

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 17, 2004*

¹ We are grateful to seminar participants at the Cowles Conference on repeated game with private information whose research inspired this project, to Ichiro Obara for helpful discussion, and to the editor and associate editor for useful comments on an earlier draft. The associate editor and Satoru Takahashi helped us find an error in the main proof; Takahashi and Dam Hojman helped proof-read various versions. We are happy to acknowledge the support of 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, which makes it difficult to sustain an equilibrium that is based on common beliefs. One way to restore common belief about the aspects of past play that matter for forecasts of future play is for players to send “cheap-talk” messages to one another; if these 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 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 Obara [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 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 on the announcements. However, 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. 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 the players’ signals are highly correlated.

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 $g_i(a)$. We define $\bar{g} = \max_{i,a} |g_i(a)|$. 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

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.

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. At the start of each period, a public randomization $w \in [0, 1]$ is drawn from a uniform distribution. The *public history* at time t , $h(t)$, consists of the announcements and realization of w signals in all previous periods, and also the realization of w in period t , so

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

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

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

A *strategy* for player i is a sequence of maps $\sigma_i(t)$ mapping the public and private histories $h(t), h_i(t)$ to probability distributions over S_i . A *partial strategy* is the strategy conditional on the initial realization of the public randomization device. A *public strategy* is a strategy that depends only on $h(t)$. We denote the null private history for player i by $h_i(1)$; the initial public history is $h(1)$. Observe that for each public history $h(t)$ the public profile σ induces a partial profile over the repeated game beginning at t . We denote this by $[\sigma | h(t)]$. Given a private partial profile σ , we can use the probability distribution generated by the partial profile and the null initial history to 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 $h(t)$ and any private partial strategy $\tilde{\sigma}_i$ by any player i we have

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

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

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, ν) public information with respect to m if for all action profiles a ,

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

$$(2) \text{ if } \pi_a(z) > 0 \text{ then for all } y^* \neq m_i(z_i), \pi_a^m(y^* | z_i) \leq \pi_a^m(m_i(z_i) | z_i) - \nu$$

When this condition is satisfied for some positive ν and “small” ε we say that the game has “almost public messaging.” This condition says that most of the time, each player is fairly confident of the other player's message; in the limit case of $(0, \nu)$ -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, they 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 $\lim_{\varepsilon \rightarrow 0} \pi_a^m(m_i(z_i) | z_i) = 1$. 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.

Notice that, except in the trivial case of perfect information, condition (2) rules out the possibility that the signals z_1, z_2 are independent conditional on a . In essence, condition (2) requires that if one player receives a signal that was unlikely conditional on a , it is likely that the other player receive the corresponding unlikely signal – that is, that the errors are correlated.⁶

For a given a, m we can consider $\pi_a^m(\bullet | Y^*)$ as a row vector. For player i we can 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 $\pi_{(\alpha, \alpha')} = \pi_{(\alpha', \alpha)}$, but it is satisfied for a set of probability measures π_a of full Lebesgue measure.

⁶ 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; subsequent work by Compte [1998] and Kandori and Matsushima [1998] show that it is possible to work with independent signals by considering strategies that only report truthfully periodically.

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 and a $v > 0$ such that for each n , game n has (ε^n, ν) public information with respect to m , that $\varepsilon^n \rightarrow 0$, that these games have common diagonal probabilities $\pi_a^{mn}(\cdot | Y^*) = \pi_a^m(\cdot | Y^*)$, 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 $v > 0$ in the interior of the feasible set 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 $\|v^n - v\| < \gamma^n$.*

Proof: To begin, we suppose that players use the given message profile m , and consider public equilibria of the game whose public signals are the announcements. When players use m and their signals are highly correlated, they are likely (but not certain) to make the same announcements. Let Γ be an auxiliary game of public information without public randomizations, with the same payoff functions as in the original game, and whose distribution over public signals is the conditional distribution on Y^* generated by the common diagonal probabilities.⁷

Using the arguments from Fudenberg, Levine, and Maskin [1994] and McLean, Obara and Postlewaite [2003], one can show that for any $\hat{v} > 0$ strictly interior to the feasible set there is an η such that for any interior $v \geq \hat{v}$ there exists $\hat{\delta} < 1$, a $K > 0$ and a set of payoffs V with $v \in V$, and $v_i' \geq \hat{v}_i$ for all i and 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:⁸

- a) the continuation payoffs $w(y)$ lie in V for all $y \in Y^*$,
- b) $|w(y) - w(y')| \leq K(1 - \bar{\delta})$

⁷ That is, in the auxiliary game we ignore the possibility that the players might make different announcements.

⁸ This follows from the facts that any compact convex subset of the interior of the positive feasible region is self-generating for all sufficiently large δ and that the self-generation can be constructed to satisfy conditions a) through e).

- 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 $h(t)$, the continuation payoffs $w(y) = w$ are constant.
- e) whenever the current profile is other than a static Nash equilibrium, any deviation from the current period equilibrium action costs at least $(1 - \bar{\delta})\eta$ in current period average present value.

For any $\bar{\delta}$ and an equilibrium satisfying the conditions above; we refer to the strategies in this equilibrium as the base strategy. Our method of proof will be to use base strategies to construct equilibrium strategies in the game of interest.

Fix a δ . Recall from the definition of public information that ν is the minimum probability (for any n) of disagreement given that one player lies. Define

$$q(\delta) = (1 - \delta) \frac{K}{\nu \hat{v}},$$

which is a probability for δ sufficiently large, and recalling that $\bar{\pi}_a^{mn}$ is the probability of agreement. We now define

$$\begin{aligned} Q^n(\delta) &= \min_a \{ \bar{\pi}_a^{mn} + (1 - \bar{\pi}_a^{mn})(1 - q(\delta)) \} \\ &= 1 - q(\delta) \max_a \{ 1 - \bar{\pi}_a^{mn} \}, \end{aligned}$$

which is also a probability if $q(\delta)$ is.

We now consider the public information game with discount factor

$$\bar{\delta} = Q^n(\delta)\delta.$$

We construct a map χ from public histories $h^n(t)$ in the game of interest to the set of histories of the same length in the public-information game Γ and the symbol P (for punishment). This map induces strategies in the original game by assigning the profile from the base game strategy when the history is mapped to a history in Γ , and by assigning the static Nash equilibrium when the history is mapped to P . We will show that these strategies form an equilibrium in the original game.

Formally, 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 $h^n(t) = (h^n(t - 1), y_{t-1}^n, w_t^n)$.

When the history is understood, we will abuse notation and write a_t to be the action taken in the base strategy given the history $\chi(h^n(t-1))$. It is convenient also to define

$$p^n(a, \delta) = \frac{Q^n(\delta) - (1 - \bar{\pi}_a^{mn})(1 - q(\delta))}{\bar{\pi}_a^{mn}}.$$

Notice that we have chosen $q(\delta), Q^n(\delta)$ so that $p^n(a, \delta) \leq 1$. Notice also that $Q^n(\delta) \geq \min_a \bar{\pi}_a^n$, so if $\min_a \bar{\pi}_a^{mn} > 1/2$, we also have $p^n(a, \delta) \geq 0$.

Case 0: If the current base strategy profile is a static Nash equilibrium use the public randomization w_t^n choosing each $y^* \in Y^*$ with probability $Q^n(\delta)\pi_{a_t}^m(y^* | Y^*)$ and setting $\chi(h^n(t-1), y_{t-1}^n, w_t^n) = (\chi(h^n(t-1)), y^*)$ and with the remaining probability $\chi(h^n(t-1), y_{t-1}^n, w_t^n) = P$.

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 $y_{t-1,1}^n = y_{t-1,2}^n$, use the public randomization w_t^n so that with probability $p^n(a_t, \delta)$ we have

$$\chi(h^n(t-1), y_{t-1}^n, w_t^n) = (\chi(h^n(t-1)), y_{t-1}^n)$$

and $\chi(h^n(t-1), y_{t-1}^n, w_t^n) = P$ otherwise.

Case 3: If $\chi(h^n(t-1)) \neq P$ and $y_{t1}^n \neq y_{t2}^n$, use the public randomization w_t^n choosing each $y^* \in Y^*$ with probability $[1 - q(\delta)]\pi_{a_t}^m(y^* | Y^*)$ and setting $\chi(h^n(t-1), y_{t-1}^n, w_t^n) = (\chi(h^n(t-1)), y^*)$ and with the remaining probability $\chi(h^n(t-1), y_{t-1}^n, w_t^n) = P$.

The remainder of the proof, which is in the appendix, verifies that these strategies form an equilibrium in game n with average payoffs v^n , and that there is a sequence $\gamma^n \rightarrow 0$ such that $\|v^n - v\| < \gamma^n$.

Remark: The proof can easily be adjusted so that the equilibrium, except when the static Nash equilibrium is played, is strict of order $2(1 - \delta)\bar{g}$, where $\bar{g} = \max_{i,a} |g_i(a)|$. Moreover if we perturb the game so that $\|g - g'\| \leq \varepsilon$ and $\|\pi - \pi'\| \leq \varepsilon$, then for any

fixed δ the average discounted value of particular strategies changes by at most $2\varepsilon\bar{g}/(1-\delta)$. Hence if $\varepsilon \leq (1-\delta^*)^2$, and the perturbed game still has a static equilibrium with payoffs 0 the theorem continues to hold for the perturbed game and discount factor δ^* . Tracing out the equilibrium payoffs corresponding to the various histories, it follows that the convex hull of this set is self-generating for δ^* , and hence by the results of Fudenberg, Levine and Maskin [1994] these are payoff of perfect public equilibria for all $\delta' \geq \delta^*$. This establishes 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, ν) public information with respect to m , that $\varepsilon^n \rightarrow 0$, that $\pi_a^m(\cdot | Y^*)$, has pure-strategy pairwise full rank with respect to m , and that each g^n has a static equilibrium with payoffs converging to 0. Then there is a sequence $\gamma^n \rightarrow 0$ such that for any feasible interior 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 $\|v^n - v\| < \gamma^n$.*

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. Discussion

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. 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.

Turning to applications of the main theorem, one feature is that it allows the possibility that players aggregate information by making the same announcement for several different private signals. 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. The working paper version of this paper gives an illustrative example.

Appendix: Completion of the proof of the Theorem

We must verify that the strategies constructed in the text form an equilibrium in game n with average payoffs v^n , and that there is a sequence $\gamma^n \rightarrow 0$ such that $\|v^n - v\| < \gamma^n$.

Let $v_i(h(t))$ be the average present values from following the base strategies in Γ when the discount factor is $\bar{\delta} = Q^n(\delta)\delta$. It is convenient to do calculation in terms of total present value, which we denote by $\tilde{\cdot}$.

Part (0): The Static Nash Equilibrium

Notice that whenever these strategies call upon players to mix, future outcomes do not depend on current play, so incentive compatibility holds. The total present value payoff is

$$\begin{aligned} & \tilde{v}_i^n(h(t-1)) \\ &= g_i(\alpha_t) + \bar{\delta} \left(\sum_{y^*} \pi_{a_i}^{m,n}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) = \\ & \bar{\delta} \left(\sum_{y^*} \pi_{a_i}^{m,n}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \end{aligned}$$

since we have normalized the payoff of the static equilibrium to be 0.

The remainder of the proof considers payoffs and incentive compatibility when the strategies call for a pure profile a_t .

Recall that the probability that the two players make the same announcement in game n is $\bar{\pi}_a^{mn}$. Thus the total present values from following the corresponding strategies in the game n when the discount factor is δ and the history $h^n(t-1)$ has not yet reached the state P

$$\begin{aligned}
& \tilde{v}_i^n(h^n(t-1)) \\
&= g_i(a_t) \\
&+ \delta \left(\sum_{y^*} \pi_{a_i}^{mn}(y^*) p^n(a_t, \delta) \tilde{v}_i^n(h(t-1), y^*) \right) \\
&+ \delta (1 - \bar{\pi}_{a_i}^{mn}) \left(\sum_{y^*} [1 - q(\delta)] \pi_{a_i}^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
&= g_i(a_t) \\
&+ \delta p^n(a_t, \delta) \bar{\pi}_{a_i}^{mn} \left(\sum_{y^*} \pi_{a_i}^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
&+ \delta (1 - \bar{\pi}_{a_i}^{mn}) [1 - q(\delta)] \left(\sum_{y^*} \pi_{a_i}^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
&= g(a_t) + \\
&+ \delta Q^n(\delta) \left(\sum_{y^*} \pi_{a_i}^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
&= g_i(a_t) + \bar{\delta} \left(\sum_{y^*} \pi_{a_i}^{m,n}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right)
\end{aligned}$$

where we have made use of the fact that the continuation payoff in the punishment state is 0. Since this equation holds for every starting point, including the case in which the current profile is static Nash, this implies the total present value payoff to using the proposed strategies in the game n with discount factor δ is the same as in the base game with discount factor $\bar{\delta}$. Thus the average present value $v_i^n(h^n(t-1))$ in the n -game is related to the average present value in the base game according to

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

We now calculate the factor on the RHS:

$$\begin{aligned}
\frac{1 - \bar{\delta}}{1 - \delta} &= \frac{1 - Q^n(\delta)\delta}{1 - \delta} = \\
&= \frac{1 - \delta - \delta q(\delta) \max_a \{1 - \bar{\pi}_a^{mn}\}}{1 - \delta} \\
&= 1 - \delta \max_a \{1 - \bar{\pi}_a^{mn}\} \frac{K}{\nu \hat{v}}
\end{aligned}$$

Recall that the theorem asserts the existence of equilibria in the n -games with average payoffs v^n and a sequence $\gamma^n \rightarrow 0$ such that $\|v^n - v\| < \gamma^n$. If we take

$$\gamma^n = \max_a \{1 - \bar{\pi}_a^{mn}\} \frac{K}{\nu \hat{v}} \bar{g}$$

then since $\max_a \{1 - \bar{\pi}_a^{mn}\} \rightarrow 0$, so does γ^n as desired.

It remains to show that these strategies are indeed an equilibrium for the n -games. To prove this, we will show that (1) no player has an incentive to choose a non-equilibrium action when all players tell the truth, and that (2) it is optimal to use m regardless of the action chosen.

Part (1): No player has an incentive to choose a non-equilibrium action given truth-telling.

Let $a = (a^i, a_t^{-i})$ represent a deviation for player i , that is $a^i \neq a_t^i$. We compute the payoff to changing actions, while telling the truth

$$\begin{aligned}
& g_i(a) \\
& + \delta \left(\sum_{y^*} \pi_a^{mn}(y^*) p^n(a_t, \delta) \tilde{v}_i^n(h(t-1), y^*) \right) \\
& + \delta (1 - \bar{\pi}_a^{mn}) \left(\sum_{y^*} [1 - q(\delta)] \pi_{a_t}^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
& = g_i(a) + \\
& + \delta p^n(a_t, \delta) \bar{\pi}_a^{mn} \left(\sum_{y^*} \pi_a^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
& + \delta (1 - \bar{\pi}_a^{mn}) [1 - q(\delta)] \left(\sum_{y^*} \pi_{a_t}^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
& = g_i(a) + \\
& + \delta Q^n(\delta) \left(\sum_{y^*} \pi_a^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
& + \delta [p^n(a_t, \delta) + q(\delta) - 1] [\bar{\pi}_a^{mn} - \bar{\pi}_{a_t}^{mn}] \left(\sum_{y^*} \pi_a^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \\
& + \delta p^n(a_t, \delta) [1 - \bar{\pi}_a^{mn}] [1 - q(\delta)] [\bar{\pi}_a^{mn} - \bar{\pi}_{a_t}^{mn}] \left(\sum_{y^*} [\pi_{a_t}^{mn}(y^* | Y^*) - \pi_a^{mn}(y^* | Y^*)] \tilde{v}_i^n(h(t-1), y^*) \right)
\end{aligned}$$

Because the base equilibrium is η -strict, the first two terms lose at least η (in total present value) compared to playing the equilibrium action. Hence the gain to deviating is at most

$$\begin{aligned}
& -\eta \\
& +\delta[p^n(a_t, \delta) + q(\delta) - 1][\bar{\pi}_a^{mn} - \bar{\pi}_{a_t}^{mn}]\left(\sum_{y^*} \pi_a^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*)\right) \\
& +\delta p^n(a_t, \delta)[\bar{\pi}_a^{mn} - \bar{\pi}_{a_t}^{mn}]\left(\sum_{y^*} [\pi_{a_t}^{mn}(y^* | Y^*) - \pi_a^{mn}(y^* | Y^*)] \tilde{v}_i^n(h(t-1), y^*)\right)
\end{aligned}$$

First we compute

$$\begin{aligned}
& p^n(a_t, \delta) + q(\delta) - 1 = \\
& \frac{Q^n(\delta) - (1 - \bar{\pi}_a^{mn})(1 - q(\delta))}{\bar{\pi}_a^{mn}} + q(\delta) - 1 = \\
& \frac{Q^n(\delta)}{\bar{\pi}_a^{mn}} - (1 - q(\delta)) \left[\frac{1 - \bar{\pi}_a^{mn}}{\bar{\pi}_a^{mn}} + 1 \right] = \\
& \frac{Q^n(\delta) - (1 - q(\delta))}{\bar{\pi}_a^{mn}} = \\
& \frac{\min_{a'} \{ \bar{\pi}_a^{mn} + (1 - \bar{\pi}_{a'}^{mn})(1 - q(\delta)) \} - (1 - q(\delta))}{\bar{\pi}_a^{mn}} = \\
& \frac{\min_{a'} \{ \bar{\pi}_a^{mn} + (1 - \bar{\pi}_{a'}^{mn})(1 - q(\delta)) - (1 - q(\delta)) \}}{\bar{\pi}_a^{mn}} = \\
& \frac{\min_{a'} \{ \bar{\pi}_a^{mn} - \bar{\pi}_{a'}^{mn}(1 - q(\delta)) \}}{\bar{\pi}_a^{mn}} = \\
& \frac{q(\delta) \min_{a'} \{ \bar{\pi}_{a'}^{mn} \}}{\bar{\pi}_a^{mn}} = \\
& (1 - \delta) \frac{K}{\nu \hat{v}} \frac{\min_{a'} \{ \bar{\pi}_{a'}^{mn} \}}{\bar{\pi}_a^{mn}}
\end{aligned}$$

Then we observe

$$\begin{aligned}
& \left| \sum_{y^*} [\pi_{a_t}^{mn}(y^* | Y^*) - \pi_a^{mn}(y^* | Y^*)] \tilde{v}_i^n(h(t-1), y^*) \right| \leq K \\
& \left(\sum_{y^*} \pi_a^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right) \leq \bar{U} / (1 - \delta)
\end{aligned}$$

Hence the gain to deviating is bounded by

$$-\eta + \left| \max_a \bar{\pi}_a^{mn} - \min_a \bar{\pi}_a^{mn} \right| \left[\frac{K}{\nu \hat{v}} \bar{U} + K \right].$$

In other words there is no incentive to deviate, provided that n is large enough that

$$\left| \max_a \bar{\pi}_a^{mn} - \min_a \bar{\pi}_a^{mn} \right| \leq \eta \frac{\nu \hat{v}}{K \bar{U} + K \nu \hat{v}}.$$

When this is true, we take γ^n as specified above, otherwise, take $\gamma^n = \bar{U}$.

Part (2): It is optimal to use m regardless of the action chosen.

Suppose the message z_i is received and that the action a^i is chosen. Let $a = (a^i, a_t^{-i})$. Also let $\pi_y^n = \pi_a^{mn}(y | z_i)$, and $y^* = m_i(z_i)$. Finally, set

$$\bar{v}_i = \left(\sum_{y^*} \pi_{a_i}^{mn}(y^* | Y^*) \tilde{v}_i^n(h(t-1), y^*) \right).$$

Playing a^i and truth-telling results in total present value utility of

$$g_i(a) + \delta \pi_{y^*}^n \tilde{v}_i(h(t-1), y^*) + \delta(1 - \pi_{y^*}^n)(1 - q(\delta))\bar{v}_i.$$

Playing a^i and telling the lie y results in

$$g_i(a) + \delta \pi_y^n \tilde{v}_i(h(t-1), y) + \delta(1 - \pi_y^n)(1 - q(\delta))\bar{v}_i,$$

so the gain to lying is

$$\begin{aligned} & \delta[\pi_y^n \tilde{v}_i(h(t-1), y) - \pi_{y^*}^n \tilde{v}_i(h(t-1), y^*)] + \delta[\pi_{y^*}^n - \pi_y^n](1 - q(\delta))\bar{v}_i \\ & = \delta[\pi_y^n [\tilde{v}_i(h(t-1), y) - \bar{v}_i] - \pi_{y^*}^n [\tilde{v}_i(h(t-1), y^*) - \bar{v}_i]] - \delta[\pi_{y^*}^n - \pi_y^n]q(\delta)\bar{v}_i \\ & \leq \delta K - \delta \nu q(\delta) \hat{v} / (1 - \delta) \end{aligned}$$

This is less than or equal zero by the definition of $q(\delta)$.

☑

References

- Aoyagi, M. [2002] “Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication,” *Journal of Economic Theory* 102, 229-248.
- Ben-Porath, E. and M. Kahneman, [1996] “Communication in Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 70, 281-297.
- Bhaskar, V. and I. Obara [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. Välimäki [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.
- McLean, R., I. Obara and A. Postlewaite [2003] "Informational Smallness and Private Monitoring in Repeated Games," mimeo UCLA.
- Myerson, R. [1986] "Multistage Games with Communication," *Econometrica*, 54, 323-358.
- Piccione, M. [2002] "The Repeated Prisoner's Dilemma with Imperfect Private 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.