

THE EXPONENTIAL CONVERGENCE
OF BAYESIAN LEARNING IN NORMAL FORM GAMES

by

James S. Jordan

Discussion Paper No. 259, July 1990

Center for Economic Research
Department of Economics
University of Minnesota
Minneapolis, Minn 55455

The Exponential Convergence of Bayesian Learning in Normal Form Games

J.S. Jordan*

Department of Economics
University of Minnesota
1035 Management and Economics
271 19th Avenue South
Minneapolis, MN 55455

July 1990

*The support of the National Science Foundation is gratefully acknowledged. I would also like to thank Julio Escolano for numerous corrections to an earlier draft of this paper. Of course, any errors that remain are the sole responsibility of the author.

Abstract

This paper continues the study of Bayesian learning processes for general finite-player, finite-strategy normal form games. Bayesian learning was introduced in an earlier paper by the present author as an iterative mechanism by which players can learn Nash equilibria. The main result of the present paper is that if prior beliefs are sufficiently uniform and expectations converge to a "regular" Nash equilibrium, then the rate of convergence is exponential.

1. Introduction

This paper continues the study of Bayesian learning processes for general finite-player, finite-strategy normal form games. Bayesian learning was introduced in Jordan (1990) as an iterative mechanism by which players can learn Nash equilibria. Initially, each player is presumed to know his own payoff function but not the payoff functions of the other players. The initial strategies are determined as a Bayesian Nash equilibrium of the incomplete information game in which each player's private characteristic is his payoff function. These strategies are then observed by all players, causing a revision of beliefs. The new beliefs determine a new Bayesian Nash equilibrium and so on. Assuming that the common prior distribution of payoff functions satisfies independence across players, it was proved in Jordan (1989, Theorem 3.2) that the conditional distributions on strategies converge to a subset of the set of Nash equilibria with probability one. Under a further assumption that the common prior distributions are sufficiently uniform, convergence to the set of Nash equilibria was proved for every profile of payoff functions, that is, for every normal form game (Jordan (1989, Theorem 3.8)). If a game has multiple Nash equilibria, the sequence of expectations may contain subsequences converging to different Nash equilibria, which is why the two convergence results only assert convergence to a set of Nash equilibrium set.

The focus of the present paper is the rate at which expectations converge. Our main result, Theorem 3.6, states that if a subsequence of expectations converges to a "regular" Nash equilibrium, the rate of convergence is exponential, again under the assumption that prior beliefs are sufficiently uniform. This result is in contrast with the behavior of learning by "fictitious play" (Robinson (1951)) and somewhat similar that learning mechanisms studied recently by Fudenberg and Kreps (1988), Milgrom and Roberts (1989), and others, which converge, if at all, at a rate no faster than t^{-1} .

The rate of convergence of a learning mechanism is directly relevant to the interpre-

tation of equilibrium. Since a Nash equilibrium requires considerable coordination among the players, one would not expect the equilibrium conditions to hold exactly in any informationally decentralized setting. However, the existence of a learning mechanism which converges to equilibrium fairly rapidly provides some theoretical justification for interpreting the Nash equilibrium concept as a good approximation.

Of course, the comparatively rapid convergence of Bayesian learning should not be interpreted as an assertion of computational superiority. In fact, since each iteration of a Bayesian learning process is a Bayesian equilibrium of an incomplete information game, each iteration is computationally more difficult than solving for the Nash equilibrium directly. Instead, the conclusion to be drawn from the exponential convergence result appears to be that if players make substantial use of deductive reasoning in forming expectations, convergence may be much more rapid than if expectations are formed in a quasi-statistical inductive fashion.

The exponential convergence result also depends on the absence of "noise". If the players' payoff functions or strategy choices were subject to random perturbations each period, then Bayesian learning would necessarily proceed more slowly.

The remainder of the paper is organized as follows. Section 2 gives the definition of regular Nash equilibrium, and Section 3 defines Bayesian strategy processes and gives the statement and proof of the main result. Since Bayesian strategy processes are discussed in Jordan (1989), only the necessary definitions are provided below. The reader is referred to the earlier paper for examples and motivation. However, Section 3 below describes an abstract example which motivates the exponential convergence result. A number of analytical results needed in the proof are collected in an appendix to this paper..

2. Regular Nash Equilibrium

2.1 Definitions: There are n players, $n \geq 2$, indexed by the subscript p . For each

$1 \leq p \leq n$, let S_p be a finite set, and let $S = \prod_{p=1}^n S_p$. For each p , let $S_{-p} = \prod_{q \neq p} S_q$, with generic element $s_{-p} = (s_1, \dots, s_{p-1}, s_{p+1}, \dots, s_n)$. Each player p is characterized by a payoff function $\pi_p : S \rightarrow R$. The space of possible payoff functions for each player is the unit ball in R^S , that is, $B = \{\pi \in R^S : (\sum_{s \in S} \pi(s)^2)^{\frac{1}{2}} \leq 1\}$. Given the strategy sets S_p , a normal form game is completely specified by an n -tuple of payoff functions $(\pi_p)_p \in B^n$, so let $G = B^n$ denote the space of all games, with generic element $(\pi_p)_p$.

A probability distribution σ on S is a *product distribution* if $\sigma = \sigma_1 \times \dots \times \sigma_n$, where each σ_p is the marginal distribution on S_p . That is, $\sigma(s) = \sigma_1(s_1) \cdots \sigma_n(s_n)$ for each $s \in S$. A *Nash equilibrium* for a game $(\pi_p)_p \in G$ is a product distribution σ^* on S such that for each p , σ_p^* maximizes $\sum_{s_p} \sigma_p(s_p) \sum_{s_{-p}} \pi_p(s_p, s_{-p}) \sigma_{-p}^*(s_{-p})$ over the set of probability measures σ_p on S_p , where $\pi_p(s_p, s_{-p}) = \pi_p(s_p, \dots, s_{p-1}, s_p, s_{p+1}, \dots, s_n)$, and σ_{-p}^* is the distribution $\prod_{q \neq p} \sigma_q^*$ on S_{-p} . For each $(\pi_p)_p \in G$, let $N((\pi_p)_p)$ denote the set of Nash equilibria.

2.2 Remarks: It might seem more natural to normalize payoff functions to the unit sphere, since the extra dimension present in B is redundant, but the convexity of the unit ball simplifies some of the analysis. The space of unnormalized payoff functions, R^S , is also convex, but the compactness of B is also convenient. Of course, the Theorem is independent of the normalization.

The definition of “regular Nash equilibrium”, below, is fairly intuitive. In a Nash equilibrium, the strategies which have positive probability are characterized by equations. All strategies s_p with $\sigma_p(s_p) > 0$ must yield the same expected payoff to player p . Hence it is natural to require that these equations, as polynomial equations in these positive probabilities, have full rank. This ensures the local uniqueness of the positive probabilities. The strategies which have zero probability are characterized by inequalities, that is, they yield lesser expected payoffs, so it is natural to require that these inequalities be strict. This ensures the local uniqueness of the zero probabilities and thus the local uniqueness of

regular Nash equilibria. These two requirements are embodied in conditions (a) and (b), respectively, below.

2.3 Definitions: Let $(\pi_p)_p \in G$ and let σ^* be a Nash equilibrium for $(\pi_p)_p$. For each p , let $L_p \leq \#S_p$ and suppose that S_p is indexed as $\{s_{ip}\}_{i=1}^{\#S_p}$, where $\sigma_p^*(s_{ip}) > 0$ if $i \leq L_p$ and $\sigma_p^*(s_{ip}) = 0$ if $i > L_p$. For each p and each $i \leq \#S_p$, define the function $\pi_{ip} : \prod_{q=1}^n R^{S_q} \rightarrow R$ by

$$\pi_{ip}(\sigma) = \sum_{\{s_{-p} \in \prod_{q \neq p} S_q\}} \pi_p(s_{ip}, s_{-p}) \sigma_{-p}(s_{-p})$$

for all $\sigma \in \prod_{q=1}^n R^{S_q}$, where

$$\sigma_{-p}(s_{-p}) = \sigma_1(s_1) \cdots \sigma_{p-1}(s_{p-1}) \cdot \sigma_{p+1}(s_{p+1}) \cdots \sigma_n(s_n)$$

for each

$$s_{-p} = (s_1, \dots, s_{p-1}, s_{p+1}, \dots, s_n) \in \prod_{q \neq p} S_q.$$

For each p with $L_p > 1$, define $g_p : \prod_{q=1}^n R^{S_q} \rightarrow R^{L_p-1}$ by

$$g_p(\sigma) = (\pi_{ip}(\sigma) - \pi_{1p}(\sigma))_{i=2}^{L_p}.$$

Define $h : \prod_{p=1}^n R^{S_p} \rightarrow R^{\sum_{p=1}^n \#S_p}$ by

$$h(\sigma) = ((g_p(\sigma))_{\{p: L_p > 1\}}, (\sigma_{ip})_{\{i, p: i > L_p\}}), (1 - \sum_{i=1}^{\#S_p} \sigma_{ip})_{p=1}^n).$$

The Nash equilibrium σ^* is *regular* if

- a) $Dh(\sigma^*)$ is nonsingular; and
- b) $\pi_{ip}(\sigma^*) < \pi_{1p}(\sigma^*)$ for all p and all $i > L_p$.

2.4 Remarks: Since σ^* is a Nash equilibrium, $h(\sigma^*) = 0$. The nonsingularity of $Dh(\sigma^*)$ enables us to consider h as a local coordinate system near a regular Nash equilibrium.

3. The Main Result

3.1 Definitions: Suppose that X and Y are metric spaces and Z is the product space, $Z = X \times Y$. Let φ be a Borel probability measure on Z . The support of φ , written “supp φ ”, is defined to be the smallest closed subset $C \subset Z$ with $\varphi(C) = 1$. The support of the marginal probability distribution on X is the smallest closed subset $C \subset X$ such that $\varphi(C \times Y) = 1$. If $x \in X$, the statement “ $x \in \text{supp } \varphi$ ” will be understood to mean that x is in the support of the marginal distribution on X . If E is a Borel subset of X , $\varphi(E \times Y)$ will often be written simply $\varphi(E)$, and if $x \in X$, $\varphi(x)$ denotes $\varphi(\{x\})$, that is, $\varphi(\{x\} \times Y)$.

For each $1 \leq t < \infty$, let $S^t = \prod_{\tau=1}^t S$, with generic elements $s^t = (s_1, \dots, s_t)$ and let $S^\infty = \prod_{\tau=1}^\infty S$. We topologize each S^t and S^∞ as products of the discrete space S . Let G have the Euclidean topology, and let φ be a Borel probability measure on $G \times S^\infty$. For each p, t , each $s^t \in \text{supp } \varphi$, and each $s_p \in S_p$, let $\varphi_p(s_p|s^t)$ be the probability that $s_{pt+1} = s_p$ conditional on s^t , that is, $\varphi_p(s_p|s^t) = \varphi(s^t, s_p)/\varphi(s^t)$. For each $s_{-p} \in S_{-p}$, let $\varphi_{-p}(s_{-p}|s^t)$ be the probability that $s_{-pt+1} = s_{-p}$ conditional on s^t , that is $\varphi_{-p}(s_{-p}|s^t) = \varphi(s^t, s_{-p})/\varphi(s^t)$. If $t = 0$, we adopt the convention that $(s^t, s_p) = s_p$, $(s^t, s_{-p}) = s_{-p}$, and $\varphi_p(s_p|s^t)$ (resp. $\varphi_{-p}(s_{-p}|s^t)$) represents the unconditional probability that $s_{p1} = s_p$ (resp. $s_{-p1} = s_{-p}$). Finally, for each $t \geq 1$, and each $s^t \in \text{supp } \varphi$, $\varphi(s|s^t)$ denotes the probability that $s_{t+1} = s$ conditional on s^t , with the analogous convention for $t = 0$.

3.2 Bayesian Strategy Processes: For each p , let μ_p be a Borel probability distribution on G , and let μ denote the product measure, $\mu = \mu_1 \times \dots \times \mu_n$ on G . A Borel probability measure φ on $G \times S^\infty$ is a *Bayesian strategy process* (BSP) for μ if

- i) the marginal distribution on G agrees with μ ;
- ii) for each p , each $t \geq 0$, and each $(\pi_p; s^t, s_p) \in \text{supp } \varphi$, s_p maximizes $\sum_{s_{-p}} \pi_p(\cdot, s_{-p}) \varphi_{-p}(s_{-p}|s^t)$ on S_p ; and

iii) for each $s^t \in \text{supp } \varphi$, $\varphi(\cdot|s^t)$ is a product distribution, that is,

$$\varphi(\cdot|s^t) = \varphi_1(\cdot|s^t) \times \cdots \times \varphi_p(\cdot|s^t) \times \cdots \times \varphi_n(\cdot|s^t).$$

3.3 Remarks: Before stating the exponential convergence result, we record a useful preliminary result which is proved in (Jordan 1989, Proposition 3.6). As mentioned in the introduction, the exponential convergence result is stated for subsequences of expectations since different subsequences can converge to different Nash equilibria. Moreover, an iteration in one subsequence may yield no information relevant to another subsequence, so convergence to equilibrium is exponential only in the number of iterations along a given subsequence.

3.4 Definitions: Let m denote Lebesgue measure on B . For $1 \leq p \leq n$, the notation $\mu_p \ll m$ states that μ_p is absolutely continuous with respect to Lebesgue measure. For each p , we will write $\mu_p \sim m$ if there is some number $\alpha > 0$ such that for every Borel set $E \subset B$, $\alpha m(E) \geq \mu_p(E) \geq m(E)/\alpha$.

3.5 Proposition: Suppose that $\mu_p \ll m$ for each p and that φ is a BSP for μ . Let $s^t \in \text{supp } \varphi$ with $t \geq 1$, and for each p , define $C_p(s^t) = \{\pi_p \in B : \text{for each } \tau < t, s_{p\tau+1}^t \text{ maximizes } \Sigma_{s_{-p}} \pi_p(\cdot, s_{-p}) \varphi_{-p}(s_{-p}|s^\tau) \text{ on } S_p, \text{ where } s^\tau < s^t\}$, and define $C_p^\circ(s^t) = \{\pi_p \in B : \text{for each } \tau < t, s_{p\tau+1}^t \text{ is the unique maximizer of } \Sigma_{s_{-p}} \pi_p(\cdot, s_{-p}) \varphi_{-p}(s_{-p}|s^\tau), \text{ where } s^\tau < s^t\}$. Then for each p ,

- i) $C_p(s^t)$ is a compact convex set, and $C_p^\circ(s^t)$ is a convex set which is open relative to B ;
- ii) $\mu_p(C_p(s^t) \setminus C_p^\circ(s^t)) = m(C_p(s^t) \setminus C_p^\circ(s^t)) = 0$;
- iii) $\{\pi_p \in B : (\pi_p, s^t) \in \text{supp } \varphi\} \subset C_p(s^t)$; and
- iv) $\varphi(s^t) = \mu_1(C_1(s^t)) \times \cdots \times \mu_n(C_n(s^t))$.

3.6 Theorem: Suppose that $\mu_p \sim m$ for each p , and that φ is a BSP for μ . Let

$((\pi_p^\circ)_p, s^\infty) \in \text{supp } \varphi$, and let σ° be a regular Nash equilibrium for $(\pi_p^\circ)_p$, and suppose there is a subsequence indexed by $\{t(u)\}_{u=1}^\infty$ such that $\varphi(\cdot|s^{t(u)}) \rightarrow \sigma^\circ$. Then there are positive constants k and K such that $\|\varphi(\cdot|s^{t(u)}) - \sigma^\circ\| < K(1+k)^{-u}$ for all u .

3.7 Remarks: The proof of the Theorem is rather long, but the essential insight is contained in the following transparent mathematical problem. Suppose that a real variable π is distributed on the interval $(-1, 1)$ according to a probability measure μ which has a density function which is bounded and also bounded away from zero. First, π° is drawn from this distribution and remains fixed throughout the following process. Next, a number $\varphi_1 \in [0, 1]$ is chosen according to the equation

$$\varphi_1 = \mu(\{\pi : \pi > \varphi_1\}).$$

After φ_1 is chosen, we find out whether or not $\pi^\circ > \varphi_1$, and then choose φ_2 according to the same equation but with the probability conditioned on the new information, and so on. More precisely, define

$$s(\pi, \varphi) = \begin{cases} 1 & \text{if } \pi > \varphi \\ 0 & \text{otherwise;} \end{cases}$$

and for each t , let

$$\varphi_{t+1} = \mu(\{\pi : \pi > \varphi_{t+1}\} | s(\pi^\circ, \varphi_1), \dots, s(\pi^\circ, \varphi_t)),$$

where the right hand side is a conditional probability. Of course, the information $s(\pi^\circ, \varphi_1), \dots, s(\pi^\circ, \varphi_t)$ simply reveals that π° lies in the interval between the largest φ_i with $s(\pi^\circ, \varphi_i) = 1$ and the smallest φ_i with $s(\pi^\circ, \varphi_i) = 0$. If $s(\pi^\circ, \varphi_i) = 0$ for all $i \leq t$, then the left endpoint remains -1 , or if $s(\pi^\circ, \varphi_i) = 1$ for all $i \leq t$, then the right endpoint remains 1 .

The asymptotic behavior of φ_t depends on whether π° is positive. If $\pi^\circ \in (0, 1)$, then $\varphi_t \rightarrow \pi^\circ$, and it is fairly obvious that the convergence is exponential. Since $\pi^\circ \in (0, 1)$, it is clear that φ_t is bounded away from both 0 and 1 . Since φ_t is a

conditional probability, and since the density function for μ is bounded and bounded away from 0, φ_t must split the revealed interval into fractions which are bounded away from 0 and 1, so the revealed interval decreases exponentially to the single point π° . This case is analogous to the convergence of a BSP to a mixed strategy equilibrium.

If $\pi^\circ \leq 0$, then $\varphi_t \rightarrow 0$ exponentially, but the reason for exponential convergence is more subtle. In this case, φ_t decreases monotonically to zero. Because of our assumption on the density function for μ , the conditional probability $\mu(\{\pi : \pi > \varphi_{t+1}\} | \pi \leq \varphi_t)$ is bounded above by a constant multiple of the conditional Lebesgue measure, $(\varphi_t - \varphi_{t+1}) / (1 + \varphi_t)$, so for some $k > 0$, $\varphi_{t+1} \leq k[(\varphi_t - \varphi_{t+1}) / (1 + \varphi_t)]$. If we write $\varphi_{t+1} = \alpha_t \varphi_t$, it follows that $\alpha_t \leq k / (1 + k) < 1$, so $\varphi_t \rightarrow 0$ exponentially. This case is analogous to the convergence of a BSP to a pure strategy equilibrium.

This example indicates the role of our assumption that μ is sufficiently uniform. In the example, the function $\varphi_t \mapsto \pi^\circ - \varphi_t$ plays the role of the function h , so regularity is satisfied automatically. The proof of the Theorem involves additional complications caused by a multidimensional nonlinear coordinate system and equilibria which may be neither completely mixed nor completely pure, so the two types of asymptotic behavior described above occur simultaneously and interactively.

Proof of the Theorem: For each t , let $\sigma_{t-1} = \varphi(\cdot | s^{t-1})$, and, to conserve notation, suppose that the subsequence $\{\sigma_{t(u)}\}$ is the sequence $\{\sigma_t\}_t$ itself. We will now apply the technical results contained in the Appendix. Let $x^\circ = \sigma^\circ$, and for each player p and each t , let $C_{pt} = C_p(s^{t-1})$, as defined in Proposition 3.5, and let $x_{pt} = \varphi_{-p}(\cdot | s^{t-1})$. We will use a subscript p to denote an object derived from C_{pt} and x_{pt} in the same fashion as the respective object is derived from C_t and x_t in the Appendix; e.g., $\delta_{pt}(\cdot)$ is derived from C_{pt} and x_{pt} just as $\delta_t(\cdot)$ is derived from C_t and x_t in the Appendix.

Given p , we have to verify that the sequences $\{C_{pt}, x_{pt}\}$ satisfy assumptions A.4 (i - v). We will simply assume, until the last step of the proof, that $x_{pt} \neq x^\circ$ for all

t. Proposition 3.5 implies that the sequence $\{C_{pt}, x_{pt}\}$ satisfies assumptions A.4 (i) and (ii). As in the definition of regularity, index S_p so that for some $L_p \leq \#S_p$, $\{i : \sigma_p^\circ(s_{ip}) > 0\} = \{1, \dots, L_p\}$. Then, for t sufficiently large, $\sigma_{ipt} = \varphi_p(s_{ip}|s^t) > 0$ for all $i \leq L_p$. Since $\mu_p \sim m$, Proposition 3.5 implies that A.4 (iii) is satisfied for large t . Assumption A.4 (iv) is implied by Proposition 3.5 also. Since $\pi_p^\circ \in \cap_t C_{pt}$, $\delta_{pt}(i, j) \geq (\pi_{ip}^\circ - \pi_{jp}^\circ)x_{pt}$, so regularity condition (b) implies Assumption A.4(v).

For each p and each t , let

$$h_{pt} = \begin{cases} ((\sigma_{ipt})_{\{i > 1\}}, (1 - \sum_{i=1}^{\#S_p} \sigma_{ipt})) & \text{if } L_p = 1; \text{ and} \\ (g_p(\sigma_t), (\sigma_{ipt})_{\{i > L_p\}}, (1 - \sum_{i=1}^{\#S_p} \sigma_{ipt})) & \text{if } L_p > 1. \end{cases}$$

Then $(h_{pt})_{p=1}^n$ is simply a permutation of the coordinates of $h(\sigma_t)$. Also $\sum_{i=1}^{\#S_p} \sigma_{ipt} = 1$ for all p, t , so the last coordinate of h_{pt} is always zero. By Lemma A.8, there is some k'_p such that $\sigma_{ipt} \leq k'_p \delta_{pt}(i, j(t))$ for all i , and all i, p with $i > L_p$. For any p, t and any $1 \leq i, j \leq L_p$, $(\pi_{pi}^\circ - \pi_{pj}^\circ)x_{pt} \leq \delta_{pt}(i, j)$, and since $\pi_p^\circ \in C_{pt+1}$, $(\pi_{pj(t)}^\circ - \pi_{pi}^\circ)x_{pt} \geq 0$ for all $i \leq \#S_p$. Hence, for all $1 \leq i, j \leq L_p$, $(\pi_{pi}^\circ - \pi_{pj}^\circ)x_{pt} \leq \delta_{pt}(j(t), j)$. Also, by Lemma A.5(4), there is some k''_p such that $\delta_{pt}(j(t), j) \leq k''_p \delta_{pt}(j, j(t))$ for all $j \leq L_p, j \neq j(t)$. Therefore $\|g_p(\sigma_t)\| \leq k''_p \max\{\delta_{pt}(i, j(t)) : i \leq L_p\}$ if $L_p > 1$. Hence $\|h_{pt}\| \leq k_p \delta_{pt}^*$, where $k_p = \max\{k'_p, k''_p\}$. For each t , let $\delta_t^{**} = \max\{k_p \delta_{pt}^* : 1 \leq p \leq n\}$. Then $\|h_{pt}\| \leq \delta_t^{**}$ for all t . Since $\sigma_t \rightarrow \sigma^\circ$ and σ° is a regular equilibrium for π° , there is some $c > 0$ such that $\|\sigma_t - \sigma^\circ\| \leq c\delta_t^{**}$ for all t . By the definition of x_{pt} , we can choose c so that $\|x_{pt} - x^\circ\| \leq c\delta_t^{**}$ for all p, t as well.

We will prove the Theorem by showing that $\delta_t^{**} \rightarrow 0$ exponentially. Since $\|x_{pt} - x^\circ\| \leq c\delta_t^{**}$ for all p, t , Lemma A.12 implies that there are positive constants T, K such that for all $u > t > T$, $\delta_u^{**} < K\delta_t^{**}$. Then we can choose K sufficiently large so that we can set $T = 1$. We will now prove the following:

Claim: Either $\delta_t^{**} \rightarrow 0$ exponentially, or there is a subsequence indexed by $\{t(v)\}_{v=1}^\infty$ such that for some p , some $i \leq \#S_p, j \leq L_p$, and some $d_p \in D_p$,

- i) $\delta_{i(v)}^{**} = \delta_{pt(v)}(i, j)$ for all v ;
- ii) $d_{pt} \rightarrow d_p$; and
- iii) $\limsup \delta_{i(v+1)}^{**} / \delta_{i(v)}^{**} \geq 1$.

To prove the Claim, fix $\varepsilon > 0$ and $0 < r < 1$. We say that u *succeeds* t if $u > t$ and

- i') $\delta_t^{**} = \delta_{pt}(i, j)$ and $\delta_u^{**} = \delta_{pu}(i, j)$ for some $i, j \leq L_p$, and $i \leq \#S_p$; and
- ii') $\|d_{pt} - d_{pu}\| < \varepsilon$.

We will show that either $\delta_t^{**} \rightarrow 0$ exponentially or there exist u and t such that u succeeds t and

- iii') $\delta_u^{**} / \delta_t^{**} > r$.

For each p , the unit sphere in which D_p lies can be covered with a finite number of open balls of radius $\varepsilon/2$, so there is some number b such that for any $t^\circ > 0$, we can find a set of indices $\{t_k\}_{k=1}^{k^\circ}$, with t_k increasing in k and $t_{k^\circ} \leq t^\circ$, such that t_{k+1} succeeds t_k for each k and $k^\circ \geq [t^\circ/b]$, where $[z]$ denotes the largest integer less than z . Suppose that whenever u succeeds t , $\delta_u^{**} / \delta_t^{**} \leq r$. Then $\delta_{t_{k^\circ}}^{**} \leq \delta_{t_1}^{**} r^{[t^\circ/b]-1} < K \delta_1^{**} r^{[t^\circ/b]-1}$, and for any $t > t_{k^\circ}$, $\delta_t^{**} < K^2 \delta_1^{**} r^{[t^\circ/b]-1}$. Choose $K' > K^2 \delta_1^{**} / r$ and $r^{1/b} < r' < 1$. Then $K'(r')^t > K^2 \delta_1^{**} r^{[t/b]-1}$ for large t . Hence $\delta_t^{**} < K'(r')^t$ for large t . Hence δ_t^{**} converges exponentially. Therefore, if δ_t^{**} does not converge exponentially, we can find t and u satisfying (i' - iii'). Letting $\varepsilon \rightarrow 0$ and $r \rightarrow 1$ provides a subsequence satisfying (i - iii), thus proving the claim.

Lemma A.9 implies that any subsequence satisfying (i) and (ii) must violate (iii), so $\delta_t^{**} \rightarrow 0$ exponentially. This proves the Theorem under the assumption that $x_{pt} \neq x^\circ$ for all p and all t .

It only remains to consider the case in which for some p, t , $x_{pt} = x_p^\circ$. Of course,

to prove the Theorem it suffices to consider only large values of t , so it suffices to consider the case in which $x_{pt} = x^\circ$ infinitely often. By Proposition 3.5(iii), if for some $u > t$, $x_{pu} = x_{pt}$, then $\varphi_p(s_{pt}|s^{u-1}) = 1$ (loosely speaking, player p revealed in period t that he plays s_{pt} against $x_{pu} = x_{pt}$). Therefore, if for some $v > u > t$, $x_{pv} = x_{pu} = x_{pt}$, then $x_{qv} = x_{qu}$ for every player q , and $\sigma_v = \sigma_u$ is a pure strategy Nash equilibrium for every game $(\pi_q)_q$ with $((\pi_q)_q, s^{u-1}) \in \text{supp } \varphi$. In particular, if $x_{pt} = x_p^\circ$ infinitely often, then $\sigma_t = \sigma^\circ$ infinitely often and σ° is a pure strategy Nash equilibrium. Suppose by way of contradiction that, in addition, $\sigma_t \neq \sigma^\circ$ infinitely often, since the conclusion holds trivially if $\sigma_t = \sigma^\circ$ for all large t . Then there is some p with $x_{pt} \neq x_p^\circ$ and $\varphi_p(\cdot|s^{t-1}) \neq \sigma_p^\circ$ infinitely often. Let $j \leq L_p$ denote the strategy with $\sigma_p^\circ(s_{jp}) = 1$. Since σ° is a regular pure strategy Nash equilibrium for $(\pi_q^\circ)_q$ and $\sigma_t \rightarrow \sigma^\circ$, regularity condition 2.3(b) implies that $s_{pt} = s_{jp}$ for all large t . Also, since $x_{pt} = x^\circ$ infinitely often, it follows that for all large t , $(\pi_{ip} - \pi_{jp})x_p^\circ \leq 0$ for all i and all $\pi_p \in C_{pt}$. Thus, along a subsequence with $x_{pt} \neq x_p^\circ$, the hypothesis of Lemma A.13 is satisfied, so $\delta_{pt}^*/\|x_{pt} - x_p^\circ\| \rightarrow 0$. Therefore, along any subsequence with $\sigma_t \neq \sigma^\circ$, $\delta_{pt}^*/\|\sigma_t - \sigma^\circ\| \rightarrow 0$ for all p . Then, by Lemma A.8, along such a subsequence, $\|\sigma_t - \sigma^\circ\| / \|\sigma_t - \sigma^\circ\| \rightarrow 0$. This contradiction proves that if, for some p , $x_{pt} = x_p^\circ$ infinitely often, then $\sigma_t = \sigma^\circ$ for all large t . ■

Appendix

This appendix contains most of the proof of the Theorem. The definitions and results below concern a decreasing sequence of compact convex sets C_t which result from the successive partitioning of B by hyperplanes determined by vectors x_t . This appendix is self-contained, and makes no reference to game theory. However, when these results are applied to prove the Theorem, the set C_t will be the set of payoff types for some player p which are associated with a history s^{t-1} , and x_t will be player p 's expectations $\varphi_{-p}(\cdot|s^{t-1})$. Player p 's strategy set will be split into strategies $i \leq L$ which have positive probability in the limit, and strategies $i > L$ which have zero probability in the limit. The results below focus on the relation between the asymptotic behavior of the distribution $\varphi_p(\cdot|s^{t-1})$ of player p 's responses, and the rate of convergence of the expectations $\varphi_{-p}(\cdot|s^{t-1})$, although the appropriate Lebesgue measure is used as a proxy for $\varphi_p(\cdot|s^{t-1})$.

It should be noted that all of the assumptions and results on the sequences C_t and x_t apply equally to subsequences.

A.1 Definitions: Let $M, N \geq 2$ and let B denote the unit ball in \mathbb{R}^{MN} . A generic element of B is written $\pi = (\pi_j)_j$, where $\pi_j \in \mathbb{R}^N$ for each $1 \leq j \leq M$. Let m denote the Lebesgue measure on \mathbb{R}^{MN} .

Remarks: An $M \times N$ matrix π represents the player's payoff matrix and each π_j represents the row associated with strategy j . The following lemma establishes an important relation between the expected payoff differences between strategies and the relative Lebesgue measure of payoff matrices, or "types", which prefer given strategies.

A.2 Lemma: Let C be a compact convex subset of B with $\text{int}C$ nonempty, and let $x \in \mathbb{R}^N$ with $x \neq 0$. For each $i, j \leq M$, let $\delta(i, j) = \max\{0, (\pi_i - \pi_j)x : \pi \in C\}$, and let $C_j = \{\pi \in C : (\pi_j - \pi_i)x > 0 \text{ for all } i \neq j\}$. Given i, j with $i \neq j$, define

$C_{ij}^+ = \{\pi \in C : (\pi_i - \pi_j)x > 0\}$ and $C_{ij}^- = \{\pi \in C : (\pi_i - \pi_j)x < 0\}$. Assume that $c_{ij}^+ \neq \emptyset$ and $c_{ij}^- \neq \emptyset$. Then

- 1) $\delta(i, j)^{MN} / \{[\delta(j, i) + \delta(i, j)]^{MN} - \delta(i, j)^{MN}\} \leq m(C_{ij}^+) / m(C_{ij}^-)$; and
- 2) $m(C_i) / m(C) \leq m(C_{ij}^+) / m(C_{ij}^-) \leq m(C) / m(C_j)$.

Proof: Since $C_j \subset C_{ij}^-$ and $C_i \subset C_{ij}^+$, (2) is clear. Let $C^0 = \{\pi \in C : (\pi_i - \pi_j)x = 0\}$. Let $\pi^\circ \in C_{ij}^+$ with $(\pi_i^\circ - \pi_j^\circ)x = \delta(i, j)$, and let $K^+ = \{\lambda\pi + (1 - \lambda)\pi^\circ : 0 \leq \lambda < 1, \text{ and } \pi \in C^0\}$. Let $K = \{\lambda\pi + (1 - \lambda)\pi^\circ : 0 \leq \lambda < 1 + [\delta(j, i) / \delta(i, j)] \text{ and } \pi \in C^0\}$. Then $K^+ \subset C_{ij}^+$. For any $\pi' \in C_{ij}^-$, $0 > (\pi'_i - \pi'_j)x \geq -\delta(j, i)$ so $\gamma\pi^\circ + (1 - \gamma)\pi' \in C^0$ for some $0 < \gamma \leq \delta(j, i) / [\delta(j, i) + \delta(i, j)]$, which implies that $\pi' \in K$. Therefore $C_{ij}^- \subset K \setminus K^+$, so $m(C_{ij}^+) / m(C_{ij}^-) \geq m(K^+) / m(K \setminus K^+)$. Let m^0 denote Lebesgue measure on the $(MN - 1)$ -dimensional space $\{\pi \in \mathbb{R}^{MN} : (\pi_i - \pi_j)x = 0\}$, and let $a = \max\{(\pi_i - \pi_j)x : \pi \in B\}$. Then we have the following volume measurements:

$$\begin{aligned} m(K^+) &= \delta(i, j)m^0(C^0) / MN; \text{ and} \\ m(K) &= \{([\delta(i, j) + \delta(j, i)]^{MN} / \delta(i, j)^{(MN-1)})m^0(C^0) / aMN\}; \text{ so} \\ m(K^+) / m(K \setminus K^+) &= \delta(i, j)^{MN} / \{[\delta(j, i) + \delta(i, j)]^{MN} - \delta(i, j)^{MN}\}, \end{aligned}$$

which proves (1). ■

A.3 Definitions: Let $x^\circ \in \mathbb{R}_+^N$ with $x^\circ \neq 0$, and let $\{x_t\}_{t=1}^\infty$ be a sequence in \mathbb{R}_+^N converging to x° , with $x_t \neq x^\circ$ for all t . For each t , let C_t be a compact convex subset of B , and for each i, j with $i \neq j$, let $\delta_t(i, j) = \max\{0, (\pi_i - \pi_j)x_t : \pi \in C_t\}$. For each i , let $C_{it} = \{\pi \in C_t : (\pi_i - \pi_j)x_t > 0 \text{ for all } j \neq i\}$. Let $L \leq M$, and let $\{j(t)\}_{t=1}^\infty$ be a sequence in $\{1, \dots, L\}$.

Remarks: The vector x_t represents the player's expectation of the probability distribution over the $(n - 1)$ -tuples of strategies to be played by the other players in period t . The set C_t represents the set of payoff types consistent with the player's previous strategy choices. The strategy $j(t)$ is the strategy which the player actually chooses in

period t . The strategies $1, \dots, L$ are the pure strategies which have positive probability in the player's limiting mixed strategy, and the pure strategies $L + 1, \dots, M$ are inferior in the limit. This interpretation is formalized by the following assumptions.

A.4 Assumptions: We assume that for each t ,

- (i) $C_{t+1} \subset C_t$;
- (ii) $\text{int } C_t$ is nonempty and $\text{int } C_{t+1} \subset C_{j(t)t}$;
- (iii) for each $i \leq L$, $m(C_{it}) > 0$;

and assume that

- (iv) for each $i \leq L$, $\liminf m(C_{it})/m(C_t) > 0$; and
- (v) for each $i \leq L$ and each $j > L$, $\liminf \delta_t(i, j) > 0$.

Remarks: In the proof of Theorem 3.6, the results of this Appendix will be applied to subsequences of expectations. Hence, in this Appendix, $t+1$ should be broadly interpreted as some period subsequent to t . Thus $\text{int } C_{t+1}$ may be only a small subset of $C_{j(t)t}$ because intervening strategy choices may further reduce the set of payoff types which are consistent with all choices made prior to period $t + 1$.

A.5 Lemma: Assumptions (i - iv) imply

- 3) $\delta_t(i, j) > 0$ for all t , all $i \leq L$, and all j ;
- 4) $\liminf \delta_t(j, i)/\delta_t(i, j) > 0$ for all $i, j \leq L$; and
- 5) $\lim \delta_t(i, j) = 0$ for all $j \leq L$ and all i .

Proof: Assertion (3) follows directly from Assumption (iii). Let $i, j \leq L$ with $i \neq j$, and for each t , let $r_t = \delta_t(j, i)/\delta_t(i, j)$. By (3), r_t is well-defined, and by (1) and (2)

$$[(1 + r_t)^{MN} - 1]^{-1} \leq m(C_t)/m(C_{jt})$$

for all t . Therefore Assumption (iv) implies that $\liminf r_t > 0$, which establishes (4).

To prove (5), let $j \leq L$ and for any i , let $\pi^t \in C_t$ with $(\pi_i^t - \pi_j^t)x_t = \delta_t(i, j)$ for each t . We need to show that $(\pi_i^t - \pi_j^t)x_t \rightarrow 0$. Taking a subsequence, if necessary, we can assume that $\pi^t \rightarrow \pi$ for some $\pi \in \bigcap_t C_t$, so it suffices to show that $(\pi_i - \pi_j)x^\circ = 0$. For each t , $(\pi_{j(t)} - \pi_i)x_t \geq 0$, so it suffices to prove the preceding equality for i such that $i = j(t)$ infinitely often. If it were the case that $(\pi_i - \pi_j)x^\circ > 0$, then (4) would imply the existence of some $\pi' \in \bigcap_t C_t$ with $(\pi'_i - \pi'_j)x^\circ < 0$. However, since $i = j(t)$ infinitely often, $(\pi'_i - \pi'_j)x_t \geq 0$ infinitely often, so $(\pi'_i - \pi'_j)x^\circ = 0$. ■

A.6 Definition: For each t , let $d_t = (x_t - x^\circ)/\|x_t - x^\circ\|$, and let D denote the set of cluster points of the sequence $\{d_t\}$.

A.7 Lemma: Assumptions (i - iv) imply that

$$6) \liminf \{\min\{(\pi_{j(t)} - \pi_i)d_t : \pi \in C_t\}\} \geq 0 \text{ for all } i.$$

Proof: Suppose by way of contradiction that for some i , (6) is false. Taking a subsequence, if necessary, we can assume that $j(t) = j \leq L$ for all t , and that $d_t \rightarrow d \in D$. Then since $\{C_t\}_t$ is a decreasing sequence of compact sets, there exists $\pi \in \bigcap_t C_t$ such that $(\pi_j - \pi_i)d < 0$. By (5), $(\pi_j - \pi_i)x^\circ \geq 0$. Therefore, since $d_t \rightarrow d$, we can find t arbitrarily large such that $(\pi_j - \pi_i)x_t < 0$. However, Assumption (ii), together with the convexity of each C_{t+1} , implies that $(\pi_j - \pi_i)x_t \geq 0$ for all t . ■

A.8 Lemma: There is some constant k such that for each $i > L$, $m(C_{it})/m(C_t) < k\delta_t(i, j(t))$ for all t .

Proof: Let $i > L$. By (1) and (2),

$$(8.1) \quad [(\delta_t(j(t), i) + \delta_t(i, j(t)))^{MN} - \delta_t(j(t), i)^{MN}] / \delta_t(j(t), i)^{MN} \geq m(C_{it})/m(C_t)$$

for all t . By Assumption (v), $\liminf \delta_t(j(t), i) > 0$, and by (5), $\delta_t(i, j(t)) \rightarrow 0$. Since

the left hand side of (8.1) is differentiable, the result follows. ■

A.9 Lemma: Suppose that

- a) $j(t) = j \leq L$ for all t ;
- b) $d_t \rightarrow d \in D$; and
- c) for some i , $\liminf \delta_t(i, j) / \|x_t - x^\circ\| > 0$.

Then $x_t \rightarrow x^\circ$ exponentially and $\delta_t(i, j) \rightarrow 0$ exponentially.

Proof: For each t , let $\pi^t \in C_t$ satisfy $(\pi_i^t - \pi_j^t)x_t = \delta_t(i, j)$. Then for each t

$$(9.1) \quad (\pi_i^t - \pi_j^t)x_t = (\pi_i^t - \pi_j^t)x^\circ + \|x_t - x^\circ\|(\pi_i^t - \pi_j^t)d + (\pi_i^t - \pi_j^t)z_t$$

where $\|z_t\| / \|x_t - x^\circ\| \rightarrow 0$ by (b). Since $\pi^t \in C_t \subset C_{j_{t-1}}$, (a) implies that $(\pi_i^t - \pi_j^t)x_{t-1} \leq 0$, so

$$(9.2) \quad (\pi_i^t - \pi_j^t)x^\circ + \|x_{t-1} - x^\circ\|(\pi_i^t - \pi_j^t)d + (\pi_i^t - \pi_j^t)z_{t-1} \leq 0.$$

Hence, by (c), there is some $k > 0$ with

$$(9.3) \quad (\pi_i^t - \pi_j^t)x^\circ / \|x_t - x^\circ\| > k + (\pi_i^t - \pi_j^t)x^\circ / \|x_{t-1} - x^\circ\|$$

for all t sufficiently large. By Lemma A.7, $\limsup (\pi_i^t - \pi_j^t)d \leq 0$ so (c) implies that $(\pi_i^t - \pi_j^t)x^\circ > 0$ for large t . Therefore (9.3) can be written

$$(9.4) \quad \|x_{t-1} - x^\circ\| / \|x_t - x^\circ\| > 1 + k \|x_{t-1} - x^\circ\| / (\pi_i^t - \pi_j^t)x^\circ.$$

By (9.2), $(\pi_i^t - \pi_j^t)x^\circ / \|x_{t-1} - x^\circ\|$ is bounded above, so (9.4) implies that $\liminf \|x_{t-1} - x^\circ\| / \|x_t - x^\circ\| > 1$, so $x_t \rightarrow x^\circ$ exponentially. Since $(\pi_i^t - \pi_j^t)x^\circ / \|x_{t-1} - x^\circ\|$ is bounded above and $x_t \rightarrow x^\circ$ exponentially, (9.1) implies that $\delta_t(i, j) \rightarrow 0$ exponentially. ■

A.10 Lemma: For any $\varepsilon > 0$, there exists $T > 0$ such that for all u, t with $u > t > T$ and all i ,

$$\delta_u(i, j(u)) < \delta_t(j(t), j(u)) + 2\|x_t - x^\circ\| + \varepsilon\|x_u - x^\circ\|.$$

Proof: By Lemma A.7, we can choose T sufficiently large that for all $\pi \in C_u$ and all $i > L$, we have $(\pi_i - \pi_{j(u)})d_u < \varepsilon$ if $u > T$. Then

$$\begin{aligned} (\pi_i - \pi_{j(u)})x_u &= (\pi_i - \pi_{j(u)})x^\circ + \|x_u - x^\circ\|(\pi_i - \pi_{j(u)})d_u \\ &< (\pi_i - \pi_{j(u)})x^\circ + \varepsilon\|x_u - x^\circ\|. \end{aligned}$$

Since $(\pi_i - \pi_{j(u)})x^\circ < (\pi_i - \pi_{j(u)})x_t + 2\|x_t - x^\circ\|$, and

$$(\pi_i - \pi_{j(u)})x_t \leq (\pi_i - \pi_{j(t)})x_t + \delta_t(j(t), j(u)) \leq \delta_t(j(t), j(u)),$$

the result follows. ■

A.11 Definition: For each t , let $\delta_t^* = \max\{\delta_t(i, j(t)) : i \leq N\}$, and let $\Delta_t = \max\{\delta_t^*, \|x_t - x^\circ\|\}$.

A.12 Lemma: For any $c > 0$ there is some $T > 0$ such that for any u, t with $u > t > T$ and $\|x_u - x^\circ\| < c\delta_u^*$, $\Delta_u < 4(1+c)\Delta_t$.

Proof: Let $c > 0$. By Lemma A.10, there is some $T > 0$ such that for all i and all $u > t > T$,

$$(12.1) \quad \delta_u(i, j(u)) < \delta_t(j(t), j(u)) + 2\|x_t - x^\circ\| + \|x_u - x^\circ\|/4c.$$

Therefore, if $c\delta_u(i, j(u)) > \|x_u - x^\circ\|$, then $\|x_u - x^\circ\|/c < 3\Delta_t + \|x_u - x^\circ\|/4c$, so $\|x_u - x^\circ\| < 4c\Delta_t$. Then by (12.1) $\delta_u(i, j(u)) < 4\Delta_t$. Since $\Delta_u < \delta_u^* + \|x_u - x^\circ\|$, the result follows. ■

A.13 Lemma: Suppose that for some $j \leq L$, $j(t) = j$ for all t , and for some t , $(\pi_i - \pi_j)x^\circ \leq 0$ for all i and all $\pi \in C_t$. Then

$$\delta_t^*/\|x_t - x^\circ\| \rightarrow 0.$$

Proof: Fix $i \leq M$, and a subsequence, if necessary, so that $\delta_t(i, j)/\|x_t - x^\circ\|$ converges to its limsup, and d_t converges to some $d \in D$. Then for each t we can write

$$x_t = x^\circ + \|x_t - x^\circ\|d + z_t$$

where $\|z_t\|/\|x_t - x^\circ\| \rightarrow 0$. Since $C_u \subset C_t$ for all $u > t$, $(\pi_i - \pi_j)x^\circ \leq 0$ for all $\pi \in C_t$ for all large t . Let $\varepsilon > 0$, and suppose that $u > t$ and that u and t are sufficiently large that $\|z_t\|/\|x_t - x^\circ\| < \varepsilon$, $\|z_u\|/\|x_u - x^\circ\| < \varepsilon$, and $(\pi_i - \pi_j)x^\circ \leq 0$ for all $\pi \in C_u$. Let $\pi \in C_u$ with $\delta_u(i, j) = (\pi_i - \pi_j)x_u > 0$. Then we have

$$(13.1) \quad \delta_u(i, j) = (\pi_i - \pi_j)x^\circ + \|x_u - x^\circ\|(\pi_i - \pi_j)d + (\pi_i - \pi_j)z_u.$$

Since $C_u \subset C_{t+1} \subset C_t$ and $j = j(t)$, we also have

$$(13.2) \quad (\pi_i - \pi_j)x^\circ + \|x_t - x^\circ\|(\pi_i - \pi_j)d + (\pi_i - \pi_j)z_t \leq 0.$$

Then (13.2) implies that

$$(13.3) \quad (\pi_i - \pi_j)x^\circ + \|x_t - x^\circ\|(\pi_i - \pi_j)d \leq 2\varepsilon\|x_t - x^\circ\|$$

because $\|z_t\|/\|x_t - x^\circ\| < \varepsilon$ and $\pi \in B$. We can assume that u is sufficiently larger than t that $\|x_u - x^\circ\| \leq \|x_t - x^\circ\|$. Then, since $(\pi_i - \pi_j)x^\circ \leq 0$, it follows from (13.1) that

$$(13.4) \quad \delta_u(i, j) \leq (\|x_u - x^\circ\|/\|x_t - x^\circ\|) \{(\pi_i - \pi_j)x^\circ + \|x_t - x^\circ\|(\pi_i - \pi_j)d\} \\ + 2\varepsilon\|x_u - x^\circ\|.$$

Finally, (13.4) and (13.3) imply

$$(13.5) \quad \delta_u(i, j) \leq 4\varepsilon\|x_u - x^\circ\|.$$

Since i and ε were chosen arbitrarily, the result follows. ■

References

- Fudenberg, Drew and David Kreps (1988), "A Theory of Learning, Experimentation, and Equilibrium in Games", Mimeo.
- Jordan, J. (1989), "Bayesian Learning in Normal Form Games", *Games and Economic Behavior*, Forthcoming.
- Milgrom, P. and J. Roberts (1989), "Adaptive and Sophisticated Learning in Repeated Normal Form Games", Mimeo.
- Robinson, J. (1951), "An Iterative Method of Solving a Game", *Annals of Mathematics*, 54, 296-301.