PERFECT EQUILIBRIA OF FINITE AND INFINITE HORIZON GAMES[1]

by

Drew Fudenberg and David Levine*

*Department of Economics, UC Berkeley and UCLA respectively

UCLA Department of Economics
Working Paper #216
June 1981

ABSTRACT

We show that perfect equilibria of infinite-horizon games arise as limits, as $T \to \infty$ and $\varepsilon^T \to 0$, of ε^T-perfect 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 perfect equilibria.

1. Introduction

The notion of a perfect equilibrium of a game in extensive form, introduced by Selten [13,14], has recently been found useful in analyzing a wide variety of economic problems. The set of perfect equilibria can be straightforwardly computed in finite-horizon games by backwards induction from the horizon. In many cases, however, specification of the horizon is arbitrary, and frequently infinite-horizon games better capture the economics of the problem (the prisoner's dilemma is an example of this). Unfortunately, perfect equilibria of infinite-horizon games are more difficult to characterize. This paper describes a device that can be used for that purpose.

Our main result is that the set of infinite-horizon perfect equilibria of any game satisfying a reasonable continuity requirement is the same as the set of limit points as $T\to\infty$ and $\varepsilon^T\to 0$ of ε^T-perfect equilibria in the game which is truncated after T periods of play.[2] Here we follow Radner [9] in defining an ε-equilibrium as a state in which each player, given his opponent's strategies, receives within ε of the maximum payoff he could obtain. Since the truncated games have a finite horizon, this reduces the problem of finding infinite-horizon perfect equilibria to finding finite-horizon ε-perfect equilibria, which can be done via backwards induction from the horizon.

To demonstrate the usefulness of this technique we consider several applications. For repeated games we explain how a multiplicity of infinite-horizon equilibria can arise from a multiplicity of finite-horizon ε-equilibria. For games with a finite number of actions in each of an infinite number of periods we show that games of perfect information have perfect equilibria, and give an easily verifiable necessary and sufficient condition for the uniqueness of perfect equilibrium. For several simple examples we explicitly compute the sets of perfect equilibria. Finally, we study a special case of Rubinstein's bargaining game [11], giving a more informative proof of uniqueness than in the original.

Section 2 introduces the notation and contains definitions and a preliminary lemma. Section 3 gives a technical analysis of continuity and the limiting behavior of equilibria. It contains our main result: a necessary and sufficient characterization of infinite-horizon equilibria in terms of finite-horizon truncated equilibria. Section 4 examines repeated games, and discusses some work of Radner. Section 5 analyzes finite-action games, and gives an existence proof for games of perfect information. Section 6 studies the uniqueness of perfect equilibria for several classes of games. Section 7 reviews our findings.

2. Games, Subgames, and Equilibria

This section defines games in extensive and normal form and the associated ε-perfect equilibria. We do not consider the most general definition of a game in extensive form: our focus is on non-stochastic environments with only mild informational imperfections. However, 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 <u>outcome</u> 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 <u>history</u> of the game is a sequence of outcomes $x = (x_1, x_2 \ldots) \in \mathbb{B} \equiv X_{t=1}^{\infty} \mathbb{R}^{MN}$. The <u>action space</u> of the game E is $E^A \subseteq \mathbb{B}$: it is a list of all possible histories of the game. An example helps illustrate this.

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

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.[4]

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:

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

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

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

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 $E^A \subset \mathbb{B}$ satisfies (2-1) and (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 we don't allow mixed strategies.) Let $g_s^i(x)$ denote this choice. Thus for s = 1 $g_s^i \in E^A(0,1)$ while for s > 1 g_s^i is a mapping

(2-3) $g_s^i:E^A(s-1) \to \mathbb{R}^M$ 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:

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

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

We denote the history that occurs when each player plays g^i from the start by $F_{01}(g)$.

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

$$g^1 = (0,0,g_3^1,0,0,\ldots)$$

$$g_3^1 = \begin{cases} 0 & x_2 = (0,1) \\ 1 & x_2 = (0,0) \end{cases}$$

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. This is Selten's [13,14] concept of a (subgame) perfect Nash equilibrium. If players can only approximately solve this optimization problem they may only be able to get within ε of the optimum. Thus a generalization of perfect equilibrium is Radner's [8] concept of a perfect ε-Nash equilibrium.

<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

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

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 $N(\infty)$ 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) :

$$(2\text{-}6)\quad w^T \equiv \sup_{\substack{1 \le i \le N \\ x,z \in E^A \\ x(T-1) = z(T-1)}} |E^{Vi}(x) - E^{Vi}(z)|$$

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

$h^* \equiv g^*(T) \equiv (g_1^*, g_2^*, \ldots, g_T^*, 0, 0, \ldots)$ is $(\varepsilon+2w^T)$-perfect in $N(T)$

proof:

(A) The point is that h^* is ε-optimal against strategies in $N(T)$ while strategies in $N(\infty)$ differ from strategies in $N(T)$ only after time T and thus by (2-6) have payoffs no more than w^T greater than truncated strategies. Formally 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

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

while since h and g differ only after T by definition

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

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

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

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

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

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

Since h* and g* differ only after T

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

and also

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

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

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

and thus h* is $(\varepsilon+2w^T)$-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 $E^A \subset \times_{T=1}^{\infty} \mathbb{R}^{MN} = \mathbb{B}$. The metric

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

induces the product topology on $\mathbb{B}$.[5] 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 game players.

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.

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 $w^T \to 0$.

<u>Definition 3-2</u>: 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.

<u>Lemma 3-1</u>: E uniformly continuous implies E continuous at infinity.

<u>proof</u>:
If, conversely, $w^T \not\to 0$, then there are sequences x^n, z^n with $T(n) \to \infty$, $x^n(T(n)) = z^n(T(n))$, and $|E^V(x^n) - E^V(z^n)| > \delta > 0$. However, from the definition of d, $d(x^n, z^n) \leq (1/(T+1))$ since $x_t^n = z_t^n$ for $1 \leq t \leq T$. Thus by uniform continuity, $|E^V(x^n) - E^V(z^n)| \to 0$, a contradiction. Q.E.D.

Many economically interesting games are uniformly continuous. An example gives a broad class of such games.

Example 3-1 [τ-Markovian Games]:

The valuation functions have the form

$$E^{Vi}(x) = \Sigma_{t=\tau}^{\infty} u_t^i (x_t, x_{t-1}, \ldots, x_{t-\tau})$$

where u_t^i (defined on a suitable domain) are real valued utility functions. Since the actions available to players at a moment of time may depend on history, many interesting games have a Markovian structure. Set

$$\bar{u}_t \equiv \sup_{y \in \text{domain } u_t(\cdot)} |u_t(y)|$$

then it is apparent that if

$$\Sigma_{t=\tau}^{\infty} \bar{u}_t < \infty$$

the game is continuous at infinity.

Example 3-2:

Suppose E is τ-Markovian with $\Sigma_{t=1}^{\infty} \bar{u}_t < \infty$ and the u_t^i are uniformly continuous. Then E is uniformly continuous. To see this let $\varepsilon > 0$ be given and choose T such that $\Sigma_{t=T+1}^{\infty} \bar{u}_t < \varepsilon/2$. Choose δ such that for $|y^1 - y^2| < \delta$ all NT functions u_t^i have $|u_t^i(y^1) - u_t^i(y^2)| < \varepsilon/2T$. Then $d(x^1, x^2) < (\delta/T)$ implies $|x_t^1 - x_t^2| < \delta$ for $t = 1, 2, \ldots, T$

and thus $|E^{Vi}(x^1) - E^{Vi}(x^2)| < \varepsilon$.

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 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

$$(3\text{-}2) \quad d(f,g) \equiv$$

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

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

<u>Lemma 3-2</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 $g^i \in N^{Si}(\infty)$,

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

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

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

contradicting g^n ε-perfect. Q.E.D.

The lemma shows the chosen topology was fine enough to guarantee that the equilibrium set is closed. Of course this would be trivial in the discrete topology, 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 3-3 [Limit Theorem]:

Suppose E is uniformly continuous. Then

(A) A necessary and sufficient condition that g* be perfect in N(∞) 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(∞)).

(B) A necessary and sufficient condition that g* be perfect in N(∞) 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 $g^*(n) \to g^*$. 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 $\varepsilon^n + w^{T(n)} \to 0$, 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.

4. Repeated Games

Much work has been done on repeated games, although for such games the perfect equilibrium concept has been disappointing in failing to isolate a small set of equilibria.[6] Here we briefly describe how our results apply to repeated games.

Recall from Example 2-2 that a repeated game has $E^A = X_{t=1}^{\infty} A$ and $E^{Vi} = \Sigma_{t=1}^{\infty} \beta_i^t u^i(a)$. These games typically have a multiplicity of perfect equilibria. For example, let a* be a Nash equilibrium of the static game and let a' satisfy for all i

(4-1) $u^i(a') \geq u^i(a^*)$ [pareto dominance]

and

(4-2) $(1+\beta_i)u^i(a') \geq \max_{a_i \in A_i} u^i(a_i, a'_{-i}) + \beta_i u^i(a^*)$

Then the strategies "play a'_i as long as everyone else played a'_{-i} in the past otherwise play a^*_i forever" are easily seen to be a perfect equilibrium in the infinite-horizon game. We conclude immediately from Lemma 2-1(B) that a' can also be supported as an $2w^T$-equilibrium in the game repeated T periods only. Alternatively we could easily show that the strategies suggested are

$2w^T$-perfect in the truncated game and conclude that a' can be supported as perfect in the infinite-horizon game.

<u>Example 4-1</u> [The Prisoner's Dilemma]:

If both prisoners confess both get -1. If one confesses and his partner does not he gets +2, his partner -2. If neither confess both get +1. The common discount factor is $0 < \beta < 1$.

For each T the game has a unique perfect equilibrium: both players confess. The reason is simple - confessing is a dominant strategy. In the final period both players must confess. Consequently they must confess in the previous period, and so forth.

Suppose instead both players confess if and only if their opponent confessed in some previous period. If a player cheats in the final period he makes β^T. If he cheats in the next to last period he makes $(1-\beta)\beta^{T-1}$ and if he cheats on the t-th period $t < T$ he makes $\beta^t(1-\beta-\beta^2-\beta^3\ldots-\beta^{T-t})$. Thus, this is an ε^T-perfect equilibrium for $\varepsilon^T \equiv \max\{(1-\beta)\beta^{T-1},\beta^T\}$. Since $\varepsilon^T \to 0$ Theorem 3-3 implies the strategy in question is perfect in the infinite horizon.

The concept of ε-perfect equilibrium was introduced by Radner as a bounded rationality explanation of how the prisoner's dilemma could be resolved in a finite horizon. Theorem 3-3 and Example 4-1 make it clear why the attempt was successful.[7]

5. 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.

This section proves 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 six we use the results of this section in conjunction with Theorem 3-3 to analyze the uniqueness of equilibrium in finite-action games. In general it ought to be possible to use our results to characterize equilibrium sets of finite-action games that arise in applied economic problems. Note that repeated finite games such as the prisoner's dilemma are finite-action games.

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.

<u>Definition 5-2</u>: $\{x^n\} \subset \mathbb{B}$ 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 $1 \leq t \leq T$ $x_t^n = x_t$).

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

<u>proof</u>:

<u>f-convergence implies convergence</u>

Observe that $x^n \overset{f}{\to} x$ iff $T(x^n - x) \to \infty$ where $T(z) \equiv \sup\{T | z_t = 0 \quad t < T\}$. We have

$$d(x,z) \equiv \sup\ [(1/t)\min\{|x_t - z_t|,1\}]$$

$$\leq 1/T(x - z) \qquad (5\text{-}1)$$

Hence $T(x^n - x) \to \infty$ implies $d(x^n,x) \to 0$ and thus f-convergence implies convergence. Q.E.D.

<u>convergence implies f-convergence</u>

Suppose $x^n \to x$ but doesn't f-converge to x. Then there is a subsequence $\{z^n\} \subset \{x^n\}$ and $T \geq 1$ such that $z^n(T) \neq x(T)$. Since $z^n(T) \in E^A(T)$ and $E^A(T)$ is a finite set $d(z^n(T),x(T)) \geq \delta > 0$. However $d(z^n,x) \geq d(z^n(T),x(T)) \geq \delta > 0$ contradicting $z^n \to x$. Q.E.D.

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 $N^S(\infty)$.

Proof of this lemma (which is omitted) parallels that of Lemma 5-1: the essential point is that $N^S(T)$ is a finite set.

We turn now to the compactness of $N^S(\infty)$. First a technical aside. We call a game closed if the action space E^A is closed in $\mathbb{B}$. This rules out degenerate situations such as the one-player with action space $E^A = \{(0,0,\ldots),(1,0,0,\ldots),(1,1,0,0,\ldots)\}$. In this game

(1,1,...) is the limit of (1,0,0,...),(1,1,0,0,...), (1,1,1,0,0,...),... but isn't in E^A. Closedness should be viewed as a technical rather than a substantive assumption.[8]

<u>Lemma 5-4:</u> In a closed finite-action game $N^S(\infty)$ is compact.

<u>proof</u>:

Since $N^S(\infty)$ is a metric space a sufficient condition for compactness is that it be complete and totally bounded.[9]

<u>$N^S(\infty)$ is complete</u>

Let $\{g^n\} \subseteq N^S(\infty)$ be a Cauchy sequence. By definition of the metric for each $t \geq 1$ and history $x \in E^A$ $F_{xt}(g^n)$ is a Cauchy sequence in E^A. Since E^A is closed in the complete space $\mathbb{B}$ it is complete and $F_{xt}(g^n) \to z(x,t) \in E^A$. Define $g \in N^S(\infty)$ by $g_t(x) = (z(x,t))_t$. Since $F_{xt}(g^n) \to z(x,t)$ and for each $T < \infty$ $N^S(T)$ is a finite set, g^n converges finitely to g and thus by Lemma 5-3 $g^n \to g$.

<u>$N^S(\infty)$ is totally bounded</u>

Fix $\varepsilon > 0$. We must find a finite subset of $N^S(\infty)$ such that every $g \in N^S(\infty)$ is within ε of some member of the subset. Choose $T > (1-\varepsilon)/\varepsilon$ or equivalently $\varepsilon > 1/(T+1)$. We claim $N^S(T)$ is the desired subset. It is certainly finite since it is a finite sequence of mappings with finite domain and range. If $g \in N^S(\infty)$ then $g(T) \in N^S(T)$. By definition for $1 \leq t \leq T$ $g_t = (g(T))_t$. Thus by definition of the metric on $N^S(\infty)$ given in (3-1) and (3-2), $d(g,g(T)) \leq 1/(T+1) < \varepsilon$. Thus $N^S(\infty)$ is totally bounded. Q.E.D.

To conclude the section we 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) closed 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.

6. 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 finitely determined (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.

proof:

f.d. implies existence

f.d. implies the existence of a sequence of $2w^T$-perfect equilibria with $T \to \infty$. By Lemma 5-4 this sequence has a convergent subsequence, and by Theorem 3-3 the limit point is perfect in $N(\infty)$.

f.d. implies uniqueness

Let g' and g^* be two infinite-horizon perfect equilibria. By Lemma 2-1 for any $T > 0$ $g^*(T)$ and $g'(T)$ are $2w^T$-perfect in $N(T)$. By f.d. this means for any $k > 0$ $g^*(k) = g'(k)$ and thus $g^* = g'$.

existence and uniqueness imply f.d.

Assume conversely that g^* is the unique perfect equilibrium of a game which isn't finitely determined. By Lemma 2-1 $g^*(T)$ is $2w^T$-perfect in $N(T)$. Since the game isn't finitely determined, there is $k > 0$ such that for $T \geq k$ we can find g^T $2w^T$-perfect in $N(T)$ with $g^T(k) \neq g^*(k)$. $N(\infty)$ is compact, so $\{g^T\}$ has a convergent subsequence $\{g^n\}$, and by Theorem 3-3 this subsequence converges to g^*. But $g^n(k) \neq g^*(k)$, implying that g^n doesn't converge finitely to g^*, a contradiction.Q.E.D.

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).

Let us show that this game is finitely determined in two cases

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

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

and that it is <u>not</u> finitely determined in the complementary cases

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

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

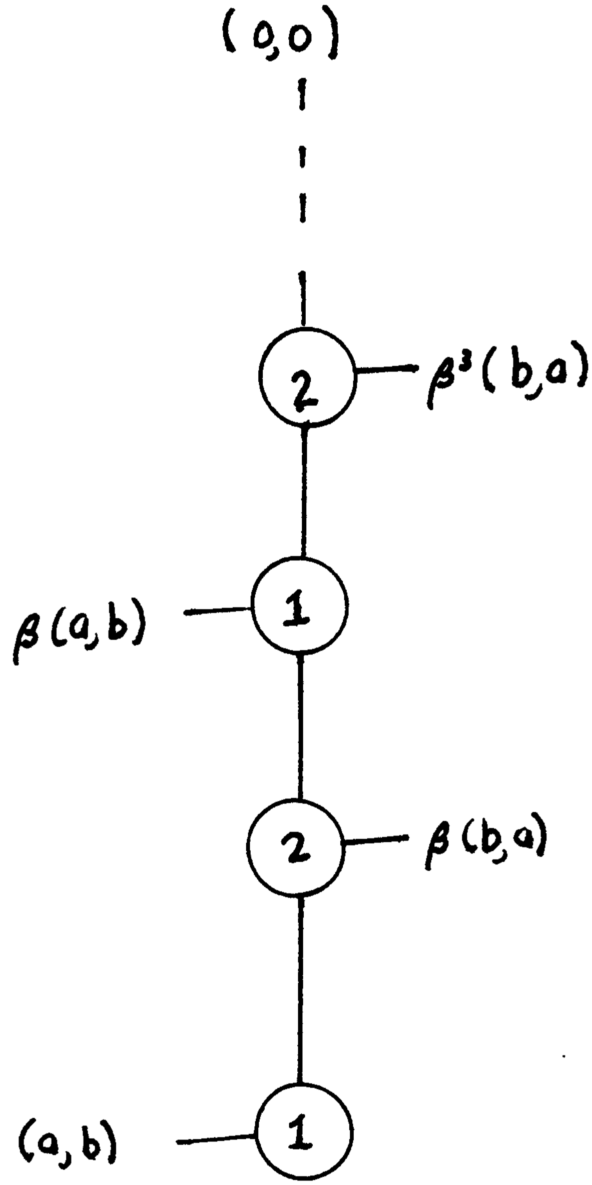

FIGURE 6-1

To do this note that there is only one history at each node which allows players a choice: if the game has stopped, both players must choose the null action. Thus 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. Note that the truncated games are assumed to continue forever if they haven't stopped by the truncation point.

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 $\varepsilon < \beta^k \min (a - \beta^2 a, a - \beta b)$ all ε-equilibria must terminate at all times before k. Since $w^T \to 0$ 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 (ii) is a game which both players wish to last as long as possible. Arguing as in case (i) the minimum loss from stopping at k is $\beta^k \min(\beta b - a, \beta^2 a - a)$. Thus for k fixed and T large enough the $2w^T$-perfect equilibria require that neither player stop before k.

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.

Case (iv) is a game in which it is pareto optimal never to stop, but if your opponent decides to stop, you'd prefer to stop first. It is like a "reverse prisoner's dilemma": the finite horizon truncated games have a unique equilibrium "never stop" (which is pareto optimal) while in the infinite-horizon game the pareto inferior "always stop" is also a perfect equilibrium. Naturally the finite-horizon truncated strategies "stop in every period before the horizon" are $2w^T$-perfect; indeed, the only sub-optimal play involved is by the last player to move - he loses no more than $\beta^{T+1}b = w^T$. Thus 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 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 example is 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 accepts 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 $\beta^{1-k}(\beta^k \underline{S}^k - 2w^T)$. 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

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

rejects proposals offering him less than

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

this implies that

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

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

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.

7. 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] The theorem may also extend to the sequential equilibria of Kreps-Wilson [4]. 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 McClellan, Ariel Rubinstein and Jean Tirole for helpful conversations. Joe Farrell and Franklin M. Fisher provided useful comments on an earlier draft.

2. A similar theorem holds for overlapping generations competitive equilibrium. See Balasko, Cass and Shell [1].

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 McClellan for providing this example, which helped clarify our thinking in the early stages of our investigation.

5. See Munkres [8] p. 123.

6. This is pointed out by Rubinstein [10] in the context of supergames. It remains true in repeated games with discounting. More details on repeated games can be found in Friedman [2].

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.

8. In finite-action games, continuity at infinity implies both uniform continuity and that the valuation functions are bounded. The valuation functions can then be extended to the closure of E^A.

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] Balasko, Y., D. Cass and K. Shell: "Existence of Competitive Equilibrium in a General Overlapping Generations Model", Journal of Economic Theory, 23 (1981), 307-322.

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

[3] Fudenberg,D. and J. Tirole: "Capital as a Commitment: 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", Stanford University Graduate School of Business R.P. No. 584, 1980.

[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.