SEQUENTIAL EQUILIBRIA OF

FINITE AND INFINITE HORIZON GAMES[1,2]

By

Drew Fudenberg and David Levine*

*Department of Economics, UC Berkeley and UCLA respectively

ABSTRACT

We show that sequential equilibria of infinite-
horizon games arise as limits, as $T \to \infty$ and $\epsilon^T \to 0$, of sequen-
tial ϵ^T-equilibria of the game which is truncated after T
periods of play. A number of applications show that this
result provides a useful technique for analyzing the
existence and uniqueness of infinite-horizon sequential
equilibria.

1. Introduction

The idea that economic agents are **sequentially rational** has been useful in understanding a **wide variety of problems**. Sequential rationality was first **formalized by Selten** [13, 14] as "sub-game" and "trembling-hand"-perfect equilibrium. A mathematically more convenient formalization is Kreps/ Wilson [4]'s "sequential" equilibrium. These equilibria are ordinarily computed by backwards induction from a finite horizon. Specification of a fixed horizon is often artificial, however, and an infinite horizon game better captures the economics of a situation. This paper describes a method for characterizing sequential equilibria in infinite horizon games.

We do not consider here the most general possible extensive-form games. We do consider a formulation which allows simultaneous moves, and, in Section Six, uncertainty and mixed strategies. We allow current options to be limited by the history of play. We also allow relatively general forms of intertemporal preference, requiring neither stationarity nor additive separability. One restriction we do impose is that agents should be impatient -- they should not be too concerned about events in the far distant future. While our model is somewhat restrictive, it covers many cases considered in the economics literature and involves substantially less notation than the most general case.

The technique we propose is to study ε^T-equilibria in the game truncated after T periods of play. Here we follow Radner [9] in defining an ε-equilibrium as a strategy selection in which each player, taking opponents strategies as given, is within ε of the largest possible payoff. Our main result says that as $T \to \infty$ and $\varepsilon^T \to 0$ the set of ε^T-equilibria in the truncated games converges to the set of sequential equilibria in the infinite horizon game. Because players are not too concerned about the distant future equilibria for a long finite horizon will "almost" be equilibria in the infinite horizon, and conversely. Characterization of infinite horizon equilibria as limit points is then made possible by finding a suitable topology on the space of strategies.

In Section 2 we introduce a model which allows simultaneous moves, but not uncertainty or mixed strategies. Section 3 contains a technical analysis of continuity and the limiting behavior of equilibria. Section 4 considers games with finitely many actions in each period. We show that with perfect information perfect equilibria exist. We also show that our results extend to the case of uncertainty and mixed strategies. In Section 5 we discuss the uniqueness of equilibrium. In the finite action case we give an easily verifiable necessary and sufficient condition for the uniqueness of pure strategy equilibrium. Using a similar technique we study a special case of Rubinstein's [11] bargaining game, giving a more informative

proof of uniqueness than in the original. Section 6
reviews our findings.

2. Games, Subgames, and Equilibria

This section defines games in extensive and normal form when there is no uncertainty. We do not consider the most general definition of a game in extensive form. Nevertheless, many economically important games are in the class we study.

For our purposes a game (in extensive form) has an infinite number of periods $t = 1,2,\ldots$. Each period all N players simultaneously choose actions from feasible sets of actions, which we take to be subsets of $\mathbb{R}^M$. When they choose an action in period t they know the entire history of the game until and including time $t-1$.[3] It is possible that the set of feasible actions is constrained by the history of play.

The _outcome_ of the game in period T lies in $\mathbb{R}^{MN}$. The way in which the outcome is made up of individual choices is discussed below. The _history_ of the game is a sequence of outcomes $x = (x_1, x_2 \ldots) \in \mathbb{B} \equiv X_{t=1}^{\infty} \mathbb{R}^{MN}$. We will regard this as a topological space with the product topology, an assumption we will explore in more detail in subsequent sections. The _action space_ of the game E is $E^A \subseteq \mathbb{B}$: it is a list of all possible histories of the game. (For technical reasons we shall always assume E^A is closed in $\mathbb{B}$.) An example helps illustrate this.

<u>Example 2-1</u> [McLennan's Termination Game[4]]:

There are two players, one and two. Play alternates with player one moving first. On his move a player may either continue or terminate the game. If a player terminates the game in period t he receives a present value of $\beta^t a$ and his opponent $\beta^t b$ where a and b are scalars and $0<\beta<1$ is the common discount factor. If play never terminates both players receive zero.

Let "0" denote the option of "doing nothing" and "1" be the option of terminating the game. Here N = 2 and M = 1: the outcome of the game is a pair (y_1, y_2) where $y_1, y_2 \in \{0,1\} \subset \mathbb{R}^2$. A player must choose 0 if it isn't his move, or if the game has already terminated. Thus the action space E^A is the set of sequences of the form $((0,0)_1, (0,0)_2, \ldots (1,0)_t, (0,0)_{t+1}, \ldots)$ where t is odd, $((0,0)_1, (0,0)_2, \ldots (0,1)_t, (0,0)_{t+1}, \ldots$ where t is even, or $((0,0)_1, (0,0)_2, \ldots)$.

It is generally useful and entails no loss of generality to designate the outcome $0 \in \mathbb{R}^{MN}$ the "null" outcome "nothing happens". We require that the null outcome always be feasible. This means that if x is feasible then the vector x(t), truncated after t by requiring that the null outcome occur in periods t+1, t+2,..., is also feasible:

(2-1) $\forall x \in E^A \quad \forall t \qquad x(t) \equiv (x_1, x_2, \ldots, x_t, 0, 0, \ldots) \in E^A$

Let $E^A(x,s)$ be the space of all possible outcomes
in period s consistent with the history $x_1, x_2, \ldots x_{s-1}$,
with the convention that $E^A(0,1)$ is the set of possible
first-period outcomes. By assumption (2-1) we may consider
this to be the space of vectors y such that
$(x_1, x_2, \ldots, x_{s-1}, y, 0, 0, \ldots) \in E^A$ since if
$z = (x_1, \ldots, x_{s-1}, y, z_{s+1}, z_{s+2}, \ldots) \in E^A$ then
$z(s) = (x_1, \ldots, x_{s-1}, y, 0, 0, \ldots) \in E^A$ as well.

If E^A is to be the action space of a game then the
choices available to player i in period t given a prior
history x , denoted $E^{Ai}(x,t)$, must not depend on what
other players do in period t. Thus, in addition to (2-1),
we must also require that the space of all feasible
outcomes $E^A(x,t)$ is the cartesian product of the individual
action spaces

(2-2) $\forall x \in E^A \; \forall t \quad E^A(x,t) = X_{i=1}^N E^{Ai}(x,t)$.

Thus in Example 2-1 the set of possible outcomes at
time 2 if the game has not yet been terminated,
$E^A(0,2) = \{(0,0),(0,1)\}$, is the cartesian product of
$E^{A1}(0,2) = \{0\}$ and $E^{A2}(0,2) = \{\, 0 \,,\, 1 \,\}$.

<u>Definition 2-1</u>: A game in extensive form E is a pair
(E^A, E^V) where $E^A \subset \mathbb{B}$ is a closed set satisfying (2-1)
(2-2) and $E^V = (E^{Vi})_{i=1}^N$ is an N-tuple of valuation functions

$E^{Vi}:E^A \to \mathbb{R}$ assigning a value to each history of the game.

In Example 2-1 where $z^1 = ((1,0)_1, (0,0)_2, \ldots)$ and $z^2 = ((0,0)_1, (0,1)_2, (0,0)_3, \ldots)$ $E^V(0) = (0,0)$; $E^V(z^1) = (\beta a, \beta b)$ and $E^V(z^2) = (\beta^2 b, \beta^2 a)$.

Example 2-2 [Repeated Games]:

Each agent i has a fixed set of actions $0 \in A^i \subset \mathbb{R}^M$, a utility function $u^i:A \to \mathbb{R}$ where $A \equiv X_{i=1}^N A^i$ and a discount factor β_i. Then in our framework the repeated game has the action space $E^A \equiv X_{t=1}^\infty A$ so that history places no constraints on behavior. The valuation functions are

$$E^{Vi}(x) = \Sigma_{t=1}^\infty \beta_i^t u^i(x_t).$$

Further examples are given in subsequent sections.

Associated with each game in extensive form are a collection of truncated games in normal form: $N(T)$ denotes the normal form of the game truncated at time T by assigning the null outcome to all periods following time T. Let us formally describe the strategy space of $N(T)$. At time s player i, knowing the history $x_1, x_2, \ldots, x_{s-1}$, must choose a feasible action in $E^A(x,s)$ to undertake in period s. (Note that for now we don't allow mixed strategies.) Let $g_s^i(x)$ denote this choice. Thus for $s = 1$ $g_s^i \in E^{Ai}(0,1)$ while for $s > 1$ g_s^i is a mapping

$$(2-3) \quad g_s^i:E^A(s-1) \to \mathbb{R}^M \text{ with } g_s^i(x) \in E^{Ai}(x,s)$$

where $E^A(t)$ denotes all possible histories to time t, i.e. all vectors $(x_1, x_2, \ldots, x_t, 0, 0, \ldots) \in E^A$. A complete set of contingent choices of this type is called a strategy and is simply a sequence $(g_1^i, g_2^i, \ldots, g_T^i, 0, 0, \ldots)$ where $g_1^i \in E^A(0,1)$ and for $s > 1$ g_s^i satisfies (2-3). The set of all such strategies is called the strategy space for player i and is denoted $N^{Si}(T)$. The strategy space for the truncated game $N(T)$ is just the cartesian product $N^S(T) \equiv X_{i=1}^N N^{Si}(T)$. Note that $N^S(1) \subseteq N^S(2) \subseteq \ldots \subseteq N^S(\infty)$. While the truncated games depend on which action is specified as the null action, we will later see that this is irrelevant for our results.

The outcome function F_{xs} assigns strategy selection $g \in N(\infty)$ the history of the game that occurs when the initial history is $x_1, \ldots, x_{s-1}$ and afterwards each player plays g^i:

$$F_{xs}(g) = z \quad \text{where for } s > 1$$

$$(2\text{-}4) \quad z_t = \begin{cases} x_t & 1 \leq t \leq s-1 \\ \\ g_t(z_1, z_2, \ldots, z_{t-1}, 0, 0, \ldots) & t \geq \max(s,2) \end{cases}$$

We denote the history that occurs when each player plays g^i from the start by $F_{01}(g)$. Note that $F_{xs} \in E^A$ follows from the fact that E^A is closed.[5]

To illustrate these definitions consider in Example 2-1 the strategy by player one "terminate in period three unless player two has already terminated, after

period three don't terminate" which has the form

$$g^1 = (0,0,g^1_3,0,0,\ldots)$$

$$g^1_3 = \begin{cases} 0 & x_2 = (0,1) \\ 1 & x_2 = (0,0) \end{cases}$$

and the strategy by player two "never terminate" which is given by

$$g^2 = (0,0,\ldots).$$

Then $F_{01}(g)$ is the outcome that actually occurs so

$$F_{01}(g) = ((0,0)_1,(0,0)_2,(1,0)_3,(0,0)_4,\ldots)$$

while $F_{04}(g)$ is the outcome which occurs if the history before time 4 is $x_1 = (0,0)$ $x_2 = (0,0)$, and $x_3 = (0,0)$ so that $F_{04}(g) = 0$. In other words if one reneges on his plan to terminate in period 3 neither player ever terminates. Finally if

$$x = ((0,0)_1,(0,1)_2,(0,0)\ldots)$$

(so that two <u>does</u> terminate in period 2) $F_{x3}(g) = x$ and one <u>must</u> (and does) choose the null action in period 3.

We turn now to equilibrium in the games $N(T)$. Complete rationality of all players implies that whatever the history of the game to date they should choose the optimal course of action. More precisely, every decision must be part of an optimal strategy for the remainder of

the game. As there is no uncertainty at the beginning of each period, this rationality requirement can be imposed using Selten [13]'s concept of a subgame-perfect Nash equilibrium, that is, the subgame perfect and sequential equilibrium coincide in this case. Radner's [8] concept of a subgame-perfect ε-Nash equilibrium generalizes perfectness by assuming players may only be able to get within ε of the optimal payoff.[7]

<u>Definition 2-2</u>: $g^* \in N^S(T)$ is a subgame perfect ε-Nash equilibrium (or simply ε-perfect) iff for each $s \geq 0$, history x, strategy $g \in N^S(T)$ and player i

$$(2-5) \quad E^{Vi}(F_{xs}(g^i, g^{*-i})) - E^{Vi}(F_{xs}(g^*)) \leq \varepsilon \ ;$$

that is, iff in no circumstance can player i improve his payoff by more than ε given the strategies of all players.

Note that g^{-i} denotes the cartesian product of all players' strategies except for that of player i. Note also that the restriction $s \leq T$ in (2-5) would be vacuous , since, with $g, g^* \in N(T)$, for $t > T$ $g_t = g_t^* = 0$. Finally, if $\varepsilon = 0$ the equilibrium is simply called perfect.

One goal of this paper is to relate ε-perfect equilibria of truncated N(T) games to perfect equilibria of the N(∞) game. To this end define the constants w^T to be the greatest variation in any player's payoff due strictly to events after (T-1) :

$(2-6)$ $\quad w^T \equiv \textbf{sup} \quad |E^{Vi}(x) - E^{Vi}(z)|$

$$1 \leq i \leq N$$

$$x, z \in E^A$$

$$x(T-1) = z(T-1)$$

At this point w^T may be infinite, but we argue later that most games of interest in economics have $w^T \to 0$ as $T \to \infty$.

The idea behind the limit theorem of the next section is revealed in

Lemma 2-1:

(A) $h*$ ε-perfect in $N(T)$ is $(\varepsilon + w^T)$-perfect in $N(\infty)$

(B) $g*$ ε-perfect in $N(\infty)$ then

$\quad$ $h* \equiv g*(T) \equiv (g_1^*, g_2^*, \ldots, g_T^*, 0, 0, \ldots)$ is

$\quad$ $(\varepsilon + 2w^T)$-perfect in $N(T)$

The point is that strategies in $N(\infty)$ differ from strategies in $N(T)$ only after time T and thus by (2-6) have payoffs within w^T of the truncated strategies.

proof:

(A) let $g \in N^S(\infty)$ and let x and s be given. Set $h = g(T) \equiv (g_1, g_2, \ldots, g_T, 0, \ldots)$. By assumption

$(2-7)$ $\quad E^{Vi}(F_{xs}(h^i, h*^{-i})) - E^{Vi}(F_{xs}(h*)) \leq \varepsilon$

11

while since h and g differ only after T by definition

(2-8) $E^{Vi}(F_{xs}(g^i, h^{*-i})) - E^{Vi}(h^i, h^{*-i})) \leq w^T.$

Adding (2-7) to (2-8) shows

(2-9) $E^{Vi}(F_{xs}(g^i, h^{*-i})) - E^{Vi}(F_{xs}(h^*)) \leq \epsilon + w^T.$

Since g, x, and s are arbitrary (2-9) implies h* is $(\epsilon + w^T)$-perfect.

(B) Let $h \in N(T)$, and x, s be given. Since g* is ϵ-perfect in $N(\infty)$

(2-10) $E^{Vi}(F_{xs}(h^i, g^{*-i})) - E^{Vi}(F_{xs}(g^*)) \leq \epsilon .$

Since h* and g* differ only after T

(2-11) $E^{Vi}(F_{xs}(g^*)) - E^{Vi}(F_{xs}(h^*)) \leq w^T.$

and also

(2-12) $E^{Vi}(F_{xs}(h^i, h^{*-i})) - E^{Vi}(F_{xs}(h^i, g^{*-i})) \leq w^T.$

Adding (2-10), (2-11), and (2-12) shows

(2-13) $E^{Vi}(F_{xs}(h^i, h^{*-i})) - E^{Vi}(F_{xs}(h^*)) \leq \epsilon + 2w^T ,$

and thus h* is $(\epsilon + 2w^T)$-perfect.   Q.E.D.

3. Continuity and Limit Equilibria

This section contains our main result: a strategy selection is perfect in $N(\infty)$ if and only if it is the limit as $T \to \infty$ and $\varepsilon^T \to 0$ of ε^T-perfect equilibria in $N(T)$. Before proving this result we must discuss the continuity of the valuation functions and the convergence of equilibria. This requires that we define topologies on E^A and $N^S(\infty)$.

Recall that $E^A \subset X_{T=1}^{\infty} \, \mathbb{R}^{MN} = \mathbb{B}$. The metric

$$(3-1) \quad d(x,z) \equiv \sup_T \; [(1/T)\min\{|x_T - z_T|, 1\}]$$

induces the product topology on $\mathbb{B}$.[6] Hereafter all statements about continuity, convergence, etc. will be with respect to this topology (relativized to E^A).

Having introduced a topology on E^A we now discuss continuity of the valuation function $E^V : E^A \to \mathbb{R}^N$, which we refer to as continuity of the game. Continuity implies events in the far distant future don't matter very much. While this may not be a good assumption in planning models, such as that of Svenson [14], it is a natural assumption about the preferences of individual economic agents.

Definition 3-1: E is *uniformly continuous* if for all $x^n, z^n \subset E^A$, $(x^n - z^n) \to 0$ implies $|E^V(x^n) - E^V(z^n)| \to 0$.

We shall only be interested in uniformly continuous

games. Many economically interesting games are of this type.

Recall that w^T is the greatest variation in any

player's payoff due solely to events after T. The idea

that the future doesn't matter very much is captured by

requiring $w^T \to 0$.

Definition 3-2: E is <u>continuous at infinity</u> iff $w^T \to 0$

as $T \to \infty$.

A supergame has w^T constant over time and is not

continuous at infinity. A repeated game (Example 2-2)

with discount factor $1 > \beta > 0$ has $w^T = \beta w^{T-1}$ and is

continuous at infinity provided $w^1 < \infty$.

An important fact is that uniform continuity implies

continuity at infinity.

Lemma 3-1: E uniformly continuous implies E continuous

at infinity.

This follows simply from unwinding the definitions.

Finally, we must extend our notion of convergence

in E^A to the strategy space $N^S(\infty)$ (and implicitly to its

subsets $N^S(T)$ $T < \infty$). We choose a topology which captures

the notion of closeness most relevant to perfect

equilibrium : two strategies f and g are close if for

every t and initial history $x \in E^A$ the histories resulting

from f and g being played are close <u>and</u> the history

14

resulting when any **one player deviates from f is close**
to that resulting from the **same deviation against g.**
This topology is generated by the metric

(3-2) $d(f,g) \equiv$

$$\sup_{x \in E^A, t} \{d(F_{xt}(f), F_{xt}(g)), \sup_{\substack{i, T \\ h^i \in N^{Si}(T)}} [d(F_{xt}(h^i, f^{-i}), F_{xt}(h^i, g^{-i}))]\}.$$

Our motivation for choosing this topology is
revealed by the following lemmas.

<u>Lemma 3-2</u>: The strategy space $N^S(\infty)$ is closed.

<u>Lemma 3-3</u>: Let g^n be ε-perfect in $N(\infty)$ and $g^n \to g$ in a
continuous game. Then g is also ε-perfect.

<u>proof</u>:
Suppose g is <u>not</u> ε-perfect so that for some t, some $x \in E^A$,
and $g^i \in N^{Si}(\infty)$,

$$E^{Vi}(F_{xt}(\tilde{g}^i, g^{-i})) - E^{Vi}(F_{xt}(g)) > \varepsilon \qquad (2-3)$$

Since $g^n \to g$, for large n $F_{xt}(\tilde{g}^i, g^{n-i})$ is close to
$F_{xt}(\tilde{g}^i, g^{-i})$; and as E^{Vi} is continuous, for N large
enough

$$(3-4) E^{Vi}(F_{xt}(\tilde{g}^i, g^{N-i})) - E^{Vi}(F_{xt}(g)) > \varepsilon$$

contradicting g^n ε-perfect.   Q.E.D.

15

The lemma shows the chosen topology was fine enough to guarantee that the ε-equilibrium sets are closed. Of course we could simply have declared these to be closed, but then we could hardly hope to characterize infinite-horizon equilibria as limit points. The interest in the lemma, and the justification of the chosen topology on $N^S(\infty)$, is

Theorem 3-4 [Limit Theorem]:

Suppose E is uniformly continuous. Then

(A) A necessary and sufficient condition that g* be perfect in $N(\infty)$ is that there be a sequence $\{g^n\}$ of $2w^{T(n)}$-perfect in $N(T(n))$ such that as $n \to \infty$, $T(n) \to \infty$ and $g^n \to g*$ (in the space $N(\infty)$).

(B) A necessary and sufficient condition that g* be perfect in $N(\infty)$ is that there be sequences ε^n, $T(n)$, and g^n such g^n is ε^n-perfect in $N(T(n))$ and as $n \to \infty$, $\varepsilon^n \to 0$, $T(n) \to \infty$, and $g^n \to g*$.

proof:

Since the hypothesis of (A) implies that of (B), it suffices to show the hypothesis of (A) necessary and that of (B) sufficient.

(A) Necessary:

We claim the sequence $\{g*(n)\}$,

$g*(n) = (g_1^*, g_2^*, \ldots g_n^*, 0, 0, \ldots)$ with $T(n) = n$ has the requisite property. First, since g*(n) and g* exactly

agree in the first n periods, $d(g*(n),g*) \leq 1/(n+1)$ (see (3-2)). Thus $g*(n) \rightarrow g*$. By Lemma 2-1(B) we also have $g*(n)$ $2w^T$-perfect in $N(n)$.

(B) Sufficient:

By Lemma 2-1(A) g^n is $(\varepsilon^n + w^{T(n)})$-perfect in $N(\infty)$. Since $\varepsilon^n + w^{T(n)} \rightarrow 0$, for each $\delta > 0$ there is an N such that $w^{T(n)} + \varepsilon^n < \delta$, whenever $n > N$. Thus by Lemma 3-2 $g*$ is δ-perfect. Since this is true for every $\delta > 0$, $g*$ is in fact perfect.
                                                                    Q.E.D.

One application of this theorem is to infinitely repeated games with discounting. These games are known to have a great multiplicity of equilibria.[8] An implication of our theorem is that there are a great multiplicity of ε-equilibria in the finite horizon. This allows a finite horizon resolution of the prisoner's dilemma, Radner introduced ε-equilibrium for precisely this reason. Although his results were for games with time averaging, our results provide insight into why his effort was successful.

4. Finite-Action Games

Finite-action games are games in which there are only a finite number of possible actions in each period.

Definition 5-1: E is a __finite-action game__ iff for each t and history $x \in E^A$ the set of feasible outcomes in period t given the history x, $E^A(x,t)$, is a finite set.

First, we prove that finite-action games have three key properties:

(1) they are uniformly continuous if and only if they are continuous at infinity;

(2) g^n converges to g iff for any T g^n and g eventually coincide for the first T periods; and

(3) the strategy space $N^S(\infty)$ is compact.

As a corollary to these results we show that in finite-action games of perfect information perfect equilibria always exist. In section five we use the results of this section in conjunction with Theorem 3-4 to analyze the uniqueness of equilibrium in finite-action games. To conclude the section we explain how these results can be extended to allow uncertainty and mixed strategies.

It is convenient to have a concrete description of convergence in finite-action games: convergent sequences in E^A must for any T eventually coincide for the first T periods.

Definition 5-2: $\{x^n\} \subset \mathbb{B}$ converges finitely to x (or f-converges) iff $\forall T \geq 1$ $\exists N > 0$ such that $n \geq N$ implies $x^n(T) = x(T)$ (i.e. for $1 \leq t \leq T$ $x_t^n = x_t$).

Lemma 5-1: In finite-action games f-convergence and convergence are equivalent on E^A.

As an immediate consequence we have

Corollary 5-2: In finite-action games uniform continuity and continuity at infinity are equivalent.

Just as convergent sequences of histories must eventually coincide in finite-action games, convergent sequences of strategies must too.

Definition 5-3: $\{g^n\} \subset N^S(\infty)$ converges finitely to g (or f-converges) iff $\forall T \geq 1$ $\exists N > 0$ such that $n \geq N$ implies for $1 \leq t \leq T$ $g_t^n = g_t$

Lemma 5-3: In finite-action games f-convergence and convergence are equivalent on $N^S(\infty)$.

We turn now to the compactness of $N^s(\infty)$. A useful way to study this problem is to observe that $N^s(\infty)$ is the space of sequences of maps $(g_1, g_2, \ldots)$. The map g_t has a finite domain, with say L_t elements, and ranges in $\mathbb{R}^{NM}$. Thus it may be viewed simply as a vector in $\mathbb{B}_t^* \equiv \mathbb{R}^{NML_t}$, and $N^s(\infty) \subset \mathbb{B}^* \equiv \underset{t=1}{\overset{\infty}{\times}} \mathbb{B}_t^*$ in a natural way. Furthermore it is easy to see that the topology on $N^s(\infty)$ given in (3-1) and (3-2) is the same as the relative product topology in $\mathbb{B}^*$: in both cases convergence means that for any fixed horizon the sequence is eventually stationary before that horizon. However $N^s(\infty)$ is the cartesian product of finite (and thus compact) subsets of the $\mathbb{B}_i^*$ implying that it is itself compact and proving

Lemma (5-4): In finite-action games $N^s(\infty)$ is compact.

We can use lemma 5-4 to prove an existence theorem. A game of perfect information has no more than one player making a decision in each period (who the player is may depend on the history). In our notation, for each t and history $x \in E^A$, there is a player i such that $E^{A-i}(x,t) = 0$; only player i faces a decision.

It is well-known and can easily be established by backwards induction from the horizon that a finite-horizon finite-action game of perfect information has a perfect equilibrium. From this we deduce

<u>Corollary 5-4</u>: Continuous (at infinity) finite-action
games of perfect information have perfect equilibria.

<u>proof</u>:

Each finite-horizon subgame $N(T)$ has a perfect equilibrium
g^T. By Lemma 2-1(A) g^T is w^T-perfect in $N(\infty)$. Since
$N^S(\infty)$ is compact there is a subsequence $\{h^T\} \subset \{g^T\}$
with $h^T \to g* \in N^S(\infty)$. By Theorem 3-3(B) this implies
$g*$ is perfect in $N(\infty)$.   Q.E.D.

We now wish to extend the analysis of finite games to
permit the possibility of uncertainty and mixed strategies.
In each period t "nature" makes a random move Θ_t. Players
are only partially aware of the result of this move. Play-
er i observes only a signal $\Theta_t^i \in \Pi_t^i \subset \mathbb{R}^M$ where ⓟ $_t^i$ is a
finite set. For simplicity we assume $\Theta_t = (\Theta_t^1, \Theta_t^2, \ldots, \Theta_t^N)$
so that players would know Θ_t if they pooled their know-
ledge. Any additional uncertainty (on the part of all
players) can be incorporated directly into payoffs via
expected utility. Naturally the payoff functions E^{Vi} are
defined on ⓟ $\times E^A$ where ⓟ $\equiv \underset{t=1}{\overset{\infty}{x}} \underset{n=1}{\overset{N}{x}}$ ⓟ $_t^i$. We give
ⓟ $\times E^A$ the product topology and define continuity at
infinity and uniform continuity over ⓟ $\times E^A$ rather than
just E^A.

Decisions by agents are based on their probability
beliefs about past and future values of Θ_t. Future beliefs
are based on <u>a priori</u> information; beliefs about past
values of Θ_t are based upon private information Θ_t^i and

upon information revealed by the play of opponents as captured by the history of the game. Letting $\circledR^i(T) \equiv X_{t=1}^T \circledR_t^i$ we define a system of beliefs for player i to be a sequence of mappings M_s^i with domain $\circledR^i(s-1) \times E^A(s-1)$ and ranging over the space of probability measures on $\circledR(s-1)$; that is it represents beliefs about past outcomes given current information. Since $\circledR(s-1)$ and $E^A(s-1)$ are finite μ_s^i may be viewed as a vector in a finite dimensional vector space $B_s^{\mu_i}$.

An agent's play is characterized both by his beliefs and by the strategy he plays. Analogous to our previous definition a strategy for player i is a sequence of mappings g_s^i with domain $\circledR^i(s-1) \times E^A(s-1)$. Now, however, we wish to allow mixed strategies so that g_s^i ranges over the space of <u>probability measures</u> on $E^{Ai}(x,s)$. Since $\circledR^i(s-1) \times E^A(s-1)$ and $E^{Ai}(x,s)$ are finite sets g_s^i may be viewed as a vector in a finite dimensional vector space $B_s^{g_i}$.

The overall play of an agent is called an <u>assessment</u>: it is a system of beliefs $(\mu_1^i, \mu_2^i, \dots)$ and a strategy $(g_1^i, g_2^i, \dots)$ for each agent i. The space of all possible assessments is denoted by $M(\infty)$. Just as $N(\infty) \subset X_{t=1}^\infty \mathbb{B}_t^*$ so $M(\infty) \subset X_{t=1}^\infty X_{i=1}^N (B_t^{\mu i} \times B_t^{g i}) \equiv \mathbb{B}^*$. The product topology on $\mathbb{B}^*$ then introduces a corresponding topology on $M(\infty)$. Note that we could have introduced an economically meaningful topology along the lines of the metric in (3-2). However, due to the finiteness of the game this will be identical to the product topology. We also have the notion of

22

truncated assessment $M(T) \subset M(\infty)$ in which g_s^i places unit probability weight on action zero for $s > T$. Finally we may define $U_{\Theta \times s}^i(M)$ as the expected utility accruing to player i at time s when M is an assessment selection and the expectation is taken according to i's probability beliefs conditional on the history x and the private information available from Θ.

With this set-up we can define a sequential ε-equilibrium, following Radner [9] and Kreps-Wilson [4].

Definition (5-4): A sequential ε-equilibrium is an assessment selection (μ^*, g^*) such that

(1)   The strategy g^{*i} is ε-optimal for each player given his beliefs and the play of opponents for all i, Θ, x and s

$$U_{\Theta \times s}^i(\mu^*, g^i, g^{*-i}) - U_{\Theta \times s}^i(\mu^*, g^*) \leq \varepsilon.$$

(2)   Agents beliefs are consistent with Bayes law in the sense that there is a sequence (μ^n, g^n) converging to (μ^*, g^*) with g^n placing positive weight on every possible outcome and μ^n derived from Bayes law.

Our goal is to show that the limit theorem (3-3) holds for this new model with "assessments" replacing "strategies." To do this we must reprove the truncation lemma (2-1), lemma (3-3) showing that the set of ε-equilibria is closed and the proof theorem (3-4) itself. With the exception of lemma (3-3) all proofs go through verbatim by merely

23

changing the notation to replace "strategies" by "assessments." Lemma (3-3) follows quite easily from part (2) of definition (5-4): sequential equilibria are well-behaved with respect to limiting operations. Note that this would not be the case had we chosen to work with "trembling-hand perfect" equilibria.

We can now prove an existence result. Since $M(\infty)$ is the product of compact sets in the product topology it is itself compact. From Kreps-Wilson we know there is a sequential equilibrium m^T in each $M(T)$. Since $M(\infty)$ is compact these have a subsequent converging to $m^* \in M(\infty)$. By the limit theorem m^* is a sequential equilibrium. Thus we have demonstrated

Theorem (5-5): Continuous (at infinity) finite action games with imperfect information have mixed-strategy sequential equilibria.

5. Uniqueness of the Infinite-Horizon Perfect Equilibrium

This section uses the limit theorem of section three
to study the uniqueness of infinite-horizon perfect
equilibrium. The limit theorem implies that there will
be a unique equilibrium if and only if all convergent
sequences of truncated $2w^T$-perfect equilibria have the
same limit as $T \to \infty$. As an aside, note that a necessary
condition for uniqueness is that every convergent
sequence of truncated perfect equilibria have the same
limit.

The first class of games we consider are the
finite-action games of section five. Recall that in
such games a sequence of strategies converges if and only
if it converges finitely (Lemma 5-3). This means that
there will be a unique infinite-horizon perfect
equilibrium if and only if by taking the horizon ,T, large
enough, we can ensure both that a $2w^T$-perfect equilibrium
exists, and that all $2w^T$-perfect equilibria exactly agree
in the first k periods. Formally we have

Definition 6-1: A game is <u>finitely determined</u> (f.d.)
iff for any k > 0 there is T $\geq$ k such that

(a) there is g $2w^T$-perfect in N(T)

(b) if g' is $2w^T$-perfect in N(T) and k $\geq$ t > 0 $g_t = g'_t$.

Proposition 6-1: There exists a unique infinite-horizon perfect equilibrium in a closed finite-action game that is continuous at infinity if and only if it is finitely determined.

Thus uniqueness in finite-action games requires that changes in strategies at the horizon not affect play in the early periods. As an illustration, consider McClellan's terminating game of Example 2-1 shown in Figure 6-1. At each node, the indicated player chooses whether to "terminate" or "continue". If the game terminates at node k, k odd, the payoffs are $\beta^{k-1}(a,b)$; if k is even, they are $\beta^{k-1}(b,a)$; and if no player chooses to terminate, they are $(0,0)$.

This game is finitely determined in two cases

case (i) $a > 0$ $a > \beta b$
case (ii) $a < 0$ $a < \beta b$

and it is not finitely determined in the complementary cases

case (iii) $a \geq 0$ $a \leq \beta b$
case (iv) $a \leq 0$ $a \geq \beta b$.

We show this for cases (i) and (iii). Note that a strategy may be viewed as a choice of which nodes to stop at (if the game hasn't stopped already). For example, if T is even, "stop at T, T-2, T-4,..." is a strategy for player two: it means that if the game hasn't stopped before T, two will stop it, otherwise he chooses the null action.

$(0,0)$

$\beta^2(b,a)$

$\beta(a,b)$

$\beta(b,a)$

(a,b)
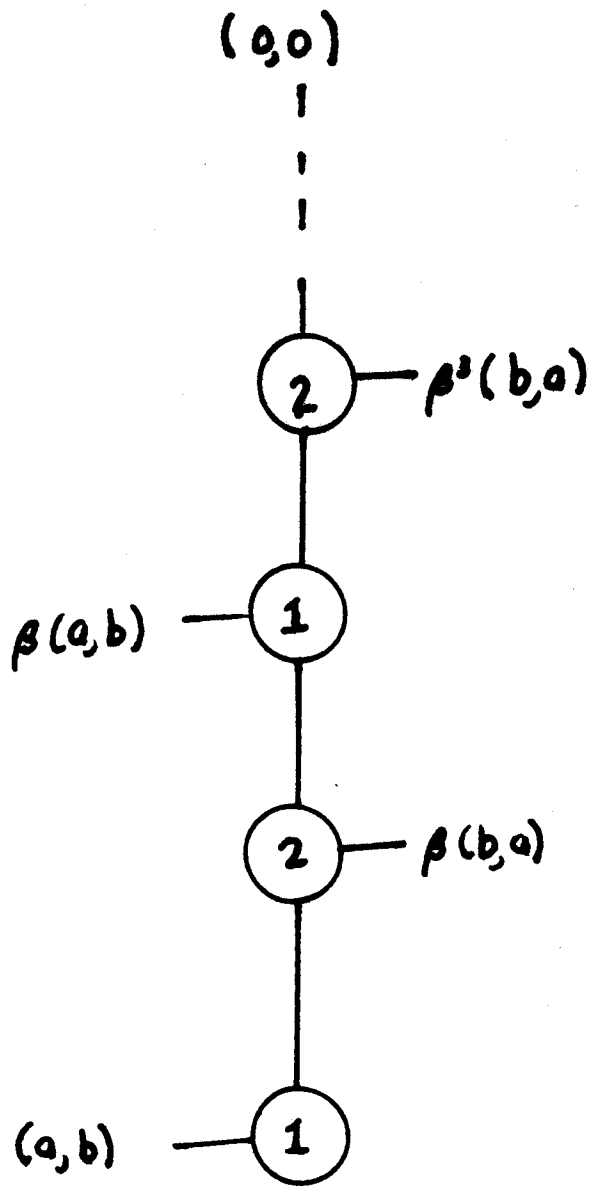
FIGURE   6-1

Case (i) is a game which both players want to stop
as quickly as possible. Indeed, in the perfect equilibria
of the truncated game the last player to move must stop,
and in every previous period the moving player stops.
In a $2w^T$-perfect equilibrium the last player to move can choose
to continue. However, at earlier nodes k, the minimum
loss from continuing is $\beta^k \min (a - \beta^2 a, a - \beta b)$. Thus
if $\epsilon < \beta^k \min (a - \beta^2 a, a - \beta b)$ all ϵ-equilibria must
terminate at all times before k. Since $w^T \rightarrow 0$ with T we
can always choose T large enough that $2w^T$-perfect equilibria have
both players stopping before T. Thus the game is finitely
determined and both players always stop.

Case (iii) is a game of "chicken": each player wants
the game to stop, but doesn't want to end it himself. In
the game truncated at an even time T the unique perfect
equilibrium is for two always to stop and one always to
continue. In the game truncated at an odd time T the
unique perfect equilibrium is for one always to stop and
two always to continue. Thus the period one action by
player one isn't uniquely determined and the game isn't
finitely determined.

In finite-action games, uniqueness of the
infinite-horizon perfect equilibrium is equivalent to
the condition that changes in strategies at the horizon
have no effect on (equilibrium) play earlier. In
continuous-action games we need not require that such

28

changes have no effect on earlier play but only that the effect is damped out as we work backwards from the horizon.

We illustrate this point with an example.

Example 6-1 [Rubinstein's Bargaining Game with Discounting]:
This is a special case of a game due to Rubinstein [11].
Two players, one and two, must decide how to partition a
pie of size one. Both players have a common discount
factor β and a utility function linear in pie. In odd
periods player one proposes a partition which player two
accepts or rejects. Similarly, in even periods, two makes
proposals. Play begins with player one in period one.
Play ends when a proposal is accepted. Thus if a partition
s is accepted in period k, player one gets a present value
of $\beta^k s$ and two $\beta^k(1-s)$.

We will show that this game has a unique infinite horizon
perfect equilibrium. To do so we will demonstrate that, for
any history x and time t, if T is big enough all $2w^T$-equilibria
have the player moving at t making an offer his opponent ac-
cepts in the same period. We then use this fact to show that
the offer by player one on an odd move k converges to $1/(1+\beta)$
as $T \to \infty$ and $w^T \to 0$. By symmetry this is also true of player
two's offers. It follows directly that the acceptance sets of
both players converge. The convergence of offers and acceptance
sets implies that the corresponding strategies (when properly
written out in the formalism of this paper) must converge.
Thus the infinite horizon equilibrium is unique.

We recall the convention that a partition is the amount of pie going to player one. Let $\varepsilon(k) \equiv \beta^k (1-\beta)/3$. If $T > k$ we claim all $\varepsilon(k)$-equilibria in $N(T)$ stop immediately. Assume without loss of generality k is odd so that one proposes the partition at k. If two doesn't accept one's proposal either no agreement is reached or two gets $1-s$ in period $k+j$. So two must accept any proposal promising him a present value of more than $\beta^{k+j}(1-s) + \varepsilon(k)$. In other words if one proposes a partition of $1-\beta^j(1-s) - \beta^{-k}\varepsilon(k)$ it will be accepted. If he is to make a proposal that is refused he must ultimately get more than this:

(6-1) $\beta^k[1-\beta^j(1-s)-\beta^{-k}\varepsilon(k)] \leq \beta^{k+j}s + \varepsilon(k),$

This implies

(6-2) $\varepsilon(k) \geq \beta^k (1-\beta)/2$

which contradicts our assumption.

Since $w^T \to 0$ when T is big enough $2w^T < \varepsilon(k)$ and at time k player one must make two an offer he can't refuse.

We continue to consider a $2w^T$ - perfect equilibrium. Let $\bar{s}^k$ be the largest (sup) proposal one makes at k and $\underline{s}^k$ the smallest (inf). If $2w^T$ is small enough these proposals will be accepted by two and the game ends. Thus at k one gets a present value of at least $\beta^k \underline{s}^k$ and no more than $\beta^k \bar{s}^k$. Now consider one's decision in period $k-1$ to accept or reject two's offer. If two proposes more than $\beta^{1-k}(\beta^k \bar{s}^k + 2w^T)$ one must accept since he can't get more than $\beta^k \bar{s}^k$ by continuing. Similarly he'll reject proposals of less than $\beta^{1-k}(\beta^k \underline{s}^k - 2w^T)$.

30

Since two's proposals must be ~~irresistable they won't be less~~
than $\beta^{1-k}(\beta^k \underline{s}^k - 2w^T)$ and two certainly won't be ~~offer~~ more than
$\beta^{1-k}(\beta^k \bar{s}^k + 4w^T)$. Reasoning as above, this means that at $k-2$
two accepts proposals offering him more than

$$\beta^{2-k}\left\{\beta^{k-1}\left[(1-\beta^{1-k}(\beta^k\underline{s}^k - 2w^T)\right] + 2w^T\right\} \quad \text{and}$$

rejects proposals offering him less than

$$\beta^{2-k}\left\{\beta^{k-1}\left[(1-\beta^{1-k}(\beta^k\bar{s}^k + 4w^T)\right] - 2w^T\right\}. \quad \text{As before}$$

this implies that

$$\bar{s}^{k-2} = 1 - \beta^{2-k}\left\{\beta^{k-1}\left[(1-\beta^{1-k}(\beta^k\bar{s}^k + 4w^T) - 2w^T\right]\right\}$$

$$(6\text{-}3) \quad \underline{s}^{k-2} = 1 - \beta^{2-k}\left\{\beta^{k-1}\left[(1-\beta^{1-k}(\beta^k\underline{s}^k - 2w^T)\right] + 4w^T\right\}$$

The claim we wish to establish is that as $T \to \infty$
$\bar{s}^k, \underline{s}^k \to 1/(1+\beta)$. Since the mapping in (6-3) is a contraction
as we work it backwards from period $k+j$, j large,
$\bar{s}^k$ approaches $\left[1/(1+\beta)\right] + c_j^k w^T$ and
$\underline{s}^k$ approaches $\left[1/(1+\beta)\right] - c_j^k w^T$.
Letting $w^T \to 0$ and noticing that c_j^k
doesn't depend on T yields the desired conclusion.

31

6. Conclusion

In games which satisfy an economically appealing continuity requirement, infinite-horizon equilibria coincide with the limits (as $T \to \infty$) of ε^T-equilibria of the finite-horizon truncated games. Because finite-horizon equilibria are easier to work with than infinite-horizon ones, this theorem provides a powerful tool for analyzing infinite-horizon games. It can be used to compute answers to such questions as the existence and uniqueness of infinite-horizon equilibria.

While our analysis examines only simultaneous-move extensive form games, it can easily be extended to cover other economic models such as strong perfect equilibrium, and "state space" games, in which payoffs and strategies depend not on all history but on a finite vector of "state" variables.[10,11] As a technical matter all that is required is to prove an analog of Lemma 3-1 and to find some reasonable notion of the convergence of strategies.

Footnotes

[1]It is our pleasure to thank Timothy Kehoe, Eric Maskin, Andreu Mas-Colell, Andrew McLennan, David Kreps, Ariel Rubinstein and Jean Tirole for helpful conversations. Joe Farrell and Franklin M. Fisher provided useful comments on an earlier draft.

[2]For conciseness some propositions are without proof. The missing proofs are in [2].

[3]More general definitions involving information sets can be found in Luce and Raiffa [7] or Kreps and Wilson [4].

[4]We are grateful to Andrew McLennan for providing this example, which helped clarify our thinking in the early stages of our investigation.

[5]Consider the one-player game with action space $0 = (0,0,...)$, $Z_1 = (1,0,0,...)$, $Z_2 = (1,1,0,...)$ ... Then $\lim_{n \to \infty} Z_n = (1,1,...) = Z_\infty \notin E^A$ so the game is not closed. One possible strategy, however, is for $g_n^1(Z_{n-1}) = 1$, that is, if 1 has always been played before play it again. Then $F_{01}(g) = Z_\infty \notin E^A$. This pathology is avoided by assuming E^A is closed.

[6]See Munkres [8], p. 123.

[7]As a model of bounded rationality ε-perfect equilibrium combines almost-optimization with perfect knowledge

of the game and perfect foresight. Levine [5] presents an alternative formulation.

[9]See, e.g., Munkres [8], p. 275.

[10]We thank A. Rubinstein for pointing this out. See [12] for a treatment of strong perfectness in supergames.

[11]For examples of such games see Fudenberg and Tirole [3] or Levine [6].

References

[1]   Friedman, J.: **Oligopoly and the Theory of Games**.
      Amsterdam: North-Holland, 1977.

[2]   Fudenberg, D. and D. Levine: "Prefect Equilibrium of
      Finite and Infinite Horizon Crimes," UCLA Working
      Paper #216.

[3]   Fudenberg, D. and J. Tirole: "Capital as a Commit-
      ment: Strategic Investment in Continuous Time,"
      essay 2, in Jean Tirole, **Essays in Economic
      Theory**, MIT Ph.D. dissertation, 1981.

[4]   Kreps, D. M. and R. Wilson: "Sequential Equilibria,"
      forthcoming, **Econometrica**, 1982.

[5]   Levine, D.: "Local Almost Perfect Equilibrium in a
      Game with Adjustment Costs," essay 3 in **The
      Enforcement of Collusion in Oligopoly**,MIT Ph.D.
      dissertation, 1981.

[6]   _____: "The Enforcement of Collusion in a
      Quadratic Symmetric Oligopoly," essay 2 in ibid.

[7]   Luce, R. P. and H. Raiffa: **Games and Decisions**. New
      York: Wiley, 1957.

[8]   Munkres, J. R.: **Topology: A First Course**. Englewood
      Cliffs, New Jersey: Prentice-Hall, 1975.

[9]   Radner, R.: "Collusive Behavior in Non-Cooperative
      Epsilon-Equilibria of Oligopolies with Long but
      Finite Lives," **Journal of Economic Theory**, 22
      (1980), 136-154.

[10]   Rubinstein, A.: "Equilibrium in Supergames with
       the Overtaking Criterion", Journal of Economic
       Theory, 21 (1979), 1-9.

[11]   _____: "Perfect Equilibrium in a Bargaining
       Model", mimeo, The Hebrew University, 1980.

[12]   _____: "Strong Perfect Equilibrium in Supergames",
       International Journal of Game Theory, 9 (1979),1-12.

[13]   Selten, R. "Spieltheoretische Behandlung eines
       Oligopolmodells mit Nachfrageträheit",
       Zeitschrift für die gesamte Staatswissenschaft,
       12 (1965), 301-324.

[14]   _____: "Re-examination of the Perfectness
       Concept for Equilibrium Points in Extensive
       Games", International Journal of Game Theory,
       1 (1975), 25-55.

[15]   Svenson, L-G.:"Equity Among Generations", Econometrica,
       48 (1980), 1251-1256.