



Limit Games and Limit Equilibria

Citation

Fudenberg, Drew, and David Levine. 1986. Limit games and limit equilibria. *Journal of Economic Theory* 38, no. 2: 261-279.

Published Version

[http://dx.doi.org/10.1016/0022-0531\(86\)90118-3](http://dx.doi.org/10.1016/0022-0531(86)90118-3)

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3350443>

Terms of Use

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

Share Your Story

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

[Accessibility](#)

LIMIT GAMES AND LIMIT EQUILIBRIA

by

Drew Fudenberg and David Levine

Departments of Economics

UC Berkeley and UCLA respectively

UCLA Department of Economics
Working Paper #289
April 1983

Abstract

We provide a sufficient condition for equilibria of a game to arise as limits of ϵ -equilibria of games with smaller strategy spaces. As the smaller games are frequently more tractable, our result facilitates the characterization of the set of equilibria.

We would like to thank Ed Green, Jean-Charles Rochet, Sylvain Sorain, and participants at the IMSSS Summer Workshop on Oligopoly. We are particularly indebted to Bob Anderson who encouraged us to find a more general formulation. Financial support from the NSF is gratefully acknowledged.

1. Introduction

Although game-theoretic models play an important role in economic theory, in many cases of interest it is difficult to characterize the set of noncooperative equilibria. We provide, as a tool for this purpose, a sufficient condition for equilibria to arise as the limits of ϵ -equilibria in games with smaller (and more tractable) strategy spaces. We then extend our result to mixed-strategy equilibrium. As illustrations we consider finite-horizon approximations of infinite-horizon games and discrete-time approximations of continuous-time games. We considered the first application in Fudenberg-Levine [1982]; the framework we use here allows a clearer proof and a weaker continuity requirement.

The idea of the theorem is straightforward: if a sequence of "restricted games" approximates the game of interest in the appropriate sense then any convergent sequence of ϵ -equilibria of the restricted games with $\epsilon \rightarrow 0$ as the approximation improves converges to an equilibrium of the original game, and every equilibrium of the underlying game is a limit of ϵ -equilibria of the restricted games with $\epsilon \rightarrow 0$. Thus the set of equilibria can be characterized by computing the set of limit points of the ϵ -equilibria of the restricted games.

Section 2 briefly reviews some related work. Section 3 gives some basic definitions and gives a sufficient condition for equilibrium. Section 4 takes a more topological approach and provides a necessary condition, as well. Section 5 extends our results to mixed strategies. Section 6 briefly reviews the finite to infinite horizon limit. Section 7 treats discrete-time to continuous time, and Section 8 analyzes some games of timing. Section 9 concludes.

2. Related Work

The work most similar to ours is that concerned with the existence of minimax solutions for two-player zero-sum games¹ (e.g., Wald [1947], Tjoe-Tie [1963]). Wald defined an "intrinsic metric" on the strategy space and proved that if it is compact in the intrinsic topology then the game has a minimax value in mixed strategies. Our goal of characterizing all the equilibria is irrelevant in two-person zero-sum games because when the minimax value exists it is unique. Thus this literature has been concerned with only "one direction" of our limit theorem, and then only for the case in which the strategy space is compact so that existence of the restricted equilibria implies existence in the limit.

A number of authors have observed that some, but not all, equilibria of continuous-time games are limits of equilibria of discrete-time games, among them Kreps-Wilson [1982] and Stokey [1981]. Dasgupta-Maskin [1982] use finite-action approximations of games with a continuum of actions to investigate the existence of mixed-strategy equilibrium.

Green [1982] considers finite-player approximations of games with many players. He focuses on the upper hemi-continuity of the equilibrium correspondence.

Radner [1980] showed that cooperation could arise as an ϵ -equilibrium of the finitely-repeated Prisoner's dilemma. As cooperation is an equilibrium with an infinite horizon, Radner's result corresponds to the second direction of our theorem.

3. Basic Notions

We begin by introducing the basic concepts and definitions. Players are elements i of a finite set I . Each individual player has a strategy

space S_i . The overall strategy space is the Cartesian product $S \equiv \prod_{i \in I} S_i$; elements of S will be called strategy combinations. Notice that S may be a space of mixed strategies. The strategy combination derived from S by replacing its i^{th} component by h_i is denoted by (h_i, s_{-i}) . Player i 's payoff is a bounded function $\pi^i: S \rightarrow \mathbb{R}$.

Definition (3.1): A strategy combination s is an ϵ -equilibrium if for all i and $h_i \in S_i$

$$\pi^i(h_i, s_{-i}) < \pi^i(s) + \epsilon.$$

Thus each player gets within ϵ of the maximum. (The notion of an epsilon equilibrium is due to Radner [1980]). One rationale for ϵ -equilibrium as a solution concept is that if players have sufficient inertia they will not bother to realize small gains. When $\epsilon = 0$ we refer simply to equilibria: this is the usual noncooperative Nash equilibrium of a game in strategic (normal) form. We are not primarily interested in ϵ -equilibria themselves; we will use them as a tool to characterize the equilibria of S .

We are particularly interested in games in which players are restricted to a proper subset of the strategies available to them in S .

Definition (3.2): R is a restriction or restricted game if $R = \prod_{i \in I} R_i$ with $R_i \subseteq S_i$.

If R is a restriction of S we have the notion of an ϵ -equilibrium relative to R .

Definition (3.3): A strategy combination $r \in R$ is an ϵ -equilibrium relative to R if for all i and $h_i \in R_i$

$$\pi^i(h_i, r_{-i}) < \pi^i(r) + \epsilon.$$

Thus each player gets within ϵ of the best he can do within his restricted strategy space R_i .

Now we consider a sequence of strategy spaces $\{R^n\}$. If ϵ -equilibrium in R^n are to be related to ϵ -equilibria in S it must be possible to do nearly as well in R^n as in S . Thus we are led to define

Definition (3.4): $\{R^n\}$ is thick in S if for any $\delta > 0$, $s \in S$, $i \in I$ and large enough n there is a $\tilde{s}_i^n \in R^n$ with

$$\pi^i(s) < \pi^i(\tilde{s}_i^n, s_{-i}) + \delta.$$

Next suppose that r^n are ϵ^n -equilibria in R^n with $\epsilon^n \rightarrow 0$, and that s is in S . Suppose further that the r^n approximate s in the sense that each agent can do only slightly better by deviating against r^n rather than against s :

Definition (3.5): r^n approaches s if for all $\delta > 0$ and large enough n

$$\sup_{\tilde{s}_i^n \in R_i^n} \pi^i(\tilde{s}_i^n, s_{-i}) - \pi^i(s) - [\pi^i(\tilde{s}_i^n, r_{-i}^n) - \pi^i(r^n)] < \delta.$$

The significance of thickness is

Theorem (3.1): If $\{R^n\}$ is thick in S and r^n approaches s where r^n are ϵ^n -equilibria in R^n with $\epsilon^n \rightarrow \epsilon$ then s is an ϵ -equilibrium in S .

Proof: If not for some i there is an \bar{s}_i and a $\delta > 0$ such that

$$\pi^i(\bar{s}_i, s_{-i}) = \pi^i(s) + \epsilon + \delta$$

By thickness for large enough n we can actually find an $\tilde{s}_i^n \in R^n$ that does almost as well, so that

$$\pi^i(\tilde{s}_i^n, s_{-i}) > \pi^i(s) + \epsilon + \delta/2.$$

Also since r^n approaches s for large enough n

$$\pi^i(\tilde{s}_i^n, r_{-i}^n) > \pi^i(r) + \epsilon + \delta/4.$$

which is a contradiction since $\epsilon^n \rightarrow \epsilon$.

Q.E.D.

In particular, if $\epsilon^n \rightarrow 0$ (or if $\epsilon^n \equiv 0$) then s is an equilibrium.

4. A Topological Approach

Although the hypotheses of Theorem (3.1) are easily checked it does not prove that all ϵ -equilibria are reached as limits of ϵ^n -equilibria of restricted games. To prove this we need a stronger notion of the convergence of strategies.

Definition (4.1): The "distance" between two strategies $s, t \in S$ is

$$d(s, t) \equiv \sup_{i, h_i \in S_i} (|\pi^i(s) - \pi^i(t)| + |\pi^i(h_i, s_{-i}) - \pi^i(h_i, t_{-i})|).$$

What $d(,)$ says is that two strategy combinations s and t are close if they yield payoffs that are close and any deviation against s yields payoffs close to that of the same deviation against t . It is easy to check that d satisfies the triangle inequality. Thus $d(,)$ is a pseudo-metre.

However, d is not a metric because it is possible that $d(s, t) = 0$ even though s and t are distinct. All statements about convergence, continuity, etc. will be with respect to the topology generated by d , which we will call the "inherent" topology. It follows directly from the definition of $d(,)$ that if s is an ϵ -equilibrium then t is an $\epsilon + d(s, t)$ equilibrium. Thus

Proposition (4.1): If $(s^n) \rightarrow s$ are a sequence of ϵ^n -equilibria and $\epsilon^n \rightarrow \epsilon$ then s is an ϵ -equilibrium. Also if $s^n \rightarrow s$ and s is an ϵ -equilibrium then the s^n are ϵ^n -equilibria with $\epsilon^n \rightarrow \epsilon$.

Remark: $d(,)$ does not induce the coarsest topology in which Proposition (3.1) holds. For example the pseudometric

$$\bar{d}(s, t) = \sup_{i, h_i \in S_i} |\pi^i(h_i, t_{-i}) - \pi^i(t) - (\pi^i(h_i, s_{-i}) - \pi^i(s))|$$

induces the weak topology on S generated by the requirement that each player's best response (in terms of payoffs) be continuous, which is the coarsest topology for which Proposition (4.1) holds. We will call this the "Nash topology". As we shall see, the Nash topology is too weak for the results below.

A topological notion of thickness is

Definition (4.2): The sequence of restricted games $\{R_n\}$ approximates S if for every subsequence $\{R^{n_k}\}$, $\bigcup_{k=1}^{\infty} R^{n_k}$ is dense in S .

We now state and prove our main theorem which relates the ϵ^n -equilibria of an approximating sequence of restricted games to the ϵ -equilibria of the unrestricted game.

Proposition (4.1): [Limit Theorem]: Suppose $\{R^n\}$ approximates S and $r^n \in R^n$.

(A) If the r^n are ϵ^n -equilibria relative to R^n with $\epsilon^n \rightarrow \epsilon$ and $r^n \rightarrow s$ then s is an ϵ -equilibrium.

(B) If s is an ϵ -equilibrium and $r^n \rightarrow s$ then there is a sequence $\epsilon^n \rightarrow \epsilon$ such that r^n is an ϵ^n -equilibrium relative to R^n .

Corollary (4.2): If s is an ϵ -equilibrium there exists sequences $\{r^n\}$, $r^n \in R^n$, and $\{\epsilon^n\}$, with $r^n \rightarrow s$ and $\epsilon^n \rightarrow \epsilon$ such that r^n is an ϵ^n -equilibrium relative to R^n .

Proof:

(A) This is immediate by Theorem (3.1) since approximation implies thickness and convergence implies approaching.

(B) Since $R^n \rightarrow S$ this follows directly from Proposition (4.1).

Finally, the corollary follows from observing that since $\{R^n\}$ approximates S , there is some sequence $r^n \rightarrow s$ and applying (B). Q.E.D.

Note that thickness is not implied by the "Nash topology" of the last section which required only that deviations against nearby strategies yield similar payoffs. It is in this sense that topology is too coarse for our theorem — if (h_1^n, r_{-1}^n) approximates (h_1, s_{-1}) we need to know that the two strategy combinations themselves yield nearby payoffs. In a loose sense this observation suggests a bound on the coarseness of topologies under which Proposition (4.1) holds; we have not however established that the topology we employ is the coarsest possible.

In some cases the actual equilibrium strategies are of less interest than the equilibrium payoffs.

Corollary (4.3): If $\{R^n\}$ approximates S , $\{v^n\} \rightarrow v$ is a sequence of payoffs (one for each player) of ϵ^n -equilibria relative to R^n , and S is compact, then there is an equilibrium in S with payoffs v .

Proof: Take a convergent subsequence of the ϵ^n -equilibria and apply (4.1).

Remark: If we let R^* be the space of all restricted games, endowed with the Hausdorff metric, and define $\Omega: [0, \infty) \times R^* \rightarrow S$ to be the correspondence yielding for any (ϵ, R) the set of ϵ -equilibria relative to R , then

Proposition (4.1) (A) says that Ω is closed-valued at (ϵ, S) for every ϵ . As S is not necessarily compact, we cannot conclude that Ω is upper hemi-continuous (u.h.c.) Green [1982] gives sufficient conditions for the equilibrium correspondence to be u.h.c. in a setting similar to ours.

Remark: One may have a priori notions of a natural topology on S , and prefer to approximate the equilibria of S with restricted equilibria which are near their limits in that natural sense. For example, one might wish nearby strategies to yield nearby "outcomes". An immediate corollary of Proposition (4.1) is that any topology on S that is finer than ours will do. The drawback to using finer topologies is that the restricted games employed became less parsimonious approximations, or conversely that a "reasonable" collection of restricted games is less likely to be sufficiently good approximation.

5. Mixed Strategies

In this section we show that an analog of the limit theorem holds for mixed-strategy equilibrium as well. Of course the space of probability measured over a strategy space can itself be viewed as a strategy space, and so theorem (4.1) can be directly applied. The point of this section is that it suffices to know that the restricted pure strategy spaces approximate S to know that the set of mixed-strategy equilibria coincides with the set of limits of mixed-strategy ϵ -equilibria of the restricted games. This section also relates our work to that of Wald [1947] on the existence of minimax values.

In order to define mixed strategies, that is to introduce probability measures over the S_1 , it will be convenient to topologize the S_1 directly.

Definition (5.1): The distance between two strategies of player i , $s_i, t_i \in S_i$, is

$$\hat{d}(s_i, t_i) = \sup_{h_{-i} \in S_{-i}} \max_{j \in I} |\pi^j(s_i, h_{-i}) - \pi^j(t_i, h_{-i})|.$$

Thus $\hat{d}(s_i, t_i)$ measures the greatest difference it would make to any player if player i used s_i instead of t_i .

Definition (5.2): The inherent product topology on S is the product topology induced when each of the S_i is topologized by \hat{d} .

Remark: The inherent product topology is finer than the inherent topology.

Remark: For a two-person zero-sum game $\hat{d}(s_i, t_i)$ reduces to

$$\sup_{h_{-i} \in S_{-i}} |\pi^1(s_i, h_{-i}) - \pi^1(t_i, h_{-i})|.$$

this is Wald's "intrinsic distance".
All topological statements in this section only will be with respect to the inherent product topology. For this section only we will also make

Assumption (5.1): Each S_i is separable, that is, each S_i has a countable dense subset.

Let B_i be the Borel algebra on S_i , and let \mathcal{S}_i be the space of probability measures on B_i . S is given the product algebra B , and $\mathcal{S} = \times_i \mathcal{S}_i$.² We endow \mathcal{S} with the topology of weak convergence.³ Let $\pi^i(s_1, \dots, s_n)$ denote player i 's expected payoff to an n -tuple of mixed strategies.

$$\pi^i(s_1, \dots, s_n) = \int_{S_1} \dots \int_{S_n} \pi^i(s_1, \dots, s_n) ds_1, \dots, ds_n.$$
⁴

Corresponding to a restricted game R^n S we have the restricted mixed-strategy game \mathcal{R}^n \mathcal{P} created by replacing \mathcal{I}_1 with the space of probability measures over R_1^n .

Proposition (5.1): Suppose $\{R^n\}$ approximates S , each R^n is countable, and $r^n \in R^n$. Then

(A) If the r^n are mixed strategy ϵ^n -equilibria relative to \mathcal{R}^n with $\epsilon^n \rightarrow \epsilon$ and $r^n \rightarrow s \in S$ then s is a mixed strategy ϵ -equilibrium.

(B) If $s \in S$ is a mixed strategy ϵ -equilibrium and $r^n \rightarrow s$ then there is a sequence $\epsilon^n \rightarrow \epsilon$ such that the r^n are mixed strategy ϵ^n -equilibria relative to \mathcal{R}^n .

Proof: See Appendix.

Proposition (5.2): Suppose each S_1 is compact — then S has a (mixed-strategy) equilibrium.

Proof: if S_1 is compact (w.r.t \hat{d}) it is totally bounded.

That is, for any ϵ^n there is a finite set R_1^n of elements of S_1 such that every element of S_1 is within ϵ^n of some element of R_1^n . Clearly the R^n approximate S . Every R^n is finite so it has a mixed-strategy equilibrium r^n . As S is a compact metric space S is compact and we can take a convergent subsequence and apply Proposition (5.1)(A).⁵

As Nash equilibrium payoffs are minimax in two-person zero-sum games, we have

Proposition (5.3) (Wald): Let S be a two-person zero-sum game. If S_1 is compact (w.r.t. \hat{d} .) then S has a minimax value.

6. Finite-Horizon Approximations to Infinite Horizon Games

Let A_1 be a set of actions for player 1. A strategy for player 1 is a sequence of mappings $s_1 = (s_1^1, s_1^2, \dots)$ where $s_1^1 \in A_1$ and for $t > 1$ $s_1^t: (x_{A_1})^{t-1} \rightarrow A_1$. The strategy space for 1, S_1 , is a subset of the space of all such sequences of mappings which may incorporate various restrictions such as if player two placed a_2 last period then player one can't play a_1 . However, we assume that there is a designated "null" action \hat{a}_1 for each player which is always feasible so that if $s_1 \in S_1$ then $(s_1^1, \dots, s_1^t, \hat{a}_1, \hat{a}_1, \dots) \in S_1$.

We approximate S with a collection of finite horizon games. Define $R_1^n \equiv \{s_1 \in S_1 \mid s_1^t = \hat{a}_1 \text{ for } t > n\}$, so that players must play the null action in all periods after n . Clearly R^n will only be a good approximation of S if events after period n are relatively unimportant.

Definition (6.1): π^1 is continuous at infinity if

$$\lim_{T \rightarrow \infty} \sup_{s, \tilde{s} \text{ s.t. } (s^1, \dots, s^T) = (\tilde{s}^1, \dots, \tilde{s}^T)} |\pi^1(s) - \pi^1(\tilde{s})| \rightarrow 0.$$

While continuity at infinity is a strong requirement it is satisfied by many games of interest to economists.

Proposition (6.1): In an infinite-horizon game S , if the π^1 are continuous at infinity then for every subsequence of finite horizon games

$\{R^{n_k}\}$, $\bigcup_{k=1}^{\infty} R^{n_k}$ is dense in S . Thus the limit theorem and its corollaries apply.

Proof: Obvious.

This proposition generalizes our earlier result, which required an additional continuity assumption on the payoffs, because we employed a finer topology.

As an example consider the repeated Prisoner's dilemma with discount factor $0 < \rho < 1$ with two players and action spaces $A_i = \{\text{cheat, don't cheat}\}$ and payoff 1 to both if both cheat; 2 to both if neither cheat and 3 to a player who cheats when his rival doesn't and 0 to the rival. The strategy spaces S_i are unrestricted. The null action $\hat{\alpha}_i$ can be taken to be "cheat". If $\rho > 1/2$ so that it doesn't pay to cheat for a single period it is well known (and trivial to show) that in S the "cooperative" strategy "don't cheat if cheating has never occurred, otherwise cheat forever" is an equilibrium. However, in R^n $n < \infty$ every equilibrium requires that both players cheat in every period. What is true is that the "cooperative" strategy is an ϵ^n -equilibrium with $\epsilon^n \rightarrow 0$; indeed ϵ^n is the cost of failing to cheat when it becomes optimal to do so. In addition the limit of the finite equilibria "cheat no matter what" is an equilibrium in the infinite game.

This example illustrates the content of the limit theorem. Every infinite horizon equilibrium is the limit of ϵ^n -equilibria with $\lim_{n \rightarrow \infty} \epsilon^n \rightarrow 0$; every limit of ϵ^n -equilibria satisfying this condition is an infinite horizon equilibrium. However, as the example clearly shows, there may not be any 0-equilibria in R^n which converge to a given equilibrium in S ; $\epsilon^n > 0$ may be required.

A deeper discussion of the finite to infinite horizon case can be found in Fudenberg/Levine [1982] in which we discuss sequential equilibria with incomplete information, and give applications to the existence and uniqueness of infinite horizon equilibrium.

7. Open-Loop Equilibrium

We now suppose that S_1 are Lebesgue measurable functions on $[0,1]$ into R^m representing a time path of actions. Payoffs $\pi^1 = S + R$ are assumed continuous in the L_1 norm. We define R^n to be functions which are constant between lattice points -- that is, on $[k/n, (k+1)/n]$. Since these functions are dense in the L_1 norm $\{R^n\}$ approximates S and the limit theorem holds.

The space S demonstrates a possible limitation of the limit theorem: S isn't compact and as a result a sequence of restricted equilibria may not converge to anything at all. For example the sequence of restricted strategies

$$(7.1) \quad s^1(k/n) = \begin{matrix} 1 & k \text{ even} \\ -1 & k \text{ odd} \end{matrix}$$

simply does not have a limit. This is the "chattering" problem and can arise even in continuous time control problems. One solution that works for control problems is to define "chattering" controls -- simply define "idealized" or "generalized" functions as the limit of sequences such as (6.1).

Unfortunately this "solution" doesn't appear to work when there is more than one player and the chattering problem continues to pose an unsolved challenge for game theory.

Footnotes

¹We would like to thank Sylvain Sorain for bringing this literature to our attention.

²Because each S_1 is separable $B = XB_1$. (See, e.g., Hildebrand [1974]).

³ $\mu_n \rightarrow \mu$ "weakly" if $\int f d\mu_n \rightarrow \int f d\mu$ for every bounded and uniformly continuous function f .

⁴As π^1 is continuous and bounded, the expected payoff is well-defined by Fubini's theorem. See Hildenbrand [1974].

⁵While this is a consequence of the Alaoglu theorem, in the present context it is interesting to note that Wald [1947] gives a direct proof.

References

- Dasgupta, P. and E. Maskin [1982] "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory," London School of Economics, D.P. 82/54.
- Fudenberg, D. and D. Levine [1982] "Subgame-Perfect Equilibria of Finite and Infinite Horizon Games," Journal of Economic Theory, forthcoming.
- Green, E. [1982] "Continuum and Finite-Player Non-Cooperative Models of Competition," Cal Tech W.P. 418.
- Hildenbrand, W. [1974] Core and Equilibria of a Large Economy, Princeton: Princeton University Press.
- Kreps, D. and R. Wilson [1982] "Reputation and Imperfect Information," Journal of Economic Theory 27, 253-279.
- Radner, R. [1980] "Collusive Behavior in Oligopolies with Long but Finite Lives," Journal of Economic Theory, 22, 136-156.
- Stokey, N. [1981] "Rational Expectations and Durable Goods Pricing," Bell Journal of Economics, 12, 112-128.
- Tjoe-Tie, Teh [1963] "Minimax Theorems on Conditionally Compact Sets," Annals of Mathematical Statistics, 1536-1540.
- Wald, A. [1947] "Sequential Decision Functions," Econometrica, 15, 279-313.
- Young, L.C. [1980] Lectures on the Calculus of Variations and Optimal Control Theory, New York: Chelsea Publishing Co.

Appendix

Proof of Proposition (5.1)

Let $(r_{i1}^n, \dots, r_{ij}^n, \dots)$ be the elements of R_i^n , but let $w_n = \sup_{s_i \in S_i} \inf_{r_i \in R_i^n} d(x_i, r_i)$. As R^n approximates S $w_n \rightarrow 0$. Let $(T_{i1}^n, \dots, T_{ij}^n, \dots)$ be a collection of measurable mutually disjoint subsets of S_i such that T_{ij}^n contains the j^{th} element of R_i^n , the diameter of each T_{ij}^n is at most w_n , and $\bigcup_j T_{ij}^n = S_i$. [Because R^n is countable such a partition exists.] For any $h_i \in R_i^n$ let $h_i(n)$ denote the probability measure defined by $h_i(n)(r_{ij}^n) = h_i(T_{ij}^n)$. Clearly

$$(A.1) \quad |\pi^i(h_i(n), s_{-i}) - \pi^i(h_i, s_{-i})| < w_n \forall s_{-i} \in S_{-i}.$$

Proof of part (A):

Suppose s were not an ϵ -equilibrium. Then there would be a player i , $h_i \in S_i$, on $\delta > 0$ such that $\pi^i(h_i, s_{-i}) > \pi^i(s) + \epsilon + \delta$. Since $r^n \rightarrow s$, and π^i is bounded and uniformly continuous on S , we have

$$(A.2) \quad |\pi^i(h_i(n), s_{-i}) - \pi^i(h_i(n), r_{-i}^n)| \rightarrow 0$$

$$(A.3) \quad |\pi^i(r^n) - \pi^i(s)| \rightarrow 0.$$

Combining (A.1), (A.2) and (A.3), we have

$$\pi^i(h_i(n), r_{-i}^n) - \pi^i(r^n) + \pi^i(s) > \epsilon \delta.$$

This contradicts r^n on ϵ^n -equilibrium with $\epsilon^n \rightarrow 0$.

Proof of part (B):

Let s be an ϵ -equilibrium and $r^n \rightarrow s$. Then

$$(A.4) \quad \forall_{i, h_i} \pi^i(h_i, r_{-i}^n) \rightarrow \pi^i(h_i, s_{-i}).$$

Combining (A.3) and (A.4), \forall_{i, h_i} :

$$\pi^i(h_i, r_{-i}^n) - \pi^i(r^n) \rightarrow \pi^i(h_i, s_{-i}) - \pi^i(s) < \epsilon.$$

Thus there is a sequence $\epsilon^n \rightarrow \epsilon$ such that r^n is an ϵ^n -equilibrium.