Perfect Public Equilibrium When Players Are Patient

Citation

Fudenberg, Drew, David K. Levine, and Satoru Takahashi. 2007. Perfect public equilibrium when players are patient. *Games and Economic Behavior* 61, no. 1: 27-49.

Published Version

<http://dx.doi.org/10.1016/j.geb.2006.11.002>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3196336>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

HIER

Harvard Institute of Economic Research

Discussion Paper Number 2051

Perfect Public Equilibrium When Players Are Patient

by

Drew Fudenberg, David K. Levine and
Satoru Takahashi

November 2004

Harvard University
Cambridge, Massachusetts

This paper can be downloaded without charge from:
<http://post.economics.harvard.edu/hier/2004papers/2004list.html>

The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=619283>

PERFECT PUBLIC EQUILIBRIUM WHEN PLAYERS ARE PATIENT

DREW FUDENBERG, DAVID K. LEVINE AND SATORU TAKAHASHI
ECONOMICS DEPARTMENTS: HARVARD, FEDERAL RESERVE BANK
OF MINNEAPOLIS AND HARVARD

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: Economic Theory, Game Theory, Repeated Games, Perfect Public Equilibrium, Folk Theorem.

JEL Classification: X15; X16

Date: First Version: 3rd November 2003, This Version: 4th August 2004.

We thank National Science Foundation Grant SES-0314713.

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 note extends the 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 two such restrictions from the literature. The first 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 second application is to the restriction that all payoffs lie on a linear 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 $i = 1, \dots, n$ 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 player i . For each profile $\alpha = (\alpha_1, \dots, \alpha_n)$ of mixed actions $\alpha_i \in \mathcal{A}_i$, we can induce $\pi_y(\alpha)$ and $g_i(\alpha)$.

For $i \in LR = \{1, \dots, L\}$, $L \leq n$, 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 $j \in SR = \{L+1, \dots, n\}$ represent short-run players, each of whom plays only once. Let

$$B: \mathcal{A}_1 \times \dots \times \mathcal{A}_L \rightarrow \mathcal{A}_{L+1} \times \dots \times \mathcal{A}_n$$

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

Let $\mathcal{A}^0$ be a subset of $\text{graph}(B)$. 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 define the sequence $X^0, Q^0, X^1, Q^1, X^2, Q^2, \dots$ 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 $\delta \in (0, 1)$:

$$\begin{aligned} k^m(\alpha, \lambda, \delta) &= \max_{v, w} \lambda \cdot v \quad \text{subject to} \\ \text{(a)} \quad v_i &= (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) w_i(y) \\ &\quad \text{for } i \in LR \text{ and } a_i \in A_i \text{ s.t. } \alpha_i(a_i) > 0, \\ \text{(b)} \quad v_i &\geq (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) w_i(y) \\ &\quad \text{for } i \in LR \text{ and } a_i \in A_i \text{ s.t. } \alpha_i(a_i) = 0, \\ \text{(c)} \quad \lambda \cdot v &\geq \lambda \cdot w(y) \quad \text{for } y \in Y, \\ \text{(d)} \quad w(y) &\in X^m \quad \text{for } y \in Y. \end{aligned}$$

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

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

where $H(\lambda, k) = \{v \in \mathbb{R}^L \mid \lambda \cdot v \leq k\}$. 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 $Q^* = \emptyset$.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 $Q^* = Q^m$. 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} .

Due to the following lemma, we see that the algorithm stops in finite (at most L) steps.

Lemma. *X^m and Q^m are decreasing in m in the set-inclusion ordering. Hence the dimension of X^m is decreasing in m .*

Proof. By construction, we have $Q^m \subseteq X^m$. Since X^{m+1} is the smallest affine space including Q^m and X^m is an affine space including Q^m , it follows that $X^{m+1} \subseteq X^m$. Each λ parallel to X^m has the orthogonal projection $\lambda = \lambda^\parallel + \lambda^\perp$, where $\lambda^\parallel$ is parallel to X^{m+1} and $\lambda^\perp$ is orthogonal to X^{m+1} . $\lambda^\perp \cdot v$ is constant for any $v \in X^{m+1}$, and denoted by k . Since the linear programming problem with constraint X^m is no more restricted than with constraint X^{m+1} , we have $k^m(\alpha, \lambda) \geq k^{m+1}(\alpha, \lambda^\parallel) + k$ for any $\alpha \in \mathcal{A}^0$ with $g_{LR}(\alpha) \in X^{m+1} \subseteq X^m$ and any λ parallel to X^m . Hence $k^m(\lambda) \geq k^{m+1}(\lambda^\parallel) + k$ for any λ parallel to X^m , $H^m(\lambda) \cap X^m \supseteq H^{m+1}(\lambda^\parallel) \cap X^{m+1}$, and $Q^m \supseteq Q^{m+1}$. $\square$

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

¹We 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.

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 $Q^* = Q$.

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.

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

Proof. Let $E^*(\mathcal{A}^0, \delta)$ be the convex hull of $E(\mathcal{A}^0, \delta)$. 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, we may find a half-space $H(\lambda, k)$ and a point $v \in E(\mathcal{A}^0, \delta)$ such that λ is parallel to X^m , $\lambda \cdot v = k > k^m(\lambda)$, and $E^*(\mathcal{A}^0, \delta) \subseteq H(\lambda, k)$. Then (α, v) is enforceable with continuation payoffs in $E(\mathcal{A}^0, \delta) \subseteq H(\lambda, k) \cap X^m$ for some $\alpha \in \mathcal{A}^0$. Since v and the continuation payoffs are in X^m , we have $g_{LR}(\alpha) \in X^m$, contradicting the definition of $k^m(\lambda)$.

Now we show $K \subseteq E(\mathcal{A}^0, \delta)$ for sufficiently large δ . If Q^* is a singleton, then $E(\mathcal{A}^0, \delta) = Q^*$ for any δ . Otherwise, let X^* be the affine hull of Q^* . 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 For any compact set K in the relative interior of Q^* , there exists a smooth, convex, and compact set $W \supseteq K$ in the relative interior of Q^* . We show that $W \subseteq E(\mathcal{A}^0, \delta)$ for sufficiently large δ . Since W is compact, 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 $H = H(\lambda, k)$

³The relative interior of Q^* is the interior of Q^* under the topology induced on the affine hull of Q^* .

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

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^* , it follows that $H \cap X^*$ is a proper subset of $H^*(\lambda) \cap X^*$, where $H^*(\lambda)$ is the maximal half-space in direction λ with continuation payoffs in X^* . Let $\alpha \in \mathcal{A}^0$ be a strategy profile with $g_{LR}(\alpha) \in X^*$ that generates a relative boundary point of $H^*(\lambda) \cap X^*$ using continuation payoffs in $H^*(\lambda) \cap X^*$. Since v is a relative boundary point of W , which is a proper subset of $H^*(\lambda) \cap X^*$, 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 up any $\alpha \in \mathcal{A}^0$ and λ parallel to X^* such that $g_{LR}(\alpha) \in X^*$ generates a relative boundary point of $H^*(\lambda) \cap X^*$ using continuation payoffs in $H^*(\lambda) \cap X^*$. Let $k = \lambda \cdot v$. Again, for some $\delta' < 1$ and $\varepsilon > 0$, (α, v) can be enforced with respect to $H(\lambda, k - \varepsilon) \cap X^*$.

From the incentive constraints, for $\delta'' \geq \delta'$, 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 one there exists $\tilde{\kappa} > 0$ such that the difference between $H \cap X^*$ and 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 . $\square$

4. APPLICATIONS

4.1. Fudenberg and Maskin's Example. Here we consider an example which violates the full dimensionality condition on the set of feasible payoffs. We compute the limit set of equilibrium payoffs by our algorithm.

Let the stage game be a three-player two-action game, and consider its repeated game with perfect monitoring in which all the three players are long-run players. That is, $L = n = 3$, $Y = A = \{0, 1\}^3$, and $\pi_y(a) = 1$ if and only if $y = a$. Stage game payoffs are depicted in Figure 4.1. This game is given by Fudenberg and Maskin [1986].

Let $\mathcal{A}^0 = \mathcal{A}_1 \times \mathcal{A}_2 \times \mathcal{A}_3$ and $X^0 = \mathbb{R}^3$, and solve the first step of our algorithm. By a simple computation, we have $Q^0 = \{(x, x, x) \mid 0 \leq x \leq$

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

1}. Since Q^0 has a lower dimension, 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 exists a player i such that $g_i(a_i, \alpha_{-i}) \geq 1/4$ for some action a_i . Since (v, w) in the linear programming problem satisfies constraints (a) and (b), we have

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

for some player i and some pure action a_i . 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 $-3v_i = (-\mathbf{1}) \cdot v \geq (-\mathbf{1}) \cdot w(y) = -3w_i(y)$ for any outcome y by constraints (c), we have

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

and hence $v_i \geq 1/4$. Therefore, we have $k^1(\alpha, -\mathbf{1}) \leq -3/4$ 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 $Q^1 = H^1(\mathbf{1}) \cap H^1(-\mathbf{1}) \cap X^1 = \{(x, x, x) \mid 1/4 \leq x \leq 1\}$ and $\dim Q^1 = 1 = \dim X^1$, we stop the algorithm and conclude that $Q^* = Q^1$ is the limit set of equilibrium payoffs as $\delta \rightarrow 1$.

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. Symmetry Assumptions.

4.2.1. *Strongly Symmetric Equilibria.* Assume that the static game is symmetric for long-run players, i.e., $A_1 = \dots = A_L$ 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 $\mathcal{A}^s$ of symmetric mixed action profiles for the long-run players in $\text{graph}(B)$, and denote by Q^s the result Q^* of our algorithm under the restriction of $\mathcal{A}^s$. Our Theorem can characterize the limit of $E(\mathcal{A}^s, \delta)$ by Q^s . Similarly to Lemma 3.1 (ii) in FL, we have

$$k^0(\alpha, \lambda) \leq \lambda \cdot g_{LR}(\alpha)$$

for any $\alpha \in \mathcal{A}^s$ in the first step of the algorithm. Let e_i be the vector whose i th component equals 1, and whose other components are all 0. For any $i, j \in LR$ with $i \neq j$ and $\alpha \in \mathcal{A}^s$, we have

$$k^0(\alpha, e_i - e_j) \leq g_i(\alpha) - g_j(\alpha) = 0,$$

and the equality holds when α is a symmetric static equilibrium. So we have $k^0(e_i - e_j) = 0$ for any $i \neq j$. Therefore, Q^0 is a subset of $\bigcap_{i \neq j} H^0(e_i - e_j) = \bigcap_{i \neq j} \{v \in \mathbb{R}^L \mid v_i - v_j \leq 0\} = \{(x, \dots, x) \mid x \in \mathbb{R}\}$, and its dimension is at most 1. Since $\mathcal{A}^s$ contains at least one static equilibrium, we have $Q^0 \neq \emptyset$. If Q^0 is a singleton (which must be a unique symmetric static equilibrium payoff), then we have $Q^s = Q^0$. Otherwise, we move to the second step of our algorithm. Let $X^1 = \{(x, \dots, x) \mid x \in \mathbb{R}\}$ and compute $Q^1 = H^1(\mathbf{1}) \cap H^1(-\mathbf{1}) \cap X^1$, where $\mathbf{1} = (1, \dots, 1)$. Since continuation payoffs have to be symmetric, Q^1 is smaller than or equal to Q^0 . This second step corresponds to Abreu et al's [1986] analysis for large δ . Since Q^1 is a singleton or $\dim Q^1 = 1 = \dim X^1$, we have $Q^s = Q^1$.

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

Corollary 4.1. $Q^s = \lim_{\delta \rightarrow 1} E(\mathcal{A}^s, \delta)$. That is, Q^s is the limit as δ goes to 1 of strongly symmetric equilibrium payoffs with discount factor δ .

4.2.2. *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, Q^0 is a set of payoff profiles symmetric within groups. Then, we can move to the second step, 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 method is applicable and Q^* is relatively easy to compute.

4.3. Exact Achievability of First-Best Outcomes. 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 $\{g_{LR}(\alpha) \mid \alpha \in \text{graph}(B)\}$, 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 $E(\mathcal{A}^h, \delta)$.⁵ At the first step of our algorithm, we have $Q^0 \subset h$, and hence $\dim Q^0 \leq$

⁵Imposing 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

$L - 1$. Therefore, we need to continue the algorithm in at least one more step until we finally obtain Q^* . We denote by Q^h the set Q^* under the restriction of $\mathcal{A}^h$.

Our Theorem implies the following.

Corollary 4.2. *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 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 $(\beta_1, \beta_2, \gamma) = (g_2^1 - g_2^2, g_1^1 - g_1^2, g_1^1 g_2^1 - g_1^2 g_2^2)$ be the line connecting g^1 and g^2 . Note that $\beta_1, \beta_2 > 0$. 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 .⁶ 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 $\ell = \{v \in h \mid g_1^{11} \leq v_1 \leq g_1^{22}\}$. 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

static equilibrium, we can easily modify our argument by analyzing on- and off-schedule deviations separately. See Athey and Bagwell [2001]. Note also that allowing off-path play not in $\mathcal{A}^h$ destroys the renegotiation-proofness property of the equilibria.

⁶ $\mathcal{A}^h = \{a^1, a^2\}$ is satisfied for generic stage games where no three pure action payoff profiles lie on a line. However, this assumption excludes all duopoly games where both firms have the same constant marginal cost. If $\mathcal{A}^h \supsetneq \{a^1, a^2\}$, then the condition of $g_1^{11} \geq g_1^{22}$ in Part 1 of Theorem 4.3 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.

that for

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

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

1. $\pi^{ii}(a^i) < 1$ and $\pi^i(a^i)\omega^i + \pi^{ij}(a^i)\omega^{ij} = g_i^{ii} - g_i^i$,
2. $\pi^{ii}(a_i, a_j^i) = 1$ for any $a_i \neq a_i^i$, and
3. $\pi^i(a_j, a_i^i)\omega^i + \pi^{ij}(a_j, a_i^i)\omega^{ij} > g_i^{ii} - g_i^{ij}$ for any $a_j \neq a_j^i$.

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 $Y^i = \{a^i\}$, $Y^{ii} = \{(a_i, a_j^i) \mid a_i \neq a_i^i\}$, $Y^{ij} = \{a \mid a_j \neq a_j^i\}$, $\omega^i = g_i^{ii} - g_i^i$, and $\omega^{ij} = \max(g_i^{ii} - g_i^{ij} + 1, 0)$.

Theorem 4.3. *Suppose that neither a^1 nor a^2 is a static Nash equilibrium. There are two cases:*

1. $g_1^{11} \geq g_1^{22}$. We have $Q^h = \emptyset$ under any signal structure. Hence there is no $\mathcal{A}^h$ -perfect public equilibrium for any δ .
2. $g_1^{11} < g_1^{22}$. 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 $\max_{y,a} |\pi_y(a) - \tilde{\pi}_y(a)| < \varepsilon$. Then there exists an $\mathcal{A}^h$ -perfect public equilibrium for sufficiently large δ .

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. If Q^0 is the empty set or a singleton, then $Q^h = \emptyset$ since neither a^1 nor a^2 is a static Nash equilibrium.

Otherwise, Q^0 is a line segment in h . Let Q^1 be the result of the second 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 $\lambda^1 \cdot v \geq \lambda^1 \cdot w(y)$, we have $v_2 \leq w_2(y)$. Then, by constraint (b), we have

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

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

$\lambda^1 \cdot g^1$, we have

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

Similarly, we have $k^1(\lambda^2) \leq \max(\lambda^2 \cdot g^{11}, \lambda^2 \cdot g^2)$. Therefore,

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

Since $g_1^{11} \geq g_1^{22}$, $g_1^{11} \geq g_1^1$, $g_1^2 \geq g_1^{22}$, and $g_1^2 > g_1^1$, we have $\min(g_1^{11}, g_1^2) \geq \max(g_1^{22}, g_1^1)$. 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^1 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 $\{v' \in h \mid v'_1 \geq v_1\}$ 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 $\omega^1, \omega^{12} \geq 0$ and $\zeta \rightarrow \eta/(1 - \pi^{11}(a^1)) > 0$ as $\tilde{\pi} \rightarrow \pi$, all $w(y)$ satisfy $w_1(y) \geq v_1$ 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

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

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 $\tilde{\pi} \rightarrow \pi$. 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 $\beta_1, \beta_2 > 0$, this deviation is unprofitable for player 2 if $v_1' \geq v_1$ for $\tilde{\pi}$ close to π . Taking $\tilde{\pi} \rightarrow \pi$, we have

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

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 $\{v' \in h \mid v_1' \leq v_1\}$ 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^1 includes the line segment, and the algorithm stops with $Q^h = Q^1$. $\square$

REFERENCES

- [1] Abreu, D. [1986] “Extremal equilibria of oligopolistic supergames,” *Journal of Economic Theory* 39, 191-228.
- [2] Abreu, D., D. Pearce, and E. Stacchetti [1986] “Optimal cartel equilibria with imperfect monitoring,” *Journal of Economic Theory* 39, 251-269.
- [3] Abreu, D., D. Pearce, and E. Stacchetti [1990] “Toward a theory of discounted repeated games with imperfect monitoring,” *Econometrica* 58, 1041-1063.
- [4] Athey, S. and K. Bagwell [2001] “Optimal collusion with private information,” *RAND Journal of Economics* 32, 428-465.
- [5] Dellarocas, C. [2003] “Efficiency and robustness of eBay-like online reputation mechanisms in environments with moral hazard,” working paper, MIT.
- [6] Ely, J., D. Fudenberg, and D. K. Levine [2003] “When is reputation bad?,” mimeo.
- [7] Farrell, J. and E. Maskin [1989] “Renegotiation in repeated games,” *Games and Economic Behavior* 1, 327-360.
- [8] Fudenberg, D. and D. K. Levine [1994] “Efficiency and observability with long-run and short-run players,” *Journal of Economic Theory* 62, 103-135.
- [9] Fudenberg, D., D. K. Levine, and E. Maskin [1994] “The folk theorem in repeated games with imperfect public information,” *Econometrica* 62, 997-1039.
- [10] Fudenberg, D. and E. Maskin [1986] “The folk theorem in repeated games with discounting or with incomplete information,” *Econometrica* 54, 533-554.
- [11] Kandori M. and H. Matsushima [1998] “Private observation, communication and collusion,” *Econometrica* 66, 627-652.
- [12] Wen, Q. [1994] “The ‘folk theorem’ for repeated games with complete information,” *Econometrica* 62, 949-954.