

Limit Games and Limit Equilibria*

DREW FUDENBERG

University of California, Berkeley, California 94720

AND

DAVID LEVINE

University of California, Los Angeles, California 90024

Received July 25, 1984; revised September 17, 1985

We provide a necessary and 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. *Journal of Economic Literature* Classification Number: 022. © 1986 Academic Press, Inc.

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 non-cooperative 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 equilibria. 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 and Levine [3]; 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 $\varepsilon \rightarrow 0$

* We would like to thank Ed Green, Jean-Charles Rochet, Sylvain Sorain, and participants at the IMSSS Summer Workshop and the Caltech Theory Workshop. We are particularly indebted to Bob Anderson, who encouraged us to find a more general formulation, and to an associate editor and a referee who made detailed comments on earlier versions of the paper. Financial support from the NSF and the UCLA Academic Senate Committee on Research is gratefully acknowledged.

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 $\varepsilon \rightarrow 0$. Thus the set of equilibria can be characterized by computing the set of limit points of the ε -equilibria of the restricted games. A related issue is the *appropriate* definition of equilibrium in games that are defined as limits. For example, it is difficult to give a general formulation of continuous-time games. Instead, one may simply define equilibrium as the (abstract) limit of those occurring in discrete time approximations. However, if the definition is to include all equilibria of games whose continuous-time limit is well-defined, our results show that it is necessary to take limits, not only of equilibria, but of ε -equilibria with $\varepsilon \rightarrow 0$.

Section 3 introduces a product topology on the strategies and uses it to prove a first version of our limit result. Section 4 applies this topology to finite-horizon approximations of infinite-horizon games, and Section 5 applies it to discrete-time approximations of open-loop equilibria of continuous-time games with continuous payoffs. Section 6 extends our results to mixed strategies.

Section 7 explains that the product topology is too restrictive, in that for some games of interest the only admissible sequence of restricted games is the original game itself. We are thus led to construct the coarsest topology for which our limit result obtains, because such a topology admits the largest possible class of approximating games. Section 8 applies this topology to games of timing.

2. RELATED WORK

Most discussions of the relationships between the equilibria of various games have focused on whether limits of equilibria are equilibria of the limit game, that is, whether the equilibrium correspondence is upper hemicontinuous. Walker [14] provides general conditions for a family of games to have this property. Green [5] establishes it for finite player approximations of games with a continuum of players. 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 and Wilson [8] and Stokey [11]. Dasgupta and Maskin [1] use finite-action approximations of games with a continuum of actions to investigate the existence of mixed-strategy equilibrium. Earlier papers by Wald [13] and Tjoe-Tie [12] demonstrate this for zero-sum games.¹ All these papers are concerned with only one direction of our limit theorem.

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

Radner [9] 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.

Harris [6] extends our earlier work on the finite to infinite-horizon limit. He introduces two new and more tractable topologies, each of which permits the characterization of infinite-horizon equilibria by finite-horizon ones if the game is continuous at infinity. One of his topologies is in fact the finest one for which this result obtains; his other topology is coarser, and, he argues, is "nearly" the coarsest possible.

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 $S \equiv \prod_i S_i$; elements of S will be called strategy profiles. Notice that S may be a space of mixed strategies. The profile derived from S by replacing its i th component by h_i is denoted by (h_i, s_{-i}) . Player i 's payoff π_i is a *bounded* real valued function on S .

DEFINITION (3.1). A strategy profile s is an ε -equilibrium if for all i and $h_i \in S_i$

$$\pi_i(h_i, s_{-i}) \leq \pi_i(s) + \varepsilon.$$

Thus each player gets within ε of the maximum. (See, e.g., Radner [9].) 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 $\varepsilon = 0$ we refer simply to equilibria: this is the usual noncooperative Nash equilibrium. 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 subset of the strategies available to them in S .

DEFINITION (3.2). $R \subseteq S$ 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 profile $r \in R$ is an ε -equilibrium relative to R if for all i and $h_i \in R_i$

$$\pi_i(h_i, r_{-i}) \leq \pi_i(r) + \varepsilon.$$

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

Let $\{R^n\}$ be a sequence of restricted games and ε^n a sequence converging to ε . We will be concerned with the conditions under which the ε^n -equilibria of R^n converge to the ε -equilibria of S . These conditions of course depend on the topology of S : each topology will generate conditions a sequence of restricted games must satisfy for the limit result to obtain. For example, if we endowed S with the discrete topology, the only sufficiently good approximation of S would be S itself. Coarser topologies on S allow more approximations.

We will begin with a fairly simple topology on S that we call the "inherent product" or p -topology. This topology will prove adequate for finite-horizon approximations of infinite-horizon games, and for finite-action approximations of games with payoffs continuous in the ordinary topology. Moreover, the product structure is essential for introducing mixed strategies. Later, however, we shall see that the p -topology is too fine to permit discrete-time approximations of discontinuous games of timing, which will motivate us to introduce an alternative topology that is the coarsest possible.

We now define the inherent product topology.

DEFINITION (3.4). The distance between two strategies of player i , $s_i, t_i \in S_i$, is

$$p_i(s_i, t_i) = \sup_{h_{-i} \in S_{-i}} \max_{j \in I} |\pi_j(s_j, h_{-i}) - \pi_j(t_j, h_{-i})|.$$

Thus $p_i(s_i, t_i)$ measures the greatest difference it would make to any player if player i used s_i instead of t_i . It is easy to check that p_i satisfies the triangle inequality. Thus p_i is a pseudometric, and so generates a topology. Notice that p_i may not be a metric, because it is possible that $p_i(s_i, t_i) = 0$ yet $s_i \neq t_i$. If p_i is not a metric, the topology it generates does not separate points.²

Remark. For a two-person zero-sum game $p_i(s_i, t_i)$ reduces to

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

This is Wald's [13] "intrinsic distance."

DEFINITION (3.5). The inherent product topology on S is the product topology induced when each of the S_i is topologized by p_i .

² While the term topology is used correctly here, most topologies with which economists are familiar do separate points. Topologies which do not have unfamiliar properties such as many limit points for a convergent sequence.

Until Section 7, all statements about convergence, continuity, etc., will be with respect to the inherent product topology, unless otherwise noted. This topology is generated by the pseudometric $p(s, t) = \sum_{i=1}^n p_i(s_i, t_i)$.

PROPOSITION (3.1). *If s is an ε -equilibrium then t is an $[\varepsilon + 2p(s, t)]$ -equilibrium. Thus if $s^n \rightarrow s$ are a sequence of ε^n -equilibria and $\varepsilon^n \rightarrow \varepsilon$, then s is an ε -equilibrium, and if $s^n \rightarrow s$ and s is an ε -equilibrium then the s^n are ε^n -equilibria with $\varepsilon^n \rightarrow \varepsilon$.*

Proof. Suppose s is an ε -equilibrium. Then unravelling the definitions,

$$\begin{aligned} \sup_{i, h_i \in S_i} \pi_i(h_i, t_{-i}) - \pi_i(t) &\leq \sup_{i, h_i \in S_i} \pi_i(h_i, t_{-i}) - \pi_i(s) + p(s, t) \\ &\leq \sup_{i, h_i \in S_i} \pi_i(h_i, s_{-i}) - \pi_i(s) + 2p(s, t) \\ &\leq \varepsilon + 2p(s, t). \end{aligned} \quad \text{Q.E.D.}$$

DEFINITION (3.6). 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 , or equivalently if for every $s \in S$ and subsequence $\{R^{n_k}\}$ there is a sequence contained in R^{n_k} , r^{n_k} which converges to s .

We now state and prove the first version of 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 (3.2) (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 $\varepsilon^n \rightarrow \varepsilon$ 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 $\varepsilon^n \rightarrow \varepsilon$ such that the r^n are ε^n -equilibria relative to R^n .*

COROLLARY (3.3). *If s is an ε -equilibrium there exists sequences $\{r^n\}$, $r^n \in R^n$, and $\{\varepsilon^n\}$, with $r^n \rightarrow s$ and $\varepsilon^n \rightarrow \varepsilon$ such that r^n is an ε^n -equilibrium relative to R^n .*

Proof.

$$(A) \quad \forall h_i \in S_i, \quad \pi_i(h_i, s_{-i}) - \pi_i(s) \leq \pi_i(h_i, r^n_{-i}) - \pi_i(r^n) + 2p(s, r^n).$$

Since $\{R^n\}$ approximates S , $h_i^n \in R_i^n$ such that $p_i(h_i^n, h_i) \rightarrow 0$. Since $\pi_i(h_i^n, r_{-i}^n) - \pi_i(h_i, r_{-i}^n) \leq p(h_i^n, h_i)$,

$$\begin{aligned} \pi_i(h_i, s_{-i}) - \pi^i(s) &\leq \pi_i(h_i^n, r_{-i}^n) - \pi_i(r^n) + 2p(s, r^n) + p_i(h_i^n, h_i) \\ &\leq \varepsilon^n + 2p(s, r^n) + p_i(h_i^n, h_i) \rightarrow \varepsilon. \end{aligned}$$

(B) Since $R^n \subseteq S$ this follows from Proposition (3.1).

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

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

COROLLARY (3.4). *If $\{R^n\}$ approximates S , $\{v^n\} \rightarrow v$ is a sequence of payoffs (one for each player) of ε^n -equilibrium relative to R^n , with $\varepsilon^n \rightarrow \varepsilon$, and S is compact, then there is an ε -equilibrium in S with payoffs v .*

Proof. Take a convergent subsequence of the ε^n -equilibria and apply Proposition (3.2).

Remark. Let N be the topological space consisting of the positive integers and $+\infty$ with the metric $\eta(n, m) = |(1/n) - (1/m)|$ (and $1/\infty = 0$). Fix a sequence R^n which approximates S and set $R^\infty = S$. Define the correspondence $\Omega: [0, \infty) \times N \rightrightarrows S$ to yield for any (ε, n) the set of ε -equilibria relative to R^n . Proposition (3.2A) says that Ω is closed valued at (ε, ∞) for every ε . As S is not necessarily compact, we cannot conclude that Ω is upper hemi-continuous. Green [5] gives conditions for upper hemi-continuity in a setting similar to ours. In addition the closed valuedness of Ω may be inferred directly from Walker's [14] generalization of the maximum theorem. We may view ε -equilibrium as an equilibrium relative to intransitive preferences, and these are clearly continuous with respect to ε . Since the mapping from N to the subsets of S is lower hemi-continuous at ∞ relative to the Hausdorff topology on subsets of S , Proposition (3.2A) follows directly from Walker's theorem. The lower hemi-continuity of Ω has not previously been discussed for games. There is a literature on lower hemi-continuity in general equilibrium theory, but this requires differentiability and regularity, which are not useful for discontinuous problems.

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 (3.2) is that any topology on S that is finer than ours will do.

In particular, any *product* topology on S such that the π^i are all uniformly continuous will do, because such topologies are at least as fine as the inherent-product topology.

PROPOSITION (3.5). *The inherent-product topology p is the coarsest product uniformity³ such that the π_i are uniformly continuous.*

Proof. That the payoffs are uniformly continuous w.r.t. p is obvious. To see that p is the coarsest such uniformity let τ_j be the components of a uniformity in which the π_i are uniformly continuous. We must show that this implies that p_j is uniformly continuous in τ_j . Let $\varepsilon > 0$ be given. Since all the π_i are uniformly continuous in τ there is a set $\bigcap_j T_j$ with $T_j \in \tau_j$ and such that $(r, t) \in T$ implies $|\pi_i(r) - \pi_i(t)| < \varepsilon$. Thus since $(s_k, s_k) \in T_k$ by the definition of a uniformity,

$$|\pi_i(r_j, s_{-j}) - \pi_i(t_j, s_{-j})| < \varepsilon. \quad \text{Q.E.D.}$$

4. FINITE- TO INFINITE-HORIZON LIMITS

We illustrate the idea of our limit theorem with a simple treatment of finite-horizon approximations of infinite-horizon games. Let A_i be a set of *actions* for player i . A strategy for player i is a sequence of mappings $s_i = (s_i^1, s_i^2, \dots)$ where $s_i^t \in A_i$ and for $t > 1$ $s_i^t: (\prod_{i \neq i} A_i)^{t-1} \rightarrow A_i$. The strategy space for i , S_i , is a subset of the space of all such sequences of mappings. This allows the incorporation of various restrictions, such as if player two played a_2 last period then player one cannot play a_1 . However, we assume that there is a designated "null" action \hat{a}_i for each player which is always feasible so that if $s_i \in S_i$ then $(s_i^1, \dots, s_i^t, \hat{a}_i, \hat{a}_i, \dots) \in S_i$.

We approximate S with a collection of finite-horizon games. Define $R_i^n \equiv \{s_i \in S_i \mid s_i^t = \hat{a}_i \text{ for } t > n\}$, 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 (4.1). π^i is *continuous at infinity* if

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

³ A uniformity is a type of topological space in which uniform continuity can be defined. See Kelley [7, Chap. 6] for details.

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

PROPOSITION (4.1). *In an infinite-horizon game S , if the π_i 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. Harris [6] provides a similar generalization for perfect equilibria.

As an example consider the repeated two-player Prisoner's dilemma with discounting. If the discount factor is not too small, it is well known that with an infinite horizon the "cooperative" strategy "don't cheat if cheating has never occurred, otherwise cheat forever" is an equilibrium. However, in the finite-horizon games the unique equilibrium requires that both players cheat in every period. What is true is that the "cooperative" strategy is an ε^n -equilibrium with $\varepsilon^n \rightarrow 0$; indeed $\varepsilon^n \rightarrow 0$; indeed ε^n is the cost of failing to cheat in the later periods 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 $\lim_{n \rightarrow \infty} \varepsilon^n \rightarrow 0$; 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 the restricted games which converge to a given equilibrium in S ; $\varepsilon^n > 0$ may be required.

A deeper discussion of the finite to infinite-horizon case can be found in Fudenberg and Levine [3] in which we discuss sequential equilibria with incomplete information and given applications to the existence and uniqueness of infinite-horizon equilibrium. Harris' topology is more useful for determining uniqueness.

5. OPEN-LOOP EQUILIBRIUM

We now suppose that time paths of actions are represented by S_i , the Lebesgue measurable functions from $[0, 1]$ to a finite-dimensional Euclidean space. Payoffs are assumed uniformly 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 is

not compact and as a sequence of restricted equilibria may not converge to anything at all. For example, the sequence of restricted strategies

$$\begin{aligned} s^i(k/n) &= 1, & k \text{ even} \\ &= -1, & k \text{ odd} \end{aligned} \quad (5.1)$$

does not have a limit. This is the "chattering" problem and can arise even in continuous-time control problems (see, for example, Davidson and Harris [2]). One solution that works for control problems is to define "chattering" controls—simply define "idealized" or "generalized" strategies as the limit of sequences such as (5.1). The limiting "strategy" is simply a functional which assigns each player the limiting payoff (see Young [15] for details). Unfortunately this "solution" does not appear to work when there is more than one player.

6. MIXED STRATEGIES

In this section we consider the relationship between an underlying space of pure strategies and mixed-strategy equilibrium. The point of this section is that if the restricted *pure*-strategy spaces approximate the unrestricted pure-strategy space *uniformly* then the restricted mixed-strategy spaces approximate the unrestricted mixed-strategy space, and so our limit theorem (Proposition (3.2)) applies. This section also relates our work to that of Wald [13] on the existence of minimax values.

Let B_i be the Borel algebra on S_i relative to the p -topology and let \tilde{S}_i be the space of probability measures on B_i . We endow S with the product algebra B and define \tilde{S} to be the family of probability measures which are products of measures of the S_i . Since π_i is continuous with respect to the p -topology it is measurable, and thus for $\tilde{s} \in \tilde{S}$ $\pi_i(\tilde{s})$ represents the random variable induced by π_i .

When \tilde{s} is played, player i 's payoff is the expectation of $\pi_i(\tilde{s})$, which we denote $\pi_i^*(\tilde{s})$. Since π_i is bounded, this expectation exists and is bounded. We call (\tilde{S}, π^*) the mixed-strategy game corresponding to (S, π) . To introduce the restricted mixed-strategy games, we assume that R_i^n are measurable subsets of S , and define \tilde{R}_i^n and \tilde{R}^n analogously to \tilde{S}_i and \tilde{S} .

The intrinsic product topology on \tilde{S}_i is generated by the pseudometric \tilde{p}_i given by

$$\tilde{p}_i(\tilde{s}_i, \tilde{t}_i) \equiv \sup_{\tilde{h}_{-i} \in \tilde{S}_{-i}} \max_{j \in I} |\pi_j^*(\tilde{s}_i, \tilde{h}_{-i}) - \pi_j^*(\tilde{t}_i, \tilde{h}_{-i})| \quad (6.1)$$

The key to extending our result to mixed strategies is the relationship

between \tilde{p}_i and the expectation Ep_i of the pure-strategy pseudometric p_i . Observe first that

$$\tilde{p}_i(\tilde{s}_i, \tilde{t}_i) = \sup_{h_{-i} \in S_{-i}} \max_{j \in I} |\pi_j^*(\tilde{s}_i, h_{-i}) - \pi_j^*(\tilde{t}_i, h_{-i})|, \quad (6.2)$$

because the largest difference in j 's payoff caused by opponents' mixed strategies is no larger than that caused by their pure strategies. From (6.2) we see that for pure strategies s_i and t_i , $\tilde{p}_i(s_i, t_i) = E(p_i(s_i, t_i)) = p_i(s_i, t_i)$. However, for mixed strategies \tilde{s}_i and \tilde{t}_i , the payoff difference measured by \tilde{p}_i allows only a single choice of the opponents' strategies h_{-i} , while the (expected) difference measured by Ep_i allows a different selection h_{-i} for each pair of pure strategies in the support of $\tilde{s}_i \times \tilde{t}_i$. We conclude that

$$\tilde{p}_i(\tilde{s}_i, \tilde{t}_i) \leq E(p_i(\tilde{s}_i, \tilde{t}_i)). \quad (6.3)$$

To extend our results to mixed-strategy equilibria we must strengthen our notion of what it means for restricted games to approximate S . We do this by requiring that for fixed n , sufficiently large, points in S can *all* be approximated by points in R^n , strengthening Definition (3.6), which did not require the approximation to be uniform.

DEFINITION (6.1). The sequence of restricted strategy spaces $\{R_i^n\}$ approximates S_i uniformly if

$$\omega_i^n \equiv \sup_{s_i \in S_i} \inf_{r_i \in R_i^n} p_i(s_i, r_i) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

In other words, all points in S_i can be approximated equally well by points in R_i^n for n sufficiently large. Obviously if $\{R_i^n\}$ approximate S_i uniformly for all i then $\{R^n\}$ approximates S .

We can use Eq. (6.3) to show that uniform approximation is inherited from pure to mixed strategies. The proof maps mixed strategies in \tilde{S} to mixed strategies in \tilde{R}^n by assigning each point in R^n all the probability weight "near" that point.

PROPOSITION (6.1). If $\{R_i^n\}$ approximate S_i uniformly and each R_i^n is countable then $\{\tilde{R}_i^n\}$ approximate \tilde{S}_i uniformly.

Proof. Find a measurable partition of S_i with the property that the diameter of each member does not exceed ω_i^n , and such that each member of the partition contains an element of R_i^n . Since the R_i^n are countable, such a partition may be constructed by ordering the ω_i^n -balls around the points in R_i^n and deleting from each ball the union of preceding balls. Thus the function $a_i^n: S_i \rightarrow R_i^n$ assigning a point in S_i the point in R_i^n representing its

equivalence class in S_i is measurable and satisfies $p_i(s_i, a_i^n(s_i)) \leq \omega_i^n$. To each $\tilde{s}_i \in \tilde{S}_i$ assign the probability measure $a_i^n(\tilde{s}_i)$ of the R_i^n -valued random variable induced by a_i^n . In other words, $a_i^n(\tilde{s}_i)$ assigns points in R_i^n the probability of their equivalence class. Then by (6.2)

$$\tilde{p}_i(\tilde{s}_i, a_i^n(\tilde{s}_i)) \leq E p_i(\tilde{s}_i, a_i^n(\tilde{s}_i)) \leq \omega_i^n. \quad (6.4)$$

Thus $\tilde{\omega}_i^n \leq \omega_i^n \rightarrow 0$.

Q.E.D.

Remark. If the R_i^n did not approximate S_i uniformly, we could still partition S_i using balls around points in R_i^n , but we could not conclude that their expected diameter converged to zero as $n \rightarrow \infty$.

The consequence of Proposition (6.1) is that when approximation is strengthened to uniform approximation the limit theorems carry over immediately from pure to mixed strategies. This is the case, for example, in the finite- to infinite-horizon limit discussed in Section 4.

PROPOSITION (6.2). *Suppose $\{R_i^n\}$ approximates S_i uniformly, each R_i^n is countable, and $\tilde{r}^n \in \tilde{R}^n$.*

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

(B) *If \tilde{s} is an ε -equilibrium and $\tilde{r}^n \rightarrow \tilde{s}$ then there is a sequence $\varepsilon^n \rightarrow \varepsilon$ such that \tilde{r}^n are ε^n -equilibria relative to \tilde{R}^n .*

We can also use Proposition (6.1) to show that compact games have mixed-strategy equilibrium.

PROPOSITION (6.3). *If each S_i is compact \tilde{S} has a mixed-strategy equilibrium.*

We prove this in two steps. Since S_i is a compact metric space it is totally bounded.

LEMMA (6.4). *If each S_i is totally bounded then for every $\varepsilon > 0$ \tilde{S} has an ε -equilibrium.*

Proof. Totally bounded means (by definition) the existence of a sequence of finite sets $\{R_i^n\}$ which uniformly approximate S_i . Since the R_i^n are finite, (\tilde{R}^n, π^*) has an equilibrium; since the approximation is uniform the type of inequalities used in Section 3 shows that this equilibrium is a ω^n -equilibrium in \tilde{S} . As $\omega^n \rightarrow 0$ we have the desired conclusion.

Proof of Proposition (6.3). Since S is compact \tilde{S} is weakly compact.⁴ Thus the sequence of ε^n -equilibria with $\varepsilon^n \rightarrow 0$ ensured by Lemma (6.4) has a weakly convergent subsequence. From Eq. (6.3) this is also a \tilde{p} convergent subsequence, and by the limit theorem Proposition (6.2) the limit point is an equilibrium of \tilde{S} . Q.E.D.

Remark. In proving Proposition (6.3) we make use of the implication of Eq. (6.3) that weak convergence implies \tilde{p} -convergence.

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

PROPOSITION (6.5) (Wald). *Let S be a two-person zero-sum game. If S_i is compact then \tilde{S} has a minimax value.*

7. THE INHERENT TOPOLOGY

While the inherent product topology is fairly natural, and is adequate for a number of applications, it is too fine to admit parsimonious approximations of some games. In particular, if payoffs are discontinuous "along the diagonal" in some natural topology, the use of a product structure on S is inconvenient. This leads us to pose the problem of finding the coarsest topology for which a limit theorem like Proposition (3.2) could hold. We find the coarsest topology under which a slightly stronger result obtains.

Let us begin by considering two different two-player games on the unit square illustrated in Fig. 1. In each game both players' payoffs are equal, and can be either 0 or 1. In game I, $\pi_i(s_i, s_2) = 1$ if $s_1 \geq \frac{1}{2}$, 0 otherwise; while in game II, $\pi_i(s_i, s_2) = 1$ if $s_2 \geq s_1$. In either game any point with a payoff of 1 is an equilibrium. In game I the strategy space can be approximated by a grid of points with coordinates $(j/n, k/n)$; indeed, any two points with the same payoff are identified by the p -topology (that is, their distance is zero). In game II, however, all points are at distance 1 from each other, that is, the p -topology for this game is discrete. Thus the only approximation of game II in the p -topology is the original game itself. However, it is obvious that the ε -equilibria of the discretized versions of game II are sufficient to characterize its ε -equilibria. It is equally obvious that the appropriate

⁴ μ_n converges to μ "weakly" if $\int f d\mu_n \rightarrow \int f d\mu$ for every bounded and uniformly continuous function f . That \tilde{S} is weakly compact is a consequence of the Alaoglu theorem. It is interesting to note that Wald [13] does not cite this theorem and gives a direct proof that \tilde{S} is compact.




FIGURE 1

topology to prove such a result will not have a product structure—two points on opposite sides of the diagonal are “far apart,” but this cannot be deduced by comparing each coordinate separately. Discontinuities along “diagonals” are not uncommon in economic games. They arise, for example, in games of timing in which it is substantially better to be first than second. (For other examples, see Dasgupta and Maskin [1].) For this reason we would like to have a way of proving a limit result for such games. As the appropriate topology will not have a product structure, it will seem less “natural” than the p -topology. We thus abandon “naturalism” as a criterion, and instead search for the coarsest topology possible.

First, though, we point out that the original limit result, Proposition (3.2), is not in fact the strongest result available. Proposition (3.2A) is consistent with the r^n being ε^n -equilibria only for ε^n bounded away from zero, $\varepsilon^n \rightarrow \varepsilon$, $r^n \rightarrow s$, yet s being a zero-equilibrium, because if s is a zero-equilibrium it is also an ε -equilibrium for any epsilon. Proposition (3.2A) requires that the required size of epsilon not “jump up” in the limit, but it is consistent with the required epsilon “jumping down.”

However, we can prove a stronger version of the limit theorem in which all the epsilons used are the smallest possible, thus ruling out jumps down in the required size of epsilon. This result brings the relationship between the restricted and unrestricted games into sharper focus.

DEFINITION (7.1). For any restricted game R and strategy selection $r \in R$, let $\varepsilon(r, R)$ be the smallest ε such that r is an ε -equilibrium relative to R ; that is, for $r \in R$, $\varepsilon(r, R) = \max_i \sup_{s_i \in R_i} (\pi_i(s_i, r_{-i}) - \pi_i(r))$.

PROPOSITION (7.1). Suppose $\{R^n\}$ approximates S , $r^n \in R^n$, and $r^n \rightarrow s$. Then $\varepsilon(r^n, R^n) \rightarrow \varepsilon(s, S)$.

Proof. By definition

$$\begin{aligned}
 \varepsilon(s, S) &= \sup_{i, h_i \in S} \pi_i(h_i, s_{-i}) - \pi_i(s) \\
 &\geq \sup_{n, j, h_j^n \in R_j^n} \pi_j(h_j^n, s_{-j}) - \pi_j(s) \\
 &\geq \sup_{n, j, h_j^n \in R_j^n} \pi_j(h_j^n, r^n_{-j}) - \pi_j(r^n) - 2p(r^n, s) \\
 &\geq \limsup \varepsilon(r^n, R^n) - 2p(r^n, s)
 \end{aligned}$$

where the first inequality follows from $R^n \subseteq S$. As $r^n \rightarrow s$ we conclude that

$$\varepsilon(s, S) \geq \limsup \varepsilon(r^n, R^n).$$

From Proposition (3.2A)

$$\varepsilon(s, S) \leq \liminf \varepsilon(r^n, R^n)$$

and thus it follows that all three values are equal to each other and to $\lim \varepsilon(r^n, R^n)$. Q.E.D.

An obvious candidate for a coarse topology on S that will be consistent with (7.1) is that generated by using ε .

DEFINITION (7.2). The "inherent" pseudometric is

$$m(r, s) = |\varepsilon(r, S) - \varepsilon(s, S)|,$$

which generates the "inherent" or m -topology.

Under m , all strategy configurations with the same ε are at zero distance from each other. For example, two strategies with different payoffs can be close, which is not the case with p . Therefore, unlike p , m is not likely to be a metric. Moreover, in contrast to p and to the work of Green and Walker, players' preferences need not be open in the m -topology (that is, the payoffs need not be continuous). This allows topologies which will admit discrete approximations to discontinuous games such as game II above.

Associated with any topology μ on S we have the following notion of subset convergence:

DEFINITION (7.3). A sequence of restricted games $\{R^n\}$ " μ -approaches" S if $r^n \in R^n$ and $r^n \rightarrow^\mu S$ implies $|\varepsilon(r^n, S) - \varepsilon(r^n, R^n)| \rightarrow 0$.

Consequently, R^n μ -approaches S if for all convergent sequences r^n , the best deviation against r^n in R^n is almost as good as the best deviation

against r^n in S . As we shall see, the proviso that the sequences r^n be convergent is frequently irrelevant.

The reason for our interest in the m -topology and m -approaching is that these are the weakest conditions which yield "epsilon-continuity." Consider a pair (μ, σ) where μ is a topology on S and σ is a list of sequences of subsets of S which approach S . Although σ need not be a topology, we do require that $(S, S, \dots) \in \sigma$; that is, S approaches itself.

DEFINITION (7.4). A game is *epsilon-continuous* if $r^n \rightarrow^\mu s$ and $R^n \rightarrow^\sigma S$ imply $\varepsilon(r^n, R^n) \rightarrow \varepsilon(s, S)$.

In this case since $S \rightarrow^\sigma S$ by assumption, $r^n \rightarrow^\mu s$ implies that $\varepsilon(r^n, S) \rightarrow \varepsilon(s, S)$, or, by definition, that $r^n \rightarrow^m s$. In other words, m has the most convergent sequences of any topology in which ε is continuous. On the other hand, if $r^n \in R^n$, $r^n \rightarrow^\mu s$, and $R^n \rightarrow^\sigma S$, then $\varepsilon(r^n, S) \rightarrow \varepsilon(s, S)$ and $\varepsilon(r^n, R^n) \rightarrow \varepsilon(s, S)$. We conclude that $[\varepsilon(r^n, S) - \varepsilon(r^n, R^n)] \rightarrow 0$, and by definition this means that R^n μ -approaches S . Consequently for fixed μ the largest possible collection σ consistent with epsilon-continuity is the set of μ -approaching sequences. In summary

PROPOSITION (7.2). A pair (μ, σ) are *epsilon-continuous* if and only if μ -convergence implies m -convergence, and every sequence in σ μ -approaches S .

In practice, we are interested in a fixed sequence of approximations R^n . In order to be able to approximate all equilibria in S with equilibria in R^n , we must require that the sequence $\{R^n\}$ is dense (in the sense of Definition (3.6)) in the chosen topology. By choosing the m -topology we get the most possible convergent sequences, and thus the best possible chance of density. If, in addition to density, R^n m -approaches S the game is also epsilon-continuous. In this case we have the limit theorem (Proposition (3.2)) strengthened to epsilon-continuity (Proposition (7.1)). In practice, rather than testing that R^n m -approaches S , it is often easier to check simply that $\sup_{r^n \in R^n} [\varepsilon(r^n, S) - \varepsilon(r^n, R^n)] \rightarrow 0$, in which case we say that R^n approaches S uniformly.

If R^n does not m -approach S there is a tradeoff. We must examine topologies μ which are strictly finer than m (have fewer convergent sequences). This makes it more likely that R^n μ -approaches S , as only those sequences $r^n \in R^n$ which μ -converge need be checked to see if $[\varepsilon(r^n, S) - \varepsilon(r^n, R^n)] \rightarrow 0$. In this way we may get epsilon-continuity. On the other hand, weakening μ reduces the chance that R^n will be dense, so a tradeoff is involved. If no topology μ can be found in which R^n is both dense and μ -approaching, then the necessity of Proposition (7.2) implies that the game is not epsilon-continuous.

Unlike the intrinsic product topology, there are no easy theorems relating the intrinsic topology in a pure-strategy game to the intrinsic topology of the corresponding mixed-strategy game. Consider the game in Fig. 2, in which player 1 chooses a number between 0 and 1 on the horizontal axis and player 2 on the vertical axis. If $s_1 + 2s_2$ is strictly less than $\frac{1}{2}$ or strictly bigger than $2\frac{1}{2}$, player 1 gets -1 , otherwise he gets $+1$. For simplicity, player 2 always gets 0. Let \tilde{s}_2 be the mixed strategy in which player 2 plays 0 and $+1$ with equal probability of $\frac{1}{2}$, and \tilde{s}_1 be the pure strategy in which player 1 plays 0. Then $\varepsilon(\tilde{s}, \tilde{S})$ is 1 since player 1 can gain this amount by switching to the pure strategy $\frac{1}{2}$. On the other hand, if R_1^n does not include the strategy $\frac{1}{2}$ for player 1 then $\varepsilon(\tilde{s}, \tilde{R}^n) = 0$. If R_2^n includes both 0 and 1, and R_1^n includes 0, then $\tilde{s} \in \tilde{R}^n$ and \tilde{R}^n does not μ -approach \tilde{S} in any topology. On the other hand, in pure strategies, if the R^n are increasingly fine grids, it is clear that R^n approaches S uniformly: for any pure strategy by player 2, player 1 can always get $+1$ provided the points on his grid are spaced no more than $\frac{1}{2}$ apart. The difficulty is that for any given pure strategy by player 2, either moving a little to the right or a little to the left of one-half will be as good for player 1 as playing one-half exactly, even though the other direction may make him much worse off. With mixed strategies player 2 may randomize between a pure strategy which requires 1 to approximate to the left and one which requires him to approximate to the right. In the example, no approximate strategy works nearly as well as one-half does.

Insofar as the limit theorem is to be used only to prove the existence of mixed-strategy equilibria in the limit game, some progress is possible. Dasgupta and Maskin [1] show that if R^n are any sequence of increasingly fine grids, and the discontinuities satisfy some very stringent conditions, then limits of zero-equilibria are zero-equilibria. Simon [10] is able to weaken these conditions somewhat by restricting the set of grids on which the limit theorem is to hold: in Fig. 2, for example, if R_1^n eventually contained $\frac{1}{2}$, our limit theorem would hold, although not for other grids.




FIG. 2. Payoffs to player 1.

8. TWO-PERSON TIMING GAMES

In a game of timing each player must choose when to undertake a particular investment or some other activity. Furthermore, if his opponent moves first, a player may wish to revise his choice of time. Thus a strategy for player i is to choose a time to move first T_i and also a response if his opponent moves first, $\rho_i(T_{-i}) > T_{-i}$. If $T_1 = T_2$ both players move simultaneously at this time; if $T_1 < T_2$ one moves at $\bar{T}_1 = T_1$ and two at $\bar{T}_2 = \rho_2(T_1)$, and conversely. Reaction functions are restricted to be piecewise continuous. The payoff to player i is then $\pi_i(\bar{T}_1, \bar{T}_2)$. This we also take to be piecewise continuous, allowing the possibility that there is a substantial advantage (or disadvantage) to moving first. We also choose to normalize time so that $0 \leq \bar{T}_i \leq 1$. This then defines π and S . Note that we restrict attention to pure-strategy equilibrium: some but not all timing games have such equilibria.

We approximate S with finite move games—the games R^n at which the feasible choices of time are the lattice points $[k/n, (k+1)/n]$. We now wish to check that in the m -topology R^n approaches S and that the limit of R^n is dense in S . In this case Proposition (7.2) and its converse will hold: continuous-time equilibria can be approximated by discrete-time equilibria and conversely.

Let $s = (T_1, T_2, \rho_1, \rho_2)$ be a given profile in S . We must show two things:

- (a) density: find a sequence r^n such that $\varepsilon(r^n, S) \rightarrow \varepsilon(s, S)$;
- (b) m -approaching: if $\varepsilon(r^n, S) \rightarrow \varepsilon(s, S)$ then $[\varepsilon(r^n, S) - \varepsilon(r^n, R^n)] \rightarrow 0$.

Turning first to density let s be given. Define r^n to be the strategy defined by rounding off times to adjacent lattice times. Then if h_i is a deviation against s , since the payoffs depend only on the times at which players move, it is clearly possible to find a deviation h_i^n in R^n that is close to h_i in that the times at which players move at (h_i^n, r_{-i}^n) and at (h_i, s_{-i}) are approximately the same. In Fig. 3 the result A of 2 playing ρ_2 and 1 deviating to T_1 can be approximated by player 1 deviating to T_1^n against the approximate reaction function ρ_2^n with the result A^n . Furthermore, with a little care, when the grid is fine enough h_i^n may be chosen so that these times lie in the same component as each other (with respect to the discontinuities in π_i and π_{-i}). Note that in Fig. 3 it is essential that the deviation T_1^n be chosen so that the result A^n lies in the same component of the square (relative to the line of discontinuity) as A . Similarly for a deviation h_i^n against r^n we can find a deviation h_i against s yielding nearly the same movement times. Thus $\varepsilon(r^n, S)$ and $\varepsilon(s, S)$ must be close and as the grid is refined $\varepsilon(r^n, S) \rightarrow \varepsilon(s, S)$. This establishes density.

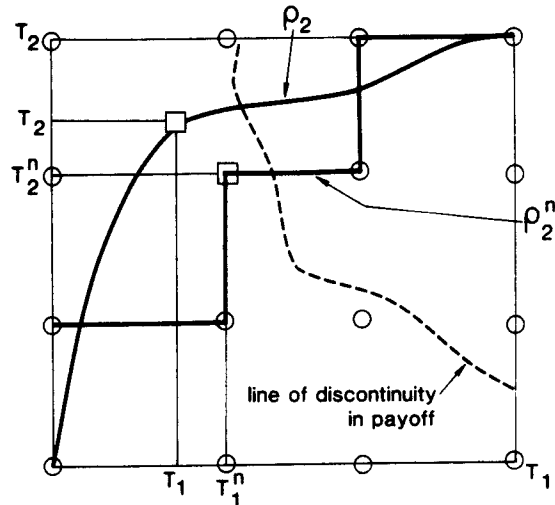


FIG. 3. In the figure in continuous time 1 chooses T_1 and 2 responds with T_2 . In discrete time 2 responds with the step function ρ_2^n . To approximate the result that he gets with T_1 , player 1 should choose T_1^n resulting in a response of T_2^n . Note that (T_1^n, T_2^n) lies in the same component as (T_1, T_2) , ensuring that the payoffs are close.

To show that R^n m -approaches S we can simply show that R^n approaches S uniformly, that is, the bound on $[\varepsilon(r^n, s) - \varepsilon(r^n, R^n)]$ depends only on n . Since π^i is piecewise continuous, the loss from moving at lattice times depends only on the fineness of the grid and can be bounded independently of the particular lattice strategy followed by the opposing player. The argument proceeds as above by showing that the times at which the moves take place nearly the same, and since a piecewise continuous function is (by definition) uniformly continuous on components this has only an effect on payoffs bounded by the modulus of continuity. Thus R^n approaches S uniformly.

Note that even if the π_i are continuous, the product topology is useless since the *responses* can be discontinuous, and these discontinuities do not respect the product structure of the unit square.

Consider the following simplification of the game in Fudenberg and Tirole [4]:

$$\begin{aligned}\pi_1(\bar{T}_1, \bar{T}_2) &= (\bar{T}_1 + \bar{T}_2)/2 + \bar{T}_2 - \bar{T}_1 \\ \pi_2(\bar{T}_1, \bar{T}_2) &= (\bar{T}_1 + \bar{T}_2)/2 + \bar{T}_1 - \bar{T}_2, \\ 0 &\leq \bar{T}_1, \quad \bar{T}_2 \leq 1.\end{aligned}\tag{8.1}$$

The symmetric Pareto optimum outcome is $\bar{T}_1 = \bar{T}_2 = 1$, which gives each player a payoff of one-half. This outcome can be sustained as an equilibrium in continuous time with reaction functions $\rho_1(T_2) = T_2$ and $\rho_2(T_1) = T_1$. (These reactions are in fact best responses.) While each player

would prefer to move earlier than this opponent if his opponent's time were fixed, the threat of immediate retaliation enforces the cooperative outcome. In discrete time, a player who is moving at the same time or later than this opponent always gains at least $\delta/2$ (where δ is the period length) by undercutting his opponent by one period. The cooperative equilibrium unravels in the manner of the finitely repeated Prisoners' dilemma and the unique equilibrium is $\bar{T}_1 = \bar{T}_2 = 0$. Consequently, the equilibrium correspondence is not lower hemi-continuous in the limit of shorter time periods. However, the cooperative outcome can be sustained as an ε -equilibrium in discrete time, and the required epsilon converges to zero in the continuous-time limit.

REFERENCES

1. P. DASGUPTA AND E. MASKIN, The existence of equilibrium in discontinuous economic games. I. Theory, London School of Economics, D. P. 82/54, 1982.
2. R. DAVIDSON AND R. HARRIS, Nonconvexities in continuous-time investment, *Rev. Econ. Stud.* **48** (1981), 235-253.
3. D. FUDENBERG AND D. LEVINE, Subgame-perfect equilibria of finite and infinite horizon games, *J. Econ. Theory*, **31** (1983), 251-268.
4. D. FUDENBERG AND J. TIROLE, Preemption and rent equalization in the adoption of new technology, *Rev. Econ. Stud.* **52**, 383-401.
5. E. GREEN, Continuum and finite-player non-cooperative models of competition, *Econometrica* **52** (1984), 975-994.
6. C. HARRIS, Perfect equilibrium in infinite horizon games, *J. Econ. Theory* **37**, 99-125.
7. J. L. KELLEY, "General Topology," Springer-Verlag, New York, 1955.
8. D. KREPS AND R. WILSON, Reputation and imperfect information, *J. Econ. Theory* **27** (1982), 253-279.
9. R. RADNER, Collusive behavior in oligopolies with long but finite lives, *J. Econ. Theory* **22** (1980), 136-156.
10. L. SIMON, Games with discontinuous payoffs. Part I. Theory, mimeo, University of California at Berkeley, November, 1984.
11. N. STOKEY, Rational expectations and durable goods pricing, *Bell J. Econ.* **12** (1981), 112-128.
12. T. TJOE-TIE, Minimax theorems on conditionally compact sets, *Ann. Math. Statist.*, (1963), 1536-1540.
13. A. WALD, Sequential decision functions, *Econometrica* **15** (1947), 279-313.
14. M. WALKER, A generalization of the maximum theorem, *Int. Econ. Rev.* **20** (1979), 267-271.
15. L. C. YOUNG, "Lectures on the Calculus of Variations and Optimal Control Theory," Chelsea, New York, 1980.