

Folk Theorems in Overlapping Generations Games*

LONES SMITH

*Department of Economics, Building E52, Massachusetts Institute of Technology,
50 Memorial Drive, Cambridge, Massachusetts 02139*

Received January 31, 1990

This paper characterizes perfect folk theorems for repeated games with overlapping generations of finite-lived players. We prove two *uniform* folk theorems that admit arbitrarily long-lived players for a given discount factor; the result for $n > 2$ players requires a full-dimensional payoff space. Under no assumptions whatsoever, a *nonuniform* n -player folk theorem obtains in which the discount factor must covary with the players' lifespans. Our focus on the *overlap* rather than the *generation* makes possible compact and explicit descriptions of all equilibria. We later synthesize our results in a more general setting with some finite- and some infinite-lived players. *Journal of Economic Literature* Classification Numbers: C72, C73. © 1992 Academic Press, Inc.

1. INTRODUCTION

A "folk theorem" asserts that any individually rational payoff of a stage game can be attained on average in an equilibrium of the corresponding repeated game. The wealth of known folk results underscores one dramatic difference between static and dynamic behavior. Although the earliest circulating folk theorems¹ may have been conceived with the *Nash* equilibrium concept in mind, more recent efforts, including this one, insist upon *subgame perfection* (cf. Selten, 1975). Further remarks shall therefore be confined to the latter arena of perfect folk results.

* This paper is the second chapter of my Ph.D. dissertation in economics at the University of Chicago. My advisor, In-Koo Cho, was very instrumental early on in this project. I later benefited from the advice of Ariel Rubinstein and the remarks of several referees. All remaining errors are my own. Financial assistance from the Social Sciences and Humanities Research Council of Canada, the Searle Foundation, and the Jacob K. Javits Fellowship Program is gratefully acknowledged.

¹ Aumann (1985) sketches some of the earlier folk theorem literature.

By using a suitably defined payoff objective—the “overtaking criterion”—Rubinstein (1979) obtained his folk theorem for supergames without discounting. When payoffs are discounted, Fudenberg and Maskin (1986) (henceforth F–M) additionally require that the feasible payoff space have *full dimension*. Extension to the real world of mortal players has encountered yet another caveat. The folk theorem for finitely repeated games in Benoit and Krishna (1985) also requires at least two distinct Nash payoffs for each player in the stage game. Indeed, it is well known that in the finitely repeated Prisoners’ Dilemma, cooperation does not arise in equilibrium. The forward thrust of positive economic theory has sadly floundered upon this simple game.

Combining elements of finite and infinite horizon games, this paper is a game theoretic foray into relatively new territory: overlapping generations (OLG) games. By this, we mean games played by n teams of finite-lived individuals, each of whom “lives” (i.e., plays) for n consecutive “overlaps”; at the end of each T period overlap, one player “dies,” to be replaced by the next younger fellow team member. We establish three separate folk theorems for OLG games. Each obtains for suitable selections of the overlap length and discount factor, and the interplay between these two quantities shares centerstage in this paper. All results hold even when the stage game has only one Nash equilibrium.

Recent work by Crémer (1986), Cooper and Daughety (1988), and Salant (1991) throws light on the flavor of results in this new field. But it is Kandori (1989) who—independently of this work—obtains the first general folk theorem for OLG games. Although our approach and extent of results differ somewhat from that of Kandori (1989), we do touch on a few common themes. We first recall a simple example of cooperation in the repeated Prisoners’ Dilemma by unboundedly rational finite-lived players. The insights acquired in this exercise then motivate our two-player folk theorem. As often occurs in game theory, for want of a mutual minimax point, the n -player results do not come so easily. But faced with two degrees of freedom (the discount factor and the overlap length) we must also decide *which* folk theorem we want. This predicament typifies the many surprises offered by this medley of finite and infinite repeated game theory.

Our first n -player folk theorem obtains without even insisting upon a full-dimensional feasible payoff space—the very condition which drives the n -player results of both F–M and Benoit and Krishna (1985). Because the discount factor must co-vary with the players’ lifespans, we have suggestively labeled this a *nonuniform* folk theorem. A simple corollary slightly improves upon the folk theorem discovered by Kandori (1989). Since a nonuniform folk theorem *always* holds for general OLG games, one might suppose that the special revolving-door roster of players in our

model rather lends itself to the support of out-of-equilibrium threats. This intuition is not misguided. But the flexibility is to our advantage, as we derive a stronger *uniform* n -player folk theorem under the added assumption of full-dimensionality. Unlike our nonuniform folk theorem, this second result admits arbitrarily long-lived players for a *given* discount factor. Why is this? As it turns out, a full-dimensional payoff space affords us the luxury of immediate punishment for deviations; in its absence, some punishments must await the “death” of one or more players.

In light of these somewhat bipolar results, it is natural to question the exact role played by the dimensionality of the payoff space. For instance, when $n = 3$, what is the incremental value of having a two- versus a one-dimensional payoff space? A nonuniform OLG folk theorem holds in either case, but can anything more be said? To this end, it is helpful to note that both of the above folk theorems must essentially tackle the same problem: How does one linearly disentangle the players’ payoff streams so that deviants are punished and punishers rewarded? Full-dimensionality allows us to construct player-specific punishments at the stage game level. This, in a nutshell, is what F–M and Benoit and Krishna (1985) have done. In an OLG game, however, we have more options. To deal with stage games in which all players receive identical payoffs, the equilibria for our nonuniform folk theorem systematically exploit the cascading age hierarchy, sometimes deferring a given punishment or reward until one or more of the current players has died. Now, it so happens that we may combine the two approaches in a hybrid OLG–supergame model in which exactly $m < n$ of the players are *infinite*-lived. The paper concludes on this note, with the following profitable extension of F–M and our earlier result: *A nonuniform folk theorem² obtains when the projection of the payoff space onto the coordinates corresponding to the infinite-lived players has dimension at least m .* This implies (for nondegenerate stage games) that our earlier nonuniform folk theorem holds when exactly one of the teams in the OLG game is actually a single infinite-lived player!

Much insight into the structure of our equilibria is gained by an appreciation of their *resilient* nature. That is, after any finite sequence of deviations, the play always returns to the principal equilibrium path once all current players have “died.” On the one hand, resiliency may enhance the intuitive appeal of our equilibria, and will perhaps blunt some of the criticisms now motivating the renegotiation-proofness literature.³ But there is a much more far-reaching implication of resiliency—namely, that players need not know the entire history of the OLG game. In fact, it suffices that each player condition his actions on a fixed finite truncation of the action

² By nonuniform, we now mean that the discount factor must covary with the *mortal* players’ lifespans.

³ See Abreu *et al.* (1989) for references.

history, which Kandori (1989) has christened *informational decentralization*. Yet if behavior in one overlap does not depend at all on the previous history, then each overlap is just an isolated finitely repeated game, and the analysis (and proviso) of Benoit and Krishna (1985) applies. Thus, strategies for our uniform folk theorem do make use of the last portion of the previous overlap. In our nonuniform folk theorem, however, where the payoff space can be one-dimensional, players must condition their actions (albeit in a remarkably simple fashion) on the previous $n - 1$ overlaps.⁴ In other words, each player need only know the action history associated with the oldest living player; anything earlier is irrelevant.

As already alluded to, our approach throughout is constructive and not existential. Explicit compact descriptions of the (occasionally Byzantine) strategies are provided. Interested readers can “see” for themselves the punishment routines at work. In particular, the equilibrium algorithm discovered for the nonuniform folk theorem is new, and is an important contribution of this analysis. In all equilibria, the discounting of future payoffs and the finite lifespans of the players have forced us to “make the punishment fit the time,” so to speak: During any overlap, “early” deviants from an equilibrium face a much different lot than do “late” ones. For instance, transgressions by younger players near the end of an overlap—in the “buffer zone”—provoke a temporary Nash respite; only when the oldest player dies do we proceed to punish the culprit. On the other hand, for deviations prior to the buffer zone, our nonuniform folk theorem employs a nested hierarchy of minimax punishment phases based upon a player’s “age.” In effect, “elderly” deviants are punished at once, whereas criminal records are maintained for the “youth.” Nevertheless, we do believe that our proofs and strategies exhibit a simplicity that comes from treating the *overlap* as the basic atom of analysis. The focus in Kandori (1989) on the *generation* (i.e., n consecutive overlaps, during which n players “die”) is nonessentially more general, and unnecessarily complicates matters and conceals deeper results.

Section 2 outlines the model and provides two motivational examples. In Section 3, we present the general results, and draw the distinction between uniform and nonuniform folk theorems. We then synthesize the intuition and analysis behind these two results in Section 4, and prove one last folk theorem. Most constructions of the equilibria, which represent a major contribution of this paper, are included in the text. An appendix contains the verification of each equilibrium.

⁴ In the corresponding result in Kandori (1989), knowledge of the previous $2n - 1$ overlaps is needed. We conjecture (but have been unable to prove) that our $n - 1$ is best possible.

2. THE MODEL

Let $G = (A_1, A_2, \dots, A_n; U_1, U_2, \dots, U_n)$ be an n -person (one-shot) normal form game. Blurring the distinction between pure and mixed strategies, we simply assume throughout that each A_i consists of player i 's mixed strategies; therefore, A_i is a compact and convex strategy space. Every player i maximizes his (continuous) payoff function $U_i: A \rightarrow R$, where $A \equiv \prod_{i=1}^n A_i$. We sometimes interchangeably refer to an element $a \in A$, or its associated payoff vector, as an *outcome* of G .

As we make liberal use of *correlated strategies*,⁵ the *feasible* payoff space V is the convex hull in R^n of the set of attainable payoff vectors. Note that we only consider nondegenerate games G , i.e., where $\dim V \geq 1$. (All our results are vacuously true when $\dim V = 0$.) Finally, we assume that deviations from mixed strategies are observable.⁶

For convenience, set $N = \{1, 2, \dots, n\}$. Let M^i be a minimax strategy against player i in the game G . Not wishing to break ranks with the entire repeated game literature, we do not permit players $j \neq i$ to (privately among themselves) correlate their possibly mixed minimax strategies against i . By this, we mean that $M^i \equiv (M^i_1, M^i_{-i})$, where

$$M^i_{-i} \equiv (M^i_1, \dots, M^i_{i-1}, M^i_{i+1}, \dots, M^i_n) \in \arg \min_{a_{-i}} \max_{a_i} U_i(a_i, a_{-i}),$$

and M^i_i is i 's best response to M^i_{-i} . Intuitively, among all publicly preannounced and independently executed conspiracies (a fair description of an equilibrium) by all players $j \neq i$, M^i_{-i} minimizes i 's payoff. Hence, no rational player i would willingly accept an average payoff below $U_i(M^i)$. We assume WLOG that $U_i(M^i) = 0$ for all $i \in N$. The feasible and *strictly individually rational* payoff set is therefore the *positive* orthant of V . Finally, when $n = 2$, we often resort to the *mutual minimax* outcome $M \equiv (M^2_1, M^2_2)$.

We begin with a simple observation first made by Hammond (1975), and later resurrected in Crémer (1986).

⁵ This is not WLOG. The use of correlated strategies is for expositional ease only. Without the requisite public randomizing devices, arbitrarily long deterministic cycles would be needed. This, however, would lead to messier equilibrium routines, and might well require that players "live" longer. For instance, ensuring compliance with several nonindividually rational cyclical outcomes would necessitate a slightly more potent punishment mechanism than is employed below.

⁶ F-M discuss this issue on page 536, and later consider unobservable mixed strategies. We avoid any such (admittedly nontrivial) complications.

	C	D
C	2, 2	-1, 3
D	3, -1	0, 0

FIGURE 1

EXAMPLE 1. Consider the game G_1 (see Fig. 1), an example of the classic Prisoners' Dilemma. The unique Nash equilibrium *and* minimax point of G_1 is (0, 0). The object is to explain "cooperation"—choosing C—when this game is played repeatedly by finite-lived rational individuals. Consider an overlapping sequence of players, each of whom lives two periods. When a player is "young," he plays his "father"; when "old," he plays his "son." Clearly, when one is old and playing his last game, he will choose his dominant strategy D. In our subgame perfect equilibrium, we insist that he play C when young. We must ask, "What if he deviates to D when young?" Should this occur, the threat is that his son will play D (and not C) against him; also, all subsequent plays (by future players) will be the one-shot Nash equilibrium outcome (D, D). Not only is this threat credible, it even remains so when future payoffs are discounted at any rate $\delta > \frac{1}{3}$, for then $-1 + 3\delta > 0 + 0$.

Let there be n teams of players $S_k = \{k, k + n, k + 2n, \dots\}$, $k \in \mathbf{N}$. Returning to our earlier arbitrary game G , define OLG (G ; δ , T) as the following repeated game:

1. Every period, G is played by n finite-lived players, one from each team. Initially, the players are $1, 2, \dots, n$.
2. Every T periods the roster changes. Player $k + tn$ from team S_k dies and is replaced by player $k + (t + 1)n$ just after time $(k + tn)T$, for $k \in \mathbf{N}$ and $t = 0, 1, 2, \dots$.
3. Each player is fully apprised of the history of play, and seeks to maximize his lifetime average payoff. Future payoffs are discounted at some fixed and common rate $\delta \leq 1$.

Much as with Kandori (1989), it suffices to consider only equilibria in which fellow team members use identical strategies. We thus simply refer to the players by their team numbers $1, 2, \dots, n$; hence, k always denotes the *current* player from team S_k , $k \in \mathbf{N}$. Note that each player lives for exactly nT periods, and that the identities of all n players remain unchanged for T periods—the *overlap* duration—and then one is replaced. We later relax the insistence on identical overlaps to deduce the result of Kandori

	E	F
E	1, 1	0, -2
F	-2, 0	-1, -1

FIGURE 2

(1989). Also, all equilibrium strategies remain valid when the time domain is doubly infinite.

Example 1 might suggest that we may approximate any feasible and strictly individually rational payoff vector by a subgame perfect average outcome of some OLG (G_1 ; δ , T). In other words, a folk theorem holds. We further investigate this possibility in the next example.

EXAMPLE 2. As vexing a game as is the Prisoners' Dilemma, it still has one redeeming virtue: The Nash equilibrium coincides with the mutual minimax point. Thus, in a repeated setting, we can always threaten to revert to a subgame perfect zero payoff, as in Friedman (1971). We accordingly shift our attention to the game G_2 (see Fig. 2), taken from F-M. Note that (1, 1)—the only Nash equilibrium of G_2 —strictly Pareto-dominates the (mixed) minimax point (0, 0). We claim that, for instance, $(\frac{1}{2}, \frac{1}{2})$ can be arbitrarily approximated by an average subgame perfect outcome of OLG (G_2 ; δ , T) by choosing $\delta < 1$ and $T < \infty$ sufficiently large. Let $a^* \equiv \frac{3}{4}(E, E) \oplus \frac{1}{4}(F, F)$ be a correlated strategy generating $(\frac{1}{2}, \frac{1}{2})$. Then the typical equilibrium overlap assumes the form

$$(E, E), \dots, (E, E); a^*, \dots, a^*; (E, E), (E, E).$$

Here, the initial (E, E) *compensation* phase lasts P periods, while a^* is played $T - P - 2$ times. As is now the wont, we insist that a deviant participate in his own punishment. Ignoring simultaneous deviations, the threats which support the equilibrium are as follows:

1. For deviations prior to the third last play of an overlap,
 - (a) one immediate mutual minimax $M \equiv (F, F)$ follows;
 - (b) after any deviation from (a), start (1), (2), or (3) as appropriate.
2. If the older player deviates in his third last round,
 - (a) his younger opponent plays F until the end of the overlap;
 - (b) if the younger player deviates from (a), start (3).
3. If the younger player deviates in any of the last three rounds of an

overlap, then (E, E) is played until the end of the current overlap, and a^* 's replace the next overlap's (E, E) compensation phase.

Computation should reveal that these punishments strictly deter any deviations exactly when $7\delta^2 > 6 + \delta^{P+2}$. It is therefore sufficient that $\delta = 0.96$ and $P = 18$. Finally, for $\delta < 1$ and $T (>20)$ large enough, the discounted average payoff can be chosen arbitrarily close to $(\frac{1}{2}, \frac{1}{2})$.

It is useful at this point to highlight the basic design of the above equilibrium. A mutual punishment mechanism is operative for most of the overlap. But when the older player has only one or two periods to go, we cannot credibly demand that he punish his opponent. Thus, a three-period *buffer zone* rounds out the overlap, during which we capitalize on the OLG structure. From here on, the nature of the punishments is decidedly different. We threaten the oldest player with his minimax payoff if he deviates, and the consequent loss of his final (E, E) *reward* phase. Should the younger player deviate (for instance, by not punishing the older player), we finish the overlap with the one-shot Nash outcome, but deny the younger player his long compensation period at the outset of the next overlap.

This equilibrium—and all others in this paper—is somewhat *resilient*: After any finite sequence of deviations, play always eventually returns to the principal equilibrium path. The two-player equilibrium described in F–M is resilient, while their n -player equilibrium can be made so. Given the changing identities of players in our model, and in light of the burgeoning renegotiation-proofness literature, this new property of infinite-horizon games is especially useful. Simply put, players and their successors are not forever condemned to an inferior outcome path.

Much of our basic intuition flows from Examples 1 and 2. In the Prisoners' Dilemma, the (one-shot) Nash equilibrium is the worst outcome for either player; therefore, it serves as an *out-of-equilibrium* “stick” to *deter* deviators. The flip side to this scenario is captured in Example 2. Here, the Nash equilibrium is the unique Pareto optimal outcome, and thus is used *in equilibrium* as a reward for *not* deviating. In general, however, the Nash outcome may lie anywhere within the feasible and individually rational payoff space, and may very well function simultaneously as a “carrot” for one player and a “stick” for another.

3. THE FOLK THEOREMS

In this section we consolidate the ideas of Examples 1 and 2. Our first folk theorem works for sufficiently large discount factors $\delta \geq \delta_0$ and overlap lengths $T \geq T_0$. Because δ and T can vary independently over this range, we call this a *uniform* folk result.

3.1 A UNIFORM TWO-PLAYER FOLK THEOREM. *Let G be a two-player normal form game. Let $u = (u_1, u_2)$ be feasible and strictly individually rational for G . Then $\forall \varepsilon > 0 \exists \delta_0 < 1$ and $T_0 < \infty$ so that $\delta \in [\delta_0, 1]$ and $T \geq T_0 \Rightarrow \text{OLG}(G; \delta, T)$ has a subgame perfect discounted average payoff within ε distance of u .*

Proof. Let e^* be a Nash equilibrium of G , and set $v_k = U_k(e^*)$ for $k = 1, 2$. Let a^* be a correlated strategy generating u , so that $u_k = U_k(a^*)$ for $k = 1, 2$. We may assume WLOG that $u_1 \neq v_1$ and $u_2 \neq v_2$. For if not, we can choose some feasible and strictly individually rational payoff vector \tilde{u} within $\varepsilon/2$ of u with $\tilde{u}_k \neq v_k$ ($k = 1, 2$), and then proceed with $\varepsilon/2$ and \tilde{u} in place of ε and u . Finally, let b^k be player k 's best outcome in G , for $k = 1, 2$.

The equilibrium path for the overlap when k is "oldest" assumes one of two forms,

$$e^*, \dots, e^*; a^*, \dots, a^*; b^k, \dots, b^k \quad \text{or} \quad a^*, \dots, a^*; b^k, \dots, b^k,$$

according as $v_k > u_k$ or $v_k < u_k$, respectively. Consider the first overlap above. A crucial and recurring feature is that the initial e^* compensation phase and the final b^k reward phase are each of fixed duration⁷—lasting P and S periods, respectively—whereas a^* can be played an arbitrarily large number $T - P - S$ times. A simple argument establishes the following result:

Fact. Given P and S , there exist $T_1 < \infty$ and $\delta_1 \in (0, 1)$ so that the average equilibrium payoff vector in any overlap lies within ε of u whenever $T \geq T_1$ and $\delta \geq \delta_1$.

We now proceed to find $P, S, T \geq T_1$, and $\delta_0 \in [\delta_1, 1)$ for which the above suggested path is in fact subgame perfect. The strategies involved are rather complex, and will accordingly be presented in the form of a computer program. We make use of a few dummy variables: $\lambda \in \{ \text{"good," "bad"} \}$ registers any deviations "late" in an overlap; the next overlap begins with the oldest player's favored outcome among a^* and e^* exactly when $\lambda = \text{"good."}$ The counter t marks the current period in the overlap; t equals $1, 2, \dots, T$ in succession, and then is reset to 1; we assume (but do not write) that t is automatically incremented after every play of G . Actions repeating "until $t = x$ " occur up to and including time x . We choose the discount factor and the variables P, Q , and S afterward to

⁷ This observation corresponds to the spirit of Theorem 2 of Kandori (1989): Here, the compensation phase acts as a "terminal payment" to all remaining players.

ensure that no deviation is profitable. To avoid confusion, we employ the notation $x \leftarrow y$ to mean “assign x the value y .” Steps follow sequentially, unless otherwise prescribed. Instructions following deviations are in square brackets, and are executed immediately.

The equilibrium program begins with $\lambda =$ “good.” For simplicity, let i be “old” and j “young,” with arbitrary players denoted k . For ease of reference, we always refer to steps 1, 2, and 3 as the *main path*.

⊙ Reset $t \leftarrow 1$. If $\lambda =$ “good” and $v_i > u_i$ or if $\lambda =$ “bad” and $v_i < u_i$, start 1; else start 2.

1. Play e^* until $t = P$.
 2. Play a^* until $t = T - S$. [If k deviates at $t \leq T - Q - S$, start 4; if k deviates at $t > T - Q - S$, start 5 if $k = i$ and start 6 if $k \neq i$.]
 3. Play b^i until $t = T$. [If j deviates, start 6.] Then reset $\lambda \leftarrow$ “good,” and go to (★).
 4. Play M for Q periods. [If k deviates at $t \leq T - Q - S$, restart 4; if k deviates at $t > T - Q - S$, start 5 if $k = i$ and start 6 if $k \neq i$.] Then return to the main path.
 5. Play M^i until $t = T$. [If j deviates, start 6.] Then reset $\lambda \leftarrow$ “good,” and go to (★).
 6. Play e^* until $t = T$. Set $\lambda \leftarrow$ “bad,” and go to (★).
- (★) Interchange $i \leftrightarrow j$, and restart the program.

We can finally determine $Q, S, P, T_0 \geq T_1$, and $\delta_0 \geq \delta_1$ (in that order) for which the above program is subgame perfect. The methodology is first to provide strictly positive penalties—at least 1 below—for any deviation when there is *no* discounting. Then, by continuity, there is some $\delta_0 \in [\delta_1, 1)$ for which all penalties are still positive, as required.

Let β and ω be the best and worst payoffs for any player in G , and suppose $\delta = 1$. For each step, we consider the “worst-case scenario,” where the incentive to deviate is greatest. Then step 4 deters deviations from step 2 if we choose Q so that

$$\omega + Qu_k > \beta + 1 \tag{1}$$

for $k = 1, 2$. Since $u \geq 0$, the threat to restart step 4 dissuades any deviation from 4. Next, step 5 deters deviations from steps 2 and 4 by the older player i if S is chosen to satisfy

$$Q\omega + SU_k(b^k) > \beta + 1$$

for $k = 1, 2$. Also, step 6 will inhibit the younger player j from deviating from steps 2, 3, 4, and 5 if P satisfies

$$(Q + S)\omega + P \max \langle u_k, v_k \rangle > \beta + (Q + S - 1)V_k + P \min \langle u_k, v_k \rangle + 1$$

for $k = 1, 2$. Let $T_0 = \max \langle T_1, P + Q + S \rangle$. Finally, observe that after any deviation, all punishments and rewards are concluded within $P + Q + S$ periods. Thus, by continuity, there is some $\delta_0 \in [\delta_1, 1)$ for which all threats remain credible for any $\delta \geq \delta_0$ and $T \geq T_0$.⁸ Q.E.D.

It is worthwhile at this stage to point out a convenient feature of the above strategies which unfortunately does *not* extend to general n -person games. At the outset of each overlap, the Nash equilibrium of G has been variously used as a carrot or a stick, depending upon circumstances: Previously “good” older players begin their final overlap with their choice of e^* (the Nash outcome) or a^* ; their “bad” counterparts must play whatever outcome they most detest. This use of the Nash outcome as a two-edged sword saves considerably on time and effort. But with three or more players, in equilibrium the Nash outcome might be preferred by one of the “good” younger players and not by the other. Hence, either (i) personalized (out-of-equilibrium) punishment outcomes are needed, or (ii) there must be unanimous preference for one particular outcome over the status quo. In the absence of a full-dimensional payoff space, we must opt for (i); therefore, the compensation phase is hereafter replaced by a *punishment* phase.

Note also that a sufficient statistic for the entire history of the game is captured in the one simple binary “flag” λ . Given that the cardinality of a length τ history in any nondegenerate repeated game is at least 2^τ , it should be especially appealing that players need only condition their actions on such a basic capsule summary of just the last overlap. All equilibrium strategies in this paper are parameterized by a small number of such state variables. But due to the explicit time dependence in our routines, the equilibria are unavoidably *nonstationary*.

When we attempt to generalize the above result to arbitrary n -player games, our first effort falls short. Indeed, we encounter the same game theoretic leap needed to move from $n = 2$ to $n = 3$ as first discovered in (an earlier version of) F–M. Both F–M and Benoit and Krishna (1985) must resort to the following assumption to get any meaningful folk theorem:

⁸ A referee has pointed out that uniformity also holds for all u within any compact set $K \subset \text{int } V$. Indeed, $\delta_1 < 1$ and $T_1 < \infty$ obtain just as before, with minor alterations. Next, we can clearly choose Q so that (1) is true for all $u \in K$, and do likewise for S and P . That is, we can use the same parameters for all $u \in K$, which establishes the claim. This remark applies with equal force to Theorem 3.4 below.

(FD) The feasible and individually rational payoff set for all players has dimension n , or is of *full dimension*.

Now in any *one-shot* game, a given punishment or reward may have externalities. For instance, punishing i may necessarily inflict damage on j ; or conversely, j may reap a windfall from i 's misfortune. But when the game has three or more players, we cannot take advantage of a natural duality inherent in two-player punishment strategies. In effect, two players can simultaneously punish one another (at the mutual minimax point)—a feat not possible when there are three or more players. A full-dimensional payoff space basically sidesteps this impasse by eliminating the externality. That is, (FD) allows us to tailor a given punishment to hurt only its intended victim. As it turns out, however, we may dispense with (FD) altogether if we are willing to sacrifice uniformity.

3.2. A NONUNIFORM n -PLAYER FOLK THEOREM. *Let $u = (u_1, u_2, \dots, u_n)$ be feasible and strictly individually rational for the n -player normal form game G . Then $\forall \varepsilon > 0 \exists T_0 < \infty$ so that $\forall T \geq T_0 \exists \delta_0 < 1$ so that $\delta \in [\delta_0, 1] \Rightarrow \text{OLG}(G; \delta, T)$ has a subgame perfect discounted average payoff within ε distance of u .*

Proof. We assume the entire setting and preliminaries established in Theorem 3.1. That $u_k \neq v_k \forall k \in \mathbf{N}$ is crucial. We may thus choose a correlated strategy c^* so that $U_i(c^*) = (u_i + v_i)/2 > 0 \forall i \in \mathbf{N}$. By design, no player is indifferent between a^* and c^* , although all prefer a^* or c^* to being minimaxed. We may also select correlated strategies a^1, \dots, a^n yielding positive payoffs for all players, and such that $U_k(a^k) < \min \langle (u_k + v_k)/2, u_k \rangle \forall k \in \mathbf{N}$ and $U_k(a^k) \leq U_k(a^l) \forall k, l \in \mathbf{N}$. We should emphasize that a^k is fashioned to hurt k and still leave all players with strictly individually rational (though possibly diminished) payoffs.

At the heart of the program below is the evolution of the parameters λ , γ , and π . When positive, $\lambda \in \mathbf{N} \cup \{0\}$ now records the identities of players deviating late in an overlap, and targeted for the punishment a^λ at the outset of the next overlap. Next, $\gamma \in \mathbf{N} \cup \{0\}$ is an overlap counter, initially set at its equilibrium value of 0; if $\gamma \neq 0$, a change in equilibrium paths ($\pi = 0$ vs $\pi = 1$) is scheduled to occur when player γ is replaced. In equilibrium, a^* [resp. c^*] is played on path $\pi = 0$ [resp. $\pi = 1$]. We denote arbitrary players by j and k , and use the notation $j >_i k$ to mean, literally, “player j is older than k when player i is oldest”; that is, either $i \leq j < k$ or $j < k < i$ or $k < i \leq j$, implying that j will die before k does. Finally, because the strategies of fellow team members are identical, we identify players 1 and $n + 1$ (so that $\gamma = n + 1$ is interpreted as $\gamma = 1$, for instance).

In the overlap depicted below, player i is oldest, and the equilibrium

outcome sequence is $a^*, \dots, a^*; b^i, \dots, b^i$. Once more, the b^i reward sequence lasts S periods. We begin on the path $\pi = 0$ with $\lambda = \gamma = 0$, and $r = 1$.

⊙ Reset $t \leftarrow 1$. If $\lambda = 0$, start 2; else start 1.

1. Play a^λ until $t = P$. [If j deviates, start 4.]
 2. Play a^* if $\pi = 0$ and c^* if $\pi = 1$ until $t = T - S$. [If j deviates at $t \leq T - Q_1 - S$, start 4; if k deviates at $t > T - Q_1 - S$, start 7 if $k = i$ and start 8 if $k \neq i$.]
 3. Play b^i until $t = T$. [If $k \neq i$ deviates, start 8.] Then reset $\lambda \leftarrow 0$, and go to (★).
 4. Play M^j for Q_r periods. [If $k \neq j$ deviates at $t \leq T - Q_n - R - S$: Start 5 if $v_k < u_k$ and $\pi = 0$ or if $v_k > u_k$ and $\pi = 1$, but start 6 if $v_k > u_k$ and $\pi = 0$ or if $v_k < u_k$ and $\pi = 1$. If $k \neq j$ deviates at $t > T - Q_n - R - S$, start 7 if $k = i$ and start 8 if $k \neq i$.] Then reset $r \leftarrow 1$, set $\gamma \leftarrow j$, and return to the main path.
 5. Switch paths $\pi \leftarrow 1 - \pi$, and return to the main path.
 6. (a) If $j >_i k$, set $\gamma \leftarrow k$. Then return to the main path.
(b) If $k >_i j$, set $j \leftarrow k$. Then increment $r \leftarrow r + 1$, and restart 4.
 7. Play M^i until $t = T$. [If $k \neq i$ deviates, start 8.] Then reset $\lambda \leftarrow 0$, and go to (★).
 8. Play e^* until $t = T$. Then set $\lambda \leftarrow k$, and go to (★).
- (★) If $\gamma = i$, switch paths $\pi \leftarrow 1 - \pi$ and reset $\gamma \leftarrow 0$. Increment $i \leftarrow i + 1$. Then restart the program.

See the Appendix for selections of suitable parameter values, T_0 , and, when $T \geq T_0$, a minimal discount factor $\delta_0(T)$ so that the program is subgame perfect.⁹ Q.E.D.

We now briefly review the essence of the program. It is the creation of a dual-track equilibrium path that permits us to exploit the players' strict preferences between a^* and c^* . Focus on the many plays of a^* [or c^* , its off-the-equilibrium-path counterpart]. First note that we need not worry about any deviations "late" in an overlap (in the buffer zone), for we can minimax the oldest player i until he dies, effectively denying him his coveted reward phase. This is buttressed by the threat to impose the next overlap's punishment phase should any younger player j deviate (for instance, by opting not to administer the above lengthy punishment). Now consider the play prior to the buffer zone. We discourage any early

⁹ To ensure that strategies are resilient, we could have used an optional flag $\rho \in \mathbf{N} \cup \{0\}$, so that when $\rho = i$ the program is set to return to the principal equilibrium path $\pi = 0$ exactly $n - 1$ overlaps after player i dies. This measure is later incorporated into Theorem 4.1.

deviations from a^* [or c^*] by threatening to minimax the culprit Q_1 times. Here is the rub: Is this threat credible? Indeed, the punishers may suffer more than the criminal. Because $\delta < 1$, we must at all costs avoid an upward spiral of progressively longer minimax phases. We therefore exploit the natural age asymmetry of the overlapping generations setting. For earlier deviants from the minimax phase who prefer the path $\pi = 0$ [or $\pi = 1$], we can credibly threaten to abandon the minimax phase and switch paths. But what if a punisher who prefers the other path $\pi = 1$ [or $\pi = 0$] deviates? Well, if he is younger than the original culprit, we tell him, "Had you obeyed the equilibrium, you would have enjoyed your preferred path $\pi = 1$ [or $\pi = 0$] as soon as the *culprit* died; however, we shall now defer this path change until *you* die." Otherwise, if the deviant punisher is older than the culprit, we proceed to minimax the would-be punisher for Q_2 periods. We choose Q_1, Q_2, \dots, Q_n inductively to ensure that progressively older players will opt not to deviate from an ongoing minimax phase.

It is also important to observe that the program and its verification make no essential use of the *identical* length of the overlaps; it only requires that each exceed the computed value T . Thus, our folk theorem actually applies to the slightly more general class of games considered in Kandori (1989). Setting $\mathbf{T} \equiv (T_1, T_2, \dots, T_n)$, we may define OLG $(G; \delta, \mathbf{T})$ exactly as was OLG $(G; \delta, T)$, except that point 2 of its definition is replaced by

2*. Player $k + tn$ from team S_k dies and is replaced by player $k + (t + 1)n$ just after time $(T_1 + \dots + T_k) + t(T_1 + \dots + T_n)$, for $k \in \mathbf{N}$ and $t = 0, 1, 2, \dots$

This yields a quick constructive proof of

3.3. THEOREM (Kandori, 1989). *Let $u = (u_1, u_2, \dots, u_n)$ be feasible and strictly individually rational for the n -player normal form game G . Then $\forall \varepsilon > 0 \exists T_0 < \infty$ so that $\forall \mathbf{T} \gg (T_0, \dots, T_0) \exists \delta_0 < 1$ so that $\delta \in [\delta_0, 1] \Rightarrow$ OLG $(G; \delta, \mathbf{T})$ has a subgame perfect discounted average payoff with ε of u .*

The original result of Kandori (1989) is actually weaker, as T_1 depends on (T_2, \dots, T_n) . It proves illuminating to contrast his result and strategies with ours. To a large extent, the two analyses have fundamentally different focuses: Kandori emphasizes not the overlap, but the *generation*, where generation $t \geq 0$ consists of players $\{1 + tn, 2 + tn, \dots, n + tn\}$. Instead of the n -fold "ratcheted" minimax mechanism that is the keystone of our approach, Kandori devotes the *entire* overlaps when players $2 + tn, 3 + tn, \dots, n + tn$ are the oldest to maintaining equilibrium discipline over generation t ; it is here that his "terminal payments" are apportioned, irrespective of the target equilibrium payoff. Unfortunately, this procedure

1, 1, 1	0, 0, 0
0, 0, 0	0, 0, 0

0, 0, 0	0, 0, 0
0, 0, 0	1, 1, 1

FIGURE 3

compromises any possible application of his result to all but the most asymmetrical age hierarchies. Indeed, his proof requires $(T_2 + \dots + T_n)/T_1 \rightarrow 0$ as $\varepsilon \rightarrow 0$. Essentially, his folk theorem is best suited to economic organizations with periodic mass (though staggered) departures of all players.

Note that we often opt to wait as many as $n - 1$ overlaps until a deviator “dies” before throwing a party for his punishers. Unfortunately, the very nature of this approach limits the utility of our folk theorem. For such threats can remain credible only so long as $\delta_0^{(n-1)T}$ exceeds some constant $\theta \in (0, 1)$. It is in this sense that the overlap length T and the minimum possible δ_0 must covary.

Even more telling is a simple continuous time implication of the covariance condition. That is, let $\delta_0 = e^{-r\Delta}$, where $r > 0$ is the interest rate and Δ the time interval between periods, and set $\tau = \Delta T$. Then our equilibrium is subgame perfect so long as $r(n - 1)\tau < -\log \theta < \infty$. In other words, players cannot live “too long” or *our* folk theorem might not obtain. Whether this result is true in general is an open question; however, in discrete time, we can say something more definite: For a *given* discount factor, the overlap T cannot be too large, or would-be punishers may view the (heavily discounted) future rewards as insufficient compensation for repeatedly minimaxing another player. We illustrate this hitch with the simple coordination game in Fig. 3 devised by F–M.

EXAMPLE 3. Consider G_3 . In this game, player 1 chooses rows, 2 chooses columns, and 3 chooses matrices. It is crucial that all three players receive the same payoff, yielding a one-dimensional individually rational set of feasible payoffs. Thus (FD) is *not* satisfied. Despite the fact that the minimax point is $(0, 0, 0)$, we have the following result¹⁰:

CLAIM. For all $0 < \delta < 1$ and $0 < \varepsilon < \frac{1}{4}$, $\exists T_0 < \infty$ so that $T \geq T_0 \Rightarrow$ OLG $(G_3; \delta, T)$ has no subgame perfect discounted average payoff less than $\frac{1}{4} - \varepsilon$.

¹⁰ The claim is stronger than it seems, as its proof in no way depends on our standard assumption that team members use identical strategies. I am grateful to a referee for inquiring on this point.

Proof. Let the overlap length be T . Let $w_k(t)$ be the least subgame perfect discounted *total* payoff for a player with t periods to go in his k^{th} last overlap. Then $w_3(T)$ is the worst punishment we may inflict on any player. We remark that since $T < \infty$, each $w_k(t)$ does indeed exist. We now shift gears. Decompose $w_2(t)$ and $w_3(t)$ into payoffs accruing in each overlap, i.e., $w_2(t) \equiv w_2^1(t) + \delta^t w_2^2(T; t)$ and $w_3(t) \equiv w_3^1(t) + \delta^t w_3^2(T; t) + \delta^{t+T} w_3^3(T; t)$.

Next, as F–M argue, for any given play of the game, the Pigeon-Hole Principle yields some player who can gain at least $\frac{1}{4}$ by deviating. This is the linchpin to the whole analysis. Define $w_k(0) = 0 \forall k$. In order that no one-shot deviations by the oldest, middle, or youngest players (respectively) be profitable, we must have for $t = 1, 2, \dots, T$

$$w_3(t) - \delta^t w_3^2(T; t) - \delta^{t+T} w_3^3(T; t) \equiv w_3^1(t) \geq \frac{1}{4} + \delta w_1(t-1),$$

or

$$w_3(t) - \delta^{t+T} w_3^3(T; t) \equiv w_3^1(t) + \delta^t w_3^2(T; t) \geq \frac{1}{4} + \delta w_2(t-1),$$

or

$$w_3(t) \geq \frac{1}{4} + \delta w_3(t-1);$$

and

$$w_2(t) - \delta^t w_2^2(T; t) \equiv w_2^1(t) \geq \frac{1}{4} + \delta w_1(t-1),$$

or

$$w_2(t) \geq \frac{1}{4} + \delta w_2(t-1),$$

or

$$w_2(t) + \delta^{t+T} \frac{1 - \delta^T}{1 - \delta} \geq \frac{1}{4} + \delta w_3(t-1);$$

and

$$w_1(t) \geq \frac{1}{4} + \delta w_1(t-1),$$

or

$$w_1(t) + \delta^t \frac{1 - \delta^T}{1 - \delta} \geq \frac{1}{4} + \delta w_2(t - 1),$$

or

$$w_1(t) + \delta^t \frac{1 - \delta^{2T}}{1 - \delta} \geq \frac{1}{4} + \delta w_3(t - 1).$$

Because each $w_k^j(T; t) < 1/(1 - \delta)$, we can choose a minimum overlap length $T_0 < \infty$ sufficiently large so that if $T \geq T_0$, then the above sets respectively imply that

$$w_3(t) \geq \frac{1}{4} + \min \langle \delta w_1(t - 1), \delta w_2(t - 1), \delta w_3(t - 1) \rangle,$$

and

$$w_2(t) \geq \frac{1}{4} + \min \langle \delta w_1(t - 1), \delta w_2(t - 1), \delta w_3(t - 1) - \varepsilon/2 \rangle,$$

and

$$w_1(t) \geq \frac{1}{4} + \min \langle \delta w_1(t - 1), \delta w_2(t - 1) - \varepsilon/2, \delta w_3(t - 1) - \varepsilon/2 \rangle,$$

where $t = T, T - 1, \dots, [2\varepsilon T] + 1$. (Here, $[x]$ is the integer part of x .) By iterating these inequalities over the relevant range of t 's, and using $w_k(t - 1) \geq 0$, we have $\forall k$

$$\begin{aligned} w_k(T) &\geq \left(\frac{1}{4} - \varepsilon/2\right) + \delta \min \langle w_1(T - 1), w_2(T - 1), w_3(T - 1) \rangle \\ &\geq \dots \geq \left(\frac{1}{4} - \varepsilon/2\right)(1 + \delta + \delta^2 + \dots + \delta^{l(1-2\varepsilon)T-1}) \\ &\geq \left(\frac{1}{4} - \varepsilon/2\right)(1 - 2\varepsilon)(1 + \delta + \dots + \delta^T) \\ &= \left(\frac{1}{4} - \varepsilon + \varepsilon^2\right)(1 + \delta + \dots + \delta^T) \\ &> \left(\frac{1}{4} - \varepsilon\right)(1 + \delta + \dots + \delta^{3T}), \end{aligned}$$

where the last inequality is assured for T_0 large enough. Normalizing $w_3(T)$, we find that the worst subgame perfect discounted *average* payoff is at least $\frac{1}{4} - \varepsilon$, as required.

The absence of a full-dimensional payoff space thus implies an upper bound on the overlap length, and consequently on the players' lifespans. A nonuniform folk theorem is simply the best we can hope for with this game.¹¹

This naturally leads us to consider an n -player uniform folk theorem.

3.4. A UNIFORM n -PLAYER FOLK THEOREM. *Let $u = (u_1, \dots, u_n)$ be a feasible and strictly individually rational payoff vector for \mathbf{G} . Suppose that (FD) holds. Then $\forall \varepsilon > 0 \exists T_0 < \infty$ and $\delta_0 < 1$ so that $T \geq T_0$ and $\delta \in [\delta_0, 1] \Rightarrow \text{OLG}(\mathbf{G}; \delta, T)$ has a subgame perfect discounted average payoff within ε of u .*

Proof. We now use much of the framework of Theorem 3.2. For instance, the typical equilibrium overlap is the same. The principal point of departure in this analysis is that (FD) allows us to design each personal punishment vector a^λ so that only player λ suffers (yet with a positive payoff), and all others receive their equilibrium payoffs. Although not required for the initial punishment phase, this feature does allow us to forego the labyrinthine minimax mechanism. We illustrate the equilibrium program for the overlap when i is oldest, and begin with $\lambda = 0$.

⊙ Reset $t \leftarrow 1$. If $\lambda = 0$, start 2; else start 1.

1. Play a^λ until $t = P$. [If j deviates at $t \leq T - Q - R - S$, start 4; if k deviates at $t > T - Q - R - S$, start 6 if $k = i$ and start 7 if $k \neq i$.]

2. Play a^* until $t = T - S$. [If j deviates at $t \leq T - Q - R - S$, start 4; if k deviates at $t > T - Q - R - S$, start 6 if $k = i$ and start 7 if $k \neq i$.]

3. Play b^i until $t = T$. [If $k \neq i$ deviates, start 7.] Then reset $\lambda \leftarrow 0$, and go to (★).

4. Play M^j for Q periods. [If $l \neq j$ deviates, start 5.] Then set $l \leftarrow j$.

5. Play a^l for R periods. [If j deviates at $t \leq T - Q - R - S$, restart 4; if k deviates at $t > T - Q - R - S$, start 6 if $k = i$ and start 7 if $k \neq i$.] Then return to the main path.

6. Play M^i until $t = T$. [If $k \neq i$ deviates, start 7.] Then reset $\lambda \leftarrow 0$, and go to (★).

7. Play e^* until $t = T$. Then set $\lambda \leftarrow k$, and go to (★).

(★) Increment $i \leftarrow i + 1$, and restart the program.

See the Appendix for selection of T_0 and δ_0 so that the program is subgame perfect for suitable parameter values. Q.E.D.

¹¹ In point of fact, Smith (1990) shows that (FD) can be replaced here and in the finitely and infinitely repeated contexts with a less onerous requirement: The payoff space must have two-dimensional projections onto every coordinate plane.

We comment briefly on the above program. The minimax and subsequent R period *recovery* phases are a slight twist on the two-tier punishment scheme in F–M: The deviant is first minimized, and then (he alone) suffers through the recovery phase. Should anyone deviate from minimizing another, the recovery period commences at once (*unlike* F–M)—only now it is the deviant punisher who experiences his personalized punishment vector. As an aside, such a punishment mechanism does not resort to “punishment loops,” i.e., where j must minimize k because k fails to minimize j .

4. SUMMARY AND EXTENSIONS

Thus far, the uniform and nonuniform folk theorems might appear distant cousins at best. Although similar in statement, the different premises lead to somewhat disparate proofs. And yet in the search for subgame perfection, both results must essentially tackle the same problem: How does one disentangle the payoff streams of all n players? Indeed, when $n > 2$, this is the principal difficulty faced. There must be proper compensation for engaging in the potentially self-lacerating punishment of a fellow player. But if, in a standard supergame with discounting, all players receive identical payoffs, this may not be possible without also rewarding the original deviant. F–M is unequivocal on this point. At the very least, players’ payoff sequences must be linearly independent.

The uniform and nonuniform folk theorems represent two distinct resolutions of this quandary. In the uniform case, full-dimensionality allows us to fully disentangle the payoff streams at the stage game level; consequently, the OLG structure is really not needed for this purpose. Not surprisingly, F–M and Benoit and Krishna (1985) exercise this option in their models. In the nonuniform case, however, to deal with games like G_3 , we can repeatedly exploit the fact that the younger players will continue to receive payoffs long after the older ones are gone. Intuitively then, to ensure that all deviations are eventually punished, players must condition their actions in the current overlap on happenings over the past $n - 1$ overlaps—that is, for as long as any player now in the game has been alive. Prior history can be safely ignored.

In the sequel, we consider a cross-breed of our OLG model and a standard supergame. In particular, we define $OLG^m(G; \delta, T)$ analogously to $OLG(G; \delta, T)$, except that the first $m < n$ of the n teams are rather degenerate, each consisting of just a single *infinite*-lived player.¹² Thus, at the end of exactly m of the “overlaps” in each “generation,” no player

¹² This differs from the work of Fudenberg *et al.* (1990) and Fudenberg and Levine (1989), where the “short-run” players each live for one period only.

will die. Note that $OLG^0(\mathbf{G}; \delta, T) \equiv OLG(\mathbf{G}; \delta, T)$, while $OLG^n(\mathbf{G}; \delta, T)$ is rather awkward notation for a standard supergame. Next, let P^mV be the orthogonal projection of V onto its first m coordinates (in R^m). For instance, $P^nV \equiv V$, whereas P^1V is simply the range of feasible payoffs for team 1. We then have the following natural extension of both Theorem 3.2 and the perfect information folk theorem in F–M.

4.1. NONUNIFORM FOLK THEOREM FOR A HYBRID OLG-SUPERGAME MODEL. *Let $u = (u_1, u_2, \dots, u_n)$ be feasible and strictly individually rational for the n -player normal form game \mathbf{G} . If $\dim P^mV = m$, then $\forall \varepsilon > 0 \exists T_0 < \infty$ so that $\forall T \geq T_0 \exists \delta_0 < 1$ so that $\delta \in [\delta_0, 1] \Rightarrow OLG^m(\mathbf{G}; \delta, T)$ has a subgame perfect discounted average payoff within ε of u .*

Proof. See the Appendix.

Note that Theorem 3.2 considers the case $m = 0$, whereas F–M proved that uniformity obtains when $m = n$. A remarkable corollary of this result is that Theorem 3.2 also holds when exactly one of the teams is a single infinitely lived player.¹³ In retrospect, the logic is most compelling: Indeed, the earlier algorithm for our nonuniform folk theorem is valid so long as the life expectancy of any two current players differs by at least one overlap.

APPENDIX

Verification of the Program for Theorem 3.2. As with Theorem 3.1, we first suppose $\delta = 1$, and show that each punishment can be made an effective deterrent. Since $U_k(a^k) > 0$ for all $k \in \mathbf{N}$, we may choose Q_1 so that

$$\omega + Q_1U_k(a^k) > \beta + 1 \tag{2}$$

for all $k \in \mathbf{N}$. This ensures that step 4 deters early deviations from steps 1 and 2. We then select the Q 's inductively so that given Q_l, Q_{l+1} satisfies

$$Q_l\omega + (Q_{l+1} - Q_l + 1)U_k(a^k) > \beta + 1$$

for all $k \in \mathbf{N}$ and $l = 1, 2, \dots, n - 1$. This guarantees that the progressively longer minimax phases promised by step 6(b) prevent (appropriate) deviations from step 4. Next, for step 5 to discourage (appropriate) deviations from step 4, we must choose R large enough that

¹³ These demographics correspond, for instance, to those in Atkeson (1991).

$$Q_n\omega + R \max \langle u_k, (u_k + v_k)/2 \rangle > \beta + (Q_n + R - 1) \min \langle u_k, (u_k + v_k)/2 \rangle + 1$$

for all $k \in \mathbf{N}$. Step 7 deters late deviations from steps 2 and 4 by the oldest player if S satisfies

$$Q_n\omega + SU_k(b^k) > \beta + 1$$

for all $k \in \mathbf{N}$. Continuing, step 8 discourages all other players' late deviations from steps 2, 3, 4, and 7 if we select P so that

$$(Q_n + S)\omega + (P + R) \min \langle u_k, (u_k + v_k)/2 \rangle > \beta \\ + (Q_n + R + S - 1)v_k + PU_k(a^k) + 1$$

for all $k \in \mathbf{N}$. Note that we can define δ_1 and T_1 as was done in Theorem 3.1. Next, we select a minimum overlap length $T_0 \geq \max \langle T_1, P + Q_n + R + S \rangle$ so that step 6(a) deters (appropriate) deviants from step 4:

$$Q_n\omega + (T_0 - S) \max \langle u_k, (u_k + v_k)/2 \rangle > \beta + (Q_n - 1) \max \langle u_k, (u_k + v_k)/2 \rangle \\ + (T_0 - S) \min \langle u_k, (u_k + v_k)/2 \rangle + 1$$

for all $k \in \mathbf{N}$. Finally, given $T \geq T_0$, there is some level of discounting $\delta_0 \in [\delta_1, 1)$ for which the program remains credible.

Verification of the Program for Theorem 3.4. Assume $\delta = 1$. Then step 4 deters early deviations from steps 1, 2, and 5 if Q satisfies (2). Step 5 will prevent any deviations from step 4 if R is large enough that

$$Q\omega + Ru_k > \beta + RU_k(a^k) + (Q - 1)u_k + 1$$

for all $k \in \mathbf{N}$. Next, step 6 discourages late deviations from steps 1, 2, and 5 by the oldest player if we select S so that

$$Q\omega + RU_k(a^k) + SU_k(b^k) > \beta + 1$$

for all $k \in \mathbf{N}$. Finally, step 7 deters all younger players from deviating late during steps 1, 2, 3, 5, and 6 if P satisfies

$$(Q + R + S)\omega + Pu_k > \beta + (Q + R + S - 1)v_k + PU_k(a^k) + 1$$

for all $k \in \mathbf{N}$. As this program is viable for any $T \geq P + Q + R + S$, we may let $T_0 \equiv \max \langle T_1, P + Q + R + S \rangle$. Finally, because all punishments

are completed within $P + Q + R + S$ periods, the equilibrium remains subgame perfect for discount factors exceeding some $\delta_0 \in [\delta_1, 1)$.

Q.E.D.

Proof of Theorem 4.1. We first note that if $\dim P^m V = m$, then $\dim V = d \geq m$. Hence, V is locally the intersection of $n - d$ hyperplanes in R^n . In particular, we may assume WLOG that the players are ordered so that the payoff vector $x = (x_1, \dots, x_n)$ close to u is feasible exactly when

$$x_r = \alpha_{r,0} + \alpha_{r,1}x_1 + \dots + \alpha_{r,d}x_d \tag{3}$$

for $r = d + 1, \dots, n$. Then, as with Theorem 3.2, we may select n correlated strategies a^1, \dots, a^n yielding positive payoffs to all players, with a^k designed to give player k a payoff below $\min \langle u_k, (u_k + v_k)/2 \rangle$; however, by (3), we may insist that for $k = 1, 2, \dots, m$, the punishment a^k leaves all other players $j \neq k, j = 1, 2, \dots, m$ with their equilibrium payoffs.

We now carefully amalgamate the programs of Theorems 3.2 and 3.4. The idea is the following: For deviations by the infinitely lived players $1, 2, \dots, m$, the threat to shift into a recovery phase deters deviations from the minimax phase; for all other players, we must use the punishment mechanism of our nonuniform folk theorem.

The typical equilibrium overlap is once more the same as in Theorem 3.2. We begin on the path $\pi = 0$ with $\lambda = \gamma = \rho = 0$, and $r = 1$. Player i is the oldest.

⊙ Rest $t \leftarrow 1$. If $\lambda = 0$, start 2; else start 1.

1. Play a^λ until $t = P$. [If j deviates, start 4.]
2. Play a^* if $\pi = 0$ and c^* if $\pi = 1$ until $t = T - S$. [If j deviates at $t \leq T - Q_{n-m+1} - R - S$, start 4; if k deviates at $t > T - Q_{n-m+1} - R - S$, start 8 if $k = i$ and start 9 if $k \neq i$.]
3. Play b^i until $t = T$. [If $k \neq i$ deviates, start 9.] Then reset $\lambda \leftarrow 0$, and go to (★).
4. Play M^j for Q_r periods. [If $k \neq j$ deviates at $t \leq T - Q_{n-m+1} - R - S$: Set $l \leftarrow k$ and start 5 if $k \leq m$; else start 6 if $v_k < u_k$ and $\pi = 0$ or if $v_k > u_k$ and $\pi = 1$, but start 7 if $v_k > u_k$ and $\pi = 0$ or if $v_k < u_k$ and $\pi = 1$. If $k \neq j$ deviates at $t > T - Q_{n-m+1} - R - S$, start 8 if $k = i$ and start 9 if $k \neq i$.] Then set $\gamma \leftarrow j, l \leftarrow j$, and reset $r \leftarrow 1$. Return to the main path if $j > m$.
5. Reset $r \leftarrow 1$, and play a^l for R_1 periods. [If j deviates at $t \leq T - Q_{n-m+1} - R - S$, restart 4; if k deviates at $t > T - Q_{n-m+1} - R - S$, start 8 if $k = i$ and start 9 if $k \neq i$.] Then return to the main path.

6. Switch paths $\pi \leftarrow 1 - \pi$, and return to the main path.
7. (a) If $j >_i k$, set $\gamma \leftarrow k$, $\rho \leftarrow i$, and return to the main path.
 (b) If $k >_i j$, set $j \leftarrow k$ and $\rho \leftarrow i$. Then increment $r \leftarrow r + 1$, and restart 4.
8. Play M^i until $t = T$. [If $k \neq i$ deviates, start 9.] Then reset $\lambda \leftarrow 0$, and go to (\star).
9. Play e^* until $t = T$. Then set $\lambda \leftarrow k$, and go to (\star).
- (\star) If $\gamma = i$, switch paths $\pi \leftarrow 1 - \pi$ and reset $\gamma \leftarrow 0$. Increment $i \leftarrow i + 1$. If $\rho = i$, reset $\pi \leftarrow 0$ and $\rho \leftarrow 0$. Then restart the program.

We leave it as an exercise to verify that there are values of $Q_1, \dots, Q_{n-m+1}, R > R_1, S, P$, and T_0 so that all threats in the above program are strict deterrents. The theorem follows. Q.E.D.

REFERENCES

- ABREU, D., PEARCE, D., AND STACCHETTI, E. (1989). "Renegotiation and Symmetry in Games," mimeo, Harvard University, May.
- ATKESON, A. (1991). "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica* **59**, 1069–1089.
- AUMANN, R. J. (1985). "Repeated Games," in *Issues in Contemporary Microeconomics and Welfare* (G. R. Feiwel, Ed.), pp. 209–242. Albany: State Univ. of New York Press.
- BERNOIT, J.-P., AND KRISHNA, V. (1985) "Finitely Repeated Games," *Econometrica* **53**, 905–922.
- COOPER, R., AND DAUGHETY, A. (1988). "Cooperation in Infinite Horizon Games with Finitely Lived Players," mimeo, University of Iowa, December.
- CRÉMER, J. (1986). "Cooperation in Ongoing Organizations," *Quart. J. Econ.*, **101**, 33–49.
- FRIEDMAN, J. (1971). "A Noncooperative Equilibrium for Supergames," *Rev. Econ. Stud.* **38**, 1–12.
- FUDENBERG, D., KREPS, D., AND MASKIN, E. (1990). "Repeated Games with Long-Run and Short-Run Players," *Rev. Econ. Stud.* **57**, 555–573.
- FUDENBERG, D., AND LEVINE, D. (1989). "Equilibrium Payoffs with Long-Run and Short-Run Players and Imperfect Information," M.I.T. Economics working paper 524, June.
- FUDENBERG, D., AND MASKIN, E. (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica* **54**, 533–554.
- HAMMOND, P. (1975). "Charity: Altruism or Cooperative Egoism?" in *Altruism, Morality and Economic Theory* (E. Phelps, Ed.), pp. 115–131. New York: Russell Sage Foundation.
- KANDORI, M. (1989). "Repeated Games Played by Overlapping Generations of Players," *Rev. Econ. Stud.*, to appear.
- RUBINSTEIN, A. (1979). "Equilibrium in Supergames with the Overtaking Criterion," *J. Econ. Theory* **21**, 1–9.

- SALANT, D. (1991). "A Repeated Game with Finitely Lived Overlapping Generations of Players," *Games Econ. Behav.* **3**, 244–259.
- SELTEN, R. (1975). "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *Int. J. Game Theory* **4**, 25–55.
- SMITH, L. (1990). "Folk Theorems: Two-Dimensionality Is (Almost) Enough," mimeo, University of Chicago, October.