

Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games*, †

DREW FUDENBERG

Department of Economics, U. C. Berkeley, Berkeley, California 94720

AND

DAVID LEVINE

Department of Economics, U.C.L.A., Los Angeles, California 90024

Received May 26, 1982; revised February 17, 1983

We show that subgame-perfect equilibria of infinite-horizon games arise as limits, as the horizon grows long and epsilon small, of subgame-perfect epsilon-equilibria of games which are truncated after a finite horizon. A number of applications show that this result provides a useful technique for analyzing the existence and uniqueness of infinite-horizon equilibria. We extend our result to the sequential equilibrium concept. *Journal of Economic Literature* Classification Numbers: 022, 026, 611.

1. INTRODUCTION

The concept of sequential rationality has been useful in understanding a wide range of economic problems. The idea is that the agents will not be misled by opponents' threats, but will instead compute their opponents' future actions from their knowledge of the structure of the game. This rationality requirement was first formalized by Selten [12] as subgame-perfect equilibrium.¹ With a finite horizon, such equilibria are ordinarily

* It is our pleasure to thank Robert Anderson, Timothy Kehoe, Eric Maskin, Andrew McLennan, Ariel Rubinstein and Jean Tirole for helpful conversations. Joe Farrell, Franklin M. Fisher, David Kreps and the referees provided useful comments on an earlier draft.

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

¹ Selten [13] introduced the more restrictive notion of "trembling-hand" perfection. Kreps-Wilson's [4] "sequential equilibrium" is a mathematically more convenient version of the trembling hand, since the sequential and trembling-hand equilibria coincide for generic payoffs. While we title our paper "Subgame-Perfect Equilibria...", due to the information structures we consider *all* of the equilibria we discuss will be sequential (in particular *not* just those of Section 6).

computed by backwards induction; with an infinite horizon determining the perfect equilibria is more difficult. However, specification of a fixed horizon is often artificial, and frequently an infinite-horizon game better captures the economics of a situation. This paper describes a method for characterizing equilibria of infinite-horizon games.

We do not consider the most general extensive-form games. Our formulation does allow simultaneous moves, and, in Section 6, 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 games considered in the economics literature.

The technique we propose is to study ε^T -equilibria in the game truncated after T periods of play. Here we follow Radner [8] 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 $T \rightarrow \infty$ and $\varepsilon^T \rightarrow 0$ the set of ε^T -equilibria in the truncated games converges to the set of 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. 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 [10] bargaining game, giving a more informative proof of uniqueness than in the original. In Section 6 we extend our results to allow for mixed strategies and unobservable moves by nature. We characterize infinite-horizon sequential equilibria as limit points and show that they exist. Section 7 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 $t = 1, 2, \dots$. Each period all N players simultaneously choose actions from feasible sets of actions, which we take to be subsets of R^M . When they choose an action in period t they know the entire history of the game until and including time $t - 1$.² It is possible that the set of feasible actions is constrained by the history of play.

The *period- t outcome* of the game lies in R^{MN} . The way in which the period- t outcome is made up of individual actions is discussed below. The *outcome of the game* is a sequence of period- t outcomes $x = (x_1, x_2, \dots) \in B \equiv X_{t=1} R^{MN}$. The *outcome space* of the game is $X \subseteq B$: it is a list of all possible outcomes of the game. Note that not *every* sequence of period t outcomes is in X ; i.e., X is not necessarily a product space. An example helps illustrate this.

EXAMPLE 2.1 (McLENNAN'S TERMINATION GAME³). 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 $y_1, y_2 \in \{0, 1\} \subset R$. A player must choose 0 if it isn't his move, or if the game has already terminated. Thus the outcome space X is the set of sequences of the form $((0, 0)_1, (0, 0)_2, \dots, (1, 0)_t, (0, 0)_{t+1}, \dots)$, where t is odd, $((0, 0)_1, (0, 0)_2, \dots, (0, 1)_t, (0, 0)_{t+1}, \dots)$, where t is even, or $((0, 0)_1, (0, 0)_2, \dots)$.

It is generally useful and entails no loss of generality to designate the outcome $0 \in R^{MN}$ 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 $t + 1, t + 2, \dots$, is also feasible:

$$\forall x \in X \forall t \ x(t) \equiv (x_1, x_2, \dots, x_t, 0, 0, \dots) \in X. \quad (2.1)$$

² Thus while we do not restrict attention to games of perfect information, the information structure is "almost" perfect in that at the end of each period there is no uncertainty. In Section 6 we extend our results to games in which Nature may make moves that are not completely observable; we will, however, continue to assume that the past actions of the players are common knowledge. For more general definitions of extensive-form games see Luce-Raiffa [6] or Kreps-Wilson [4].

³ We are grateful to Andrew McLennan for providing this example, which helped clarify our thinking in the early states of our investigation.

Let $X(x, s)$ be the space of all possible outcomes in period s consistent with the history x_1, x_2, \dots, x_{s-1} , with the convention that $X(0, 1)$ 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 $(x_1, x_2, \dots, x_{s-1}, y, 0, 0, \dots) \in X$ since if $z = (x_1, \dots, x_{s-1}, y, z_{s+1}, z_{s+2}, \dots) \in X$ then $z(s) = (x_1, \dots, x_{s-1}, y, 0, 0, \dots) \in X$ as well.

Furthermore, although past outcomes may restrict current choices we require that a sequence feasible at each point in time is actually feasible:

$$\text{If } \forall t \ x(t) \in X \quad \text{then} \quad x \in X. \quad (2.2)$$

(See footnote 4 below.)

If X 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 $X^i(x, t)$, 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 $X(x, t)$ is the cartesian product of the individual action spaces

$$\forall x \in X \ \forall t \ X(x, t) = X_{i=1}^N X^i(x, t). \quad (2.3)$$

Thus in Example 2.1 the set of possible outcomes at time 2 if the game has not yet been terminated, $X(0, 2) = \{(0, 0), (0, 1)\}$, is the cartesian product of $X^1(0, 2) = \{0\}$ and $X^2(0, 2) = \{0, 1\}$.

DEFINITION 2.1. A game in extensive form is a pair (X, V) , where $X \subset B$ satisfies (2.1), (2.2), (2.3) and $V = (V^i)_{i=1}^N$ is an N -tuple of valuation functions $V^i : X \rightarrow R$ assigning a value to each history of the game.

In Example 2.1 where $z^1 = ((1, 0)_1, (0, 0)_2, \dots)$ and $z^2 = ((0, 0)_1, (0, 1)_2, (0, 0)_3, \dots)$, $V(0) = (0, 0)$, $V(z^1) = (\beta a, \beta b)$, and $V(z^2) = (\beta^2 b, \beta^2 a)$.

EXAMPLE 2.2 (REPEATED GAMES). Each agent i has a fixed set of actions $0 \in A^i \subset R^M$, a utility function $U^i : A \rightarrow R$, 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 $X \equiv X_{i=1}^\infty A$ so that history places no constraints on behavior. The valuation functions are $V^i(x) = \sum_{t=1}^\infty \beta_i^t U^i(x_t)$.

Further examples are given in subsequent sections.

Associated with each game in extensive form is a collection of truncated games in normal form. The normal form of the game truncated at time T by assigning the null outcome to all following periods is specified by giving its strategy space $S(T)$. Let us formally describe this space. At time s player i , knowing the history x_1, x_2, \dots, x_{s-1} , must choose a feasible action in $X(x, s)$ to undertake in period s . (Note that for now we do not allow mixed

strategies.) Let $g_s^i(x)$ denote this choice. Thus for $s = 1$, $g_s^i \in X^i(0, 1)$, while for $s > 1$, g_s^i is a mapping

$$g_s^i : X(s-1) \rightarrow R^M \quad \text{with} \quad g_s^i(x) \in X(x, s)$$

where $X(t)$ denotes all possible histories to time t , i.e., all vectors $(x_1, x_2, \dots, x_t, 0, 0, \dots) \in X$. A complete set of contingent choices of this type is called a strategy and is simply a sequence $(g_1^i, g_2^i, \dots, g_T^i, 0, 0, \dots)$, where $g_1^i \in X(0, 1)$ and for $s > 1$ g_s^i is as above. The set of all such strategies is called the strategy space for player i and is denoted $S^i(T)$. The strategy space for the game truncated at time T is just the cartesian product $S(T) \equiv \prod_{i=1}^N S^i(T)$. Note that $S(1) \subseteq S(2) \subseteq \dots \subseteq S(\infty)$. This will allow us to use the valuation functions of the untruncated games to assign payoffs to the truncated games. 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 $x^s(g)$ assigns a strategy selection $g \in S(\infty)$ the outcome of the game that occurs when the initial history is x_1, \dots, x_{s-1} and afterwards each player plays g^i :

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

$$z_t = \begin{cases} x_t & 1 \leq t \leq s-1 \\ g_t^i(z_1, z_2, \dots, z_{t-1}, 0, 0, \dots) & t \geq \max(s, 2). \end{cases} \quad (2.4)$$

We denote the outcome that occurs when each player plays g^i from the start by $x^0(g)$. Note that $x^s(\cdot) \in X$ follows from (2.2).⁴

To illustrate these definitions consider in Example 2.1 the strategy by player one to "terminate in period three unless player two has already terminated, after period three don't terminate" which has the form

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

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

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

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

⁴ Consider the one-player game with outcome space X consisting of any sequence of *finitely* many "1's" followed by "0's." Then $x_1 = (1, 0, 0, \dots) \in X$, $x_2 = (1, 1, 0, 0, \dots) \in X$, but $\lim_{n \rightarrow \infty} x_n = (1, 1, \dots) \equiv x_\infty \notin X$. One possible strategy, however, is for $g_n^1(x_{n-1}) = 1$, that is, if 1 has always been played before then play it again: the outcome of this strategy would be a sequence of all 1's, that is, $x_0(g) = x_\infty \notin X$. This pathology is avoided since (2.2) requires that $x_\infty \in X$.

Then for any x , $x^0(g)$ is the outcome that actually occurs, so

$$x^0(g) = ((0, 0)_1, (0, 0)_2, (1, 0)_3, (0, 0)_4, \dots),$$

while if the history before time 4 is $y_1 = (0, 0)$, $y_2 = (0, 0)$, and $y_3 = (0, 0)$, then $y^4(g) = 0$. In other words if one reneges on his plan to terminate in period 3 neither player ever terminates. Finally, for

$$z = ((0, 0)_1, (0, 1)_2, (0, 0)_3, \dots)$$

(so that two *does* terminate in period 2) $z^3(g) = z$ and one *must* (and does) choose the null action in period 3.

We turn now to equilibrium in the games $S(T)$. 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's [13] concept of a subgame-perfect Nash equilibrium. (Note that the subgame perfect and sequential equilibria coincide with the given information structure.) 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.⁵

DEFINITION 2.2. $g^* \in S(T)$ is a subgame-perfect ε -Nash equilibrium (or simply ε -perfect) if for each $s \geq 0$, history x , strategy $g \in S(T)$ and player i ,

$$V^i(x^s(g^i, g^{*-i})) - V^i(x^s(g^*)) \leq \varepsilon; \quad (2.5)$$

that is, if after no initial history 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 $g, g^* \in S(T)$, 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 games to perfect equilibria of the infinite game. To this end define the constants w^T to be the greatest variation in any player's payoff due strictly to events after $(T-1)$:

$$w^T \equiv \sup_{\substack{1 \leq i \leq N \\ x, z \in X \\ x(T-1) = z(T-1)}} |V^i(x) - V^i(z)|. \quad (2.6)$$

⁵ 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.

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

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

LEMMA 2.1.

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

(B) g^* ε -perfect in $S(\infty)$ then $h^* \equiv g^*(T) \equiv (g_1^*, g_2^*, \dots, g_T^*, 0, 0, \dots)$ is $(\varepsilon + 2w^T)$ -perfect in $S(T)$.

The point is that strategies in $S(\infty)$ differ from strategies in $S(T)$ only after time T and thus by (2.6) have payoffs within w^T of the truncated strategies. Formally we just add inequalities.

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

$$V^i(x^s(h^i, h^{*-i})) - V^i(x^s(h^*)) \leq \varepsilon \quad (2.7)$$

while since h and g differ only after T , by definition

$$V^i(x^s(g^i, h^{*-i})) - V^i(x^s(h^i, h^{*-i})) \leq w^T. \quad (2.8)$$

Adding (2.7) to (2.8) shows

$$V^i(x^s(g^i, h^{*-i})) - V^i(x^s(h^*)) \leq \varepsilon + w^T. \quad (2.9)$$

Since g , x , and s are arbitrary (2.9) implies h^* is $(\varepsilon + w^T)$ -perfect.

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

$$V^i(x^s(h^i, g^{*-i})) - V^i(x^s(g^*)) \leq \varepsilon. \quad (2.10)$$

Since h^* and g differ only after T

$$V^i(x^s(g^*)) - V^i(x^s(h^*)) \leq w^T, \quad (2.11)$$

and also

$$V^i(x^s(h^i, h^{*-i})) - V^i(x^s(h^i, g^{*-i})) \leq w^T. \quad (2.12)$$

Adding (2.10), (2.11), and (2.12) shows

$$V^i(x^s(h^i, h^{*-i})) - V^i(x^s(h^*)) \leq \varepsilon + 2w^T, \quad (2.13)$$

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 $S(\infty)$ if and only if it is the limit as $T \rightarrow \infty$ and $\varepsilon^T \rightarrow 0$ of ε^T -perfect equilibria in $S(T)$. 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 X and $S(\infty)$.

Recall that $X \subset X_{T=1}^\infty R^{MN} = B$. The metric

$$d(x, z) \equiv \sup_T [(1/T) \min\{|x_T - z_T|, 1\}] \quad (3.1)$$

induces the product topology on B .⁶ Hereafter all statements about continuity, convergence, etc., will be with respect to this topology (relativized to X).

Having introduced a topology on X we now discuss continuity of the valuation function $V: X \rightarrow R^N$, which we refer to as continuity of the game. Continuity implies that 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. V is *uniformly continuous* if for all $x^n, z^n \subset X$, $(x^n - z^n) \rightarrow 0$ implies $|V(x^n) - V(z^n)| \rightarrow 0$.

Although we shall only be interested in uniformly continuous games, this restriction may not be necessary for our limit theorem. However, many games of interest to economists are uniformly continuous.

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 $w^T \rightarrow 0$.

DEFINITION 3.2. (X, V) is *continuous at infinity* iff $w^T \rightarrow 0$ as $T \rightarrow \infty$.

An important fact is that uniform continuity implies continuity at infinity.

LEMMA 3.1. (X, V) *uniformly continuous implies (X, V) continuous at infinity.*

This follows simply from unwinding the definitions.

A supergame has w^T constant over time and is not continuous at infinity. Thus our analysis will not apply to super-games. A repeated game (Example

⁶ See Munkres [7, p. 123].

2.2) with discount factor $1 > \beta > 0$ has $w^T = \beta w^{T-1}$ and is continuous at infinity provided $w^1 < \infty$.

Finally, we must extend our notion of convergence in X to the strategy space $S(\infty)$ (and implicitly to its subsets $S(T)$ $T < \infty$). 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 $x \in X$ the histories resulting from f and g being played are close and 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

$$d(f, g) \equiv \sup_{x \in X, t} \{d(x^t(f), x^t(g)), \sup_{h^i \in S^i(\infty)} [d(x^t(h^i, f^{-i}), x^t(h^i, g^{-i}))]\}. \quad (3.2)$$

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

LEMMA 3.2. *Let g_n be ε -perfect in $S(\infty)$ and $g_n \rightarrow g$ in a continuous game. Then g is also ε -perfect.*

Proof. Suppose g is not ε -perfect so that for some t , some $x \in X$, and some $\tilde{g}^i \in S^i(\infty)$,

$$V^i(x^t(\tilde{g}^i, g^{-i}) - V^i(x^t(g)) \geq \varepsilon + 3\delta. \quad (3.3)$$

Since $g_n \rightarrow g$, for large n , $x^t(g)$ is near $x^t(g_n)$, and $x^t(\tilde{g}^i, g^{-i})$ is near $x^t(\tilde{g}^i, g_n^{-i})$. As V^i is continuous, for any $\delta > 0$ we have that for large enough n ,

$$V^i(x^t(g_n)) - V^i(x^t(g)) < \delta \quad (3.4)$$

$$V^i(x^t(\tilde{g}^i, g^{-i})) - V^i(x^t(\tilde{g}^i, g_n^{-i})) < \delta. \quad (3.5)$$

Combining (3.3), (3.4), and (3.5), we have, for large enough n ,

$$V^i(x^t(\tilde{g}^i, g_n^{-i})) - V^i(x^t(g_n)) > \varepsilon + \delta. \quad (3.6)$$

As δ can be taken to be arbitrarily small this contradicts g_n ε -perfect.

Q.E.D.

The lemma shows that chosen topology was fine enough to guarantee that the ε -equilibrium sets are closed. Of course we could simply have declared them 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 the strategy, is

THEOREM 3.3 (LIMIT THEOREM). *Suppose V is uniformly continuous. Then*

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

(B) *A necessary and sufficient condition that g^* be perfect in $S(\infty)$ is that there be sequences ε_n , $T(n)$, and g_n such that g_n is ε_n -perfect in $S(T(n))$ and as $n \rightarrow \infty$, $\varepsilon_n \rightarrow 0$, $T(n) \rightarrow \infty$, and $g_n \rightarrow g^*$.*

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 $\{g^*(n)\}$, $g^*(n) = (g_1^*, g_2^*, \dots, g_n^*, 0, 0, \dots)$ 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 (3.1) and (3.2)). Thus $g^*(n) \rightarrow g^*$. By Lemma 2.1(B) we also have $g^*(n)$ $2w^{T(n)}$ -perfect in $S(n)$.

(B) *Sufficient:* By Lemma 2.1(A) g_n is $(\varepsilon_n + w^{T(n)})$ -perfect in $S(\infty)$. 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 repeated games. Our analysis applies to such games if payoffs of the single-period game are continuous in the single-period actions (in particular to repeated games with finitely many actions) and if the discount factor is strictly less than one. In the infinite horizon these games are known to have a plethora of equilibria when the discount factor is sufficiently close to one. Thus we can conclude there are a great multiplicity of ε -equilibria in the finite horizon. Since in the *finite* horizon case (for generic payoffs) the ε -equilibrium set is well behaved with respect to the discount factor, we can conclude that in the limit as the discount factor reaches one (finite time-averaging) there continue to be almost the same plethora of ε -equilibria. This allows, for example, a finite-horizon resolution of the prisoner's dilemma. Such a resolution was the reason Radner originally introduced the concept of ε -equilibrium.

While our theorem can thus be applied via a second limit to finite-horizon time-averaging, this trick will not work for infinite-horizon time-averaging. These are games without impatience, in which only the long-run matters—the opposite of our case where the future is unimportant. The problems as the discount factor approaches one in the infinite horizon are suggested by the example of the “cake-eating” problem—a single agent must choose a consumption path $\{x_t\}$ such that $\sum_{t=1}^{\infty} x_t = 1$, $x_t \geq 0$. The agent's payoff is $V(x) = \sum_{t=1}^{\infty} \beta^t U(x_t)$, with U concave. For $\beta < 1$ optimal

consumption declines geometrically. However, with time averaging no solution exists. (Note, however, that with a finite horizon the limit is well behaved.)

4. FINITE-ACTION GAMES

Finite-action games are games in which there are only a finite number of possible actions in each period. This section introduces finite-action games and proves that with perfect information perfect equilibria exist. In Section 5 we use the results of this section in conjunction with Theorem 3.3 to analyze the uniqueness of equilibrium in these games. Section 6 analyzes sequential equilibrium of finite-action games with uncertainty and mixed strategies.

DEFINITION 4.1. (X, V) is a *finite-action game* iff for each t and history $x \in X$ the set of outcomes in period t given the history x , $X(x, t)$ is a finite set.

Convergence in finite-action games is easily described. First, a sequence of realizations converges if and only if the constituent outcomes eventually coincide for the first T periods for any T . As an immediate consequence we have

LEMMA 4.1. *In finite-action games uniform continuity and continuity at infinity are equivalent.*

Similarly, a sequence of strategies converges if and only if its components eventually coincide.

We turn now to the compactness of $S(\infty)$. A useful way to study this problem is to observe that $S(\infty)$ is the space of sequences of maps (g_1, g_2, \dots) . The map g_t has a finite domain, with, say, L_t elements, and ranges in R^{NM} . Thus it may be viewed simply as a vector in $B_t^* \equiv R^{NML_t}$, and $S(\infty) \subset B^* \equiv \prod_{t=1}^{\infty} B_t^*$ in a natural way. Furthermore it is easy to see that the topology in $S(\infty)$ given in (3.1) and (3.2) is the same as the relative product topology in B^* : in both cases convergence means that for any fixed horizon the sequence is eventually stationary before that horizon. However $S(\infty)$ is the cartesian product of finite (and thus compact) subsets of the B_t^* implying that it is itself compact and proving

LEMMA 4.2. *In finite-action games $S(\infty)$ is compact.*

We can use Lemma 4.2 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 $x \in X$, there is a player i such that $X^{-i}(x, t) = 0$; only player i faces a decision.

It is well known and can 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

COROLLARY 4.2. *Continuous (at infinity) finite-action games of perfect information have perfect equilibria.*

Proof. Each finite-horizon subgame $S(T)$ has a perfect equilibrium g^T . By Lemma 2.1(A) g^T is w^T -perfect in $S(\infty)$. Since $S(\infty)$ is compact there is a subsequence $\{h^T\} \subset \{g^T\}$ with $h^T \rightarrow g^* \in S(\infty)$. By Theorem 3.3(B) this implies g^* is perfect in $S(\infty)$. Q.E.D.

5. UNIQUENESS OF THE INFINITE-HORIZON PERFECT EQUILIBRIUM

This section uses the limit theorem of Section 3 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 $T \rightarrow \infty$. Note that this implies that a necessary condition for uniqueness is that every convergent sequence of perfect equilibria of the truncated games have the same limit.

The first class of games we consider are the finite-action games of Section 4. Recall that in such games a sequence of strategies converges if and only if they eventually agree prior to each fixed finite horizon. 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 5.1. 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 $S(T)$,
- (b) if g' is $2w^T$ -perfect in $S(T)$ and $k \geq t > 0$, $g_t = g'_t$.

PROPOSITION 5.1. *There exists a unique infinite-horizon perfect equilibrium in a finite-action game that is continuous at infinite if and only if it is finitely determined.*

Proof. Omitted, see [2].

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. Player one moves in odd periods, player two in even periods. Each period the player moving

chooses whether to “terminate” or “continue.” If the game terminates in period 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) $a > 0 \quad a > \beta b$

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

and it is *not* finitely determined in the complementary cases

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

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

We show this for cases (i) and (iii). Note that a strategy may be viewed as a choice the period to stop at (if the game hasn't stopped already). For example, if T is even, “stop at $T, T-2, T-4, \dots$ ” 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.

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, in earlier periods 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 $w^T \rightarrow 0$ 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 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 5.1 (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 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 $T \rightarrow \infty$ and $w^T \rightarrow 0$. 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 $S(T)$ 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:

$$\beta^k[1 - \beta^j(1-s) - \beta^{-k}\varepsilon(k)] \leq \beta^{k+j}s + \varepsilon(k). \quad (5.1)$$

This implies

$$\varepsilon(k) \geq \beta^k(1-\beta)/2 \quad (5.2)$$

which contradicts our assumption. Since $w^T \rightarrow 0$ 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 irresistible they won't be

less than $\beta^{1-k}(\beta^k \underline{S}^k - 2w^T)$ and two certainly won't be offered 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 $\beta^{2-k}\{\beta^{k-1}[1 - \beta^{1-k}(\beta^k \underline{S}^k - 2w^T)] + 2w^T\}$ and rejects proposals offering him less than

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

As before this implies that

$$\begin{aligned} \bar{S}^{k-2} &= 1 - \beta^{2-k}\{\beta^{k-1}[1 - \beta^{1-k}(\beta^k \bar{S}^k + 4w^T)] - 2w^T\} \\ \underline{S}^{k-2} &= 1 - \beta^{2-k}\{\beta^{k-1}[1 - \beta^{1-k}(\beta^k \underline{S}^k - 2w^T)] + 4w^T\}. \end{aligned} \tag{5.3}$$

The claim we wish to establish is that as $T \rightarrow \infty$ $\bar{S}^k, \underline{S}^k \rightarrow 1/(1 + \beta)$. Since the mapping in (6.3) is a contraction as we work it backwards from period $k + j, j$ large, \bar{S}^k approaches $[1/(1 + \beta)] + C_j^k w^T$ and \underline{S}^k approaches $[1/(1 + \beta)] - C_j^k w^T$. Letting $w^T \rightarrow 0$ and noticing that C_j^k is independent of T yields the desired conclusion.

6. SEQUENTIAL EQUILIBRIA

In this section we extend our analysis to allow for uncertainty and mixed strategies in finite-action games. We use Kreps–Wilson's [4] concept of a sequential equilibrium to model rationality in this setting. In each period t "nature" makes a random move θ_t . Players are only partially aware of the result of this move. Player i observes only a signal $\theta_t^i \in \Theta_t^i \subset R^M$, where Θ_t^i is a finite set. For simplicity we assume $\theta_t = (\theta_t^1, \theta_t^2, \dots, \theta_t^N)$ so that players would know θ_t if they pooled their knowledge. Any additional uncertainty (on the part of all players) can be incorporated directly into payoffs via expected utility. Naturally the payoff functions V^i are defined on $\Theta \times X$, where $\Theta \equiv \prod_{i=1}^N \prod_{t=1}^{\infty} \Theta_t^i$. We give $\Theta \times X$ the product topology and define continuity at infinity and uniform continuity over $\Theta \times X$ rather than just X .

Decisions by agents are based on their probabilistic beliefs about past and future values of θ_t ; these beliefs reflect private information θ_t^i and information revealed by the play of opponents as revealed by the history of the game. Let $\Theta^i(T) \equiv \prod_{t=1}^T \Theta_t^i$. A system of beliefs for player i is sequence of mappings μ_s^i with domain $\Theta^i(s - 1) \times X(s - 1)$ and ranging over the space of probability measures on $\Theta(s - 1)$; it represents beliefs about past outcomes given current information. Since $\Theta(s - 1)$ and $X(s - 1)$ are finite μ_s^i may be viewed as a vector in a finite-dimensional vector space $B_s^{\mu^i}$.

To characterize an agent's play requires specifying both his beliefs and the strategy he chooses. Analogous to our previous definition a strategy for player i is a sequence of mappings g_s^i with domain $\Theta^i(s - 1) \times X(s - 1)$.

Now, however, we wish to allow mixed strategies so that g_s^i ranges over the space of *probability measures* on $X^i(x, s)$. Since $\Theta^i(s-1) \times X(s-1)$ and $X^i(x, s)$ 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 *assessment*: 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 $A(\infty)$. Just as $S(\infty) \subset X_{t=1}^\infty B_t^*$ so $A(\infty) \subset X_{t=1}^N X_{t=1}^\infty (B_t^{\mu^i} \times B_t^{g^i}) \equiv B^*$. The product topology on B^* then introduces a corresponding topology on $A(\infty)$. 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 truncated assessment $A(T) \subset A(\infty)$ in which g_s^i places unit probability weight on action zero for $s > T$. Finally, we may define $U_{\theta, x, s}^i(s)$ as the expected utility accruing to player i at time s when a 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 setup we can define a sequential ε -equilibrium, following Radner [8] and Kreps–Wilson [4].

DEFINITION 6.1. A sequential ε -equilibrium is an assessment selection (μ^*, g^*) such that

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

$$U_{\theta, x, s}^i(\mu^*, g^i, g^{*-i}) - U_{\theta, x, s}^i(\mu^*, g^*) < \varepsilon.$$

(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 3.2 showing that the set of ε -equilibria is closed, and Theorem 3.3 itself. With the exception of Lemma 3.2 all proofs go through verbatim by merely changing the notation to replace “strategies” by “assessments.” Lemma 3.2 follows quite easily from part (2) of Definition 6.1: 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 $A(\infty)$ is the product of compact sets in the product topology it is itself compact. From Kreps–Wilson we know there is a sequential equilibrium a^T in each $A(T)$. Since

$A(\infty)$ is compact these have a subsequence converging to $a^* \in A(\infty)$. By the limit theorem a^* is a sequential equilibrium. Thus we have demonstrated

THEOREM 6.1. *Continuous (at infinity) finite-action games with imperfect information have mixed-strategy sequential equilibria.*

7. CONCLUSION

In games which satisfy an economically appealing continuity requirement, infinite-horizon equilibria coincide with the limits (as $T \rightarrow \infty$) 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.^{7,8} As a technical matter all that is required is to prove an analog of Lemma 2.1 and to find some reasonable notion of the convergence of strategies.

REFERENCES

1. J. FRIEDMAN, "Oligopoly and the Theory of Games," North-Holland, Amsterdam, 1977.
2. D. FUDENBERG AND D. LEVINE, Perfect equilibria of finite and infinite horizon games, UCLA Working Paper No. 216.
3. D. FUDENBERG AND J. TIROLE, Capital as a Commitment: Strategic Investment to Deter Mobility, *J. Econ. Theory* **31** (1983), 227–250.
4. D. M. KREPS AND R. WILSON, Sequential equilibria, *Econometrica* **50** (1982), 863–894.
5. D. LEVINE, Local almost perfect equilibrium in a game with adjustment costs, Essay 3 in "The Enforcement of Collusion in Oligopoly," Ph.D. dissertation, Massachusetts Institute of Technology, 1981.
6. R. P. LUCE AND H. RAIFFA, "Games and Decisions," Wiley, New York, 1957.
7. J. R. MUNKRES, "Topology, A First Course," Prentice-Hall, Englewood Cliffs, N. J., 1975.
8. R. RADNER, Collusive behavior in non-cooperative epsilon-equilibria of oligopolies with long but finite lives, *J. Econ. Theory* **22** (1980), 136–154.

⁷ We thank A. Rubinstein for pointing this out. See [11] for a treatment of strong perfectness in supergames.

⁸ For examples of such games see Fudenberg and Tirole [3] or Levine [5].

9. A. RUBINSTEIN, Equilibrium in supergames with the overtaking criterion, *J. Econ. Theory* **21** (1979), 1–9.
10. A. RUBINSTEIN, Perfect equilibrium in a bargaining model, *Econometrica* **50** (1982), 97–110.
11. A. RUBINSTEIN, Strong perfect equilibrium in supergames, *Internat. J. Game Theory* **9** (1979), 1–12.
12. R. SELTEN, Spieltheoretische behandlung eines oligopolmodells mit nachfragetraheit, *Z. Gesamte Staatwissenschaft* **12** (1965), 301–324.
13. R. SELTEN, Re-examination of the perfectness concept for equilibrium points in extensive games, *Internat. J. Game Theory* **1** (1975), 25–55.
14. L-G. SVENSON, Equity among generations, *Econometrica* **48** (1980), 1251–1256.