilon) \rightarrow 0$. Since $\bar{x} \leq g_i(\bar{\alpha}_{-i}^h) + \eta + \epsilon$ and $\bar{r} > g_i(\bar{\alpha}_{-i}^h) + 2\epsilon$ implies $\bar{x} - \bar{r} < \eta - \epsilon$, this gives the desired conclusion. ■

Lemma A.2: Suppose player j is in the punishment state, and that (j, k, i) is an active communication phase. If all players $k' \neq i$ follow their equilibrium strategies then for every $1 > \beta$ there exists an ϵ^{jki} and ℓ^{jki} such that for $\epsilon' \leq \epsilon^{jki}$ and $\ell^R \geq \ell^{jki}$

$$\text{Prob}(|\bar{\pi}_k^{jki}(\cdot) - \pi_k(\cdot, \alpha^{jki})| > \epsilon') \geq \beta.$$

Proof: This is similar to, but simpler than Lemma A.1. Observe that the set of vectors $\pi_k = \pi_k(\cdot, \hat{\alpha}^{jki}, \alpha_i, \alpha_{-j-i}^{jki})$ for different α_i is compact convex, and by assumption does not contain $\pi_k(\cdot, \alpha^{jki})$. Consequently, we may find a m_k -vector of weights λ and a scalar $\eta > 0$ such that $\lambda \pi_k \geq \eta$ and $\lambda \pi_k(\cdot, \alpha^{jki}) = 0$. Set $x(t) = \lambda(z_k(t))$. Again, $\lambda \bar{\pi}_j^{jik}(\cdot) = \bar{x}$, so there exists ϵ^{jki} such that if $\bar{x} > \eta/2$, $|\bar{\pi}_k^{jik} - \pi_j(\cdot, \bar{\alpha}^h)| > \epsilon^{jki}$. Again $\tilde{x}(t) = x(t) - E[x(t) | x(t-1), \dots]$ has $\bar{\tilde{x}}$ converging uniformly to zero in


probability. Since $E[x(t)|x(t-1), \dots] = \lambda \pi_k(\cdot, \hat{\alpha}^{jki}, \alpha_i(t), \alpha_{-j-i}^{jki}) \geq \eta$, we get the desired conclusion. ■

REFERENCES

- Abreu, D., D. Pearce and E. Stacchetti (1987), "Toward a Theory of Discounted Repeated Games With Imperfect Monitoring," unpublished ms., Harvard.
- _____, (1986), "Optimal Cartel Monitoring With Imperfect Information," Journal of Economic Theory, 39, 251-69.
- Friedman, J. (1972), "A Non-Cooperative Equilibrium for Supergames," Review of Economic Studies, 38, 1-12.
- Fudenberg, D., D. Kreps and E. Maskin (1989), "Repeated Games With Long Run and Short Run Players," mimeo, MIT.
- Fudenberg, D., D. Levine, and E. Maskin (1989), "The Folk Theorem With Imperfect Public Information," mimeo.
- Fudenberg, D., and E. Maskin (1986a), "Folk Theorems for Repeated Games With Discounting and Incomplete Information," Econometrica, 54, 533-54.
- _____, (1986b), "Discounted Repeated Games With One-Sided Moral Hazard," unpublished ms., MIT.
- _____, (1989), "The Dispensability of Public Randomizations in Discounted Repeated Games," unpublished ms., MIT.
- Fudenberg, D., and D. Levine [1983], "Subgame Perfect Equilibria of Finite- and Infinite-Horizon Games," Journal of Economic Theory, 31 (December), 251-63.
- Green, E., and R. Porter (1984), "Noncooperative Collusion Under Imperfect Price Information," Econometrica, 52, 87-100.
- Kreps, D., and R. Wilson (1982), "Sequential Equilibrium," Econometrica, 50, 863-94.

- Lehrer, E. (1986), "Two-Player Repeated Games With Non-Observable Actions and Observable Payoffs," mimeo, Northwestern.
- _____ (1988a), "Extensions of the Folk Theorem to Games With Imperfect Monitoring," mimeo, Northwestern.
- _____ (1988b), "Internal Correlation in Repeated Games," Northwestern, CMSEMS Discussion Paper #800.
- _____ (1988c), "Nash Equilibria of n-Player Repeated Games With Semi-Standard Information," mimeo, Northwestern.
- _____ (1988d), "Two Players' Repeated Games With Non-Observable Actions," mimeo, Northwestern.
- _____ (1989), "Lower Equilibrium Payoffs in Two-Players Repeated Games With Nonobservable Actions," forthcoming, International Journal of Game Theory.
- Porter, R. (1983), "Optimal Cartel Trigger-Price Strategies," Journal of Economic Theory, 29, 313-38.
- Radner, R. (1981), "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," Econometrica, 49, 1127-48.
- _____ (1985), "Repeated Principal Agent Games With Discounting," Econometrica, 53, 1173-98.
- _____ (1986), "Repeated Partnership Games With Imperfect Monitoring and No Discounting," Review of Economic Studies, 53, 43-58.
- 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.
- Rubinstein, A. and M. Yaari (1983), "Repeated Insurance Contracts and Moral Hazard," Journal of Economic Theory, 30, 74-97.

- Sorin, S. (1988), "Asymptotic Properties of a Non-Zero Sum Stochastic Game," International Journal of Game Theory, 15, 101-107.
- _____ (1986), "Supergames," mimeo, CORE.
- Stigler, G. (1961), "The Economics of Information," Journal of Political Economy, 69, 213-25.



Date Due

MAY 24 1991

MIT LIBRARIES DUPL 1



3 9080 00567414 5

