



MIT LIBRARIES



3 9080 00567416 0






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

<http://www.archive.org/details/folktheoremwithi00fude>

HB31
.M415

no. 523

**working paper
department
of economics**



THE FOLK THEOREM WITH IMPERFECT
PUBLIC INFORMATION

Drew Fudenberg
David K. Levine
Eric Maskin

No. 523

May 1989

**massachusetts
institute of
technology**

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

THE FOLK THEOREM WITH IMPERFECT
PUBLIC INFORMATION

Drew Fudenberg
David K. Levine
Eric Maskin

No. 523

May 1989

THE FOLK THEOREM WITH IMPERFECT PUBLIC INFORMATION*

By

DREW FUDENBERG**

DAVID K. LEVINE***

AND

ERIC MASKIN****

MAY 1989

* This paper combines and expands two earlier ones, "Discounted Repeated Games with One-Sided Moral Hazard" and "The Folk Theorem with Unobservable Actions". We would like to thank Andrew Atkinson, Patrick Kehoe, Andreu Mas Colell, Abraham Neyman, and Neil Wallace for helpful comments. NSF grants 87-08616, 86-09697, 85-20952 and a grant from the UCLA Academic Senate provided research support.

** Department of Economics, Massachusetts Institute of Technology

*** Department of Economics, UCLA; part of this research was carried out at the University of Minnesota.

**** Department of Economics, Harvard University

MIT LIBRARIES
AUG 23 1989
RECEIVED

ABSTRACT

The Folk Theorem obtains in repeated games with imperfect public information if for each pair of agents there is at least one action profile where the information revealed by the publicly observed outcome permits deviations by one of the agents to be statistically distinguished from deviations by the other, and the dimension of the set of feasible payoffs equals the number of players. Under somewhat stronger conditions, we obtain a Folk Theorem for strict equilibria. Without pairwise full rank, a "Nash-threats" Folk Theorem obtains if the observed outcomes statistically give independent information about each player's actions. We give applications of our results to repeated agency models, the Green-Porter oligopoly model and repeated mechanism design.

Introduction

The Folk Theorem for infinitely repeated games with discounting asserts that any feasible payoff vector that Pareto dominates the minmax point of a stage game can arise as the average discounted payoffs in a Nash (or perfect) equilibrium of the repeated game if the discount factor is sufficiently close to one. This means that patient players can obtain a large set of payoffs without using formal, contractually enforced agreement by playing an equilibrium with the same payoffs. The proof of the Folk Theorem exhibits strategies that yield the desired payoffs and that punish deviations sufficiently strongly that all players wish to conform. An important hypothesis of the theorem is that the players' actions are observed at the end of each period, so that deviations outside the support of the equilibrium strategies are sure to be detected.

This paper provides conditions under which the conclusion of the Folk Theorem holds when players do not directly observe their opponents' actions, but receive imperfect statistical information about them. In keeping with previous work, we assume this information takes the form of publicly observed "outcomes" at the end of each period, so that all players receive the same information about their opponents' play. Obviously, the Folk Theorem fails if the outcomes are completely uninformative about players' actions; the relevant question is how much information the outcomes must convey.

In each period t of a repeated game equilibrium, each player's action maximizes a weighted sum of his payoff at period t and his continuation payoff from period $t+1$ on. If the continuation payoffs are independent of the realized period t outcome, each player will choose an action that maximizes his period t payoff, so the action profile will correspond to a

Nash equilibrium of the stage game. The key to the Folk Theorem is constructing strategies in which the continuation payoffs vary with the outcomes in such a way as to induce players to choose sequences of actions that have any desired payoff. It is easy to determine which actions can be enforced with arbitrary continuation payoffs, but such considerations leave open the question of how these continuation payoffs are themselves to be enforced. This is why we look for a "perfect Folk Theorem", meaning that we require the continuation payoffs themselves to be sustainable by a perfect equilibrium.

To induce the players to choose a specified profile of actions, there must exist continuation payoffs that make each player's action a best response. In that case, we say that the profile is enforceable. Enforceability is ensured by the "individual full rank" condition introduced by Fudenberg-Maskin [1986b], which requires that the probability distributions over outcomes corresponding to different choices of action by a single player be linearly independent. When this condition fails for a given player, different probability mixtures over the player's actions can yield the same distributions over outcomes, and so the mixtures cannot be distinguished. As we will see, individual full rank is essentially sufficient for the Folk Theorem in principal-agent games (see Section 8); it is not, however, sufficient in general.

Rather, we require the stronger condition of pairwise full rank: For every pair of players i and j , there must exist a strategy profile for which the probability distributions over outcomes corresponding to deviations by i and j are linearly independent. Individual full rank ensures that the actions of a single player can be (statistically) distinguished from each other; pairwise full rank extends this condition to the actions of any pair

of players. Pairwise full rank is important because it implies "enforceability on hyperplanes": Any profile of pure actions can be enforced by continuation payoffs lying on (almost) any hyperplane. The Folk Theorem follows from enforceability on hyperplanes and the "full-dimensionality" condition introduced by Fudenberg-Maskin [1986a].

To develop intuition for the significance of enforceability on hyperplanes, consider a two-player partnership game where each player can either "work" or "shirk" and there are two observed outcomes, "good" and "bad". Assume that the payoffs $v^* = (v_1^*, v_2^*)$ when both work are on the efficient frontier, but that shirking is a dominant strategy in the stage game. Now ask whether v^* can be approximated in a repeated-game equilibrium where both players work in the first period.

Fix such an equilibrium, and let $v(G) = (v_1(G), v_2(G))$ and $v(B) = (v_1(B), v_2(B))$ be the continuation payoffs after good and bad outcomes respectively. If we let v denote the average payoffs in this equilibrium, then

$$(1.1) \quad v = (1-\delta)v^* + \delta[p_G v(G) + p_B v(B)],$$

where p_G and p_B are the probabilities of the good and bad outcome when both player works and δ is the discount factor.

Since each player has a short-run incentive not to work, the differences $v_1(G) - v_1(B)$ and $v_2(G) - v_2(B)$ must be positive and sufficiently large to induce the players to work in the first period. When δ is close to 1, even a slight difference between the continuation payoffs has large incentive effects, and we can take $v(G)$ and $v(B)$ to be close together. However, the distance between the efficient first-period payoff

v^* and the equilibrium payoff v cannot converge to zero or δ tends to 1. The reason is that in order to maintain incentives, the value of the difference in continuation payoffs in units of stage-game utility, $[\delta/(1-\delta)](v(G) - v(B))$, must be of the same order of magnitude as the one-shot gains the players could obtain by deviating, and so the "cost" of the difference -- the wedge between v and v^* -- remains constant as $\delta \rightarrow 1$. (This is essentially the Radner-Myerson-Maskin [1986] counterexample; see their paper for the complete argument that equilibria are uniformly inefficient for all discount factors.)

The example above does not satisfy pairwise full rank, and it fails enforceability on hyperplanes, as both players must be "punished" when the bad outcome occurs, and rewarded after the good outcome, so that only hyperplanes with positive slope can be used (see Figure 1.1).

The idea behind enforceability on hyperplanes is that to construct approximately efficient equilibria the continuation payoffs must be arranged to avoid punishing all players at once: If one player's payoff falls, some other's payoff should increase. Such discrimination requires that the outcome "provide statistical information about which player deviated": Otherwise everyone must be punished, which will lead to inefficiency. The pairwise full rank condition ensures that deviations by player i can be (statistically) distinguished from deviations by player j , and thus permits enforceability on hyperplanes.

To illustrate this idea modify the partnership game above so that there are two "bad" outcomes, B_1 and B_2 . The analog of equation (1.1) is

$$(1.2) \quad v = (1-\delta)v^* + \delta(p_G v(G) + p_{B_1} v(B_1) + p_{B_2} v(B_2))$$




Figure 1.1 Decomposition of equilibrium payoffs v when pairwise full rank fails

For $i = 1, 2$, suppose that, if player i deviates from the action profile that sustains v^* , there is a higher probability of outcome B_i than of B_j , $j \neq i$. The modified game satisfies enforceability on hyperplanes, and so we can arrange the continuation payoffs to provide the necessary incentives while lying in a line that is parallel to the Pareto frontier (see Figure 1.2). The earlier argument that the equilibrium payoffs are bounded away from v^* now loses its force, and, as we will show, the Folk Theorem obtains in the modified game.

So far we have stressed the role of full rank condition in ensuring enforceability on hyperplanes. However, we will show that any Pareto efficient vector of actions is enforceable regardless of whether or not individual full rank holds. We use this observation to show that, in the class of games where enforceability implies enforceability on hyperplanes, all payoffs that Pareto dominate a Nash equilibrium of the stage game can arise in an equilibrium of the repeated game.

A case of particular importance is when the distribution over outcomes has a "product structure", with different components of the outcome identified with the actions of different players, and each component's distribution depending only on the action of the associated player. For these games, enforceability implies enforceability on hyperplanes. We exploit this observation to explain how "full insurance" can be provided in Green [1987]'s model of repeated insurance contracts for privately observed risks.

Another class of games where enforceability implies enforceability in hyperplanes is games where the actions of all but one player are observable, including the familiar principal-agent model. Indeed, as we show in Fudenberg-Levine-Maskin [1989], it often applies to games with several




Figure 1.2 Decomposition of equilibrium payoffs when pairwise full rank holds.

"agents" (whose actions are not observable), provided there exists a "principal" who can make observable transfers.

There has been considerable previous work related to the Folk Theorem, including Aumann-Shapley [1976], Friedman [1971], Porter [1983], Green-Porter [1984], Fudenberg-Maskin [1986a], Rubinstein [1979], Radner [1981], [1985], [1986], Rubinstein-Yaari [1983], and Radner-Myerson-Maskin [1986].

In the observable action case, Aumann-Shapley [1976] and Rubinstein [1979] considered the Folk Theorem in games with no discounting at all. Friedman [1971] proved a "Nash threat" version of the result for the discounting case, and Fudenberg-Maskin [1986a] established the general Folk Theorem for discounted games with observable actions under the condition that the dimension of the feasible set be equal to the number of players. They showed by example that with more than two players, some such non-degeneracy condition is required. (They also introduced one special case of unobservability to the Folk Theorem literature by considering mixed-strategy punishments.)

Porter [1983] and Green-Porter [1984] considered repeated oligopoly games with unobservable outputs and a stochastic price. They showed that there are equilibria that Pareto-dominate the static equilibrium.

Radner [1981], Rubinstein-Yaari [1983], and Radner [1985] established results similar to the Folk Theorem in a class of principal-agent games where only one player's (the agent's) action is not observable; the first two of these papers treated the case of no discounting. We generalize Radner [1985], who assumed discounting, in Section 8. Radner [1986] showed that a Folk Theorem obtains in a class of repeated partnership games with time-averaging (i.e. no discounting) payoffs that does not satisfy our informational conditions. The no-discounting assumption is essential:

Radner-Myerson-Maskin [1986] construct an example to show the result fails with discounting.

The papers by Radner and Rubinstein-Yaari used techniques based on statistical inference and "review strategies." These strategies avoid the need to provide period-by-period incentives not to deviate. Instead, at the end of each "review phase" a test is conducted to determine whether it is likely that some player has deviated a non-negligible fraction of the time. This approach may give the impression that the key to obtaining the Folk Theorem is the improvement in information over time. We feel that this intuition is misleading, as increasing the number of periods during reviews also increases the number of profitable opportunities for the players to deviate.

In contrast to this statistical approach, we use dynamic programming arguments and geometry, and focus on the relationship between the information revealed by outcomes and the enforceability conditions. This permits us to provide more general theorems. It also provides an alternative understanding of when and why the theorems obtain.

Several authors have exploited the link between discounted repeated games and dynamic programming, including Abreu [1988], Fudenberg-Maskin [1986a], and Radner-Myerson-Maskin [1986]. Our methods are closest to those of Abreu-Pearce-Stacchetti [1988], who develop the technique of "self-generation," which extends the optimality principle of stochastic dynamic programming to pure strategy equilibria of games with unobservable actions. We apply their technique to a class of mixed-strategy sequential equilibria that we call "perfect public equilibria," where the players' strategies depend only on the publicly observable information, and not on private information about their own past actions. The set of perfect public

equilibrium payoffs is stationary, which greatly simplifies the analysis.

After a draft of this paper was completed, we became aware of the prior contribution of Matsushima [1988] that anticipated some of our key ideas. In particular, the idea of enforceability on hyperplanes is used, and differentiability of the boundary of a set of equilibrium payoffs plays a key role in the proof. Our work differs from that of Matsushima in several important ways. First, where Matsushima makes assumptions about the endogenous limit set of equilibrium payoffs, we instead impose conditions directly on the structure of the stage game, which enables us to develop the link between enforceability on tangent hyperplanes and the information revealed by the outcomes. Second, Matsushima assumes that there is a continuum of actions, and defines an equilibrium by means of first-order conditions for an optimum, ignoring the second-order conditions. We consider games with finitely many actions, and study the full set of incentive compatibility constraints. Finally, we consider a "full" Folk Theorem, whereas Matsushima shows that a limit of points efficient relative to the set of equilibria is efficient relative to the set of feasible payoffs leaving open the question of the "width" of the equilibrium set.

This paper uses enforceability on hyperplanes to give sufficient conditions for the Folk Theorem. By modifying the techniques developed here, and considering enforceability on half-spaces rather than hyperplanes, one can obtain an exact characterization of the limit set of perfect public equilibrium payoffs for any information structure. Fudenberg-Levine [1989b] use this approach to analyze repeated games with long-run and short-run players, where the Folk Theorem fails.

Section 2 introduces our model and defines perfect public equilibria. Section 3 reviews the way that self-generation extends one-player dynamic programming, and establishes several useful lemmas. Section 4 introduces enforceability on hyperplanes and shows how it can be used to prove a Folk Theorem. Section 5 investigates the conditions on a player's information that imply that enforceability on hyperplanes is satisfied. Section 6 treats games where the outcomes have a product structure. Section 7 applies the results of 6 to adverse selection models, such as that of Green [1987]. Section 8 considers principal-agent games. We conclude with some more general remarks in Section 9.

2. The Model

In the stage game, each player $i = 1, \dots, n$ simultaneously chooses an action a_i from a finite set A_i with m_i elements. We let $a \in A = \times A_i$. A profile of action $a \in A = \times A_i$ induces probability distributions on a public outcome $y \in Y$ and on n private outcomes $z_i \in Z_i$, where Y has \bar{m} elements and Z_i has \bar{m}_i elements. We let $Z = \times Z_i$. Each player i 's realized payoff $r_i(z_i, y)$ depends on his private outcome z_i and the public outcome y , but not on z_j , $j \neq i$. Until Section 7 below, the private outcome z_i can be identified with player i 's action a_i ; we require the extra generality only to deal with repeated games of adverse selection. Let $\pi_{yz}(a)$ be the distribution over y and z induced by a . We let $\pi_y(a) = \sum_{z \in Z} \pi_{yz}(a)$ be the marginal distribution of y . Player i 's expected payoff from an action profile a is

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

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

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

and the expected payoffs

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

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 analogously.

Our formulation embraces several models considered in the repeated games literature. All but the last example identify z_i with the action a_i .

First, if the public outcome is simply the profile of actions, we have the conventional framework of perfect observability ex post, as in Fudenberg-Maskin [1986a]. Second, with a judicious interpretation of the variables, we can obtain the standard principal-agent model. In this model; one player, the principal, moves first and selects one of a finite number of monetary transfer rules, where the transfer to the other player, the agent, depends on output. Next, the agent selects an effort level (unobservable to the principal), which stochastically determines output. This is a model of one-sided moral hazard. If we designate player 2 as the principal, then we can identify his action a_2 with a transfer rule. Because we have modeled the players as moving simultaneously--whereas actually the agent moves after the principal-- we must interpret a_1 as a contingent effort level, that is, a function dependent on the principal's move. A public outcome is then a

realized output level together with the principal's monetary transfer rule. (By including the principal's action as part of the outcome, we make it observable to the agent).¹

Third, we can think of the players as a partnership in which each player provides an unobservable input a_i to a production process (see Radner [1986]). Output y depends stochastically on the quantities of all inputs and is divided among the partners in a predetermined way. This is a model of multi-sided moral hazard.

Fourth, we can imagine that the players are Cournot oligopolists whose outputs a_i are not publicly observable. Demand is stochastic, and so the market price, y , is a stochastic function of output (see Green-Porter [1984]).

Finally, if a_i is no longer identified with z_i , we can imagine that, each period, "agents" receive private information about their endowments that their actions are "reports" (possibly untruthful) of their information, and

¹ Although this interpretation of our model has the same normal form as the standard principal-agent model, it has more subgame-perfect equilibria because of the simultaneity of moves. For example, one subgame-perfect equilibrium of our model (but not of the standard model) entails the agent playing the weakly dominated strategy of refusing all transfer rules. To ensure that the equilibrium sets in the two models are the same, we could either (a) invoke solution concepts such as trembling hand-perfection (Selten [1975]) or sequential equilibrium (Kreps-Wilson [1982]) (once we repeat the game, however, we would then have to modify these concepts to deal with the infinity of strategies) or (b) perturb the game slightly so that the principal literally makes small "mistakes", in which case no strategy for the agent that is ruled out by backward induction in the sequential (standard) principal-agent game can be a weak best response in our simultaneous formulation. We should note that many of our results guarantee strict equilibrium, so that the issue of refinements does not arise.

that a mechanism chooses transfers between players as a function of the reports. Here a_i is a report, z_i is the true endowment level, and the public outcome y is a profile of reports.

In the above examples each player's payoff depends on other players' actions only through the public outcome y , and not through his private outcome z_i . Formally, for y with $\pi_y(\alpha) > 0$, we can define the conditional expected utility

$$\bar{r}_i(y, \alpha) = \sum_{z \in Z} \frac{r_i(z_i, y) \pi_{yz}(\alpha)}{\pi_y(\alpha)}.$$

One case of particular interest is when the following condition is satisfied.

Definition 2.1: Player i 's payoff satisfies the factorization condition if for all (α, y) with $\pi_y(\alpha) > 0$,

$$\bar{r}_i(y, \alpha) = \bar{r}_i(y_i, \alpha'_{-i}), \text{ for any } \alpha'_{-i}.$$

that is, it depends on the play of other players only through the public outcome y . When the factorization condition holds, we can write $\bar{r}_i(y, \alpha_i) = \bar{r}_i(y, \alpha)$. Naturally, $g_i(\alpha) = \sum_{y \in Y} \pi_y(\alpha) \bar{r}_i(y, \alpha_i)$, where $\bar{r}_i(y, \alpha_i)$ is undefined only for values of y for which $\pi_y(\alpha) = 0$.

Note that when factorization is not satisfied, a player's stage-game payoff provides information about his opponents' actions that is not provided by the public outcome. We will not impose the factorization condition for the time being. We will assume factorization in Section 6 through 8, to obtain results for games that do not satisfy the informational conditions we develop in Section 5.

Turning now to the repeated game, in each period $t = 0, 1, \dots$, the stage game is played and the corresponding outcome is then revealed. The public history at the end of period t is $h(t) = (y(0), \dots, y(t))$. The private history for player i at the end of period t is $h_i(t) = (a_i(0), z_i(0), \dots, a_i(t), z_i(t))$. A strategy for player i is a sequence of maps $\sigma_i(t)$ mapping public and private histories $(h(t-1), h_i(t-1))$ to probability distributions over A_i .

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. Player i 's objective is to maximize the average discounted expected value of per-period payoffs using the common discount factor δ . Thus, if $\{g_i(t)\}$ is a sequence of stage payoffs for player i , player i 's overall average payoff is

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

(We use average payoffs, obtained by multiplying total discounted payoffs by $1-\delta$, in order to measure the repeated game payoffs in the same units as payoffs in the stage game.)

Let V be the convex hull of the set of vectors $g(a)$ for $a \in A$. Sorin [1986] shows that all payoffs in V are feasible in the repeated game if $\delta > 1-1/d$, where d is the number of pure strategy profiles of the stage game. The extremal points of V are those which are not convex combinations of other points in V ; extremal actions are those which generate extremal payoffs. Note that all extremal strategies are pure. Let

$v_i^* = \min_{\alpha_{-i}} \max_{a_i} g_i(a_i, \alpha_{-i})$ be player i 's minmax value, and let

m_{-i}^i be strategies for i 's opponents that attain this minimum, and m_i^i a

best response by i to m_{-i}^i . We call m_{-i}^i the minmax profile against player i .

The payoff vector v is individually rational if $v_i \geq v_i^*$ for all players i ; it is strictly individually rational if this inequality is strict. Let $V^* = \{v \in V \mid v_i \geq v_i^* \text{ for all } i\}$ be the set of feasible, individually rational payoffs. The Folk Theorem for games with observable actions asserts that any strictly individually rational payoff in V^* can be supported by an equilibrium of the repeated game if the discount factor is sufficiently close to one.

Our focus is on a special class of the Nash equilibria that we call perfect public equilibria. A strategy $\{\sigma(t)\}$ for player i is public if, at each time t , $\sigma(t)$ depends only on the public information $h(t-1)$ and not on the private history $h_i(t-1)$. A perfect public equilibrium is a profile of public strategies such that beginning at any date t and for every history $h(t-1)$ the strategies are a Nash equilibrium from that date on. This definition allows the players to contemplate deviating to a non-public strategy, but such deviations are irrelevant: If player i 's opponents use public strategies, it follows that player i cannot profit from a non-public strategy. Note also that in a public equilibrium the players' beliefs about each others' past actions 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.

Every perfect public equilibrium is clearly a perfect Bayesian equilibrium.² It is easy to see that the set of perfect public equilibrium payoffs is stationary, meaning that the set of perfect public equilibrium continuation payoffs in any period t and for any public history $h(t-1)$ is independent of t and $h(t-1)$. Note also that every perfect Bayesian equilibrium payoff must lie in V^* . Consequently when the Folk Theorem obtains it exactly characterizes the perfect public equilibrium payoffs as $\delta \rightarrow 1$.

However, there can be perfect Bayesian equilibria whose payoffs are not in V^* . The reason is that the players may be able to correlate their play without using an external "correlating device" by combining their public and private information. (This was pointed out to us by A. Neyman.) Since the minmax value for player i can be strictly lower when his opponents are allowed to use correlated strategies, this endogenous correlation can lead to equilibria where some players have lower payoffs than v_i^* . We provide an example of this possibility in the Appendix.

3. Dynamic Programming and Local Generation

Let $E(\delta) \subset V^*$ be the set of average payoffs that can arise in perfect public equilibria when the discount factor is δ , and let $V(\delta)$ denote the set of average payoffs that are feasible when the discount factor is δ , that is, the values that can be generated by some sequence of repeated-game strategies that depend only on public information.

² Here we use the weakest notion of perfect Bayesian equilibrium: actions are required to be sequentially rational given beliefs, and beliefs after any history that the equilibrium assigns positive probability are obtained using Bayes rule.

Lemma 3.1: $E(\delta)$, $V(\delta)$, V and V^* are compact sets, and V and V^* are convex.

Proof: The compactness of $E(\delta)$ is well known, see for example Fudenberg and Levine [1983]. The sets $V(\delta)$, V and V^* are obviously compact; V is convex by definition, and V^* is the intersection of two convex sets. ■

Definition 3.1: Let δ , $W \subseteq \mathbb{R}^n$ and $v \in \mathbb{R}^n$ be given. Profile α is enforceable with respect to v , W and δ if there exist vectors $w(y) \in \mathbb{R}^n$, for each $y \in Y$, such that for all i and $a_i \in A_i$,

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

and

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

We say that the $w(y)$'s enforce α with respect to v, W and δ . If for given α , there exists v satisfying (3.1), we say that α is enforceable with respect to δ and W .

Formula 3.1 requires that if the players other than i play α_{-i} and player i 's continuation payoff contingent on outcome y is $w_i(y)$, then player i obtains the same average payoff from all actions in the support of α_i , and no action yields a higher average payoff.

Definition 3.2: If, for given v , W , and δ , there exists an α enforceable with respect to v, W and δ , we say that v is generated by δ and W . $B(\delta, W)$ is the set of all points generated by δ and W .

The following lemma exploits the stationary nature of the set of perfect public equilibria.

Lemma 3.2: (Self-Generation) If $W \subseteq B(\delta, W)$ and W is bounded then $W \subseteq E(\delta)$.

Sketch of Proof: The idea behind this observation and its proof are due to Abreu, Pearce, and Stacchetti [1987], who considered pure-strategy equilibria of games with a continuum of outcomes; our finite-outcome model avoids questions of measurability. The proof mimics that of the sufficiency of the Principle of Optimality in dynamic programming: First, decompose $w \in W$ as in (3.1). The decomposition defines the players' first period actions. Next decompose each continuation payoff $w(y)$, thereby defining the second period actions, etc. By construction and the boundedness of W , the sequence of actions thus defined has average payoff w , so w is feasible. Next we verify that no player can gain by deviating. As in the dynamic programming case, this is a consequence of condition (3.1), which implies that no one-period deviation is profitable, and the fact that payoffs are uniformly bounded. ■

We can now prove a simple consequence of (3.1) and (3.2).

Lemma 3.3: (Monotonicity on convex sets) If W is convex and $\delta \leq \delta'$, then $[B(\delta, W) \cap W] \subseteq [B(\delta', W) \cap W]$.

Remark: A related result is proved by Abreu, Pearce, and Stacchetti [1987]; in their model the set of equilibria is itself convex.

Proof: If $v \in [B(\delta, W) \cap W]$, then since v can be enforced by continuation payoffs in W there exist α and $w(y) \in W$ satisfying (3.1) and (3.2). We will construct a $w'(y) \in W$ that enforces α when the discount factor is δ' and such that the corresponding average payoff is still v .

Set

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

Since $\delta' > \delta$, $w'(y)$ is a convex combination of v and $w(y)$, so $w'(y) \in W$.

For each player i and profile α' , i 's average payoff from (α'_i, α_{-i}) is

$$\begin{aligned} & (1 - \delta')g_i(\alpha'_i, \alpha_{-i}) + \delta' \sum_y \pi_y(\alpha'_i, \alpha_{-i}) w'_i(y) - \\ & \quad [(\delta' - \delta)/(1 - \delta)]v_i + \\ & (1 - \delta')/(1 - \delta)[(1 - \delta)g_i(\alpha'_i, \alpha_{-i}) + \delta \sum_y \pi_y(\alpha'_i, \alpha_{-i})w'(y)]. \end{aligned}$$

In particular, if player i plays α_i , he obtains

$[(\delta' - \delta)/(1 - \delta)] + (1 - \delta')/(1 - \delta)v_i = v_i$. It then follows that since (3.1) was satisfied for α , δ and $w(y)$, it is satisfied for α , δ' , and $w'(y)$ as well. ■

It will be convenient for the analysis below to rewrite (3.1) by subtracting out the expected values. Thus define

$$\begin{aligned}
& \hat{g}_i(a_i, \alpha) = g_i(a_i, \alpha_{-i}) - g_i(\alpha), \\
(3.3) \quad & \bar{w}_i = \sum_{y \in Y} \pi_y(\alpha) w_i(y), \\
& \text{and } \hat{w}_i(y) = w_i(y) - \bar{w}_i.
\end{aligned}$$

Formula (3.1) can be rewritten as

$$\begin{aligned}
(3.1') \quad & - \hat{g}_i(a_i, \alpha)(1-\delta)/\delta = \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) \hat{w}_i(y) \quad \text{for } a_i \text{ with } \alpha_i(a_i) > 0 \\
& - \hat{g}_i(a_i, \alpha)(1-\delta)/\delta \geq \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) \hat{w}_i(y) \quad \text{for } a_i \text{ with } \alpha_i(a_i) = 0,
\end{aligned}$$

and the requirement that $w(y) \in W$ is

$$(3.2') \quad \bar{w} + \hat{w}(y) \in W.$$

To establish that a set W is self-generating for a fixed discount factor, one must show that every point in W can be generated using continuation payoffs lying in W . We will find it convenient to work instead with the weaker property of local generation.

Definition 3.3: A subset W of \mathbb{R}^n is locally generated if for each $v \in W$ there exists $\delta < 1$ and an open set U containing v such that $U \cap W \subseteq B(\delta, W)$.

Note that local generation allows a different δ 's to correspond to different points $v \in W$. However, if the set W is compact, we can replace this multiplicity of δ 's with a single discount factor. More precisely, we have

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

Proof: The open neighborhoods U of the definition of local generation form a cover of the compact set W , and thus there is a finite collection (U^k) that cover W . Let δ^k be a discount factor such that $Z^k = W \cap U^k \subseteq B(\delta^k, W)$, and let δ' be the largest of the δ^k . From monotonicity, (lemma 3.3), if $\delta \geq \delta'$, $[B(\delta^k, W) \cap W] \subseteq [B(\delta, W) \cap W]$. Thus $Z^k \subseteq B(\delta, W)$. Since the union of the Z^k equals W , we conclude that $W \subseteq B(\delta, W)$. ■

Lemma 3.5 (Enforcement with Small Variation): Suppose H is a linear subspace of \mathbb{R}^n . If α is enforceable with respect to H for some $\delta' > 0$, then there exists a constant κ such that for all $v \in \mathbb{R}^n$ and $\delta > 0$, there are continuation payoffs $w(y)$ that enforce α with respect to $v + H$ and δ , such that

$$\sum_{y \in Y} \pi_y(\alpha) w(y) = v, \text{ and } \|w(y) - v\| \leq \kappa(1 - \delta)/\delta.$$

Remark: If $H = \mathbb{R}^n$, we say simply that α is enforceable. The lemma shows that by taking δ sufficiently close to one, an enforceable profile can be enforced with continuation payoffs that are arbitrarily close together, and that if a profile is enforceable in a hyperplane, it is enforceable in

any translate of that hyperplane. As we will see, the reason that many payoffs can be supported by equilibria when δ is near one is precisely that the variation in continuation payoffs can be made small. Conversely, it is clear that if the required variation is large, there may be no way to enforce certain actions with continuation payoffs that lie in the feasible set.

Proof: Let $\hat{w}'(y)$ and \bar{w}' satisfy (3.1') and (3.2') for $H+v'$ and δ' . Since

$\sum_{y \in Y} \pi_y(\alpha) \hat{w}'(y) = 0$, $\hat{w}'(y) \in H$ for each y , and $\bar{w}' = 0$. One can check that

$\hat{w}(y) - [\delta'(1-\delta)/\delta(1-\delta')] \hat{w}'(y)$ satisfies (3.1') for δ , and clearly $\hat{w} \in H$.

Take $w(y) = v + \hat{w}(y) \in v + H$. Then $\|w(y) - v\| \leq [\delta'(1-\delta)/\delta(1-\delta')] \max_y \|w'(y) - v'\|$.

So the conclusion of the lemma holds for $\kappa = \delta'/(1-\delta') \max_y \|w'(y) - v'\|$. ■

4. Enforceability and Local Generation

This section introduces a condition that is sufficient for a set of payoffs to be locally generated. Section 5 proves a version of the Folk Theorem by providing assumptions on the information structure that imply the sufficient condition is satisfied by sets that are arbitrarily good approximations to the set V^* , the set of feasible, individually rational payoffs. We offer a counter example in Section 5 to illustrate the need for the informational assumptions.

Recall that V is the convex hull of socially feasible payoffs in the one-shot game. Let \hat{V} be an arbitrary closed convex subset of V . Such a set satisfies enforceability on tangent hyperplanes if for every extremal point $\hat{v} \in \hat{V}$ and every $n-1$ -dimensional subspace H such that $\hat{v} + H$ is tangent to \hat{V}

at \hat{v} there exist v' weakly separated from \hat{V} by $\hat{v} + H$, a mixed strategy α with $g(\alpha) = v'$ and continuation payoffs $w(y) \in H$ for each $y \in Y$ that enforce α . A somewhat weaker condition is

Condition 4.1: A closed convex subset $\hat{V} \subset V$ satisfies weak enforceability on tangent hyperplanes if for every extremal point $\hat{v} \in \hat{V}$ and every $n-1$ -dimensional subspace H such that $\hat{v} + H$ is tangent to \hat{V} at \hat{v} , there is a sequence of mixed profiles $\{\alpha^k\}$ and a point v' weakly separated from \hat{V} by $\hat{v} + H$, with $\lim g(\alpha^k) = v'$ and for each α^k , continuation payoffs with associated $w^k(y) \in H$ that enforce α^k .

Note that Lemma 3.5 (enforcement with small variation) implies that, under Condition 4.1, each extremal point can be enforced with continuation payoffs that vary little across actions if δ is near 1. Note also that the $\hat{\alpha}$ hypothesized in Condition 4.1 need not have payoffs that are weakly separated from \hat{V} .

Theorem 4.1: If \hat{V} is a closed convex subset of V^* that satisfies condition 4.1, then for every closed set W in the interior of \hat{V} there is a $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$, $W \subseteq E(\delta)$.

To prove Theorem 4.1, we will show that each closed W in the interior of \hat{V} is contained in a smooth set W' that is locally generated.

Definition 4.1: A subset $W \subseteq \mathbb{R}^n$ is smooth if it is closed and has non-empty interior with respect to \mathbb{R}^n , and if its boundary is a C^2 submanifold of \mathbb{R}^n .

Lemma 4.1: If W is a closed convex subset of the interior of \hat{V} , there exists a smooth convex set W' such that $W \subset W' \subset \text{interior } \hat{V}$.

Proof: Standard. ■

Theorem 4.1 now follows once we show that each smooth convex W in the interior of \hat{V} is locally generated.

Lemma 4.2: If \hat{V} satisfies condition (4.1), and $W \subset \text{interior } (\hat{V})$ is a smooth convex set, then W is locally generated.

Outline of Proof

Fix v on the boundary of W , and consider how it might be generated by W . We know that there must be some extremal payoff \hat{v} of \hat{V} , which is strictly separated from W by the tangent plane $v+H$ to W at v . Moreover, we can clearly pick the extremal point \hat{v} so that a translate of $v+H$ is also tangent to \hat{V} there. Consequently, from condition 4.1, there exist a v' weakly separated from \hat{V} by T_v and a sequence (α^k) such that $g(\alpha^k) \rightarrow v'$ and; for any translate of H , α^k can be enforced with continuation payoffs that lie on that translate. Since $W \subset \text{interior } (\hat{V})$, v' is strictly separated from W by T_v , and we may pick k so that $g(\alpha^k)$ is also strictly separated from W' by $v+H$. (See Figure 4.2) ~~where the tangent plane to \hat{V} at \hat{v} is also tangent to W at v .~~
 If α is enforceable with respect to v , then

$$v = (1-\delta)g(\alpha) + \delta[v - (v - \sum \pi_y(\alpha)w(y))]$$

$$= v + (1-\delta)(g(\alpha) - v) - \delta(v - \ell),$$

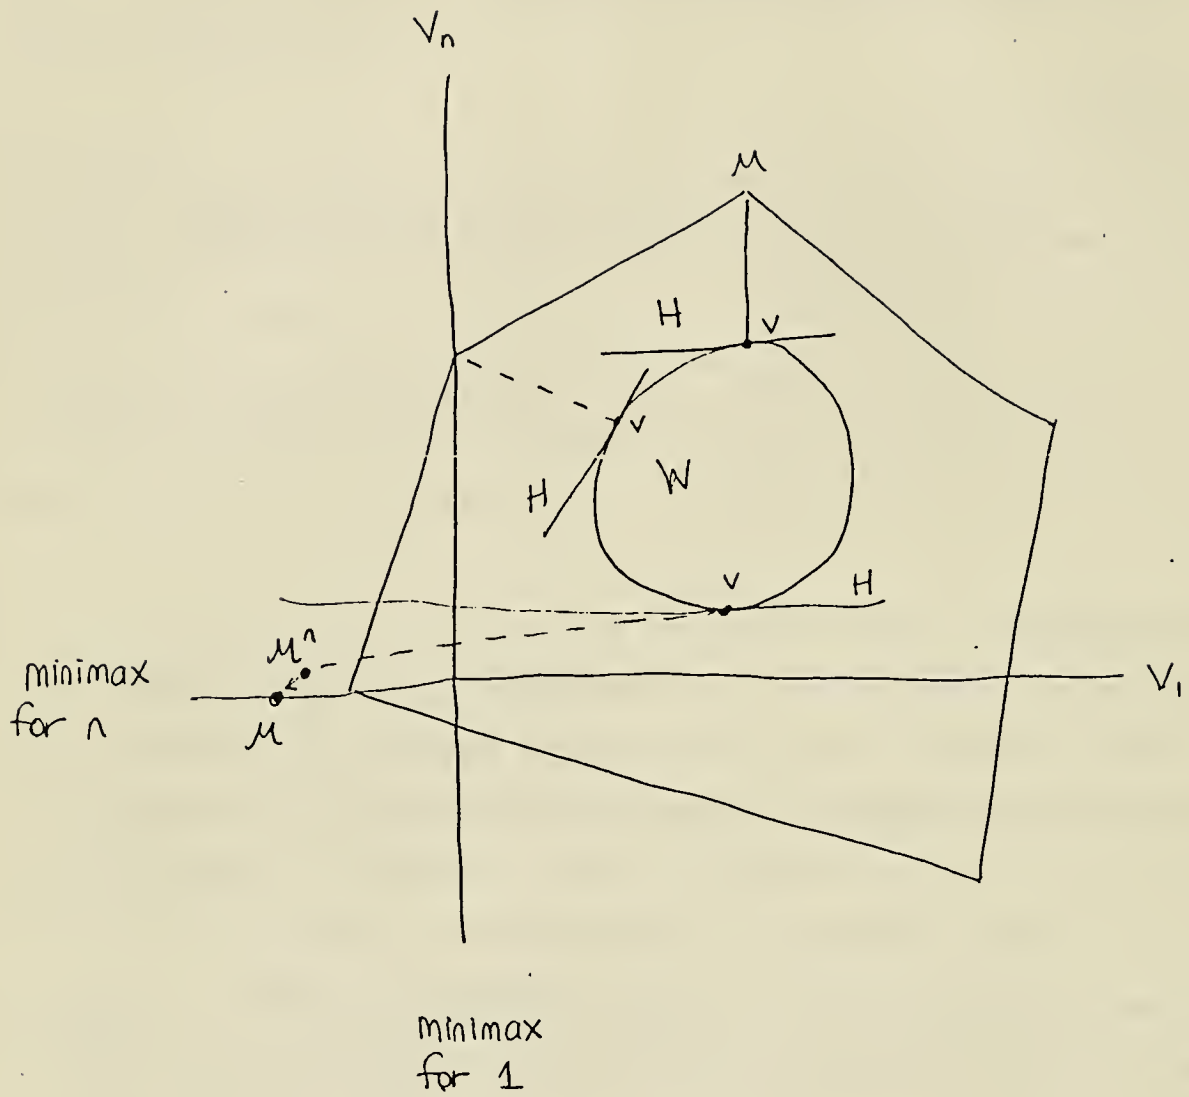


Figure 4.1: Selection of Extreme Points

where

$$\ell = v - \int \pi_y(\alpha)w(y).$$

is the average distance of the continuation payoffs from the boundary at W .

Because W is smooth, the distance from \bar{w} to the boundary of W along the hyperplane $\bar{w}+H$ is on the order of $\sqrt{\|\ell\|}$ (see figure ^{4.2}₆). From Lemma 3.5, we can make do with continuation payoffs whose variation is proportional to $(1-\delta)/\delta$. But $\|\ell\|$ is also proportional to $(1-\delta)/\delta$ and so, because $\|\ell\| < \sqrt{\|\ell\|}$ for δ near 1, we conclude that the continuation payoffs belong to W for such discount factors. The role of Condition 4.1 in obtaining local generation is to ensure that we can find continuation payoffs lying on a translate of H . If, instead, the continuation payoffs were constrained to lie on a hyperplane orthogonal to the boundary of W at v (as in the first example of the introduction), then the distance $\|\ell\|$ would be of the same order as the variation in the continuation payoffs, so that the continuation payoffs could lie outside W even for δ near 1.

Proof of Lemma 4.2:

Step 1: Suppose $v \in \text{interior}(W)$. Let $U' \subset W$ be an open ball of radius r that contains v , and let U be an open ball of radius $r/5$. By condition 4.1, there exists a profile α that is enforceable by some continuation payoffs $\{w(y)\}$. By lemma 3.5, we may assume $\|v-w(y)\| \leq \kappa(1-\delta)/\delta$. Choose δ so that $[(1-\delta)/\delta] \|v-g(\alpha)\| < r/5$, and $\kappa(1-\delta)/\delta \leq r/5$. For given $v' \in U$, Lemma 3.5 implies that there exist continuation payoffs $\{w'(y)\}$ satisfying $\bar{w}' = [v' - (1-\delta)g(\alpha)]/\delta$. This implies that $v' = (1-\delta)g(\alpha) + \delta w'(y)$. Moreover, $\|w'(y) - v\| = \|w'(y) - \bar{w}' + \bar{w}' - v' + v' - v\| \leq \frac{3}{5} r$. Hence $w'(y) \in U'$, and so $U \subset B(\delta, W)$.




Figure ~~4.2~~ Ensuring that continuation payoffs lie in w

4.2

Step 2: Consider $v \in \text{boundary}(W)$. Since W is smooth, there is a unique hyperplane $v+H$ tangent to W at v . We argued above that, from condition 4.1, there is an action profile α such that $g(\alpha)$ is strictly separated from W by $v+H$, and such that α can be enforced with continuation payoffs that lie on $v+H$. Lemma 3.5 then implies that there exists a κ such that for any $\bar{w} \in \mathbb{R}^n$, $g(\alpha)$ is enforceable with respect to $\bar{w} + H$ using continuation payoffs $w(y)$ that differ from \bar{w} by at most $\kappa(1-\delta)/\delta$.

Step 3: It is now convenient to choose a new coordinate system. As shown in Figure 4, we take v to be the origin, choose the vertical (1st) axis to be the line connecting v and $g(\alpha)$, and the remaining axes to lie in $v+H$. Since $g(\alpha) \notin v+H$ this transformation is possible. In the new coordinates, W is still smooth and convex, and, while the norm is changed by this transformation of coordinates, κ may be replaced by a new constant κ' since all norms on \mathbb{R}^n are equivalent. Note that $\bar{w} = (-\|\ell\|, 0)$, where, as in the outline of the proof, $\ell = v - \bar{w}$.

For $x \in \mathbb{R}^n$ we now write $x = (x^0, x^H)$, where x^0 is the component on the vertical axis and x^H the component in $v+H$. Since W is smooth and convex there is a $c > 0$ such that, within the set $X = \{x \mid \|x^0\| \leq c, \|x^H\| \leq c\}$, the boundary of W can be represented by a C^2 concave function $x^0 = f(x^H)$.

Step 4: Let $\bar{s} = \max \{D^2 f(x^H) \mid \|x^H\| \leq c\}$. By Taylor's Theorem, $f(x^H) \geq -\bar{s} \|x^H\|^2$. Thus if $\|x^H\| < \kappa'(1-\delta)/\delta$, $f(x^H) \geq -\bar{s} [\kappa'(1-\delta)/\delta]^2 = -\kappa'' [(1-\delta)/\delta]^2$. Set $\bar{w}_\delta = (-\kappa'' [(1-\delta)/\delta]^2, 0)$ to be a point that lies along the first coordinate axis; $\bar{w}_\delta \in W$ if δ is sufficiently close to 1. From Lemma 3.5 $g(\alpha)$, is enforceable with continuation payoffs $w(y)$ such that

$\|w(y) - \bar{w}_\delta\| \leq \kappa' (1-\delta)/\delta$, $w(y) \in \bar{w} + H$, and $\sum \pi_y(\alpha) w(y) = \bar{w}$. By construction, this implies that $w(y) \in W$. The corresponding average payoffs are:

$$v_\delta = (1-\delta)g(\alpha) + \delta \sum_{y \in Y} \pi_y(\alpha) w(y) - (1-\delta)g(\alpha) + \delta \bar{w}_\delta - \\ [(1-\delta)(g(\alpha))^0 - \kappa'' [(1-\delta)/\delta]^2, 0] = (v_\delta^0, 0),$$

where v_δ^0 is defined by the final equality. Since $(g(\alpha))^0$, the first component of $g(\alpha)$ in the new coordinate system, is positive, there is a $\delta < 1$ such that $0 < v_\delta^0 < c/4$. Choose a $\delta > 1/2$ satisfying that condition and also $\bar{w}_\delta > -c/4$. We conclude that $(w(y) - v_\delta/\delta) \in \text{interior}(W)$ for all y . Thus there is an $\epsilon > 0$ such that for any $v'_\delta \in \mathbb{R}^n$, $\|v_\delta - v'_\delta\| < \epsilon$ implies that $(w(y) - v'_\delta/\delta) \in \text{interior}(W)$. Let U be a ball about the origin of radius ϵ . For $v' \in U$, set $v'_\delta = v_\delta - v'$. Then $w'(y) = w(y) - v'_\delta/\delta \in W$. Since α is enforced by continuation payoffs $w(y)$ for discount factor δ , α is also enforced by $w'(y)$: subtracting off a constant from all of the continuation payoffs does not change the incentive to deviate. And by construction, the average payoff from playing α followed by $w'(y)$ is

$$(1-\delta)g(\alpha) + \delta \sum \pi_y(\alpha) w'(y) = (1-\delta)g(\alpha) + \delta \sum \pi_y(\alpha) w(y) - v_\delta + v' = v'$$

from the definition of v_δ . Thus $U \subseteq B(\delta, W)$. ■

5. Informational Conditions

In this section we develop sufficient conditions for a "full" Folk Theorem to hold, that is, conditions under which we may take $\hat{V} = V^*$ in Condition 4.1. Recall that m_i is the number of actions in A_i , and that m denotes the number of outcomes in Y .

Definition 5.1: Given a pair of players $i \neq j$, $\pi_y(\cdot)$ satisfies pairwise full rank at α for (i,j) , if the $(m_i + m_j)$ vectors

$$\left\{ \left\{ \pi_y(a_i, \alpha_{-i}) \right\}_{a_i \in A_i}, \left\{ \pi_y(a_j, \alpha_{-j}) \right\}_{a_j \in A_j} \right\}$$

considered as rows of a matrix, (where the columns correspond to outcomes y) have rank $(m_i + m_j - 1)$.

This is a full rank condition because the vectors must always admit one linear dependency $\sum \alpha_i(a_i) \pi_y(a_i, \alpha_{-i}) - \pi_y(\alpha) - \sum \alpha_j(a_j) \pi_y(a_j, \alpha_{-j})$. The pairwise full rank condition obviously requires that the number of outcomes m be at least $(m_i + m_j - 1)$.

Condition 5.1: The game satisfies the pairwise full rank condition if, for all pairs i,j , $i \geq j$, there exists a profile α such that π satisfies pairwise full rank at α .

Lemmas 5.1 and 5.4 show that condition 5.1 is sufficient for V^* to satisfy weak enforceability on tangent hyperplanes. Note that the condition is much weaker than the requirement that π should satisfy pairwise full rank at all profiles α . Indeed, in a "strongly symmetric" game, where permuting

the actions of the players does not change the distribution of outcomes, the outcome function π cannot satisfy pairwise full rank at any profile where all players use the same action. Nevertheless, in such games, pairwise full rank may be satisfied at certain asymmetric profiles (see Example 5.2 below). This explains why there are many strongly symmetric games in which the Folk Theorem applies, but in which the set of payoffs to symmetric equilibria are bounded away from efficiency.

Lemma 5.1: Under condition 5.1, every profile α can be approximated by a sequence $\alpha^k \rightarrow \alpha$ such that π satisfies pairwise full rank at each α^k .

Proof: We defer the proof until later in this section.

Remark: This lemma is an extension of a result of Legros [1988], who studies a static partnership problem with a balanced-budget constraint. He shows that while there is no balanced-budget sharing rule that induces both partners to work with probability one, if one player randomizes between working and shirking while the other works with probability one, then the distributions over outcomes will be different if one player shirks than if the other does, and so work incentives can be provided while satisfying the balanced budget-constraint.

In addition to condition 5.1, which requires that pairwise full rank be satisfied by some profile, we will require that the weaker property of individual full rank hold at all pure strategy profiles.

Definition 5.2: $\pi_y(\cdot)$ has individual full rank at α if for all players i , the m_i vectors $\{\pi_y(a_i, \alpha_{-i})\}_{a_i \in A_i}$ are linearly independent.

Note that in order for $\pi_y(\cdot)$ to have individual full rank, the number of observable outcomes \bar{m} must be at least $\max_i m_i$.

Lemma 5.2: If $\pi_y(\cdot)$ has individual full rank at α , then α is enforceable.

Proof: This follows immediately from inspection of the linear system of inequalities (3.1').

Condition 5.2: The game satisfies the individual full rank condition if π has individual full rank for all pure action profiles.

Condition 5.3: The dimension of V^* is equal to n , the number of players.

Obviously the dimension of V^* is no greater than that of V ; it may however be less. Fudenberg-Maskin [1986a] give an example of a three-player game with observable actions and $\dim V^* = 1$ where the Folk Theorem does not obtain. Fudenberg-Maskin [1986b] construct a similar counter example in a two player game with $\dim V^* = 1$ where one player's actions are unobservable. Note that Theorem 5.1 holds vacuously for sets \hat{V} of dimension less than n , for such sets have an empty interior.

Theorem 5.1. (Folk Theorem) Under conditions 5.1, 5.2 and 5.3 for every closed set W in the relative interior of V^* there is a $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$, $W \subseteq E(\delta)$.

Proof: The proof invokes a series of lemmas developed below to verify that V^* satisfies weak enforceability on tangent hyperplanes (condition 4.1). Thus, any closed set in the interior of V^* consists of equilibrium payoffs for sufficiently large δ by theorem 4.1. Condition 5.3 implies that the interior and relative interior of V^* coincide.

It will be convenient to distinguish hyperplanes that are orthogonal to a coordinate axis, the "coordinate" hyperplanes, from the "non-coordinate" hyperplanes. Any non-coordinate hyperplane that is tangent to V^* weakly separates an extremal point $g(a)$ from V^* . From lemma 5.1, under condition 5.1 any profile can be approximated by a sequence α^k that satisfies pairwise full rank, and lemma 5.5 shows that profiles satisfying pairwise full rank can be enforced on non-coordinate hyperplanes, hence, weak enforceability is satisfied for the non-coordinate hyperplanes.

Next consider the hyperplane \bar{H}^i orthogonal to the i th coordinate axis and tangent to V^* at a point where player i 's payoff is maximal in V^* . This hyperplane weakly separates V^* from a point that is best for player i among all points in V . The corresponding profile a^i has the property that a^i_i is a best response to a^i_{-i} in the stage game, and so a^i is enforceable on \bar{H}^i by lemma 5.3 below.

Finally, consider the tangent hyperplane H^i that is orthogonal to the i th coordinate axis and tangent to V^* where i 's payoff is minimized over V^* . From lemma 5.4, there is a sequence $\{\alpha^k\}$ of enforceable profiles with $g_i(\alpha^k) \rightarrow v_i^*$ and α^k_i a best response to α^k_{-i} in the stage game. This implies that $\{\alpha^k\}$ has a subsequence whose payoffs converge to $g(m^i_1, m^i_{-1})$, which weakly separated from V^* by H^i , and lemma 5.2 implies that the α^k are enforceable on H^i . ■

Remark: An important step in the proof above is the inference that if profile a maximizes player i 's payoff over the feasible set, then a_i is a static best response to a_{-i} (that is, a best response in the stage game). This property fails in repeated games where some players are short-run (that is, play for only one period), as in Fudenberg-Kreps-Maskin [1989], because there the definition of the feasible set must incorporate the restriction that the short-run players play static best responses. Thus, in order to induce the short-run players to play the profile a_{-i} that maximizes player i 's payoff, player i may need to play an action a_i that is not a best response to a_{-i} . Not only does our proof fail in this case, but the Folk Theorem fails as well, unless the players' choices of randomizing probabilities are publicly observed, and not simply their choices of actions. Fudenberg-Kreps-Maskin characterize the equilibrium payoffs as $\delta \rightarrow 1$ with observed probabilities, and with unobserved probabilities and a single long-run player. Fudenberg-Levine [1989b] extend the techniques of this paper to obtain a characterization for the remaining case of unobserved probabilities and several long-run players.

Proof of Lemma 5.1:

For any profile α and each pair of players $i \neq j$, let $\pi^{i,j}$ be the matrix whose rows are the $m_i + m_j - 1$ vectors,

$$\left\{ \pi_y(a_i, \alpha_{-i})_{a_i \in A_i}, \pi_y(a_j, \alpha_{-j})_{a_j \in A_j} \right\}.$$

Fix a profile α' and a pair of players $i \neq j$, and let $\alpha^{i,j}$ be a profile such that $\pi_y(\alpha^{i,j})$ satisfies pairwise full rank for i and j . Then for each $\epsilon \in$

[0,1], define $\alpha(\epsilon)$ by $\alpha(\epsilon) = (1-\epsilon)\alpha'_k + \epsilon\alpha_k^{i,j}$ for all k . Thus $\alpha(\epsilon)$ is a polynomial in ϵ , and the rows of $\pi^{i,j}(\alpha(\epsilon))$ are polynomials in ϵ as well. Thus the determinant of any $m_i + m_j - 1$ square submatrix of $\pi^{i,j}(\alpha(\epsilon))$ is itself a polynomial in ϵ , and since $\pi^{i,j}(\alpha(1)) = \pi^{i,j}(\alpha^{i,j})$ has full rank, its determinant is not zero for some such submatrix.³ Now every polynomial is either identically zero or has a set of zeros of measure zero so the determinant of $\pi^{i,j}(\alpha(\epsilon))$ is nonzero for almost all ϵ . Moreover, if $\pi^{i,j}(\alpha(\epsilon))$ has full rank, $\pi^{i,j}(\alpha)$ has full rank for all α in an open neighborhood of $\alpha(\epsilon)$. By considering ϵ near zero, we conclude that there is an open neighborhood of the fixed profile α' that satisfies pairwise full rank for player i,j , and since α was arbitrary we conclude there is an open, dense set of profiles that satisfy pairwise full rank for i and j . As this is true for each pair i,j , and the intersection of open dense sets is open and dense, we conclude that there is an open dense set of profiles that satisfies pairwise full rank for all pairs i,j . ■

Lemma 5.3: If α_i is a static best response to α_{-i} and α is enforceable, then α can be enforced with continuation payoffs such that $w_i(y)$ is equal to any constant. That is, the $w(y)$ can be chosen to lie on any $(n-1)$ dimensional hyperplane orthogonal to the i th coordinate axis.

Proof: Since α is enforceable it is enforceable by some continuation* payoffs $w(y)$. Let $\bar{w}_i = \sum_{y \in Y} \pi_y(\alpha) w_i(y)$, and define $w'(y)$ by $w'_j(y) = w_j(y)$, $j \neq i$,

We thank Andreu Mas-Colell for suggesting this argument using polynomials.

and $w'_i(y) = \bar{w}_i$. Clearly α is enforceable by the $w'(y)$, which lie on a hyperplane orthogonal to the i th axis. Enforcement with small variation (Lemma 3.4) then implies α can be enforced on any such hyperplane. ■

We now show that each player i 's minmax value can be weakly enforced on a hyperplane orthogonal to the i th coordinate axis.

Lemma 5.4: (Weak Enforceability of the Minmax) Under Condition 5.2 for each player i there is a sequence of enforceable mixed strategy profiles α^k such that $g_i(\alpha^k) \rightarrow v_i^*$ and α_i^k is a static best response to α_{-i}^k .

Proof: Let m_i^1 be the minmax actions against player i , and let m_i^1 be a best response for player i to m_{-i}^1 that is not weakly dominated by an alternative response. Define $\alpha = (m_i^1, m_{-i}^1)$, and for each player $j \neq i$, let $\pi^j(\alpha_{-j})$ be the matrix whose rows are the vectors $\pi_y(a_j, \alpha_{-j})$. Then $\pi^j(\alpha_{-j})$ is a convex combination of the matrices whose rows are $\pi_y(a_j, a'_{-j})$, for $a'_{-j} \in \text{support}(\alpha_{-j})$, and each of these matrices has full rank from the individual full rank assumption. As in the proof of Lemma 5.1, the fact that each matrix has full rank implies that almost every convex combination of them does. Moreover, there is an open dense neighborhood of U of α_{-i} such that m_i^1 is a best response to all α'_{-i} in U . Therefore, there is a sequence $\alpha^k \rightarrow \alpha$ such that $\pi^j(\alpha_{-j}^k)$ has full rank for each player $j \neq i$, and m_i^1 is a best response to each of the α_{-i}^k . Therefore each α^k is enforceable, and $g_i(\alpha^k) \rightarrow g_i(m_i^1)$. ■

The remaining case involves hyperplanes that are not orthogonal to a coordinate axis, which we call non-coordinate hyperplanes.

Lemma 5.5: If π satisfies the pairwise full rank condition at α , then α is enforceable on non-coordinate hyperplanes.

Proof: Consider an $n-1$ dimensional linear subspace H that is not orthogonal to a coordinate axis. Without loss of generality, we may assume this has the form $\hat{w}_n(y) = \sum_{i=1}^{n-1} \beta_i \hat{w}_i(y)$. Moreover, since H is not orthogonal to a coordinate axis, we may assume $\beta_1 \neq 0$. By lemma 3.4, it suffices to consider the case $\delta = 1/2$. Because α is enforceable, solutions $(\hat{w}_i(y))$ to (3.1') and (3.2') exist for players $i=2, \dots, n-1$. Let $\hat{w}^* = \sum_{i=2}^{n-1} \beta_i \hat{w}_i(y)$. We must find $\hat{w}_1(y)$ and $\hat{w}_n(y)$ such that:

$$(i) \quad \sum_y \pi_y(\alpha) \hat{w}_1(y) = 0$$

$$(ii) \quad \sum_y \pi_y(a_1, \alpha_{-1}) \hat{w}_1(y) = -\hat{g}_1(a_1, \alpha) \text{ for } a_1 \in \text{support}(\alpha_1)$$

$$\sum_y \pi_y(a_1, \alpha_{-1}) \hat{w}_1(y) \leq -\hat{g}_1(a_1, \alpha) \text{ for } a_1 \notin \text{support}(\alpha_1)$$

$$(iii) \quad \hat{w}_n(y) = \beta_1 \hat{w}_1(y) + \hat{w}^*(y)$$

$$(iv) \quad \sum_y \pi_y(\alpha) \hat{w}_n(y) = 0$$

$$(v) \quad \sum_y \pi_y(a_n, \alpha_{-n}) \hat{w}_n(y) = -\hat{g}_n(a_n, \alpha) \text{ for } a_n \in \text{support}(\alpha_n).$$

$$\sum_y \pi_y(a_n, \alpha_{-n}) \hat{w}_n(y) \leq -\hat{g}_n(a_n, \alpha) \text{ for } a_n \notin \text{support}(\alpha_n).$$

Constraints (i),(ii),(iii) and (iv) require that $\hat{w}_i(y)$ satisfy condition (3.1) for $i=1,n$, and constraint (iii) ensures $\hat{w}(y) \in H$. Constraint (iv) is redundant given (i), (iii), and the fact $\sum_y \pi_y(\alpha) \hat{w}^*(y) = 0$.

Substituting (iii) into (v), we have

$$\begin{aligned} \text{(vi)} \quad & \beta_1 \sum_y \pi_y(a_n, \alpha_{-n}) \hat{w}_1(y) = -g_n(a_n, \alpha) - \sum_y \pi_y(a_n, \alpha_{-n}) \hat{w}^*(y) \text{ for } a_n \in \text{support}(\alpha_n) \\ & \beta_1 \sum_y \pi_y(a_n, \alpha_{-n}) \hat{w}_1(y) \leq -g_n(a_n, \alpha) - \sum_y \pi_y(a_n, \alpha_{-n}) \hat{w}^*(y) \text{ for } a_n \notin \text{support}(\alpha_n) \end{aligned}$$

Thus it suffices to find $\hat{w}_1(y)$ satisfying (i), (ii), and (vi). Moreover, given any solution $\tilde{w}_1(y)$ to constraints (ii) and (vi), we can obtain a solution that satisfies (i) as well by setting $\hat{w}_1(y) = \tilde{w}_1(y) - \sum_y \pi_y(\alpha) \tilde{w}_1(y)$. Finally, observe that a solution to (ii) that solves (vi) with equality for all but one $a'_n \in A_n$ solves (vi) for a'_n as well, since

$$\begin{aligned} & \alpha(a'_n) \beta_1 \sum_y \pi_y(a'_n, \alpha_{-n}) \hat{w}_1(y) = -\beta_1 \sum_{a_n \neq a'_n} \alpha(a_n) \sum_y \pi_y(a_n, \alpha_{-n}) \hat{w}_1(y) \\ & - \sum_{a_n \neq a'_n} \alpha(a_n) \left[\hat{g}_n(a_n, \alpha_{-n}) + \sum_y \pi_y(a_n, \alpha) \hat{w}^*(y) \right] \\ & = -\alpha(a'_n) \hat{g}_n(a'_n, \alpha) - \alpha(a'_n) \sum_y \pi_y(a'_n, \alpha) \hat{w}^*(y) \end{aligned}$$

That is, one of the $(m_1 + m_n)$ constraints in (ii) and (vi) is redundant. The assumption of pairwise full rank is precisely that for each pair of players this system is has rank $(m_1 + m_n - 1)$. Hence (ii) and (vi) can be solved with exact equality. ■

It may be helpful to compare examples where the hypotheses of Theorem 5.1 are and are not satisfied.

Example 5.1: Players $i = 1$ and 2 can either work ($a_i = 1$) or not work ($a_i = 0$). The work decisions are unobservable and jointly result in output (which is observable) $y = 0$ or 12 . The probability of output 12 is $\frac{2}{3}$ if both players work, $\frac{1}{3}$ if only one does, and 0 if neither does. Player i 's utility of output $= y$ is $\frac{1}{2}y - 3a_i$. Because this game has only two observable outcomes, the pairwise full rank condition (condition 5.1) fails (although the individual full rank condition -- condition 5.2 -- is satisfied). Indeed, the Folk Theorem does not apply to this game: although the payoffs $(1,1)$ (corresponding to both players working) are feasible and Pareto dominate the minmax point $(0,0)$, any payoffs (v_1, v_2) belonging to $E(\delta)$ must satisfy $v_1 + v_2 \leq 1$ regardless of the value of δ . (This can be shown using arguments similar to those in Radner-Myerson-Maskin [1986]).

By contrast, consider the following game:

Example 5.2: Players have the same choices and payoffs as in Example 5.1. Now, however, there are three output levels: 12 , 6 , and 0 . If both players work, the probabilities of these outcomes are $\frac{1}{2}$, $\frac{1}{3}$, and $\frac{1}{6}$ respectively. If only one works, the probabilities are $\frac{1}{6}$, $\frac{1}{3}$, and $\frac{1}{3}$, and if neither works the probability of 0 is 1 . Notice that at the profile where one player works and the other does not, the probability matrix corresponding to the definition of pairwise full rank is

$$\begin{bmatrix} \frac{1}{6} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{6} \\ \frac{1}{6} & \frac{1}{2} & \frac{1}{3} \end{bmatrix}$$

Because this matrix has full rank, condition 5.1 is satisfied, and so Theorem 5.1 applies. We conclude that any feasible payoffs (v_1, v_2) with $v_i > 0$, belong to $E(\delta)$ for δ near enough 1 (despite the fact that the action profile where both players work violates pair-wise full rank).

We can modify the proof of lemma 5.5 to obtain a sharper version of theorem 5.1. Call an equilibrium strict if each player strictly prefers his equilibrium strategy to any other.

Theorem 5.2: Suppose that every outcome has positive probability under every action profile, that π satisfies pairwise full rank for all (i, j) at every pure action profile, that, for each player, there is a minmax profile in pure actions, and that condition 5.3 holds. Then for every closed set W in the interior of V^* there exists $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$, W is contained in the set of strict equilibria of $E(\delta)$.

Remark: Even if some minmax profiles entail mixed actions, we can obtain as strict equilibria all payoffs above the pure-action minmax levels as long as the other two hypotheses hold.

Proof: When pairwise full rank holds, the inequalities of conditions (ii) and (vi) in the proof of lemma 5.5 can be solved to hold strictly for any pure action profile. When all outcomes have positive probability, any deviation then induces a strict loss in payoff if the continuation payoffs are as specified. ■

In finite games, strict Nash equilibria are stable as singleton sets in the sense of Kohlberg-Mertens [1986]. Although there is not yet an established definition of stability for infinite-horizon games, it seems likely that strict equilibria would satisfy any such concept. Thus, theorem 5.2 casts doubt on conjectures that stability might restrict the force of the Folk Theorem. Moreover, strict equilibria are robust to small changes in the information structure of the game.

To obtain strict equilibria requires both that players do not randomize and that, for each player, every information set that could be reached if some opponent deviated is assigned positive probability by the equilibrium strategies. These conditions cannot be met in repeated games with observable actions. However, in such games we can solve (ii) and (vi) with exact equality for all actions, and thus induce players to play all their actions with positive probability in every period. We conclude that in such games all feasible, individually rational payoffs can be approximated by the payoff to a totally mixed equilibria. This observation is of interest because totally mixed equilibria, like strict ones, are stable in finite games. (However, it is less clear that all extensions of stability to infinite-horizon games stability will make the totally mixed equilibria stable.)

Theorems 5.1 and 5.2 show that any point w in the interior of V^* can be obtained as an equilibrium for large δ , but the theorems do not cover points on the efficient frontier. That is, under the hypotheses of the theorem efficient payoffs can be approximated by equilibria but may not be exactly attainable. This contrasts with the Folk Theorem for repeated games with observable actions, where efficient payoffs can be exactly attained for discount factors sufficiently close to one. The following theorem explains why in general theorem 5.1 cannot be strengthened to exactly obtain efficient payoffs.

Theorem 5.3 Let v be an extremal payoff of V . If for each a with $g(a) = v$ there exists player i and action a'_i such that $g_i(a'_i, a_{-i}) > g_i(a)$ and support $\pi_y(a'_i, a_{-i}) \subseteq \text{support } \pi_y(a)$, then for all $\delta < 1$ $v \notin E(\delta)$.

Remark: The inclusion of the supports hypothesized in theorem 5.2 is clearly satisfied if the support of $\pi_y(\cdot)$ is independent of a . When the hypothesis fails, player i can be deterred from playing a'_i by continuation payoffs that are very low outside the support $\pi_y(a)$.

Proof: Because v is extremal, the only sequence of feasible payoffs with average value v is the sequence where v is repeated each period. Consider an equilibrium giving rise to this sequence. Then the first period strategies must specify a profile a with $g(a) = v$, and the continuation payoff $w(y)$ must satisfy $w(y) = v$ for all $y \in \text{support } \pi_y(a)$. Since $g_i(a'_i, a_{-i}) > g_i(a)$, and support $\pi_y(a'_i, a_{-i}) \subseteq \text{support } \pi_y(a)$, player i will prefer a'_i to a_i . ■

6. Games with a Product Structure

So far we have focused on full rank hypotheses to ensure enforceability of actions. Radner [1985], however, obtained positive results in repeated principal-agent games without considering the kinds of informational conditions developed in Section 5. These games have special structure that can take the place of some of our informational conditions. The distribution over outcomes has a "product structure", meaning that the outcome y is a vector (y_1, \dots, y_n) , where the y_i are statistically independent and the distribution of each y_i does not depend on the actions of other players than i . Theorem 6.1 below shows that this condition is sufficient for a "Nash-threats" Folk Theorem, pertaining to payoffs that Pareto-dominate a static Nash equilibrium.

We first show that profiles corresponding to Pareto-efficient payoffs are enforceable regardless of the information structure under the standard assumption that payoffs "factor" as in Definition 2.1. As we have emphasized, enforceability alone is not sufficient for the Folk Theorem. In games with a product structure, though, enforceability implies enforceability on hyperplanes and so the Pareto-efficient payoffs are enforceable on hyperplanes. Of course, the minmax profiles are not necessarily Pareto optimal, and will not be enforceable without additional assumptions. But any static equilibrium is trivially enforceable on any hyperplane, and so we can use theorem 4.1 to conclude that all payoffs Pareto-dominating a static equilibrium can be attained for δ close to 1.

Theorem 6.1 covers the principal-agent games considered by Radner, where the principal's only action is a transfer to the agent. Section 7 applies theorem 6.1 to a class of repeated mechanism problems that also have

a product structure. Section 8 extends the results to games with a weaker kind of product structure when the actions of some players are observable.

We will now restrict attention to games that satisfy the factorization condition given in definition 2.1, that each player's payoff is completely determined by the public outcome y and his own action a_i .

Lemma 6.1: If payoffs satisfy the factorization condition, then all external payoffs on the Pareto frontier of v are enforceable.

Proof: If v is external and Pareto efficient then there are positive weights λ_i such that v solves $\max_{v \in V} \lambda \circ v$. Choose a profile a such that $g(a) = v$, and for all i , set

$$w_i(y) = [(1-\delta)/\delta] \sum_{j \neq i} [\lambda_j/\lambda_i] \bar{r}_j(y, a_j).$$

Then

$$(6.1) \quad (1-\delta)g_i(a'_i, a_{-i}) + \delta \sum_y \pi_y(a'_i, a_{-i}) w_i(y) =$$

$$[(1-\delta)/\lambda_i] \lambda_i g_i(a'_i, a_{-i}) + [(1-\delta)/\lambda_i] \sum_{j \neq i} \sum_{y \in Y} \lambda_j \pi_y(a'_i, a_{-i}) \bar{r}_j(y, a_j) =$$

$$[(1-\delta)/\lambda_i] \sum_{j=1}^n \lambda_j g_j(a'_i, a_{-i}).$$

Because $g(a) = v$, a solves $\max_{a \in A} \lambda \circ g(a')$, and thus for all players i , a_i solves $\max_{a'_i \in A_i} \lambda \circ g(a'_i, a_{-i})$. Thus a_i is maximized (6.1), implying a is enforced by $\{w_i(y)\}$. ■

Definition 6.1: The game has a product structure if we can write

$y = (y_1, \dots, y_n)$, where, for all i and $a = (a_i, a_{-i})$ and a'_i , the marginal distributions of y_i , $\pi_{y_i}(a)$, satisfies

$$(6.2) \quad \left\{ \begin{array}{l} \pi_{y_i}(a_i, a_{-i}) = \pi_{y_i}(a_i, a'_{-i}) \text{ for all } a_{-i}, a'_{-i} \in A_{-i}, \text{ and} \\ \pi_y(a) = \prod_{i=1}^n \pi_{y_i}(a_i) \end{array} \right.$$

The first condition says that the marginal distribution of y_i depends on a_i alone, and the second that the joint distribution over y is the product of the marginals. When the game has a product structure, one might expect that an enforceable profile can be enforced with continuation payoffs $(w(y))$ such that $w_i(y)$ depends only on y_i , that is, that the incentive constraints for the different players are not linked. The following lemma shows that this expectation is justified.

Lemma 6.2: If the game has a product structure, then any pure action that is enforceable is enforceable on non-normal hyperplanes.

Proof of Lemma 6.2: Fix a , and as in the proof of Lemma 5.5, pick an $n-1$ dimensional linear subspace H that is not orthogonal to a coordinate axis, having the form $\sum_{i=1}^n \beta_i \hat{w}_i(y) = 0$, with $\beta \in \mathbb{R}$, and where, without loss of generality we can assume that $\beta_1 \neq 0 \neq \beta_2$. Also set $\delta = 1/2$. We must solve

$$\begin{aligned}
(6.3) \quad (a) \quad & \sum_y \pi_y(a) \hat{w}_i(y) = 0 \\
(b) \quad & \sum_y \pi_y(a'_i, a_{-i}) \hat{w}_i(y) \leq -\hat{g}_i(a'_i, a_{-i}) \quad \text{for all } a'_i \\
(c) \quad & \sum_{i=1}^n \beta_i \hat{w}_i(y) = 0
\end{aligned}$$

Since profile a is enforceable, for each player i there are continuation payoffs $\tilde{w}_i(y)$ satisfying (6.3) (a) and (b). Moreover, since the game has a product structure, we can write these formulae as

$$\begin{aligned}
(6.3) \quad (a') \quad & \sum_{y_{-i}} \pi_{y_{-i}}(a_{-i}) \sum_{y_i} \pi_{y_i}(a_i) \tilde{w}_i(y_i, y_{-i}) = 0. \\
(b') \quad & \sum_{y_{-i}} \pi_{y_{-i}}(a_{-i}) \sum_{y_i} \pi_{y_i}(a'_i) \tilde{w}_i(y_i, y_{-i}) \leq -\hat{g}_i(a'_i, a_{-i}) \quad a'_i \neq a_i
\end{aligned}$$

Thus if we set

$$(6.4) \quad w_i^*(y_i) = \sum_{y_{-i}} \pi_{y_{-i}}(a_{-i}) \tilde{w}_i(y_i, y_{-i}),$$

then the $w_i^*(y_i)$ are continuation payoffs satisfying (6.3) (a) and (b).

Define

$$(6.5) \quad \hat{w}_1(y) = w_1^*(y_1) - \sum_{j=2}^n \beta_j w_j^*(y_j) / \beta_1,$$

$$\hat{w}_2(y) = w_2^*(y_2) - \beta_1 / \beta_2 w_1^*(y_1),$$

$$\hat{w}_j(y) = w_j^*(y_j) \quad \text{for } j \neq 1, 2$$

The \hat{w}_i satisfy 6.3 (c) by construction:

$$\sum_{i=1}^n \beta_i w_i - \beta_1 w_1^*(y_1) - \sum_{j=2}^n \beta_j w_j^*(y_j) + \beta_2 w_2^*(y_2) - \beta_1^* w_1^*(y_1) + \sum_{j=3}^n \beta_j w_j^*(y_j) = 0.$$

Moreover, for $i \neq 1, 2$ the \hat{w}_i clearly satisfy (6.3) a and b. For players 1 and 2, \hat{w}_i differs from w_i^* by the addition of terms whose expected value under $\pi_y(a)$ is zero, so 6.3 a is satisfied. Further, since y_i and y_{-i} are independent, for each a'_i the expected value of $\hat{w}_i(y)$ under $\pi_y(a'_i, a_{-i})$ is the same as the expected value of $w_i^*(y)$, so that \hat{w}_i satisfies (6.3) b. ■

Theorem 6.1 In a game with a product structure that satisfies the factorization condition, let \tilde{V} be the convex hull of the payoffs of a static Nash equilibrium and all Pareto-efficient payoffs that strictly Pareto-dominate this equilibrium. For any $v \in \text{interior}(\tilde{V})$ there is a discount factor $\underline{\delta}$ such that for all $\delta > \underline{\delta}$ $v \in E(\delta)$.

Proof: From lemma 6.1 and 6.2, extremal efficient profiles are enforceable on all non-coordinate hyperplanes. Both the static equilibrium and a profile yielding the maximal payoff for player i are enforceable on hyperplanes where i 's payoff is constant. Thus the set \tilde{V} satisfies enforceability on tangent hyperplanes, and the conclusion follows from theorem 4.1. ■

One class of games with a product structure is the principal-agent games considered by Radner [1985] where the principal's only action is a publicly observed transfer payment, and the distribution of each period's output depends only on the agent's effort. Theorem 6.1 thus extends

Radner's results to general specifications of the payoff functions. In particular, the principal can be allowed to be risk-averse, and the agent's utility function need not be separable in effort and income. Section 8 extends the conclusion of theorem 6.1 to games where the principal can, in addition to making transfers, also influence the distribution of output.

7. Incomplete Information

We now apply theorem 6.1 to a class of games of incomplete information. We then observe that this class includes mechanism design problems.

We first consider the following class of games. Each player first receives private information determining his type $z_i \in Z_i$, where Z_i is finite. The distribution of types is independent across players, and player i 's type has distribution $\pi_i(z_i)$, which is common knowledge. After learning their types, players choose actions $y_i \in Y_i$, which are publicly observed. Player i 's payoff $r_i(z_i, y)$ depends on the profile of actions and on his type but not on the types of others.

A pure strategy for player i is a map a_i from Z_i to Y . The public information is $y \in Y = Y_1 \times \dots \times Y_n$, and the distribution of outcomes corresponding to strategy profile a is

$$\pi_{yz}(a) = \begin{cases} \prod_{i=1}^n \pi_i(z) & y = a(z) \\ 0 & y \neq a(z) \end{cases}$$

Because the marginal distribution of y_i satisfies

$$\pi_{y_i}(a_i) = \sum_{z_i \in a_i^{-1}(y_i)} \pi_i(z_i),$$

it is clear that π has a product structure in the sense of definition 6.1. It is also clear that the payoffs satisfy factorization.

In the repeated version of the game, players' types in each period are independent draws than the fixed distribution $\pi_i(z_i)$. In each period, players observe the outcome y but not their opponents' types.

Note that these games do not satisfy the individual full rank condition, and a fortiori fail the stronger condition of pairwise full rank: Let z'_1, z''_1 satisfy $\pi_1(z'_1) \geq \pi_1(z''_1)$, and fix a pure strategy a_1 with $a_1(z'_1) = y'_1$, $a_1(z''_1) = y''_1$. This strategy yields the same distribution over outcomes as (mixed) strategy $\hat{\alpha}_1$, where $\hat{\alpha}_1(z''_1) = y'_1$, and $\hat{\alpha}_1(z'_1)$ takes on the value y''_1 with probability $\pi_1(z''_1)/\pi_1(z'_1)$, and y'_1 with probability $(\pi_1(z'_1) - \pi_1(z''_1))/\pi_1(z'_1)$. Thus many strategy profiles cannot be enforced. However, since the outcomes have a product structure, theorem 6.1 applies, and we obtain the Nash threat version of the Folk Theorem.

An example of a repeated game in this class is mechanism design under repeated adverse selection. To see this, imagine that the players' actions y_i are reports \tilde{z}_i of their types. Let S be a finite set of social states; and fix a mechanism $m: T \rightarrow S$. Player i 's payoff is $r_i^m(y, z_i) = u_i(m(\tilde{z}), z_i)$. Each player's payoff depends on his own type and on the social state, but not on the types of the other players.

One example of this kind of repeated adverse selection problem has been considered by Green [1987]. The types z_i are the players' (privately observed) endowments, and the mechanism m redistributes endowments to provide insurance against fluctuations. Bewley [1983], and Levine [1988], and Scheinkman-Weiss [1988] consider related models to study monetary theory.

Call a vector of payoffs efficient relative to m if it can be attained by some reporting strategy α and no other strategy profile $\tilde{\alpha}$ yields payoffs that are Pareto superior. Theorem 6.1 shows that some payoffs that are efficient relative to m can be approximated by equilibria of the repeated game for δ close enough to 1. However, such payoffs might be Pareto-dominated by payoffs that are attainable under a different mechanism m' .

This motivates a "universal mechanism game" which we define as follows. First, a status-quo state $\hat{s} \in S$ is specified. Each period, player i 's action y_i is a pair (\tilde{z}_i, m) where \tilde{z}_i is a report of his type and m is a mechanism, that is, a map from T to S . (Note that since S and T are finite, so is the space of mechanisms.) If all players announce the same map m , each player i 's payoff is $r_i(y, z_i) - u_i(m(\tilde{z}), z_i)$; otherwise, the status-quo state \hat{s} occurs, and player i 's payoff is $r_i(y, z_i) - u_i(\hat{s}, z_i)$.

In any perfect public equilibrium of this game, player i 's expected continuation payoff, public history h and current type z_i , must be at least what he would obtain under the status quo rule. Furthermore, the efficient payoffs of the universal game are first-best payoffs: No choice of mechanism and reporting rule, incentive compatible or not, yields Pareto-superior payoffs. Consequently if these payoffs can be approximated by equilibria when δ is close to one, we can conclude that repeated adverse selection has negligible efficiency costs when players are sufficiently patient. Since Theorem 6.1 applies to these games, this is indeed the case.

To help interpret this result, consider Green's [1987] repeated insurance game discussed above. Imagine that there are two identical agents, one good, that utility of consumption is $u_i(c_i) = \ln(c_i)$, that c_i must equal 2, 3, or 4, and that each agent's endowment z_i takes on the

values 2 and 4 with equal probability. The best symmetric allocation is for each agent to consume his endowment when endowments are identical, and for each to consume the average endowment of 3 when one endowment is high and the other is low. Theorem 6.1 says that the corresponding expected utility $1/4 \ln(2) + 1/2 \ln(3) + 1/4 \ln(4)$ can be approximated by an equilibrium when δ is close to 1. The conclusion is not that risk and/or risk aversion per se become unimportant, or that the players can do as well as if they had access to a productive storage technology: No mechanism can approximate the utilities corresponding to each player consuming his average endowment. Repeated play does not reduce the social risk corresponding to a random total endowment. Rather, repeated play reduces, and for large δ virtually eliminates, the incentive costs of inducing truthful revelation.

8. Principal-Agent Games

In this section we consider the class of games where the actions of at least one player are observable. This class includes the standard principal-agent model, where the agent is subject to moral hazard but the principal's actions -- which take the form of transfers to the agent -- can be observed.

As we noted in the introduction, the additional freedom provided by the observability of one player's actions often permits many (nearly) efficient points to be sustained as equilibria in the repeated game even when our full rank conditions are not satisfied. This is true, in particular, in two-player games (with one player's actions observable), where any payoff vector that Pareto dominates a static equilibrium can arise in equilibrium provided that the part of the boundary lying above that equilibrium is downward-sloping. Let the agent be player 1 and the principal

be player 2. We say that player 2's actions are observable if for all a the support $(\pi_y(a_1, a_2)) \cap \text{support}(\pi_y(a_1, a'_2)) = \emptyset$ for all $a_2 \neq a'_2$, so that any observed outcome is consistent with exactly one action by player 2.

As in theorem 6.1, let \tilde{V} be the convex hull of payoffs from a static equilibrium and the Pareto efficient points that Pareto-dominate it.

Theorem 8.1: In a two-player game where player 2's actions are observable. For any payoff vector $v \in \text{interior}(\tilde{V})$, there exists $\underline{\delta}$ such that for all $\delta > \underline{\delta}$, $v \in E(\delta)$.

Proof: Because the outcome reveals player 2's actions, the continuation payoffs corresponding to deviations by player 2 have probability zero so long as player 2 never deviates. An inspection of the proof of theorem 4.1 shows that the condition of enforceability on tangent hyperplanes can be weakened to the condition that all continuation payoffs that have positive probability under the profile being enforced lie on the hyperplane, and that all continuation payoffs either lie on the hyperplane or correspond to the static equilibrium. Now fix a static equilibrium, with payoffs v^e , and consider a smooth set w in the interior of the convex hull of the payoffs v^e of the static equilibrium and the Pareto frontier. We will show that for all $\hat{v} \in w$ there is a $v' \in V$ and an α with $g(\alpha) = v'$ such that v' is separated from w by the tangent line $H(\hat{v})$ to w at \hat{v} , and such that the continuation payoffs $w(y)$ are in $H(\hat{v})$ if $y \in \bigcup_{a_1} \text{support}(\pi_y(a_1, \alpha_2)) = \text{support}(\pi(\alpha_2))$ and $w(y) = v^e$ otherwise.

Consider $\hat{v} \in W$ that maximizes player 1's payoff. Let (\bar{a}_1, \bar{a}_2) be the "maximax" actions that sustain player 1's maximal payoff and let $(\bar{v}_1, \bar{v}_2) = g(\bar{a}_1, \bar{a}_2)$. Clearly $H(\hat{v})$ separates w and (\bar{v}_1, \bar{v}_2) . Moreover, because (\bar{a}_1, \bar{a}_2) is 1's maximal payoff, we can take $w(y) = \hat{v}$ for all $y \in (\text{support } \pi(\alpha_2))$.

Next consider any other $\hat{v} \in W$. If $H(\hat{v})$ separates w and the static equilibrium payoffs v , then we can take $w(y) = \hat{v}$ for all $y \in (\text{support } (\alpha_2))$. If not, then because w lies in the interior of the convex hull v^e and the Pareto frontier, there exists a Pareto optimum v' which is separated from W by $H(\hat{v})$. Let (a'_1, a'_2) be the actions associated with v' . From Lemma 6.1 there exist continuation payoffs $\{w_1(y)\}$ for player 1 that enforce a'_1 . Because $H(\hat{v})$ is not parallel to 2's axis we can choose $w_2(y)$ such that $(w_1(y), w_2(y)) \in H(\hat{v})$ for all $y \in (\text{support } (\alpha_2))$. Q.E.D.

A boundary point of V is on the outer boundary with respect to v^e if it belongs to a ray emanating from the equilibrium payoffs (v_1^e, \dots, v_n^e) with positive slope in all directions. The outer boundary is downward-sloping if it consists only of (strict) Pareto optima. An immediate implication of Theorem 7.1 is the following "Folk Theorem-like" result.

Corollary 8.1: If we add the hypothesis of a downward-sloping outer boundary to Theorem 8.1, then for any vector v in the interior of V^* that Pareto dominates a static Nash equilibrium there exists $\underline{\delta}$ such that, for all $\delta > \underline{\delta}$, $v \in E(\delta)$.

The corollary applies, in particular, to games where the observed player's action is a monetary transfer to the agent, as in Radner [1985]. Both theorem 8.1 and its corollary clearly extend to multiplayer games where the actions of only one player are unobservable. In fact, we can extend them to games with several players whose actions are unobservable, as long as the observable outcomes have a product structure.

Theorem 8.2: Consider an n-player game where the actions of players $\ell+1, \dots, n$ are observable and those of $1, \dots, \ell$ are not. Suppose that we can write outcomes y and $y = (y_1, \dots, y_\ell)$, where for all $i=1, \dots, \ell$, π_y depend only on a_i and $(a_{\ell+1}, \dots, a_n)$ and $\pi_y(a) = \prod_{i=1}^m \pi_{y_i} a_i$. Then the conclusion of Theorem 8.1 obtains.

When the product structure hypothesis fails, not all efficient payoffs Pareto dominating a static equilibrium can, in general, be approximated by repeated game equilibria, even when one player's actions are observable and the outer boundary is downward sloping. From Lemma 6.1, any Pareto optimal actions are enforceable; indeed, from the proof of Theorem 8.1, they are enforceable on any hyperplane that is not parallel to the observable player's axis. The difficulty is enforcing them on the coordinate hyperplanes. In the proof of theorem 8.1, we used the maximax action profile (for the unobservable player) for one of the two coordinate hyperplanes. But such a profile may not be enforceable on a coordinate hyperplane once there are several players subject to moral hazard.

Nonetheless, as we show in Fudenberg, Levine, and Maskin [1989], the efficient payoffs Pareto-dominating a static equilibrium are attainable as separated game equilibria in a large subclass of games where the observable player can make transfers of money or some other private good.

9. Discussion

Our method of constructing equilibria is based on the dynamic-programming decomposition of equilibrium payoffs into "payoffs today" $g(\alpha)$ and "continuation payoffs" $w(y)$, the point being that if the continuation payoffs are known to be equilibrium values, and profile α is enforceable by the $w(y)$, then $(1-\delta)g(\alpha) + \delta \sum \pi_y(\alpha)w(y)$ is an equilibrium payoff as well. This approach prompts the question of how the equilibrium set $E(\delta)$ is related to the payoffs that can be enforced by some continuation payoffs in the feasible set $V(\delta)$, whether or not these payoffs are themselves equilibria. This latter set, which we denote $C(\delta)$, consists of the first-period payoffs that the players could attain if they could commit themselves (say, through binding contracts) to the continuation payoffs $w(y) \in V(\delta)$ as a function of the outcome y . Let $N(\delta)$ be the set of Nash equilibrium payoffs; Nash equilibrium requires that each $w(y)$ be a Nash equilibrium payoff if y has positive probability under the first period profile α . Since the definition of $C(\delta)$ imposes fewer constraints on the continuation payoffs than that of $N(\delta)$, it is clear that $E(\delta) \subseteq N(\delta) \subseteq C(\delta)$. From the study of repeated games with observable actions we know that for fixed δ less than one the first inclusion can be strict, and brief reflection should convince the reader that the second inclusion can be strict as well.

Let us compare the behavior of these sets as $\delta \rightarrow 1$. The first thing to notice is that if δ is too small, many payoffs in V may not even be feasible, whereas $V(\delta) = V$ if $\delta > 1 - 1/d$, where d is the number of vertices of V (Sorin [1986]). This discrepancy reflects the fact that one effect of repeated play with patient players is to allow "intertemporally transferable utility" in the interior of V , even if the stage game does not provide much flexibility for utility tradeoffs. A second effect of sending δ towards one is to make the possible gain from a one-period deviation smaller relative to a given change in the expected continuation payoff. This is the basis for intuition that the Folk Theorem "should" obtain in repeated games with observable actions.

Let $E(1)$, $N(1)$ and $C(1)$ denote the limiting values of the corresponding sets as $\delta \rightarrow 1$. (A point is in $E(1)$ if it is in $E(\delta)$ for all sufficiently large δ .) The perfect Folk Theorem asserts that $E(1) \supseteq \text{interior}(V^*)$; in the Nash version $N(1) \supseteq \text{interior}(V^*)$. Clearly, a necessary condition for either Theorem is $C(1) \supseteq \text{interior}(V^*)$. This inclusion obtains if every pure action profile is enforceable, and thus if $\pi_y(a)$ satisfies individual full rank at all pure strategies a ; it fails if for some individually rational vertex v of V the corresponding profile a is not enforceable.

Even when all the pure action profiles are enforceable, the Folk Theorems may still fail. For example, in repeated games with observable actions, the perfect Folk Theorem fails without the full-dimensionality condition (5.3). Radner-Myerson-Maskin [1986] consider a game with two-sided moral hazard where $C(1) \supseteq \text{interior}(V^*)$ and condition 5.3 is satisfied and yet even the Nash Folk Theorem fails. Although individual full rank implies that any Pareto-efficient profile can be enforced with nearly continuation

payoffs when δ is close to one, these continuation payoffs may not be equilibria without the stronger condition of pairwise full rank (or some alternative assumption such as a product structure or the existence of an observable player).

To conclude, let us mention some limitations of our results. The Folk Theorem shows that many payoffs are equilibria in the limiting case of extreme patience. One consequence of looking at limit results is that even a "small amount" of the right kind of information is sufficient. For example, Condition 5.1 requires that at some profile α , the matrix $\pi^{i,j}(\alpha)$ have full rank, but the matrix is allowed to be arbitrarily close to one that is singular. For a fixed discount factor, δ , the equilibrium set is smaller when the outcomes reveal less information in the sense of garbling. (Kandori [1988]). Intuitively, as the information deteriorates, the required variation in the continuation payoffs grows, and for a fixed δ the necessary variation may not be feasible. When the outcomes are almost uninformative, the δ required to approximate efficient payoffs may be very near to 1.

Furthermore, as Abreu, Pearce, and Milgrom [1988] have stressed, it is misleading to interpret $\delta \rightarrow 1$ in our model as a consequence of the time interval between periods growing shorter. Our results hinge on δ being large relative to a given information structure while quite plausibly, the information revealed by the outcomes is a function of the period length. Abreu, Pearce, and Milgrom give an example where the Folk Theorem does not obtain in the limit of shorter time periods because the informativeness of the outcomes decays too quickly.

REFERENCES

- Abreu, D. (1988), "On the Theory of Infinitely Repeated Games with Discounting", Econometrica, 56, 383-396.
- Abreu, D., D. Pearce, and P. Milgrom (1988), "Information and Timing in Repeated Partnerships", mimeo.
- Abreu, D., D. Pearce and E. Stacchetti (1986), "Optimal Cartel Monitoring with Imperfect Information", JET, 39, 251-69.
- Abreu, D., D. Pearce and E. Stacchetti (1988), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", Harvard.
- Aumann, R., and L. Shapley (1976), "Long-Term Competition: A Game-Theoretic Analysis", mimeo, Hebrew University.
- Bewley, T. (1983), "A Difficulty with the Optimum Quantity of Money", Econometrica 51, 1485-1504.
- Friedman, J. (1971), "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.
- Fudenberg, D., and D. Levine (1989a). "An Approximate Folk Theorem for Repeated Games with Private Information," mimeo.
- Fudenberg, D., and D. Levine (1989b). "Perfect Public Equilibrium with Long and Short Run Players, mimeo."
- Fudenberg, D., and D. Levine (1983), "Subgame-perfect Equilibria of Finite and Infinite Horizon Games", Journal of Economic Theory 31, 251-268.
- Fudenberg, D., and D. Levine (1988), "The Folk Theorem in Repeated Games with Unobservable Actions", mimeo.
- Fudenberg, D. and E. Maskin (1986a), "The Folk Theorem for Repeated Games with Discounting and Incomplete Information", Econometrica, 54, 533-54.
- Fudenberg, D. and E. Maskin (1986b), "Discounted Repeated Games with One-Sided Moral Hazard", MIT.
- Fudenberg, D., and E. Maskin [1988], "On the Dispensability of Public Randomizations," mimeo.
- Green, E. (1987), "Lending and the Smoothing of Uninsurable Income", in E. Prescott and N. Wallace (eds.), Contractual Arrangements for Intertemporal Trade, Minneapolis: University of Minnesota Press.

- Green, E. and R. Porter (1984), "Noncooperative Collusion under Imperfect Price Information", Econometrica, 52, 87-100.
- Kandori, M. (1988), "Monotonicity of Equilibrium Payoff Sets with Respect to Observability in Repeated Games with Imperfect Monitoring", mimeo.
- Kohlberg, E. and J.-F. Mertens (1986), "On the Strategic Stability of Equilibrium," Econometrica, 54, 1003-1034.
- Legros, P. (1988), "Sustainability in Partnerships," mimeo, California Institute of Technology
- Levine, D. (1988), "Asset Trading Mechanisms and Expansionary Policy", mimeo, UCLA.
- Matsushima, H. (1988), "Efficiency in Repeated Games with Imperfect Monitoring", mimeo.
- 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, 1127-1148.
- Radner, R., (1985), "Repeated Principal Agent Games with Discounting", Econometrica, 53, 1173-1198.
- Radner, R., (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. (1979), "Equilibrium in Supergames with the Overtaking Criterion", Journal of Economic Theory 21(1), 1-9.
- Rubinstein, A. and M. Yaari (1983), "Repeated Insurance Contracts and Moral Hazard", Journal of Economic Theory, 30, 74-97.
- Scheinkman, J., and L. Weiss (1988), "Borrowing Constraints and Aggregate Economic Activity", Econometrica 54, 23-45.
- Sorin, S. (1986), "On Repeated Games with Complete Information", Mathematics of Operations Research, 11, 147-160.

APPENDIX

We show that there can be sequential equilibria with payoffs, not in V^* . Consider the following three person game. Players 1 and 2 choose pairs, the first element being Up (U) or Down (D), the second being Heads (H) or Tails (T). Player 3 chooses Right or Left. Players 1 and 2 receive zero regardless of what happens, while player 3's payoff is -1 if 2 and 3 both chose Up and he went Right, or if 2 and 3 both chose Down and he went Left; otherwise he gets zero. Notice that the choice of H or T is payoff irrelevant.

The payoffs to player 3 may be summarized as

	3		
	L	R	
1U, 2U	-1	0]
1U, 2D	0	0	
1D, 2U	0	0	
1D, 2D	0	-1	

3's payoff

The minmax for player 3 is $-1/4$, which is achieved by 1 and 2 randomizing 50-50 between Up and Down. If, however, players 1 and 2 could jointly correlate so as to both play U $1/2$ of the time both D $1/2$ of the time, player 3 would only get $-1/2$.

Now suppose player 3's choices are observable, that whether 1 and 2 play Up or Down is observable, but that the common public information about their choice of H or T is only the total number of H chosen by both players. Formally, we have $y(t) = \sum_i I(a_i(t))$, where $I(\cdot)$ is the indicator function $I(H) = 1$, $I(T) = 0$. Notice that if $y(t)$ is 2 or 0, then the period- t actions of players 1 and 2 are public information. If $y(t) = 1$, then the period- t actions of players 1 and 2 are common knowledge for players 1 and 2 (since it is common knowledge that they each know their own action), but player 3 does not know which player played H.

Now consider the following strategies for players 1 and 2. "Randomize 1/2-1/2 between H and T in every period. If $y(t-1) = 1$ and player 2 played H, play D, if $y(t-1) = 1$ and player 2 played T, then play U. If $y(t-1) = 0$ or 2, play U with randomize 1/2-1/2 between U and D." Since players 1 and 2 are indifferent between all their strategies, they are willing to follow the strategies just described. As these strategies are independent of player 3's play, player 3's best response is to play in each period to maximize that period's expected payoff. Doing so yields an expected payoff of $-1/2$ when $y(t-1) = 1$. Thus player 3's expected normalized payoff facing these strategies converges to $-3/8$ as $\delta \rightarrow 1$, which is less than $v_3^* = -1/4$.

Date Due

JUL 11 1991

APR 14 2004

MIT LIBRARIES DUPL 1



3 9080 00567416 0

