The Folk Theorems for Repeated Games: A Synthesis

Jean-Pierre Benoît*
Department of Economics
New York University

Vijay Krishna*
Department of Economics
Penn State University

February 9, 1996

Abstract

We present a synthesis of various folk theorems for repeated games.

1 Introduction

The theory of repeated games occupies a central place in noncooperative game theory as it forms a relatively simple platform from which to study dynamic aspects of strategic interaction. The key results concerning repeated games, often called 'folk theorems,' delineate the set of equilibrium outcomes in situations where the future looms large in players' assessments of their prospects. Typically, the folk theorems show that under these circumstances the set of equilibrium outcomes is essentially unrestricted.

There are numerous results of this genre, varying along many dimensions: whether the horizon is finite or infinite, whether the equilibria are perfect or not, whether there is discounting or not, whether there is complete or incomplete

*We are grateful to the C. V. Starr Center at NYU for research support and the Institute of Economics at the University of Copenhagen for its generous hospitality during a visit while part of this paper was written.

information, whether players are maximizers or 'satisficers,' and whether the set of players remains the same or consists of overlapping generations.[1]

Perhaps the most debated of these aspects concerns the time horizon of the repeated game: whether it is finite or infinite. As is well-known, the set of equilibrium outcomes of a game repeated a large but finite number of times, may be radically different from the equilibrium outcomes of its infinitely repeated counterpart. This is sometimes referred to as the 'finite horizon paradox' and the numerous attempts to resolve it have been responsible for much of the work alluded to above. In addition, there has been some debate on whether a finite or infinite horizon is the appropriate choice for modelling repeated interaction among economic agents. For instance, Rubinstein (1991) has taken the position that players generally perceive repeated games, including those with a known, fixed, and even short finite horizon as a game of infinite duration, and thus infinite horizon games are the appropriate model. However, the reasons underlying Rubinstein's position are difficult to fathom. The assumption that *otherwise rational* players view the twenty-fold repetition of the prisoners' dilemma as being infinitely repeated is rather curious. And while the predictions of the infinite horizon model are certainly consistent with observed behavior in the finitely repeated prisoner's dilemma, there is little in the model to particularly suggest the sort of end-game play that is typically observed. Indeed, the predictions of some finite horizon models, such as those with incomplete information or satisficing players, are more in tune with observed behavior.

In our opinion, generally modelling finitely repeated games as infinitely repeated ones is unwarranted. Perhaps any attempt to reach an unequivocal conclusion on this issue is futile. In any case, it is unnecessary. The theory of repeated games seeks to identify circumstances in which the set of equilibrium outcomes of repeated games is larger than that of the one-shot game. Postulating an infinite horizon is neither necessary nor sufficient for this. That it is not necessary is well-known; if the constituent game has multiple equilibrium payoffs folk theorems for games with long but finite horizons are available.

To see that it is not sufficient, consider the following example. Suppose the discount rate, which may also be interpreted as the probability of continuation, is time dependent. In particular, suppose that the sequence of per-period discount

[1] Detailed references are given below. See Aumann (1980), Pearce (1992) and Sorin (1993) for surveys of the area.

factors, δ_t, declines and approaches $\frac{1}{2}$ as t increases.[2] Specifically, given a $\gamma \in (0,1)$ let

$$\delta_t = \frac{1}{2} + \frac{\gamma^t}{2}.$$

Notice that for all γ, $\lim_{t\to\infty} \delta_t = \frac{1}{2}$ and that for all t, $\lim_{\gamma\to 1} \delta_t = 1$. The latter property implies that as $\gamma \to 1$, players become arbitrarily patient. Now, consider the following game

$0,\ 0$	$x, -3$
$-3,\ x$	$2,\ \ 2$

which is a prisoners' dilemma when $x > 2$. It may be verified that if $x = 5$, for all γ, the unique equilibrium payoff in the infinitely repeated game is $(0,0)$. Although the horizon is infinite, there is a unique equilibrium outcome even when players are arbitrarily patient.

One might be tempted to argue that since the discount factors decline this is 'just like' a finite horizon situation rather than an infinite one. However, such a contention is refuted by the fact that if $x = 4$, then for large γ any feasible individually rational payoff can be obtained in a perfect equilibrium. Thus, an argument that this is like or unlike a finite horizon could not be based upon the underlying time structure, that is the sequence $\langle \delta_t \rangle$, but rather, would have to depend on the payoffs. The example illustrates that the finite-infinite distinction is of limited use in analyzing or classifying the strategic possibilities. Moreover, even if one favors the modeling fiction of an unbounded time horizon, there is little reason to further assume that the continuation probability, or discount rate, is constant.[3]

In this paper we attempt a synthesis of the various folk theorems by adopting a point of view which de-emphasizes the choice of horizon. Instead, we choose to view any repeated game as consisting of a finitely repeated game followed by an 'end-game' whose exact form and duration we leave unspecified. We show that a folk theorem like result can be established whenever the end-game has enough threat potential to discipline players; this requires that the end-game have multiple equilibria. From this perspective the various resolutions of the finite horizon paradox may simply be viewed as devices which create or enhance this threat potential. Indeed, we show that all of the disparate folk theorems for finite

[2]Thus payoffs from period t are discounted by a factor of $\prod_{\tau=1}^{t} \delta_\tau$.

[3]In Section 6 we discuss the paper of Bernheim and Dasgupta (1995), where the continuation probability is not constant.

horizon games are rather simple corollaries of a single central result. Furthermore, infinite horizon games may also be treated in the same way: an infinitely repeated game may be thought of as a finitely repeated game followed by an end-game of infinite duration.

Our approach offers some advantages. First, it demonstrates the essential unity of the various folk theorems and their proofs. The proof of each theorem may be decomposed into two parts: The construction of the threat in the end-game and an application of our central result.

Second, we are able to derive some new results and stronger versions of existing results. For instance, our methodology yields a strong version of Radner's (1980) ϵ-equilibrium folk theorem that is 'uniform' in the needed variance from optimizing behavior. Similarly, we obtain a folk theorem with incomplete information that is uniform in the type of incomplete information needed.

Finally, rather than focusing attention on the horizon of the game our approach isolates what is essential for the folk theorems to hold: whether there is an end-game in which players' may be threatened sufficiently severely.

This paper is organized as follows. Section 2 contains basics and some preliminary results. We consider perfect equilibria of repeated games with common discounting and define the central concept: a two-part game that consists of a finitely repeated game followed by an 'end-game.' We adopt the effective minmax methodology of Wen (1994) as it leads to the most general results. In Section 3 we derive the main result (Theorem 1) which identifies circumstances under which a folk theorem like result may be derived for the two-part game.

Section 4 derives the various folk theorems for finitely repeated games by invoking the main result of Section 3. Theorem 2 is the folk theorem for finitely repeated games in which each player has distinct equilibrium payoffs (Benoît and Krishna (1985)). Theorem 3 is the recent generalization to the case where the players have recursively distinct equilibrium payoffs (Smith (1995)). We then consider epsilon equilibria of finitely repeated games and derive two folk theorems (Theorems 4 and 5) that originate in the work of Radner (1980) and Chou and Geanakoplos (forthcoming). Finally, we consider games with incomplete information and derive a result (Theorem 6) along the lines of Fudenberg and Maskin (1986).

Section 5 concerns infinitely repeated games. We consider the Fudenberg and Maskin (1986) result and its generalizations by Abreu, Dutta and Smith (1994) and Wen (1994) (Theorem 7). We then consider the recent model of an infinitely repeated game with a declining discount factor studied by Bernheim and Dasgupta

(1995) (Theorem 8).

Section 6 derives a folk theorem for a model with overlapping generations similar to results of Kandori (1992) and Smith (1992) (Theorem 9).

Section 7 introduces the notion of a 'frequent response game' in which the horizon is fixed but players revise moves rapidly. Section 8 indicates some possible extensions and Section 9 concludes.

While all of the above results are derived for the case of pure strategies or alternatively for the case when mixed strategies are observable, in an appendix we show that our main result, and thus all of the results that stem from it, remain true even if mixed strategies are not observable.

2 Preliminaries

Let $G = (A_1, A_2, ..., A_n;\ U_1, U_2, ..., U_n)$ be a game where A_i is i's strategy space and U_i is his payoff function. We assume that the A_i's are compact and the U_i's are continuous. As usual we write $A \equiv \prod_{j=1}^{n} A_j$ and $A_{-i} \equiv \prod_{j \neq i} A_j$ with generic elements a and a_{-i} respectively.

We also assume that G has at least one equilibrium.

If the A_i's are convex subsets of a Euclidean space we call G a *continuous game.* If the A_i's are finite sets we call G a *finite game*.

Let v_i be player i's *minmax* payoff defined as:

$$v_i = \min_{a} \max_{a_i} U_i\left(a_i, a_{-i}\right).$$

Let F denote the feasible and *individually rational* payoffs in G, that is,

$$F = \left\{u \in co\, U\left(A\right) :\ u \geq v\right\}.$$

Two players i and j are said to have *equivalent utilities* if i's payoff function is an increasing affine transformation of j's (see Abreu, Dutta and Smith (1994)). Let $N(i)$ denote the set of players j such that i and j have equivalent utilities. $N(i)$ defines an equivalence class. Following Wen (1994) let v_i^* be player i's *effective minmax* payoff defined as:

$$v_i^* = \min_{a} \max_{k \in N(i)} \max_{a_k} U_i(a_k, a_{-k}). \tag{1}$$

As in Wen (1994), for all i let m^i be the solution to (1) so that we have that for all $j \in N(i)$:

$$\max_{a_j} U_j(a_j, m^i_{-j}) \leq \max_{k \in N(i)} \max_{a_k} U_j(a_k, m^i_{-k}) = v^*_j.$$

If m^i is played, each player $j \in N(i)$ obtains v^*_j and is at a best response. Clearly, $v^*_i \geq v_i$.

Normalize the game so that for all i, $v^*_i = 0$ and define:

$$F^* = \{u \in co\, U(A) : u \geq 0\}.$$

to be the set of feasible and *effectively rational* payoffs in G. We assume that there exists a $u \in F^*$, $u \gg 0$.

Note that if no two players have equivalent utilities, then $F^* = F$. Wen (1994) also shows that $F^* = F$ for all two player games.

Let $\overline{u}$ denote the maximum payoff of any player in G and let $\underline{u}$ denote the minimum payoff.

$G^\delta(T)$ will denote the game which consists of T repetitions of G and in which players use a discount factor of $\delta \in (0, 1)$ to evaluate payoffs.

If $(a^1, a^2, ..., a^T)$ is a path in $G^\delta(T)$ the resulting discounted average payoff vector is

$$\frac{(1-\delta)}{\delta(1-\delta^T)} \sum_{t=1}^{T} \delta^t U(a^t).$$

$G^\delta(\infty)$ denotes the infinitely repeated game with discount factor δ.

In both $G^\delta(T)$ and $G^\delta(\infty)$ we assume that players can coordinate their actions by means of a public randomizing device. Initially, we assume that mixed strategies are observable (alternatively, we can say that the A_i's themselves subsume all mixing possibilities). In an appendix we show that all of our results are true even if the mixed strategies themselves are not observable; rather only the pure strategy choices of the players may be observed.

Let $P^\delta(T)$ be the set of (subgame) perfect equilibrium discounted average payoffs of $G^\delta(T)$. Similarly, $P^\delta(\infty)$ denotes the set of perfect equilibrium discounted average payoffs of $G^\delta(\infty)$.

Wen (1994) easily establishes that for all δ and T, $P^\delta(T) \subseteq F^*$ and that for all δ, $P^\delta(\infty) \subseteq F^*$.

3 The Main Result

Let $H^\delta = (S_1, S_2, ..., S_n;\ V_1^\delta, V_2^\delta, ..., V_n^\delta)$ be a game with parameter $\delta \in (0, 1)$.

Definition 1 *Given a finitely repeated game $G^\delta(T)$ and another game H^δ, the* ***conjunction*** *of the two games, written $\langle G^\delta(T), H^\delta \rangle$, is a $T+1$ period game that consists of $G^\delta(T)$ followed by H^δ. H^δ is then referred to as the* ***end-game.***

If $(a^1, a^2, ..., a^T, s)$ is a path in $\langle G^\delta(T), H^\delta \rangle$, the resulting discounted average payoff is

$$\frac{(1-\delta)}{\delta\,(1-\delta^{T+1})}\left[\sum_{t=1}^{T} \delta^t U(a^t) + \delta^{T+1} V^\delta(s)\right].$$

If $(\sigma^1, \sigma^2, ..., \sigma^T, \sigma^{T+1})$ is a perfect equilibrium of the game $\langle G^\delta(T), H^\delta \rangle$, then for all $(a^1, a^2, ..., a^T) \in A^T$, we must have that $\sigma^{T+1}(a^1, a^2, ..., a^T)$ is a perfect equilibrium of H^δ.[4] Notice that if H^δ is itself a repeated game $G^\delta(T')$ (where T' may be finite or infinite) then a perfect equilibrium payoff of $\langle G^\delta(T), H^\delta \rangle$ is a perfect equilibrium payoff of $G^\delta(T+T')$.

Definition 2 *A game $H = (S_1, ..., S_n, V_1, ..., V_n)$ has a* ***threat*** *of M if there exist $n+1$ perfect equilibria of H : $s, s^1, s^2, ..., s^n$ satisfying for all i,*

$$[V_i(s) - V_i(s^i)] \geq M. \tag{2}$$

$s, s^1, s^2, ..., s^n$ will be referred to as the ***threat strategies*** *which yield M.*

Let $\Pi^\delta(T)$ be the set of (subgame) perfect equilibrium payoffs of the conjoined game $\langle G^\delta(T), H^\delta \rangle$. Our main result concerns the Hausdorff (double) limit of $\Pi^\delta(T)$ as $\delta \to 1$ and $T \to \infty$. In the statement below, F^* is derived from the game G.

Theorem 1 *There exists an M such that if for all large δ, H^δ has a threat of M then*

$$\lim_{\substack{\delta \to 1 \\ T \to \infty}} \Pi^\delta(T) = F^*.$$

[4] Of course, if H^δ has no proper subgames, the set of perfect equilibria of H^δ is the same as the set of Nash equilibria of H^δ.

Proof. Let $u \in F^*$, $u \gg 0$. From Abreu, Dutta and Smith (1994) there exist n vectors $u^1, u^2, ..., u^n$ in F^*, $u^i \gg 0$, such that for all $i \in N$,

$$u_i^i < u_i \tag{3}$$

for all $j \notin N(i)$,

$$u_i^i < u_i^j \tag{4}$$

and for all $j \in N(i)$,

$$u_i^i = u_i^j. \tag{5}$$

Let $a \in A$ be such that $U(a) = u$. Similarly, for all i let $a^i \in A$ be such that $U(a^i) = u^i$.

Suppose Q, R and T are given, $R < Q < T$. Let s, $s^1, s^2, ..., s^n$ be perfect equilibria of H^δ . (Note that s and s^i depend on δ).

Let π_1^0 be a path from period 1 to $T + 1$ described by:

$$\pi_1^0 = (a, a, ..., a, s).$$

For $i = 1, 2, ..., n$, and $\tau \leq T - Q$ let π_τ^i be a path that begins in period τ and ends in period $T + 1$ described by:

$$\pi_\tau^i = (\overbrace{m^i, m^i, ..., m^i}^{R \text{ periods}}, a^i, a^i, ..., a^i, s).$$

Given a strategy combination σ, we say that *i deviates from a path π_τ^i in period t* if σ calls for i to play $\pi_\tau^i(t)$ but i plays something else. Consider the following (recursively defined) strategies:

- Follow π_1^0 (until someone deviates). If player i is the (lowest indexed) player to deviate from π_τ^j in period $t \leq T - Q$, follow π_{t+1}^i.
- If player i is the first player (with the lowest index) to deviate from π_τ^j in some period t, $T - Q < t \leq T$, then play to some equilibrium e of G in each subsequent period through period T, and play s^i in period $T + 1$. Any deviation after the first is ignored.

We now show that for large enough Q, R, T and δ these are perfect equilibrium strategies. It is sufficient to verify that no player wants to deviate from these strategies just once and conform thereafter.

First, consider $t \leq T - Q$ and deviations by player $i \in N(j)$ from π_1^0 or π_τ^j. Clearly i cannot gain by deviating from π_τ^j while being (effectively) minmaxed. For other periods, if i deviates his remaining payoff stream is bounded above by:

$$(\overline{u}, \overbrace{0, 0, ..., 0}^{R \text{ periods}}, u_i^j, ..., u_i^j, V_i^\delta(s)) \tag{6}$$

On the other hand, if i does not deviate his remaining payoff stream will be at worst:

$$(u_i^j, \overbrace{u_i^j, u_i^j, ..., u_i^j}^{R \text{ periods}}, u_i^j, ...u_i^j, V_i^\delta(s)) \tag{7}$$

(recalling that for $i \in N(j)$, $u_i^j < u_i$).

For large enough R and δ the discounted value of (7) is greater than that of (6). Fix R so that this inequality holds for all i.

Next, consider $t \leq T - Q$ and deviations by player $i \notin N(j)$ from π_τ^j. Deviating once yields a remaining payoff stream bounded above by:

$$(\overline{u}, \overbrace{0, ..., 0}^{R \text{ periods}}, \overbrace{u_i^i, ..., u_i^i}^{Q-R \text{ periods}}, \overbrace{u_i^i, ..., u_i^i}^{T-t-Q}, V_i^\delta(s)) \tag{8}$$

Not deviating yields i at worst:

$$(\overbrace{\underline{u}, ..., \underline{u}}^{R \text{ periods}}, \overbrace{u_i^j, ..., u_i^j}^{Q-R \text{ periods}}, \overbrace{u_i^j, u_i^j, ..., u_i^j}^{T-t-Q+1}, V_i^\delta(s)) \tag{9}$$

Choose Q so that the discounted value of (9) is greater than that of (8) for all i, j, $i \notin N(j)$, when $t = T - Q$ and $\delta = 1$. The inequality then also holds for all large δ and $t < T - Q$.

Finally, consider deviations by a first deviator i in a period $t > T - Q$. If i deviates his remaining payoff stream is bounded above by:

$$(\overbrace{\overline{u}, \overline{u}, ..., \overline{u}}^{Q \text{ periods}}, V_i^\delta(s^i)) \tag{10}$$

If i does not deviate his remaining payoff stream will be at worst:

$$(\overbrace{\underline{u}, \underline{u}, ..., \underline{u}}^{Q \text{ periods}}, V_i^\delta(s)) \tag{11}$$

Let $M(u)$ be a number satisfying $M(u) > Q(\overline{u} - \underline{u})$ and suppose $s, s^1, s^2, ..., s^n$ are perfect equilibria of H^δ such that for all i

$$\left[V_i^\delta(s) - V_i^\delta(s^i)\right] \geq M(u).$$

For large enough δ, the discounted value of (11) is greater than that of (10), for all i.

Thus, we have shown that for all $u \in F^*$, $u \gg 0$, there exists an $M(u)$ such that if for all large δ, H^δ has a threat of $M(u)$ then for all large T and δ, $(a, a, ..., a, s)$ is an equilibrium path of $\langle G^\delta(T), H^\delta \rangle$.

Now fix $n+1$ points $\widehat{u}, \widehat{u}^1, \widehat{u}^2, ..., \widehat{u}^n$ in F^*, $\widehat{u} \gg 0$, $\widehat{u}^i \gg 0$, with corresponding outcomes $\widehat{a}, \widehat{a}^1, \widehat{a}^2, ..., \widehat{a}^n$, such that for all $i \in N$,

$$\widehat{u}_i^i < \widehat{u}_i \tag{12}$$

and for all $j \notin N(i)$,

$$\widehat{u}_i^i < \widehat{u}_i^j. \tag{13}$$

Let $M = \max\left\{M(\widehat{u}), M\left(\widehat{u}^1\right), M\left(\widehat{u}^2\right), ..., M(\widehat{u}^n)\right\}$.

Suppose that for all large δ, H^δ has a threat of M. We use this threat M to construct an arbitrarily large threat M'. In fact, for all M' there exists a $T(M')$ such that for all $T > T(M')$ and large δ, $\langle G^\delta(T), H^\delta \rangle$ has a threat of M'. This follows from the fact that for large δ, $(\widehat{a}, \widehat{a}, ..., \widehat{a}, s)$ and $(\widehat{a}^i, \widehat{a}^i, ..., \widehat{a}^i, s)$ are perfect equilibrium paths of $\langle G^\delta(T), H^\delta \rangle$, and the difference in player i's payoffs from these paths exceeds M' when T is large.

Consider an arbitrary $u' \in F^*$, $u' \gg 0$, and the game $G^\delta(T)$ followed by an end-game. As above, u' can be obtained as a perfect equilibrium payoff in every period but the last, if the end-game has a large enough threat M' $(= M(u'))$. While H^δ has a threat of (only) M, we have shown that for all large δ, the conjoined game $\langle G^\delta(T(M')), H^\delta \rangle$ has a threat M'. Therefore, u' can be obtained in every period but the last of the game $G^\delta(T)$ followed by the end-game $\overline{H}^\delta \equiv \langle G^\delta(T(M')), H^\delta \rangle$. But this game, $\langle G^\delta(T), \overline{H}^\delta \rangle$, is the same as $\langle G^\delta(T + T(M')), H^\delta \rangle$.

Since for large T and δ the payoff u' can obtained in all but the last $T(M') + 1$ periods of $\langle G^\delta(T), \overline{H}^\delta \rangle = \langle G^\delta(T + T(M')), H^\delta \rangle$, the resulting discounted average payoff is approximately u'.

When T is large, the threat M required for Theorem 1 is small relative to the total payoffs in $G^\delta(T)$. The following corollary gives a sufficient condition for M to be small relative to the payoffs in G.

Given an outcome a define the maximum gain from deviating from a :

$$d(a) = \max_i \{U_i(b_i(a), a_{-i}) - U_i(a)\}.$$

where $b_i(a)$ is i's best response to a_{-i}.

Corollary 1 *Suppose G has an equilibrium e that is inefficient. If for all large δ, H^δ has a threat of $M > \inf\{d(a) : U(a) \gg U(e)\}$ then*

$$\lim_{\substack{\delta \to 1 \\ T \to \infty}} \Pi^\delta(T) = F^*.$$

Proof. Choose a such that $U(a) \gg U(e)$ and $d(a) < M$. Let $s, s^1, s^2, ..., s^n$ be the threat strategies which yield M. Then for all T and large δ the path $(a, a, ..., a, s)$ is a perfect equilibrium path of $\langle G^\delta(T), H^\delta \rangle$ since deviations can be punished with $(e, e, ..., e, s^i)$. Since the payoff difference between $(a, a, ..., a, s)$ and $(e, e, ..., e, s)$ can be made arbitrarily large, the game $\overline{H}^\delta = \langle G^\delta(T'), H^\delta \rangle$ has an arbitrarily large threat for large δ and T'. Apply Theorem 1 to $\langle G^\delta(T), \overline{H}^\delta \rangle$.

Note that for continuous games $\inf\{d(a) : U(a) \gg U(e)\} = 0$, so that if G is a continuous game with an inefficient equilibrium it is enough for H^δ to have any positive threat.[5]

4 Finitely Repeated Games

As noted earlier, if H^δ is itself a repeated game $G^\delta(T')$ then a perfect equilibrium payoff of $\langle G^\delta(T), H^\delta \rangle$ is a perfect equilibrium payoff of $G^\delta(T + T')$, that is, $\Pi^\delta(T + 1) = P^\delta(T + T')$. Furthermore, if there exists a set of perfect equilibrium strategies $s, s^1, s^2, ..., s^n$ of $G^\delta(T')$ which, for all $\delta \geq \delta'$, yield a positive threat for $G^\delta(T')$, then by choosing k large enough, the game $G^\delta(kT')$ can be made to have an arbitrarily large threat for large δ. This threat can be obtained simply by 'patching' together the relevant threat strategies of $G^\delta(T')$ k times.[6] Then, applying Theorem

[5]Chou and Geanakoplos (forthcoming) show that for "smooth commitment" games a folk theorem may obtain with any positive commitment. For more on this see the *extensions* section at the end of this paper.

[6]Suppose $(a^1, a^2, ..., a^T)$ is a perfect equilibrium path of $G^\delta(T)$ and $(\overline{a}^1, \overline{a}^2, ..., \overline{a}^{\overline{T}})$ is a perfect equilibrium path of $G^\delta(\overline{T})$. Then $(a^1, a^2, ..., a^T, \overline{a}^1, \overline{a}^2, ..., \overline{a}^{\overline{T}})$ is a perfect equilibrium path of $G^\delta(T + \overline{T})$. This process is referred to as 'patching' the two paths together.

1 to $\langle G^{\delta}(T), G^{\delta}(kT')\rangle$ we have that there exists a k such that

$$\lim_{\substack{\delta\to 1\\ T\to\infty}} \Pi^{\delta}(T) = \lim_{\substack{\delta\to 1\\ T\to\infty}} P^{\delta}(T+kT') = F^{*}.$$

Thus, we have the following:

Proposition 1 *Suppose that for some T' and δ' there exists a set of strategies which, for all $\delta \geq \delta'$, yield a positive threat in $G^{\delta}(T')$. Then*

$$\lim_{\substack{\delta\to 1\\ T\to\infty}} P^{\delta}(T) = F^{*}.$$

4.1 Games with Distinct Equilibrium Payoffs

Say that G has *distinct* equilibrium payoffs if every player has two equilibrium payoffs. The following result is due to Benoît and Krishna (1985).[7]

Theorem 2 *If G has distinct equilibrium payoffs then*

$$\lim_{\substack{\delta\to 1\\ T\to\infty}} P^{\delta}(T) = F^{*}.$$

Proof. For all i, let $U_i(\overline{e}^i) > U_i(\underline{e}^i)$ be the best and worst equilibrium payoff, respectively, for player i. Set $T' = 2n$, and let $s = \left(\overline{e}^1, \underline{e}^1, \overline{e}^2, \underline{e}^2, ..., \overline{e}^n, \underline{e}^n\right)$ and $s^i = (\underline{e}^i, \underline{e}^i, ..., \underline{e}^i)$ be T' period paths. Note that for all δ, $s, s^1, s^2, ..., s^n$ forms a positive threat in $G^{\delta}(T')$. Now apply Proposition 1.

As in Smith (1995) we say that G has *recursively distinct* equilibrium payoffs if there exists an ordering of the players $1, ..., n$ such that (a) player 1 has two equilibrium payoffs, and (b) for all i, $1 \leq i < n$, there exist strategy combinations h^{i+1} and l^{i+1} such that $u_{i+1}\left(h^{i+1}\right) > u_{i+1}\left(l^{i+1}\right)$ and each player $i+1, ..., n$ is at a best response.[8] The following generalization of Theorem 2 is due to Smith (1995).

Theorem 3 *If G has recursively distinct equilibrium payoffs then*

$$\lim_{\substack{\delta\to 1\\ T\to\infty}} P^{\delta}(T) = F^{*}.$$

[7] We remind the reader that this and all other the folk theorems are given in Wen's (1994) 'effective minmax formulation.'

[8] This definition is equivalent to the definition in Smith (1995).

Proof. We proceed inductively. Suppose that for some i, $1 \leq i < n$, there exists a Q_i such that $G^1(Q_i)$ has two equilibrium payoffs for players $1, ..., i$ in which, in every period, each player either strictly prefers to follow the equilibrium path than to deviate from it or is at a single period best response. For all i let $\overline{\pi}^i$ and $\underline{\pi}^i$ be such perfect equilibrium paths of $G^1(Q_i)$ yielding i his highest and lowest payoff among such paths, respectively. Since G has recursively distinct equilibrium payoffs there exists a k such that

$$(h^{i+1}, \overbrace{\overline{\pi}^1, ..., \overline{\pi}^1}^{k \text{ times}}, \overbrace{\underline{\pi}^1, ..., \underline{\pi}^1}^{k \text{ times}}, \ \dots \ , \overbrace{\overline{\pi}^i, ..., \overline{\pi}^i}^{k \text{ times}}, \overbrace{\underline{\pi}^i, ..., \underline{\pi}^i}^{k \text{ times}})$$

and

$$(l^{i+1}, \overbrace{\overline{\pi}^1, ..., \overline{\pi}^1}^{k \text{ times}}, \overbrace{\underline{\pi}^1, ..., \underline{\pi}^1}^{k \text{ times}}, \ \dots \ , \overbrace{\overline{\pi}^i, ..., \overline{\pi}^i}^{k \text{ times}}, \overbrace{\underline{\pi}^i, ..., \underline{\pi}^i}^{k \text{ times}})$$

are perfect equilibrium paths of $G^1(2ikQ_i + 1)$ with different payoffs for player $i + 1$, and again in every period each player either strictly prefers to follow the equilibrium path than to deviate or is at a single period best response . These paths are supported by threatening player $j = 1, ..., i$ with a punishment of going to $\underline{\pi}^j$ ($2ik$ times) for a deviation in period 1; note that players $i + 1, ..., n$ are all at best responses in period 1. Thus, continuing inductively, we can find a Q such that every player has two equilibrium payoffs for all δ sufficiently close to 1. By patching these equilibria together as in the proof of Theorem 2, the game $G^\delta(T')$ has a positive threat, where now $T' = 2nQ$. Now apply Proposition 1.

4.2 ϵ-Equilibria

Radner (1980) introduced the following notion of approximate equilibrium behavior.

Definition 3 *An* ϵ***-equilibrium*** *of* $G^\delta(T)$ *is a strategy combination* σ *such that for all* i *and* σ_i'

$$\frac{1-\delta}{\delta(1-\delta^T)}\left[\sum_{t=1}^{T}\delta^t U(a^t(\sigma))\right] \geq \frac{1-\delta}{\delta(1-\delta^T)}\left[\sum_{t=1}^{T}\delta^t U(a^t(\sigma_i', \sigma_{-i}))\right] - \epsilon$$

where $a^t(\sigma)$ *and* $a^t(\sigma_i', \sigma_{-i})$ *are the outcomes at time* t *resulting from* σ *and* (σ_i', σ_{-i}) *respectively. A* ***perfect*** ϵ***-equilibrium*** *is an* ϵ*-equilibrium in every subgame of* $G^\delta(T)$.

Below we establish a folk theorem for perfect ϵ-equilibria along the lines of Radner (1980). However, the notion of perfect ϵ-equilibrium may be criticized on the grounds that while a player's payoff is close to his best-response payoff in terms of *averages*, in long horizon game the discrepancy may be quite large in terms of *totals.* The following definition is intended to address this issue directly.

Definition 4 *An* ***Ω-satisficing equilibrium*** *of $G^{\delta}(T)$ is a strategy combination σ such that for all i and σ_i'*

$$\sum_{t=1}^{T} \delta^t U(a^t(\sigma)) \geq \left[\sum_{t=1}^{T} \delta^t U(a^t(\sigma_i', \sigma_{-i}))\right] - \Omega.$$

A ***perfect Ω-satisficing equilibrium*** *is an Ω-satisficing equilibrium in every subgame of $G^{\delta}(T)$.*

The notion of perfect Ω-satisficing equilibrium may be criticized on the grounds that while in a long horizon game a player's *total* loss from not optimizing is small relative to his overall payoff, this loss may be large relative to the remaining payoff towards the end of the game.

We now show how Theorem 1, and hence Proposition 1, may be applied to obtain folk theorems for both notions.

Let $P_{\epsilon}^{\delta}(T)$ be the set of perfect ϵ-equilibrium average payoffs of $G^{\delta}(T)$. Let $P_{\Omega}^{\delta}(T)$ be the set of perfect Ω-satisficing equilibrium average payoffs of $G^{\delta}(T)$. To apply Theorem 1 and Proposition 1 in the present context recall that a strategy combination σ in $\langle G^{\delta}(T), H^{\delta}\rangle$ such that (1) no player can profitably deviate in any subgame which starts in one of the first T periods, and (2) for any history σ induces a subgame perfect equilibrium of H^{δ}, is a subgame perfect equilibrium of $\langle G^{\delta}(T), H^{\delta}\rangle$. A strategy combination σ in $\langle G^{\delta}(T), H^{\delta}\rangle$ such that (1) no player can profitably deviate in any subgame which starts in one of the first T periods, and (2′) for any history σ induces a perfect ϵ-equilibrium of H^{δ}, is a perfect ϵ-equilibrium of $\langle G^{\delta}(T), H^{\delta}\rangle$. Similarly, for perfect Ω-satisficing equilibria. This reasoning can be carried through to Proposition 1 so that:

Remark 1 *If the threat strategies in Definition 2 are perfect ϵ-equilibria, then in Proposition 1, $P^{\delta}(T)$ can be replaced by $P_{\epsilon}^{\delta}(T)$. Similarly, if the threat strategies are perfect Ω-satisficing equilibria, $P^{\delta}(T)$ can be replaced by $P_{\Omega}^{\delta}(T)$.*

In light of the two criticisms mentioned above, Theorem 4 below requires a strategy combination to be both a perfect ϵ-equilibrium and a perfect Ω-satisficing equilibrium.

Theorem 4 *There exists an Ω such that for all $\epsilon > 0$,*

$$\lim_{\substack{\delta \to 1 \\ T \to \infty}} \left[P_{\Omega}^{\delta}(T) \cap P_{\epsilon}^{\delta}(T) \right] \supseteq F^{*}.$$

Proof. Let e be an equilibrium of G, and let $\widehat{u}, \widehat{u}^1, \widehat{u}^2, ..., \widehat{u}^n$ be elements of F satisfying for all i:

$$\widehat{u}_i > \widehat{u}_i^i.$$

Clearly, for all outcomes a there exists a T' and δ' such that for all $\delta > \delta'$, $(a, e, e, ..., e)$ is a perfect ϵ-equilibrium path of $G^{\delta}(T')$. Using the outcomes corresponding to the $n + 1$ above points yields a positive threat in $G^{\delta}(T')$ for all $\delta > \delta'$. Now apply Remark 1.

Notice that if Ω is very small, then the notion of a perfect Ω-satisficing equilibrium is satisfactory on its own. The following theorem is due to Chou and Geanakoplos (forthcoming).

Theorem 5 *Suppose G is a continuous game with an inefficient equilibrium. Then for any $\Omega > 0$,*

$$\lim_{\substack{\delta \to 1 \\ T \to \infty}} P_{\Omega}^{\delta}(T) \supseteq F^{*}.$$

Proof. For any Ω choose a such that $U(a) \gg U(e)$ and the maximal deviation $d(a) < \Omega$. This is an Ω-satisficing equilibrium of $G^{\delta}(1)$. Set $H^{\delta} = G^{\delta}(1)$ and apply Corollary 1 and Remark 1.

4.3 Games with Incomplete Information

Following Kreps *et al* (1982), Fudenberg and Maskin (1986) have shown that even if a game has a unique equilibrium, a folk theorem may obtain if a small amount of incomplete information is introduced. Specifically, they showed that for every $u \in F^{*}$, there is a game of incomplete information in which u is a sequential

equilibrium payoff. Note that in this statement the type of incomplete information used may vary with the outcome u being sustained. Indeed, Fudenberg and Maskin (1986) state that "One may argue for or against certain equilibria on the basis of the type of irrationality needed to support them."

In this section we present a stronger version of their result in which all the outcomes can be sustained with the *same* type of irrationality. Thus, the (optimistic) claim that different outcomes can be distinguished on the basis of the 'needed' irrationality is mistaken. In fact, Theorem 1 shows that this must be the case: once two Pareto comparable equilibria, for instance, have been identified, these can be used in a suitable end-game to immediately establish a folk theorem.

To apply Theorem 1 to games of incomplete information recall that a strategy combination σ in $\langle G^{\delta}(T), H^{\delta}\rangle$ such that (1) no player can profitably deviate in any subgame which starts in one of the first T periods, and (2) for any history, σ induces a subgame perfect equilibrium of H^{δ}, is a subgame perfect equilibrium of $\langle G^{\delta}(T), H^{\delta}\rangle$. Similarly, a strategy combination σ in $\langle G^{\delta}(T), H^{\delta}\rangle$ such that (1) no player can profitably deviate in any subgame which starts in one of the first T periods, and (2″) for any history σ induces a sequential equilibrium of H^{δ}, is a sequential equilibrium of $\langle G^{\delta}(T), H^{\delta}\rangle$. Thus:

Remark 2 *If the threat strategies in Definition 2 are sequential equilibria, then in Theorem 1* $\Pi^{\delta}(T)$ *can be replaced by the set of sequential equilibrium payoffs of* $\langle G^{\delta}(T), H^{\delta}\rangle$.

Theorem 6 *For all T there exists a game of incomplete information $G^{\delta}(T,\epsilon)$ in which with probability $(1-\epsilon)$ each player's payoffs are the same as in $G^{\delta}(T)$ and*

$$\lim_{\epsilon\to 0}\lim_{\substack{\delta\to 1\\ T\to\infty}} Seq^{\delta}(T,\epsilon) = F^{*},$$

where $Seq^{\delta}(T,\epsilon)$ is the set of sequential equilibrium payoffs of $G^{\delta}(T,\epsilon)$.

Proof. Let $G^{\delta}(Q,\epsilon)$ be the Q period game of incomplete information in which each player is 'rational' with probability $(1-\epsilon)$ and with probability ϵ is an 'irrational' player who is completely indifferent among all outcomes.

Fix some equilibrium e of G satisfying $U(e) \gg 0$.[9] Let $u \in F$ and let $a \in A$ be such that $U(a) = u$. Consider the following behavior strategies.

[9]The assumption that G has an equilibrium such that $U(e) \gg 0$ is made only for convenience.

Irrational player i.

I1. Play a_i in period 1. If each player j plays a_j in period 1, then play e_i for the remaining $Q-1$ periods of the game at all information sets.

I2. Suppose that some player j does not play a_j in period 1, and indeed let j be the lowest indexed such player. Then:

- In period $t = 2$: (effectively) minmax j
- In period $t > 2$:
  - if for all τ such that $2 \leq \tau < t$, all players have minmaxed j then minmax j in period t;
  - if there is a τ, $2 \leq \tau < t$, such that some player did not minmax j, then play e_i in period t.

Rational player i.

R1. Play a_i in period 1. If each player j plays a_j in period 1, then play e_i for the remaining $Q-1$ periods of the game at all information sets.

R2. Suppose that some player j does not play a_j in period 1, and let j be the lowest indexed such player.

- In period $t > 2$, if there is a τ, $2 \leq \tau < t$, such that some player did not minmax j, then play e_i in period t.
- Otherwise for $t \geq 2$, all rational players play to a particular sequential equilibrium derived under the restriction that all players' strategies conform to the above restrictions.

We now verify that the above constitutes a sequential equilibrium.

We need only check that rational players are optimizing.

Following period 1, all players are either at a single period best response, or by construction cannot gain by deviating. By construction also, all beliefs are sequentially consistent.

We now verify that for Q large enough playing a_i is strictly optimal in period 1 when $\delta = 1$. Playing a_i yields

$$u_i(a) + (Q-1)\, u_i(e)\,. \tag{14}$$

Deviating (once) yields a payoff which is bounded above by

$$\overline{u} + \epsilon^{n-1}(Q-1)\,0 + \left(1-\epsilon^{n-1}\right)\left[\overline{u} + (Q-2)\,u_i(e)\right]. \tag{15}$$

(15) < (14) for large enough Q.

Thus, we have shown that for all $u \in F$, there exists a $\overline{Q}$ such that for all $Q > \overline{Q}$, there is a sequential equilibrium of $G^{\delta}(Q,\epsilon)$ which results in a stream of payoffs $(u, U(e), ..., U(e))$.

Now let $\widehat{u}, \widehat{u}^1, \widehat{u}^2, ..., \widehat{u}^n$ be elements of F satisfying for all i:

$$\widehat{u}_i > \widehat{u}_i^i.$$

We have found $n+1$ sequential equilibrium payoffs of $G^{\delta}(Q,\epsilon)$ which show that $G^{\delta}(Q,\epsilon)$ has a positive 'sequentially perfect threat.'

An arbitrarily large threat can be obtained by patching together sequential equilibria of $H^{\delta} \equiv \langle G^{\delta}(Q,\epsilon), ..., k \text{ times}, ..., G^{\delta}(Q,\epsilon)\rangle$, for some k. The resulting strategy combination is a sequential equilibrium. Remark 2 shows that Theorem 1 may be applied. Finally, note that a sequential equilibrium of $\langle G^{\delta}(T), H^{\delta}\rangle$ is a sequential equilibrium of a $T+kQ$ period game of incomplete information (in this game of incomplete information, the irrationality manifests itself independently every kQ periods).

5 Infinitely Repeated Games

5.1 Stationary Discounting

For infinitely repeated games with discounting, the following generalization of the theorem of Fudenberg and Maskin (1986) is due to Wen (1994).

Theorem 7

$$\lim_{\delta \to 1} P^{\delta}(\infty) = F^*$$

Theorem 7 follows from Theorem 1 if $G^{\delta}(\infty)$ can be shown to have an arbitrarily large threat as $\delta \to 1$. To see this set $H^{\delta} = G^{\delta}(\infty)$, and note that a perfect equilibrium payoff of $\langle G^{\delta}(T), H^{\delta}\rangle$ is a perfect equilibrium payoff of $G^{\delta}(\infty)$.

If G has (recursively) distinct equilibrium payoffs for each player then an arbitrarily large threat can be constructed by patching together single-shot equilibria in much the same manner as in the finite case. If G has an inefficient equilibrium e this threat can be constructed as follows: Let a be such that $u(a) \gg u(e)$. Let $\overline{s}$ be the strategies: each player i plays a_i so long as all others do; if anyone deviates play e_i forever. Let $\underline{s}$ be e repeated forever. For large δ these strategies form an arbitrarily large threat.

If G has a single equilibrium which is efficient Theorem 1 may still be applied, however it appears that it is then no easier to establish that $G^{\delta}(\infty)$ has an arbitrarily large threat than it is to establish Theorem 7 directly (the latter can be done by setting $T = \infty$ and $Q = 0$ in the proof of Theorem 1).

Theorems 2 and 7 together confirm a conjecture of Pearce (1992) that if G has (recursively) distinct equilibrium payoffs then

$$\lim_{T\to\infty} P^{1}(T) = \lim_{\delta\to 1} P^{\delta}(\infty).$$

Note, however, that an example in Benoît and Krishna (1987) shows that even with distinct equilibrium payoffs, in general

$$\lim_{T\to\infty} P^{\delta}(T) \neq P^{\delta}(\infty).$$

5.2 Non-Stationary Discounting

The standard model of an infinitely repeated game with a common discount factor of δ can be reinterpreted to represent a situation where the horizon of the game is uncertain and δ represents the probability that the game will be played in period $t+1$ conditional on it being played in period t. With probability $(1-\delta)$ the game ends in period t. Notice that in this formulation players are assumed to maximize their expected payoff and the probability of continuation is independent of the period.

In a recent paper, Bernheim and Dasgupta (1995) have examined repeated games where the probability of continuation is time dependent; in particular, it declines over time.[10] Thus let δ_t represent the probability that the game will be played in period t given that it was played in period $t-1$. Let $\langle\delta_t\rangle$ be the sequence of such continuation probabilities and given such a sequence, let $G^{\langle\delta_t\rangle}(\infty)$ represent

[10] Alternatively, the discount factor declines with time.

a repeated game in which the total (expected) payoff vector from a path $(a^1, a^2, ...)$ is,

$$\sum_{t=1}^{\infty} \left(\prod_{\tau=1}^{t} \delta_\tau \right) U(a^t).$$

$G^{\langle \delta_t \rangle}(\infty)$ is said to be a repeated game with an *asymptotically finite horizon* if for all t, $\delta_t > 0$ and $\lim_{t\to\infty} \delta_t = 0$.

Let γ_t be a monotonically declining sequence satisfying $\lim_{t\to\infty} \gamma_t = 0$ and for fixed δ and T consider the game $\langle G^\delta(T), G^{\langle \delta\gamma_t \rangle}(\infty)\rangle$ which consists of a T period repeated game followed by an infinitely repeated game with an asymptotically finite horizon in which the sequence of continuation probabilities is $\langle \delta\gamma_t \rangle$.

First, suppose that the constituent game G is finite (that is, all the A_i's are finite) and has a *unique* equilibrium, say e. Then for all δ, the game $G^{\langle \delta\gamma_t \rangle}(\infty)$ also has a unique perfect equilibrium. This is because there exists a $c > 0$ such that from any $a \neq e$ each player can gain at least c by deviating. Since $\delta_t \to 0$, there exists a T_c such that for all $t \geq T_c$, no deviations can be punished in the subsequent game and thus no $a \neq e$ can be played after period T_c. The fact that this is the only perfect equilibrium path in the overall game now follows from backwards induction.

Second, if the constituent game G has (recursively) distinct equilibrium payoffs, then the arguments of the section on finitely repeated games may be applied to derive a folk theorem for $\langle G^\delta(T), G^{\langle \delta\gamma_t \rangle}(\infty)\rangle$ as $\delta \to 1$ and $T \to \infty$.

Finally, suppose that the game G has a continuum of strategies and that G has a unique equilibrium e. When can a folk theorem like result be derived for the game $\langle G^\delta(T), G^{\langle \delta\gamma_t \rangle}(\infty)\rangle$ as defined above? The answer depends on how fast the continuation probabilities are declining. Say that the sequence of continuation probabilities $\langle \delta_t \rangle$ declines *slowly* if

$$\lim_{T\to\infty} \sum_{t=1}^{T} 2^{-t} \ln \delta_t > -\infty.$$

Suppose that e is inefficient. Bernheim and Dasgupta (1995) then show that $G^{\langle \delta_t \rangle}(\infty)$ has a non-trivial equilibrium (that is, other than playing e repeatedly) *only if* $\langle \delta_t \rangle$ declines slowly.[11]

They also prove the following folk theorem:

[11] Some additional conditions are also needed for this result: the payoff functions and the best-response functions must be twice continuously differentiable and the equilibrium must be regular. See Bernheim and Dasgupta (1995) for details.

Theorem 8 *Suppose that G is a continuous game with an inefficient equilibrium. If $\langle \gamma_t \rangle$ declines slowly then,*

$$\lim_{\substack{\delta \to 1 \\ T \to \infty}} \Pi^{\delta}(T) = F^*,$$

where $\Pi^{\delta}(T)$ is the set of perfect equilibrium average payoffs in the game $\langle G^{\delta}(T), G^{\langle \delta \gamma_t \rangle}(\infty) \rangle$.

Proof. By Theorem 1 in Bernheim and Dasgupta (1995) the game $G^{\langle \delta \gamma_t \rangle}(\infty)$ has a perfect equilibrium whose payoffs Pareto dominate $U(e)$. Thus the game $G^{\langle \delta \gamma_t \rangle}(\infty)$ has a positive threat. Apply Corollary 1.

6 Overlapping Generations

Crémer (1986), Kandori (1992) and Smith (1992) have examined models in which each player is finitely lived but there is an infinite population of players who interact in 'overlapping generations.'

Consider a model with *n types* of players (indexed by i) of different generations (indexed by $r = 1, 2, ...$). Player (i, r) is the player of type i in the rth generation. Let $K > 0$ be fixed and suppose that each player (i, r) lives for $T > nK$ periods. For $r > 1$, player (i, r) is assumed to be born in period $(i-1)K + (r-1)T + 1$ and die in period $(i-1)K + (r-1)T + T$. Thus, all n players of a given generation r overlap for exactly $T - (n-1)K$ periods. We will refer to such a game as $OLG^{\delta}(T, K)$.

Define

$$\begin{aligned} P_r^{\delta}(T, K) \quad = \quad & \{u \in F^* : \text{there is a perfect equilibrium of } OLG^{\delta}(T, K) \\ & \text{in which the payoffs of all generations 1 through } r \text{ is } u\}. \end{aligned}$$

The following theorem is similar to results derived by Kandori (1992) and Smith (1992).

Theorem 9 *There exists a K such that for all $u \in F^*$,*

$$u \in \lim_{\substack{\delta \to 1 \\ T \to \infty}} P_r^{\delta}(T, K)$$

for every r.

Proof. Fix some inefficient equilibrium e of G.[12] For ease of exposition suppose that $n = 3$. We first find two equilibria of $OLG^{\delta}(T, K)$. The first is simply $\underline{s} = (e, e, \ldots)$, that is a path in which all players play e. $\overline{s}$ is defined as follows.

Let a^i be the outcome that is best for a player of type i, and let a be such that $U(a) \gg U(e)$. For $k_1 < K$ and $k_2 < K$ consider the T period path beginning with the birth of a type 1 player:

$$(\overbrace{a^2, ..., a^2}^{k_2 \text{ periods}}, \overbrace{e, ..., e}^{K-k_2}, \overbrace{a^3, ..., a^3}^{K}, \overbrace{a, a, ..., a}^{T-3K}, \overbrace{a^1, ..., a^1}^{k_1}, \overbrace{e, ..., e}^{K-k_1}). \tag{16}$$

This is repeated in successive generations with the birth of each new player of type 1.

Any deviation is met by a play of e forever.

For every k_1, we can choose k_2 to be large enough so that a player of type 2 cannot gain by deviating in the k_1 periods that a^1 is played. For every k_1 and k_2, K can be chosen large enough so that there is no gain for a player of type 3 from deviating either in the k_1 periods that a^1 is played or in the k_2 periods that a^2 is played.

Let M be defined as in the statement of Theorem 1 and choose k_1, k_2 and K large enough so that

$$\begin{aligned} k_1 \left[U(a^1) - U_1(e)\right] &> M \\ k_1 \left[U_2(a^1) - U_2(e)\right] + k_2 \left[U_2(a^2) - U_2(e)\right] &> M \\ k_1 \left[U_3(a^1) - U_3(e)\right] + k_2 \left[U_3(a^2) - U_3(e)\right] + K \left[U_3(a^3) - U_3(e)\right] &> M \end{aligned}$$

and the deviations of the previous paragraph are not profitable.

Finally, choose T large enough so that no player can profitably deviate in the remaining periods.

Fix a generation r and let $G_r^{\delta}(T - 3K)$ be a finitely repeated game among the players of generation r. For sufficiently large T and δ, let H^{δ} be the $OLG^{\delta}(T, K)$ that begins when a type 1 player of generation r is $T - K$ periods old. H^{δ} has a threat of M. Applying Theorem 1 to $\langle G_r^{\delta}(T - 3K), H^{\delta} \rangle$ shows that any payoff $u > 0$ and a payoff $u' \ll u$ can be sustained in this game when T and δ are large. Let c be the outcome which results in u, and c' be the outcome which results in c'.

We can now construct threats which yield the players in generation $r - 1$ a payoff of (approximately) u, and in which the play for generation r is like (16), but

[12] We assume that G has an inefficient equilibrium only for convenience.

with c replacing a along the path, and with c' replacing e as the punishment for a deviation from c. Continuing in this manner for earlier generations establishes the theorem.

7 Frequent Response Games

One recommendation for a finite horizon model over an infinite horizon model is that agents typically do not live forever. However, neither do they have extremely long albeit finite lives. Nevertheless, even with a short horizon, say, fixed at one year, the players may move frequently, say daily. We now establish a folk theorem for such situations.

Consider a game G. Suppose the payoff function U refers to annual flow payoffs; that is, $U(a)$ is the payoff when all players play a throughout the year. Similarly, let δ refer to the annual discount rate.

If players move once at the beginning of the year and are not allowed to revise their moves, the payoff received at the end of the year is $U(a)$ and its value discounted to the beginning of the year is $\delta U(a)$.

Now suppose that players can move more frequently in this *same* one year game. For instance, if they moved twice a year, and there were no discounting then their semi-annual payoff from playing a in any period would be $\frac{1}{2}U(a)$. With discounting their semi-annual payoff would be the discounted average $\frac{\sqrt{\delta}}{1+\sqrt{\delta}}U(a)$.

More generally, if players move K times during the course of the year and the path $(a^1, a^2, ..., a^K)$ results, the discounted total payoff at the beginning of the year is

$$\sum_{k=1}^{K} \delta^{\frac{k}{K}} \frac{(1-\delta^{\frac{1}{K}})}{\delta^{\frac{1}{K}}(1-\delta)} \delta U(a^k).$$

By writing $\widehat{\delta}_K \equiv \delta^{\frac{1}{K}}$ and $\widehat{U}_{K,\delta}(a^k) \equiv \frac{(1-\delta^{\frac{1}{K}})}{\delta^{\frac{1}{K}}(1-\delta)}\delta U(a^k)$ this can be rewritten in a familiar form as,

$$\sum_{k=1}^{K} \widehat{\delta}_K^k \widehat{U}_{K,\delta}(a^k).$$

Notice that with this specification, the one-period ($\frac{1}{K}$th of a year) payoff function is $\widehat{U}_{K,\delta} \equiv \frac{(1-\delta^{\frac{1}{K}})}{\delta^{\frac{1}{K}}(1-\delta)}\delta U$, the discounted average of an annual payoff function U.

Call the game described above a game with *K-responses* denoted by $G_K^\delta(1)$, and let Γ_K^δ be the game which consists of $G_K^\delta(1)$ followed by an end-game H^δ. The total payoff from the path $(a^1, a^2, ..., a^K, s)$ is,

$$\sum_{k=1}^{K} \delta^{\frac{k}{K}} \frac{(1-\delta^{\frac{1}{K}})}{\delta^{\frac{1}{K}}(1-\delta)} \delta U(a^k) + \delta^{\frac{K+1}{K}} V^\delta(s)$$

or, more compactly,

$$\sum_{k=1}^{K} \widehat{\delta}_K^k \widehat{U}_{K,\delta}(a^k) + \widehat{\delta}_K^{K+1} V^\delta(s).$$

Let $\Pi_K^\delta(1)$ denote the set of perfect equilibrium payoffs from Γ_K^δ. We obtain the following result:

Theorem 10 *If there exists an $M > 0$ such that for all large δ, H^δ has a threat of M then*

$$\lim_{\substack{\delta \to 1 \\ K \to \infty}} \Pi_K^\delta(1) = F^*.$$

Proof. The proof is virtually the same as the proof of Theorem 1, once it is recognized that the rescaling of the payoffs from U to $\widehat{U}_{K,\delta}$ implies that $M(u)$, the threat needed to sustain u, goes to 0 as K increases.

Note that in Theorem 10, H^δ can have an arbitrarily small threat.

Section 4 derived theorems for finitely repeated games as the horizon T increased. Using Theorem 10, one can derive an analogue of each of these theorems for a fixed horizon as the frequency of response K increases. An overlapping generations folk theorem for short-lived frequently responding agents can similarly be derived.

8 Extensions

In the game $\langle G^\delta(T), H^\delta \rangle$, $G^\delta(T)$ was followed by an end-game H^δ and in applying Theorem 1, H^δ was itself taken to be a repeated game. In this section we present some alternative formulations.

8.1 Games of Division

Let H^δ be a *game of division* in which players 'split a pie' of size L. That is, the players simultaneously announce an $x \in R^n : \sum_{i-1}^n x_i = L$. If the players make the same announcement they receive this payoff, otherwise they receive 0. Call such an end-game a *game of division with a pie of size L.*

As an example consider two players engaged in an enterprise which has the features of a prisoners' dilemma D:

$0, 0$	$5, -3$
$-3, 5$	$2, \quad 2$

Suppose the players' discount factor comes from the market interest rate and let each player contribute $3\delta^{T+1}$ at the beginning of the game $D^\delta(T)$. For large T, these contributions are infinitesimal; they grow to a pie of size 6 at the end of the game. Corollary 1 implies that for large δ and T, any feasible individually rational payoff is sustainable in the game $D^\delta(T)$ ending with this pie.

For a game with a fixed horizon, but frequent responses Theorem 10 yields the following strong result.

Proposition 2 *Let $G_K^\delta(1)$ be a game with K responses. If the end-game H^δ is a game of division with a pie of positive size then*

$$\lim_{\substack{\delta \to 1 \\ K \to \infty}} \Pi_K^\delta(1) = F^*.$$

In particular, if players move frequently enough in a prisoner's dilemma, cooperation can be sustained with any positive pie.

For continuous games we also have a strong result, similar to Theorem 5.

Proposition 3 *Let G be a continuous game with an inefficient equilibrium. If the end-game H^δ is a game of division with a pie of positive size then*

$$\lim_{\substack{\delta \to 1 \\ T \to \infty}} \Pi^\delta(T) = F^*.$$

Recall the standard linear Cournot oligopoly model. This has a unique equilibrium if repeated any finite number of times. Nonetheless, Proposition 3 implies that if the firms have a penny to split at the end, then for large δ any feasible individually rational payoff is sustainable with enough repetitions, a result similar to that of Conlon (1994).

8.2 Exogenous Payoffs

Suppose that after the end of the repeated game $G^{\delta}(T)$ players receive additional payoffs according to the exogenously specified function $h : A^T \to R^n$. Specifically, the total payoff from a path $(a^1, a^2, ..., a^T)$ is:

$$\left[\sum_{t=1}^{T} \delta^t U(a^t)\right] + \delta^{T+1} h(a^1, a^2, ..., a^T).$$

Some possible interpretations of the function h are given below.[13]

Games with a Bonus Let h be a *bonus* payoff contingent on actions taken during the course of a finitely repeated game. For any $u \in F^*$, there is a contract specifying a strategy combination σ and a payment of size M to the players contingent on σ being followed, such that u is sustainable. Although this does not follow directly from Theorem 1 as stated, it is immediately clear from its proof, since h functions exactly as H^{δ}.

Employment and other contracts which specify bonus payments (such as pensions and buyouts) contingent on certain behavior may be functioning in this role. Note that Corollary 1 and Theorem 10 imply that for many games this payment M may be quite small.

Commitments Chou and Geanakoplos (forthcoming) propose a model in which players can exogenously commit to behave in a particular manner in the last few periods of a finitely repeated game. The function h can be viewed as resulting from such a commitment.[14]

Psychological Costs Now suppose h represents a (small) psychological cost. Consider the above prisoners' dilemma D. Suppose that after T periods of cooperation each player feels a psychological cost of 3 from defecting. Notice that this cost is small relative to the total payoff in the repeated game. Nonetheless, it is sufficient to permit cooperation to be sustained.

[13] Kandori's (1992) model of a finitely repeated game with 'bonding' is quite similar. However, his model cannot be used to derive the various folk theorems of previous sections. The reasons for this are the same as those discussed in Smith (1992).

[14] Starting with this formulation they derive Theorem 4 and rederive Theorem 2. Notice that in a similar vein we could have used Theorem 2 as the starting point from which to derive all the theorems of Section 4. Indeed, Proposition 1 is a simple consequence of Theorem 2.

8.3 Nash Equilibria

A simple modification of our framework can also accommodate *Nash* equilibrium folk theorems. Instead of requiring that H^δ have a sufficiently large (perfect equilibrium) threat, the weaker condition that H^δ have an equilibrium whose payoffs are sufficiently greater than minmax levels (of H^δ) is enough to establish a Nash equilibrium analog of Theorem 1 (along the lines of Benoît and Krishna (1987)).

9 Conclusion

Our earlier work established that if the stage game has distinct equilibrium payoffs, a folk theorem can be derived (Benoît and Krishna (1985)). This paper extends that idea: The distinct payoffs in the stage game enables the construction of sufficiently severe threats in an 'end-game,' and our main result (Theorem 1) essentially takes these end-game threats as a starting point.

The utility of Theorem 1 should be apparent: All the major folk theorems can be recast so that they are simple consequences of this result.

A Appendix

A.1 Unobservable Mixed Strategies

In this appendix we show that our main result, Theorem 1, continues to hold if mixed strategies are unobservable. The basic idea of the proof is similar to that given by Fudenberg and Maskin (1986) for infinitely repeated games: punishing players are kept indifferent on the support of the minmax (mixed) strategy. For this result we suppose that the game is finite.

Let $\Pi_u^\delta(T)$ denote the set of perfect equilibrium average payoffs when mixed strategies are unobservable.

Theorem A1 *There exists an M such that if for all large δ, H^δ has a threat of M then*

$$\lim_{\substack{\delta \to 1 \\ T \to \infty}} \Pi_u^\delta(T) = F^*.$$

Proof. Let $u, u^1, u^2, ..., u^n$ in F^*, $u \gg 0, u^i \gg 0$ *be* such that for all $i \in N$,

$$u_i^i < u_i$$

and for all $j \notin N(i)$,

$$u_i^i < u_i^j.$$

For every player i define

$$L(i) = \left\{ j \in N : U_j = \alpha^{ji} U_i + \beta^{ji},\ \alpha^{ji} \neq 0 \right\}.$$

(Notice that here we allow $\alpha^{ji} < 0$ and thus $L(i)$ is different from $N(i)$). Consider a partition of N into K equivalence classes:

$$N = L(i_1) \cup L(i_2) \cup ... \cup L(i_K).$$

>From each member of the partition choose exactly one player. Suppose the players so chosen are $i_1, i_2, ..., i_K$ and without loss of generality, rename these $1, 2, ..., K$.

Fix a player i and without loss of generality, in all that follows suppose that $i = 1$. For each $k = 2, 3, ..., K$ and $j \in L(k)$, choose $x^{ij} \in F^*$ and x^{ij} close enough to u_j^i so that

$$\begin{array}{ll} x_j^{ij} \neq u_j^i \text{ and } x_j^{ij} > u_j^j & \text{if } j \in L(k) \\ x_h^{ij} = u_h^i & \text{if } 1 < h \leq K,\ h \neq k \end{array} \tag{17}$$

since $j \notin L(i)\ (= L(1))$.

Let $a \in A$ be such that $U(a) = u$. Similarly, for all i let $a^i \in A$ be such that $U(a^i) = u^i$ and let $a^{ik} \in A$ be such that $U(a^{ik}) = x^{ik}$.

Given a strategy combination σ, say that i is *observed to deviate* from a path π_τ^j in period t if σ calls for the players to play the mixed strategy $\pi_\tau^j(t)$, but i's (pure) action is not in the support of the ith component of $\pi_\tau^j(t)$.

Consider the following paths:

π_τ^0 :

- In periods $\tau + 1, \tau + 2, ..., T$: play a; and
- In period $T + 1$: play s.

$\pi^i_{\tau+1}$:

- In periods $\tau+1, \tau+2, ..., \tau+R$: play m^i;
- In periods $\tau+R+1, \tau+R+2, ..., \tau+R+R'$: if the observed outcome path in the first R periods is $\alpha' \in (supp\, m^i)^R$, for each $k = 2, 3, ..., K$:
  - $*$ play a^{ik} with probability $p^{ik}(\alpha')$, and
  - $*$ play a^i with probability $1 - \sum_{k=2}^{K} p^{ik}(\alpha')$;
- In periods, $\tau+R+R'+1, \tau+R+R'+2, ..., T$: play a^i; and
- In period $T+1$: play s.

Now consider the following strategies.

- Start π^0_1 and continue to follow π^0_1 if no one deviates.
- If i is observed to deviate from π^j_τ in period $t \leq T - Q$, start π^i_{t+1}.
- If player i is the first player observed to deviate from π^j_τ in some period t, $T - Q < t \leq T$, then play e in each of the periods $t+1, t+2, ..., T$ and play s^i in period $T+1$.

We now show that for large enough δ, Q, R, R' and T, these are perfect equilibrium strategies. It is sufficient to verify that no player wants to deviate from these strategies just once and conform thereafter.

Choose R so that for all i :

$$(R+1)\, u_i > \overline{u}$$

Such an R exists since $u_i > 0$. Choose δ_R so that for all $\delta > \delta_R$, for all i :

$$\frac{\left(1-\delta^{R+1}\right)}{(1-\delta)} u_i > \overline{u}$$

and for all i :

$$\left(1-\delta^{R+1}\right) \underline{u} + \delta^{R+1} u^i_i > 0.$$

As before fix a player i and suppose that $i = 1$.

Claim: *There exists an R' and a $\delta_{R'}$ such that for all $\delta > \delta_{R'}$, for all $k = 2, 3, ..., K$ and for all $\alpha', \alpha'' \in (supp\, m^i)^R$ there exist $p^{ik}(\alpha')$ and $p^{ik}(\alpha'')$ such that:*

$$\begin{aligned} & U_k^R(\alpha') + \delta^{R+1}\frac{(1-\delta^{R'})}{(1-\delta)}\left(p^{ik}(\alpha')x_k^{ik} + \left(1 - p^{ik}(\alpha')\right)u_k^i\right) \\ = \; & U_k^R(\alpha'') + \delta^{R+1}\frac{(1-\delta^{R'})}{(1-\delta)}\left(p^{ik}(\alpha'')x_k^{ik} + \left(1 - p^{ik}(\alpha'')\right)u_k^i\right) \end{aligned} \tag{18}$$

$$\text{and } \sum_{k=2}^{K} p^{ik}(\alpha') < 1.$$

where $U_k^R(\alpha') \equiv \sum_{r=1}^{R} \delta^r U_k(a'(r))$ is player k's total payoff from the R period path α'. $U_k^R(\alpha'')$ is defined similarly.

Proof of claim: Let $\overline{\alpha}^k \in (supp\, m^i)^R$ be the R period path that is best for player k and let α be an arbitrary path in $(supp\, m^i)^R$. For these paths rewrite the equality in (18) as:

$$\frac{(1-\delta)\left[U_k^R(\overline{\alpha}^k) - U_k^R(\alpha)\right]}{\delta^{R+1}\left(1-\delta^{R'}\right)} = \left[p^{ik}(\alpha) - p^{ik}(\overline{\alpha}^k)\right]\left(x_k^{ik} - u_k^i\right) \tag{19}$$

By choosing δ and R large enough, the left hand side of (19) can be made arbitrarily small. Therefore, $p^{ik}(\alpha)$'s can be chosen such that for all i, for all $k = 2, 3, ..., K$ and for all α, (19) holds and

$$\sum_{k=2}^{K} p^{ik}(\alpha') < 1. \quad \square$$

Once again fix a player $i = 1$ (say). For every $k \leq K$, $k \neq 1$, choose the $p^{ik}(\alpha)$ as in the claim.

Suppose $i = 1$ deviates. If $\alpha' \in (supp\, m^i)^R$ is the path in the first R periods of π_t^i when m^i is played, the payoff to player $k = 2, 3, ..., K$ is the left hand side of the equation in (18), whereas if $\alpha'' \in (supp\, m^i)^R$ is the path, the payoff is the right hand side of (18). By construction these are equal. Now consider a player $j \in L(k)$, $k \neq i$. Since j's payoff function is just an affine transformation of k's, j is indifferent between these paths. Thus every player $j \notin L(i)$ is indifferent among all the paths that lie in $(supp\, m^i)^R$.

Players in $j \in L(i)$ are all at single-period best responses when m^i is played.

The verification that the strategies given above form an equilibrium is now routine. (The fact that the x_j^{ij}'s satisfy (17) is important here). The remainder of the proof may be completed by following the proof of Theorem 1 exactly.

In the proof given above we have made use of the fact that players can coordinate their actions by means of public randomization (correlation). Gossner (1995) demonstrates a folk-theorem for finitely repeated games without discounting, $G^1(T)$, when mixed strategies are not observable and public randomization is not allowed. The non-discounting assumption appears to be crucial to his argument as he makes use of Blackwell's approachability theorem. Fudenberg and Maskin (1991) show that the use of public randomization is not needed for the folk theorem for infinitely repeated games with discounting. Note also that public randomization plays

References

[1] Abreu, D., P. K. Dutta and L. Smith (1994): "The Folk Theorem for Repeated Games: A NEU Condition," *Econometrica*, 62, pp. 939-948.

[2] Benoît, J-P. and V. Krishna (1985): "Finitely Repeated Games," *Econometrica*, 53, pp. 905-922.

[3] Benoît, J-P. and V. Krishna (1987): "Nash Equilibria of Finitely Repeated Games," *International Journal of Game Theory*, 16, pp. 197-204.

[4] Bernheim, B. D. and A. Dasgupta (1995): "Repeated Games with Asymptotically Finite Horizons," *Journal of Economic Theory*, 67, pp. 129-152.

[5] Chou, C-F. and J. Geanakoplos (forthcoming) : "The Power of Commitment," *Journal of Economic Theory*, Yale University.

[6] Conlon, J. (1994): "Cooperation for Pennies," mimeo, University of Mississippi.

[7] Crémer, J. (1986): "Cooperation in Ongoing Organizations," *Quarterly Journal of Economics*, 101, pp. 33-49.

[8] Fudenberg, D. and E. Maskin (1986): "The Folk Theorem for Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, pp. 533-554.

[9] Fudenberg, D. and E. Maskin (1991): "On the Dispensability of Public Randomization in Discounted Repeated Games," *Journal of Economic Theory*, 53, pp. 428-438.

[10] Gossner, O. (1995): "The Folk Theorem for Finitely Repeated Games with Mixed Strategies," *International Journal of Game Theory*, 24, pp. 95-107.

[11] Kandori, M. (1992): "Repeated Games Played by Overlapping Generations of Players," *Review of Economic Studies*, 59, pp. 81-92.

[12] Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982): "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27, pp. 245-252.

[13] Pearce, D. (1992): "Repeated Games: Cooperation and Rationality" Chapter 4 in *Advances in Economic Theory: Sixth World Congress*, Vol. 1 edited by J-J. Laffont, Econometric Society Monographs, Cambridge University Press, Cambridge, pp. 132-174.

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

[15] Rubinstein, A. (1991): "Comments on the Interpretation of Game Theory," *Econometrica*, 59, pp. 909-924.

[16] Smith, L. (1992): "Folk Theorems in Overlapping Generations Games," *Games and Economic Behavior*, 4, pp. 426-449.

[17] Smith, L. (1995): "Necessary and Sufficient Conditions for the Perfect Finite Horizon Folk Theorem," *Econometrica*, 63, pp. 425-430.

[18] Sorin, S. (1993): "Repeated Games with Complete Information," Chapter 4 in *Handbook of Game Theory*, Volume 1, edited by R. Aumann and S. Hart, North-Holland, pp. 71-107.

[19] Wen, Q. (1994): "The 'Folk Theorem' for Repeated Games with Complete Information," *Econometrica*, 62, pp. 949-954.