



Repeated Games with Long-run and Short-run Players

Citation

Fudenberg, Drew, David M. Kreps, and Eric S. Maskin. 1990. Repeated games with long-run and short-run players. *Review of Economic Studies* 57, no. 4: 555-573.

Published Version

<http://dx.doi.org/10.2307/2298086>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3226950>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

REPEATED GAMES

WITH

LONG-RUN AND SHORT-RUN PLAYERS

BY

DREW FUDENBERG, DAVID M. KREPS, AND ERIC MASKIN

JANUARY 1989

M.I.T., DEPARTMENT OF ECONOMICS; STANFORD UNIVERSITY, GRADUATE SCHOOL OF BUSINESS; AND HARVARD UNIVERSITY, DEPARTMENT OF ECONOMICS, RESPECTIVELY. WE WOULD LIKE TO THANK MATT McCARTER FOR POINTING OUT AN ERROR IN AN EARLIER DRAFT.

ABSTRACT

This paper studies the set of equilibrium payoffs in games with long- and short-run players and little discounting. Because the short-run players are unconcerned about the future, equilibrium outcomes must always lie on their static reaction (best response) curves. The obvious extension of the Folk Theorem to games with this constraint would simply include the constraint in the definitions of the feasible payoffs and of the minmax values. This extension does obtain under the assumption that each player's choice of a mixed strategy for the stage game is publicly observable, but, in contrast to standard repeated games, the limit value of the set of equilibrium payoffs is different if players can observe only their opponents' realized actions.

1. Introduction

The "folk theorem" for repeated games with discounting says that (under mild conditions) each individually-rational payoff can be attained in a perfect equilibrium for a range of discount factors close to one. It has long been realized that results similar to the folk theorem can arise if some of the players play the constituent game infinitely often and others play the constituent game only once, so long as all of the players are aware of all previous play. A standard example is the infinitely-repeated version of Selten's [1977] chain-store game, where a single incumbent faces an infinite sequence of short-run entrants in the game depicted in Figure 1. Each entrant cares only about its one-period payoff, while the incumbent maximizes its net present value. For discount factors close to one there is a perfect equilibrium in which entry never occurs, even though this is not a perfect equilibrium if the game is played only once or even a fixed finite number of times. In this equilibrium, each entrant's strategy is "stay out if the incumbent has fought all previous entry; otherwise, enter;" and the incumbent's strategies is "fight each entry as long as entry has always been fought in the past, otherwise acquiesce." Other examples of games with long and short run players are the papers of Dybvig-Spatt [1980] and Shapiro [1982] on a firm's reputation for producing high-quality goods and the papers of Simon [1951] and Kreps [1984] on the nature of the employment relationship.

This paper studies the set of equilibrium payoffs in games with long- and short-run players and little discounting. This set differs from what it would be if all players were long-run, as demonstrated by the prisoner's dilemma with one enduring player facing a sequence of short-run opponents.

Because the short-run players will fink in every period, the only equilibrium is the static one, no matter what the discount factor. In general, because the short-run players are unconcerned about the future, equilibrium outcomes must always lie on their static reaction (best response) curves. This is also true off of the equilibrium path, so the values of the long-run players are higher when some of their opponents are short-run, because their punishments must be drawn from a smaller set.

The perfect folk theorem for discounted repeated games (Fudenberg-Maskin [1986]) shows that, under a mild full-dimensionality condition, any feasible payoffs that give all players more than their minmax values can be attained by a perfect equilibrium if the discount factor is near enough to one.¹ The obvious extension of this result to games with the constraint that short-run players always play static best responses would simply include that constraint in the definitions of the feasible payoffs and of the minmax values. Propositions 1 and 2 of Section 2 shows that this extension does obtain under the assumption that each player's choice of a mixed strategy for the stage game is publicly observable.

We then turn to the more realistic case in which players observe only their opponents' realized actions and not their opponents' mixed strategies. While in standard repeated games the folk theorem obtains in either case, when there are some short-run players the set of equilibria can be strictly smaller if mixed strategies are not observed. The explanation for this difference is that in ordinary repeated games, while mixed strategies may be needed during punishment phases, they are not necessary along the equilibrium path. In contrast, with short-run players some best responses, and thus some of the feasible payoffs, can only be obtained if the long-run players use

mixed strategies. If the mixed strategies are not observable, inducing the long-run players to randomize may require that "punishments" occur with positive probability even if no player has deviated, and the set of equilibrium payoffs may be bounded away from the frontier of the feasible set. This fact in itself would not preclude a folk theorem for the same is true in repeated games with moral hazard, where the folk theorem does indeed obtain under the conditions provided by Fudenberg-Levine-Maskin [1989]. However, one of their conditions is that the highest feasible payoff for each player i be attainable with a strategy m^i where that player uses a static best response against the strategies m_{-i}^i of his opponent. This condition is not satisfied in games with short-run players, as we explain in Section 3.

Proposition 3 of Section 3 provides a complete characterization of the limiting value of the set of equilibrium payoffs for a single long-run player. This characterization, and the results of Section 2, assume that players have access to a publicly observable randomizing device. The device is used to implement strategies of the form: "If player i deviates, then players jointly switch to a "punishment equilibrium" with some probability $p < 1$."

While the assumption of public randomizations is not implausible, it is interesting to know whether it leads to a larger limit set of equilibrium payoffs. Proposition 4 in Section 4 shows that it does not. We construct equilibria in "target strategies," in which a player is punished whenever his discounted payoff to date exceeds a target value. Whenever the player's payoff to date is below the target, he is given a stochastic "reward" whose expected value exceeds the target level. We show that no strategy can yield a player more than his target, and that the specified strategies do yield the target payoffs so long as these targets are below the critical level of

Proposition 3. Thus the same set of equilibria can be obtained without public randomizations.

One nice feature of the target strategies is that they have a simple form. This simplicity helps us provide additional insight into the question of why some feasible payoff cannot be attained by equilibria. Inspection of the strategies shows that they do not form an equilibrium for target payoffs that are higher than what the long-run player can obtain with probability one given that the short run player play static best responses. For payoffs this high, there is a positive probability that player 1 will suffer a run of "bad luck" after which no possible sequence of payoffs could draw his discounted normalized value up to the target.

As this problem does not arise under the criterion of time-average payoffs, one might wonder if the set of equilibrium payoffs is larger under time-averaging. Proposition 5 shows that the answer is yes. In fact, any payoffs that are feasible given the restriction that short-run players play static best responses can arise as equilibria with time-averaging, so that we obtain the same set of payoffs as in the case where the player's privately mixed strategies are observable. We prove this with "target strategies" of the kind cited in Proposition 4. The fact that the equilibrium set expands discontinuously in passing from discounting to time averaging is reminiscent of a similar discontinuity that has been established for the equilibria of repeated partnership games (Radner [1986], Radner-Myerson-Maskin [1986]). Our simple construction may help shed light on the reason for discontinuities there as well. The relationship between the two models is discussed further in Section 5.

The case of several long-run players and unobservable mixed strategies

is more difficult, as shown by example 2 in Section 6. This case has recently been solved by Fudenberg-Levine [1989a], who extend the techniques of Fudenberg-Levine-Maskin [1989] to games where the folk theorem does not obtain.

2. Observable Mixed Strategies

Consider a finite n-player game in normal form,

$$g: S_1 \times \dots \times S_n \rightarrow \mathbb{R}^n .$$

We denote player i's mixed strategies by $\sigma_i \in \Sigma_i$, and write $g(\sigma)$ for the expected value of g under distribution σ .

In this section we assume that a player can observe the others' past mixed strategies. This assumption (or a restriction to pure strategies) is standard in the repeated games literature, but as Fudenberg-Maskin [1986] [1987a] have shown, it is not necessary there. (Here it matters -- see the next section!) We will also assume that the players can make their actions contingent on the outcome of a publicly observable randomizing device.

Label the players so that players 1 to ℓ are long-run and $\ell+1$ to n are short-run. Let

$$B: \Sigma_1 \times \dots \times \Sigma_\ell \rightarrow \Sigma_{\ell+1} \times \dots \times \Sigma_n$$

be the correspondence which maps any strategy selection $(\sigma_1, \dots, \sigma_\ell)$ for the long-run players to the corresponding Nash equilibria strategy selections for the short-run players. That is, for each $\sigma \in \text{graph}(B)$, and each $i \geq \ell+1$,

σ_i is a best response to σ_{-i} . If there is only one short-run player, $B(\sigma)$ is his best response correspondence.

For each i from 1 to j , choose $m^i = (m_1^i, \dots, m_n^i)$ so that m^i solves

$$m^i \min_{m^i \in \text{graph}(B)} \max_{\sigma_i} g_i(\sigma_i, m_{-i}^i),$$

and set

$$v_i = \max_{\sigma_i} g_i(\sigma_i, m_{-i}^i).$$

(This minimum is attained because the constraint set $\text{graph}(B)$ is compact and the function $\max_{\sigma_i} g_i(\sigma_i, m_{-i}^i)$ is continuous in m_{-i}^i .)

The strategies m_{-i}^i minimize long-run player i 's maximum attainable payoff over the graph of B . The restriction to this set reflects the constraint that the short-run players will always choose actions that are short-run optimal. The short-run players could force player i 's payoff even lower using strategies that are not short-run optimal, but this cannot occur in equilibrium. Note that m^i specifies player i 's strategy m_i^i , which need not be a best response to m_{-i}^i : Player i must play in a certain way to induce the short-run players to attain the minimum in the definition of m^i . To illustrate the definition of the minmax strategies, consider repeated play of the game in Figure 1, with a long-run player 1 facing a sequence of short-run player 2's.

		Player 2		
		L	M	R
Player 1	L	0,0	2,3	-2,-1
	D	-1,2	1,1	-2,-1

Figure 1

Here player 2's minmax strategy against player 1 is $m_2^1 = L$, which holds player 1's payoff to at most $v_1 = 0$. Playing R would hold player 1's payoff even lower, but R is a strictly dominated strategy, and so it will never be used by a short-run player 2. (If player 2 were a long-run player, we would have $v_1 = -2$.) In order to induce player 2 to play L, player 1 must put probability at least 1/2 on D, so any choice of m^1 must have $g_1(m^1) \leq -1/2$, which is less than player 1's minmax value of $0 = v_1$. Thus, in contrast to standard repeated games, equilibrium strategies that hold a long-run player close to his minmax value may need to provide him with incentives to cooperate in his own punishment. We will see that such strategies can indeed be constructed for discount factors near one provided the players' choice of mixed strategies are observable.

In the repeated version of g , we suppose that long-run players maximize the discounted normalized sum of their single-period payoffs, with common discount factor δ . That is, long-run player i 's payoff is $(1-\delta) \sum_{t=0}^{\infty} \delta^t g_i(\sigma(t))$. Short-run players in each period act to maximize that period's payoff. All players, both long- and short-run, can condition their play on all previous choices of mixed strategies.

More formally, the history at time t is $h(t) = \{\sigma(\tau)\}_{0 \leq \tau < t}$, where $\sigma(\tau)$

$= (\sigma_1(\tau), \sigma_2(\tau), \dots, \sigma_n(\tau))$ is the vector of mixed strategies chosen in period τ . A strategy for player 1 is then a sequence of maps $S_1^t = H^t \rightarrow \Sigma_1$, and a strategy for the period t player 2 maps H^t to Σ_2 .

One interpretation of this definition of the information structure is that players implement their mixed strategies using a "randomizing device," which has a fixed, known density on the unit interval. A mixed strategy is then a map from the outcome of the device to actions. The assumption that players observe each others' choice of randomizing probability corresponds to the assumption that both the outcome of the device and the map from outcome to action are observable.

If players observed only the outcome of the randomizing devices and not the maps used, they would not be able to infer the exact randomizing probabilities used by their opponents, so the mixed strategies themselves would not be observable. However, players would still be able to detect any deviation from the action prescribed by a given mixed strategy. As we will explain, this weaker form of observability will be sufficient for our results.

Let $U = \{ v = (v_1, \dots, v_i) \mid \exists \sigma \text{ in graph } (B) \text{ with } g_i(\sigma) = v_i$
for all i from 1 to ℓ

Let $V = \text{convex hull of } U;$

and let $V^* = \{ v \in V \mid v_i > \underline{v}_i \}.$

Note that points in $U, V,$ and V^* specify payoffs for the long-run players only. We will not be concerned with the possible payoffs of the short-run players.

We call payoffs in V^* attainable payoffs for the long-run players. We

begin with the case of a single long-run player.

Proposition 1: For no δ is there an equilibrium in which player i 's payoff is less than v_i .

Proof: To see that player i 's payoff cannot be below v_i in any equilibrium, note first that since all players begin each period with the same information, in equilibrium player i correctly anticipates the mixed strategy $\sigma_{-i}(h(t))$ that his opponents will play in period t . One feasible (not necessarily optimal) strategy for player i is to play in each period the strategy that maximizes that period's expected payoff against $\sigma_{-i}(h(t))$. Since the other players are all short-run, they will only use actions in the range of B , so $\max_{a_i} g_i(a_i, \sigma_{-i}(H(t))) \geq v_i$.

Since players $\ell+1$ through n are short-run, each $\sigma(h(t))$ must lie in the graph of B in any equilibrium. Thus $\max_{a_i} g_i(a_i, \sigma_{-i}(h(+1))) \geq v_i$, and so player i can ensure himself at least v_i each period. Q.E.D.

Proposition 2: If only player 1 is a long-run player, then for any $v_1 \in V^*$ there exists a $\underline{\delta} \in (0, 1)$ such that for all $\delta \in (\underline{\delta}, 1)$, there is a subgame-perfect equilibrium of the infinitely repeated game with discount factor δ in which player 1's discounted normalized payoff is v_1 .

Proof: Fix a $v_1 \in V^*$ and consider the following strategies. Begin in Phase A, where players play a $\sigma \in \text{graph}(B)$ (or a public randomization over such σ 's) that gives player 1 payoff v_1 . Deviations by the short-run players

are ignored. If player one deviates, he is punished by players switching to the punishment strategy m^1 for $T(\delta)$ periods, after which play returns to Phase A; if $T(\delta)$ is large enough, deviations in Phase A are unprofitable. Now m^1_1 need not be a best response against m^1_{-1} , so we must insure that player one does not prefer to deviate during the punishment phase. This is done by specifying that a deviation in this phase restarts the punishment. Since the most that player 1 can obtain in any period of the punishment phase is \underline{v}_1 , he will prefer not to deviate so long as $T(\delta)$ is short enough that player 1's normalized payoff at the start of the punishment phase is at least \underline{v}_1 . Let $\bar{v}_1 = \max_{\sigma \in \text{graph}(B)} g_1(\sigma)$. The two constraints on $T(\delta)$ will be satisfied if:

$$(1) \quad (1-\delta)\bar{v}_1 + \delta(1-\delta^{T(\delta)})g_1(m^1) + \delta^{T(\delta)+1}\underline{v}_1 \leq \underline{v}_1, \text{ or equivalently}$$

$$(1') \quad \delta^{T(\delta)+1} \leq (\underline{v}_1 - \delta g_1(m^1) + (1-\delta)\bar{v}_1) / (\underline{v}_1 - g_1(m^1)), \text{ and}$$

$$(2) \quad (1-\delta^{T(\delta)})g_1(m^1) + \delta^{T(\delta)}\underline{v}_1 > \underline{v}_1, \text{ or equivalently}$$

$$(2') \quad \delta^{T(\delta)} > (\underline{v}_1 - g_1(m^1)) / (\underline{v}_1 - g_1(m^1)).$$

The right-hand side of (2') is less than one, and the right-hand side of (1') converges to 1 as σ goes to 1. Thus for σ sufficiently close to 1, we can find σ and $\tau(\delta)$ that satisfy (1') and (2'). Q.E.D.

Note that the strategies constructed in the proof of Proposition 2 treat all deviations from the prescribed mixed strategies in the same way. Thus to implement the strategies it suffices that players can detect whether

a deviation occurred; it is not necessary that they observe the way in which the opponent deviated.

For this reason, Proposition 2 would obtain if players could only observe the outcomes of their opponents' randomizing devices; it is not necessary that the choices of mixed strategy themselves be observed. A similar observation applies to Proposition 3. below'

In repeated games with three or more players, a full-dimensionality condition is required for all feasible individually rational payoffs to be enforceable when δ is near enough to one. The corresponding condition here is that the dimensionality of V^* equals the number of long-run players.

Proposition 3: Assume that the dimensionality of $V^* = \ell$, the number of long-run players. Then for each v in V^* , there is a $\underline{\delta} \in (0,1)$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a subgame-perfect equilibrium of the infinitely repeated game with discount factor δ in which player i 's normalized payoff is v_i .

Remark: The proof of Proposition 2 follows that of Fudenberg-Maskin's Theorem 2: If a (long-run) player deviates, he is punished long enough to wipe out the gain from deviation. To induce the other (long-run) players to punish him, they are given a "reward" at the end of the punishment phase. One small complication not present in Fudenberg-Maskin is that, as in Proposition 1, the strategies we construct cannot in general allow player i to play a static best response in the periods when he is being minmaxed. This, however, can be arranged with essentially the same strategies as before.

Proof: Choose a σ (or a public randomization over several σ 's) so that $g(\sigma) = v$. Also choose v' in the interior of V^* and an $\epsilon > 0$ so that for all i from 1 to j $(v'_1 + \epsilon, \dots, v'_{i-1} + \epsilon, v'_i, v'_{i+1} + \epsilon, \dots, v'_j + \epsilon)$ is in V^* and $v'_i + \epsilon < v_i$. Let T^i be a joint strategy that yields $v'_i + \epsilon$ to all the long-run players but i , and yields v'_i to i . Let $w_i^j = g_i(m^j)$ be player i 's period payoff when j is being punished with the strategies m^j . For each i , choose an integer N_i so that

$$\bar{v}_i + N_i v_i < (N+1) v'_i$$

where $\bar{v}_i = \max g_i$ is i 's greatest one-period payoff.

Consider the following repeated-game strategy for player i :

(0) Begin in Phase (A).

(A) Play σ_i each period as long as all long-run players played σ last period, or if σ had been played until last period and two or more long-run players failed to play σ last period.

If long-run player j deviates from (A), then

(B_j) Play m_i^j for N_j periods, and then

(C) Play T_i^j thereafter.

If long-run player k deviates in phase (B_j) or (C), then begin phase

(B_j) again with $j = k$. (As in phase A, players ignore simultaneous deviations by two or more long-run players.)

As usual, it suffices to check that in every subgame no player can gain

from deviating once and then conforming. The condition on N_i ensures that for δ close to one, the gain from deviating in Phase A or Phase C is outweighed by Phase B's punishment. If player j conforms in B_j (i.e. when she is being punished) her payoff is at least $q_j = (1-\delta)^{N_j} w_j^j + \delta^{N_j} v_j^j$, which exceeds v_j if δ is close enough to one. If she deviates once and then conforms, she receives at most v_j the period she deviates, and postpones the payoff $q_j > v_j$, which lowers her payoff. If player k deviates in Phase B_j , she is minmaxed for the next N_k periods and Phase-C play will give her v_k' instead of $v_k' + \epsilon$. Thus it is easy to show that such a deviation is unprofitable. (See Fudenberg-Maskin for the missing computations.) Q.E.D.

3. Unobservable Mixed Strategies

We now drop the assumption that players can observe their opponents' mixed strategies, and instead assume they can only observe their opponents' realized actions. In ordinary repeated games, (privately) mixed strategies are needed during punishment phases, because in general a player's minmax value is lower when his opponents use mixed strategies. However, mixed strategies are not required along the equilibrium path, since desired play along the path can be enforced by the threat of future punishments. Fudenberg-Maskin showed that, under the full-dimension condition of Proposition 2, players can be induced to use mixed strategies as punishments by making the continuation payoffs at the end of a punishment phase dependent on the realized actions in that phase in such a way that each action in the support of the mixed strategy yields the same overall payoff.

In contrast, with short-run players some payoffs (in the graph of B) can only be obtained if the long-run players privately randomize, so that mixed

strategies are in general required along the equilibrium path. As a consequence, the set of equilibrium payoffs in the repeated game can be strictly smaller when mixed strategies are not observable. This is illustrated by the game in Figure 2, with one long-run player, Row, and one short-run player, Col.

Let p be the probability that Row plays D. Col's best response is M if $0 \leq p \leq 1/2$, L if $1/2 \leq p \leq 100/101$, and R if $p \geq 100/101$. There are three static equilibria: the pure strategy equilibrium (D,R), a second in which $p = 1/2$ and Col mixes between M and L, and a third in which $p = 100/101$ and Col mixes between L and R. Row's maximum attainable payoff is 3, which occurs when $p = 1/2$ and Col plays L.

	L	M	R
U	4, 0	0, 1	-1, -100
D	2, 2	1, 1	0, 3

Figure 2

If Row's mixed strategy is observable, she can attain this payoff in the infinitely repeated game if δ is near enough to 1. If however Row's mixed strategy is not observable, her highest equilibrium payoff is at most 2 regardless of δ .

To see this, fix a discount factor δ , and let $v^*(\delta)$ be the supremum over all Nash equilibria of ROW's equilibrium payoff. If for some δ $v^*(\delta) =$

$2 + \epsilon' > 2$, then for all sufficiently small $\epsilon > 0$ there is an equilibrium Δ such that player 1's payoff is $v(\Delta) = v^*(\delta) - \epsilon > 2$. It is easy to see that the set of equilibrium payoffs is stationary: Any equilibrium payoff is an equilibrium payoff for any subgame, and conversely. Thus, the highest payoff player 1 can obtain starting from period 2 is also bounded by $v^*(\delta)$. Since $v(\Delta)$ is the weighted average of player 1's first-period payoff and her expected continuation payoff, player 1's first-period payoff must be at least $v^*(\delta) - \epsilon/(1-\delta)$. For ϵ sufficiently small, this implies that player 1's first period payoff must exceed 2.

In order for Row's first-period payoff to be at least 2, Col must play L with positive probability in the first period. As Col will only play L if Row randomizes between U and D, Row must be indifferent between her first period choices, and in particular must be willing to play D. Let $v_D(\Delta)$ be Row's expected payoff (according to the strategies Δ) from period 2 on if she plays D in the first period. Then we must have

$$(3) \quad 2(1-\delta) + \delta v_D(\Delta) = v(\Delta) = v^*(\delta) - \epsilon.$$

Since $v_D(\Delta) \leq v^*(\delta)$, and for all sufficiently small ϵ there is an Δ so that (3) holds, we conclude that $v^*(\delta) \leq 2$.

While Row cannot do as well as if her mixed strategies were observable, she can still gain by using mixed strategies. For δ near enough to one there is an equilibrium which gives Row an normalized payoff of 2, while Row's best payoff when restricted to pure strategies is the static equilibrium yielding 1. To induce Row to mix between U and D, specify that following periods when Col expects mixing and Row plays U, play switches with probability p to (D,R)

for ten periods and then reverts to Row randomizing and Col playing L. The probability p is chosen so that Row is just indifferent between receiving 2 for the next eleven periods, or receiving 4 today and risking punishment with probability p . This construction works quite generally, as shown in the following proposition.

Proposition 4: Consider a game with a single long-lived player, player 1, and let

$$v_1^* = \max_{\sigma \in \text{graph } B} \min_{s_1 \in \text{supp } \sigma_1} g_1(s_1, \sigma_{-1}).$$

Then for any $v_1 \in (v_1, v_1^*)$ there exists a $\delta' < 1$ such that for all $\delta \in (\delta', 1)$, there is a subgame-perfect equilibrium in which player one's normalized payoff is v_1 . For no δ is there an equilibrium where player one's payoff exceeds v_1^* .

Proof: We begin by constructing a "punishment equilibrium" in which player one's normalized payoff is exactly v_1 . If v_1 is player 1's payoff in a static equilibrium this is immediate, so assume all the static equilibria given player 1 more than v_1 . Fix a particular static equilibrium $\hat{\sigma}$, and let $e_1 = g_1(\hat{\sigma}) > v_1$ be player 1's payoff in this equilibrium.

We will construct strategies with two phases, A and B; play begins in phase A. Here the strategies are

(A) Play m^1 . If player 1's realized action was s_1 , then (using a public

randomization) remain in phase (A) with probability $1-p(s_1)$, and switch to phase (B) with probability $p(s_1)$, where $p(s_1)$ satisfies

$$(4) \quad p(s_1) = \frac{(1-\delta)(v_1 - g_1(s_1, m_{-1}^1))}{\delta(e_1 - v_1)} .$$

(If δ is near enough to one, $p(s_1)$ is between 0 and 1.)

In phase (B), the strategies are to play the static equilibrium $\hat{\sigma}$ forever.

The switching probability has been constructed so that player one's normalized profit in phase (A) is v_1 for all actions, including those in the support of m_1^1 . Thus player 1 cannot do better than to use strategy m_1^1 in phase (A), nor can she gain by deviating in phase (B). So these strategies are subgame-perfect if δ is large enough that $p(s_1)$ is less than 1 for all $s_1 \in S_1$.

Now we assume $v_1^* > v_1$ and construct equilibrium strategies yielding v_1^* . Let $\sigma^* = (\sigma_1^*, \sigma_{-1}^*)$ be mixed strategies that attain the maximum in the definition of v_1^* .

(0): Begin in Phase (A).

(A): Play σ^* . If $s_1 \notin \text{support}(\sigma_1^*)$, go to phase B. If $s_1 \in \text{support}(\sigma_1^*)$, then switch to phase B with probability $p^*(s_1)$, and otherwise remain in phase A. Here $p^*(s_1)$ satisfies

$$(5) \quad p^*(s_1) = \frac{(1-\delta)(g_1(s_1, \sigma_{-1}^*) - v_1^*)}{\delta(v_1^* - v_{-1})}$$

Equation (5) is non-negative for all $s_1 \in \text{support}(\sigma_1^*)$, and it is less than one for δ sufficiently close to 1. The probability $p^*(s_1)$ makes player 1's normalized present value of all $s_1 \in \text{support}(\sigma_1^*)$ equal to v_1^* .

(B): Play the "punishment equilibrium" constructed above.

Since player 1's payoff in the punishment equilibrium is $v_{-1} < v_1^*$, these strategies are subgame-perfect for sufficiently large δ .

Equilibrium payoffs between v_{-1} and v_1^* are obtained by using public randomizations between those two values. The argument that player one's payoff cannot exceed v_1^* is exactly as in the example. Q.E.D.

4. No Public Randomizations

The equilibria that we constructed in the proofs of Theorems 1 through 3 relied on our assumption that players can condition their play on the outcome of a publicly observed random variable. While that assumption is not implausible, it is also of interest to know whether the assumption is necessary for our results. For this reason, Proposition 4 below extends Proposition 3 to games without public randomizations. (We have not thought about the possible extension of Propositions 1 and 2 because we think the situation without public randomizations but where private randomization can be verified ex-post is without interest.) The intuition, as explained in Fudenberg-Maskin [1988b], is that public randomizations serve to convexify the set of attainable payoffs, and when δ is near to 1 this convexification can be

achieved by sequences of play which vary over time in the appropriate way. Fudenberg-Maskin [1988b] shows that public randomizations are not necessary for the proof of the perfect Folk Theorem. However, as we have already seen, there are important differences between classic repeated games and repeated games with some short-run players, so the fact that public randomizations are not needed for the folk theorem should not be thought to settle the question here.

Proposition 5: Consider a game with a single long-run player, player 1, where public randomizations are not available. As in Proposition 3, let

$$v_1^* = \max_{\sigma \in \text{graph}(B)} \min_{s_1 \in \text{supp } \sigma} g_1(s_1, \sigma_{-1}),$$

and let σ^* be a strategy that attains this max. Then, for any $v_1 \in (v_1, v_1^*)$ there exists a $\delta' < 1$ such that for all $\delta \in (\delta', 1)$ there is a subgame-perfect equilibrium where player 1's discounted normalized payoff is v_1 .

Remark: Fix a static Nash equilibrium $\hat{\sigma}$ with payoffs \hat{v} . For each v_1 the proof constructs strategies that keep track of the agent's total realized payoff to date t and compares it to the "target" value of $(1-\delta^t) v_1$, which is what the payoff to date would be if the agent received v_1 in every period. If v_1 exceeds the payoff in a static equilibrium, then play initially follows the (possibly mixed) strategy σ^* , and whenever the realized total is sufficiently greater than the target value, the agent is "punished" by reversion to the static equilibrium. If the target is less than the static equilibrium, then

play starts out at the (possibly mixed) strategy m^1 , with intermittent "rewards" of the static equilibrium whenever the realized payoff drops too low.

Proof: Let \hat{v}_1 be player 1's payoff in a static equilibrium.

(A) It is trivial to obtain \hat{v}_1 as an equilibrium payoff of the repeated game.

(B) To attain any payoff v_1 between \hat{v}_1 and v_1^* we proceed as follows. Renormalize the payoffs so that $\hat{v}_1 = 0$, and take δ large enough that $(1-\delta)\hat{v}_1 < v_1^*$. Define $J_0 = 0$ and $\sigma^*(0) = \sigma^*$, and for each time $t > 0$ define the strategies $\sigma^*(h_t)$ and an index J_t recursively as follows. First we define J_t :

$$J_t = J_{t-1} + (1-\delta)\delta^{(t-1)} g_1(s_1(t-1), \sigma_{-1}^*(h_{t-1})),$$

where $s_1(t-1)$ is player 1's action in period $t-1$, as opposed to his choice of mixed strategy. (Note that J_t is common knowledge at the start of time t .) Define $R_t = (1-\delta^t)v_1$. If player 1's payoff were v_1 each period, then J_t would equal R_t . The equilibrium strategies will "punish" the agent whenever J_t exceeds R_t by too large a margin. More precisely, we define

$$\sigma^*(h_t) = \begin{cases} \hat{\sigma} & \text{if } J_t \geq R_{t+1} \text{ and } J_\tau \geq R_\tau \text{ for all } \tau \leq t \\ \sigma^* & \text{if } J_t < R_{t+1} \text{ and } J_\tau \geq R_\tau \text{ for all } \tau \leq t \\ \hat{\sigma} & \text{if } J_\tau < R_\tau \text{ for any } \tau \leq t \end{cases}$$

Note that since J_t is a discounted sum, for each infinite history h_∞ , J_t converges to a limit J_∞ . Moreover, as long as the other players use strategy σ_{-1}^* , player 1's payoff to any strategy is simply the expected value of J_∞ , and his expected payoff in any subgame starting at time t is $\delta^{-t}(J_\infty - J_t)$ (in time- t units).

We will now argue that

(i) if player 1 uses strategy σ_1^* , then $J_t \geq R_t$ for all times t and histories h_t , which implies that $J_\infty \geq v_1$

(ii) that regardless of how player 1 plays, $J_\infty \leq v_1$, so player 1's payoff in the subgame starting at time t is bounded by $\delta^{-t}(v_1 - J_t)$ for all histories h_∞ , and

(iii) that in any subgame where at some $\tau \leq t$, $J_\tau < R_\tau$, it is a best response for all players to follow the prescribed strategy of always playing the static equilibrium $\hat{\sigma}$, and so $J_\infty < v_1$.

Conditions (i) and (ii) imply that it is a best response for player 1 to play σ^* in every subgame where J_τ has never dropped below R_τ , and that player 1's equilibrium payoff is v_1 . Condition (iii), whose proof is immediate, says that following σ^* is also a Nash equilibrium in subgames where J_τ has dropped below R_τ , so that σ^* is a subgame-perfect equilibrium. (The condition that the short-run players not wish to deviate is incorporated in the construction of σ^* .)

Proof of (i): We must show that if player 1 follows σ_1^* then for all t , $J_t \geq (1 - \delta^t)v_1 = R_t$. Since $J_0 = 0$, this is true for $t = 0$. Assume it is true for $t = \tau$. At period τ , either (a) $J_\tau \geq R_{\tau+1}$ or (b) $J_\tau < R_{\tau+1}$. In case (a),

$\Delta^*(h_r) = \hat{\sigma}$. Since $g_1(s_1, \hat{\sigma}_{-1}) = 0$ for every pure strategy s_1 in the support of $\hat{\sigma}_1$, we have $J_{r+1} = J_r \geq R_{r+1}$ from case (a). In case (b), $\Delta^*(h_r) = \hat{\sigma}^*$, so $\min \{g_1(s_1(\tau), \Delta_{-1}^*(h_r)) \mid s_1(\tau) \in \text{support}(\Delta_1^*(h_r))\} = v_1^*$ and

$$J_{r+1} \geq J_r + (1-\delta)\delta^r v_1^* \geq (1-\delta^r)v_1 + (1-\delta)\delta^r v_1 = (1-\delta^{r+1})v_1 = R_{r+1},$$

where the second inequality comes from the inductive hypothesis and the fact $v_1^* \geq v_1$. Thus $J_t \geq (1-\delta^t)v_1$ for all t , and so if player 1 follows $\hat{\Delta}_1^*$ then $J_\infty \geq v_1$.

Proof of (ii): Next we claim that regardless of how player 1 plays, $J_\infty \leq v_1$.

We will show inductively that $J_t \leq v_1$ for all t . This is clearly true for $t = 0$; assume it is true for an arbitrary t . If $J_t \geq R_{t+1}$ then since $\Delta_{-1}^*(h_t) = \hat{\sigma}_{-1}$, and the most player 1 can obtain against $\hat{\sigma}_{-1}$ is zero,

$J_{t+1} \leq J_t \leq v_1$ by inductive hypothesis. If $J_t < R_{t+1}$, then

$$\begin{aligned} J_{t+1} &\leq (1-\delta^{t+1})v_1 + \delta^{t+1}(1-\delta)\bar{v}_1 \leq \\ &(1-\delta^{t+1})v_1 + \delta^{t+1}v_1 = v_1 \end{aligned}$$

from our bound on δ and the fact that $v_1 \leq v_1^*$.

(C): Next we show how to construct equilibria (for large enough δ) that yield payoffs v_1 between v_{-1} and the static equilibrium payoff of zero. Pick a $v_1 \in (v_{-1}, 0)$, and choose $\delta > 1/2$ and large enough that $(1-\delta)\min_{\sigma} g_1(\sigma) > v_1$. Then set $J_0 = 0$, and $\Delta^*(0) = m^1$. Now define $J_t(h_t)$ and $\Delta(h_t)$ recursively as follows for $t = 1, 2, \dots$. Set $J_t = J_{t-1} + (1-\delta)\delta^{t-1}g_1(s_1(t-1), \Delta_{-1}^*(h_{t-1}))$, and

set

$$\Delta^*(h_t) = \begin{cases} \hat{\sigma} & \text{if } J_t < R_t \\ m^1 & \text{if } J_t \geq R_t \end{cases}$$

Proceeding as above, we claim that

- (i) if player 1 uses strategy Δ_1^* , then $J_t \geq v_1$ for all times t and histories h_t , which implies that $J_\infty \geq v_1$,
- (ii) that regardless of how player 1 plays, $J_\infty \leq v_1$, so player 1's payoff in the subgame starting at time t is no greater than $\delta^{-t}(v_1 - J_t)$, and
- (iii) that in subgames where $J_t < R_t$ it is a best response for player one to play $\Delta^*(h_t) = \hat{\sigma}_1$.

Proof of (i): We must show that if player 1 follows Δ_1^* then for all t , $J_t \geq v_1$. Since $J_0 = 0$, this is true for $t=0$. Assume it is true for $t=\tau$. At period τ , either (a) $J_\tau \geq R_\tau$ or (b) $J_\tau < R_\tau$.

In case (a),

$$\begin{aligned} J_{\tau+1} &\geq J_\tau + \delta^\tau (1-\delta)\min(g_1) \geq J_\tau + \delta^\tau v_1 \quad (\text{from our bound on } \delta) \\ &\geq R_\tau + \delta^\tau v_1 \quad (\text{from case (a)}) \\ &= v_1. \end{aligned}$$

In case (b), $\Delta^*(h_\tau) = \hat{\sigma}$, so $g_1(s_1(\tau), \Delta_{-1}^*(h_\tau)) = 0$ for all $s_1(\tau) \in \text{support } \Delta_1^*(h_\tau)$,

and $J_{\tau+1} - J_\tau \geq v_1$.

Proof of (ii): We claim that for all strategies of player 1 and all times t

and histories h_t , $J_t \leq (1-\delta^{t-1})v_1 = R_{t-1}$. Since $J_0 = 0$, $\Delta^*(h_0) = m^1$, and so $J_1 \leq (1-\delta)v_1 < 0 = R_0$. Assume it is true for $t = \tau$. Then at period τ , either (a) $J_\tau < R_\tau$ or (b) $J_\tau \geq R_\tau$. In case (a), $\Delta^*(t) = \hat{\sigma}$, so $J_{\tau+1} \leq J_\tau < R_{\tau-1}$. In

case (b), $\Delta^*(h_\tau) = m^1$, so $\max_{s_1} g_1(s_1, \Delta_{-1}^*(h_\tau)) = v_1$, and

$$J_{\tau+1} < J_\tau + (1-\delta)\delta^\tau v_1 < (1-\delta^{\tau-1})v_1 + (1-\delta)\delta^\tau v_1 \text{ (by the inductive hypothesis)}$$

$$\leq (1-\delta^{\tau+1} + \delta^\tau)v_1 < (1-\delta^\tau)v_1 = R_\tau \text{ as } \delta > 1/2.$$

Thus $J_t < (1-\delta^t)v_1$ for all t , and so regardless of how player 1 plays, $J_\infty \leq v_1$.

(iii) Conditions (i) and (ii) show that in any subgame with $J_t \geq R_t$ player 1 can attain the upper bound of v_1 by following Δ_1^* . Now we consider subgames with $J_t < R_t$. If $J_t \leq v_1$, then regardless of how player 1 plays, we will have $J_\tau \leq R_\tau$ for all $\tau \geq t$, so player 1's opponents will play $\hat{\sigma}_1$ for the remainder of the game. Here it is clearly a best response for player 1 to play $\hat{\sigma}_1 = \Delta_1^*$. If $J_t > v_1$, then by playing $\hat{\sigma}_1$ player 1 can ensure that $J_\tau \geq R_\tau$ at some $\tau > t$, which ensures that player 1 attains a continuation payoff of $(v_1 - J_t)$ in the subgame starting at t . If player 1 instead chooses a strategy which assigns positive probability to the event that $J_\tau < R_\tau$ for all $\tau > t$, he can only lower his payoff: The payoff for histories with $J_\infty < R_\infty$ is less than v_1 , and the payoff for the histories with $J_\infty \geq R_\infty$ is bounded above by $v_1 - J_t$. Q.E.D.

Proposition 4 shows how to attain any payoffs between v_1 and v_1^* by means of "target strategies." From Proposition 3 we know that such strategies cannot be used to attain higher payoffs. We think that it is interesting to

note where an attempted proof would fail.

In part (A), we proved that if player 1 followed σ^* then for every sequence of realizations player 1's payoff is at least v_1 . Imagine that we try to attain a payoff $v_1 > v_1^*$ by setting the target $R_t = (1-\delta^t)v_1$. Then in the "reward" phases where σ^* is played, it might be that player 1's realized payoff is less than v_1 . (Recall that by definition it cannot be lower than v_1^*). After a sufficiently long sequence of these outcomes, player 1's realized payoff J_t would be so much lower than v_1 that even receiving the best possible payoff at every future date would not bring his discounted normalized payoff up to the target.

This problem of going so far below the target that a return is impossible does not arise with the criterion of time-average payoffs, since the outcomes in any finite number of periods are then irrelevant. For this reason we can attain payoffs above v_1^* under time averaging, as we show in the next section.

(After this paper was written, Fudenberg-Levine [1989b] provided another way to obtain this result. They show that any strategy profile that is an ϵ_n -equilibrium for discount factor δ_n for a sequence $(\epsilon_n, \delta_n) \rightarrow (0,1)$, is a time-average equilibrium as well. With the strategies used above, the probability that player 1's discounted payoff is less than the target is small for δ close to 1, since the mean payoff in the reward phase exceeds the target. This is why the target strategies of this sector are ϵ -equilibria of the appropriate kind.)

5. Time Averaging

The reason that player 1's payoff is bounded by what he obtains when she plays her least favorite strategy in the support of σ^* is that every time she

plays a different action she must be "punished" in a way that makes all of the actions in σ^* equally attractive. A similar need for "punishments" along the equilibrium path occurs in repeated partnership games, where two players make an effort decision that is not observed by the other, and the link between effort and output is stochastic. Since shirking by either player increases the probability of low output, low output must provoke punishment, even though low output can occur when neither player shirks. This is why the best equilibrium outcome is bounded away from efficiency when the payoff criterion is the discounted normalized value. (Radner-Myerson-Maskin [1986], Fudenberg-Maskin [1987a]). However, Radner [1986] has shown that efficient payoffs can be attained in partnerships with time averaging. His proof constructed strategies so that (1) if players never cheat, punishment occurs only finitely often, and thus is negligible, and (2) an infinite number of deviations is very likely to trigger a substantial punishment. Since no finite number of deviations can increase the time-average payoff, in equilibrium no one cheats yet the punishment costs are negligible.

Since the inefficiencies in repeated partnerships and games with short-run players both stem from the need for punishments along the equilibrium path, it is not surprising that the inefficiencies in our model also disappear when players are completely patient. We prove this with a variant of the "target strategies" we used in Section 4. These strategies differ from Radner's in that even if player 1 plays the equilibrium strategy, she will be punished infinitely often with probability one. However, along the equilibrium path the frequency of punishment converges to zero, so that as in Radner the punishment imposes zero cost.

Proposition 5: Imagine that there is a single long-run player, player 1, who evaluates payoff streams with the criterion

$$\liminf_{T \rightarrow \infty} E (1/T) \sum_{t=0}^{T-1} g_1(s(t)).$$

Then for all $v_1 \in V^*$ there is a subgame-perfect equilibrium with payoffs v_1 .

Remark: The proof is based on a strong law of large numbers for martingales with independent increments,^{2/} which we extend to cover the difference between a supermartingale and its lowest value to date. The relevant limit theory is developed in the Appendix.

Proof: As in Proposition 4, we use different strategies for payoffs above and below some fixed static equilibrium $\hat{\sigma}$. Imagine that v_1 exceeds player 1's payoff in this equilibrium, and normalize $v_1 = 0$. Let $\check{\sigma}$ be the (possibly mixed) strategy in graph (B) that maximizes player one's expected payoff, and define $g_1(\check{\sigma}) = \check{v}_1$.

$$\text{Define } J_0 = 0 \text{ and } J_T = \sum_{t=0}^{T-1} g_1(s_1(t), s_{-1}(t)). \text{ (This differs from the}$$

definition in Section 4, where we used player 1's realized action and the mixed strategy of her opponents in defining J_t). Note that player 1's objective function is $\liminf E(1/T)J_T$. Set

$$\Delta^*(h_t) = \begin{cases} \check{\sigma} & \text{if } J_t \geq 0 \\ \hat{\sigma} & \text{if } J_t < 0 \end{cases}$$

We claim that (i) no matter how player 1 plays, her payoff is bounded by v_1 , and (ii) that by following s_1^* player 1 can attain payoff v_1 almost surely (and hence in expectation.)

To prove this, let $\Delta(h_t)$ be an arbitrary strategy for player 1, and fix the associated probability distribution over infinite-horizon histories. For each history, let $R_t(h_t) = \{\tau \leq t-1 \mid \Delta^*(h_\tau) = \hat{\sigma}\}$ be the "reward" periods, and let $P_t(h_t) = \{\tau \leq t-1 \mid \Delta^*(h_\tau) = \sigma\}$ be the "punishment" ones.

Then let $M_t(h_t) = \sum_{\tau \in R_t} g_1(x(\tau))$ be the sum of player 1's payoffs in the good periods, and set

$$N_t(h_t) = \sum_{\tau \in P_t} g_1(x(\tau)).$$

Note that the reward and punishment sets and the associated scores are defined path wise, i.e. they depend on the history h_t ; henceforth, though, we will omit the history h_t from the notation. Finally define $\hat{M}_t = \max_{\tau \leq t} M_\tau$, $\hat{N}_t = \min_{\tau \leq t} N_\tau$, and $\hat{v}_1 = \min g_1(\sigma)$

We claim that for all t ,

$$(5) \quad \hat{v}_1 + (M_t - \hat{M}_t) \leq J_t \leq \hat{v}_1 + (N_t - \hat{N}_t).$$

This is clearly true for $t = 0$. Assume (5) holds for all $\tau \leq t$. At the start of period t , either (a) $J_t > 0$ or (b) $J_t \leq 0$. In case (a), $J_{t+1} \geq \hat{v}_1 + J_t \geq \hat{v}_1 \geq \hat{v}_1 + (M_t - \hat{M}_t)$. Also, since $\Delta^*(h_t) = \hat{\sigma}$,

$J_{t+1} = J_t + N_{t+1} - N_t \leq \bar{v}_1 + N_{t+1} - \hat{N}_t \leq \bar{v}_1 + (N_{t+1} - \hat{N}_{t+1})$, so that (5) is satisfied.

In case (b), $J_{t+1} \leq \bar{v}_1 + J_t \leq \bar{v}_1 \leq \bar{v}_1 + (N_{t+1} - \hat{N}_{t+1})$,

and $J_{t+1} = J_t + M_{t+1} - M_t \geq \underline{v}_1 + M_{t+1} - \hat{M}_t \geq \underline{v}_1 + (M_{t+1} - \hat{M}_{t+1})$,

so once again (5) is satisfied.

Lemmas 3 in the Appendix shows that $(N_t - \hat{N}_t)/t$ converges to zero almost surely. Since the per-period payoffs are uniformly bounded, this implies that $\limsup (1/T)J_T \leq 0$ almost surely, and since the per-period payoffs are uniformly bounded, $\limsup 1/T E(J_T) \leq 0$ as well. Lemma 4 shows that if player 1 plays so that M_t is a submartingale, then the $(M_t - \hat{M}_t)$ converges to zero as well. Since this is true when player 1 follows σ_1^* , the result follows.

Q.E.D.

We can show that with our strategies, player 1 is punished infinitely often ($J_t > 0$) with probability one. This contrasts with Radner's construction of efficient equilibria for symmetric time-average partnership games, where the probability of infinite punishment is zero. It seems likely that our "target-strategy" approach provides another way of constructing efficient equilibria for those games; it would be interesting to know whether this could be extended to asymmetric partnerships. Our approach has the benefit of making more clear why the construction cannot be extended to the discounting case.

6. Several Long Run Players with Unobservable Mixed Strategies

The case of several long-run players is more complex, and we have not solved it. As before, we can construct mixed-strategy equilibria in which the long-run players do better than in any pure strategy equilibrium, and

once again they cannot do as well as if their mixed strategies were directly observable. However, we have not obtained a general characterization of the enforceable payoffs.


Figure 3 presents a 3-player version of the game in Figure 2. Row's and Col's choices and payoffs are exactly as before. The third player, DUMMY, who is a long-run player, receives 3 if Col plays L and receives 0 otherwise. The feasible payoffs for Row and DUMMY are depicted in Figure 2. Consider the feasible point at which $p = 1/2$ and Col plays L. Here Row and Dummy both receive 3. The argument of Section 3 shows that Row's best equilibrium payoff is not 3 but 2, which is the minimum of payoff over the actions in the support of her mixed strategy. Dummy is not mixing, so Dummy's minimum payoff over the support of her strategy is 3. (Indeed this is the minimum over the support of the product of the two strategies.) Thus one might hope that, by analogy to the proof of Proposition 3, we could show that the payoffs (2,3) were enforceable. But these payoffs are not even feasible! The highest Dummy's payoff can be when Row's payoff is 2 is $2 \frac{196}{205}$. (See Figure 3, which depicts the feasible set.) The problem is that an equilibrium in which Row usually randomizes must sometimes have Col play M or R to "tax away" Row's "excess gains" from playing U instead of D, and this "tax" imposes a cost on Dummy.

Fudenberg-Levine [1989a] show that the limit set of equilibria with several long-run players is exactly the intersection of the feasible, individually rational payoffs with the constraints $v_i \leq v_i^*$. As the techniques in that paper are quite different from those used here, we have chosen to present them separately.

	L	M	R
U	4 , 3 , 0	0 , 0 , 1	-1 , 0 , -100
D	2 , 3 , 2	1 , 0 , 1	0 , 0 , 3

Player 2 is a "dummy", player 3 chooses COLs.

Figure 3



Feasible set when three plays a SR best response.

Figure 4

APPENDIX

In this appendix we consider discrete-parameter martingales (x_n, F_n) $n = 0, 1, \dots$, where (F_n) is a filtration on an underlying probability space. We assume that $x_0 = 0$.

Lemma A.1. Let (x_n, F_n) be a martingale sequence with bounded increments. (That is, for some number B , $|x_n - x_{n-1}| \leq B$, almost surely.) Then $\lim_{n \rightarrow \infty} x_n/n = 0$ almost surely. A proof of this lemma can be found in Hall and Heyde (1980, page 36ff).

We also use the following standard adaptation of this strong law:

Lemma A.2 For (x_n, F_n) as above, let $X_n = \min_{1 \leq i \leq n} x_i$. Then $\lim_{n \rightarrow \infty} X_n/n = 0$ almost surely.

Proof: Since $x_0 = 0$, $X_n \leq 0$ for all n . Fix a sample of the stochastic process. Since $X_n/n \leq 0$, we only have to show that $\liminf X_n/n = 0$. Suppose, instead, that n_i is a subsequence along which the limit is less than 0. For each n_i , there is $m_i \leq n_i$ with $x_{m_i} = X_{n_i}$, and thus $0 > X_{n_i}/n_i = x_{m_i}/n_i \geq x_{m_i}/m_i$. Hence, along the subsequence (m_i) , x_{m_i}/m_i violates the strong law, which can happen only on a null set.

Lemma A.3 Let (x_n, F_n) be a supermartingale with bounded increments and with $x_0 = 0$. Let (X_n) be defined from (x_n) as in lemma 2. Then $\lim_{n \rightarrow \infty} (x_n - X_n)/n = 0$ almost surely.

Proof: Since $x_n \geq X_n$, we only need to show that the limsup of the sequence is nonpositive. For $n = 1, \dots$, let $\xi_n = x_n - x_{n-1}$, and let $\zeta_n = \xi_n - E(\xi_n | F_{n-1})$. Note that $\zeta_n \geq \xi_n$. Let $y_n = \sum_{i=1}^n \zeta_i$, and let $Y_n = \inf(y_i; i = 1, \dots, n)$. Then, immediately, (y_n, F_n) is a martingale sequence with bounded increments, and lemmas 1 and 2 tell us that $\lim y_n/n = \lim Y_n/n = 0$, and thus $\lim(y_n - Y_n)/n = 0$. We are done, therefore, once we show that $x_n - X_n \leq y_n - Y_n$ point wise. But this is easily done by induction. It is clearly true for $n = 0$ by convention. Assume it holds for $n - 1$; then since $\xi_n \leq \zeta_n$,

$$x_{n-1} - X_{n-1} + \xi_n \leq y_{n-1} - Y_{n-1} + \zeta_n, \text{ or}$$

$$x_n - X_{n-1} \leq y_n - Y_{n-1}.$$

If $X_n = X_{n-1}$, then since $Y_{n-1} \geq Y_n$, we are done. While if $X_n < X_{n-1}$, then $X_n = x_n$, and $x_n - X_n \leq y_n - Y_n$. Q.E.D.

A symmetrical argument completes the proof, and we obtain:

Lemma A.4 : Let (x_n, F_n) be a submartingale with bounded increments and $x_n \equiv 0$. Let $X_n = \max(x_i | i = 1, \dots, n)$. Then $\lim_{n \rightarrow \infty} (X_n - x_n)/n = 0$ almost surely.

FOOTNOTES

1. The required discount factor can depend on the payoffs to be attained.
2. We thank Ian Johnstone for pointing us to this result.

Partnership Game with Discounting and With Uniformly Inefficient
Equilibria." Review of Economic Studies 53, 59-70.

Selten, R. [1977] "The Chain-Store Paradox.", Theory and Decision, 9, 127-159.

Shapiro, C. [1982] "Consumer Information, Product Quality, and Seller
Reputation," Bell Journal of Economics, 13, 20-35.

Simon, H. [1951] "A Formal Theory of the Employment Relationship,"
Econometrica, 19, 293-305.

REFERENCES

- Abreu, D. [1986] "On the Theory of Infinitely Repeated Games with Discounting", mimeo.
- Dybvig, P. and C. Spatt [1980] "Does it Pay to Maintain a Reputation?", mimeo.
- Fudenberg, D. and D. Levine [1983] "Subgame-Perfect Equilibria of Finite and Infinite Horizon Games", Journal of Economic Theory, 31, 227-256.
- _____ [1989a] "The Characterization of the Public Equilibria of Repeated Games," in preparation.
- _____ [1989b] "An Approximate Folk Theorem for Repeated Games with Private Information," in preparation.
- Fudenberg, D., and E. Maskin [1986] "The Folk Theorem in Repeated Games with Discounting or With Incomplete Information," Econometrica, 54, 533-556.
- _____ [1987] "Discounted Repeated Games with One-Sided Moral Hazard",
Mimeo, Harvard University.
- _____ [1988a] "Nash and Perfect Equilibria of Discounted Repeated Games,"
Mimeo.
- _____ [1988b] "On the Dispensability of Public Randomizations in Discounted Repeated Games," mimeo.
- Fudenberg, D., D. Levine, and E. Maskin [1989] "The Folk Theorem in Repeated Games with Imperfect Public Information," mimeo.
- Hall, P. and C.C. Heyde, [1980] Martingale Theory and Its Applications,
Academic Press, New York.
- Kreps, D. [1984] "Corporate Culture." Mimeo, Stanford Business School.
- Radner, R. [1986] "Repeated Partnership Games with Imperfect Monitoring and no Discounting." Review of Economic Studies 53, 43-58.
- Radner, R., R. Myerson, and E. Maskin [1986] "An Example of a Repeated