

Perfect public equilibrium when players are patient [☆]

Drew Fudenberg ^a, David K. Levine ^b, Satoru Takahashi ^{a,*}

^a Department of Economics, Harvard University, Cambridge, MA 02138, USA

^b Department of Economics, Washington University in St. Louis, MO 63130, USA

Received 29 September 2005

Available online 12 February 2007

Abstract

We provide a characterization of the limit set of perfect public equilibrium payoffs of repeated games with imperfect public monitoring as the discount factor goes to one. Our result covers general stage games including those that fail a “full-dimensionality” condition that had been imposed in past work. It also provides a characterization of the limit set when the strategies are restricted in a way that endogenously makes the full-dimensionality condition fail, as in the strongly symmetric equilibrium studied by Abreu [Abreu, D., 1986. Extremal equilibria of oligopolistic supergames. *J. Econ. Theory* 39, 191–228] and Abreu et al. [Abreu, D., Pearce, D., Stacchetti, E., 1986. Optimal cartel equilibria with imperfect monitoring. *J. Econ. Theory* 39, 251–269]. Finally, we use our characterization to give 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.

© 2006 Elsevier Inc. All rights reserved.

JEL classification: C72; C73

Keywords: Repeated games; Perfect public equilibrium; Folk theorem

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

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

* Corresponding author.

E-mail address: stakahas@fas.harvard.edu (S. Takahashi).

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 that payoffs satisfy a “non-equivalent utilities” condition, and of Wen (1994), who assumed that mixed strategies are observed. The second application is to the strongly symmetric equilibria of symmetric games, studied by Abreu (1986) and Abreu et al. (1986). This equilibrium concept requires that all players take the same action, 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 face of the Pareto frontier, which we use to derive a sufficient condition for the exact achievability of first-best outcomes. Equilibria 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.

A recent paper by Tomala (2005) uses our result to analyze repeated games with communication. Here the “full-dimensionality” condition fails because the mediator’s payoff is always zero.

To incorporate exogenous restrictions on the strategies such as symmetry or efficiency, we use the concept of \mathcal{A}^0 -perfect public equilibria: these are perfect public equilibria in which players choose action profiles from \mathcal{A}^0 after all public histories. Our method yields a characterization of the limit of \mathcal{A}^0 -perfect public equilibrium payoffs for an arbitrary specification of the set \mathcal{A}^0 .

2. Model

We consider a repeated game with imperfect public monitoring played by long-run and short-run players.¹ Our notation follows FL. In the stage game, players $i = 1, \dots, n$ simultaneously choose pure actions a_i from finite sets A_i . An action profile $a \in A := \prod_{i=1}^n A_i$ 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 $\alpha \in \mathcal{A} := \prod_{i=1}^n \mathcal{A}_i$, we can define $\pi_y(\alpha)$ by $\sum_{a \in A} \pi_y(a)\alpha(a)$ and $g_i(\alpha)$ by $\sum_{a \in A} g_i(a)\alpha(a)$, where $\alpha(a) = \prod_{i=1}^n \alpha_i(a_i)$; $\pi_y(a_i, \alpha_{-i})$ and $g_i(a_j, \alpha_{-j})$ are defined similarly.

In each period $t = 1, 2, \dots$ of the repeated game, the stage game is played, with each player \tilde{i} choosing action $a_i(t)$. The resulting action profile $a(t)$ induces the distribution $\pi_y(a(t))$ on public outcomes y , and the realized public outcome $y(t)$ is observed at the end of the period. Thus at the end of period \tilde{t} , player \tilde{i} has observed a public history $h(t) = (y(1), \dots, y(t)) \in Y^t$. Because players know their own actions, each player has also observed a private history

¹ The examples we present have only long-run players. Because allowing for short-run players does not complicate the proofs, we have chosen to present the more general results.

$h_i(t) = (a_i(1), \dots, a_i(t)) \in A_i^t$. Thus a strategy for player \tilde{i} is a sequence of maps $\sigma_i(t) : Y^{t-1} \times A_i^{t-1} \rightarrow A_i$. A strategy for player i is *public* if it does not depend on private histories: $\sigma_i(t)(h(t-1), h_i(t-1)) = \sigma_i(t)(h(t-1), h'_i(t-1))$ for any $t, h(t-1), h_i(t-1)$, and $h'_i(t-1)$. With a slight abuse of notation, we denote it by $\sigma_i(t)(h(t-1))$. If all players use public strategies, then players' beliefs about other players' private histories are irrelevant.

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, $\text{graph}(B) := \{(\alpha_{LR}, \alpha_{SR}) \in \mathcal{A} \mid \alpha_{SR} \in B(\alpha_{LR})\}$ is the set of mixed action profiles α such that $g_j(\alpha) \geq g_j(a_j, \alpha_{-j})$ for any $j \in SR$ and any $a_j \in A_j$.²

A strategy profile $\{\sigma(t)\}$ is a *perfect public equilibrium (PPE)* if every player i 's strategy $\{\sigma_i(t)\}$ is public, and, after any public history $h(t-1)$, every long-run player i maximizes his average discounted payoff given the other players' strategy profile $\{\sigma_{-i}(t)\}$, and every short-run player j maximizes his one-shot payoff given the other players' action profile $\sigma_{-j}(t)(h(t-1))$, that is, $\sigma(t)(h(t-1)) \in \text{graph}(B)$.

Let \mathcal{A}^0 be a subset of $\text{graph}(B)$. We focus on \mathcal{A}^0 -*perfect public equilibria (\mathcal{A}^0 -PPE)*: perfect public equilibria in which players choose action profiles from \mathcal{A}^0 after all public histories, that is, $\sigma(t)(h(t-1)) \in \mathcal{A}^0$ for any $h(t-1)$. 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 . Let $E(\mathcal{A}^0, \delta)$ be the set of average present values for the long-run players in \mathcal{A}^0 -perfect public equilibria when all long-run players use a common discount factor δ . We will characterize the limit of $E(\mathcal{A}^0, \delta)$ as $\delta \rightarrow 1$ without the "full-dimensionality" condition.³

As an example, let $\mathcal{A}^P := \{\alpha \in \text{graph}(B) \mid \alpha(a) = 1 \text{ for some } a \in A\}$, the set of all pure action profiles in $\text{graph}(B)$. Then \mathcal{A}^P -PPEs mean pure-strategy PPEs. As another example, let $\mathcal{A}^s := \{\alpha \in \text{graph}(B) \mid \alpha_1 = \dots = \alpha_L\}$. Then the \mathcal{A}^s -PPEs are the PPEs in which the long-run players use strongly symmetric strategies.

We also consider repeated games with public randomization. In these games, an additional signal $\omega(t)$ is publicly observed at the beginning of each period t ; we suppose that these signals are independently and identically distributed according to the uniform distribution on $[0, 1]$. Thus in these games the public history is given by $h^*(t) = (\omega(1), y(1), \dots, \omega(t), y(t), \omega(t+1))$, and a strategy for player i is a sequence of measurable maps $\sigma_i^*(t) : Y^{t-1} \times A_i^{t-1} \times [0, 1]^t \rightarrow A_i$. We define \mathcal{A}^0 -PPE in a repeated game with public randomization by replacing $h(t)$ with $h^*(t)$ in the previous definition. Let $E^*(\mathcal{A}^0, \delta)$ be the set of \mathcal{A}^0 -PPE payoff profiles when public randomization devices are available. $E^*(\mathcal{A}^0, \delta)$ is a bounded convex set that contains $E(\mathcal{A}^0, \delta)$.

² Note that when there are no short-run players, $\text{graph}(B)$ is taken to be the set of all mixed action profiles \mathcal{A} .

³ Note that for some choices of \mathcal{A}^0 and δ , $E(\mathcal{A}^0, \delta)$ may be empty. This does not matter for our analysis: for example, our results cover the trivial case where $E(\mathcal{A}^0, \delta)$ is empty for all δ .

3. The algorithm

We fix \mathcal{A}^0 throughout this section. For each affine subspace X of \mathbb{R}^L with $\dim X \geq 1$, we consider a linear programming problem for given $\alpha \in \mathcal{A}^0$ with $g_{LR}(\alpha) \in X$, $\lambda \in \mathbb{R}^L \setminus \{\mathbf{0}\}$ parallel to X , and $\delta \in (0, 1)$:

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

If there is no (v, w) that satisfies constraints (a)–(d), then we set $k(\alpha, \lambda, \delta, X) := -\infty$. Note that $k(\alpha, \lambda, \delta, \mathbb{R}^L)$ corresponds to $k^*(\alpha, \lambda, \delta)$ in FL.

As is standard in this literature, payoff profile v is the target that will be supported by some equilibrium, and the function w gives continuation payoff $w(y)$ 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.

This linear programming problem differs from FL’s only in constraints (d). Constraints (d) require that all of the continuation payoffs are included in X .

To help interpret and motivate the problem, recall that the literature says that w enforces (α, v) under discount factor δ if

$$v_i = (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i, \alpha_{-i})w_i(y)$$

for $i \in LR$ and $a_i \in A_i$ such that $\alpha_i(a_i) > 0$, and

$$v_i \geq (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i, \alpha_{-i})w_i(y)$$

for $i \in LR$ and $a_i \in A_i$ such that $\alpha_i(a_i) = 0$. The pair (α, v) is *enforceable with respect to W under discount factor δ* if there exists $w : Y \rightarrow W$ that enforces (α, v) . Further, v is *generated by \mathcal{A}^0 and W under discount factor δ* if there exists $\alpha \in \mathcal{A}^0$ such that (α, v) is enforceable with respect to W under discount factor δ . Let $P(\mathcal{A}^0, \delta, W)$ be the set of payoff profiles generated by \mathcal{A}^0 and W under discount factor δ . By applying the principle of optimality of dynamic programming, Abreu et al. (1990) showed that if W is bounded and $W \subseteq P(\mathcal{A}^0, \delta, W)$, then $W \subseteq E(\mathcal{A}^0, \delta)$. Let $H(\lambda, k) := \{v \in \mathbb{R}^L \mid \lambda \cdot v \leq k\}$; then constraints (a)–(d) say that (α, v) is enforceable with respect to the (unbounded) set $H(\lambda, \lambda \cdot v) \cap X$.

Lemma 3.1.

(1) If (α, v) is enforced by w under discount factor δ , then it is enforced by

$$w'(y) = \frac{\delta' - \delta}{\delta'(1 - \delta)} v + \frac{\delta(1 - \delta')}{\delta'(1 - \delta)} w(y)$$

under discount factor δ' .

(2) If $g_{LR}(\alpha) \in X$, then $k(\alpha, \lambda, \delta, X)$ is independent of δ and positively homogeneous of degree 1 in λ .

(3) If W is convex and $\delta' \geq \delta$, then $P(\mathcal{A}^0, \delta, W) \cap W \subseteq P(\mathcal{A}^0, \delta', W)$.⁴

Proof. [Proof of (1)] Note that the coefficients on v and $w(y)$ add up to one. Moreover, for all $i \in LR$, a_i , and α_{-i} ,

$$\begin{aligned} & (1 - \delta')g_i(a_i, \alpha_{-i}) + \delta' \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) w'_i(y) \\ &= \frac{\delta' - \delta}{1 - \delta} v + \frac{1 - \delta'}{1 - \delta} \left[(1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i, \alpha_{-i}) w_i(y) \right], \end{aligned}$$

where the term in the square brackets is the payoff to playing a_i when the opponents play α_{-i} , the continuation payoffs are given by w , and the discount factor is δ . So the fact that w enforces (α, v) when the discount factor is δ implies that w' enforces (α, v) when the discount factor is δ' .

[Proof of (2)] Since we have $g_{LR}(\alpha) \in X$ and $w(y) \in X$ by (d), it follows from (a) that $v \in X$. Since X is affine, it follows then that $w'(y) \in X$ for all y . And using $\lambda \cdot v \geq \lambda \cdot w(y)$, we see that also $\lambda \cdot v \geq \lambda \cdot w'(y)$ so that these new continuation payoffs also satisfy (c). Thus given feasible $w(y)$ for any δ , we can construct feasible $w'(y)$ for any δ' that give the same value of v , so the solution of the linear programming problem must be independent of δ .

In the linear programming problem, the objective function is positively homogeneous of degree 1 in λ and the constraints are not affected by any positive scalar multiplication of λ . So the solution is positively homogeneous of degree 1 in λ .

[Proof of (3)] Pick any $v \in P(\mathcal{A}^0, \delta, W) \cap W$. There exist $\alpha \in \mathcal{A}^0$ and w such that (α, v) is enforced by w under discount factor δ and $w(y) \in W$ for all y . Since $\delta' \geq \delta$, $w'(y)$ defined in part 1 is a convex combination of v and $w(y)$, which belongs to W because W is convex. So we have $v \in P(\mathcal{A}^0, \delta', W)$. \square

Since $k(\alpha, \lambda, \delta, X)$ is independent of δ , we will therefore denote it by $k(\alpha, \lambda, X)$.

The algorithm characterizes the limit set of equilibrium payoffs using extremal halfspaces. We define these by setting

$$\begin{aligned} k(\mathcal{A}^0, \lambda, X) &:= \sup_{\alpha \in \mathcal{A}^0, g_{LR}(\alpha) \in X} k(\alpha, \lambda, X), \\ H(\mathcal{A}^0, \lambda, X) &:= H(\lambda, k(\mathcal{A}^0, \lambda, X)), \quad \text{and} \end{aligned}$$

⁴ The first two parts of Lemma 3.1 was originally proven as FL Lemma 3.1. 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 asserted that the condition is necessary as well. The condition is only necessary under the additional assumption that all outcomes have positive probability under α .

$$Q(\mathcal{A}^0, X) := \bigcap_{\lambda \in \mathbb{R}^L \setminus \{0\} : \text{parallel to } X} H(\mathcal{A}^0, \lambda, X) \cap X.$$

If $k(\mathcal{A}^0, \lambda, X) = -\infty$, then we set $H(\mathcal{A}^0, \lambda, X) = \emptyset$. $k(\mathcal{A}^0, \lambda, X)$ is the supremum of $\lambda \cdot v$ such that (α, v) is enforceable with respect to $H(\lambda, \lambda \cdot v) \cap X$ for some $\alpha \in \mathcal{A}^0$ with $g_{LR}(\alpha) \in X$. Note that, since $g_{LR}(\alpha) \in X$, by Lemma 3.1, whether (α, v) is enforceable with respect to $H(\lambda, \lambda \cdot v) \cap X$ does not depend on δ .

Lemma 3.2. *If $E^*(\mathcal{A}^0, \delta) \subseteq X$, then $E^*(\mathcal{A}^0, \delta) \subseteq Q(\mathcal{A}^0, X)$.*

Proof. The proof is the same as in FL, except that we use public randomizations and $E^*(\mathcal{A}^0, \delta)$ may not be closed.

Suppose that $E^*(\mathcal{A}^0, \delta) \not\subseteq Q(\mathcal{A}^0, X)$. Then there exists a point $v \in E^*(\mathcal{A}^0, \delta) \setminus Q(\mathcal{A}^0, X)$. By the definition of $Q(\mathcal{A}^0, X)$, we can find a vector of length one, λ , parallel to X that satisfies $\lambda \cdot v > k(\mathcal{A}^0, \lambda, X)$. Let ε be a positive number such that $\varepsilon < \lambda \cdot v - k(\mathcal{A}^0, \lambda, X)$.

Given such v, λ , and ε , since $E^*(\mathcal{A}^0, \delta)$ is bounded, we may find a point $v' \in E^*(\mathcal{A}^0, \delta)$ with $k' := \lambda \cdot v'$ that satisfies

$$k' + (1 - \delta)\varepsilon \geq \lambda \cdot x \quad \text{for any } x \in E^*(\mathcal{A}^0, \delta).$$

In particular, by plugging $x = v$, we have $k' + (1 - \delta)\varepsilon \geq \lambda \cdot v > k(\mathcal{A}^0, \lambda, X) + \varepsilon$. Thus we have $k' - \delta\varepsilon > k(\mathcal{A}^0, \lambda, X)$.

Since $v' \in E^*(\mathcal{A}^0, \delta)$, v' is written as

$$v' = \int_0^1 v'(\omega) d\omega,$$

where $v'(\omega)$ is the payoff profile after ω is realized at the beginning of the first period. For every $\omega \in [0, 1]$, let $\alpha(\omega) \in \mathcal{A}^0$ be the current action profile and $w'(y, \omega) \in E^*(\mathcal{A}^0, \delta)$ be the continuation payoff profile if the current outcome is y . Note that $w'(\cdot, \omega)$ enforces $(\alpha(\omega), v'(\omega))$ under discount factor δ . Since $v'(\omega), w'(y, \omega) \in E^*(\mathcal{A}^0, \delta) \subseteq X$, we have $g_{LR}(\alpha(\omega)) \in X$.

Pick $\omega \in [0, 1]$ such that $\lambda \cdot v'(\omega) \geq \lambda \cdot v' = k'$. Shifting payoff profiles independently of y , we define

$$v'' := v'(\omega) - \delta\varepsilon\lambda$$

and

$$w''(y) := w'(y, \omega) - \varepsilon\lambda.$$

Since $w'(\cdot, \omega)$ enforces $(\alpha(\omega), v'(\omega))$ under discount factor δ , $w''(\cdot)$ enforces $(\alpha(\omega), v'')$ under discount factor δ . Moreover, we have

$$\lambda \cdot v'' = \lambda \cdot v'(\omega) - \delta\varepsilon \geq k' - \delta\varepsilon,$$

$$\lambda \cdot w''(y) = \lambda \cdot w'(y, \omega) - \varepsilon \leq k' + (1 - \delta)\varepsilon - \varepsilon = k' - \delta\varepsilon.$$

Since $w''(y) \in H(\lambda, \lambda \cdot v'') \cap X$, $(\alpha(\omega), v'')$ is enforceable with respect to $H(\lambda, \lambda \cdot v'') \cap X$ under discount factor δ . This contradicts the definition of $k(\mathcal{A}^0, \lambda, X)$ because $\lambda \cdot v'' \geq k' - \delta\varepsilon > k(\mathcal{A}^0, \lambda, X)$. \square

For $v \in \mathbb{R}^L$, let $|v|$ denote the l^2 -norm of v , $|v| = \sqrt{\sum_{i \in LR} v_i^2}$. This norm induces the usual topology on \mathbb{R}^L . 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 . *Relative boundary* and *relatively open neighborhood* are defined similarly; a *relative interior (boundary) point* of a set is a point in its relative interior (boundary).

For any affine subspace X of \mathbb{R}^L with $\dim X \geq 1$, we say that a convex and compact subset W of X is *smooth in X* if for any relative boundary point v of W , there exists a unique vector of length one, λ , parallel to X such that $W \subseteq H(\lambda, \lambda \cdot v)$.⁵

Lemma 3.3. *Suppose that $\dim Q(\mathcal{A}^0, X) = \dim X \geq 1$.*

- (1) *For any compact set K in the relative interior of $Q(\mathcal{A}^0, X)$, there exists $\bar{\delta} < 1$ such that $K \subseteq E(\mathcal{A}^0, \delta)$ for any $\delta > \bar{\delta}$.*
- (2) *If $E^*(\mathcal{A}^0, \delta) \subseteq X$ for all δ , then*

$$\lim_{\delta \rightarrow 1} E(\mathcal{A}^0, \delta) = \lim_{\delta \rightarrow 1} E^*(\mathcal{A}^0, \delta) = Q(\mathcal{A}^0, X)$$

in the Hausdorff metric.

Remark. The role of the full-dimensionality assumption is to allow us to find W smooth in X that approximates Q from inside. Lemma A.2 in the appendix shows that for any set W smooth in X and any $\kappa, \kappa' > 0$, there exists $\varepsilon \in (0, 1)$ such that every $v' \in X$ that satisfies $\lambda \cdot v' \leq \lambda \cdot v - \kappa \varepsilon$ and $|v' - v| \leq \kappa' \varepsilon$ is a relative interior point of W . We use this in the proof at the end of the discussion of case 1 to conclude that when δ is close enough to 1, certain actions and payoffs can be enforced with continuation payoffs in the relative interior of W .

Proof. [Proof of (1)] The proof differs from FL's in two respects. First, we use the *relative topology* induced on X instead of the standard topology on \mathbb{R}^L . Second, because we do not assume the existence of static equilibria in \mathcal{A}^0 , we use a different technique to generate relative interior points.

Claim 1: Smooth Approximation. To begin the proof, we claim that there is a convex and compact set $W \supseteq K$ in the relative interior of $Q(\mathcal{A}^0, X)$ that is smooth in X . To see this, let

$$W := \left\{ v \in X \mid \min_{x \in \text{co}(K)} |v - x| \leq \varepsilon \right\}$$

with $\varepsilon > 0$, where $\text{co}(K)$ is the convex hull of K . Because K is a compact set, Lemma A.1 in the appendix shows that W is convex, compact, and smooth in X . Moreover, because K is a compact set in the relative interior of $Q(\mathcal{A}^0, X)$ and $\dim Q(\mathcal{A}^0, X) = \dim X$, W is in the relative interior of $Q(\mathcal{A}^0, X)$ for sufficiently small ε . For such W , we will show that $W \subseteq E(\mathcal{A}^0, \delta)$ for sufficiently large δ .⁶

Claim 2: Local Generation. We claim next that it is enough to show that the set W is “locally generated” by \mathcal{A}^0 and W , meaning that for each $v \in W$, there exist $\delta_v < 1$ and a relatively open neighborhood U_v of v with $U_v \subseteq P(\mathcal{A}^0, \delta_v, W)$. When this condition is satisfied, since $\{U_v\}_{v \in W}$

⁵ This definition is equivalent to the set W having a differentiable relative boundary.

⁶ Earlier work used an approximation W that has a C^2 boundary. With a C^2 boundary, we can get a stronger approximation whose error term is bounded by a quadratic function. However, as our proof shows, this stronger bound is not needed.

is an open cover of W and W is compact, we can choose a finite subcover $\{U_{v^m}\}_{m=1}^{m^*}$ of W . Let $\bar{\delta} := \max_m \delta_{v^m} < 1$. Since W is convex, it follows from Lemma 3.1 that for any $\delta > \bar{\delta}$, we have

$$U_{v^m} \cap W \subseteq P(\mathcal{A}^0, \delta_{v^m}, W) \cap W \subseteq P(\mathcal{A}^0, \delta, W)$$

for every m , which implies $W \subseteq P(\mathcal{A}^0, \delta, W)$. Since W is bounded, we have $W \subseteq E(\mathcal{A}^0, \delta)$. This proves the second claim: it is enough to show that W is locally generated by \mathcal{A}^0 and W .

To show that W is locally generated by \mathcal{A}^0 and W , we first show that for each $v \in W$, there exists δ_v such that $v \in P(\mathcal{A}^0, \delta_v, W)$, and then extend this from the points v to neighborhoods U_v .

Case 1: Relative Boundary. For any relative boundary point v of W , let λ be the unique vector of length one parallel to X and normal to W at v . Let $k := \lambda \cdot v$. Then we have $W \subseteq H(\lambda, k) \cap X$. Fix an arbitrary $\delta < 1$. Since W is in the relative interior of $Q(\mathcal{A}^0, X)$, there exist an action profile $\alpha \in \mathcal{A}^0$ with $g_{LR}(\alpha) \in X$ and a point $v' \in X$ with $k' := \lambda \cdot v' > k$ that are enforced by some $w'(y) \in H(\lambda, k') \cap X$ under discount factor δ .

Translating the continuation payoffs, we can find $\varepsilon > 0$ and $w(y) \in H(\lambda, k - \varepsilon) \cap X$ such that $w(y)$ enforce (α, v) under discount factor δ .

For any $\delta' < 1$, by Lemma 3.1, we may find $w(y, \delta')$ that enforce (α, v) under discount factor δ' :

$$w(y, \delta') = \frac{\delta' - \delta}{\delta'(1 - \delta)} v + \frac{\delta(1 - \delta')}{\delta'(1 - \delta)} w(y).$$

We have

$$w(y, \delta') \in X$$

and

$$\lambda \cdot w(y, \delta') \leq k - \kappa(1 - \delta')$$

with $\kappa := [\delta/(1 - \delta)]\varepsilon$. We also have

$$|w(y, \delta') - v| \leq \kappa'(1 - \delta')$$

with $\kappa' := [\delta/(1 - \delta)] \max_y |w(y) - v|$. Since W is smooth in X , by Lemma A.2 in the appendix, there exists $\delta_v < 1$ such that $w(y, \delta_v)$ are in the relative interior of W .

Case 2: Relative Interior. For any relative interior point v of W , fix an arbitrary $\delta < 1$. Since $Q(\mathcal{A}^0, X) \neq \emptyset$, v is generated by \mathcal{A}^0 and X under discount factor δ . Let α be a mixed action profile in \mathcal{A}^0 with $g_{LR}(\alpha) \in X$ such that (α, v) is enforced by some $w(y) \in X$ under discount factor δ .

Similarly to the relative boundary case, for any δ' , by Lemma 3.1, we define $w(y, \delta')$ as above so that $w(y, \delta')$ enforce (α, v) under discount factor δ' . Since

$$w(y, \delta') \in X$$

and

$$|w(y, \delta') - v| \leq \kappa'(1 - \delta')$$

with $\kappa' := [\delta/(1 - \delta)] \max_y |w(y) - v|$, there exists $\delta_v < 1$ such that $w(y, \delta_v)$ are in the relative interior of W .

In both cases, since $w(y, \delta_v)$ are in the relative interior, they may be translated by a small constant independent of y , generating incentive compatible payoffs in a relatively open neighborhood U_v of v .

[Proof of (2)] Since $E^*(\mathcal{A}^0, \delta) \subseteq X$, by Lemma 3.2, we have $E(\mathcal{A}^0, \delta) \subseteq E^*(\mathcal{A}^0, \delta) \subseteq Q(\mathcal{A}^0, X)$. Since we showed in (1) that any relative interior of $Q(\mathcal{A}^0, X)$ is included in $E(\mathcal{A}^0, \delta)$ for sufficiently large δ , we have

$$\lim_{\delta \rightarrow 1} E(\mathcal{A}^0, \delta) = \lim_{\delta \rightarrow 1} E^*(\mathcal{A}^0, \delta) = Q(\mathcal{A}^0, X)$$

in the Hausdorff metric. \square

Remark. It is easy to extend Lemmas 3.2 and 3.3 to games with infinitely many pure actions.⁷ However, allowing infinitely many signals would involve measure-theoretic complications that are beyond the scope of this paper.

Now we provide an algorithm to find X that satisfies the assumptions in Lemma 3.3. Let $X^0 := \mathbb{R}^L$. For each X^m , we compute the linear programming problem, and obtain $Q^m := Q(\mathcal{A}^0, X^m)$. 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 $Q^*(\mathcal{A}^0) := \emptyset$.⁸ 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 the 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 one, so the algorithm stops in a finite number of steps.

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, and Q^0 is equal to what FL called Q . If $\dim Q = L$, then the algorithm stops at the first step, and we have $Q^*(\text{graph}(B)) = Q$. However, if $\dim Q < L$, then Lemma 3.3 does not apply to Q because we cannot find a subset W of Q that is smooth in \mathbb{R}^L and the last step in the proof of case 1 fails: even if δ is close to 1, the continuation payoffs constructed in the proof need not lie in W . This is why we need to iterate the algorithm.

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

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 subset K of the relative interior of $Q^*(\mathcal{A}^0)$, there exists $\bar{\delta} < 1$ such that $K \subseteq E(\mathcal{A}^0, \delta)$ for any $\delta > \bar{\delta}$. Hence

$$\lim_{\delta \rightarrow 1} E(\mathcal{A}^0, \delta) = \lim_{\delta \rightarrow 1} E^*(\mathcal{A}^0, \delta) = Q^*(\mathcal{A}^0)$$

in the Hausdorff metric.

Proof. Since $E^*(\mathcal{A}^0, \delta) \subseteq \mathbb{R}^L = X^0$, we have $E^*(\mathcal{A}^0, \delta) \subseteq Q(\mathcal{A}^0, X^0) = Q^0$ by Lemma 3.2. Applying Lemma 3.2 inductively, we have $E^*(\mathcal{A}^0, \delta) \subseteq Q^{m^*}$, where m^* is the number of steps we need until our algorithm stops.

⁷ The proofs carry 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 $k(\alpha, \lambda, \delta, X)$, 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 .

⁸ Every static equilibrium in \mathcal{A}^0 is contained in Q^m for each m , and hence $Q^m = \emptyset$ is possible only if there is no static equilibrium in \mathcal{A}^0 . The converse is not true. For example, in the case of $\mathcal{A}^0 = \mathcal{A}^P$, the repeated game may have a pure-strategy equilibrium for large δ even if the stage game has no pure-strategy equilibrium. See Section 4.4 for another example.

If 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, since we set $Q^*(\mathcal{A}^0) = Q^{m^*}$, we have $E^*(\mathcal{A}^0, \delta) \subseteq Q^*(\mathcal{A}^0)$. Thus we have shown the first part of the claim.

The second part is shown as follows. If Q^{m^*} is a singleton whose element corresponds to a static equilibrium in \mathcal{A}^0 , then we have $E(\mathcal{A}^0, \delta) = Q^{m^*}$ for any δ , and hence $\lim_{\delta \rightarrow 1} E(\mathcal{A}^0, \delta) = Q^{m^*} = Q^*(\mathcal{A}^0)$.

If Q^{m^*} is neither the empty set nor a singleton, then we have $\dim Q^{m^*} = \dim X^{m^*} \geq 1$, and we can apply Lemma 3.3. \square

Remark. Our theorem shows that $\lim_{\delta \rightarrow 1} E(\mathcal{A}^0, \delta) = \lim_{\delta \rightarrow 1} E^*(\mathcal{A}^0, \delta)$, i.e., 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 $1 \leq \dim Q^m < \dim X^m$, then we can move to the next step with any affine subspace X^{m+1} of X^m that contains Q^m .

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 Fig. 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 $Q^0 = \{(x, x, x) \mid 0 \leq x \leq 1\}$. 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 scalar multiplication), $\mathbf{1} := (1, 1, 1)$ 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 $g_i(a_i, \alpha_{-i}) \geq 1/4$. Since (v, w) in the linear programming problem satisfies constraints (a) and (b), we have

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

Since $g(\alpha) \in X^1$ 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

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

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

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

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

Fig. 4.1. A three-player game in Fudenberg and Maskin (1986).

Since

$$Q^1 = H(A, \mathbf{1}, X^1) \cap H(A, -\mathbf{1}, X^1) \cap X^1 = \{(x, x, x) \mid 1/4 \leq x \leq 1\}$$

and

$$\dim Q^1 = 1 = \dim X^1,$$

we stop the algorithm and conclude that $Q^*(A) = Q^1$ is the limit set of subgame-perfect equilibrium payoffs as $\delta \rightarrow 1$.

The same result was obtained by Fudenberg and Maskin (1986) and Wen (1994). Fudenberg and Maskin determined the limit set by a direct computation in this specific game, whereas Wen used effective minmax values. Wen’s method is applicable to all repeated games with perfect monitoring without the full dimensionality condition. Our algorithm is even more general, as we allow 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 $L = n$. We assume that $\mathcal{A}^0 \supseteq \mathcal{A}^P = \{\alpha \in \mathcal{A} \mid \alpha(a) = 1 \text{ for some } a \in A\}$, i.e., \mathcal{A}^0 contains all pure action profiles.

We also 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 $g_j(a) = c + dg_i(a)$ 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 $-g_i$.

Recall the standard definition of the minmax payoff:

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

The previous example shows that the lowest equilibrium payoffs can be bounded away from these values; intuitively this is because it can be difficult to induce players with identical payoffs to punish each other when this punishment is costly. We will show that the correct lower bound on equilibrium payoffs is what we call the “effective minmax values.”

Player i ’s effective minmax payoff is given by

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

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

Wen (1994) defined the effective minmax payoff to be

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

In our setting, with actions constrained to lie in $A^0 \subseteq A$, the obvious extensions of the standard minmax and of Wen’s definition are respectively

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

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

The effective minmax we defined above differs from this extension of Wen’s definition because unlike Wen, we do not assume that mixed strategies are observable, so our analysis needs to consider the incentives of players in I_i^- to randomize their actions. Inducing players in I_i^- to randomize when mixing probabilities are not observed requires the use of continuation payoffs to make each player indifferent among all actions chosen with positive probabilities. Since players i and j have opposite payoff functions, player i is also indifferent among such actions. This may induce a higher reward than Wen’s minmax payoff.⁹

Recall that a game satisfies the *nonequivalent utilities (NEU) condition* of Abreu et al. (1994) if $I_i^+ = \{i\}$ for all i .

Proposition 4.1. *We have the following relations among the standard minmax, Wen’s effective minmax, and our effective minmax:*

- (1) $\underline{v}_i^s(\mathcal{A}^0) \leq \underline{v}_i^{\text{Wen}}(\mathcal{A}^0) \leq \underline{v}_i(\mathcal{A}^0)$.
- (2) $\underline{v}_i^{\text{Wen}}(\mathcal{A}^0) = \underline{v}_i(\mathcal{A}^0)$ if $\mathcal{A}^0 = \mathcal{A}^P$ or $I_i^- = \emptyset$.
- (3) $\underline{v}_i^s(\mathcal{A}^0) = \underline{v}_i(\mathcal{A}^0)$ if $(\mathcal{A}^0 = \mathcal{A}^P \text{ or } \mathcal{A})$ 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 minmax 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 minmax 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 $g_i(a_i, \alpha_{-i}^*) \leq g_i(\alpha^*)$ 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 $\underline{v}_i(\mathcal{A}) \leq g_i(\alpha^*) = \underline{v}_i^s(\mathcal{A})$. \square

Example. We may have $\underline{v}_i^{\text{Wen}}(\mathcal{A}^0) < \underline{v}_i(\mathcal{A}^0)$. Consider the stage game in Fig. 4.2. Note that $I_1^+ = \{1, 2\}$ and $I_1^- = \{3\}$. We have $\underline{v}_1^{\text{Wen}}(\mathcal{A}) = 5/2$, where the solution α to Wen’s minmax 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 minmax 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.

⁹ A similar payoff bound that arises from indifference conditions has been investigated in repeated games with short-run players. Fudenberg et al. (1990) showed that

$$\underline{v}_i^* := \max_{\alpha \in \text{graph}(B)} \min\{g_i(a_i, \alpha_{-i}) \mid a_i \in A_i \text{ s.t. } \alpha_i(a_i) > 0\}$$

is an upper bound for long-run player i ’s equilibrium payoffs if mixing probabilities are unobservable.

0, 0, 0	3, 3, -3	4, 4, -4	2, 2, -2
2, 2, -2	4, 4, -4	3, 3, -3	4, 4, -4

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

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

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

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

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 $(\mathcal{A}^0 = \mathcal{A}^p \text{ or } \mathcal{A})$ and the NEU condition is satisfied. Wen (1994) showed the pure-strategy folk theorem, which corresponds to Proposition 4.2 for $\mathcal{A}^0 = \mathcal{A}^p$.¹⁰ These classical results are stronger than Proposition 4.2 in the following sense. They show that $E(\mathcal{A}^0, \delta) \subseteq V(\mathcal{A}^0)$ for any δ , and that, for any $v \in V^*(\mathcal{A}^0)$, there exists $\underline{\delta} < 1$ such that $v \in E(\mathcal{A}^0, \delta)$ (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 $v \in V(\mathcal{A}^0)$ is approximately attained as an equilibrium payoff profile. See Section 4.4 for a discussion of the exact attainability of efficient payoffs.

We will show Proposition 4.2 by applying our algorithm. Let X be the affine hull of V . We have $\dim X \geq 1$ 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 $H(\lambda, \lambda \cdot v) \cap X = \{v' \in X \mid v'_i \geq v_i\}$.

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

Proof. Let λ be the orthogonal projection of $-e^i$ to X , where e^i is the vector whose i th component is one and whose other components are zeros. λ is nonzero since player i is not universally indifferent. By construction, λ is a punishment direction for player i . \square

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 orthogonal 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 $X_i = \mathbb{R}$.

¹⁰ 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 minmax may be lower than ours. In this case, the effective minmax value in his definition is not a tight lower bound for mixed-strategy subgame-perfect equilibrium payoffs. Thus our results show that the assumption that mixed strategies are observable is not innocuous in cases where the NEU condition is not satisfied. 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 $X_{ij} = \mathbb{R}^2$. \square

For a mixed action profile α and player i 's payoff v_i , we say that (α, v_i) is *enforceable with respect to W* if there exists a payoff profile v' with $v'_i = v_i$ such that (α, v) is enforceable with respect to W .

Lemma 4.5. (α, v_i) is enforceable with respect to $\{v' \in X \mid v'_i \geq v_i\}$ if and only if $v_i \geq g_i(a_j, \alpha_{-j})$ 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. The main difficulty to prove the “if” part is how to induce players to randomize. To control incentives of players in $I_i^+ \cup I_i^-$, we use Lemma A.3 in the appendix. Intuitively speaking, this lemma claims that if $v_i \geq g_i(a_j, \alpha_{-j})$ for any $j \in I_i^+ \cup I_i^-$ and any $a_j \in A_j$ such that $\alpha_j(a_j) > 0$, then we can find continuation payoffs to make all players in $I_i^+ \cup I_i^-$ indifferent simultaneously. Other players' incentives are easy to control without affecting player i 's payoff because we have shown $X_{ij} = \mathbb{R}^2$ for $j \notin I_i^+ \cup I_i^-$ in Lemma 4.4. \square

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 $v_i \leq x$.

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 and $v^2 \in X$ such that (i) $\lambda \cdot v^1 \geq \lambda \cdot v^2$ 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 $v_i \leq x$ 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) . \square

Lemma 4.7. If λ is not a punishment direction for any player, then $(\alpha, g(\alpha))$ is enforceable with respect to $H(\lambda, \lambda \cdot g(\alpha)) \cap X$ for any $\alpha \in \mathcal{A}^P$.

Proof. Let $a \in A$ be the pure strategy profile such that $\alpha(a) = 1$. Define $w(a') \in H(\lambda, \lambda \cdot g(\alpha)) \cap X$ 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 $w(a') \in H(\lambda, \lambda \cdot g(\alpha)) \cap X$ such that $w_i(a') \leq [g_i(\alpha) - (1 - \delta)g_i(a'_i, a_{-i})]/\delta$.
- If $a' = a$ or $a'_j \neq a_j$ for at least two players j , then let $w(a') = g(\alpha)$.

Then $(\alpha, g(\alpha))$ is enforced by w . \square

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 an 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 interior of V . This contradicts the fact that the closure of the interior of compact and convex set V is equal to V . Since $V \setminus V^*(\mathcal{A}^0)$ is closed, v is also an interior point of $V^*(\mathcal{A}^0)$, so $V^*(\mathcal{A}^0)$ and X have the same dimension. \square

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 $\mathcal{A}^0 \supseteq \mathcal{A}^p$, it follows from Lemmas 4.3, 4.5, and 4.7 that we have $Q^*(\mathcal{A}^0) \subseteq Q^0 = V(\mathcal{A}^0)$.

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 $Q^*(\mathcal{A}^0) = V(\mathcal{A}^0)$. \square

4.3. Symmetry assumptions

4.3.1. Strongly symmetric equilibria

Assume that the static game is symmetric for long-run players, i.e., $A_1 = \dots = A_L$ 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 of symmetric mixed action profiles for the long-run players in graph (B) , $\mathcal{A}^s := \{\alpha \in \text{graph}(B) \mid \alpha_1 = \dots = \alpha_L\}$, and denote by Q^s the result $Q^*(\mathcal{A}^s)$ 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 . 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}^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 $Q^s = Q^0$. 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. $Q^s = \lim_{\delta \rightarrow 1} E(\mathcal{A}^s, \delta)$. That is, Q^s is the limit as δ goes to one 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 does not apply to partially symmetric equilibria when there are $L - 1$ or less groups because Q^0 does not satisfy the full dimensionality condition. In principle one could 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 \mathcal{A}^0 and 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 subset K of 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, in a game without short-run players, 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 one, 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).

We focus on \mathcal{A}^h -PPE, assuming that equilibrium actions are in \mathcal{A}^h even at off-path histories. Every \mathcal{A}^h -PPE has 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).¹¹

Here we sketch how to obtain a sufficient condition for exact achievability. Let $X^0 = h$ and $\mathcal{A}^0 = \mathcal{A}^h$. Using our algorithm, we compute $Q^*(\mathcal{A}^h)$, which we denote 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 $\delta > \bar{\delta}$.*

This gives a sufficient condition for the exact attainability of Pareto-efficient payoffs; if hyperplane h is such that $V \cap h$ is a set of Pareto-efficient payoff profiles and $Q^h \neq \emptyset$, then every relative interior point of Q^h is a Pareto-efficient outcome that can be exactly attained by some \mathcal{A}^h -PPE payoff for sufficiently large δ .

¹¹ 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, even if the full support condition is not satisfied, we can modify our algorithm to analyze perfect public equilibria with payoff profiles in h without assuming off-path outcomes in \mathcal{A}^h . See the end of this subsection. Note also that allowing off-path play not in \mathcal{A}^h may destroy the renegotiation-proofness property of the equilibria; some equilibria which use non- \mathcal{A}^h actions in off-path play are still renegotiation-proof, but others not.

So far we have imposed the restriction of \mathcal{A}^h not only on on-path play but also on off-path play. In the case where \mathcal{A}^h corresponds to efficient payoffs, this can be motivated by the relationship with renegotiation-proofness. However, one can also follow Athey and Bagwell (2001), investigating possibility that there are exactly efficient equilibria that rely on the ability to use inefficient continuation payoffs after some outcomes that have zero probability along the equilibrium path. (Outcomes that trigger inefficient continuation payoffs must have zero probability along the equilibrium path for the overall equilibrium to be efficient.) For example, in Athey and Bagwell (2001), in each period, firms report their current period marginal costs and then choose prices. The equilibrium strategies assign a unique price to each cost report, so that many report/price combinations have zero probability in equilibrium, and the efficient equilibria that Athey and Bagwell constructed specify inefficient reversion to a static Nash equilibrium following these “off-path” deviations.

To analyze when these sorts of efficient equilibria exist, we now modify our algorithm by only imposing the restriction that actions lie in \mathcal{A}^h along the path of play. Action a_i is an *on-path deviation from α* if $\pi_y(a_i, \alpha_{-i}) > 0$ implies $\pi_y(\alpha) > 0$ for any $y \in Y$. An on-path deviation is a deviation which cannot be detected perfectly. a_i is an *off-path deviation from α* if a_i is not an on-path deviation from α .

We modify our algorithm as follows. Let $\mathcal{A}^0 = \mathcal{A}^h$. At each step, in the linear programming problem for $\alpha \in \mathcal{A}^h$, we impose incentive constraints (b) only for on-path deviations from α . Then denote the final output of the algorithm by Q_{on}^h instead of Q^h .

Proposition 4.11. (1) $E^*(\text{graph}(B), \delta) \cap h \subseteq Q_{\text{on}}^h$ for any δ . Hence if $Q_{\text{on}}^h = \emptyset$, then $E^*(\text{graph}(B), \delta) \cap h = \emptyset$ for any δ . (2) If $Q_{\text{on}}^h \neq \emptyset$ and a relative interior point v of Q_{on}^h strictly dominates some point in $Q^*(\text{graph}(B))$, then there exists $\bar{\delta} < 1$ such that $v \in E(\text{graph}(B), \delta)$ for any $\delta > \bar{\delta}$.

Proof. [Sketch of the proof of (1)] If a PPE achieves a payoff profile in h , all continuation payoffs on the equilibrium path lie in h . Thus, as far as on-path deviations are concerned, we can use linear programming problems with continuation payoffs in h to bound equilibrium payoffs, so we can modify Lemma 3.2 and our main theorem to obtain this result.

[Sketch of the proof of (2)] Let v' be a relative interior point of $Q^*(\text{graph}(B))$ strictly dominated by v . Modifying the argument in Lemma 3.3, we can show that, for sufficiently large δ , there exists a public strategy profile such that the long-run players earn v after the initial history, all continuation payoff profiles strictly dominates v' , and there is neither a profitable deviation for short-run players nor profitable on-path deviations for long-run players. We can punish off-path deviations for long-run players by reverting to a PPE that achieves v' . \square

Since Q_{on}^h may not be empty even if Q^h is, Proposition 4.11 may provide a sufficient condition for the exact achievability of first-best outcomes even when Corollary 4.10 is not applicable. The Athey and Bagwell result corresponds to the case where a point in the relative interior of Q_{on}^h strictly Pareto dominates a static Nash equilibria; static Nash equilibrium payoffs are always in $Q^*(\text{graph}(B))$.

Appendix A. Proofs

A.1. Smooth sets in the proof of Lemma 3.3

We assume throughout the non-trivial case that $\dim X \geq 1$. First we examine the set W constructed in the proof of Lemma 3.3.

Lemma A.1. *For any compact subset K of X and any $\varepsilon > 0$, $W := \{v \in X \mid \min_{x \in \text{co}(K)} |v - x| \leq \varepsilon\}$ is convex, compact, and smooth in X .*

Proof. For any $v, v' \in W$ and $0 \leq \alpha \leq 1$, there exist $x, x' \in \text{co}(K)$ such that $|v - x| \leq \varepsilon$ and $|v' - x'| \leq \varepsilon$. Then, since we have

$$|(\alpha v + (1 - \alpha)v') - (\alpha x + (1 - \alpha)x')| \leq \alpha|v - x| + (1 - \alpha)|v' - x'| \leq \varepsilon$$

and $\alpha x + (1 - \alpha)x' \in \text{co}(K)$, we have $\alpha v + (1 - \alpha)v' \in W$. So W is convex.

Since K is compact, $\text{co}(K)$ is also compact, and hence it is bounded. Then there exists $M > 0$ such that $|x| \leq M$ for any $x \in \text{co}(K)$. For any $v \in W$, there exists $x \in \text{co}(K)$ such that $|v - x| \leq \varepsilon$. Then we have $|v| \leq |x| + |v - x| \leq M + \varepsilon$. So W is bounded.

For any sequence $\{v^m\}$ on W that converges to v , there exists a sequence $\{x^m\}$ on $\text{co}(K)$ such that $|v^m - x^m| \leq \varepsilon$. Since $\text{co}(K)$ is compact, there exists a limit $x \in \text{co}(K)$ of $\{x^m\}$ (take a subsequence if necessary). In the limit, we have $|v - x| \leq \varepsilon$, and hence we have $v \in W$. So W is closed.

For any relative boundary point v of W , there exists $x \in \text{co}(K)$ such that $|v - x| = \varepsilon$. In order for a vector λ to be normal to W at v , $H(\lambda, \lambda \cdot v)$ has to contain W , and hence it has to contain the closed ball with center x and radius ε . Then a vector of length one parallel to X and normal to W at v is uniquely determined by $(v - x)/\varepsilon$. So W is smooth. \square

Lemma A.2. *Let W be a convex and compact subset of \mathbb{R}^L that is smooth in X . For a relative boundary point v of W , let λ be the unique vector of length one parallel to X and normal to W at v . For any $\kappa, \kappa' > 0$, there exists $\varepsilon \in (0, 1)$ such that every $v' \in X$ that satisfies $\lambda \cdot v' \leq \lambda \cdot v - \kappa\varepsilon$ and $|v' - v| \leq \kappa'\varepsilon$ is a relative interior point of W .*

Proof. If not, we can find positive numbers κ and κ' and a sequence $\{v^m\}_{m=2}^\infty$ of points in X such that each v^m is not in the relative interior of W , $\lambda \cdot v^m \leq \lambda \cdot v - \kappa/m$, and $|v^m - v| \leq \kappa'/m$. We have $v^m \rightarrow v$. For each m , by the separating hyperplane theorem, there exists a vector of length one, λ^m , parallel to X such that $W \subseteq H(\lambda^m, \lambda^m \cdot v^m)$. Since $v \in W$, we have $\lambda^m \cdot v \leq \lambda^m \cdot v^m$. Then we have

$$|\lambda^m - \lambda| \times |v^m - v| \geq (\lambda^m - \lambda) \cdot (v^m - v) \geq \kappa/m.$$

Since $|v^m - v| \leq \kappa'/m$, we have $|\lambda^m - \lambda| \geq \kappa/\kappa' > 0$. Since the set of vectors of length one parallel to X is compact, there exists a limit λ^* of $\{\lambda^m\}$ (take a subsequence if necessary). Then we have $\lambda^* \neq \lambda$ and $W \subseteq H(\lambda^*, \lambda^* \cdot v)$, which contradicts the definition of smoothness. \square

A.2. 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) proved the

same result for the case of two players. For notational convenience, for a given mixed action profile α , take $S_j := \{a_j \in A_j \mid \alpha_j(a_j) > 0\}$, $S := \prod_{j \in J} S_j$, $S_{-j} := \prod_{k \in J, k \neq j} S_k$, $\sigma := \alpha_J$, and $u(\sigma) := g_i(\sigma, \alpha_{-j})$.

Lemma A.3. *If $\sigma_j(s_j) > 0$ and $x \geq u(s_j, \sigma_{-j})$ for all $j \in J$ and all $s_j \in S_j$, then there exists $f : S \rightarrow \mathbb{R}$ such that $f(s) \geq x$ 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 construct the desired f explicitly. Let $S_j^0 := S_j$, $p_{-j}^0 = 1$, and $r_j^0(s_j) := [x - (1 - \delta)u(s_j, \sigma_{-j})]/\delta$ for each $j \in J$ and each $s_j \in S_j$. Here $r_j^0(s_j)$ is the amount of continuation payoff player j needs after choosing action s_j in order to achieve the target payoff x . Note that we express $r_j^0(s_j)$ in terms of player i 's utilities. We can do so because players in J have perfectly (positively or negatively) correlated payoffs. We now define a series of steps $m = 0, 1, \dots$. At step 0, we pick all pairs (j, s_j) that minimize $r_j^0(s_j)$. For any such pair, if player j takes action s_j , give him the continuation payoff $f(s_j, s_{-j})$ equal to $x^0 = r_j^0(s_j)$ independently of realized action profile s_{-j} of the other players. Notice that if $x = \max\{u(s_j, \sigma_{-j}) \mid j \in J, s_j \in S_j\}$, then this is the only possible solution with $f(s) \geq x$. We now have left the set S_j^1 of actions that remain.

Now (in general, from step $m - 1$), let S_j^m be the set of player j 's remaining actions at step m . We define

$$p_{-j}^m := \prod_{k \in J, k \neq j} \sum_{s_k \in S_k^m} \sigma_k(s_k)$$

to be the probability that $f(s)$ has been undetermined if player j chooses one of his remaining actions. Note that even if action s_j remains at step m , $f(s_j, s_{-j})$ was already determined as x^{m-1} at the previous step based on other players' actions s_{-j} . Since this occurs with probability $p_{-j}^{m-1} - p_{-j}^m$, the remaining amount of continuation payoff player j needs after choosing action s_j is given by

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

We consider the following minimization problem:

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

If pair (j, s_j) solves this minimization problem, we determine $f(s_j, s_{-j})$ as x^m at this step independently of s_{-j} . The set of remaining actions at the next step is given by

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

Observe that $S_j^{m+1} \subseteq S_j^m$ for all $j \in J$ and $S_j^{m+1} \subsetneq S_j^m$ for some $j \in J$ so that this iteration stops in a finite number of steps. Let m^* be the first step at which $S_j^{m^*+1} = \emptyset$ for some $j \in J$. Then f is completely determined at the end of step m^* . We will show that this f is the desired solution. For each $s \in S$, let $m(s) \leq m^*$ be the step at which $f(s)$ is determined.

First we show that $\sum_{s_j \in S_j^m} \sigma_j(s_j)r_j^m(s_j)$ is independent of j . The key property of $u(s_j, \sigma_{-j})$ is that $\sum_{s_j \in S_j} \sigma_j(s_j)u(s_j, \sigma_{-j}) = u(\sigma)$ is independent of j . It follows that $\sum_{s_j \in S_j^0} \sigma_j(s_j)r_j^0(s_j) = [x - (1 - \delta)u(\sigma)]/\delta$ is also 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 we have

$$\begin{aligned} \sum_{s_j \in S_j^{m+1}} \sigma_j(s_j)r_j^{m+1}(s_j) &= \sum_{s_j \in S_j^{m+1}} \sigma_j(s_j)[r_j^m(s_j) - (p_{-j}^m - p_{-j}^{m+1})x^m] \\ &= \sum_{s_j \in S_j^{m+1}} \sigma_j(s_j)(r_j^m(s_j) - p_{-j}^m x^m) + \sum_{s_j \in S_j^{m+1}} \sigma_j(s_j)p_{-j}^{m+1}x^m \\ &= \sum_{s_j \in S_j^m} \sigma_j(s_j)r_j^m(s_j) - \sum_{s_j \in S_j^m} \sigma_j(s_j)p_{-j}^m x^m \\ &\quad - \sum_{s_j \in S_j^m \setminus S_j^{m+1}} \sigma_j(s_j)(r_j^m(s_j) - p_{-j}^m x^m) \\ &\quad + \sum_{s_j \in S_j^{m+1}} \sigma_j(s_j)p_{-j}^{m+1}x^m. \end{aligned}$$

The first term is independent of j by the induction hypothesis. The second and the fourth terms are independent of j since they respectively are equal to $(\prod_{k \in J} \sum_{s_k \in S_k^m} \sigma_k(s_k))x^m$ and to $(\prod_{k \in J} \sum_{s_k \in S_k^{m+1}} \sigma_k(s_k))x^m$. As S_j^{m+1} is defined precisely so that $r_j^m(s_j) = p_{-j}^m x^m$ for any $s_j \in S_j^m \setminus S_j^{m+1}$, the third term is zero.

Next, we observe that by the definition of m^* , we have $S_j^{m^*+1} = \emptyset$ for some $j \in J$. Our second step is to show that this holds for all $j \in J$. If $S_k^{m^*+1} = \emptyset$ for some $k \in J$ but $S_{k'}^{m^*+1} \neq \emptyset$ for some other $k' \in J$, then we have

$$\begin{aligned} \sum_{s_k \in S_k^{m^*}} \sigma_k(s_k)r_k^{m^*}(s_k) &= \sum_{s_k \in S_k^{m^*}} \sigma_k(s_k)p_{-k}^{m^*}x^{m^*} = \left(\prod_{l \in J} \sum_{s_l \in S_l^{m^*}} \sigma_l(s_l)\right)x^{m^*} \\ &= \sum_{s_{k'} \in S_{k'}^{m^*}} \sigma_{k'}(s_{k'})p_{-k'}^{m^*}x^{m^*} < \sum_{s_{k'} \in S_{k'}^{m^*}} \sigma_{k'}(s_{k'})r_{k'}^{m^*}(s_{k'}). \end{aligned}$$

This contradicts the player-independence property of $\sum_{s_j \in S_j^{m^*}} \sigma_j(s_j)r_j^{m^*}(s_j)$. Note that the strict inequality holds because $r_{k'}^{m^*}(s_{k'})/p_{-k'}^{m^*} \geq x^{m^*}$ for all $s_{k'} \in S_{k'}^{m^*}$, strictly for all $s_{k'} \in S_{k'}^{m^*+1} \neq \emptyset$, and $\sigma_{k'}(s_{k'}) > 0$ for all $s_{k'} \in S_{k'}^{m^*}$. We conclude that $S_j^{m^*+1} = \emptyset$ for all $j \in J$.

It remains to show that f is in fact the desired solution. First we must show that $f(s) \geq x$ for all $s \in S$. Since $f(s) = x^{m(s)}$, this amounts to showing $x^m \geq x$ for all $m \leq m^*$. We do this recursively. Since $x \geq \max\{u(s_j, \sigma_{-j}) \mid j \in J, s_j \in S_j\}$, we have

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

For any $m < m^*$, then there exists a pair of $j \in J$ and $s_j \in S_j^{m+1}$ such that $r_j^{m+1}(s_j)/p_{-j}^{m+1} = x^{m+1}$. Then we have

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

Thus, if $x^m \geq x$, then $x^{m+1} > x^m \geq x$.

Finally, we show that f as constructed in fact solves the equation in question. For any $j \in J$, since $S_j^{m^*+1} = \emptyset$, it follows that for any $s_j \in S_j$, there exists a unique $m_j(s_j) \leq m^*$ such that $s_j \in S_j^{m_j(s_j)} \setminus S_j^{m_j(s_j)+1}$. Then we have

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

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

Lemma 4.5. (α, v_i) is enforceable with respect to $\{v' \in X \mid v'_i \geq v_i\}$ if and only if $v_i \geq g_i(a_j, \alpha_{-j})$ 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. By Lemma 4.4, fix an arbitrary $v \in X$ whose i th coordinate is v_i . Define $w(a) \in X$ with $w_i(a) \geq v_i$ 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 any vector such that

$$v_i \leq w_i(a) \leq \frac{1}{\delta} v_i - \frac{1 - \delta}{\delta} g_i(a_j, \alpha_{-j}).$$

It follows from Lemma 4.4 and $v_i \geq g_i(a_j, \alpha_{-j})$ that at least one 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 any vector such that $w_i(a) \geq v_i$ and

$$w_j(a) \leq \frac{1}{\delta} v_j - \frac{1 - \delta}{\delta} g_j(a_j, \alpha_{-j}).$$

It follows from Lemma 4.4 that at least one such $w(a)$ exists.

- Let $S_j := \{a_j \in A_j \mid \alpha_j(a_j) > 0\}$, $J := I_i^+ \cup I_i^-$, $S_J := \prod_{j \in J} S_j$, and $S := \prod_{j=1}^n S_j$. We simultaneously define $w(a) \in X$ for all $a \in S$ as follows. By setting $x := v_i$ in Lemma A.3, we obtain a function f such that $f(a_J) \geq v_i$ for all $a_J \in S_J$ and

$$v_i = (1 - \delta)g_i(a_j, \alpha_{-j}) + \delta \sum_{a_{-j} \in A_{-j}} \alpha_{-j}(a_{-j})f(a_J)$$

for all $j \in J$ and all $a_j \in S_j$. For each $a_J \in S_J$, let $w^J(a_J) \in X$ be any vector such that $w_i^J(a_J) = f(a_J)$. It follows from Lemma 4.4 that at least one such $w^J(a_J)$ exists. If the continuation payoffs are given by $w^J(a_J)$, then each player $j \in J$ is indifferent among actions in S_j , and player i 's total payoff is equal to v_i .

By modifying w^J , we can make each player $j \notin J$ indifferent among actions in S_j , maintaining incentives of players in J . For example, define $w^0 \in X$ by

$$n'w^0 := (n' - 1)g(\alpha) + \frac{1}{1 - \delta}v - \frac{\delta}{1 - \delta} \sum_{a \in A} \alpha(a)w^J(a_J),$$

where $n' := n - |J|$. (If $n' = 0$, then we choose w^0 arbitrarily.) Finally, define

$$w(a) := w^J(a_J) + \frac{1 - \delta}{\delta} \left(n'w^0 - \sum_{j \notin J} g(a_j, \alpha_{-j}) \right) \in X$$

for each $a \in S$. Then each player j is indifferent among actions in S_j and the total payoff profile is equal to v .

- If $\alpha_j(a_j) = 0$ for at least two players j , let $w(a) := v \in X$.

Then (α, v) is enforced by w .

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

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

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

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

$$v_i \geq (1 - \delta)g_i(a_j, \alpha_{-j}) + \delta \sum_{a_{-j} \in A_{-j}} \alpha_{-j}(a_{-j})w_i(a) \geq (1 - \delta)g_i(a_j, \alpha_{-j}) + \delta v_i,$$

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

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

$$v_i = (1 - \delta)g_i(a_j, \alpha_{-j}) + \delta \sum_{a_{-j} \in A_{-j}} \alpha_{-j}(a_{-j})w_i(a) \geq (1 - \delta)g_i(a_j, \alpha_{-j}) + \delta v_i,$$

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

References

- Abreu, D., 1986. Extremal equilibria of oligopolistic supergames. *J. Econ. Theory* 39, 191–228.
- Abreu, D., Pearce, D., Stacchetti, E., 1986. Optimal cartel equilibria with imperfect monitoring. *J. Econ. Theory* 39, 251–269.
- Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58, 1041–1063.
- Abreu, D., Dutta, P.K., Smith, L., 1994. The folk theorem for repeated games: a NEU condition. *Econometrica* 62, 939–948.
- Athey, S., Bagwell, K., 2001. Optimal collusion with private information. *RAND J. Econ.* 32, 428–465.
- Dellarocas, C., 2003. Efficiency and robustness of eBay-like online reputation mechanisms in environments with moral hazard. Working paper. MIT.
- Ely, J., Fudenberg, D., Levine, D. K., 2003. When is reputation bad? *Games Econ. Behav.* In press.
- Farrell, J., Maskin, E., 1989. Renegotiation in repeated games. *Games Econ. Behav.* 1, 327–360.
- Fudenberg, D., Levine, D.K., 1994. Efficiency and observability with long-run and short-run players. *J. Econ. Theory* 62, 103–135.
- Fudenberg, D., Maskin, E., 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54, 533–554.
- Fudenberg, D., Maskin, E., 1990. Nash and perfect equilibria of discounted repeated games. *J. Econ. Theory* 51, 194–206.
- Fudenberg, D., Kreps, D.M., Maskin, E., 1990. Repeated games with long-run and short-run players. *Rev. Econ. Stud.* 57, 555–573.
- Fudenberg, D., Levine, D.K., Maskin, E., 1994. The folk theorem in repeated games with imperfect public information. *Econometrica* 62, 997–1039.
- Kandori, M., Matsushima, H., 1998. Private observation, communication and collusion. *Econometrica* 66, 627–652.
- Tomala, T., 2005. On subgame-perfect communication equilibria in repeated games with imperfect monitoring. Mimeo.
- Wen, Q., 1994. The ‘folk theorem’ for repeated games with complete information. *Econometrica* 62, 949–954.