



HIER

Harvard Institute of Economic Research

Discussion Paper Number 1961

The Nash Threats Folk Theorem With Communication and Approximate
Common Knowledge in Two Player Games

by

Drew Fudenberg and David K. Levine

June 2002

Harvard University
Cambridge, Massachusetts

This paper can be downloaded without charge from the:

<http://post.economics.harvard.edu/hier/2002papers/2002list.html>

The Nash Threats Folk Theorem With Communication and Approximate Common Knowledge in Two Player Games¹

Drew Fudenberg

David K. Levine²

First Version: May 24, 2002

This Version: June 18, 2002*

¹ We are grateful to seminar participants at the Cowles Conference on repeated game with private information whose research inspired this project, to Dan Hojman for careful proofreading, and to National Science Foundation Grants SES99-86170 and SES-0112018.

² Departments of Economics, Harvard University and UCLA.

1. Introduction

In repeated games with private monitoring, the players' beliefs about past play will gradually drift apart as the game goes on, so that an equilibrium based on common forecasts will gradually break up. One way to restore common belief about aspects of past play that matter for forecasts of opponents' strategies is for players to send "cheap-talk" messages to one another; if the messages are truthful they will form a public state that can be used to govern the players' strategies. However, one of the first results on using communication in this way is negative: Matsushima [1991] proved that payoffs of equilibria with truthful, incentive compatible revelation of the signals every period are bounded away from efficiency in two-player games with independent signals. Subsequently, Compte [1998] and Kandori and Matsushima [1998] proved a folk theorem for two-player games with independent signals, by considering strategies that only report truthfully every T periods, where T goes to infinity as the discount factor goes to 1.³

This paper shows how communication can yield to a Nash-threats folk theorem in two-player games with "almost public" information but without independent signals.⁴ Our proof is a combination of the idea that communication provides a public signal and an idea from Mailath and Morris [2002]. They provided a sufficient condition for a perfect public equilibrium of a game with public information to remain equilibrium when the information structure is perturbed to be almost public. We build on these ideas by introducing the possibility that the messages sent are coarser than the underlying private signals, which extends the class of games where our information conditions are satisfied.

A key hypothesis of the Mailath and Morris folk theorem is that the equilibrium strategies for the public information game depend only on a finite history of play. This

³ In addition, these papers, and also Ben-Porath and Kahneman [1996], proved folk theorems for games with at least three players. With three or more players the report of a third player can be used to tell who is misreporting. We should also note that there are a number of folk theorem and related results in the two-player case without communication: Sekiguchi [1997], Ely and Valimaki [2002], Piccione [2002] and Bhaskar and Obarra [2002] have all studied the prisoner's dilemma game without communication.

⁴ By "communication" we mean communication between the players, without the benefit of an intermediary. Aoyagi [2002] proves a folk theorem for games with a third-party mediator who receives private reports from the players and sends them non-binding instructions, as in the communication games of Forges [1986] and Myerson [1986]. See Kandori [2002] for a survey of studies of repeated games with private monitoring, both with and without communication.

implies that the strategies have a finite-automaton description. Consequently, when the information structure is close enough to public information, the current states of the automata are almost common knowledge. Since repeated games with perfect information have efficient equilibria with finite memory, the Mailath-Morris result yields a folk theorem for games of almost-perfect, almost-public information, but the hypothesis of this theorem need not be satisfied for general games of almost-public information.

Our starting point is a game of with publicly observed signals, and the Fudenberg, Levine and Maskin [1994] (FLM) result that a folk theorem holds if there is “sufficient” public information. When players observe private signals but make public announcements, there is the possibility of constructing "FLM-like" equilibria in which the players' actions depend only the announcements, but the FLM techniques cannot be immediately applied, because it is necessary for the equilibrium to provide incentives for the players to “report truthfully.”

The basic contribution of this paper is to show how this can be done when players receive signals that are highly but not perfectly correlated. As in Mailath and Morris [2002], this leads to the notion of “almost public” information – the information is not common knowledge, but players can be fairly confident that their opponents received the same information that they did. Although the results generalize from the two-person case, in many respects it is an advantage to have more than two players, because it is possible to build equilibria by comparing the reports of different players, and using “third parties” to effectively enforce contracts. For this reason we focus here on the two-player case, and show that even without third parties, we have a folk theorem when information is “almost public.”

2. The Model

In the *stage game*, each of two players $i = 1, 2$ simultaneously chooses an *action* a_i from a finite set A_i . We refer to vectors of actions, one for each player, as *profiles*. Player i 's payoff to an action profile a is $g_i(a)$. In addition, each player observes a private signal $z_i \in Z_i$ a finite set. We let $z = (z_1, z_2)$, $Z = Z_1 \times Z_2$. Each action profile $a = (a_1, a_2) \in A = A_1 \times A_2$ induces a probability distribution π_a over *outcomes* z . At the end of each stage of the game, players make *announcements* $y_i \in Y^*$, where

Y^* is a finite set that is the same for each player.⁵ A *profile of announcements* is a $y \in Y = Y^* \times Y^*$. A stage game strategy for player i , $s_i = (a_i, m_i)$, is a choice of stage game action a_i and a map $m_i : Z_i \rightarrow Y^*$ from private signal to announcements. We refer to $m = (m_1, m_2)$ as a message profile. We let S_i denote the space of player i 's stage game strategies.

In the repeated game, in each period $t = 1, 2, \dots$, the stage game is played. Both prior to the game, and following the announcements in each period, a public randomization $w \in [0, 1]$ is drawn from a uniform distribution. The *public history* at time t is

$$h(t) = (y(1), w(1), y(2), w(2), \dots, y(t), w(t))$$

The *private history* for player i at time t is

$$h_i(t) = (a_i(1), z_i(1), a_i(2), z_i(2), \dots, a_i(t), z_i(t)).$$

A *partial strategy* for player i is a sequence of maps $\sigma_i(t)$ mapping the public and private histories $h(t-1), h_i(t-1)$ to probability distributions over S_i . A *strategy* is an assignment of a partial strategy to each initial value of the public randomization. A *public strategy* is a strategy that depends on $h_i(t)$ only through $h(t)$. We denote by $h(0) = w(0)$ the public history consisting only of the initial public randomization, and $h_i(0)$ the null private history.

Observe that for each public history $h(t-1)$ the public profile σ induces a partial profile over the repeated game beginning at t . We denote this by $[\sigma | h(t-1)]$. Given a private partial profile σ , we can compute for each player i an expected average present value with discount factor $0 \leq \delta < 1$. We denote this by $G_i(\sigma, \delta)$. A *perfect public equilibrium* is a public strategy profile σ such that for any public history $h(t-1)$ and any private partial strategy $\tilde{\sigma}_i$ by any player i we have

$$G([\sigma | h(t-1)], \delta) \geq G((\tilde{\sigma}_i, [\sigma_{-i} | h(t-1)]), \delta).$$

Note that by standard dynamic programming arguments it is sufficient to consider deviations to public strategies.

⁵ The assumption that the two players have the same sets of possible announcements is motivated by our focus on almost-public information. Although we can always increase the size of the smaller message space, doing so involves a loss of generality because increasing the possible messages can increase the set of potential equilibria.

3. The Structure of Information

Given our assumption of a common announcement space, it is convenient to think of players agreeing if they make the same announcement as each other. We can think of Y^* as being the subset of Y in which $y_1 = y_2$, and refer to this as the *diagonal* of Y . Given a message profile m the information structure π induces a distribution over the diagonal of announcement profiles. We denote by

$$\pi_a^m(y^*) = \sum_{z|m_1(z_1)=m_2(z_2)=y^*} \pi_a(z),$$

the probability of the diagonal point (y^*, y^*) , and by

$$\pi_a^m(y^* | Y^*) = \frac{\pi_a^m(y^*)}{\sum_{y \in Y^*} \pi_a^m(y)}$$

the probability conditional on the diagonal of the joint announcement y^* . It is also convenient to define the probability of the opponent's message given a player's signal

$$\pi_a^m(y_{-i} | z_i) = \sum_{z_{-i}|m_{-i}(z_{-i})=y_{-i}} \pi_a(z_{-i} | z_i).$$

Note that for this to be well defined there must be positive probability of i receiving the signal z_i when the players play a .

Definition 1: *A game has ε public information with respect to m if for all action profiles a ,*

$$(1) \bar{\pi}_a^m \equiv \sum_{y \in Y^*} \pi_a^m(y) \geq 1 - \varepsilon$$

$$(2) \text{ if } \pi_a(z) > 0 \text{ then for all } \tilde{y}_2 \neq y_2 = m_1(z_1), \pi_a^m(y_2 | z_1) > \pi_a^m(\tilde{y}_2 | z_1)$$

When this condition is satisfied for some m and “small” ε we say that the game has “almost public information.” This condition says that most of the time, each player is fairly confident of the other player's message; in the limit case of 0-public information, the two players' messages are perfectly correlated, so that they are public information. This condition is closely related to the Mailath and Morris definition of “ ε -close to public monitoring,” but it is weaker in two ways. First of all, Mailath and Morris suppose that each player i 's private signal z_i lies in the same set as do the signals in the limiting

public-information game; in our setting this corresponds to $\#Y^* = \#Z_i$. Second, suppose that in the public information limit, every signal has strictly positive probability under every action profile, and that the distribution of each player's private signals is close to this limit. These conditions imply condition (1) above, and a stronger version of condition (2), namely that for $\tilde{y}_2 \neq y_2 = m_1(z_1)$, $\lim_{\varepsilon \rightarrow 0} \pi_a^m(y_2 | z_1) = 1$. Given the assumption that $\#Y^* = \#Z_i$, their conditions are equivalent to ours, but when there are many private signals corresponding to a given public message, our condition (2) is significantly weaker, as it allows the private signals to differ in how informative they are about the message the opposing player will send.

Note that our condition is easier to satisfy with coarse message maps m , and indeed it is vacuously satisfied if m_1 and m_2 are equal to the same constant; the condition will have force when combined with the assumption that the messages “reveal enough” about the action profile that generated the underlying signals. We will examine more closely below an example showing how coarser message maps lead to greater common knowledge.

To better relate our results to the literature, we will need the following auxiliary definitions:

Definition 2: *A game is ε -close to perfect monitoring if for all action profiles a there is a signal z such that $\pi_a(z) \geq 1 - \varepsilon$.*

Definition 3: *A game has independent monitoring if for all action profiles a , $\pi_a(z_1, z_2) = \pi_a(z_1)\pi_a(z_2)$.*

In the case of the “truthful reporting” studied by previous papers on communication in game with private monitoring, (1) of Definition 1 is satisfied whenever the game is ε -close to perfect monitoring. However, independent monitoring is inconsistent with condition (2) of Definition 1 except for the case of perfect monitoring, so in particular a game does not have ε public information with respect to truthful reporting when it has independent and almost-perfect monitoring.

Remark: Kandori and Matsushima consider an example where players observe binary signals. The players “babble” most of the time, and in every T th period they send the

message “Pass” or Fail”, with the condition for passing being that the fraction of “good” signals in the last T periods has been sufficiently high. If the players are constrained to play the same action in the T periods between meaningful messages, then the law of large numbers implies that for T large the T -period reporting game has almost perfect monitoring, and this precision would allow the construction of approximately efficient strategies.⁶ Of course players are not constrained to play the same action every period; the independence assumption is useful in verifying that doing so is optimal.

For a given a, m we can consider $\pi_a^m(\cdot | Y^*)$ as a row vector. For player i we can the construct a matrix $\Pi_a^{m,i}$ by stacking the row vectors corresponding to (\tilde{a}_i, a_{-i}) as \tilde{a}_i ranges over A_i . We can further stack the two matrices corresponding to the two players to get a $(\#A_1 + \#A_2) \times \#Y^*$ matrix Π_a^m . Notice that this matrix has two rows (both corresponding to a) that are identical.

Definition 2: *A game has pure-strategy pairwise full rank with respect to m if for every pure profile a the rank of Π_a^m is $(\#A_1 + \#A_2) - 1$.*

This condition is never satisfied in games such as Green and Porter [1984], where the two players have the same sets of feasible actions-, and the distribution of signals satisfies the symmetry condition that $\pi_{(j,k)} = \pi_{(k,j)}$ is symmetric in a , but it is satisfied for a set of probability measures π_a of Lebesgue measure 1.

4. The Nash Threats Folk Theorem

Let v^* be a static Nash payoff vector. We may conveniently normalize $v^* = 0$. Fix a sequence of games with common A, Z, Y^*, g , and with signal probabilities π^n .

Theorem: *Suppose there is a message profile m such that for each n , game n has ε^n public information with respect to m , that $\varepsilon^n \rightarrow 0$, that these games have common diagonal probabilities $\pi_a^{m,n}(\cdot | Y^*) = \pi_a^m(\cdot | Y^*)$, and that $\pi_a^m(\cdot | Y^*)$ has pure-strategy pairwise full rank with respect to m . Then there is a sequence $\gamma^n \rightarrow 0$ such that for any feasible vector of payoffs $v > 0$ there exists $\delta^* < 1$ such that for any n and all $\delta \geq \delta^*$*

⁶ Note that the T -period messages are still independent, so from Matsushima’s result for any fixed T the equilibrium payoff is bounded away from efficiency even as the discount factor goes to 1. However, the distance between this bound and the frontier goes to zero as monitoring becomes perfect.

there is a perfect public equilibrium in the game n with payoffs v^n satisfying $\|v^n - v\| < \gamma^n$.

Proof: Let Γ be the public information game without public randomization with the same payoff functions as in the original game, and with public signals Y^* generated by the common diagonal probabilities. Using the arguments from Fudenberg, Maskin and Levine [1994], one can show that for any $v > 0$ there exists $\hat{\delta} < 1$, a $K > 0$ and a set of payoffs V with $v \in V$, and $v_i' \geq \hat{v} > 0$ for all $v' \in V$, such that, for all $\bar{\delta} > \hat{\delta}$ and every $v' \in V$, there is a perfect public equilibrium σ^* with payoffs v' , such that:

- a) the continuation payoffs $w(y)$ lie in V for all $y \in Y$,
- b) $|w(y) - w(y')| \leq K(1 - \delta)$
- c) at every history the profiles either prescribe pure strategies or the strategies of a static Nash equilibrium with payoff 0.
- d) if the strategies prescribe playing the static equilibrium given history h_t , they prescribe playing that static equilibrium at all subsequent histories.

Given the base strategies, we will construct corresponding public strategies in game n by constructing a map χ from public histories $h^n(t)$ to histories of the same length in the public-information game Γ , or the symbol P (for punishment). We define the strategies in the game n by the action taken by the base strategy, or the static Nash strategies respectively. The map χ is defined as follows. We map the initial null history to the null history. Given that all histories of length $t - 1$ have been mapped, we define the map for length t histories $h^n(t) = (h^n(t - 1), y_t^n, w_t^n)$.

Case 1: If $\chi(h^n(t - 1)) = P$, $\chi(h^n(t)) = P$, so punishment is absorbing.

Case 2: If $\chi(h^n(t - 1)) \neq P$ and $y_{t1}^n = y_{t2}^n$, use the public randomization w_t^n so that with probability $p_t(y_t^n)$ (to be defined below) we have

$$\chi(h^n(t - 1), y_t^n, w_t^n) = (\chi(h^n(t - 1)), y_t^n)$$

and $\chi(h^n(t - 1), y_t^n, w_t^n) = P$ with probability $1 - p_t(y_t^n)$.

Case 3: If $\chi(h^n(t-1)) \neq P$ and $y_{t1}^n \neq y_{t2}^n$, let a be the action profile specified by the base strategies given the history $\chi(h^n(t-1))$, and use the public randomization w_t^n so that $\chi(h^n(t-1), y_t^n, w_t^n) = (\chi(h^n(t-1)), y^*)$ with probability $qp_t(y^*)\pi_a^m(y^* | Y^*)$ and $\chi(h^n(t-1), y_t^n, w_t^n) = P$ otherwise, where the value of q will be chosen below.

Notice that whenever these strategies call upon players to mix, the prescribed actions both at that time and in the future are those of the static Nash equilibrium, so incentive compatibility (for both actions and messages) is trivial, and the average present value to both players is 0. We therefore need consider only payoffs and incentive compatibility when the strategies call for a pure profile a .

Let $v_i(h(t))$ be the average present values from following the base strategies in Γ when the discount factor is $\bar{\delta}$. Then

$$v_i(h(t-1)) = (1 - \bar{\delta})g(a) + \bar{\delta} \sum_{y^*} \pi_a^m(y^* | Y^*) v_i(h(t-1), y^*).$$

The average present values from following the corresponding strategies in the game n when the discount factor is δ is

$$v_i^n(h(t-1)) = (1 - \delta)g(a) + \delta \left(\sum_{y^*} [\bar{\pi}_a^m + (1 - \bar{\pi}_a^m)q] p_t(y^*) \pi_a^m(y^* | Y^*) v_i^n(h(t-1), y^*) \right)$$

Suppose that p, q are chosen so that for some constant β^n independent of history and the player

$$(*) \quad \sum_{y^*} [\bar{\pi}_a^m + (1 - \bar{\pi}_a^m)q_t] p_t(y^*) \pi_a^m(y^* | Y^*) v_i(h(t-1), y^*) = \beta^n \sum_{y^*} \pi_a^m(y^* | Y^*) v_i(h(t-1), y^*).$$

If

$$\bar{\delta} = \delta \beta^n$$

we have

$$\frac{1 - \delta}{1 - \delta \beta^n} v_i(h(t-1)) = (1 - \delta)g(a) + \delta \beta^n \sum_{y^*} \pi_a^m(y^* | Y^*) \frac{1 - \delta}{1 - \delta \beta^n} v_i(h(t-1), y^*)$$

from which we conclude that

$$v_i^n(h(t-1)) = \frac{1 - \delta}{1 - \delta \beta^n} v_i(h(t-1)).$$

Now consider a sequence $\beta^n \rightarrow 1$. Then for $\delta > \hat{\delta}$ and n sufficiently large, we can find an equilibrium in the base game for discount factor $\bar{\delta} = \delta / \beta^n$. We then choose p, q consistent with β^n so as to preserve the incentive to play a_i , and to create an incentive to send the messages m_i . Suppose

$$p_t(y^*) \geq \underline{p} > 0.$$

We want to show that no combination of action and message choice in period t can improve player i 's payoff, assuming that the player conforms to the specified strategy in future periods.⁷

Step 1: First we show that for an appropriate choice of q it is optimal to use m regardless of the action chosen. Observe that m maximizes the probability of agreement by the condition 2 of the definition of ε^n -public information. Let us refer to choosing an announcement other than the one chosen by m as *sending a false announcement*. Let z^n denote the least decrease in the probability of agreement due to a false announcement. Since the set of announcements is finite, $z^n > 0$. Sending a false announcement may result in a better continuation equilibrium; however, since the continuation payoffs in the public information game satisfy $|w(y) - w(y')| \leq K(1 - \delta)$, the increase in average present value is at most

$$\frac{1 - \delta}{1 - \delta\beta^n} K(1 - \delta).$$

On the other hand, sending a false announcement increases the probability of punishment by $(1 - q)\mu^n$, resulting in an average present value loss of at least

$$\frac{1 - \delta}{1 - \delta\beta^n} (1 - q)z^n p_t \hat{v}.$$

We conclude that it is sufficient to prevent false announcements that

$$1 - q = \frac{K(1 - \delta)}{z^n \underline{p} \hat{v}}.$$

⁷ This is the familiar application of the principle of optimality. Although this is not a multistage game with observed actions, the same argument that works there suffices to establish a “one stage deviation principle” for the perfect public equilibria.

Step 2: When it is optimal to use m regardless of the action chosen, it is sufficient to show that the base strategy choice of action is optimal when m will be used.

The incentive constraints in the n game are

$$\delta \sum_{y^*} \{ [(\bar{\pi}_a^m + (1 - \bar{\pi}_a^m)q)\pi_a^m(y^* | Y^*) - (\bar{\pi}_{a'}^m + (1 - \bar{\pi}_{a'}^m)q)\pi_a^m(y^* | Y^*)] \times p_t(y^*)v_i^n(h(t-1), y^*) \} \geq (1 - \delta)(g(a') - g(a))$$

Substituting for v_i^n we get

$$\bar{\delta} \{ [(\bar{\pi}_a^m + (1 - \bar{\pi}_a^m)q_t)\pi_a^m(y^* | Y^*) - (\bar{\pi}_{a'}^m + (1 - \bar{\pi}_{a'}^m)q_t)\pi_a^m(y^* | Y^*)] \times \frac{p_t(y^*)}{\beta^n} v_i(h(t-1), y^*) \} \geq (1 - \bar{\delta})(g(a') - g(a))$$

and in addition (*) must hold.

From the base game, we know that the incentive constraint

$$\bar{\delta} \sum_{y^*} \{ [\pi_a^m(y^* | Y^*) - \pi_{a'}^m(y^* | Y^*)] \times v_i(h(t-1), y^*) \} \geq (1 - \bar{\delta})(g(a') - g(a))$$

is satisfied. By substituting (*) and using this inequality, it is sufficient that

$$(**) \quad \sum_{y^*} (\bar{\pi}_{a'}^m + (1 - \bar{\pi}_{a'}^m)q)\pi_{a'}^m(y^* | Y^*) \frac{p_t(y^*)}{\beta^n} v_i(h(t-1), y^*) \leq \sum_{y^*} \pi_{a'}^m(y^* | Y^*) v_i(h(t-1), y^*)$$

for all profiles a' , with equality for a .

By the pairwise full rank condition, for each player i and strategy profile a , we may find a vector $w_i(a, y^*)$ such that

$$\sum_{y^*} \pi_{a'}^m(y^* | Y^*) w_i(a, y^*) = e(a')$$

for all $a' = (a'_i, a_{-i})$, $e(a) = 0$ and for $a' \neq a$ $e(a) = -1$. If we take

$$\frac{p_t(y^*)}{\beta^n} = \frac{1}{\bar{\pi}_a^m + (1 - \bar{\pi}_a^m)q} + \lambda^n \frac{w_i(a, y^*)}{v_i(h(t-1), y^*)}$$

then we have equality in (**) for $a' = a$, and for $a' \neq a$

$$\begin{aligned}
& \sum_{y^*} \frac{\bar{\pi}_{a'}^m + (1 - \bar{\pi}_{a'}^m)q}{\bar{\pi}_a^m + (1 - \bar{\pi}_a^m)q} \pi_{a'}^m(y^* | Y^*) v_i(h(t-1), y^*) \\
& - \lambda^n (\bar{\pi}_{a'}^m + (1 - \bar{\pi}_{a'}^m)q) \leq \\
& \sum_{y^*} \pi_{a'}^m(y^* | Y^*) v_i(h(t-1), y^*)
\end{aligned}$$

or,

$$\begin{aligned}
& \sum_{y^*} \frac{(1-q)(\bar{\pi}_{a'}^m - \bar{\pi}_a^m)}{(1-q)\bar{\pi}_a^m + q} \pi_{a'}^m(y^* | Y^*) v_i(h(t-1), y^*) \\
& \leq \lambda^n (\bar{\pi}_{a'}^m + (1 - \bar{\pi}_{a'}^m)q_t)
\end{aligned}$$

and we may choose

$$\lambda^n = \max_{a,a'} \frac{(1-q)(\bar{\pi}_{a'}^m - \bar{\pi}_a^m)}{\bar{\pi}_a^m \bar{\pi}_{a'}^m} \max_{a,i} g_i(a).$$

Moreover,

$$\frac{p_t(y^*)}{\beta^n} \leq \frac{1}{(1-q)\min_a \bar{\pi}_a^m + q} + \lambda^n \frac{\max_{a,y^*} \{w_i(a, y^*), 0\}}{\hat{v}}$$

Call the right-hand side of this expression μ^n . If we choose $\beta^n = 1/\mu^n$, we know that $p_t(y^*) \leq 1$. Specifically,

$$p_t(y^*) = \left(\frac{1}{(1-q)\bar{\pi}_a^m + q} + \lambda^n \frac{w_i(a, y^*)}{v_i(h(t-1), y^*)} \right) \left(\frac{1}{(1-q)\min_a \bar{\pi}_a^m + q} + \lambda^n \frac{\max_{a,y^*} \{w_i(a, y^*), 0\}}{\hat{v}} \right)^{-1}$$

As $\varepsilon^n \rightarrow 0$ we have $\lambda^n \rightarrow 0$ and $\bar{\pi}_a^m \rightarrow 1$, so $p_t(y^*) \rightarrow 1$, and so for some \underline{p} and sufficiently large n we have $p_t(y^*) > \underline{p} > 0$.

Putting this together

$$\beta^n \geq 1 - \frac{K(1-\delta)}{z^n \underline{p} \hat{v}} (1 - \min_a \bar{\pi}_a^m) = 1 - L^n (1 - \delta)$$

where $L^n \rightarrow 0$. Consequently we have an equilibrium with average present value

$$\frac{1-\delta}{1-\delta\beta^n} v = \frac{1}{1+\delta L^n} v,$$

which is the desired result.

□

Remark: Notice that the proof can easily be adjusted so that the equilibrium, except when the static Nash equilibrium is played, is strict of order $(1 - \delta)\bar{U}$, where $\bar{g} = \max_{i,a} |g_i(a)|$. Observe also that if we perturb the game so that $\|g - g'\| \leq \varepsilon$ and $\|\pi - \pi'\| \leq \varepsilon$, then the average discounted value of particular strategies changes by at most $\varepsilon\bar{g}$.

Hence if $\varepsilon \leq (1 - \delta^*)^2$ the theorem continues to hold for the perturbed game and $\delta = \delta^*$. Tracing out the equilibrium payoffs corresponding to the various histories, it follows that the convex hull of this set is self-generating for $\delta = \delta^*$, and hence by the results of Fudenberg, Levine and Maskin [1994] for these are payoff of perfect public equilibria for all $\delta \geq \delta^*$.

In other words, we have the following corollary

Corollary: *Fix a message profile m , and suppose that $g^n \rightarrow g$, $\pi^n \rightarrow \pi$, that game n has ε^n public information with respect to m , that $\varepsilon^n \rightarrow 0$ and that $\pi_a^m(\cdot | Y^*)$, has pure-strategy pairwise full rank with respect to m . Then there is a sequence $\gamma^n \rightarrow 0$ such that for any feasible vector of payoffs $v > 0$ there exists $\delta^* < 1$ such that for any n and all $\delta \geq \delta^*$ there is a perfect public equilibrium in the game n with payoffs v^n satisfying $\|v^n - v\| < \gamma^n$.*

Remark: In particular, this covers the case in which the payoffs have the form $r_i(a, z)$ and the probabilities $\pi^n \rightarrow \pi$.

5. The Role of Public Information

A crucial element of these results is the fact that the announcements are public information. Since there is already a folk theorem for games of public information, the question arises as to whether or not it can be applied directly to the game with messages, or whether in fact a separate proof is needed. Here we briefly indicate why the Fudenberg, Levine and Maskin (FLM) result does not apply to the announcement game.

The FLM folk theorem is limited to the convex hull of the set of profiles that satisfy enforceability plus pairwise identifiability. Enforceability requires that there be some continuation payoffs, feasible or not, that make it optimal for each player to play his portion of the profile. Pairwise identifiability is a weakening of the pairwise full-rank

condition that specifies that it is possible to statistically discriminate between a deviation by player 1 from one by player 2.

Our goal is to establish that in the private information game with announcements, the only profiles that satisfy enforceability plus pairwise identifiability are the static Nash equilibrium, and so the Folk Theorem with public information does not apply.

Fix a profile, including a strategy for sending messages. This determines for each player a probability distribution over messages sent. We refer to this as the marginal. One thing a player could do is to randomize his announcements independent of his private information, in such a way that the marginal distribution of messages is preserved; call this “faking the marginal.” Unless the given profile called for players to ignore their private signals, pairwise identifiability fails, because player one faking his marginal and player two faking hers are observationally equivalent, so the FLM result does not apply, and if we restrict attention to strategies where players make meaningless reports, the only public equilibria will have a static equilibrium outcome in every period.

6. Information Aggregation

One feature of the main theorem is that it allows the possibility that players aggregate information by making the same announcement corresponding to several different private signals. For example, we might imagine two partners who each receive stochastic output depending upon the effort taken by each of them. If the output is perfectly correlated, we have a game of public information; if output is imperfectly correlated we have a game of private information. If output occurs in discrete units, then observing that output is 100 may not be a terribly reliable indication that the partner also received an output of 100, but may be a good indication that output of the partner is between 90 and 110.

A simple example with 4 levels of output shows how combining several signals results in a better degree of public information as measured by the ε in Definition 1.

	1	2	3	4
1	3/48	1/48	3/48	5/48
2	1/48	3/48	5/48	3/48
3	3/48	5/48	3/48	1/48
4	5/48	3/48	1/48	3/48

Table 1

Suppose that there are four levels of output 1-4, and their probabilities at some particular action profile are given by Table 1. Then if players “report their signals,” the probability of Y^* is only 5/12. If players instead use a two point message space to report ranges of output 1 to 2 or 3 to 4, then, as shown in Table 2, the probability of the diagonal becomes 2/3.

	1-2	3-4
1-2	8/48	16/48
3-4	16/48	8/48

Table 2

Notice that aggregating signals has two effects: first, it increases the degree to which each player can forecast the other player’s message, which reduces the role of private information. Second, it reduces the informativeness of the messages, making it less likely that the assumption of pairwise full-rank is satisfied. However, when each player’s signal space Z_i is at least as large as $(\# A_1 + \# A_2) - 1$ (the minimum size consistent with pairwise full rank) it is possible to aggregate the signals while still allowing the messages to carry a substantial amount of information on the actions that were played. This observation has some importance when we notice that the proof of the Theorem remains valid even if we allow the space of private signals Z to vary, provided that the set Y^* remains fixed. For example in the partners case, we can allow the set of output levels to be measured on a finer scale as we approach the limit, while continuing to aggregate reports into a limited number of categories. Notice that when we refine the grid, if private information is received near the boundary of the category, the probability

the other player receives a signal in the same category may be close to $\frac{1}{2}$. However, this is consistent with part (2) of our Definition 1.

References

- Aoyagi, M. [2002] "Collusion in Dynamic Bertrand Oligopoly with Correlated Private Information," *Journal of Economic Theory* 102, 229-248.
- Ben-Porath, E. and M. Kahneman, "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70, 229-248.
- Bhaskar, V. and I. Obarra [2002] "Belief-Based Equilibria in the Repeated Prisoners' Dilemma with Private Monitoring," *Journal of Economic Theory*, 102, 40-69.
- Compte, O. [1998] "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica* 66, 597-626.
- Ely, J. and J. Valimaki [2002] "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102, 84-105.
- Forges, F. [1986] "An Approach to Communication Equilibrium," *Econometrica* 54, 1375-1385.
- Fudenberg, D., D.K. Levine, and E. Maskin [1994] "The Folk Theorem with Imperfect Public Information," *Econometrica* 62, 997-1040.
- Green, E. J. and R. H. Porter [1984]: "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, 52: 87-100.
- Kandori, M. [2002] "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory* 102, 1-15.
- Kandori, M. and H. Matsushima [1998] "Private Observation, Communication, and Collusion," *Econometrica* 66, 627-652.
- Mailath, G. and S. Morris [2002] "Repeated Games with Almost-Public Monitoring," *Journal of Economic Theory* 102, 189-228.
- Matsushima, H. [1991] "On the Theory of Repeated Games with Private Information, Part II: Revelation Through Communication," *Economics Letters* 35, 257-261.
- Myerson, R. [1986] "Multistage Games with Communication," *Econometrica*, 54, 323-358.
- Piccione, M. [2002] "The Repeated Prisoner's Dilemma with Imperfect Public Monitoring," *Journal of Economic Theory* 102, 70-83.
- Sekiguchi, T. [1997] "Efficiency in Repeated Prisoner's Dilemma with Private Monitoring," *Journal of Economic Theory* 76, 345-361.