PERFECT PUBLIC EQUILIBRIUM WHEN PLAYERS ARE PATIENT

DREW FUDENBERG, DAVID K. LEVINE AND SATORU TAKAHASHI

ABSTRACT. The limit set of perfect public equilibrium payoffs of a repeated game as the discount factor goes to one is characterized, with examples, even when the full-dimensionality condition fails.

Keywords: Repeated Games, Perfect Public Equilibrium, Folk Theorem.

JEL Classification: C72; C73

Date: First Version: 3rd November 2003, This Version: 12th September 2005.

This material is based upon work supported by the National Science Foundation under Grants No. SES-0314713 and SES-0426199.

Corresponding Author: David K. Levine, Department of Economics, UCLA, Los Angeles, CA 90095, USA. Phone/Fax: 310-825-3810. Email: david@dklevine.com.

1. INTRODUCTION

Fudenberg and Levine [1994] (FL) showed that the limit of the set of perfect public equilibrium payoffs of a repeated game as the discount factor goes to one can be characterized by the solution of a family of static linear programming problems. This result has been applied and extended by a number of subsequent authors, including Kandori and Matsushima [1998], Dellarocas [2003], and Ely et al [2003].

The FL result requires that the set of payoff vectors obtained by the algorithm should have "full dimension," that is, the dimension is equal to the number of long-run players in the game. This paper extends the linear programming characterization to cases where this "full-dimensionality" condition fails, either because of the payoff structure of the stage game, or because of a restriction to equilibrium strategies whose continuation payoffs are on a lower-dimensional set. We apply our result to three such restrictions from the literature. The first application is to repeated games with all long-run players and observed actions, where the feasible payoffs in the stage game lie in a lower-dimensional set. The linear programming characterization allows us to generalize the results of Abreu et al [1994], who assumed a condition called *NEU condition* and of Wen [1994], who assumed that mixed strategies are observed. The second application is to the strongly symmetric equilibrium studied by Abreu [1986] and Abreu et al [1986], which restricts the continuation payoffs to the one-dimensional set where all players' payoffs are identical. The third application is to the restriction that all payoffs lie on a line segment of the Pareto frontier, which we use to derive a sufficient condition for the exact achievability of first-best outcomes. Equilibria of this type, for which all continuation payoffs lie on the Pareto frontier, have a strong renegotiation-proofness property: regardless of the history, players can never unanimously prefer another equilibrium.

2. MODEL

We consider a repeated game with imperfect public monitoring played by long-run and short-run players. We follow FL in the notation. In the stage game, each player $i = 1,\ldots,n$ simultaneously chooses a pure action a_i from a finite set A_i. $a \in A \equiv \prod_{i=1}^{n} A_i$ induces a publicly observed outcome $y \in Y$ with probability $\pi_y(a)$. Player i's payoff to an action profile a is $g_i(a)$. For each mixed action profile $\alpha \in \mathcal{A} \equiv \prod_{i=1}^{n} \mathcal{A}_i$, we can induce $\pi_y(\alpha)$ and $g_i(\alpha)$.

For $i \in LR \equiv \{1,\ldots,L\}$, $L \leqslant n$, i is a long-run player whose objective is to maximize the average discounted value of per-period payoffs $\{g_i(t)\}$,

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

The remaining players $j \in SR \equiv \{L+1, \ldots, n\}$ represent short-run players, each of whom plays only once. Let

$$B\colon \mathcal{A}_1 \times \cdots \times \mathcal{A}_L \to \mathcal{A}_{L+1} \times \cdots \times \mathcal{A}_n$$

be the correspondence that maps any mixed action profile $\alpha_{LR} = (\alpha_1, \ldots, \alpha_L)$ for the long-run players to the corresponding static equilibria $\alpha_{SR} = (\alpha_{L+1}, \ldots, \alpha_n)$ for the short-run players. That is, for each $\alpha \in \mathrm{graph}(B) \equiv \{(\alpha_{LR}, \alpha_{SR}) \in \mathcal{A} \mid \alpha_{SR} \in B(\alpha_{LR})\}$ and each $j = L+1, \ldots, n$, α_j maximizes $g_j(\alpha'_j, \alpha_{-j})$.

Let $\mathcal{A}^0$ be a subset of graph(B). We focus on $\mathcal{A}^0$-*perfect public equilibria*: strategy profiles in which all players choose action profiles from $\mathcal{A}^0$, depending only on the public history, and in which following every public history the remaining public strategy profile forms a Nash equilibrium. Note that an action profile specified by an equilibrium belongs to $\mathcal{A}^0$ even after an off-path history, but that each player's deviations from the equilibrium need not be in $\mathcal{A}^0$. $E(\mathcal{A}^0, \delta)$ is the set of average present values for the long-run players in $\mathcal{A}^0$-perfect public equilibria. We will characterize the limit of $E(\mathcal{A}^0, \delta)$ without the "full-dimensionality" condition.

3. ALGORITHM

We fix $\mathcal{A}^0$ throughout this section. We define the sequence $X^0, Q^0, X^1, Q^1, X^2, Q^2, \ldots$ where X^m are affine subspaces of $\mathbb{R}^L$ and Q^m are compact convex subsets of X^m by the following procedure. Let $X^0 = \mathbb{R}^L$. Let $g_{LR}(\alpha)$ denote the vector of payoffs for long-run players only. For given X^m, we consider a linear programming problem for given $\alpha \in \mathcal{A}^0$ with $g_{LR}(\alpha) \in X^m$, $\lambda \in \mathbb{R}^L \setminus \{\mathbf{0}\}$ parallel to X^m, and $\delta \in (0,1)$:

$$k^m(\alpha, \lambda, \delta) = \max_{v,w} \lambda \cdot v \qquad \text{subject to}$$

$$\text{(a)} \quad v_i = (1-\delta) g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) w_i(y)$$
$$\text{for } i \in LR \text{ and } a_i \in A_i \text{ s.t. } \alpha_i(a_i) > 0,$$

$$\text{(b)} \quad v_i \geqslant (1-\delta) g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) w_i(y)$$
$$\text{for } i \in LR \text{ and } a_i \in A_i \text{ s.t. } \alpha_i(a_i) = 0,$$

$$\text{(c)} \quad \lambda \cdot v \geqslant \lambda \cdot w(y) \qquad \text{for } y \in Y,$$

$$\text{(d)} \quad w(y) \in X^m \qquad \text{for } y \in Y.$$

If there is no (v, w) that satisfies constraints (a)-(d), then we set $k^m(\alpha, \lambda, \delta) = -\infty$. Note that $k^0(\alpha, \lambda, \delta)$ corresponds to $k^*(\alpha, \lambda, \delta)$ in FL. Similarly to Lemma 3.1 (i) in FL, $k^m(\alpha, \lambda, \delta)$ is independent of δ, and thus denoted by

$k^m(\alpha,\lambda)$.[1] Set

$$\begin{aligned} k^m(\lambda) &= \sup_{\alpha\in\mathcal{A}^0,\, g_{LR}(\alpha)\in X^m} k^m(\alpha,\lambda), \\ H^m(\lambda) &= H(\lambda, k^m(\lambda)), \\ Q^m &= \bigcap_{\lambda\in\mathbb{R}^L\setminus\{\mathbf{0}\}:\ \text{parallel to } X^m} H^m(\lambda)\cap X^m, \end{aligned}$$

where $H(\lambda,k)=\{v\in\mathbb{R}^L \mid \lambda\cdot v\leqslant k\}$. If $Q^m=\emptyset$ or Q^m is a singleton whose element does not correspond to a static equilibrium in $\mathcal{A}^0$, we stop the algorithm and define $Q^*(\mathcal{A}^0)=\emptyset$.[2] If Q^m is a singleton consisting of a static equilibrium payoff profile in $\mathcal{A}^0$ or we have $\dim Q^m=\dim X^m$, we stop the algorithm and define $Q^*(\mathcal{A}^0)=Q^m$. Otherwise, let X^{m+1} be the affine hull of Q^m, which is the smallest affine space including Q^m, and we again solve a linear programming problem after X^m is replaced by X^{m+1}.

Note that every time the algorithm continues, the dimension of X^m decreases by at least 1, so the algorithm stops in a finite number of steps.

As is standard in this literature, payoff profile v is the target that will be supported by some equilibrium, and the function w gives the continuation payoffs $w(y)$ starting tomorrow if the current outcome is y. Constraints (a) are the accounting identities that define the expected payoff profile v, and constraints (b) are the incentive constraints, requiring that playing α maximizes expected payoff provided that continuation payoffs are given by w. Constraints (c) require that all of the continuation payoffs are included in the half-space defined by v and λ; loosely speaking, the continuation payoffs are not allowed to be "better" (in the λ direction) than v is.

Each step of this algorithm differs from FL's only in constraints (d) and $\mathcal{A}^0$. Constraints (d) require that all of the continuation payoffs are included by the affine hull of Q^{m-1} in the previous step, and $\mathcal{A}^0$ is a restriction on equilibrium action profiles. In the case of $\mathcal{A}^0=\text{graph}(B)$, the first step of the algorithm is exactly the same as FL's linear programming problem. Actually, Q^0 is equal to what FL call Q. If we assume the full dimensionality of Q, that is, $\dim Q=L$, then the algorithm stops at the first step, and we have $Q^*(\text{graph}(B))=Q$.

By this algorithm, we obtain the limit of $\mathcal{A}^0$-perfect public equilibrium payoffs, which is a generalization of Theorem 3.1 in FL.

[1] We should point out that the condition given in FL Lemma 3.1 (iii) is sufficient but not necessary for $k^*(\alpha,\lambda)=\lambda\cdot g_{LR}(\alpha)$; FL incorrectly assert that the condition is necessary as well. The condition is only necessary under the additional assumption that all outcomes have positive probability under α.

[2] $Q^m=\emptyset$ is possible only if $\mathcal{A}^0$ contains no static equilibrium.

Let $E^*(\mathcal{A}^0,\delta)$ be the set of $\mathcal{A}^0$-perfect public equilibrium payoff profiles when public randomization devices are available at the beginning of each period. $E^*(\mathcal{A}^0,\delta)$ is a bounded convex set that contains $E(\mathcal{A}^0,\delta)$.

Theorem. *$E^*(\mathcal{A}^0,\delta) \subseteq Q^*(\mathcal{A}^0)$ for any δ. If $Q^*(\mathcal{A}^0) \neq \emptyset$, then for any compact set K in the relative interior of $Q^*(\mathcal{A}^0)$, there exists $\bar{\delta} < 1$ such that $K \subseteq E(\mathcal{A}^0,\delta)$ for any $\delta > \bar{\delta}$.*[3] *Hence $\lim_{\delta\to 1} E(\mathcal{A}^0,\delta) = \lim_{\delta\to 1} E^*(\mathcal{A}^0,\delta) = Q^*(\mathcal{A}^0)$.*

Proof. We show that $E^*(\mathcal{A}^0,\delta) \subseteq Q^m$ for any m by induction. Suppose that $E^*(\mathcal{A}^0,\delta) \subseteq X^m$, and we show that $E^*(\mathcal{A}^0,\delta) \subseteq Q^m \subseteq X^{m+1}$. If not, since $E^*(\mathcal{A}^0,\delta)$ is bounded, we may find a positive number $\varepsilon > 0$, a point $v \in E^*(\mathcal{A}^0,\delta)$, and a unit vector $\lambda \in \mathbb{R}^L \setminus \{\mathbf{0}\}$ parallel to X^m such that $\lambda \cdot v = k > k^m(\lambda) + [(1-\delta)/\delta]\varepsilon$ and $E^*(\mathcal{A}^0,\delta) \subseteq H(\lambda, k+\varepsilon)$. Then v is written as

$$v = \int_0^1 v(\omega)\,\mathrm{d}\omega,$$

where, for almost every realization $\omega \in [0,1]$ of a public randomization device, $v(\omega)$ is enforceable with some current action profile $\alpha(\omega) \in \mathcal{A}^0$ and continuation payoff profiles $w(y,\omega)$ in $E^*(\mathcal{A}^0,\delta) \subseteq H(\lambda,k+\varepsilon) \cap X^m$. Since $v(\omega) \in E^*(\mathcal{A}^0,\delta) \subseteq X^m$, we have $g_{LR}(\alpha(\omega)) \in X^m$. Pick $\omega \in [0,1]$ such that $\lambda \cdot v(\omega) \geqslant k$. For this ω, by shifting payoff profiles independently of y, we can enforce the shifted target payoff profile $v'(\omega) = v - [(1-\delta)/\delta]\varepsilon\lambda$ by $\alpha(\omega)$ and $w'(y,\omega) = w(y,\omega) - (1/\delta)\varepsilon\lambda \in H(\lambda,k') \cap X^m$, where $k' = k - [(1-\delta)/\delta]\varepsilon > k^m(\lambda)$. So the score $\lambda \cdot v'(\omega) \geqslant k'$ is attained with continuation payoff profiles in $H(\lambda,k') \cap X^m$, which contradicts the definition of $k^m(\lambda)$.

If $Q^*(\mathcal{A}^0)$ is set to be the empty set because, at some step of the algorithm, Q^m is a singleton whose element does not correspond to a static equilibrium in $\mathcal{A}^0$, then, since there is no static equilibrium in $\mathcal{A}^0$ and continuation payoffs need to be constant, we have $E^*(\mathcal{A}^0,\delta) = \emptyset = Q^*(\mathcal{A}^0)$. Otherwise, we have $E^*(\mathcal{A}^0,\delta) \subseteq Q^m = Q^*(\mathcal{A}^0)$ for some m.

Now suppose that $Q^*(\mathcal{A}^0) \neq \emptyset$, and let K be a compact set in the relative interior of $Q^*(\mathcal{A}^0)$. We will show that $K \subseteq E(\mathcal{A}^0,\delta)$ for all sufficiently large δ. If $Q^*(\mathcal{A}^0)$ is a singleton, then $E(\mathcal{A}^0,\delta) = Q^*(\mathcal{A}^0)$ for any δ. Otherwise, let X^* be the affine hull of $Q^*(\mathcal{A}^0)$. Then the proof differs from FL's original one mainly in that we use the *relative* topology induced on X^* instead of the standard topology on $\mathbb{R}^L$.[4] Since K is a compact set in the relative interior of $Q^*(\mathcal{A}^0)$, there exists a smooth, convex, and compact set $W \supseteq K$ in the

[3] The relative interior of a subset S of $\mathbb{R}^L$ is the interior of S under the topology induced on the affine hull of S.

[4] Our proof also differs because it does not assume the existence of static equilibria in $\mathcal{A}^0$.

relative interior of $Q^*(\mathcal{A}^0)$. We show that $W \subseteq E(\mathcal{A}^0, \delta)$ for sufficiently large δ. Since W is compact and convex, it is enough to show that for each $v \in W$, there exist $\delta < 1$ and a relatively open neighborhood U of v with $U \subseteq P(\mathcal{A}^0, \delta, W)$, where $P(\mathcal{A}^0, \delta, W)$ is the set of payoff profiles generated by some $\alpha \in \mathcal{A}^0$ and W (Fudenberg et al [1994, Lemma 4.2]).

First, suppose that v is on the relative boundary of W. Let λ be parallel to X^* and normal to W at v. Let $k = \lambda \cdot v$, and let $H = H(\lambda, k)$ be the unique half-space in the direction of λ such that $H \cap X^*$ contains W and its relative boundary is tangent to W at v. Since W is in the relative interior of $Q^*(\mathcal{A}^0)$, there exists an action profile $\alpha \in \mathcal{A}^0$ with $g_{LR}(\alpha) \in X^*$ that generates a point $v' \in X^*$ with $\lambda \cdot v' > k$ using continuation payoffs in $H(\lambda, \lambda \cdot v') \cap X^*$. Then, for some $\delta' < 1$ and $\varepsilon > 0$, (α, v) can be enforced with respect to $H(\lambda, k - \varepsilon) \cap X^*$.

Second, suppose that v is in the relative interior of W. Pick any λ parallel to X^*. Let $k = \lambda \cdot v$ and $H = H(\lambda, k)$. Similarly to the above argument, there exists $\alpha \in \mathcal{A}^0$ such that, for some $\delta' < 1$ and $\varepsilon > 0$, (α, v) can be enforced with respect to $H(\lambda, k - \varepsilon) \cap X^*$.

For any $\delta'' \geqslant \delta'$, we may find $w(y, \delta'')$ that enforce (α, v) and $\bar{\kappa} > 0$ such that

$$w(y, \delta'') \in H\left(\lambda, k - \frac{\delta'(1-\delta'')}{\delta''(1-\delta')}\varepsilon\right) \cap X^*,$$

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

Consider the ball $U(\delta'')$ around v of radius $\bar{\kappa}(1-\delta'')$ in X^*. Since W is smooth in X^*, for δ'' sufficiently close to 1 there exists $\tilde{\kappa} > 0$ such that the difference between $H \cap X^*$ and $H \cap W$ in $U(\delta'')$ is at most $\tilde{\kappa}(1-\delta'')^2$. It follows that there exists $\delta < 1$ such that (α, v) can be enforced by continuation payoffs $w(y, \delta)$ in the relative interior of W. Since $w(y, \delta)$ are in the relative interior, they may be translated by a small constant independent of y generating incentive compatible payoffs in a relative neighborhood U of v. □

Remark. Our Theorem shows that allowing public randomizations does not change the limit set. For a fixed δ, however, $E^*(\mathcal{A}^0, \delta)$ may be larger than $E(\mathcal{A}^0, \delta)$.

Several other choices of how to determine the sets X^m lead to the same result $Q^*(\mathcal{A}^0)$. For example, at the beginning of the first step, we can choose X^0 to be any affine subspace of $\mathbb{R}^L$ that contains $g_{LR}(\alpha)$ for every $\alpha \in \mathcal{A}^0$. If $1 \leqslant \dim Q^m < \dim X^m$, then we can move to the next step with any affine subspace X^{m+1} of X^m that contains Q^m.

FIGURE 4.1. A three-player game in Fudenberg and Maskin [1986]

1,1,1	0,0,0
0,0,0	0,0,0

0,0,0	0,0,0
0,0,0	1,1,1

It is easy to extend this theorem to games with infinitely many pure actions.[5] However, allowing infinitely many signals would involve measure-theoretic complications that are beyond the scope of this paper.

4. APPLICATIONS

4.1. **Fudenberg and Maskin's Example.** To illustrate the algorithm, we apply it to the example Fudenberg and Maskin [1986] used to motivate the full dimensionality condition. We set $L = n = 3$, so that there are three long-run players and no short-run players, set $Y = A = \{0,1\}^3$, and set $\pi_y(a) = 1$ if and only if $y = a$, so that the signal perfectly reveals the action profile. Stage game payoffs are depicted in Figure 4.1.

Let $\mathcal{A}^0 = \mathcal{A}$ and $X^0 = \mathbb{R}^3$, and solve the first step of our algorithm. By a simple computation, we have $Q^0 = \{(x,x,x) \mid 0 \leqslant x \leqslant 1\}$. Since Q^0 has a lower dimension than X^0, we set $X^1 = \{(x,x,x) \mid x \in \mathbb{R}\}$ and move to the second step of our algorithm.

In the second step, we have two directions parallel to X^1 (up to positive constants), $\mathbf{1} = (1,1,1)$ and $-\mathbf{1} = (-1,-1,-1)$. We first consider the case of $\lambda = -\mathbf{1}$. Fix any α. As Fudenberg and Maskin show, for any α, there exist a player i and an action a_i such that $g_i(a_i, \alpha_{-i}) \geqslant 1/4$. Since (v,w) in the linear programming problem satisfies constraints (a) and (b), we have

$$v_i \geqslant (1-\delta) \times \frac{1}{4} + \delta \sum_y \pi_y(a_i, \alpha_{-i}) w_i(y).$$

Since $g_{LR}(\alpha) \in X^1$ and $w(y) \in X^1$ for any outcome y by constraints (d), it follows from constraints (a) that $v \in X^1$ as well. Then, since $-3v_i = (-\mathbf{1}) \cdot v \geqslant (-\mathbf{1}) \cdot w(y) = -3w_i(y)$ for any outcome y by constraints (c), we have

$$v_i \geqslant (1-\delta) \times \frac{1}{4} + \delta \sum_y \pi_y(a_i, \alpha_{-i}) w_i(y) \geqslant (1-\delta) \times \frac{1}{4} + \delta v_i,$$

[5] The proof carries over verbatim as long as stage-game payoff function g_i is bounded for every player i, "max" is replaced by "sup" in the definition of $k^m(\alpha, \lambda, \delta)$, and constraints (a) are required not only for every a_i with positive point mass but also for almost every a_i with respect to α_i.

and hence $v_i \geqslant 1/4$. Therefore, we have $k^1(\alpha, -\mathbf{1}) \leqslant -3/4$ for any α. Since the equality holds when each player mixes the two actions with equal probability, we have $k^1(-\mathbf{1}) = -3/4$ and $H^1(-\mathbf{1}) = H(-\mathbf{1}, -3/4)$. We also have $H^1(\mathbf{1}) = H(\mathbf{1}, 3)$ by a simple computation.

Since $Q^1 = H^1(\mathbf{1}) \cap H^1(-\mathbf{1}) \cap X^1 = \{(x,x,x) \mid 1/4 \leqslant x \leqslant 1\}$ and $\dim Q^1 = 1 = \dim X^1$, we stop the algorithm and conclude that $Q^*(\mathcal{A}) = Q^1$ is the limit set of subgame-perfect equilibrium payoffs as $\delta \to 1$.

The same result is obtained by Fudenberg and Maskin [1986] and Wen [1994]. Fudenberg and Maskin determine the limit set by a direct computation in this specific game, whereas Wen uses effective minimax values. Wen's method is applicable to repeated games with perfect monitoring without the full dimensionality condition. Note that our algorithm is even more general, as we admit imperfect public monitoring and short-run players.

4.2. **Characterization of the Limit Payoffs in General Stage Games with Observed Actions and All Long-Run Players.** Consider repeated games with perfect monitoring and without short-run players, that is, $Y = A$, $\pi_y(a) = 1$ if and only if $y = a$, and $n = L$. We assume that $\mathcal{A}^0 \supseteq \mathcal{A}^{\mathrm{p}} \equiv \{\alpha \in \mathcal{A} \mid \alpha(a) = 1 \text{ for some } a \in A\}$, i.e., $\mathcal{A}^0$ contains all pure action profiles.

We assume that no player is *universally indifferent*: for every player i, there exist two action profiles a, $a' \in A$ such that $g_i(a) \neq g_i(a')$. Players i and j have *equivalent utility functions* if there exist $c \in \mathbb{R}$ and $d > 0$ such that $g_j(a) = c + d g_i(a)$ for all $a \in A$. Denote by I_i^+ the set of players whose utility functions are equivalent to g_i. Similarly, denote by I_i^- the set of players whose utility functions are equivalent to $-g_i$. The stage game satisfies the *nonequivalent utilities (NEU) condition* if $I_i^+ = \{i\}$ for all i (Abreu et al [1994]).

Player i's *effective minimax payoff* is given by

$$\underline{v}_i(\mathcal{A}^0) = \inf_{\alpha \in \mathcal{A}^0} \max\{g_i(a_j, \alpha_{-j}) \mid j \in I_i^+, a_j \in A_j, \text{ or } j \in I_i^-, a_j \in A_j \text{ s.t. } \alpha_j(a_j) > 0\}.$$

If $\mathcal{A}^0$ is compact, then the infimum operator can be replaced by the minimum operator because the objective function is lower semi-continuous in α.

Here we compare our effective minimax payoff with the standard minimax payoff

$$\underline{v}_i^{\mathrm{s}}(\mathcal{A}^0) = \inf_{\alpha \in \mathcal{A}^0} \max\{g_i(a_i, \alpha_{-i}) \mid a_i \in A_i\},$$

and Wen's [1994] effective minimax payoff

$$\underline{v}_i^{\mathrm{Wen}}(\mathcal{A}^0) = \inf_{\alpha \in \mathcal{A}^0} \max\{g_i(a_j, \alpha_{-j}) \mid j \in I_i^+, a_j \in A_j\}.$$

Proposition 4.1. *We have the following relations between the effective minimax and Wen's effective minimax:*

FIGURE 4.2. A game in which $\underline{v}_1^{\mathrm{Wen}}(\mathcal{A}) < \underline{v}_1(\mathcal{A})$

0,0,0	3,3,−3
2,2,−2	4,4,−4

4,4,−4	2,2,−2
3,3,−3	4,4,−4

(1) $\underline{v}_i^{\mathrm{s}}(\mathcal{A}^0) \leqslant \underline{v}_i^{\mathrm{Wen}}(\mathcal{A}^0) \leqslant \underline{v}_i(\mathcal{A}^0)$.
(2) $\underline{v}_i^{\mathrm{Wen}}(\mathcal{A}^0) = \underline{v}_i(\mathcal{A}^0)$ if $\mathcal{A}^0 = \mathcal{A}^{\mathrm{p}}$ or $I_i^- = \emptyset$.
(3) $\underline{v}_i^{\mathrm{s}}(\mathcal{A}^0) = \underline{v}_i(\mathcal{A}^0)$ if ($\mathcal{A}^0 = \mathcal{A}^{\mathrm{p}}$ or $\mathcal{A}$) and the NEU condition is satisfied.

Proof. Parts 1 and 2 are obvious. Part 3 is also obvious, except for the case in which $\mathcal{A}^0 = \mathcal{A}$, the NEU condition is satisfied, and $I_i^- \neq \emptyset$. Since the NEU condition is satisfied and $I_i^- \neq \emptyset$, I_i^- is a singleton $\{j\}$. Let α^*_{-i} be a minimax action profile against player i, and α^*_i is a maximin action of player i against player j when the other players' action profile is fixed to be α^*_{-ij}. By the minimax theorem, (α^*_i, α^*_j) is a Nash equilibrium of the game between players i and j when the other players play α^*_{-ij}. Since α^*_i is a best response to α^*_{-i} for player i, we have $g_i(a_i, \alpha^*_{-i}) \leqslant g_i(\alpha^*)$ for any $a_i \in A_i$. Also, since α^*_j is a best response to α^*_{-j} for player j, player j is indifferent among all pure actions taken with positive probabilities under α^*_j, i.e., we have $g_j(a_j, \alpha^*_{-j}) = g_j(\alpha^*)$ for any $a_j \in A_j$ such that $\alpha^*_j(a_j) > 0$. Since $j \in I_i^-$, we have $g_i(a_j, \alpha^*_{-j}) = g_i(\alpha^*)$ for any $a_j \in A_j$ such that $\alpha^*_j(a_j) > 0$. Therefore, we have $\underline{v}_i(\mathcal{A}) \leqslant g_i(\alpha^*) = \underline{v}_i^{\mathrm{s}}(\mathcal{A})$. □

Example. We may have $\underline{v}_i^{\mathrm{Wen}}(\mathcal{A}^0) < \underline{v}_i(\mathcal{A}^0)$. Consider the stage game in Figure 4.2. Note that $I_1^+ = \{1,2\}$ and $I_1^- = \{3\}$. We have $\underline{v}_1^{\mathrm{Wen}}(\mathcal{A}) = 5/2$, where the solution α to Wen's minimax problem is such that players 1 and 2 choose the first actions, and player 3 mixes the two actions with equal probability. We also have $\underline{v}_1(\mathcal{A}) = 3$, where the solution α to our minimax problem is such that players 1 and 3 choose the first actions, and player 2 chooses the first action with probability more than or equal to $1/2$.

Let V be the set of feasible payoff profiles, i.e., the convex hull of $\{g(a) \in \mathbb{R}^n \mid a \in A\}$. Let

$$V(\mathcal{A}^0) = \{v \in V \mid v_i \geqslant \underline{v}_i(\mathcal{A}^0) \text{ for every player } i\},$$
$$V^*(\mathcal{A}^0) = \{v \in V \mid v_i > \underline{v}_i(\mathcal{A}^0) \text{ for every player } i\}$$

be the sets of feasible payoff profiles that weakly and strongly dominate $\underline{v}(\mathcal{A}^0)$, respectively.

Proposition 4.2. *$Q^*(\mathcal{A}^0) \subseteq V(\mathcal{A}^0)$. If $V^*(\mathcal{A}^0) \neq \emptyset$, then $Q^*(\mathcal{A}^0) = V(\mathcal{A}^0)$.*

Abreu et al [1994] showed the folk theorem under the NEU condition, which corresponds to Proposition 4.2 when ($\mathcal{A}^0 = \mathcal{A}^{\mathrm{p}}$ or $\mathcal{A}$) and the NEU condition is satisfied. Wen [1994] showed the pure-strategy folk theorem, which corresponds to Proposition 4.2 for $\mathcal{A}^0 = \mathcal{A}^{\mathrm{p}}$.[6] These classical results are stronger than Proposition 4.2 in the following sense. They show that $E(\mathcal{A}^{\mathrm{p}}, \delta) \subseteq V(\mathcal{A}^{\mathrm{p}})$ for any δ, and that, for any $v \in V^*(\mathcal{A}^{\mathrm{p}})$, there exists $\underline{\delta} < 1$ such that $v \in E(\mathcal{A}^{\mathrm{p}}, \delta)$ (*exactly* attained as an equilibrium payoff profile) for any $\delta > \underline{\delta}$. On the other hand, combined with our Theorem, Proposition 4.2 claims that any point $v \in V(\mathcal{A}^{\mathrm{p}})$ is *approximately* attained as an equilibrium payoff profile. See Subsection 4.4 for a discussion of the exact attainability of efficient payoffs.

We will show Proposition 4.2 by executing our algorithm explicitly. Let X be the affine hull of V. We have $\dim X \geqslant 1$ because of the absence of universal indifference. A vector $\lambda \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ parallel to X is said to be a *punishment direction for player i* if there exist $c \in \mathbb{R}$ and $d > 0$ such that $\lambda \cdot v = c - dv_i$ for every $v \in X$. If λ is a punishment direction for player i, then we have $H(\lambda, \lambda \cdot v) \cap X = \{v' \in X \mid v'_i \geqslant v_i\}$.

Lemma 4.3. *There exists a punishment direction for player i.*

Proof. Let λ be the projection of $-e^i$ to X, where e^i is the vector whose ith component is 1 and whose other components are 0. λ is nonzero since player i is not universally indifferent. By construction, λ is a punishment direction for player i. □

Let $X_i = \{v_i \in \mathbb{R} \mid v \in X\}$ and $X_{ij} = \{(v_i, v_j) \in \mathbb{R}^2 \mid v \in X\}$ be the projections of X to the i-axis and to the ij-plane, respectively.

Lemma 4.4. *$X_i = \mathbb{R}$; if $j \notin I_i^+ \cup I_i^-$, then $X_{ij} = \mathbb{R}^2$.*

Proof. X_i is a nonempty affine subspace of $\mathbb{R}$, i.e., a point or $\mathbb{R}$. Since player i is not universally indifferent, X_i contains at least two points. So we have $X_i = \mathbb{R}$.

[6] As we noted in the example, in the class of mixed-strategy subgame-perfect equilibria ($\mathcal{A}^0 = \mathcal{A}$), Wen's definition of effective minimax may be lower than ours. In this case, the effective minimax value in his definition is not a tight lower bound for mixed-strategy subgame-perfect equilibrium payoffs. Wen assumes that mixed strategies are observable and constructs equilibria with payoffs as low as $\underline{v}_i^{\mathrm{Wen}}(\mathcal{A})$ in general games. Our results show that the assumption that mixed strategies are observable is not innocuous in cases where the NEU condition is not satisfied. Intuitively, inducing players to randomize when mixing probabilities are not observed requires the use of continuation payoffs to make the player indifferent, and in the absence of the NEU condition, it may not be possible to induce players to randomize without rewarding the opponent they are trying to "punish." See footnote 11 in Abreu et al [1994].

X_{ij} is a nonempty affine subspace of $\mathbb{R}^2$, i.e., a point, a line, or $\mathbb{R}^2$, and from the previous step X_{ij} is not a point or a vertical or horizontal line. Since $j \notin I_i^+ \cup I_i^-$, X_{ij} is not a line with a nonzero slope, so $X_{ij} = \mathbb{R}^2$. □

We also have

Lemma 4.5. *A mixed action profile α and player i's payoff v_i are enforceable with respect to $\{v' \in X \mid v'_i \geqslant v_i\}$ if and only if $v_i \geqslant g_i(a_j, \alpha_{-j})$ for any $j \in I_i^+$ and any $a_j \in A_j$ and for any $j \in I_i^-$ and any $a_j \in A_j$ such that $\alpha_j(a_j) > 0$.*

Proof. See the Appendix. □

Lemma 4.6. *If λ is not a punishment direction for player i, then, for any $(x,k) \in \mathbb{R}^2$, there exists $v \in H(\lambda,k) \cap X$ such that $v_i \leqslant x$.*

Proof. Since λ is not a punishment direction for player i, $\lambda \cdot v$ and $-v_i$ are linear utility functions that represent different preference orderings on X. Then there exist v^1, $v^2 \in X$ such that (i) $\lambda \cdot v^1 \geqslant \lambda \cdot v^2$ and $v_i^1 > v_i^2$, or (ii) $\lambda \cdot v^1 > \lambda \cdot v^2$ and $v_i^1 = v_i^2$.

In case (i), pick any $v^3 \in H(\lambda,k) \cap X$, and let $v = v^3 - c(v^1 - v^2)$. Then we have $v \in H(\lambda,k) \cap X$ and $v_i \leqslant x$ for a sufficiently large c.

In case (ii), pick any v^4, $v^5 \in X$ such that $v_i^4 > v_i^5$, and let $\tilde{v}^1 = v^1 + \varepsilon(v^4 - v^5)$. It follows from Lemma 4.4 that such v^4 and v^5 exist. For a sufficiently small $\varepsilon > 0$, we have $\lambda \cdot \tilde{v}^1 > \lambda \cdot v^2$ and $\tilde{v}_i^1 > v_i^2$. Thus we can apply case (i) to the pair $(\tilde{v}^1, v^2)$. □

Lemma 4.7. *If λ is not a punishment direction for any player, then, every pure strategy profile $a \in A$ and the corresponding payoff profile $v = g(a)$ are enforceable with respect to $H(\lambda, \lambda \cdot v) \cap X$.*

Proof. Define $w(a') \in H(\lambda, \lambda \cdot v) \cap X$ for each $a' \in A$ as follows:

- If there exists a unique player i such that $a'_i \neq a_i$, then, because of Lemma 4.6, we can construct a sufficiently strong punishment for player i by setting $w(a') \in H(\lambda, \lambda \cdot v) \cap X$ such that $w_i(a') \leqslant [v_i - (1-\delta) g_i(a'_i, a_{-i})]/\delta$.
- If $a' = a$ or $a'_j \neq a_j$ for at least two players j, then let $w(a') = v$.

Then (a, v) is enforced by w. □

Lemma 4.8. *If $V^*(\mathcal{A}^0) \neq \emptyset$, then $\dim V^*(\mathcal{A}^0) = \dim X$.*

Proof. Here we use the relative topology induced to X. Suppose $V^*(\mathcal{A}^0) \neq \emptyset$. Then there exists a relative interior point v of V such that $v \in V^*(\mathcal{A}^0)$. Otherwise, $V \setminus V^*(\mathcal{A}^0)$ is a closed proper subset of V that contains the whole relative interior of V. This contradicts the fact that the closure of the relative interior of compact and convex set V is equal to V. Since $V \setminus V^*(\mathcal{A}^0)$ is closed, v is also a relative interior point of $V^*(\mathcal{A}^0)$, so $V^*(\mathcal{A}^0)$ and X have the same dimension. □

Now we can prove Proposition 4.2 as follows.

Proof. We use our algorithm with the constraint $X^0 = X$ on continuation payoff profiles at the first step. Since $\mathcal{A}^0 \supseteq \mathcal{A}^{\mathrm{p}}$, it follows from Lemmas 4.3, 4.5, and 4.7 that we have $Q^0 = V(\mathcal{A}^0)$. Therefore, we have $Q^*(\mathcal{A}^0) \subseteq V(\mathcal{A}^0)$.

If $V^*(\mathcal{A}^0) \neq \emptyset$, then, by Lemma 4.8, we have $\dim Q^0 = \dim X^0$. We stop the algorithm at the first step, and obtain $Q^*(\mathcal{A}^0) = V(\mathcal{A}^0)$. □

4.3. **Symmetry Assumptions.**

4.3.1. *Strongly Symmetric Equilibria.* Assume that the static game is symmetric for long-run players, i.e., $A_1 = \cdots = A_L$ and $g_i(a) = g_j(a')$ for any $i, j \in LR$ and $a, a' \in A$ if $a_i = a'_j$, a'_{LR} is a permutation of a_{LR}, and $a_{SR} = a'_{SR}$. The signal structure is also symmetric, i.e., $\pi_y(a) = \pi_y(a')$ if a'_{LR} is a permutation of a_{LR}, and $a_{SR} = a'_{SR}$.

A strategy profile is *strongly symmetric* (for long-run players) if all long-run players take the same action after every history. In this case we take $\mathcal{A}^0$ to be the set $\mathcal{A}^{\mathrm{s}}$ of symmetric mixed action profiles for the long-run players in graph(B), and denote by Q^{s} the result $Q^*(\mathcal{A}^{\mathrm{s}})$ of our algorithm under the restriction of $\mathcal{A}^{\mathrm{s}}$. Our Theorem can characterize the limit of $E(\mathcal{A}^{\mathrm{s}}, \delta)$ by Q^{s}. Set $X^0 = \{(x, \dots, x) \in \mathbb{R}^L \mid x \in \mathbb{R}\}$, and compute Q^0 in the first step of our algorithm. Since $\mathcal{A}^{\mathrm{s}}$ contains at least one static equilibrium, we have $Q^0 \neq \emptyset$. No matter whether Q^0 is a singleton (which must be a unique symmetric static equilibrium payoff) or one-dimensional, we have $Q^{\mathrm{s}} = Q^0$. Since continuation payoffs are restricted to be symmetric, Q^{s} may be strictly smaller than FL's Q without any restriction on continuation payoffs. This corresponds to Abreu et al's [1986] analysis for large δ.

As a corollary of our Theorem, we have the following.

Corollary 4.9. $Q^{\mathrm{s}} = \lim_{\delta \to 1} E(\mathcal{A}^{\mathrm{s}}, \delta)$. *That is, Q^{s} is the limit as δ goes to 1 of strongly symmetric equilibrium payoffs with discount factor δ.*

4.3.2. *Partially Symmetric Equilibria.* We can consider *partially symmetric* equilibria. Suppose that long-run and short-run players are divided into several groups, for example, buyers and sellers. The players' payoffs are symmetric within groups, but may be asymmetric between groups. Then we can restrict our attention to partially symmetric equilibria where the players behave symmetrically within groups. As in the case of strongly symmetric equilibria, let X^0 be the set of payoff profiles symmetric within groups.

Then, we can execute the first step of our algorithm, in which continuation payoffs are constrained to be symmetric within groups.

Note that the FL result, on the one hand, cannot apply to partially symmetric equilibria when there are $L-1$ or less groups because Q^0 does not satisfy the full-dimensionality condition. On the other hand, it is possible to apply Abreu et al's [1990] result and obtain the set of partially symmetric equilibria for any fixed δ, but when the number of groups for long-run players is 2 or more, it is difficult and sometimes practically infeasible to compute the set $P(\mathcal{A}^0, \delta, W)$ generated by W for any nonlinear constraint W on continuation payoff profiles. By contrast, our algorithm is applicable and relatively easy to carry out.

4.4. **Exact Achievability of First-Best Outcomes.** In the case of $\mathcal{A}^0 = \text{graph}(B)$ and $X^0 = \mathbb{R}^L$, FL showed that, under the assumption of $\dim Q^0 = L$, for any compact set K in the interior of Q^0, there exists $\bar{\delta} < 1$ such that $K \subseteq E(\text{graph}(B), \delta)$ for any $\delta > \bar{\delta}$. Under an identifiability condition, Q^0 is a full-dimensional set containing all payoff profiles that Pareto-dominate a static equilibrium (Fudenberg et al [1994, Theorem 6.1]). When this identifiability condition is satisfied, some efficient payoff profiles can be approximated by equilibrium payoff profiles as the discount factor tends to 1, even if the actions are imperfectly observed. However, this conclusion leaves open the question of whether a given efficient payoff vector v can be exactly attained by an equilibrium payoff for some large but fixed δ.

Recently Athey and Bagwell [2001] have provided sufficient conditions for the exact achievability of first-best payoffs in a repeated duopoly game. Our Theorem leads to the following generalization of their analysis.

Let V be the convex hull of $\{g_{LR}(\alpha) \in \mathbb{R}^L \mid \alpha \in \text{graph}(B)\}$, let h be a hyperplane tangent to V, and let $\mathcal{A}^h = \{\alpha \in \text{graph}(B) \mid g_{LR}(\alpha) \in h\}$. To achieve a payoff profile in h, it is necessary for the players to take actions in $\mathcal{A}^h$ at any on-path history (a public history which occurs with positive probability). As an extreme case, if $V \cap h$ is a singleton $\{v\}$, then exactly achieving v requires a stringent condition (Fudenberg et al [1994, Theorem 6.5]).

Here we sketch how to obtain a sufficient condition for exact achievability. By our algorithm, we can characterize the limit of $E(\mathcal{A}^h, \delta)$.[7] Let

[7]Imposing the restriction of $\mathcal{A}^h$ on off-path play does not lose much generality. If the full support condition holds for $\mathcal{A}^h$, i.e., $\pi_y(\alpha) > 0$ for any $\alpha \in \mathcal{A}^h$ and $y \in Y$, then there is no off-path public history, and hence any perfect public equilibrium which achieves a payoff profile in h is always an $\mathcal{A}^h$-perfect public equilibrium. Moreover, if the full support condition is not satisfied but there is an inefficient static equilibrium, we can easily modify our argument by analyzing on- and off-schedule deviations separately. See Athey and

$X^0 = h$. We compute the algorithm until we finally obtain $Q^*(\mathcal{A}^h)$. We denote it by Q^h.

Our Theorem implies the following.

Corollary 4.10. *If $Q^h \neq \emptyset$, then the relative interior of Q^h is nonempty, and for any relative interior point v of Q^h, there exists $\bar{\delta} < 1$ such that $v \in E(\mathcal{A}^h, \delta)$ for any $\delta > \bar{\delta}$.*

Equilibria in Q^h have the property that there is no history where players unanimously prefer some other feasible outcome to the continuation payoffs prescribed by the equilibria. This is a very strong form of renegotiation-proofness, and implies that the equilibria are strongly renegotiation-proof in the sense of Farrell and Maskin [1989].

In the case of two-player games, we can give a simple necessary and sufficient condition for Q^h to be nonempty. Fix two pure action profiles a^1 and a^2 whose payoff vectors $g^1 = g(a^1)$ and $g^2 = g(a^2)$ are on the Pareto frontier. Suppose that $g_i^i < g_i^j$ for $i \neq j$, so that a^i is worse for player i than a^j. Let $h = \{v \in \mathbb{R}^2 \mid \beta_1 v_1 + \beta_2 v_2 = \gamma\}$ with $(\beta_1, \beta_2, \gamma) = (g_2^1 - g_2^2, g_1^2 - g_1^1, g_1^2 g_2^1 - g_1^1 g_2^2)$ be the line connecting g^1 and g^2. Note that β_1, $\beta_2 > 0$. We assume $\mathcal{A}^h = \{a^1, a^2\}$ for simplicity, i.e., no payoff profile other than g^1 and g^2 attains h.[8] We focus on $\mathcal{A}^h$-perfect public equilibria.

Let $g^{ii} = (g_i^{ii}, g_j^{ii}) \in h$ be defined by $g_i^{ii} = \max_{a_i} g_i(a_i, a_j^i)$ and $g_j^{ii} = (\gamma - \beta_i g_i^{ii})/\beta_j$. Let $\ell = \{v \in h \mid g_1^{11} \leqslant v_1 \leqslant g_1^{22}\}$. ℓ is the empty set if $g_1^{11} > g_1^{22}$, and a line segment if $g_1^{11} < g_1^{22}$. (If $g_1^{11} = g_1^{22}$, then ℓ is a point.)

Similarly, let $g^{ij} = (g_i^{ij}, g_j^{ij}) \in h$ be defined by $g_j^{ij} = \max_{a_j \neq a_j^i} g_j(a_j, a_i^i)$ and $g_i^{ij} = (\gamma - \beta_j g_j^{ij})/\beta_i$.

Definition. The signal structure π has *perfect detectability for player i* if the set Y of outcomes can be partitioned into Y^i, Y^{ii}, and Y^{ij} such that for

$$\pi^i(a) \equiv \sum_{y \in Y^i} \pi_y(a), \qquad \pi^{ii}(a) \equiv \sum_{y \in Y^{ii}} \pi_y(a), \qquad \pi^{ij}(a) \equiv \sum_{y \in Y^{ij}} \pi_y(a),$$

there exist ω^i, $\omega^{ij} \geqslant 0$ such that

(1) $\pi^{ii}(a^i) < 1$ and $\pi^i(a^i)\omega^i + \pi^{ij}(a^i)\omega^{ij} = g_i^{ii} - g_i^i$,

---

Bagwell [2001]. Note also that allowing off-path play not in $\mathcal{A}^h$ destroys the renegotiation-proofness property of the equilibria.

[8] $\mathcal{A}^h = \{a^1, a^2\}$ if $a_i^1 \neq a_i^2$ for both players i and no three pure action payoff profiles lie on a line. The latter condition is satisfied in generic finite stage games, but not in duopoly games where both firms produce homogeneous goods with a common constant marginal cost. If $\mathcal{A}^h \supsetneq \{a^1, a^2\}$, then the condition of $g_1^{11} \geqslant g_1^{22}$ in Part 1 of Proposition 4.11 is not sufficient for Q^h to be empty. There may be an equilibrium which prescribes action profiles in $\mathcal{A}^h \setminus \{a^1, a^2\}$ after some history.

(2) $\pi^{ii}(a_i, a^i_j) = 1$ for any $a_i \neq a^i_i$, and

(3) $\pi^i(a_j, a^i_i)\omega^i + \pi^{ij}(a_j, a^i_i)\omega^{ij} > g^{ii}_i - g^{ij}_i$ for any $a_j \neq a^i_j$.

If π is perfect monitoring (i.e., $Y = A$ and $\pi_y(a) = 1$ if and only if $y = a$), then π has perfect detectability for player i by setting $Y^i = \{a^i\}$, $Y^{ii} = \{(a_i, a^i_j) \mid a_i \neq a^i_i\}$, $Y^{ij} = \{a \mid a_j \neq a^i_j\}$, $\omega^i = g^{ii}_i - g^i_i$, and $\omega^{ij} = \max(g^{ii}_i - g^{ij}_i + 1, 0)$.

Proposition 4.11. *Suppose that neither a^1 nor a^2 is a static Nash equilibrium.*

(1) *If $g^{11}_1 \geqslant g^{22}_1$, then $Q^h = \emptyset$ under any signal structure, and there is no $\mathcal{A}^h$-perfect public equilibrium for any δ.*

(2) *If $g^{11}_1 < g^{22}_1$, then for any signal structure π with perfect detectability for the both players and any compact line segment in the relative interior of ℓ, there exists $\varepsilon > 0$ such that Q^h is a nonempty set containing the line segment under any signal structure $\tilde{\pi}$ such that $\max_{y,a} |\pi_y(a) - \tilde{\pi}_y(a)| < \varepsilon$, so there exists $\bar{\delta} < 1$ such that $v \in E(\mathcal{A}^h, \delta)$ for any $\delta > \bar{\bar{\delta}}$.*

Proof. see Appendix. □

REFERENCES

[1] Abreu, D. [1986] "Extremal equilibria of oligopolistic supergames," *Journal of Economic Theory* 39, 191-228.

[2] Abreu, D., P. K. Dutta, and L. Smith [1994] "The folk theorem for repeated games: a NEU condition," *Econometrica* 62, 939-948.

[3] Abreu, D., D. Pearce, and E. Stacchetti [1986] "Optimal cartel equilibria with imperfect monitoring," *Journal of Economic Theory* 39, 251-269.

[4] Abreu, D., D. Pearce, and E. Stacchetti [1990] "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica* 58, 1041-1063.

[5] Athey, S. and K. Bagwell [2001] "Optimal collusion with private information," *RAND Journal of Economics* 32, 428-465.

[6] Dellarocas, C. [2003] "Efficiency and robustness of eBay-like online reputation mechanisms in environments with moral hazard," working paper, MIT.

[7] Ely, J., D. Fudenberg, and D. K. Levine [2003] "When is reputation bad?", mimeo.

[8] Farrell, J. and E. Maskin [1989] "Renegotiation in repeated games," *Games and Economic Behavior* 1, 327-360.

[9] Fudenberg, D. and D. K. Levine [1994] "Efficiency and observability with long-run and short-run players," *Journal of Economic Theory* 62, 103-135.

[10] Fudenberg, D., D. K. Levine, and E. Maskin [1994] "The folk theorem in repeated games with imperfect public information," *Econometrica* 62, 997-1039.

[11] Fudenberg, D. and E. Maskin [1986] "The folk theorem in repeated games with discounting or with incomplete information," *Econometrica* 54, 533-554.

[12] Fudenberg, D. and E. Maskin [1990] "Nash and perfect equilibria of discounted repeated games," *Journal of Economic Theory,* 51, 194-206.

[13] Kandori M. and H. Matsushima [1998] "Private observation, communication and collusion," *Econometrica* 66, 627-652.
[14] Wen, Q. [1994] "The 'folk theorem' for repeated games with complete information," *Econometrica* 62, 949-954.

APPENDIX A. PROOFS

A.1. **Proof of Lemma 4.5.** We use the following lemma in the proof of Lemma 4.5 below to deal with indifference conditions for players in $J = I_i^+ \cup I_i^-$. Fudenberg and Maskin [1990, Lemma 2] prove the same result for the case of two players. For notational convenience, for a given mixed action profile α, take $S_j = \{a_j \in A_j \mid \alpha_j(a_j) > 0\}$, $S = \prod_{j \in J} S_j$, $\sigma = \alpha_J$, and $u(\sigma) = g_i(\sigma, \alpha_{-J})$.

Lemma A.1. *If $\sigma_j(s_j) > 0$ and $x \geqslant u(s_j, \sigma_{-j})$ for all $j \in J$ and all $s_j \in S_j$, then there exists $f\colon S \to \mathbb{R}$ such that $f(s) \geqslant x$ for all $s \in S$ and*

$$x = (1-\delta)u(s_j,\sigma_{-j}) + \delta \sum_{s_{-j} \in S_{-j}} \sigma_{-j}(s_{-j}) f(s)$$

for all $j \in J$ and all $s_j \in S_j$.

Proof. We will show the existence of f algorithmically. Let $S_j^0 = S_j$ and $r_j^0(s_j) = [x - (1-\delta)u(s_j,\sigma_{-j})]/\delta$ for each $j \in J$ and each $s_j \in S_j$. For each step $m = 0, 1, \ldots$, we define

$$p_j^m = \sum_{s_j \in S_j^m} \sigma_j(s_j), \quad p^m = \prod_{j \in J} p_j^m, \quad p_{-j}^m = \frac{p^m}{p_j^m},$$

$$x^m = \min\left\{ \frac{r_j^m(s_j)}{p_{-j}^m} \mid j \in J, s_j \in S_j^m \right\},$$

$$S_j^{m+1} = S_j^m \setminus \left\{ s_j \in S_j^m \mid \frac{r_j^m(s_j)}{p_{-j}^m} = x^m \right\},$$

$$r_j^{m+1}(s_j) = r_j^m(s_j) - (p_{-j}^m - p_{-j}^{m+1})x^m.$$

Note that $\sum_{s_j \in S_j^0} \sigma_j(s_j) r_j^0(s_j) = [x-(1-\delta)u(\sigma)]/\delta$ is independent of j. Inductively, if $\sum_{s_j \in S_j^m} \sigma_j(s_j) r_j^m(s_j)$ is independent of j, then

$$\begin{aligned}\sum_{s_j \in S_j^{m+1}} \sigma_j(s_j) r_j^{m+1}(s_j) &= \sum_{s_j \in S_j^{m+1}} \sigma_j(s_j) \left[r_j^m(s_j) - (p_{-j}^m - p_{-j}^{m+1}) x^m \right] \\ &= \sum_{s_j \in S_j^{m+1}} \sigma_j(s_j) r_j^m(s_j) - p_j^{m+1}(p_{-j}^m - p_{-j}^{m+1}) x^m \\ &= \sum_{s_j \in S_j^m} \sigma_j(s_j) r_j^m(s_j) - (p^m - p^{m+1}) x^m\end{aligned}$$

is also independent of j.

Let m^* be the first step at which there exists $j^* \in J$ such that $S_{j^*}^{m^*+1} = \emptyset$. If $S_j^{m^*+1} \neq \emptyset$ for some $j \in J$, then we have the following contradiction:

$$\begin{aligned}p^{m^*} x^{m^*} &= \sum_{s_{j^*} \in S_{j^*}^{m^*}} \sigma_{j^*}(s_{j^*}) r_{j^*}^{m^*}(s_{j^*}) \\ &= \sum_{s_j \in S_j^{m^*}} \sigma_j(s_j) r_j^{m^*}(s_j) \\ &> \sum_{s_j \in S_j^{m^*}} \sigma_j(s_j) p_{-j}^{m^*} x^{m^*} = p^{m^*} x^{m^*},\end{aligned}$$

where the second equality holds because $\sum_{s_j \in S_j^{m^*}} \sigma_j(s_j) r_j^{m^*}(s_j)$ is independent of j, and the inequality holds because $r_j^{m^*}(s_j)/p_{-j}^{m^*} \geqslant x^{m^*}$ for all $s_j \in S_j^{m^*}$ with strict inequality for all $s_j \in S_j^{m^*+1}$, and $\sigma_j(s_j) > 0$ for all $s_j \in S_j^{m^*}$. Thus we have $S_j^{m^*+1} = \emptyset$ for all $j \in J$.

Since $p_j^0 = 1$ for all $j \in J$ and $x \geqslant \max\{u(s_j, \sigma_{-j}) \mid j \in J, s_j \in S_j\}$, we have

$$x^0 = \min\{r_j^0(s_j) \mid j \in J, s_j \in S_j\} = \frac{x - (1-\delta)\max\{u(s_j, \sigma_{-j}) \mid j \in J, s_j \in S_j\}}{\delta} \geqslant x.$$

For any $m < m^*$ and for any $s_j \in S_j^{m+1}$, we have

$$x^{m+1} \geqslant \frac{r_j^{m+1}(s_j)}{p_{-j}^{m+1}} = \frac{r_j^m(s_j) - (p_{-j}^m - p_{-j}^{m+1}) x^m}{p_{-j}^{m+1}} > \frac{p_{-j}^m x^m - (p_{-j}^m - p_{-j}^{m+1}) x^m}{p_{-j}^{m+1}} = x^m.$$

Therefore, we have $x^m \geqslant x$ for any $m \leqslant m^*$.

Let $S^m \equiv \prod_{j \in J} S_j^m$. For any $s \in S$, there exists a unique $m(s) \leqslant m^*$ such that $s \in S^{m(s)} \setminus S^{m(s)+1}$. Then we define $f(s) = x^{m(s)} \geqslant x$ for each $s \in S$.

For any $j \in J$ and any $s_j \in S_j$, there exists a unique $m_j(s_j) \leqslant m^*$ such that $s_j \in S_j^{m_j(s_j)} \setminus S_j^{m_j(s_j)+1}$ because $S_j^{m^*+1} = \emptyset$. Then we have

$$\begin{aligned}
&(1-\delta)u(s_j,\sigma_{-j}) + \delta \sum_{s_{-j}\in S_{-j}} \sigma_{-j}(s_{-j}) f(s) \\
&= (1-\delta)u(s_j,\sigma_{-j}) + \delta \sum_{m=0}^{m_j(s_j)-1} (p_{-j}^m - p_{-j}^{m+1})x^m + p_{-j}^{m_j(s_j)} x^{m_j(s_j)} \\
&= (1-\delta)u(s_j,\sigma_{-j}) + \delta \sum_{m=0}^{m_j(s_j)-1} (r_j^m(s_j) - r_j^{m+1}(s_j)) + r_j^{m_j(s_j)}(s_j) \\
&= (1-\delta)u(s_j,\sigma_{-j}) + \delta r_j^0(s_j) = x
\end{aligned}$$

for any $j \in J$ and any $s_j \in S_j$. □

Lemma. 4.5. *A mixed action profile α and player i's payoff v_i are enforceable with respect to $\{v' \in X \mid v'_i \geqslant v_i\}$ if and only if $v_i \geqslant g_i(a_j,\alpha_{-j})$ for any $j \in I_i^+$ and any $a_j \in A_j$ and for any $j \in I_i^-$ and any $a_j \in A_j$ such that $\alpha_j(a_j) > 0$.*

Proof. "If" part. Define $w(a) \in X$ with $w_i(a) \geqslant v_i$ for each $a \in A$ as follows:

- If there exists a unique player j such that $\alpha_j(a_j) = 0$ and $j \in I_i^+$, then let $w(a) \in X$ be such that $v_i \leqslant w_i(a) \leqslant [v_i - (1-\delta)g_i(a_j,\alpha_{-j})]/\delta$. It follows from Lemma 4.4 and $v_i \geqslant g_i(a_j,\alpha_{-j})$ that such $w(a)$ exists.
- If there exists a unique player j such that $\alpha_j(a_j) = 0$ and $j \notin I_i^+$, then let $w(a) \in X$ be such that $w_i(a) \geqslant v_i$ and $w_j(a) \leqslant [v_j - (1-\delta)g_j(a_j,\alpha_{-j})]/\delta$. The existence of such $w(a)$ follows from Lemma 4.4 .
- If we have $\alpha_j(a_j) > 0$ for all players j, let $J = I_i^+ \cup I_i^-$ and define $w(a) = w^J(a_J) + \sum_{j\notin J}(w^0 - w^j(a_j)) \in X$ as follows: By setting $x = v_i$ in Lemma A.1, we obtain a function f such that $f(a_J) \geqslant v_i$, any action in S_j is indifferent for each player $j \in J$, and player i's total payoff is equal to v_i. It follows from Lemma 4.4 that there exists $w^J(a_J) \in X$ such that $w_i^J(a_J) = f(a_J)$, and by the same lemma we can make any action in S_j indifferent for player $j \notin J$ without changing player i's payoff. For example, pick any $w^0 \in X$, and choose $w^j(a_j) \in X$ such that $w_i^j(a_j) = w_i^0$ and $w_j^j(a_j) = [(1-\delta)/\delta]g_j(a_j,\alpha_{-j})$.
- If $\alpha_j(a_j) = 0$ for at least two players j, let $w(a) \in X$ be such that $w_i(a) = v_i$. The existence of such $w(a)$ follows from Lemma 4.4.

Then (α, v_i) is enforced by w.

"Only if" part. Suppose that (α, v_i) is enforced by continuation payoff profiles $w(a) \in X$ with $w_i(a) \geqslant v_i$.

For any $j \in I_i^+$ and any $a_j \in A_j$, it follows from player j's incentive constraints that we have

$$v_j \geqslant (1-\delta)g_j(a_j,\alpha_{-j}) + \delta \sum_{a_{-j}\in A_{-j}} \alpha_{-j}(a_{-j})w_j(a).$$

Since $j \in I_i^+$, we can transform the above inequality to the following inequality about player i's payoffs:

$$\begin{aligned} v_i &\geqslant (1-\delta)g_i(a_j,\alpha_{-j}) + \delta \sum_{a_{-j}\in A_{-j}} \alpha_{-j}(a_{-j})w_i(a) \\ &\geqslant (1-\delta)g_i(a_j,\alpha_{-j}) + \delta v_i, \end{aligned}$$

thus we have $v_i \geqslant g_i(a_j,\alpha_{-j})$.

For any $j \in I_i^-$ and any $a_j \in A_j$ such that $\alpha_j(a_j) > 0$, we have

$$v_j = (1-\delta)g_j(a_j,\alpha_{-j}) + \delta \sum_{a_{-j}\in A_{-j}} \alpha_{-j}(a_{-j})w_j(a).$$

Since $j \in I_i^-$, we have

$$\begin{aligned} v_i &= (1-\delta)g_i(a_j,\alpha_{-j}) + \delta \sum_{a_{-j}\in A_{-j}} \alpha_{-j}(a_{-j})w_i(a) \\ &\geqslant (1-\delta)g_i(a_j,\alpha_{-j}) + \delta v_i, \end{aligned}$$

thus we have $v_i \geqslant g_i(a_j,\alpha_{-j})$. □

A.2. **Proof of Proposition 4.11.**

Proposition. 4.11. *Suppose that neither a^1 nor a^2 is a static Nash equilibrium.*

(1) *If $g_1^{11} \geqslant g_1^{22}$, then $Q^h = \emptyset$ under any signal structure, and there is no $\mathcal{A}^h$-perfect public equilibrium for any δ.*
(2) *If $g_1^{11} < g_1^{22}$, then for any signal structure π with perfect detectability for the both players and any compact line segment in the relative interior of ℓ, there exists $\varepsilon > 0$ such that Q^h is a nonempty set containing the line segment under any signal structure $\tilde{\pi}$ such that $\max_{y,a}|\pi_y(a) - \tilde{\pi}_y(a)| < \varepsilon$, so there exists $\bar{\bar{\delta}} < 1$ such that $v \in E(\mathcal{A}^h,\delta)$ for any $\delta > \bar{\bar{\delta}}$.*

Proof. [Proof of Part 1] We use our algorithm under the restriction of $\mathcal{A}^h = \{a^1,a^2\}$ to compute the set of $\mathcal{A}^h$-perfect public equilibrium payoff profiles. Let Q^0 be the result of the first step of the algorithm when continuation payoffs are restricted to h. Since this is a one-dimensional problem, we only need to consider two directions $\lambda^1 = (\beta_2,-\beta_1)$ and $\lambda^2 = (-\beta_2,\beta_1)$.

Consider the linear programming problem for action profile a^2 and direction λ^1. Let (v,w) be any collection of payoff profiles satisfying constraints (a)-(d). For any $y \in Y$, since v, $w(y)$ are chosen from h and $\lambda^1 \cdot v \geqslant \lambda^1 \cdot w(y)$, we have $v_2 \leqslant w_2(y)$. Then, by constraint (b), we have

$$v_2 \geqslant (1-\delta)g_2^{22} + \delta \sum_y \pi_y(a_1^2, a_2^{22}) w_2(y) \geqslant (1-\delta)g_2^{22} + \delta v_2,$$

where a_2^{22} is player 2's action that maximizes $g_2(a_1^2, a_2)$. Therefore, $v_2 \geqslant g_2^{22}$, $\lambda^1 \cdot v \leqslant \lambda^1 \cdot g^{22}$, and $k^0(a^2, \lambda^1) \leqslant \lambda^1 \cdot g^{22}$. Since $k^0(a^1, \lambda^1) \leqslant \lambda^1 \cdot g^1$, we have

$$k^0(\lambda^1) \leqslant \max(\lambda^1 \cdot g^{22}, \lambda^1 \cdot g^1).$$

Similarly, we have $k^0(\lambda^2) \leqslant \max(\lambda^2 \cdot g^{11}, \lambda^2 \cdot g^2)$. Therefore,

$$\begin{aligned} Q^0 &= H^0(\lambda^1) \cap H^0(\lambda^2) \cap h \\ &\subseteq \{v \in h \mid \min(g_1^{11}, g_1^2) \leqslant v_1 \leqslant \max(g_1^{22}, g_1^1)\} \equiv \ell'. \end{aligned}$$

Since $g_1^{11} \geqslant g_1^{22}$, $g_1^{11} \geqslant g_1^1$, $g_1^2 \geqslant g_1^{22}$, and $g_1^2 > g_1^1$, we have $\min(g_1^{11}, g_1^2) \geqslant \max(g_1^{22}, g_1^1)$. Therefore, ℓ' is the empty set or a singleton. Since neither a^1 nor a^2 is a static Nash equilibrium, we have $Q^h = \emptyset$.

[Proof of Part 2] Similarly to Part 1, let Q^0 be the result of the second step of the algorithm when continuation payoffs are restricted to lie on h. Let Y^1, Y^{11}, Y^{12}, ω^1, and ω^{12} be defined in Definition. We will show that for any η with

$$0 < \eta < \min_{a_2 \neq a_2^1} (\pi^1(a_1^1, a_2)\omega^1 + \pi^{12}(a_1^1, a_2)\omega^{12}) - (g_1^{11} - g_1^{12}),$$

action profile a^1 and payoff vector $v = g^{11} + (\eta, -(\beta_1/\beta_2)\eta)$ can be enforced for the both players with continuation payoffs on the ray $\{v' \in h \mid v_1' \geqslant v_1\}$ when $\tilde{\pi}$ is sufficiently close to π. Let the continuation payoffs be

$$\begin{aligned} w_1(y) &= \begin{cases} v_1 + [(1-\delta)/\delta](\omega^1 + \zeta) & (y \in Y^1), \\ v_1 & (y \in Y^{11}), \\ v_1 + [(1-\delta)/\delta](\omega^{12} + \zeta) & (y \in Y^{12}), \end{cases} \\ w_2(y) &= \frac{\gamma - \beta_1 w_1(y)}{\beta_2}, \end{aligned}$$

where

$$\zeta = \frac{(v_1 - g_1^1) - (\tilde{\pi}^1(a^1)\omega^1 + \tilde{\pi}^{12}(a^1)\omega^{12})}{1 - \tilde{\pi}^{11}(a^1)}.$$

Since ω^1, $\omega^{12} \geqslant 0$ and $\zeta \to \eta/(1 - \pi^{11}(a^1)) > 0$ as $\tilde{\pi} \to \pi$, all $w(y)$ satisfy $w_1(y) \geqslant v_1$ when $\tilde{\pi}$ is close to π.

Observe that the overall payoff for player 1 that is generated by a^1 and $w(y)$ is

$$\begin{aligned}
&(1-\delta)g_1^1+\delta\sum_y \tilde{\pi}_y(a^1)w_1(y) \\
&= (1-\delta)g_1^1+\delta\left[\tilde{\pi}^1(a^1)\left\{v_1+\frac{1-\delta}{\delta}(\omega^1+\zeta)\right\}+\tilde{\pi}^{11}(a^1)v_1\right. \\
&\qquad\left.+\tilde{\pi}^{12}(a^1)\left\{v_1+\frac{1-\delta}{\delta}(\omega^{12}+\zeta)\right\}\right] \\
&= v_1.
\end{aligned}$$

If player 1 deviates to $a_1 \neq a_1^1$, his payoff is at most

$$(1-\delta)g_1^{11}+\delta\sum_y \tilde{\pi}_y(a_1,a_2^1)w_1(y),$$

which converges to $(1-\delta)g_1^{11}+\delta v_1$ as $\tilde{\pi}\to\pi$. The limit is less than v_1 since $\eta>0$. Therefore, this deviation is unprofitable for player 1 when $\tilde{\pi}$ is close to π.

If player 2 deviates to $a_2 \neq a_2^1$, his payoff is at most

$$(1-\delta)g_2^{12}+\delta\sum_y \tilde{\pi}_y(a_1^1,a_2)w_2(y).$$

Since g^{12} and $w(y)$ are on the line h, this payoff is written as $(\gamma-\beta_1 v_1')/\beta_2$, where

$$v_1' = (1-\delta)g_1^{12}+\delta\sum_y \tilde{\pi}_y(a_1^1,a_2)w_1(y).$$

Since β_1, $\beta_2>0$, this deviation is unprofitable for player 2 if $v_1' \geqslant v_1$ for $\tilde{\pi}$ close to π. Taking $\tilde{\pi}\to\pi$, we have

$$v_1' \to (1-\delta)g_1^{12}+\delta v_1+(1-\delta)\left[\pi^1(a_1^1,a_2)\omega^1+\pi^{12}(a_1^1,a_2)\omega^{12}+\eta\frac{1-\pi^{11}(a_1^1,a_2)}{1-\pi^{11}(a^1)}\right].$$

Since $\pi^1(a_1^1,a_2)\omega^1+\pi^{12}(a_1^1,a_2)\omega^{12} > g_1^{11}-g_1^{12}+\eta$ and $\eta>0$, the limit is larger than

$$(1-\delta)g_1^{12}+\delta v_1+(1-\delta)(g_1^{11}-g_1^{12}+\eta)=v_1.$$

Similarly, for any small $\eta>0$, a^2 and $v=g^{22}-(\eta,-(\beta_1/\beta_2)\eta)$ can be enforced with continuation payoffs on the ray $\{v'\in h \mid v_1' \leqslant v_1\}$ when $\tilde{\pi}$ is close to π. Therefore, for any compact line segment in the relative interior of ℓ, if $\tilde{\pi}$ is sufficiently close to π, then Q^0 includes the line segment, and the algorithm stops with $Q^h=Q^0$. □

ECONOMICS DEPARTMENTS: HARVARD, FEDERAL RESERVE BANK OF MINNEAPOLIS AND HARVARD