

THE FOLK THEOREM WITH IMPERFECT PUBLIC INFORMATION¹

BY DREW FUDENBERG, DAVID LEVINE, AND ERIC MASKIN

We study repeated games in which players observe a public outcome that imperfectly signals the actions played. We provide conditions guaranteeing that any feasible, individually rational payoff vector of the stage game can arise as a perfect equilibrium of the repeated game with sufficiently little discounting. The central condition requires that there exist action profiles with the property that, for any two players, no two deviations—one by each player—give rise to the same probability distribution over public outcomes. The results apply to principal-agent, partnership, oligopoly, and mechanism-design models, and to one-shot games with transferable utilities.

KEYWORDS: Repeated games, Folk Theorem, imperfect observability, moral hazard, mechanism design.

1. INTRODUCTION

THE HALLMARK OF REPEATED GAMES is the enormous variety of behavior that is consistent with equilibrium. Indeed, the Folk Theorem establishes that any feasible payoff vector Pareto-dominating the minimax point of the stage game can arise in a subgame-perfect equilibrium of the infinitely repeated game for discount factors sufficiently near one.²

An important hypothesis of the standard Folk Theorem is that the players can observe one another's actions in each repetition, so that deviations from equilibrium strategies are detectable. In contrast, this paper considers games in which players observe only a public outcome that is a stochastic function of the actions played. Thus these are games of moral hazard. The major task of the paper is to provide conditions sufficient for the Folk Theorem to extend to such games. The most important hypotheses concern the way the probability distribution over public outcomes depends on the players' actions. These hypotheses also enable us to draw conclusions about one-shot games in which utility transfers are available.

A vector of actions (an action profile) satisfies pairwise identifiability if, for every pair of players, the distributions over public outcomes induced by one

¹ This paper subsumes two earlier manuscripts, "Discounted Repeated Games with One-Sided Moral Hazard" and "The Folk Theorem with Unobservable Actions." We would like to thank Svend Albaek, Andrew Atkeson, Patrick Kehoe, Andreu Mas-Colell, Abraham Neyman, and Neil Wallace for helpful comments. Martin Hellwig, David Kreps, and two referees offered valuable suggestions, particularly for improving the exposition. Steven Tadelis provided fine research assistance. NSF Grants 87-08616, 86-09697, 85-20952, 89-111121, and a grant from the UCLA Academic Senate provided research support.

² Throughout this paper we shall confine attention to versions of the Folk Theorem in which subgame-perfect equilibrium or one of its refinements is the solution concept (see, however, the remarks in Section 10). For more than two players, such versions require that the set of feasible payoffs satisfy a nonempty interior condition (see Fudenberg and Maskin (1986a) and Abreu, Dutta, and Smith (1994)).

player's unilateral deviations (pure or mixed) from this vector are distinct from those induced by the other's deviations (Definition 5.5), so that the two players' deviations can be distinguished statistically. If this condition holds for all pure-action profiles (profiles in which players do not randomize), a "Nash-threat" version of the Folk Theorem obtains: any payoff vector Pareto-dominating a Nash equilibrium of the stage game can be sustained in an equilibrium of the repeated game for discount factors near enough to 1 (Theorem 6.1).

In many games of interest, however, the hypothesis that all pure-action profiles are pairwise identifiable fails. For such games, we strengthen pairwise identifiability to require, in addition, that any two deviations by the *same* player induce different distributions (individual full rank; see Definition 5.1). If a game has at least one (mixed-action) profile satisfying the conjunction of pairwise identifiability and individual full rank—i.e., pairwise full rank (Definition 5.4)—then again the Nash-threat Folk Theorem applies. Generic games possess such a profile provided that the number of possible public outcomes is no less than the total number of elements in the action sets of any two players.

To obtain the conventional "minimax-threat" Folk Theorem (Theorem 6.2) requires more stringent conditions. Specifically, besides the hypotheses for the Nash-threat theorem, it suffices to assume that all pure-action profiles satisfy individual full rank.

One-shot games in which transfers of utility can be made contingent on the outcome are conceptually and technically much like repeated games. We derive a one-shot counterpart to the Nash-threat Folk Theorem in Section 6 (Theorem 6.3).

These results apply to several models of economic interest. In particular, we show in Section 6 that the minimax-threat theorem applies to partnership games (in which each player supplies an unobservable input to a production process and output depends stochastically on the inputs) and to Cournot oligopoly (in which firms sell output unobservably and the market price is a random function of total supply). We argue, furthermore, that the Nash-threat result applies to games with a product structure. These are games in which a public outcome is a vector with a different component for each player, and such that a player's action affects only his own component (Definition 7.1). Principal-agent models (Section 9) and games of pure adverse selection (Section 8) are prominent examples of games with a product structure.

Our paper contributes to a considerable literature on repeated games with unobservable actions. Rubinstein (1979), Radner (1981, 1985), and Rubinstein and Yaari (1983) showed that efficient payoffs can be attained or approximated by equilibria of repeated principal-agent games.³ The application of our Nash-threat proposition generalizes these results. For games in which no player's action is observable, the previous literature drew conflicting conclusions. On the

³ These are games in which only one player's actions are unobservable, i.e., in which one-sided moral hazard obtains.

one hand, Radner (1986) (in partnership games) and Lehrer (1989) (more generally) showed that efficient equilibria exist if players do not discount at all. On the other hand, Radner, Myerson, and Maskin (1986) argued by example that, with discounting, the set of equilibrium payoffs may be bounded away from the Pareto frontier even as the discount factor tends to 1. Moreover, the equilibria analyzed by Green and Porter (1984) and Abreu, Pearce, and Stacchetti (1986) in the repeated Cournot oligopoly game described above, exhibit the same sort of inefficiency.⁴ Our work makes clear that the Radner-Myerson-Maskin counterexample relies on there being relatively few possible public outcomes compared to the number of possible actions, so that the genericity result mentioned before does not apply (see Sections 3 and 6), and that equilibria of the Green-Porter and Abreu-Pearce-Stacchetti variety are necessarily inefficient only because they are symmetric: if (even slightly) asymmetric equilibria are admitted, the Folk Theorem is restored (Section 6).

Finally, Green (1987) examined a model of repeated insurance contracts against idiosyncratic shocks that are unobservable to insurers. We explain in Section 2 how we can reinterpret our basic moral hazard framework to incorporate this model of adverse selection and, in Section 8, why it admits (nearly) efficient equilibrium payoffs.

Our methods are based on dynamic programming and geometry. Several papers before ours have exploited the link between dynamic programming and repeated games, including Abreu (1988), Abreu, Pearce, and Stacchetti (1986, 1990), Fudenberg and Maskin (1986b), Matsushima (1989), Mertens (1987), Radner, Myerson, and Maskin (1986), and Spear and Srivastava (1987). Our methods are closest to those of Abreu, Pearce, and Stacchetti and Matsushima. We discuss the connections in Sections 4–6.

The papers by Radner, Rubinstein, and Rubinstein and Yaari used a technique based on “review strategies.” This technique divides time into a sequence of review phases, at the end of each of which a test is conducted to determine the likelihood that some player deviated from equilibrium a nonnegligible fraction of times. We compare this approach with ours in Section 6.

In Section 2, we lay out the basic model. Section 3 presents two examples that illustrate the methods and the principal concepts of this paper in some detail. We recommend this section to the reader wishing an overview of our ideas. The formal analysis begins in Section 4.

2. THE MODEL

In the *stage game*, players move simultaneously, and player $i = 1, \dots, n$ chooses an *action* a_i from a finite set A_i . Let m_i be the number of elements in A_i , i.e., $m_i = |A_i|$. We call vector $a \in A \equiv \times_{i=1}^n A_i$ a *profile* of actions. Profile a

⁴ Unlike us, these authors were not concerned with characterizing the entire set of equilibrium payoffs. The specific equilibria that they studied, however, were inefficient.

induces a probability distribution over the possible *public outcomes* $y \in Y$, where Y is a finite set with $|Y| = m$ elements. Each player i 's realized payoff $r_i(a_i, y)$ depends on his action a_i and the public outcome⁵ y , but not on a_j for $j \neq i$.⁶ Let $\pi(y|a)$ be the probability of y given a . Player i 's expected payoff from action profile a is

$$g_i(a) = \sum_{y \in Y} \pi(y|a) r_i(a_i, y).$$

A *mixed action* α_i for each player i is a randomization over A_i . Let $\alpha_i(a_i)$ be the probability that α_i assigns to a_i , and adopt the shorthand $r_i(\alpha_i, y) = \sum_{a_i \in A_i} \alpha_i(a_i) r_i(a_i, y)$. For each profile $\alpha = (\alpha_1, \dots, \alpha_n)$ of mixed actions, we can compute the induced probabilities of outcomes,

$$\pi(y|\alpha) = \sum_{a \in A} \pi(y|a) \alpha(a),$$

and the expected payoffs

$$g_i(\alpha) = \sum_{y \in Y} \sum_{a \in A} \pi(y|a) \alpha(a) r_i(a_i, y),$$

where $\alpha(a) = \alpha_1(a_1) \alpha_2(a_2) \cdots \alpha_n(a_n)$. We denote the profile in which player i plays a_i and all other players follow profile α by (a_i, α_{-i}) ; $\pi(y|a_i, \alpha_{-i})$ and $g_i(a_i, \alpha_{-i})$ are defined analogously.

Our formulation encompasses the following models considered in the repeated games literature.

First, suppose that the public outcome is simply the profile of actions. In this case, we have the conventional framework of perfect observability *ex post*, as in Fudenberg and Maskin (1986a).

Second, with a judicious interpretation of the variables, we can obtain a finite version of the standard principal-agent model. In this model, one player, the principal, moves first and selects one of a finite number of monetary transfer rules, where the transfer to the other player, the agent, depends on output. Next, the agent selects an effort level (unobservable to the principal), which stochastically determines output. This is, therefore, a model of one-sided moral hazard. We can identify the principal's action with a transfer rule. Because we have modeled the players as moving simultaneously—whereas actually the agent moves after the principal—we must interpret the agent's action as a *contingent* effort level, that is, a function dependent on the principal's move. A public outcome is then a realized output level together with the principal's

⁵ Lehrer (1988), Fudenberg and Levine (1991), and Kandori and Matsushima (1992) consider repeated games in which each player receives a *private* signal of the others' actions.

⁶ Actually, a good many of our results extend to the case in which r_i depends directly on all players' actions (see Section 6).

monetary transfer rule. (By formally including the principal's action as part of the outcome, we make it observable to the agent.)⁷

Third, we can think of the players as a partnership in which each player supplies an unobservable input a_i to a production process, as in Radner (1986). Output y depends stochastically on the quantities of all inputs and is divided among the partners in a prearranged way. This is a model of multi-sided moral hazard.

Fourth, we can imagine that the players are Cournot oligopolists whose outputs a_i are not publicly observable. Demand is stochastic, and so the market price, y , is a stochastic function of output (as in Green and Porter (1984)).

Finally, we can interpret the model as an insurance scheme under adverse selection (c.f. Green (1987)). Specifically, suppose that player i is subject to a stochastic preference shock, a_i is a function mapping the realization of the shock to a public "report" (not necessarily the truth), and y is the profile of reports together with a vector of transfers that the players make among themselves as a function of these reports.

Following convention, let V be the convex hull of the set of feasible payoff vectors $\{g(a) = (g_1(a), \dots, g_n(a)) | a \in A\}$. We shall assume that V has nonempty interior. (This is sometimes called the *full-dimensionality condition*.) As suggested in footnote 2, such a condition is typically imposed even for perfect-observability versions of the Folk Theorem. The *extremal points* of V are those that are not convex combinations of other points in V ; *extremal action profiles* are those that generate extremal payoffs. Note that for any extremal point there corresponds at least one extremal profile that is pure. Let $v_i = \min_{\alpha_{-i}} \max_{a_i} g_i(a_i, \alpha_{-i})$ be player i 's *minimax value*, and let $\underline{\alpha}_{-i}^i$ be mixed actions by the other players that attain this minimum. That is, $v_i = \max_{a_i} g_i(a_i, \underline{\alpha}_{-i}^i)$. Let $\underline{\alpha}_i^i$ be any best-response by player i to $\underline{\alpha}_{-i}^i$, that is, a (mixed) action such that $g_i(\underline{\alpha}_i^i, \underline{\alpha}_{-i}^i) = v_i$. We call $\underline{\alpha}^i = (\underline{\alpha}_i^i, \underline{\alpha}_{-i}^i)$ a *minimax profile* against player i .

The payoff vector v is *individually rational* if $v_i \geq v_i$ for all players i ; it is strictly individually rational if this inequality is strict. Let $V^* = \{v \in V | v_i \geq v_i \text{ for all } i\}$ be the set of feasible, individually rational payoffs. The Folk Theorem for the case of observable actions asserts that any strictly individually rational payoff vector in V^* can be supported by an equilibrium of the repeated game if the discount factor is sufficiently close to one.

⁷Although this model has the same normal form as the standard principal-agent game, it has more subgame-perfect equilibria because of the simultaneity of moves (in a simultaneous-move game, any Nash equilibrium is trivially subgame-perfect). To ensure that the equilibrium sets in the two models are the same, we could invoke solution concepts such as trembling hand-perfection (Selten (1975)) or sequential equilibrium (Kreps-Wilson (1982)). (Once we repeated the game, however, we would then have to modify these concepts since they apply only to finite games, whereas a repeated game has a continuum of strategies.) Alternatively we could perturb the game slightly so that the principal *literally* makes "mistakes" with small probability, in which case no strategy for the agent that is ruled out by backward induction in the sequential (standard) principal-agent game can be a weak best response in our simultaneous formulation. We should note, however, that some of our results guarantee *strict* equilibrium, in which case any of the usual refinements are satisfied.

Let us now turn in detail to the *repeated game*. In each period $t = 0, 1, \dots$, the stage game is played, resulting in a public outcome y^t . The *public history* at the end of period t is $h^t = (y^0, \dots, y^t)$. Player i 's *private history* at the end of period t is $h_i^t = (a_i^0, \dots, a_i^t)$. A *strategy* σ_i for player i is a sequence of functions $\{\sigma_i^t\}_{t=0}^\infty$, where σ_i^t maps each pair (h^{t-1}, h_i^{t-1}) to a probability distribution over A_i .

Each strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ generates a probability distribution over histories in the obvious way, and thus also generates a distribution over sequences of stage-game payoff vectors. Players discount future payoffs with a common discount factor δ , and player i 's objective in the repeated game is to maximize the expected value of the discounted sum of his stage game payoffs. This is equivalent to maximizing his expected discounted *average* payoff (the expected discounted sum multiplied by $1 - \delta$). Thus, if $\{g_i^t\}$ is player i 's sequence of stage-game payoffs, he maximizes

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t g_i^t.$$

(We use average rather than total payoffs in order to measure repeated and stage-game payoffs in the same units, that is, in "payoffs per period.")

For most of the paper we will focus on a special class of Nash equilibria called perfect public equilibria. A strategy $\sigma_i = \{\sigma_i^t\}$ for player i is *public* if, at each time t , σ_i^t depends only on the public history h^{t-1} and not on i 's private history h_i^{t-1} . A *perfect public equilibrium* (PPE) is a profile of public strategies that, beginning at any date t and given any public history h^{t-1} , form a Nash equilibrium from that point on. This definition allows a player to contemplate deviating to a non-public strategy, but such deviations are irrelevant: If the other players use public strategies, player i cannot profit from doing otherwise.

For any δ , let $E(\delta)$ be the set of discounted average payoff vectors that correspond to PPE's when the discount factor is δ . Clearly $E(\delta) \subseteq V^*$. Conversely, in our version of the Folk Theorem (Theorem 6.2) we will show that any point in the interior of V^* lies in $E(\delta)$ for δ near enough to 1.

In a PPE we can ascribe to each player probabilistic beliefs about which node pertains for each of his information sets in the repeated game. However, these beliefs play no role in the equilibrium: all nodes in any one of his information sets lead to the same probability distribution over his payoffs (since other players' private information—their past actions—does not affect their behavior). That is why even though—because of imperfect observability—there may be no proper subgames of the repeated game, we can treat each public history as if it gives rise to a distinct subgame.

It is clear that a PPE together with any beliefs consistent with Bayes' rule form a *perfect Bayesian equilibrium* (PBE):⁸ at any of his information sets a

⁸ Here we use the weakest notion of perfect Bayesian equilibrium: actions are required to be sequentially rational given beliefs, and beliefs after any history to which the equilibrium assigns positive probability are obtained using Bayes' rule. (Even this weak form of perfection is stronger than Nash equilibrium, as it rules out players using dominated strategies after zero-probability histories.)

TABLE 3.1

		Player 2	
		w	s
Player 1	w	1, 1	-1, 2
	s	2, -1	0, 0

player's equilibrium strategy maximizes his expected payoff, given others' strategies and his beliefs about nodes. A PPE, moreover, is "recursive" (Mertens (1987)) in the sense that it induces a PPE in every continuation game. This feature is important for our dynamic programming approach, as we "decompose" equilibrium payoffs for the repeated game into first-period payoffs and continuation payoffs, and it is essential that these continuation payoffs themselves be the payoffs of a PPE.

In contrast, a PBE need not be recursive, and hence need not be a PPE. (This was pointed out to us by A. Neyman.) Indeed, in games with three or more players, some PBE payoffs may lie outside V^* . We provided an example in the 1989 version of this paper; Lehrer (1988) has studied the issue in greater depth.⁹

3. EXAMPLES

To illustrate the main ideas of this paper, we consider two versions of a simple partnership game. In both versions, there are two partners, $i = 1, 2$, each of whom has two (unobservable) actions, "work" (w) and "shirk" (s). Expected output is 8 if both partners work, 4 if one works and the other shirks, and 0 if both shirk. Each player's utility is half of total output minus the disutility of his effort; work imposes a disutility of 3, and shirk costs 0. Thus the strategic form of both versions of the game is given in Table 3.1, and the unique equilibrium is (s, s) .

Although for any profile of effort levels, the corresponding expected output in the two versions is the same, the probability distributions over output differ. In the first version, there are two possible (publicly observable) output levels, $y = 0$ and $y = 12$. If both players work, the probability that $y = 12$ is $2/3$, if only one

⁹There is an important special case in which the sets of perfect Bayesian and perfect public equilibrium payoff vectors coincide, namely, a two-player game where only one player's actions are imperfectly observable (see Sections 7 and 9). In this case dynamic programming is applicable to perfect Bayesian equilibrium (PBE) because the continuation payoffs of a perfect Bayesian equilibrium must themselves correspond to a PBE. To see why, note that the difficulty that continuation equilibrium presents is that the imperfectly observable player's strategy might depend on his previous actions. We can, however, replace his strategy by one that averages over his strategies for all possible sequences of his previous actions. If we hold the other player's strategy fixed, this new strategy remains a best response (since it is an average of best responses) and leads to the same payoffs as before. If, moreover, the average is taken in a way that mimics the other player's beliefs about the previous actions, then that other player's strategy remains a best response too. Thus, we have constructed a perfect Bayesian equilibrium with the same payoffs as the continuation equilibrium. For details see Fudenberg and Maskin (1986b).

works the probability is $1/3$, and if neither works it is 0 . This example is similar to the Radner, Myerson, and Maskin (1986) counterexample in that there are only two public outcomes. As we shall see, this means that repeated-game equilibrium payoffs are bounded away from efficiency. In the second version, there are three possible output levels, $y = 12, 8,$ and 0 . If both players work the probability distribution over these levels is $(1/3, 1/2, 1/6)$; if player 1 shirks and 2 works, the distribution is $(1/3, 0, 2/3)$; if player 1 works and 2 shirks, it is $(0, 1/2, 1/2)$; and if both players shirk, it is $(0, 0, 1)$.¹⁰

Consider the first version. From Table 3.1, note that if $v = (v_1, v_2)$ is a Pareto-efficient point in V^* , $v_1 + v_2 \geq 3/2$. Nonetheless we claim that, in any equilibrium of the repeated game in which the discount factor is less than 1, the discounted average payoffs sum to no more than 1. To see this, fix δ and let $\gamma = \max\{v_1 + v_2 \mid (v_1, v_2) \in E(\delta)\}$. If, contrary to the claim, $\gamma > 1$, choose $v \in E(\delta)$ such that $v_1 + v_2 = \gamma$. Then, we can *decompose* v so that

$$(3.1) \quad v_i = (1 - \delta)g_i(\alpha) + \delta[\pi(12|\alpha)w_i(12) + \pi(0|\alpha)w_i(0)], \quad i = 1, 2,$$

where $\alpha = (\alpha_1, \alpha_2)$ is the profile of first-period actions and $w_i(y)$, $y = 0, 12$, is player i 's continuation payoff if first-period output is y in the repeated-game equilibrium corresponding to v . We note that α_i puts positive probability on "work" for $i = 1, 2$: if not, then, from Table 3.1 the first-period payoffs could sum to at most 1, implying, from (3.1), that either $(w_1(0), w_2(0))$ or $(w_1(12), w_2(12))$ would sum to more than γ , a contradiction. Moreover, because player 1 works with positive probability, his payoff from shirking cannot exceed that from working, i.e.,

$$(3.2) \quad (1 - \delta)(2\mu_2 - 1) + \delta\left[\left(\frac{1}{3}\mu_2 + \frac{1}{3}\right)w_1(12) + \left(-\frac{1}{3}\mu_2 + \frac{2}{3}\right)w_1(0)\right] \\ \geq (1 - \delta)2\mu_2 + \delta\left[\frac{1}{3}\mu_2 w_1(12) + \left(-\frac{1}{3}\mu_2 + 1\right)w_1(0)\right],$$

where μ_2 is the probability that α_2 assigns to "work." Simplifying (3.2), we obtain

$$(3.3) \quad w_1(12) - \frac{3(1 - \delta)}{\delta} \geq w_1(0).$$

Substituting for $w_1(0)$ in (3.1) using (3.3) and noting that $2\mu_2 - 1 \leq 1$ and $\pi(0|\alpha) \geq 1/3$, we have

$$(3.4) \quad v_1 \leq \delta w_1(12).$$

Similarly for player 2

$$(3.5) \quad v_2 \leq \delta w_2(12).$$

¹⁰ In this second version, the profiles (shirk, work) and (work, shirk) induce different probability distributions over outcomes. This asymmetry simplifies the analysis by allowing us to construct approximately efficient equilibria in pure actions, but it is not required for our results, as explained in footnote 13 and the discussion following Condition 6.2 in Section 6.

However, adding (3.4) and (3.5) and noting that $w_1(12) + w_2(12) \leq \gamma$ yields $v_1 + v_2 \leq \delta\gamma$, a contradiction. We conclude that the hypothesis that $\gamma > 1$ must be false.

Repeated-game equilibria are necessarily inefficient in this example despite the fact that the profile (work, work)—as well as any other profile—is *enforceable*. I.e., we can find continuation payoff vectors $w(12)$ and $w(0)$ such that $a_1 = \text{work}$ maximizes¹¹

$$(3.6) \quad (1 - \delta)g_1(a_1, \text{work}) + \delta(\pi(12|a_1, \text{work})w_1(12) + \pi(0|a_1, \text{work})w_1(0)),$$

and $a_2 = \text{work}$ maximizes the corresponding expression for player 2. The problem is that to induce both players to work, we must threaten them *both* with a low continuation payoff if $y = 0$ (i.e., we must have $w_i(0) < w_i(12)$ for $i = 1, 2$). This need to “punish” both players for low output (which occurs with positive probability even when both players work) is what creates the inefficiency.

We now turn to the three-outcome version of the partnership game. Choose a small ball W in the interior of V^* (see Figure 3.1). We claim that for δ near enough 1, $W \subseteq E(\delta)$. Hence the Folk Theorem extends to this game. From dynamic programming (see Lemma 4.2), it suffices to show that every point $v \in W$ can be decomposed so that, for some profile α and continuation payoff vectors $w(12), w(8), w(0) \in W$,

$$\begin{aligned} v_i &= (1 - \delta)g_i(\alpha) + \delta \sum_y \pi(y|\alpha)w_i(y) \\ &\geq (1 - \delta)g_i(\alpha'_i, \alpha_{-i}) + \delta \sum_y \pi(y|\alpha'_i, \alpha_{-i})w_i(y), \end{aligned}$$

for all α'_i and $i = 1, 2$.

Divide W into four parts, A, B, C , and D , as shown in Figure 3.1. Consider v on the boundary of A . To decompose v , choose $\alpha = (\text{work}, \text{work})$. To ensure that $w(12), w(8)$, and $w(0)$ belong to W , we select them to lie along a line P'

¹¹ The requirement that $a_1 = \text{work}$ maximize (3.6) can be re-expressed as

$$(*) \quad v_1 \geq 2(1 - \delta) + \delta\left(\frac{1}{3}w_1(12) + \frac{2}{3}w_1(0)\right).$$

If we set

$$(**) \quad 1 - \delta + \delta\left(\frac{2}{3}w_1(12) + \frac{1}{3}w_1(0)\right) = v_1,$$

then, (*) and (**) become

$$(***) \quad \begin{pmatrix} \frac{1}{3} & -\frac{2}{3} \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix} \begin{pmatrix} w_1(12) \\ w_1(0) \end{pmatrix} \leq \begin{pmatrix} \frac{v_1 - 2(1 - \delta)}{\delta} \\ \frac{v_1 - 1 + \delta}{\delta} \end{pmatrix}.$$

(There is a slight abuse of notation here since the second inequality in (***) is actually the equation (**).) The matrix of conditional probabilities in (***) has rank 2, equal to the number of player 1's actions. This, in our terminology, means that the profile (work, work) has *individual full rank* for player 1—and so (***) can be solved with equality.

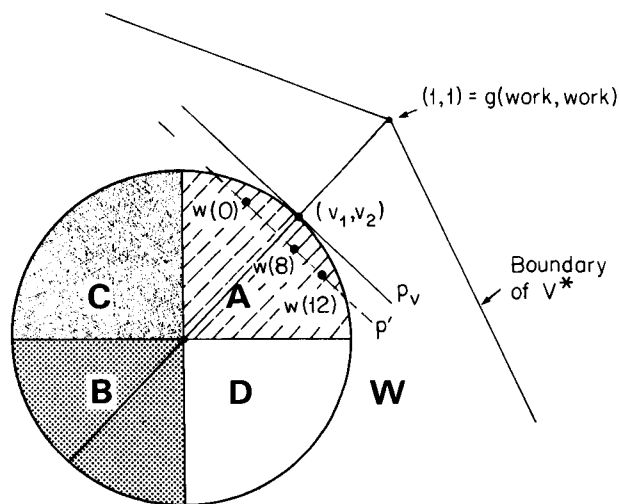


FIGURE 3.1

parallel to the tangent P_v to W at v (see Figure 3.1). Thus $(\text{work}, \text{work})$ will be *enforceable with respect to the line (hyperplane) P'* . P' takes the form $\{(v'_1, v'_2) | \beta_1 v'_1 + \beta_2 v'_2 = c\}$. Hence, selecting $w(12)$, $w(8)$, and $w(0)$ amounts to solving the system:

$$(3.7) \quad v_1 \geq 2(1 - \delta) + \delta\left(\frac{1}{3}w_1(12) + \frac{2}{3}w_1(0)\right)$$

(player 1 does not gain from shirking),

$$(3.8) \quad v_2 \geq 2(1 - \delta) + \delta\left(\frac{1}{2}w_2(8) + \frac{1}{2}w_2(0)\right)$$

(player 2 does not gain from shirking),

$$(3.9) \quad 1 - \delta + \delta\left(\frac{1}{3}w_1(12) + \frac{1}{2}w_1(8) + \frac{1}{6}w_1(0)\right) = v_1$$

(player 1's expected payoff is v_1).¹²

If δ is near enough 1, $w(12)$, $w(8)$, $w(0)$ will lie near the intersection of P' and the line through $g(\alpha)$ and v , and so will lie in W . Because these continuation payoffs lie in P' , we can replace each $w_2(y)$ in (3.8) by $(c - \beta_1 w_1(y))/\beta_2$ and

¹² We can omit the requirement that player 2's expected payoff be v_2 , since, given (3.9), and the stipulation that $w(12)$, $w(8)$, and $w(0)$ lie in P' , this is ensured automatically.

rewrite (3.7)–(3.9) as equalities to obtain

$$(3.10) \quad \begin{pmatrix} \frac{1}{3} & 0 & \frac{2}{3} \\ 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{3} & \frac{1}{2} & \frac{1}{6} \end{pmatrix} \begin{pmatrix} w_1(12) \\ w_1(8) \\ w_1(10) \end{pmatrix} = \begin{pmatrix} \frac{v_1 - 2(1 - \delta)}{\delta} \\ \frac{\beta_2(2(1 - \delta) - v_2) + c\delta}{\beta_1\delta} \\ \frac{v_1 - 1 + \delta}{\delta} \end{pmatrix}.$$

Now the matrix in (3.10) has full rank, and so (3.10) can be solved with equality. In our terminology, the profile (work,work) has *pairwise full rank*. Roughly speaking, a profile α has this property for players i and j if every deviation that i or j could make from α induces a different distribution over public outcomes. Pairwise full rank ensures that the profile (work,work) is enforceable with respect to the line (hyperplane) P' (or, for that matter, with respect to almost any other hyperplane).¹³ From Figure 3.1, we see that this means that, when low output ($y = 0$) occurs it is not necessary to punish both players, in contrast with the two-outcome partnership game. That earlier game had too few public outcomes to have profiles with pairwise full rank.

We have shown that any point v in A can be decomposed so that the continuation payoffs lie in W when δ is near 1. Points in B , C , and D can be decomposed similarly (and even more easily). Indeed, W is *decomposable on tangent hyperplanes* since each boundary point v is decomposable in such a way that, as in Figure 3.1, the continuation payoffs lie on a hyperplane parallel to the tangent to W at v . In Section 4 we will show more generally that any set W that is decomposable on tangent hyperplanes belongs to $E(\delta)$ for δ near enough to 1.

The existence of a profile having pairwise full rank was the key to our demonstration that the Folk Theorem extends to this partnership game. In Theorems 6.1 and 6.2 below we show that the existence of such a profile allows us to draw Folk-Theorem-like results in games in general. There are, however, a good many games of interest that have no profiles with pairwise full rank but to which the Folk Theorem nevertheless applies. One prominent example is a principal-agent model. Suppose that the principal (player 1) first proposes a compensation scheme in which the agent's (player 2's) monetary payment t is contingent on observable output y . Output is either "high" ($y = 12$) or "low" ($y = 0$). The principal's von Neumann-Morgenstern utility function is $y - t$. The

¹³ If we modified the three-outcome example so that the profiles (shirk,work) and (work,shirk) induced the same distribution over outcomes, then the profile (work,work) would no longer satisfy pairwise full rank. However, pairwise full rank would still hold at many mixed-action profiles in which both players work with probability near 1, and this suffices for our conclusions (see the discussion following Condition 6.2).

agent responds by choosing one of three possible (unobservable) effort levels: low ($a_2 = 0$), medium ($a_2 = 1$), or high ($a_2 = 2$). Low, medium, and high efforts result in probabilities 0, 1/2, and 2/3, respectively, of high output. The agent's utility function is $u(t) - a_2$ where u is strictly concave.

Suppose that the principal is restricted to proposing compensation schemes that yield the agent nonnegative expected utility. Let us assume that, if the agent's actions were observable, the equilibrium effort by the agent would be $a_2 = 1$. Under that assumption, lack of observability causes a loss of welfare: from the standpoint of risk-sharing the agent's compensation should be fully insured, but in that case he will choose $a_2 = 0$ rather than $a_2 = 1$. Thus equilibrium in the one-shot game with unobservable actions is Pareto-inefficient. Moreover, unlike in the three-outcome partnership game, no profile has pairwise full rank; indeed, no profile even has *individual* full rank for the agent (see footnote 11).¹⁴ Nevertheless payoffs that are arbitrarily close to the Pareto frontier are attainable when the game is repeated. This is because, despite the lack of full rank, a Pareto-efficient profile (t, a_2) is still enforceable. Specifically, from Pareto efficiency there exist weights λ_1 and λ_2 such that $a'_2 = a_2$ maximizes

$$\lambda_1 \sum_y \pi(y|a'_2)(y - t) + \lambda_2(u(t) - a'_2).$$

Hence, if we take

$$w_2(y) = \frac{1 - \delta}{\delta} \frac{\lambda_1}{\lambda_2} (y - t),$$

$a'_2 = a_2$ maximizes player 2's payoff. Because the principal's actions are observable, (t, a_2) 's enforceability means that it is, in fact, enforceable with respect to (almost) any hyperplane (refer again to the discussion before formula (3.7) for the concept of enforceability with respect to hyperplanes): once the agent's continuation payoffs $w_2(y)$ have been specified to induce him to choose a_2 , we can simply choose the corresponding $w_1(y)$ so that the resulting vector $w(y)$ lies in the desired hyperplane. Hence, the same kind of decomposition on tangent hyperplanes that we performed for the three-outcome partnership game can be done here. We show in Section 7 that there is a much broader class of games than principal-agent models, viz., games with a "product structure," to which the Folk Theorem extends without invoking any full rank conditions.

4. ENFORCEABLE ACTIONS AND DECOMPOSABLE PAYOFFS

We now introduce the concepts of enforceability and decomposability formally. The section culminates in the demonstration that if a set of payoff vectors

¹⁴ From footnote 11, we see that a profile's having individual full rank is equivalent to the corresponding matrix of conditional probabilities having rank equal to the number of the agent's actions. But the matrix's rank cannot exceed the number of observable outcomes m (since the number of columns equals m), which in this case is only 2, less than the number of actions, 3.

is decomposable on tangent hyperplanes, then the payoffs are attainable as perfect public equilibria for discount factors sufficiently near one. Readers who are more interested in the statement of the Folk Theorem and its applications than in the logic of its proof may wish to skip this section on first reading and proceed to Section 5, which develops readily verifiable sufficient conditions for decomposability on tangent hyperplanes.

LEMMA 4.1: $E(\delta)$, V , and V^* are compact sets, and V and V^* are convex.

PROOF: The compactness of $E(\delta)$ is well known (see, for example, Fudenberg and Levine (1983)). The sets V and V^* are obviously compact; V is convex by definition, and V^* is the intersection of two convex sets. Q.E.D.

DEFINITION 4.1: Let δ and $W \subseteq \mathbb{R}^n$ be given. Profile α is *enforceable with respect to W and δ* if there exist $v \in \mathbb{R}^n$ and a function $w : Y \rightarrow W$ such that, for all i ,

$$(4.1) \quad \begin{cases} v_i = (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_y \pi(y|a_i, \alpha_{-i})w_i(y), \\ \quad \text{for all } a_i \text{ with } \alpha_i(a_i) > 0, \\ v_i \geq (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_y \pi(y|a_i, \alpha_{-i})w_i(y), \\ \quad \text{for all } a_i \text{ with } \alpha_i(a_i) = 0, \end{cases}$$

where $w_i(y)$ is the i th component of $w(y)$. We say that the set $\{w(y)\}_{y \in Y}$ *enforces α with respect to v and δ* , and that v is *decomposable with respect to α , W , and δ* (v is decomposable via (4.1) into first-period payoffs $g(\alpha)$ and continuation payoffs lying in W). If α is enforceable with respect to some W and δ , we say simply that it is *enforceable*. We let $B(W, \delta)$ be the set of all payoff vectors that are decomposable as we vary the profile α with respect to fixed W and δ . If $W \subseteq B(W, \delta)$ for some δ , we say that W is *self-decomposable*.

Condition (4.1) requires that if the others play α_{-i} and player i 's expected continuation payoff contingent on outcome y is $w_i(y)$, then i obtains the same average payoff from all actions in the support of α_i , and no action yields a higher average payoff.

The concept of a self-decomposable set of payoffs was introduced by Abreu, Pearce, and Stacchetti (1986, 1990) (who used the term "self-generating"). They showed how to apply the principle of optimality of dynamic programming to a class of games with imperfect public information, and established that if W is bounded and $W \subseteq B(W, \delta)$, then W consists of equilibrium payoffs in the repeated game with discount factor δ (in particular, this means that $W \subseteq V^*$).¹⁵

¹⁵ Abreu, Pearce, and Stacchetti considered only pure-strategy equilibria, but the essence of their argument carries over to perfect public equilibria in mixed strategies (see Lemma 4.2).

To obtain results for discount factors near one, we find it convenient to work with the weaker concept of local self-decomposability.

DEFINITION 4.2: A subset W of \mathbb{R}^n is *locally self-decomposable* if for each $v \in W$ there is a $\delta < 1$ and an open set U containing v such that $U \cap W \subseteq B(W, \delta)$.

Note that in the definition of local self-decomposability (in contrast to that of self-decomposability) the discount factor may vary over W .

LEMMA 4.2: *If $W \subseteq \mathbb{R}^n$ is compact, convex, and locally self-decomposable, then there exists $\delta' < 1$ such that $W \subseteq E(\delta)$ for all $\delta \in (\delta', 1)$.*

PROOF: Since W is locally self-decomposable, there exist an open cover $\{U\}$ of W and associated discount factors $\{\delta_U\}$ such that for each $v \in U$ there exists α_v that is enforceable with respect to W and δ_U and satisfies (4.1). Since W is compact, $\{U\}$ has a finite subcover. Let δ' be the maximum of the δ_U 's on this subcover. We claim that any $v \in W$ is in $B(W, \delta)$ for any $\delta \in (\delta', 1)$. To see this, set

$$(4.2) \quad w(y; \delta) = [(\delta - \delta_U)/\delta(1 - \delta_U)]v \\ + [\delta_U(1 - \delta)/\delta(1 - \delta_U)]w(y; \delta_U),$$

where $\{w(y; \delta_U)\}_{y \in Y}$ enforces α_v with respect to v and δ_U . It is easy to check that (4.1) is satisfied for α_v, v , and any $\delta \geq \delta'$; moreover, $w(y; \delta)$ is in W since both v and $w(y; \delta_U)$ belong to W , and W is convex.

To complete the proof, we show that $W \subseteq E(\delta)$. Fix an arbitrary $v \in W$, and construct a strategy profile σ as follows. Let α_v and $w(\cdot)$ satisfy (4.1) for v, δ , and W , and let the date-0 actions be α_v . For each date-0 outcome y^0 , $w(y^0) \in W$, and so there is an action profile $\alpha_{w(y^0)}$ and continuation function $w(\cdot)$ that satisfies (4.1) for $w(y^0), \delta$, and W . We then let the date-1 actions be $\sigma(h^0) = \alpha_{w(y^0)}$. Continuing iteratively, we define action profiles in each period as a function of the public history to date. By construction and the boundedness of W , the sequence of actions thus defined gives rise to the discounted average payoff vector v . The principle of optimality implies that no player can gain by deviating, since (4.1) implies that no one-period deviation is profitable, and the payoffs are uniformly bounded. *Q.E.D.*

Instead of explicitly constructing equilibrium strategies with particular payoff vectors (see, however, the following paragraph), we give sufficient conditions for a set of payoffs to be locally self-decomposable; the existence of the corresponding equilibrium then follows from Lemma 4.2. Each set U that we identify not

only satisfies Definition 4.2, but also the properties that (i) all payoff vectors in U can be decomposed with respect to a *single* action profile α_U , and (ii) there exists a continuation function $w(\cdot; \delta_U)$ such that any $v \in U$ is decomposable with respect to α_U , the set $\{w(y; \delta_U)\}_{y \in Y}$, and discount factor δ_U .

This allows us to be a bit more concrete about the structure of the equilibrium strategies. Specifically, fix a locally self-decomposable set W . Suppose that we wish to attain the payoff vector $v \in W$ in equilibrium. As in the proof of Lemma 4.2, W is covered by finitely many regions U . (There is no harm in assuming that these regions form a partition of W , although strictly speaking, this would conflict with their being open.) We can now define a payoff statistic v^t and construct equilibrium strategies in which, in period t , α_U is played provided that v^t lies in U (recall (i) above). To define v^t , set $v^0 = v$, and take

$$v^t = [(\delta - \delta_U)/\delta(1 - \delta_U)]v^{t-1} + [\delta_U(1 - \delta)/\delta(1 - \delta_U)]w_U(y; \delta_U).$$

After updating the payoff statistic, players determine the region to which v^{t+1} belongs and play the corresponding action profile.

For the remainder of this section we consider action profiles that are enforceable with respect to particular *hyperplanes* (we omit any mention of discount factors here by virtue of Lemma 4.3 below), that is, sets of the form $P = v' + H (= \{v' + w | w \in H\})$ where v' is a fixed vector and H is an $(n - 1)$ -dimensional linear subspace of \mathbb{R}^n .¹⁶ Theorem 4.1 shows that if a set of payoffs is “decomposable on tangent hyperplanes” (which means that the corresponding action profiles are enforceable with respect to hyperplanes that are tangent to the set; see Definition 4.4), then it is locally self-decomposable. The following lemma is a useful preliminary:

LEMMA 4.3: *Suppose that P is a hyperplane in \mathbb{R}^n containing the origin. If α is enforceable with respect to P and some δ , then (i) α is enforceable with respect to any translate $P' = v' + P$ and any $\delta' > 0$; (ii) there is a constant κ such that, for all δ' and $v' \in \mathbb{R}^n$, there exists a continuation function $w': Y \rightarrow P'$ such that $\{w'(y)\}_{y \in Y}$ enforces α with respect to v' and δ' , $v' = \sum_y \pi(y|\alpha)w'(y)$, and $\|w'(y) - v'\| < \kappa(1 - \delta')/\delta'$ for all $y \in Y$, where $\|\cdot\|$ denotes Euclidean distance.*

REMARK: Part (i) of the lemma allows us henceforth to avoid any mention of discount factors when stating that a profile is enforceable with respect to some hyperplane. It also implies that enforceability with respect to P is equivalent to enforceability with respect to all translates of P , and so we will sometimes rather loosely identify a hyperplane with its translates. Part (ii) shows that when enforcing α with respect to a hyperplane, continuation payoff vectors can be found that differ from each other by no more than order $1 - \delta$ when δ is near 1.

¹⁶ P can also be expressed as $P = \{v \in \mathbb{R}^n | \sum \beta_i v_i = c\}$ for some nonzero vector $(\beta_1, \dots, \beta_n)$ and constant c .

This observation, coupled with the fact that sets with smooth boundaries can be approximated by their tangent hyperplanes, is the basis of Theorem 4.1.

PROOF: Suppose that α is enforceable with respect to P and δ using the continuation function $w(\cdot)$. Clearly α is enforceable with respect to δ and any translate $v' + P$ since we can simply add v' to all continuation payoff vectors and $\delta v'$ to the other side of the equation. To see that we can enforce α with respect to any discount factor and establish the claims of part (ii), choose an arbitrary v' and δ' , and take

$$w'(y) = v' + [\delta(1 - \delta')/\delta'(1 - \delta)][w(y) - \bar{w}],$$

where $\bar{w} = \sum_y \pi(y|\alpha)w(y)$. Computation verifies that α is enforceable with respect to δ' and P' , that $\sum_y \pi(y|\alpha)w'(y) = v'$, and that the difference between $w'(y)$ and v' equals

$$[\delta(1 - \delta')/\delta'(1 - \delta)]\|w(y) - \bar{w}\|. \quad Q.E.D.$$

DEFINITION 4.3: A set $W \subseteq \mathbb{R}^n$ is *smooth* if (i) it is closed and convex; (ii) it has a nonempty interior; and (iii) at each boundary point v there is a unique tangent hyperplane P_v , which varies continuously with v , i.e., the boundary is a C^2 -submanifold of \mathbb{R}^n .

We indicated in Section 3 that extending the Folk Theorem to games with unobservable actions depends on whether we can enforce action profiles with continuation payoffs lying in particular hyperplanes. This motivates the following definition.

DEFINITION 4.4: A smooth subset $W \subseteq V^*$ is *decomposable on tangent hyperplanes* if for every point v on the boundary of W there exists a profile α such that (i) $g(\alpha)$ is separated from W by the (unique) $(n - 1)$ -dimensional hyperplane P_v that is tangent to W at v and (ii) there exist continuation payoff vectors $\{w(y)\}_{y \in Y}$ in P_v that enforce α .¹⁷

Note that in view of the remark following Lemma 4.3 we have not specified a discount factor with respect to which the continuation payoffs $\{w(y)\}_{y \in Y}$ enforce α in part (ii). As we shall see next, the ability to decompose W on tangent hyperplanes implies that it is locally self-decomposable. In the proof of this implication, we will choose enforcing continuation payoffs to lie on a translate of P_v (which Lemma 4.3 permits us to do) so that they belong to W .

¹⁷Strictly speaking, v in this definition is not decomposable with respect to the tangent hyperplane P_v (since it actually lies in P_v), but rather with respect to one of P_v 's translates (see the paragraph following the definition). Thus the locution "decomposable on tangent hyperplanes" may be a slight misnomer. It is, however, less cumbersome than "decomposable with respect to a translate of the tangent hyperplane."




FIGURE 4.1

THEOREM 4.1: *If a smooth set $W \subseteq V^*$ is decomposable on tangent hyperplanes, then it is locally self-decomposable. Hence, from Lemma 4.2, there exists $\underline{\delta} < 1$ such that, for all $\delta > \underline{\delta}$, $W \subseteq E(\delta)$.*¹⁸

SKETCH OF PROOF: Fix v on the boundary of W and let P_v be the tangent hyperplane to W at v (see Figure 4.1). By hypothesis, there exist a profile α and continuation function $w : Y \rightarrow P_v$ such that $\{w(y)\}_{y \in Y}$ enforces α and $g(\alpha)$ is separated from W by P_v . We must show that, for δ near enough to 1, v is decomposable with respect to α , W , and δ .

For given δ choose w' so that $v = (1 - \delta)g(\alpha) + \delta w'$. For δ near enough to 1, w' is near v and so lies in the interior of W , since $g(\alpha)$ is separated from W by P_v . By Lemma 4.3, profile α is enforceable with respect to the translate P' of P_v that passes through w' , and by adding an appropriate vector we can ensure that the continuation payoffs satisfy $w' = \sum \pi(y|\alpha)w(y)$. Lemma 4.3 then guarantees that the distance between w' and $w(y)$ is no greater than of order $(1 - \delta)/\delta$.

Now because W is smooth and the distance from v to w' is of order $(1 - \delta)/\delta$, the distance from w' to the boundary of W along P' is of order no less than $\sqrt{(1 - \delta)/\delta}$ (see Figure 4.1). Because $\sqrt{(1 - \delta)/\delta}$ is bigger than $(1 - \delta)/\delta$, therefore, the continuation payoffs $\{w(y)\}$ lie in W for δ near 1, as required. Details may be found in Appendix 1.

5. ENFORCEABILITY AND IDENTIFIABILITY

Fix a smooth subset W of V^* . Theorem 4.1 showed that all payoff vectors in W can be attained as equilibria provided that any boundary point of W can be decomposed using a suitable action profile that is enforceable with respect to the tangent hyperplane at that point. This section relates enforceability with

¹⁸ Matsushima (1989) used the concept of decomposability on tangent hyperplanes to show that some efficient outcomes can be approximated by repeated game equilibria for discount factors near 1. The principal differences between his approach and ours are that (i) Matsushima *assumed* that the set of equilibrium payoffs has a smooth boundary, whereas we work with subsets of equilibrium payoffs that can be *chosen* to have smooth boundaries; (ii) Matsushima supposed that first-order conditions are sufficient for equilibrium, while we require no such supposition; and (iii) Matsushima was interested in showing only that there exist (approximately) efficient repeated-game equilibria, whereas we are more concerned with characterizing the full set of equilibria.

respect to hyperplanes to the ability of players to identify one another's actions on the basis of public outcomes.

Consider an arbitrary profile α . If α is to be enforceable with respect to some hyperplane, it must first simply be enforceable. Suppose that player i has an action α'_i such that profiles (α'_i, α_{-i}) and (α_i, α_{-i}) lead to the same probability distribution over public outcomes. If $g_i(\alpha'_i, \alpha_{-i}) > g_i(\alpha_i, \alpha_{-i})$, then *no* choice of continuation payoffs can enforce profile α , since player i does better in the initial period with α'_i than with α_i , and has exactly the same expected continuation payoff in either case. This difficulty can be ruled out if there is no $\alpha'_i \neq \alpha_i$ for which (α'_i, α_{-i}) and (α_i, α_{-i}) induce the same probability distribution. Since mixed action profiles correspond to convex combinations of pure actions (and in particular the weight on each pure action is between 0 and 1), it is sufficient that no *linear* combination of pure actions in A_i leads to the same distribution as α_i when other players play α_{-i} .

Before stating this condition more formally, we recall that m_i is the number of actions in A_i and that m is the number of outcomes in Y . Thus $\pi(\cdot | \cdot, \alpha_{-i})$ can be thought of as an $m_i \times m$ matrix whose rows correspond to actions a_i and whose columns to public outcomes y . To simplify notation, let $\Pi_i(\alpha_{-i}) = \pi(\cdot | \cdot, \alpha_{-i})$.

DEFINITION 5.1: The profile α has *individual full rank for player i* if $\Pi_i(\alpha_{-i})$ has rank equal to m_i , that is, the m_i vectors $\{\pi(\cdot | a_i, \alpha_{-i})\}_{a_i \in A_i}$ are linearly independent. If this is so for every player i , α has *individual full rank*.¹⁹

Note that if α has individual full rank, the number of observable outcomes m must be at least $\max_i m_i$.

LEMMA 5.1: *If α has individual full rank, then α is enforceable.*

PROOF: The constraints in (4.1) for player i are satisfied if

$$\sum \pi(y | a_i, \alpha_{-i}) w_i(y) = (v_i - (1 - \delta) g_i(a_i, \alpha_{-i})) / \delta,$$

for all a_i . The full rank condition for player i ensures that this system can be solved for the $w_i(y)$'s. *Q.E.D.*

This lemma establishes that if, given α_{-i} , we can identify player i 's actions probabilistically, we can choose continuation payoffs to make him exactly indifferent among all his actions.

We now turn to enforceability with respect to hyperplanes. For our purposes, we must distinguish between coordinate and regular hyperplanes. The *coordinate hyperplanes* are orthogonal to a coordinate axis, whereas any hyperplane that is not coordinate is *regular*. Algebraically, a hyperplane $\{v \in \mathbb{R}^n | \sum \beta_i v_i = c\}$

¹⁹ This condition was introduced in Fudenberg and Maskin (1986b).

is coordinate if exactly one coefficient β_i is nonzero; it is regular if it has two or more nonzero components.

Suppose that P is a coordinate hyperplane that is orthogonal to the i th axis. Then, in (4.1), all the $w(y)$'s must have the same i th component if they belong to P . Hence for profile α to be enforceable with respect to P , α_i must be a best-response to α_{-i} in the stage game (that is, a *static* best-response). Moreover, if α is enforceable, this *best-response property* for player i is sufficient for α to be enforceable with respect to P because this hyperplane places no restriction on the components other than i .

An equilibrium profile of the stage game obviously satisfies the best-response property. It is also enforceable since the constraints (4.1) are satisfied if all the $w(y)$'s are the same. Similarly, a profile a^i that maximizes player i 's payoff on the joint action set A clearly has the best-response property for player i . As we shall see in Section 7, such a profile is also enforceable under mild regularity conditions. These facts enable us to prove extensions of the Folk Theorem that use a stage-game equilibrium (rather than the minimax point) as the threat point. "Minimax-threat" extensions require additional hypotheses since, although a minimax profile \underline{a}^i against player i satisfies the best-response property for player i , it need not be enforceable.

For future reference, we summarize this discussion in the following lemma.

LEMMA 5.2: *If a profile α is enforceable and satisfies the best-response property for player i , then it is enforceable with respect to coordinate hyperplanes orthogonal to the i th coordinate axis. The best-response property is satisfied by (i) a profile that maximizes i 's payoff on A , (ii) a minimax profile against i , and (iii) a static equilibrium profile; static equilibrium profiles are enforceable with respect to all hyperplanes, coordinate or regular.*

We now consider regular hyperplanes. We provide conditions under which profile α is enforceable with respect to all regular hyperplanes.²⁰ We will also consider *strong enforceability*.

DEFINITION 5.2: The profile α is *strongly enforceable with respect to the hyperplane P* if, in applying the definition of enforceability, we can find a continuation function $w: Y \rightarrow P$ such that the inequalities (4.1) hold with equality for all players and actions.

We first show that we can limit our attention to two-player games. Specifically, consider *pairwise* hyperplanes, i.e., (regular) hyperplanes for which only two β_k 's are nonzero (that is, no constraint is imposed on the payoffs of the other $n - 2$ players).

²⁰ By Lemma 4.3 the discount factor is irrelevant when considering whether a profile is enforceable with respect to a hyperplane.

DEFINITION 5.3: For given i and j , a profile that is enforceable with respect to all pairwise hyperplanes in which only β_i and β_j are nonzero is called *enforceable with respect to all ij -pairwise hyperplanes*. If it is so enforceable for any i and j , it is called *enforceable with respect to all pairwise hyperplanes*.

Obviously, enforceability with respect to all regular hyperplanes implies enforceability with respect to all pairwise hyperplanes. The converse is only slightly less transparent.

LEMMA 5.3: *A profile is enforceable (strongly enforceable) with respect to all regular hyperplanes if and only if it is enforceable (strongly enforceable) with respect to all pairwise hyperplanes.*

PROOF: It is clear that enforceability with respect to all regular hyperplanes implies enforceability with respect to all pairwise hyperplanes. To prove the converse, consider first a regular hyperplane $\sum \beta_i v_i = c$ for which the number of nonzero components of β is even. We may assume without loss of generality that β_1 and β_2 are nonzero. Then a sufficient condition for $\sum \beta_i w_i(y) = c$ is that $\beta_1 w_1(y) + \beta_2 w_2(y) = c$, $\beta_3 w_3(y) + \beta_4 w_4(y) = 0$, and so forth. But this simply reduces to the solving of the incentive constraints (4.1) on hyperplanes with coefficients $(\beta_1, \beta_2, 0, \dots, 0, (0, 0, \beta_3, \beta_4, 0, \dots, 0), \dots$. (Recall that translating a hyperplane preserves enforceability.)

Thus, the lemma is proved if there is an even number of nonzero components of β , and if there is an odd number it suffices to consider the case in which there are three nonzero components. In this case we can separately solve the following two equations: $\beta_1 w_1(y) + (\beta_2/2)w_2'(y) = c$ and $(\beta_2/2)w_2''(y) + \beta_3 w_3(y) = 0$. If we now let $w_2(y) = (1/2)(w_2'(y) + w_2''(y))$, the continuation payoffs $\{w(y)\}$ lie in the desired hyperplane, and, moreover, satisfy the constraints (4.1), since they are convex combinations of payoffs that satisfy (4.1). *Q.E.D.*

We now examine when profile α is enforceable with respect to all ij -pairwise hyperplanes for a given pair of players i and j . The requirement that continuation payoffs lie in hyperplane $\beta_i w_i + \beta_j w_j = c$ (so that $\beta_k = 0$ for all $k \notin \{i, j\}$) amounts to imposing an “exchange rate” of β_j/β_i between player i 's and j 's continuation payoffs. Thus it is as though utility were transferable between the players. Following this analogy, we can think of an “aggregate” player who can choose from either player i 's or j 's action set, and so faces the union of their incentive constraints. To apply the individual full rank condition to this aggregate player, we construct an $(m_i + m_j) \times m$ matrix by “stacking” $\Pi_i(\alpha_{-i})$ on top of $\Pi_j(\alpha_{-j})$. Call this matrix $\Pi_{ij}(\alpha)$:

$$\Pi_{ij}(\alpha) = \begin{pmatrix} \Pi_i(\alpha_{-i}) \\ \Pi_j(\alpha_{-j}) \end{pmatrix}.$$

Matrix $\Pi_{ij}(\alpha)$ has a linear dependence among its rows. This is clearly seen when α is the profile where each player plays his first pure strategy, in which case the first rows of $\Pi_i(\alpha_{-i})$ and $\Pi_j(\alpha_{-j})$ are the same. More generally,

$$\pi(\cdot|\alpha) = \sum_{a_i \in A_i} \alpha_i(a_i)\pi(\cdot|a_i, \alpha_{-i}) = \sum_{a_j \in A_j} \alpha_j(a_j)\pi(\cdot|a_j, \alpha_{-j}).$$

In view of this linear dependence, $\Pi_{ij}(\alpha)$ has rank at most $m_i + m_j - 1$. When the maximum is attained, we say that α has pairwise full rank:

DEFINITION 5.4: Profile α has *pairwise full rank* for players i and j if the matrix $\Pi_{ij}(\alpha)$ has rank $m_i + m_j - 1$.

Note that this is a considerably weaker condition than “ n -wise full rank,” which requires that the distributions corresponding to all n players’ actions have maximal rank. When this stronger condition is satisfied, the incentive constraints (4.1) can be solved with equality in any regular hyperplane. Observe, however, that n -wise full rank demands at least $\sum_i m_i - 1$ public outcomes, which can be a good many more than $\max_{i,j}\{m_i + m_j - 1\}$.

LEMMA 5.4: *If a profile has pairwise full rank for players i and j , then it is strongly enforceable with respect to all ij -pairwise hyperplanes.*

PROOF: We must show that pairwise full rank implies that, for any constant c and nonzero β_i and β_j , there exist $v_i, v_j, w_i(\cdot)$, and $w_j(\cdot)$ such that, for all a_i and a_j ,

$$(5.1) \quad v_i = (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \sum_y \pi(y|a_i, \alpha_{-i})w_i(y),$$

$$v_j = (1 - \delta)g_j(a_j, \alpha_{-j}) + \delta \sum_y \pi(y|a_j, \alpha_{-j})w_j(y),$$

and, for all y ,

$$(5.2) \quad \beta_i w_i(y) + \beta_j w_j(y) = c.$$

We can eliminate (5.2) by solving for $w_j(y)$ (which we can do since $\beta_j \neq 0$). Formulae (5.1) must hold in particular when a_i and a_j are replaced by α_i and α_j . Therefore,

$$\begin{aligned} v_j &= (1 - \delta)g_j(\alpha) + \delta \sum_y \pi(y|\alpha)w_j(y) \\ &= (1 - \delta)g_j(\alpha) - [\beta_i(v_i - (1 - \delta)g_i(\alpha)) - \delta c] / \beta_j. \end{aligned}$$

Hence, for any choice of v_i , (5.1)–(5.2) can be rewritten as

$$(5.3) \quad \sum \pi(y|a_i, \alpha_{-i}) w_i(y) = (v_i - (1 - \delta) g_i(a_i, \alpha_{-i})) / \delta, \quad a_i \in A_i,$$

$$(5.4) \quad \sum \pi(y|a_j, \alpha_{-j}) w_i(y) = \frac{1}{\delta} \left[v_i - (1 - \delta) g_i(\alpha) - \frac{\beta_j(1 - \delta)(g_j(\alpha) - g_j(a_j, \alpha_{-j}))}{\beta_i} \right],$$

$a_j \in A_j.$

Now, multiply (5.3) by $\alpha_i(a_i)$ and sum over a_i . Subtract from this total the sum over a_j of (5.4) multiplied by $\alpha_j(a_j)$. The result is zero. Hence, although formulae (5.3) and (5.4) comprise $m_i + m_j$ equations, their rank is only $m_i + m_j - 1$, and so the full rank hypothesis enables us to solve for $\{w_i(y)\}$. *Q.E.D.*

Pairwise full rank is actually the conjunction of two weaker conditions. First, it obviously implies individual full rank. Second, it ensures that deviations by players i and j are distinct in the sense that they induce different probability distributions over public outcomes. The first condition implies that incentives can be designed to induce a player to choose a given action; the second that player i 's incentives can be designed without interfering with those of player j . This flexibility in incentive design is precisely what is required by enforceability with respect to regular hyperplanes. However, the two conditions are by no means necessary for such enforceability. In particular, we shall see in Sections 7–9 that there are prominent examples in which versions of the Folk Theorem obtain even when *individual* full rank is violated. The fact that individual full rank is not in general necessary motivates Lemma 5.5 below.

DEFINITION 5.5: Profile α is *pairwise-identifiable* for players i and j if the rank of matrix $\Pi_{ij}(\alpha)$ equals $\text{rank}(\Pi_i(\alpha_{-i})) + \text{rank}(\Pi_j(\alpha_{-j})) - 1$.

Pairwise-identifiability—the second implication of pairwise full rank above—implies that the distributions induced by linear combinations of player j 's deviations are distinct from those induced by combinations of player i 's deviations. Pairwise full rank is just pairwise-identifiability plus individual full rank.

LEMMA 5.5: *If a pure-action profile is enforceable and pairwise-identifiable for all pairs of players, then it is enforceable with respect to all regular hyperplanes.*

HEURISTIC PROOF: Without individual full rank, a player may have different actions that lead to the same probability distribution over public outcomes. If profile a is enforceable, however, then an action $a'_i \neq a_i$ leading to the same distribution over public outcomes as a_i does not increase player i 's payoff and hence can be ignored. The idea of the proof is to show that when ignored

actions are deleted, the matrix Π_{ij} corresponding to the remaining actions satisfies pairwise full rank, that is, it has rank equal to the number of rows minus one. Continuation payoffs chosen to ensure exact indifference among the remaining actions will then also deter deviations to the deleted ones. The details may be found in Appendix 2. *Q.E.D.*

REMARK: Notice that by replacing the individual full rank condition with enforceability we can no longer ensure strong enforceability on regular hyperplanes.

THEOREM 5.1: *If a profile α has pairwise full rank for every pair of players, or if it is an enforceable pure-action profile that is pairwise-identifiable for every pair of players, then α is enforceable with respect to all regular hyperplanes.*

In the partnership games of Section 3, each player has two actions, and so pairwise full rank demands that there be at least $2 + 2 - 1 = 3$ public outcomes. Thus, no profile has pairwise full rank in the two-outcome game, which implies that profiles in which both players work with positive probability cannot be enforced with respect to hyperplanes with negative slope. As we observed in Section 3, this means that, in a repeated-game equilibrium, either both players are “rewarded” or they are both “punished” in response to a public outcome. In contrast, generic profiles have pairwise full rank in the three-outcome game and, consequently, they are enforceable with respect to any regular hyperplane.

6. THE FOLK THEOREM

In this section we apply the identifiability conditions from Section 5 together with the results about decomposability on tangent hyperplanes from Section 4 to obtain analogs of the Folk Theorem. A point v on the boundary of a smooth set W can be decomposed according to Definition 4.4 provided that the tangent hyperplane P_v to W at v is *regular* and there exists a profile α for which $g(\alpha)$ is separated from W by P_v and which satisfies suitable identifiability hypotheses such as pairwise full rank (or the conjunction of pairwise-identifiability and enforceability). If P_v is not regular, then, for some i , v either maximizes or minimizes player i 's payoff on the set W . In the former case, v is decomposable with respect to a translate of P_v and profile a^i such that $a^i \in \arg \max_a g_i(a)$. If v minimizes player i 's payoff on W , finding a suitable profile α that is enforceable with respect to (a translate of) P_v may be more difficult. If the points in W Pareto-dominate the payoffs from a static Nash equilibrium α^e , then we can simply take $\alpha = \alpha^e$, since α^e is enforceable with respect to *any* hyperplane. That is why we establish a “Nash-threat” version of the Folk Theorem (Theorem 6.1). The stronger “minimax-threat” version requires an additional assumption (see Condition 6.3 below) to ensure that there exists α that is both (i) near a profile which minimaxes player i 's payoff, and (ii) enforceable with respect to P_v .

For the Nash-threat Folk Theorem, either of two hypotheses (Condition 6.1 or 6.2) suffices.

CONDITION 6.1: *Every pure-action, Pareto-efficient profile is pairwise-identifiable for all pairs of players.*

The next lemma shows that all Pareto-efficient profiles are enforceable. Hence, when Condition 6.1 holds, pure-action, Pareto-efficient profiles are enforceable with respect to all regular hyperplanes.

LEMMA 6.1: *Any Pareto-efficient profile is enforceable.*

PROOF: If α is Pareto efficient and $v = g(\alpha)$, then (because action spaces are finite) there are positive weights $\{\lambda_i\}$ such that v solves $\max_{v' \in V} \sum \lambda_i v'_i$. For all i , set

$$w_i(y) = [(1 - \delta)/\delta] \sum_{j \neq i} [\lambda_j/\lambda_i] r_j(\alpha_j, y).$$

Then for any α'_i

$$\begin{aligned} (6.1) \quad & (1 - \delta) g_i(\alpha'_i, \alpha_{-i}) + \delta \sum_y \pi(y|\alpha'_i, \alpha_{-i}) w_i(y) \\ &= [(1 - \delta)/\lambda_i] \lambda_i g_i(\alpha'_i, \alpha_{-i}) \\ & \quad + [(1 - \delta)/\lambda_i] \sum_{j \neq i} \sum_{y \in Y} \lambda_j \pi(y|\alpha'_i, \alpha_{-i}) r_j(\alpha_j, y) \\ &= [(1 - \delta)/\lambda_i] \sum_{j=1}^n \lambda_j g_j(\alpha'_i, \alpha_{-i}). \end{aligned}$$

Because $g(\alpha) = v$, α solves $\max \sum \lambda_i g_i(\alpha')$, and thus, for all i , α_i solves $\max \sum \lambda_j g_j(\alpha'_i, \alpha_{-i})$. Therefore $\alpha'_i = \alpha_i$ maximizes (6.1), implying that α is enforceable with respect to $\{w_i(y)\}$. *Q.E.D.*

REMARK: This is the only preliminary result that invokes the assumption that a player's payoff depends just on his own action and the public outcome: $r_i = r_i(a_i, y)$. Hence, our versions of the Folk Theorem that do not rely on this lemma (see Theorem 6.4 and the parts of Theorems 6.2 and 6.3 that invoke Condition 6.2) hold for the more general case in which $r_i = r_i(a, y)$.

In Section 7 we show that pure-action profiles are pairwise-identifiable for all pairs of players in any game having a "product structure," and Sections 8 and 9 provide some prominent examples. Nevertheless, there are many games of interest in which not all pure-action profiles are pairwise-identifiable. Specifically, in symmetric games for which the distribution over public outcomes is invariant to permutations of the components of an action profile, any profile in

which all players use the same action cannot be pairwise-identifiable. Such games include the Green and Porter (1984) oligopoly model and the Radner (1986) partnership model.²¹ Accordingly, we provide an alternative condition that is satisfied by many of these games.

CONDITION 6.2: *For all pairs i, j , there exists a profile α that has pairwise full rank for that pair.*

Condition 6.2 is satisfied by generic games in which the number of public outcomes exceeds $\max_{i,j}\{m_i + m_j - 1\}$.²² Note that the condition is much weaker than the requirement that all profiles have pairwise full rank. Indeed, in the oligopoly and partnership games of the preceding paragraph (and in the example of footnote 21) profiles in which all players use the same action are not even pairwise-identifiable and so *a fortiori* do not have pairwise full rank. Nevertheless, such games may well have *asymmetric* profiles with pairwise full rank, as we shall see below. This explains why the Folk Theorem applies to these games, but the set of payoffs for *symmetric* equilibria is bounded away from efficiency.

Condition 6.2 permits a different profile α for each pair of players i and j ; superficially, it does not impose the stronger requirement that a single profile have pairwise full rank for all pairs. However, the condition implies that this stronger property holds for an open and dense set of profiles.

LEMMA 6.2: *If Condition 6.2 holds, then there exists an open and dense set of profiles each of which has pairwise full rank for all pairs of players.*

PROOF: Fix a profile α' and a pair of players $i \neq j$, and let α^{ij} be a profile having pairwise full rank for i and j . For each $x \in [0, 1]$, define $\alpha(x)$ by $\alpha_k(x) = (1-x)\alpha'_k + x\alpha_k^{ij}$ for all k . Then, the determinant of any $m_i + m_j - 1$ square submatrix of $\Pi_{ij}(\alpha(x))$ is a polynomial in x , and since $\Pi_{ij}(\alpha(1)) = \Pi_{ij}(\alpha^{ij})$ has full rank, there exists such a submatrix of $\Pi_{ij}(\alpha(1))$ whose determinant is not zero. Now, any polynomial is either identically zero or else equals zero only on a finite set. Thus the determinant of the corresponding submatrix of $\Pi_{ij}(\alpha(x))$ is nonzero for almost all x , and since there is only a finite number of

²¹ For another example, suppose that we symmetrize the second partnership game in Section 3 so that if one player works and the other shirks, the probabilities of public outcomes 12, 8, and 0 are $\frac{1}{3}$, 0, and $\frac{2}{3}$, respectively, regardless of the players' identities (the model is otherwise unchanged from before). Then, deviations from the (work, work) profile will give rise to distributions over public outcomes that do not depend on the identity of the deviator, and so pairwise-identifiability fails. Note that the invariance entailed here is stronger than that implied by symmetry alone; we are supposing that permuting actions keeps the distribution on outcomes unchanged. A symmetric game with *observable* actions, for example, does not satisfy this invariance condition, since, in that case the public outcome is the action profile itself.

²² More precisely, suppose that we fix the sets Y and A , and the functions $r_i: A_i \times Y \rightarrow \mathbb{R}$, $i = 1, \dots, n$. Then Condition 6.1 is satisfied for all distributions $\pi(\cdot | \cdot)$ in an open and dense subset of the set of all probability distributions on Y .

submatrices, there exists a sequence of profiles converging to α' all of which have pairwise full rank for i and j .²³

Since (i) α' was arbitrary and (ii) if $\Pi_{ij}(\alpha)$ has full rank then, for all profiles α' in an open neighborhood of α , $\Pi_{ij}(\alpha')$ has full rank, we conclude there is an open, dense set of profiles having pairwise full rank for i and j . As this is true for each pair (i, j) and the intersection of open dense sets is open and dense, we conclude that there is an open dense set of profiles having pairwise full rank for all pairs (i, j) . *Q.E.D.*

THEOREM 6.1 (Nash-threat Folk Theorem): *Suppose that either Condition 6.1 or 6.2 holds. Consider a Nash equilibrium α^e of the stage game. Let V^0 be the convex hull of the set consisting of $g(\alpha^e)$ and the Pareto-efficient payoff vectors Pareto-dominating $g(\alpha^e)$. Let W be a smooth set in the interior of V^0 . Then there exists $\underline{\delta} < 1$ such that, for all $\delta > \underline{\delta}$, $W \subseteq E(\delta)$.*

PROOF: Thanks to Theorem 4.1, it suffices to show that W is decomposable on tangent hyperplanes. Consider v on the boundary of W , and let P_v be the hyperplane tangent to W at v . Suppose first that P_v is orthogonal to the i th axis. Then v either maximizes or minimizes player i 's payoff on the set W . Lemmas 5.2 and 6.1 imply that, in the former case, a^i (an extremal profile that maximizes player i 's payoff on A) is enforceable with respect to P_v and that, in the latter case, the same is true of α^e . Moreover in either case the profile in question is separated from W by P_v . (In the latter case this is so because v Pareto-dominates $g(\alpha^e)$.)

Next, suppose that P_v is a regular hyperplane. Assume first that Condition 6.1 holds. Because W is contained in the convex hull of the set comprising $g(\alpha^e)$ and the Pareto-optimal points, there exists profile α , with $g(\alpha)$ separated from W by P_v such that either (a) α is a pure action Pareto optimum or (b) $\alpha = \alpha^e$. In case (a), Lemma 6.1 and Theorem 5.1 imply that α is enforceable with respect to P_v , and part (iii) of Lemma 5.2 implies the same thing in case (b). Hence, W is indeed decomposable on tangent hyperplanes in this case. If instead Condition 6.2 holds, then Lemma 6.2 and Theorem 5.1 ensure that there exists α such that $g(\alpha)$ is separated from W by P_v and α is enforceable with respect to P_v . *Q.E.D.*

To obtain the minimax version of the Folk Theorem, we require, in addition to the hypotheses of Theorem 6.1, that all pure-action profiles have individual full rank. We use this assumption to show that the minimax profile against each player i can be approximated by an enforceable profile α that has the best-response property for player i . (This enables us to conclude that α is enforceable with respect to a hyperplane that is orthogonal with respect to the i th axis.)

²³ We thank Andreu Mas-Colell for suggesting this argument using polynomials.

CONDITION 6.3: *Any pure action profile has individual full rank.*

LEMMA 6.3: *Suppose that Condition 6.3 holds. Then, for any $\varepsilon > 0$ and any i , there exists an enforceable profile α having the best response property for i and such that $|g_i(\alpha) - v_i| < \varepsilon$, where v_i is player i 's minimax payoff.*

PROOF: For each $a_i \in A_i$, let $C(a_i)$ be the set of $\alpha_{-i} \in A_{-i}$ such that (a_i, α_{-i}) has individual full rank for every player $j \neq i$. By Condition 6.3 $C(a_i)$ is not empty; by an argument essentially the same as that in Lemma 6.1, it follows that $C(a_i)$ is open and dense in A_{-i} . Consequently $C \equiv \bigcap_{a_i \in A_i} C(a_i)$ is also open and dense in A_{-i} .

Now choose a minimax profile $\underline{\alpha}^i$ against player i . We can find a sequence $\{\alpha_{-i}^k\}_{k=1}^\infty$ converging to $\underline{\alpha}_{-i}^i$ with $\alpha_{-i}^k \in C$ for all k , since C is dense in A_{-i} . Let a_i^k be a best response to α_{-i}^k . By definition of C , it follows that (a_i^k, α_{-i}^k) has individual full rank for every player $j \neq i$. Choose a_i and a subsequence of $\{(a_i^k, \alpha_{-i}^k)\}$ such that $a_i^k = a_i$ for all k . Since $\alpha_{-i}^k \rightarrow \underline{\alpha}_{-i}^i$, a_i is also a best response to $\underline{\alpha}_{-i}^i$, and so $g_i(a_i, \underline{\alpha}_{-i}^i) = v_i$. Consequently, for any $\varepsilon > 0$, we can choose k so that $|g_i(a_i, \alpha_{-i}^k) - v_i| < \varepsilon$. Because (a_i, α_{-i}^k) has the best response property for i and individual full rank for $j \neq i$, it is enforceable. *Q.E.D.*

REMARK: Notice that Condition 6.3 is used only to show that $C(a_i)$ is nonempty for each a_i , that is, for every $a_i \in A_i$ there exists $\alpha_{-i} \in A_{-i}$ such that (a_i, α_{-i}) has individual full rank for all $j \neq i$. This weaker condition may be used in place of Condition 6.3. Indeed inspection of the proof shows that an even weaker condition suffices: instead of every $a_i \in A_i$, we may restrict attention to any subset $A'_i \subseteq A_i$ for which there exists a minimax profile $\underline{\alpha}^i$, such that if a_i is a best response to $\underline{\alpha}_{-i}^i$, $a_i \in A'_i$.

THEOREM 6.2 (Folk Theorem): *Suppose that Condition 6.3 and either Condition 6.1 or 6.2 hold. Let W be a smooth subset in the interior of V^* . Then there exists $\underline{\delta} < 1$ such that, for all $\delta \geq \underline{\delta}$, $W \subseteq E(\delta)$.²⁴*

PROOF: The proof is the same as that for Theorem 6.1 except that now, if v minimizes player i 's payoff on W , we invoke Lemma 6.3 to conclude that there exists an enforceable profile α having the best-response property for i and for which $g_i(\alpha)$ is near enough v_i to ensure that $g(\alpha)$ is separated from W by P_v . Hence, because P_v is orthogonal to the i th axis, α is enforceable with respect to P_v . *Q.E.D.*

²⁴ In view of the remark following Lemma 6.1, this result extends to the more general case in which a player's payoff may depend on others' actions as well as his own, provided that Condition 6.2 holds. In the proof of Theorem 6.1, we required Lemma 6.1 (and hence the hypothesis that r_i depends only on a_i) in order to conclude that a^i is enforceable. In the proof of Theorem 6.2, however, we can invoke Condition 6.3 and Lemma 5.1 to establish a^i 's enforceability.

Consider the partnership examples from Section 3. We saw in that section that the two-outcome version has too few public outcomes relative to the size of the players' action sets for any action profile to satisfy pairwise full rank. (Indeed, profiles fail even to be pairwise identifiable.) This feature—which also pertains to the example in Radner, Myerson, and Maskin (1986)—helps explain why repeated game equilibria are bounded away from efficiency in this model.

In contrast, we saw that, in the three-outcome version, the profile (work, work) has pairwise full rank, and so Condition 6.2 holds. Moreover, it is easy to verify that Condition 6.3 is satisfied as well, and so Theorem 6.2 also applies to this game.

In this latter partnership example, continuation payoffs in the repeated game provide the incentives that induce partners to work. If, however, transfers of utility between partners are possible, we can reinterpret the continuation payoffs as such transfers in the static (one-shot) version of the game.²⁵ More specifically, consider functions $\mu_i : Y \rightarrow \mathbb{R}$, $i = 1, \dots, n$, where $\mu_i(y)$ is a transfer of utility to player i contingent on public outcome y . Then, if transfers $\{\mu_i(\cdot)\}_{i=1}^n$ are made and the public outcome is y , player i 's total payoff is $r_i(a_i, y) + \mu_i(y)$. We say that the transfers $\{\mu_i(\cdot)\}_{i=1}^n$ are *balanced* if $\sum_{i=1}^n \mu_i(y) = 0$ for all y .

THEOREM 6.3 (One-shot Folk Theorem): *If a game satisfies Condition 6.1, then, for any pure-action Pareto-efficient profile a , there exist balanced transfers $\{\mu_i(\cdot)\}_{i=1}^n$ for which a is a Nash equilibrium in the one-shot game with transfers. If instead the game satisfies Condition 6.2, then the same conclusion pertain to all profiles α in an open and dense set.²⁶*

PROOF: Finding balanced transfers that induce players to choose the desired profile is equivalent to showing that this profile is enforceable with respect to the regular hyperplane $\{v \mid \sum v_i = 0\}$. *Q.E.D.*

Just as Theorems 6.1–6.3 apply to generic partnership games, provided that there are “enough” observable outcomes, they apply as well to the version of the Green and Porter (1984) oligopoly model discussed in Section 2. (As typically formulated, the Green-Porter model has a continuum of possible outcomes.) In their analysis of this model, Abreu, Pearce, and Stacchetti (1986) concentrated on symmetric equilibria, i.e., those in which all players use the same strategy. As we noted above, however, symmetric profiles are *not* pairwise-identifiable, and so it turns out that symmetric equilibria are bounded away from efficiency. Nevertheless, Theorem 6.2 implies that (near) efficiency is restored once one admits asymmetric equilibria.

²⁵ Indeed, as we will see, the analysis is a good deal easier in the one-shot case since perfection issues do not arise.

²⁶ This theorem is a generalization of results due to Legros (1988) and Williams and Radner (1988).

We can modify the proof of Lemma 5.4 to obtain a sharper version of Theorem 6.2 for some games. Call a perfect public equilibrium *strict* if each player strictly prefers his equilibrium strategy to any other strategy.

THEOREM 6.4: *Suppose that every pure action profile a has pairwise full rank for all pairs of players and that $\pi(y|a) > 0$ for all y . Assume also that, for each player, there is a minimax profile in pure actions. Then, for every smooth set W in the interior of V^* there exists $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$, each point in W corresponds to a strict perfect public equilibrium with discount factor δ .*

REMARK: Even if all minimax profiles entail mixed actions, we can obtain as strict equilibria all interior, feasible payoff-vectors that Pareto-dominate $(\underline{v}_1, \dots, \underline{v}_n)$, where $\underline{v}_i = \min_{a_{-i}} \max_{a_i} g_i(a_i, a_{-i})$, as long as the other hypotheses of Theorem 6.4 hold.

PROOF: At any pure action profile, the inequalities (4.1) can be solved to hold strictly for any deviation. When all outcomes have positive probability, any deviant strategy then induces a strict decrease in payoff if the continuation payoffs are as specified. *Q.E.D.*

In finite games, strict Nash equilibria (those for which any unilateral deviation leaves the deviator strictly worse off) are stable as singleton sets in the sense of Kohlberg-Mertens (1986). Although there is not yet an established definition of stability for infinite-horizon games, it seems likely that strict equilibria would satisfy any such concept. Thus, Theorem 6.4 casts doubt on conjectures that stability might restrict the force of the Folk Theorem. Moreover, strict equilibria are robust to small changes in the information structure of the game.

To obtain strict equilibria requires both that players do not randomize and that any finite sequence of public outcomes occurs with positive probability in equilibrium. The two requirements are clearly inconsistent in repeated games with observable actions, where outcomes are the same as actions. However, in such games we can solve (4.1) with exact equality for all actions, and thus induce players to play all their actions with positive probability in every period. We conclude that in such games every feasible, interior, individually rational payoff vector can be attained by a totally mixed equilibrium. This observation may be pertinent because totally mixed equilibria, like strict ones, are stable in finite games. (However, it is less clear whether totally mixed equilibria will satisfy natural extensions of stability to infinite-horizon games.)

Theorems 6.1, 6.2, and 6.4 apply only to interior points and so do not pertain to payoffs on the efficient frontier. This contrasts with the standard Folk Theorem for observable actions, in which efficient payoffs *can* be exactly attained. The following result, related to a finding of Madrigale (1986), explains why imperfect information can preclude exact efficiency.

THEOREM 6.5: *Let v be an extremal payoff vector of V . If for each pure action profile a with $g(a) = v$ there exist player i and action a'_i such that $g_i(a'_i, a_{-i}) > g_i(a)$ and support $\pi(\cdot | a'_i, a_{-i}) \subseteq \text{support } \pi(\cdot | a)$, then $v \notin E(\delta)$ for all $\delta < 1$.*

REMARK: The hypothesis on supports is clearly satisfied if the support of $\pi(\cdot | a)$ is independent of a . When the hypothesis fails, player i may be deterred from playing a'_i by continuation payoffs that are very low outside support of $\pi(\cdot | a)$.²⁷

PROOF: Because v is extremal, the only sequence of feasible vectors with average value v is the sequence in which v is repeated each period. Consider an equilibrium giving rise to this sequence. Then the first period strategies must specify a profile a with $g(a) = v$, and the continuation payoffs $\{w(y)\}$ must satisfy $w(y) = v$ for all $y \in \text{support } \pi(\cdot | a)$. Since $g_i(a'_i, a_{-i}) > g_i(a)$, and support $\pi(\cdot | a'_i, a_{-i}) \subseteq \text{support } \pi(\cdot | a)$, player i prefers a'_i to a_i . *Q.E.D.*

As we noted in the introduction, our methods in this paper differ markedly from those of Radner (1985), Rubinstein (1979), and Rubinstein and Yaari (1983). These authors studied principal-agent relationships in which the agent's actions are not directly observed by the principal. They showed that, when the discount factor is near enough 1, it is possible to choose a "review phase" consisting of a certain number of periods of play such that (i) if the agent conforms to efficient play throughout the phase, the "average" publicly observable outcome will be within a given tolerance level of a specified target with very high probability, and (ii) if the agent deviates from efficient play throughout the phase, the average outcome will be outside the tolerance level with high probability.

When there is zero discounting, there exists an equilibrium where the agent plays efficiently in every period for fear that if he did not, the average outcome would fall outside the tolerance level, in which case play would revert to a one-shot Nash equilibrium thereafter. When there is positive discounting, there no longer exists an equilibrium in which the agent plays efficiently in every period of the review phase. Nonetheless, Radner (1985) established that, when the discount factor is sufficiently high, there exists an equilibrium in which the agent plays efficiently in a high fraction of periods. Radner's proof did not explicitly construct the players' equilibrium strategies, which may be quite complicated; the techniques we sketched following Lemma 4.2 are one way to derive them.

There appear to be two stumbling blocks to applying the Radner-Rubinstein-Yaari techniques to the more general class of games considered in this paper. First, it is not clear how to handle the case in which there are several players whose actions are imperfectly observed. When there is only one such player, the

²⁷ If the support of $\pi(\cdot | a)$ is independent of a , then an efficient repeated-game equilibrium in a two-player game is *strongly renegotiation-proof* in the sense of Farrell and Maskin (1989).

review phase can be constructed so that that player does not “cheat” very much when the discount factor is high. But this construction does not seem to extend readily. Indeed, we know from examples as in Radner, Myerson, and Maskin (1986) that, with more than one imperfectly observed player, the construction *cannot* be extended without additional hypotheses (such as our identifiability conditions). But it is not apparent how to make use of these conditions within the Rubinstein-Radner-Yaari framework.

Second, it is not evident how to handle punishments other than reversion to a one-shot equilibrium using the Radner-Rubinstein-Yaari approach. Getting players to minimax a given player in the absence of a common punishment, for example, may be problematic.

7. GAMES WITH A PRODUCT STRUCTURE

Condition 6.2 in Theorems 6.1–6.4 required a profile with pairwise full rank for all pairs of players. Radner (1985), however, demonstrated the existence of (nearly) efficient equilibria in repeated principal-agent games without the need for even individual full rank. (Refer also to the principal-agent model in Section 3.) Such games have a special feature that can take the place of full-rank conditions: The distribution over public outcomes has a “product structure,” meaning that the outcome y is a vector (y_1, \dots, y_n) , in which the components are statistically independent and the distribution of each component y_i depends only on the actions of player i . Theorem 7.1 below shows that this feature is sufficient for the Nash-threat Folk Theorem.

We first demonstrate that in games with a product structure, pure-action profiles always satisfy pairwise identifiability. This means that Condition 6.1 holds, and so Theorem 6.1 obtains.

DEFINITION 7.1: A game has a *product structure* if we can write each public outcome y as $y = (y_1, \dots, y_n)$, where, for all i and $\alpha = (\alpha_i, \alpha_{-i})$, the marginal distribution of y_i , $\pi_i(y_i|\alpha)$, satisfies

$$(7.1) \quad \begin{cases} \pi_i(y_i|\alpha) = \pi_i(y_i|\alpha_i, \alpha'_{-i}) & \text{for all } y_i \text{ and } \alpha'_{-i}, \\ \pi(y|\alpha) = \prod_{i=1}^n \pi_i(y_i|\alpha) & \text{for all } y. \end{cases}$$

The first half of (7.1) says that the marginal distribution of y_i depends on α_i alone, and the second that the joint distribution over y is the product of the marginals (that is, the y_i 's are independent given α). When the game has a product structure, one might expect that an enforceable profile should be enforceable with respect to continuation payoffs in which player i 's continuation payoff $w_i(y)$ depends only on y_i .

LEMMA 7.1: *If the game has a product structure, then every pure-action profile is pairwise-identifiable for all pairs of players.*

IDEA OF PROOF: Since pairwise-identifiability means that deviations by one player can be distinguished statistically from deviations by another, the fact that y_i contains all information about i 's play and y_j about j 's makes this fairly intuitive. A formal proof is provided in Appendix 3. Q.E.D.

THEOREM 7.1: *In a game with a product structure, let W be a smooth set in the interior of V^0 (see Theorem 6.1). Then there exists a discount factor $\underline{\delta}$ such that, for all $\delta > \underline{\delta}$, $W \subseteq E(\delta)$.*

PROOF: From Lemma 6.1 Pareto-efficient profiles are enforceable. From Lemma 7.1, any pure action profile is pairwise-identifiable for all pairs of players. Hence, the result follows from Theorem 6.1. Q.E.D.

One example of a game with a product structure is the principal-agent model considered by Radner (1985), in which the principal's only action is to propose a (publicly-observed) output-contingent transfer-payment function, and the probability distribution of each period's output depends only on the agent's effort. (In this case a public outcome is a pair consisting of the transfer-payment function and output; the former depends only on the principal, the latter only on the agent.) Radner showed that when the discount factor is sufficiently near 1 there exist (nearly) efficient repeated game equilibria that Pareto-dominate a static equilibrium. Theorem 7.1 extends Radner's results to a general specification of the payoff functions (for example, unlike in Radner (1985), the principal can be risk-averse, and the agent's utility function need not be separable in effort and income). Section 9 extends the conclusion of Theorem 7.1 to games in which (i) the principal can, in addition to making transfers, influence the distribution of output, and (ii) there are multiple agents.

8. ADVERSE SELECTION

We next apply Theorem 7.1 to a class of games of adverse selection. We then observe that this class includes standard mechanism-design problems.

Consider the following class of games. Each player i first receives private information determining his type $z_i \in Z_i$, where Z_i is finite. The distribution of types is independent across players, and player i 's type has distribution $\hat{\pi}_i$, which is common knowledge. After learning their types, players move simultaneously, and player i 's move, $y_i \in Y_i$, is publicly observable. The public outcome is $y = (y_1, \dots, y_n) \in Y = Y_1 \times \dots \times Y_n$.

A (pure) action for player i is a mapping a_i from Z_i to Y_i . If we suppose that player i 's realized payoff $\hat{r}_i(z_i, y)$ depends on the public outcome and his own type but not on others' types (i.e., the case of *private values*), we can express his *ex ante* payoff function over actions and outcomes (which pertains before he

learns his type) as

$$r_i(a_i, y) = \frac{\sum_{z_i \in a_i^{-1}(y)} \hat{r}_i(z_i, y) \hat{\pi}_i(z_i)}{\sum_{z_i \in a_i^{-1}(y)} \hat{\pi}_i(z_i)}.^{28}$$

Note that the distribution of outcomes corresponding to action profile a is

$$\pi(y|a) = \prod_{i=1}^n \pi_i(y_i|a_i),$$

where

$$\pi_i(y_i|a_i) = \begin{cases} \sum_{z_i \in a_i^{-1}(y_i)} \hat{\pi}_i(z_i), & \text{if } a_i^{-1}(y_i) \text{ is nonempty,} \\ 0, & \text{otherwise.} \end{cases}$$

Thus player i 's ex ante expected utility from the profile a is

$$g_i(a) = \sum_y \pi(y|a) r_i(a_i, y).$$

It is clear that the game has a product structure in the sense of Definition 7.1. In the repeated version of the game, player i 's types are drawn independently in each period from the fixed distribution $\hat{\pi}_i$. A player observes the public outcome y each period but not others' types.

Notice that these games do not satisfy the individual full rank condition, and *a fortiori* fail the stronger condition of pairwise full rank: Let z'_1 and z''_1 satisfy $\hat{\pi}_1(z'_1) \geq \hat{\pi}_1(z''_1)$, and fix a pure action a_1 with $a_1(z'_1) = y'_1$, $a_1(z''_1) = y''_1$. This action yields the same distribution over outcomes as (mixed) action \hat{a}_1 , where $\hat{a}_1(z'_1) = y'_1$, and $\hat{a}_1(z'_1)$ takes on the value y''_1 with probability $\pi_1(z''_1)/\pi_1(z'_1)$, and y'_1 with probability $(\pi_1(z'_1) - \pi_1(z''_1))/\pi_1(z'_1)$. This lack of full rank means that many action profiles cannot be enforced. Nevertheless, because the game has a product structure, Theorem 7.1 applies, and so the Nash-threat version of the Folk Theorem obtains.

Examples of games in this class include social-choice mechanisms under adverse selection. In the repeated version, player i 's move y_i in period t is a report $a_i(z_i)$ (not necessarily truthful) of his period t type z_i , that is, Y is isomorphic to Z . Let S be a finite set of social alternatives and fix a social choice function $f: Z \rightarrow S$. Consider the mechanism g in which, when the reported profile of types is y , the social alternative selected is $f(y)$, so that player i 's payoff is $\hat{r}_i(z_i, y) = u_i(f(y), z_i)$. We say that the repeated version of g implements f if there exists a PPE in which the expected discounted average payoffs are $g(a^*)$, where $a^* = (a_1^*, \dots, a_n^*)$ is the "truthful reporting" profile (i.e., for all i and z_i , $a_i^*(z_i) = z_i$).

One such repeated adverse-selection problem has been considered by Green (1987). The types z_i are the players' (privately observed) endowments, and the

²⁸ Formally, $r_i(a_i, y)$ is not defined for y such that $a_i^{-1}(y_i)$ is empty. However, because such a public outcome occurs with zero probability, we can define $r_i(a_i, y)$ arbitrarily without affecting the analysis.

social choice function f redistributes endowments to provide insurance against fluctuations. Bewley (1983), and Levine (1991), and Scheinkman and Weiss (1988) studied the role of money in related models.

THEOREM 8.1: *Fix a social choice function f and a Nash equilibrium α^e of the corresponding mechanism g , and suppose that the distribution of types is independent across players. Let V^0 be the set formed by taking the convex hull of $g(\alpha^e)$ and the feasible points that Pareto-dominate it. If V^0 has a nonempty interior, then all payoff vectors in V^0 can be approximated by equilibria of the repeated mechanism for discount factors close enough to 1. (That is, for $v \in V^0$ and any ε there is a $\underline{\delta} < 1$ such that there is an equilibrium payoff within ε of v for all $\delta > \underline{\delta}$.) In particular, if $g(a^*)$ Pareto-dominates $g(\alpha^e)$, then g (approximately) implements f for all δ sufficiently near 1.*

PROOF: This follows from Theorem 7.1.

Q.E.D.

To help interpret this result, consider a variant of Green's (1987) insurance model. Each of two agents, $i = 1, 2$, is endowed with 2 units of a consumption good, and is subject to a preference shock z_i whose realization can be 2 or 4 equally likely, and which is independently distributed across agents. Agent i 's utility of consumption is $\ln(z_i + c_i)$, where the level of consumption c_i must be an integer and the sum of consumptions cannot exceed the sum of endowments. The best symmetric allocation is for agent i to consume his endowment when the shocks to each agent are the same, and to give up one unit to the other consumer when $z_i = 4$ and shocks differ. Theorem 8.1 implies that the corresponding expected utility $\frac{1}{4} \ln 4 + \frac{1}{2} \ln 5 + \frac{1}{4} \ln 6$ can be approximated by a repeated game equilibrium when δ is near 1. The conclusion is not that risk or risk aversion *per se* becomes unimportant when δ nears 1. Indeed, the best symmetric allocation does not approximate the utilities that accrue when each player is perfectly insured, i.e., when $z_i + c_i = 5$ each period. That is, repeated play does not reduce the *social risk* corresponding to random preferences. Rather, it reduces the incentive cost of inducing players to report their types truthfully.

Just as Theorem 6.3 establishes a one-shot counterpart to the Folk Theorem in the case of moral hazard where utility transfers take the place of continuation payoffs in ensuring incentives, we can obtain such a one-shot counterpart in the case of adverse selection. In our (1993) working paper, we show that generalizations of many of the results from the transferable-utility mechanism-design literature can readily be derived using our methods.

9. PRINCIPAL-AGENT MODELS

We next consider the class of games in which at least one player's actions are observable. This class includes the standard principal-agent model, which we

discussed briefly in Sections 3 and 7. We first generalize this model to consider all two-player games in which one player's actions are observable.

We say that player 2's actions are *observable* if, for all a_1 , $\text{support } \pi(\cdot | a_1, a_2) \cap \text{support } \pi(\cdot | a_1, a'_2) = \emptyset$ for $a_2 \neq a'_2$. Thus, given player 1's action, any observable outcome is consistent with only one action by player 2.

As in Theorem 7.1, we shall be concerned with V^0 , the convex hull of the set consisting of a static equilibrium payoff $v^e = g(\alpha^e)$ and the Pareto efficient points that Pareto-dominate it.

THEOREM 9.1: *In a two-player game in which player 2's actions are observable, let W be a smooth set in the interior of V^0 . Then there exists $\underline{\delta}$ such that, for all $\delta > \underline{\delta}$, $W \subseteq E(\delta)$.*

PROOF: Because the public outcome reveals player 2's actions, the continuation payoffs corresponding to deviations by player 2 have probability zero so long as player 2 never deviates. Consider v on the boundary of W . Inspection of the proof of Theorem 4.1 reveals that it suffices to show that there exists a profile α , with $g(\alpha)$ separated from W by the tangent hyperplane P_v to W at v , such that (i) α is enforceable with respect to continuation payoffs $\{w(y)\}_{y \in Y}$ and (ii) $w(y) \in P_v$ if $y \in \bigcup_{a_1} \text{support } \pi(\cdot | a_1, \alpha_2) (\equiv \gamma(\alpha_2))$ and $w(y) = v^e$ otherwise. The argument that there exists such a profile is virtually the same as in the proof of Theorem 6.1. *Q.E.D.*

A boundary point of V is on the *outer boundary* with respect to the one-shot equilibrium payoff vector v^e if it strictly Pareto-dominates v^e . The outer boundary is *downward-sloping* if it consists only of (strict) Pareto optima. An immediate implication of Theorem 9.1 is the following:

COROLLARY 9.1: *If, in addition to the hypothesis of Theorem 9.1, the outer boundary of V is downward-sloping, then for any vector v in the interior of V^* that Pareto dominates a static Nash equilibrium there exists $\underline{\delta}$ such that, for all $\delta > \underline{\delta}$, $v \in E(\delta)$.*

To see what can go wrong if the downward-sloping hypothesis in Corollary 9.1 fails, consider the game described in Table 9.1. Player 1 has three (unobservable) actions, U , M , D , and Player 2 has two observable actions L and R . If player 2 chooses action R , output (which accrues to Player 2) can either be 5 (high) or 0 (low). If he chooses L , output is always low. If output is high, Player 1's utility is an increasing function of his effort; if output is low, utility is zero. A public outcome corresponds to an action by Player 2 and an output level; hence there are four possible outcomes: (L, high) , (L, low) , (R, high) , (R, low) . The feasible payoffs are illustrated in Figure 9.2.

Notice that there is a static equilibrium (U, L) with expected payoffs $(0, 0)$; these are also the minimax payoffs. Hence, if the line connecting $(1, 2)$ and $(3, 3)$ in Figure 9.2 were downward-sloping, Corollary 9.1 would imply that all convex

TABLE 9.1

		Player 2				
		L	R			
Player 1	U	(0,0)	(0,0)	U	0	
	M	(0,0)	(1,5)		M	$\frac{2}{5}$
	D	(0,0)	(3,5)		D	$\frac{3}{5}$

		L	R
U	(0,0)	(0,0)	(0,0)
M	(0,0)	(0,0)	(1,2)
D	(0,0)	(0,0)	(3,3)

Payoffs (when the entry is divided, the upper pair refers to high output and the lower to low output)

Probability of High Output

Expected Payoffs

combinations of (0, 0), (1, 2), and (3, 3) could arise in repeated game equilibria. It can easily be shown, however, that only convex combinations of (0, 0) and (3, 3) correspond to equilibria (see Fudenberg and Maskin (1986b) for a demonstration).

Corollary 9.1 applies, in particular, to the principal-agent games discussed in Section 7, and it and Theorem 9.1 clearly extend to games of three or more players in which the actions of only one player are unobservable. In fact, they generalize to games in which several players' actions are unobservable, as long as the game has a product structure. Formally, we have the following theorem.

THEOREM 9.2: Consider an n -player game in which the actions of only players $l + 1, \dots, n$ are observable. Suppose that we can express each outcome y as $y = (y_1, \dots, y_l)$ where, for $i = 1, \dots, l$, the distribution of y_i depends only on a_i




FIGURE 9.2

and (a_{l+1}, \dots, a_n) , and $\pi(y|a) = \prod_{i=1}^l \pi_i(y_i|a)$. Then the conclusion of Theorem 9.1 obtains.

10. CONCLUDING REMARKS

Our method of constructing equilibria is based on the dynamic-programming decomposition of equilibrium payoffs into current payoffs $g(\alpha)$ and continuation payoffs $w(y)$, the point being that if the continuation payoffs are themselves in the equilibrium set, and profile α can be enforced by $w(\cdot)$, then the overall payoff vector $(1 - \delta)g(\alpha) + \delta \sum \pi(y|\alpha)w(y)$ corresponds to an equilibrium as well. This approach prompts the question of how the equilibrium set $E(\delta)$ is related to the profiles that are enforceable with respect to *some* continuation payoffs in the feasible set V , whether or not these latter payoffs themselves correspond to equilibria.

Let $C(\delta)$ consist of the overall payoff vectors corresponding to such profiles. (This set could be attained if players could commit themselves in advance to outcome-contingent continuation payoffs—say, through binding contracts.) Let $N(\delta)$ be the set of Nash equilibrium payoffs in the repeated game; Nash equilibrium requires that, for each y with positive probability under profile α , $w(y)$ be a Nash equilibrium payoff vector. Since the definition of $C(\delta)$ imposes fewer constraints on the continuation payoffs than that of $N(\delta)$, it is clear that $E(\delta) \subseteq N(\delta) \subseteq C(\delta)$. From the study of repeated games with observable actions we know that for a fixed discount factor δ the first inclusion can be strict, and brief reflection should convince the reader that the second inclusion can be strict as well.

Let us compare the behavior of these sets as $\delta \rightarrow 1$. The first thing to notice is that if δ is too small, many vectors in V may be infeasible, i.e., unattainable by any profile of repeated game strategies, whereas all points in V are feasible if $\delta > 1 - 1/d$, where d is the number of vertices of V (Sorin (1986)). This last result reflects the fact that one effect of repeated play with patient players is to allow “intertemporally transferable utility” in the interior of V , even if the stage game itself does not provide much flexibility for utility transfers. A second effect of repetition is to permit the gain from a one-period deviation to be offset by variations in continuation payoffs. This is the basis for the Folk Theorem in repeated games with observable actions.

Let $E(1)$, $N(1)$, and $C(1)$ denote the limiting values of the corresponding sets as $\delta \rightarrow 1$. (A point is in $E(1)$ if it is in $E(\delta)$ for all sufficiently large δ .) Theorem 6.2 asserts that $E(1) \supset \text{interior } V^*$; according to its Nash counterpart, $N(1) \supset \text{interior } V^*$. Clearly, a necessary condition for either result is $C(1) \supset \text{interior } V^*$. This inclusion obtains if every pure-action profile is enforceable (since we can then always use an extremal pure-action profile in the first period of a decomposition), and thus if each such profile satisfies individual full rank.

Even when all the pure action profiles are enforceable, the Folk Theorem may still fail. For example, in repeated games with observable actions, it does not apply to games of three or more players unless V^* satisfies an interiority

condition (see footnote 1). And as the Radner, Myerson, and Maskin (1986) example demonstrates, the Folk Theorem can fail even when $C(1) \supset \text{interior } V^* \neq \emptyset$, if actions are unobservable: Although individual full rank implies that any Pareto-efficient profile is enforceable with respect to nearby continuation payoffs when δ is near one, these continuation payoffs may not themselves correspond to equilibria without the stronger condition of pairwise identifiability.

To conclude, let us mention a limitation on the interpretation of our results. As we have formulated the repeated game, the statistical information conveyed by outcomes about actions in any given period does not change as the discount factor tends to 1. This feature is crucial to our versions of the Folk Theorem; as Kandori (1988) has shown, the equilibrium set contracts as public outcomes reveal less and less information. Thus, although it would be proper to interpret δ tending to 1 as corresponding to greater patience on the part of players, it might be misleading—as Abreu, Milgrom, and Pearce (1991) have stressed—to construe this convergence as a consequence of the time interval between periods growing shorter. Quite plausibly, the quality of the information revealed by outcomes is a function of period length: the longer is the period, the more information accumulates. Indeed, Abreu, Milgrom, and Pearce provide an example in which the Folk Theorem does not obtain as the time period shrinks because the informativeness of outcomes decays too quickly.

*Dept. of Economics, Harvard University, Cambridge, MA 02138, U.S.A.,
Dept. of Economics, University of California—Los Angeles, Los Angeles, CA
90024, U.S.A.,*

and

Dept. of Economics, Harvard University, Cambridge, MA 02138, U.S.A.

Manuscript received May, 1989; final revision received December, 1993.

APPENDIX 1

THEOREM 4.1: *If a smooth set W is decomposable on tangent hyperplanes, then it is locally self-decomposable.*

PROOF: It suffices to show that for all $v \in W$ there exists $\delta < 1$ and an open neighborhood U of v such that each $u \in U \cap W$ is decomposable with respect to some profile α , δ , and a subset of W .

Step 1: Suppose $v \in \text{interior } W$. Let α^e be an equilibrium of the stage game, and fix an open set U containing v whose closure is in the interior of W . Then there is a δ such that for each $u \in U$ there exists $u' \in W$ such that $u = (1 - \delta)g(\alpha^e) + \delta u'$. Because α^e is a static equilibrium, the incentive conditions (3.1) are satisfied, and so v is decomposable with respect to α^e , δ , and $\{u'\}$.

Step 2: Consider $v \in \text{boundary } W$. Since W is smooth, there is a unique hyperplane P_v tangent to W at v . We argued in the text that, from decomposability on tangent hyperplanes, there exist profile α and continuation payoff vectors lying on P_v such that $g(\alpha)$ is separated from W by P_v and the continuation payoffs enforce α . Lemma 4.3 then implies that there exists κ such that, for any δ and any $w' \in \mathbb{R}^n$, α is enforceable with respect to a translate of P_v containing w' , where the continuation payoffs $\{w(y)\}$ differ from w' by at most $\kappa(1 - \delta)/\delta$.

Step 3: It is now convenient to choose a new coordinate system. As shown in Figure 4.1, we take v to be the origin, choose the vertical (first) axis to be the line connecting v and $g(\alpha)$, and the remaining axes to lie in P_v . Since $g(\alpha) \notin P_v$, this transformation is possible. In the new coordinates, W is still smooth and convex. Although the norm is changed, κ can be replaced by a new constant κ' .

For $x \in \mathbb{R}^n$ we now write $x = (x^0, x^H)$, where x^0 is the component on the vertical axis and x^H the component in P_v . Because W is smooth, Taylor's theorem ensures that there exists $\delta^* < 1$, constant $c > 0$, and a neighborhood X of the origin such that, for all $\delta > \delta^*$, if $x \in X$, then $\|x^H\| < c\sqrt{(1-\delta)/\delta}$ and $x^0 \leq -\|g(\alpha)\|[(1-\delta)/\delta]$ imply $x \in \text{interior } W$.

In Step 2, let $w' = ((1-\delta)/\delta)(-\|g(\alpha)\|, 0)$. Then $v = (1-\delta)g(\alpha) + \delta w'$. Now each continuation payoff vector $w(y)$ has vertical component $w(y)^0 = -((1-\delta)/\delta)\|g(\alpha)\|$. Moreover, because, from Step 2, $\|w(y)^H\| < \kappa'((1-\delta)/\delta)$, $w(y) \in \text{interior } W$ for all δ greater than some δ^* (since $\kappa'((1-\delta)/\delta) < c\sqrt{(1-\delta)/\delta}$). It remains to show that for all v' in a neighborhood of v , v' is decomposable with respect to α , δ , and a subset of W . For any v' , $v' = (1-\delta)g(\alpha) + \delta \sum_y \pi(y|\alpha)w'(y)$, where $w'(y) = w(y) + (v' - v)/\delta$. Notice that for $\|v' - v\|$ small enough, $w'(y) \in \text{int } W$. Hence, there exist $\epsilon > 0$ and δ such that for all $v' \in W$ with $\|v' - v\| < \epsilon$, v' has the desired decomposability properties.

APPENDIX 2

LEMMA 5.5: *If a pure-action profile is enforceable and pairwise-identifiable for all pairs of players, then it is enforceable with respect to all regular hyperplanes.*

PROOF: Fix a pure-action profile a , and focus on player i . To simplify notation, for each $a'_i \in A_i$, define the vector $\Pi_i(a'_i) = \pi(\cdot | a'_i, a_{-i})$ and the payoff $\hat{g}_i(a'_i) \equiv g_i(a'_i, a_{-i})$. Call a'_i subordinate to set $A''_i \subset A_i$ if the vector $\Pi_i(a'_i)$ can be obtained as a linear combination of vectors corresponding to actions in A''_i , and this linear combination provides player i with at least the payoff from a'_i , that is, if there exist numbers $\{\lambda_i(a''_i)\}_{a''_i \in A''_i}$ such that $\lambda_i(a'_i) = 0$, and

$$(A2.1) \quad \sum_{a''_i \in A''_i} \lambda_i(a''_i) \Pi_i(a''_i) = \Pi_i(a'_i)$$

and

$$(A2.2) \quad \sum_{a''_i \in A''_i} \lambda_i(a''_i) g_i(a''_i) \geq g_i(a'_i).$$

Note that the weights $\lambda_i(a''_i)$ are not required to be nonnegative, and so may not correspond to a mixed action. Because the components of each vector $\Pi_i(a''_i)$ sum to 1, so, from (A2.1), must the weights sum to 1. Our goal is to reduce Lemma 5.5 to Lemma 5.4 by eliminating rows of $\Pi_{ij}(a)$ corresponding to subordinate actions.

We make three claims: (1) If, for some set A''_i , the vectors $\{\Pi_i(a''_i)\}_{a''_i \in A''_i}$ are linearly dependent, then there exists $a'_i \in A''_i$ such that a'_i is subordinate to A''_i . (2) If, moreover, there exists $a_i \in A''_i$ that is both enforceable and subordinate to A''_i , then it is not the only action subordinate to that set. (In the text, we discuss enforceability in terms of profiles; here we are abusing terminology slightly by referring to the enforceability of a component of a profile.) (3) If a'_i is subordinate to A''_i and a'_i is subordinate to $A''_i \setminus a'_i$, then a'_i is subordinate to $A''_i \setminus a'_i$. That is, if an action is subordinate to a set, it is subordinate to any subset formed by the elimination of subordinate actions.

First we establish (1). If $\{\Pi_i(a''_i)\}_{a''_i \in A''_i}$ are dependent, then there exist $\{\lambda_i(a''_i)\}$, not all zero, such that $\sum_{a''_i \in A''_i} \lambda_i(a''_i) \Pi_i(a''_i) = 0$. Choose a'_i with $\lambda_i(a'_i) \neq 0$. Then

$$\sum_{a''_i \in A''_i \setminus a'_i} [-\lambda_i(a''_i)/\lambda_i(a'_i)] \Pi_i(a''_i) = \Pi_i(a'_i).$$

Note that, since each $\Pi_i(a''_i)$ is an element of the unit simplex, $-\lambda_i(\hat{a}_i)/\lambda_i(a'_i) > 0$ for some $\hat{a}_i \in A''_i \setminus a'_i$. If

$$(A.2.3) \quad \sum_{a''_i \in A''_i \setminus a'_i} [-\lambda_i(a''_i)/\lambda_i(a'_i)] \hat{g}_i(a''_i) < \hat{g}_i(a'_i),$$

then, multiplying both sides by $\lambda_i(a'_i)/\lambda_i(\hat{a}_i)$, we get

$$\sum_{a''_i \in A''_i \setminus \hat{a}_i} [-\lambda_i(a''_i)/\lambda_i(\hat{a}_i)] \hat{g}_i(a''_i) > \hat{g}_i(\hat{a}_i)$$

and

$$\sum [-\lambda_i(a''_i)/\lambda_i(\hat{a}_i)] \Pi_i(a''_i) = \Pi_i(\hat{a}_i),$$

i.e., \hat{a}_i is subordinate to A''_i . (Of course, if the inequality (A.2.3) were reversed, then a'_i would be subordinate to A''_i .) This proves claim (1).

For claim (2), suppose that a_i is enforceable and subordinate to A''_i . Then for some weights $\{\lambda_i(a''_i)\}$ with $\lambda_i(a_i) = 0$, inequalities (A2.1) and (A2.2) must hold for $a'_i = a_i$. If (A2.2) holds strictly and $\lambda_i(a''_i) \geq 0$ for all a''_i , then the weights $\{\lambda_i(a''_i)\}$ correspond to a mixed action that yields a higher payoff than a_i , and so a_i is not enforceable, a contradiction. Thus, for some $\hat{a}_i \in A''_i \setminus a_i$, $\lambda_i(\hat{a}_i) < 0$. Proceeding as above, we find \hat{a}_i is subordinate to $A''_i \setminus \hat{a}_i$.

On the other hand, if (A2.2) holds with equality, choose $\hat{a}_i \in A''_i \setminus a_i$ such that $\lambda_i(\hat{a}_i) \neq 0$. We have

$$\sum_{a''_i \in A''_i \setminus \hat{a}_i} [-\lambda_i(a''_i)/\lambda_i(\hat{a}_i)] \hat{g}_i(a''_i) = \hat{g}_i(\hat{a}_i),$$

and so \hat{a}_i is subordinate to A''_i . This establishes the second claim.

Next, we prove (3). Suppose a'_i is subordinate to A''_i and a''_i is subordinate to $A''_i \setminus a'_i$. Then there exist weights $\{\lambda'_i(a_i)\}$ and $\{\lambda''_i(a_i)\}$ such that

$$(A2.4) \quad \sum_{a_i \in A''_i \setminus a'_i} \lambda'_i(a_i) \Pi_i(a_i) = \Pi_i(a'_i),$$

$$(A2.5) \quad \sum_{a_i \in A''_i \setminus \{a'_i, a''_i\}} \lambda''_i(a_i) \Pi_i(a_i) = \Pi_i(a''_i).$$

Now define $\hat{\lambda}_i(a_i) = \lambda'_i(a_i) + \lambda''_i(a_i)$ for $a_i \neq a'_i$ and $a_i \neq a''_i$; and $\hat{\lambda}_i(a'_i) = \hat{\lambda}_i(a''_i) = 0$. Substituting (A2.5) for $\Pi_i(a''_i)$ in (A2.4) shows that $\hat{\lambda}_i(\cdot)$ satisfies (A2.1): $\sum_{a_i \in A''_i} \hat{\lambda}_i(a_i) \Pi_i(a_i) = \Pi_i(a'_i)$. Similarly, we can verify that it satisfies (A2.2), so a'_i is subordinate to A''_i/a''_i . This proves claim (3).

It follows from the claims that if a is enforceable we can find a set A'_i containing a_i such that the vectors $\{\Pi_i(a'_i)\}_{a'_i \in A'_i}$ are independent and such that for all $a''_i \in A_i \setminus A'_i$, a''_i is subordinate to A'_i . Similarly, we can find a corresponding set A'_j for j . Let $\Pi'_i(a_{-i})$ and $\Pi'_j(a_{-j})$ be the corresponding matrices. By pairwise-identifiability and construction, the matrix

$$\begin{bmatrix} \Pi'_i(a_{-i}) \\ \Pi'_j(a_{-j}) \end{bmatrix}$$

has the same rank as $\Pi_{ij}(a)$. Lemmas 5.3 and 5.4 imply that for any regular hyperplane we can find $w_i(\cdot)$ and $w_j(\cdot)$ consistent with this hyperplane such that

$$(A2.6) \quad v_i = (1 - \delta) \hat{g}_i(a'_i) + \delta \sum_y \pi(y|a'_i, a_{-i}) w_i(y) \quad a'_i \in A'_i,$$

and similarly for j . From (A2.6) player i cannot gain by deviating to $a'_i \in A'_i$. If $a''_i \in A_i \setminus A'_i$, there exist weights such that

$$\Pi_i(a''_i) = \sum_{a'_i \in A'_i} \lambda_i(a'_i) \Pi_i(a'_i),$$

$$g_i(a''_i) \leq \sum_{a'_i \in A'_i} \lambda_i(a'_i) \hat{g}_i(a'_i),$$

and $\sum_{a'_i \in A'_i} \lambda_i(a'_i) = 1$. Thus,

$$\begin{aligned} v_i &= \sum_{a'_i \in A'_i} \lambda_i(a'_i) v_i \\ &= (1 - \delta) \sum_{a'_i \in A'_i} \lambda_i(a'_i) \hat{g}_i(a'_i) + \delta \sum_{a'_i \in A'_i} \lambda_i(a'_i) \sum_y \pi(y | a'_i, a_{-i}) w_i(y) \\ &\geq (1 - \delta) \hat{g}_i(a''_i) + \delta \sum_y \pi(y | a''_i, a_{-i}) w_i(y), \end{aligned}$$

so that (4.1) is satisfied for all $a''_i \in A_i$.

Q.E.D.

APPENDIX 3

LEMMA 7.1: *If a game has a product structure, then every pure-action profile is pairwise-identifiable for all pairs of players.*

PROOF: Let a be a pure-action profile and, for some pair (i, j) , choose $A'_i = \{a_i(1), \dots, a_i(k_i)\}$ and $A'_j = \{a_j(1), \dots, a_j(k_j)\}$ such that $B_i \equiv \{\pi(\cdot | a_i(s), a_{-i})\}_{s=1}^{k_i}$ is a basis for $\Pi_i(a_{-i})$ and $B_j \equiv \{\pi(\cdot | a_j(s), a_{-j})\}_{s=1}^{k_j}$ is a basis for $\Pi_j(a_{-j})$. Then, the ranks of $\Pi_i(a_{-i})$ and $\Pi_j(a_{-j})$ are k_i and k_j respectively. More specifically, choose A'_i to include a_i , and A'_j to include a_j and index so that $a_j(k_j) = a_j$. Now, as we observed in Section 5, $B_i \cup B_j$ is a linearly dependent set, and in particular for our choice of a_i and a_j $\pi(\cdot | a_i, a_{-i}) = \pi(\cdot | a_j, a_{-j})$. We now finish the proof by showing that the vectors corresponding to the set $A'_i \cup A'_j \setminus a_j$ are linearly independent, that is, that $\text{rank } \Pi_{ij}(a) = k_i + k_j - 1$.

Suppose not. That is, assume that there exist scalars $\{\lambda(s)\}$ and $\{\gamma(s)\}$, such that $\sum_{s=1}^{k_i} \lambda(s) \pi(y | a_i(s), a_{-i}) + \sum_{s=1}^{k_j-1} \gamma(s) \pi(y | a_j(s), a_{-j}) = 0$ for all y . Since B_i and B_j are each linearly independent we may assume at least one $\gamma(s)$ and at least one $\lambda(s)$ are nonzero. Because the game has a product structure, we can integrate out the components of y other than y_i and y_j to obtain

$$\begin{aligned} \text{(A3.1)} \quad & \sum_{s=1}^{k_i} \lambda(s) \pi_i(y_i | a_i(s)) \pi_j(y_j | a_j) \\ & + \sum_{s=1}^{k_j-1} \gamma(s) \pi_i(y_i | a_i) \pi_j(y_j | a_j(s)) = 0 \quad \text{for all } y_i \text{ and } y_j. \end{aligned}$$

For any y_i and y_j such that $\pi_i(y_i | a_i) > 0$ and $\pi_j(y_j | a_j) > 0$, we may divide (A3.1) by $\pi_i(y_i | a_i) \pi_j(y_j | a_j)$ to find

$$\text{(A3.2)} \quad \sum_{s=1}^{k_i} \lambda(s) \frac{\pi_i(y_i | a_i(s))}{\pi_i(y_i | a_i)} + \sum_{s=1}^{k_j-1} \gamma(s) \frac{\pi_j(y_j | a_j(s))}{\pi_j(y_j | a_j)} = 0.$$

Now, if the second sum in (A3.2) depends on y_j , we have a contradiction, since the first sum clearly does not. Assume, therefore, that

$$\text{(A3.3)} \quad \sum_{s=1}^{k_j-1} \gamma(s) \pi_j(y_j | a_j(s)) = c \pi_j(y_j | a_j)$$

for all y_j such that $\pi_j(y_j | a_j) > 0$. Indeed, from (A3.1), we note that (A3.3) holds even for y_j such that $\pi_j(y_j | a_j) = 0$ (simply divide (A3.1) by $\pi_i(y_i | a_i) \neq 0$). But since the $\gamma(s)$ are not all zero the fact that (A3.3) holds for all y_j contradicts the linear independence of B_j . *Q.E.D.*

REFERENCES

- ABREU, D. (1988): "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56, 383–396.
- ABREU, D., P. DUTTA, AND L. SMITH (1994): "Folk Theorems for Repeated Games: A NEU Condition," *Econometrica*, 62, 939–948.
- ABREU, D., P. MILGROM, AND D. PEARCE (1991): "Information and Timing in Repeated Partnerships," *Econometrica*, 59, 1713–1734.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1986): "Optimal Cartel Monitoring with Imperfect Information," *Journal of Economic Theory*, 39, 251–269.
- (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.
- BEWLEY, T. (1983): "A Difficulty with the Optimum Quantity of Money," *Econometrica*, 51, 1485–1504.
- FARRELL, J., AND E. MASKIN (1989): "Renegotiation in Repeated Games," *Games and Economic Behavior*, 1, 327–360.
- FUDENBERG, D., AND D. LEVINE (1983): "Subgame-perfect Equilibria of Finite and Infinite Horizon Games," *Journal of Economic Theory*, 31, 251–268.
- (1988): "The Folk Theorem in Repeated Games with Unobservable Actions," mimeo.
- (1991): "An Approximate Folk Theorem for Repeated Games with Private Information," *Journal of Economic Theory*, 54, 26–47.
- FUDENBERG, D., D. LEVINE, AND E. MASKIN (1993): "The Folk Theorem with Imperfect Public Information," mimeo.
- FUDENBERG, D., AND E. MASKIN (1986a): "The Folk Theorem for Repeated Games with Discounting and Incomplete Information," *Econometrica*, 54, 533–554.
- (1986b): "Discounted Repeated Games with One-Sided Moral Hazard," MIT.
- GREEN, E. (1987): "Lending and the Smoothing of Uninsurable Income," in *Contractual Arrangements for Intertemporal Trade*, ed. by E. Prescott and N. Wallace. Minneapolis: University of Minnesota Press.
- GREEN, E., AND R. PORTER (1984): "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 52, 87–100.
- KANDORI, M. (1988): "Monotonicity of Equilibrium Payoff Sets with Respect to Observability in Repeated Games with Imperfect Monitoring," mimeo.
- KANDORI, M., AND H. MATSUSHIMA (1992): "The Role of Communication in Repeated Games with Private Information," mimeo.
- KOHLBERG, E., AND J.-F. MERTENS (1986): "On the Strategic Stability of Equilibrium," *Econometrica*, 54, 1003–1034.
- KREPS, D., AND R. WILSON (1982): "Sequential Equilibrium," *Econometrica*, 50, 863–894.
- LEHRER, E. (1988): "Internal Correlation in Repeated Games," mimeo, Northwestern University.
- (1989): "Lower Equilibrium Payoffs in Two Player Repeated Games with Nonobservable Actions," forthcoming in *International Journal of Game Theory*.
- LEGROS, P. (1988): "Sustainability in Partnerships," mimeo, California Institute of Technology.
- LEVINE, D. (1991): "Asset Trading Mechanisms and Expansionary Policy," *Journal of Economic Theory*, 54, 148–164.
- MADRIGALE, V. (1986): "On the Non-existence of Efficient Equilibria of Repeated Principal Agent Games with Discounting," mimeo.
- MATSUSHIMA, H. (1989): "Efficiency in Repeated Games with Imperfect Monitoring," *Journal of Economic Theory*, 48, 428–442.
- MERTENS, J. F. (1987): "Repeated Games," *Proceedings of the International Congress of Mathematicians*, 1986, 1528–1571.
- RADNER, R. (1981): "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica*, 49, 1127–1148.
- (1985): "Repeated Principal Agent Games with Discounting," *Econometrica*, 53, 1173–1198.
- (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 Partnership Game with Discounting and with Uniformly Inefficient Equilibria," *Review of Economic Studies*, 53, 59–70.

- RUBINSTEIN, A. (1979): "Optimal Deterrence for Offenses that May Have Been Committed by Accident," in *Applied Game Theory*, ed. by S. J. Brams, A. Schotter, and G. Schwodiauer. Würzburg: Physica-Verlag, 406-413.
- RUBINSTEIN, A., AND M. YAARI (1983): "Repeated Insurance Contracts and Moral Hazard," *Journal of Economic Theory*, 30, 74-97.
- SCHEINKMAN, J., AND L. WEISS (1988): "Borrowing Constraints and Aggregate Economic Activity," *Econometrica*, 54, 23-45.
- SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25-55.
- SORIN, S. (1986): "On Repeated Games with Complete Information," *Mathematics of Operations Research*, 11, 147-160.
- SPEAR, S., AND S. SRIVASTAVA (1987): "On Repeated Moral Hazard with Discounting," *Review of Economic Studies*, 54, 599-618.
- WILLIAMS, S., AND R. RADNER (1988): "Efficiency in Partnership when the Joint Output is Uncertain," mimeo.

