

Renegotiation in Repeated Games*

JOSEPH FARRELL

Department of Economics, University of California, Berkeley, California 94720

AND

ERIC MASKIN

Department of Economics, Harvard University, Cambridge, Massachusetts 02138

In repeated games, subgame-perfect equilibria involving threats of punishment may be implausible if punishing one player hurts the other(s). If players can renegotiate after a defection, such a punishment may not be carried out. We explore a solution concept that recognizes this fact, and show that in many games the prospect of renegotiation strictly limits the cooperative outcomes that can be sustained. We characterize those outcomes in general, and in the prisoner's dilemma, Cournot and Bertrand duopolies, and an advertising game in particular.

© 1989 Academic Press, Inc.

1. INTRODUCTION

In an infinitely repeated game, many outcomes that are not Nash equilibria of the one-shot game can be sustained as subgame-perfect equilibria. The familiar basic idea is that players can be deterred from short-run opportunism by threats of future retaliation. Such threats sustain cooperative behavior in *Nash* equilibrium provided only that they are sufficiently dire: their credibility is never challenged. An equilibrium is *subgame-perfect* if each threat is credible in the sense that, should players be called upon to carry it out, no single player would unilaterally wish to back out.

Although there are no incentives for unilateral deviation from punishment strategies in a subgame-perfect equilibrium, the punishment may hurt the innocent as well as the guilty. For example, in the repeated prisoner's dilemma, it is a subgame-perfect equilibrium to sustain mutual

* We thank the NSF for financial support, and a referee for very helpful and prompt comments.

cooperation by threatening indefinite mutual “finking” after anyone cheats. But such a threat may be implausible if players can communicate and “renegotiate” after a deviation. If a player did cheat once, the other might propose that they overlook it and continue cooperating. Indeed, it seems inconsistent to argue that cooperation is sustainable because mutual finking is a credible outcome: if players believe that cooperation is sustainable, they will not abjure it just because one of them once cheated—even if, *ex ante*, they agreed on a subgame-perfect equilibrium in which they threatened to do so. Rather, it seems plausible that they will *renegotiate* back to mutual cooperation or to some other desirable outcome.

We are skeptical, therefore, about subgame-perfect equilibria in which punishments hurt both¹ players. We believe that, when renegotiation is possible, players are unlikely to play, or to be deterred by, a proposed continuation equilibrium (whether on or off the equilibrium path) that is strictly Pareto-dominated by another equilibrium that they believe is available to them.

It is a familiar notion that players select an equilibrium that is Pareto-efficient (in the set of subgame-perfect equilibria): most applications of repeated-game theory focus on the “good” equilibrium outcomes, and the “bad” equilibrium outcomes are of interest mainly as subgame-perfect punishments. Given the plethora of subgame-perfect equilibria in repeated games, this assumption (indeed, even the assumption that players reach an equilibrium at all) seems to postulate preplay negotiation in which the credible agreements are the subgame-perfect equilibria.

But if negotiation yields outcomes that are efficient within the available set, so will renegotiation. And, although this is a force for efficiency *ex post*, it also limits threats, so it may shrink the class of equilibrium outcomes. In this paper, we explore how far *ex post* efficient renegotiation shrinks the class of subgame-perfect equilibria, and in particular how far it limits *ex ante* efficiency.

Our paper is organized as follows. In Section 2, we review the basic concepts and notation of repeated games. In Section 3, we introduce the concept of *weakly renegotiation-proof* (WRP) equilibrium. In Section 4, we prove Theorem 1, which characterizes such equilibria for large enough discount factors. In Section 5, we apply Theorem 1 to four salient examples: the prisoner’s dilemma, Cournot and Bertrand duopolies, and an advertising game. In Section 6, we examine to what extent weak renegotiation-proofness is consistent with Pareto-efficiency (Theorems 2, 3, and 4). Section 7 considers the more demanding, but in our view more appealing, concept of *strongly renegotiation-proof* (SRP) equilibrium. Although

¹ We consider two-player games throughout. See the Conclusion for a discussion of the case of more than two players.

SRP equilibrium may fail to exist, we provide sufficient conditions for existence that cover a wide class of games. Section 8 discusses related work, especially the independent and simultaneous work of Bernheim and Ray (1989). Section 9 concludes.

2. FUNDAMENTALS

In this section we define some notation and terminology for use below. Readers familiar with repeated games may wish to skip to Section 3.

The One-Shot Game g

Consider a two-person finite game $g: A_1 \times A_2 \rightarrow R^2$, where A_i is player i 's action space and $g_i(a_1, a_2)$ is his payoff. We take A_i to be the finite-dimensional simplex consisting of player i 's mixed strategies; this assumption has force because we shall suppose that, in the repeated game, strategies can be conditioned on all past actions: thus, we suppose that players can observe each other's private randomizations *ex post*.² The set of payoffs, allowing for mixed strategies and convexification, is

$$V = \text{co}(\text{image}(g)) = \text{co}(\{(v_1, v_2) | \exists (a_1, a_2) \text{ with } g(a_1, a_2) = (v_1, v_2)\}),$$

where "co" denotes the convex hull and $g(a_1, a_2)$ is the expected payoff if mixed strategies are used. Player 1's minimax payoff, $\min_{a_2} \max_{a_1} g_1(a_1, a_2)$, is normalized to zero, as is player 2's. The set of feasible, strictly individually rational payoffs is

$$V^* = \{(v_1, v_2) \in V | v_1 > 0, v_2 > 0\}.$$

The Repeated Game

The "repeated game" g^* is defined as follows. In each of infinitely many periods 1, 2, . . . , the game g is played. In period t , player i 's choice $a_i(t)$ may depend on the entire *history* of the game through period $t - 1$: that is, on

$$h^{t-1} = ((a_1(1), a_2(1)), \dots, (a_1(t - 1), a_2(t - 1))).$$

Thus a *strategy* σ_i for player i is a function that, for every date t and every possible history h^{t-1} , defines a period- t action $a_i \in A_i$. Given a sequence of actions $\{a_1(\tau), a_2(\tau)\}$, we define player i 's *average payoff* as $(1 - \delta) \sum_{\tau=1}^{\infty} g_i(a_1(\tau), a_2(\tau)) \delta^{\tau-1}$.

² This assumption is not strictly necessary: see Section 5 of Fudenberg and Maskin (1986).

A pair of strategies $\sigma = (\sigma_1, \sigma_2)$ defines a probability distribution on infinite histories, and hence on payoffs; we write $g^*(\sigma, \delta)$ for the expected average payoffs (with discount factor δ) when the players follow the strategies σ . (Sometimes we will suppress the δ and simply write $g^*(\sigma)$.) This is the payoff function for the game g^* . The strategy pair $\sigma = (\sigma_1, \sigma_2)$ is a Nash equilibrium in g^* if σ_1 is a best response to σ_2 and conversely; it is a subgame-perfect equilibrium if, in every subgame, the subgame strategies induced by σ are a Nash equilibrium.

Because we are using average payoffs, we have the following simple relationship. Let σ be a strategy pair whose first-period actions are $a = (a_1, a_2)$ and whose continuation strategies after a are σ^c . Then

$$g^*(\sigma, \delta) = (1 - \delta)g(a) + \delta g^*(\sigma^c, \delta). \quad (1)$$

That is, the payoffs from σ are a convex combination of those from the first-period actions a and those from the continuation strategies σ^c . In particular, we will repeatedly use the fact that if $g_i^*(\sigma^c) \geq g_i^*(\sigma)$ then $g_i(a) \leq g_i^*(\sigma)$.

3. WEAKLY RENEGOTIATION-PROOF EQUILIBRIUM

As we argued in the Introduction, a natural way to think of equilibrium in repeated games is to suppose that players can engage in nonbinding "negotiations," both *ex ante* and after each round of play. Being competent negotiators, they have no trouble (in this complete-information problem) reaching a credible agreement that is not strictly Pareto-dominated by any other credible potential agreement.

If only *ex ante* negotiation, and not renegotiation, were possible, this would simply argue for picking a subgame-perfect equilibrium that is not strictly Pareto-dominated by any other—a common practice. But, as the example of the trigger-strategy equilibrium in the prisoner's dilemma shows, many efficient subgame-perfect equilibria have inefficient continuation equilibria. When renegotiation is possible, these may not be credible. What kind of solution does our argument then imply?

Evidently, for an equilibrium to be "credible," *all its continuation equilibria* must also be "credible," and our discussion suggests that, when renegotiation is possible, that requires that all those equilibria be Pareto-undominated by other "credible" equilibria. At a *minimum*, therefore, credibility requires that the equilibrium satisfy the following definition:

DEFINITION. A subgame-perfect equilibrium σ is *weakly renegotiation-proof* if there do not exist continuation equilibria σ^1, σ^2 of σ such

that σ^1 strictly Pareto-dominates σ^2 . If an equilibrium σ is WRP, then we also say that the payoffs $g^*(\sigma)$ are WRP.

Trivially, there is always at least one WRP equilibrium. To see this, let $a = (a_1, a_2)$ be a Nash equilibrium³ of the one-shot game g , and define a strategy σ_i^* for player i in the repeated game g^* as follows: always play a_i , irrespective of history. Then $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a subgame-perfect equilibrium in g^* , and has no continuation equilibria other than itself; thus it is certainly WRP. As we shall see, however, there are typically other WRP equilibria, at least for large enough δ .

A WRP equilibrium σ is intuitively “renegotiation-proof” in the sense that, if players agreed *ex ante* to play σ , and if the history of the game so far means that they should now play a continuation equilibrium σ^c of σ , they do not have a joint incentive to switch instead to another continuation equilibrium $\sigma^{c'}$ of σ . This is surely a *necessary* condition for renegotiation-proofness, but need not be *sufficient*: there might exist another “credible” equilibrium, *not* a continuation equilibrium of σ , that they both prefer. (This possibility will be addressed in Section 7.)

This question of the relevance of history is somewhat delicate, so we pause here to discuss it explicitly. Since the past is sunk in a repeated game, one could argue that it should not affect current or future behavior. This extreme view would rule out all threats, and thus cooperation could not be sustained by repetition.⁴

Just as that might well underestimate the importance of history in determining behavior in a repeated game, so the study of *Nash* equilibria in g^* surely overestimates it. Looking at Nash equilibria amounts to assuming that the history of the game can overcome players’ individual incentives to optimize—that we can specify behavior, after an out-of-equilibrium history, in which individual players do not optimize. Most game theorists find this unsatisfactory.

Studying subgame-perfect equilibrium is an intermediate assumption. In subgame-perfect equilibrium, history cannot overcome players’ individual incentives to optimize, but *can* affect what continuation equilibrium is chosen, even to the extent of leading to a Pareto-dominated continuation equilibrium.

Weakly renegotiation-proof equilibrium is intermediate between “history cannot matter” and looking at subgame-perfect equilibrium. We suppose that history *can* affect which of a set of Pareto-unranked continuation equilibria (such as the set of continuation equilibria of a WRP equilibrium) is chosen, so that there can be threats, but that it *cannot*

³ Recall that the sets A_i contain all the mixed strategies, so the existence of such an equilibrium is guaranteed.

⁴ For an elaboration and analysis of this point of view, see Maskin and Tirole (1989).

persuade players to do something that is jointly irrational for them in the face of a strictly Pareto-superior alternative.

4. CHARACTERIZING WRP PAYOFFS

Let Σ_δ be the set of all WRP equilibria for a given discount factor $\delta < 1$. Characterizing Σ_δ is difficult. As in the standard theory of repeated games, it is easier to study what payoffs are attainable in equilibrium for large enough δ . In order to do this, we define

$$W \equiv \bigcup_{\delta < 1} g^*(\Sigma_\delta, \delta).$$

Recall that, absent renegotiation, the perfect-equilibrium version of the folk theorem⁵ tells us that all of V^* is sustainable in subgame-perfect equilibrium for large enough δ ; thus the difference between V^* and W measures the constraints on cooperation imposed by the prospect of renegotiation.

Theorem 1 characterizes the payoffs that are WRP for large enough δ : that is, the set $W \subseteq V^*$. Before stating the theorem, we introduce some notation. Let $a = (a_1, a_2) \in A_1 \times A_2$ be a pair of actions. We write $c_i(a)$ for i 's cheating payoff from the action pair a : that is, his payoff from his best response to his opponent j 's move a_j . Formally,

$$c_i(a) = \max_{a'_i} g_i(a'_i, a_j).$$

THEOREM 1. *Let $v = (v_1, v_2)$ be in V^* . If there exist action pairs $a^i = (a_1^i, a_2^i)$ (for $i = 1, 2$) in g such that $c_i(a^i) < v_i$, while $g_j(a^i) \geq v_j$ for $j \neq i$, then the payoffs (v_1, v_2) are WRP for all sufficiently large $\delta < 1$. Moreover, these sufficient conditions are necessary for v to be WRP for some $\delta < 1$, if the first inequality is made weak. And, for $i = 1, 2$, the strict version is necessary if there exists no action pair a such that $g_i(a) = v = c_i(a)$ and $g_j(a) \geq v_j$.*

Proof. We first prove sufficiency, by constructing (i) a sequence of action pairs whose average payoffs are v and (ii) punishments for deviations from this path. The constructions are illustrated in Figs. 1 and 2, respectively.

If there exists an action pair a with $g(a) = v$, then constructing (i) is easy: we simply choose $a(t) = a$ for all t . In general, however, no such a exists. Because $v \in V^*$, v can be obtained as a convex combination of

⁵ See Fudenberg and Maskin (1986).

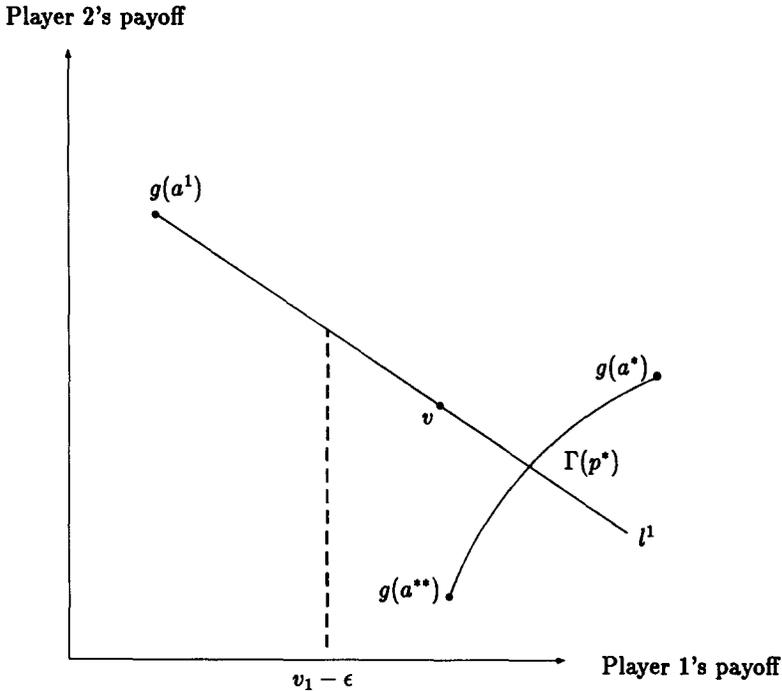


FIG. 1. Normal phase construction. All normal-phase continuation values lie on l^1 (non-strictly) to the left of v , and to the right of $v - \epsilon$.

payoffs $g(a)$, and, by Fudenberg and Maskin (1988), there exists a sequence of action pairs $\hat{a}(t)$ such that $(1 - \delta)\sum_{t=1}^{\infty} \delta^{t-1}g(\hat{a}(t)) = v$ for δ near enough to 1. Unfortunately, however, this is not enough: for renegotiation-proofness we must also ensure that the continuation payoffs starting at any two dates are Pareto-unranked. Moreover, we must also ensure that these “normal phase” continuation payoffs are not Pareto-comparable with any of the “punishment” continuation payoffs that we will construct below. In order to satisfy these constraints, we will construct normal-phase actions using the action pairs a^i given by the hypotheses of the theorem, although (as we will see below) their primary purpose is to construct punishments.

We begin, then, with three payoff vectors, $g(a^1)$, v , and $g(a^2)$ (see Fig. 1). Construct the line l^1 through $g(a^1)$ and v and extending beyond v . Conceivably, all the payoff vectors $g(\hat{a}(t))$ from the previous paragraph might lie on l^1 ; in this case, v is certainly a convex combination of two of them, and we show in the Appendix (Lemma 1) that this is enough. More probably, few or none of the $g(\hat{a}(t))$'s lie on l^1 . In this case, at least one

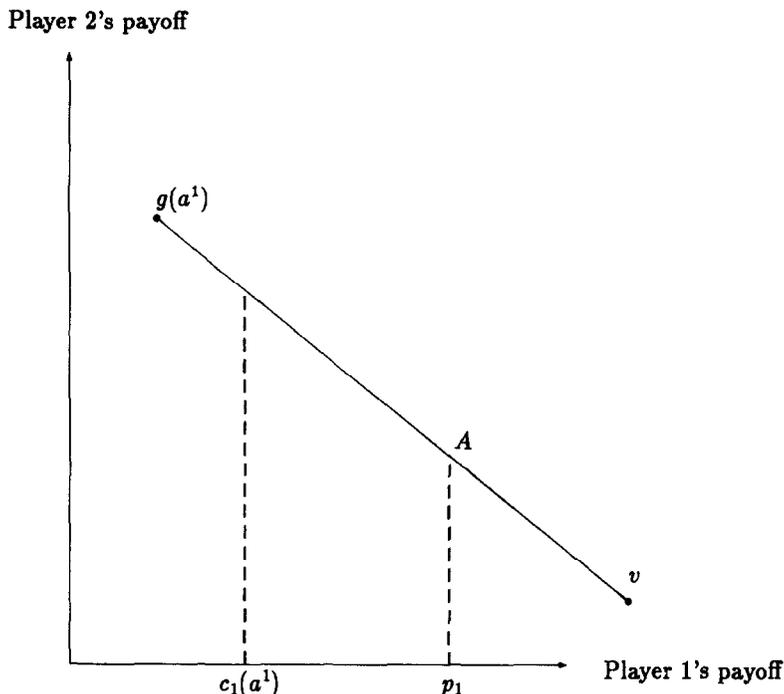


FIG. 2. Punishment phase construction. All punishment-phase continuation values (punishing player 1) lie on l^1 , between A and v .

lies above l^1 and at least one lies below; call these a^* and a^{**} , respectively.

For $p \in (0, 1)$, consider the payoffs $\Gamma(p)$ obtained when each player randomizes (independently) between a^* and a^{**} with probabilities p and $1 - p$. For small p , $\Gamma(p)$ lies above l^1 , and for large p , $\Gamma(p)$ lies below l^1 ; since $\Gamma(p)$ is continuous in p , there exists p^* such that $\Gamma(p^*)$ lies on l^1 .

If $\Gamma(p^*) = v$, then we can implement the normal phase simply by requiring randomization between a^* and a^{**} . Otherwise, suppose without loss of generality⁶ that $\Gamma(p^*)$ lies to the right of v on l^1 . Then v is a convex combination of $\Gamma(p^*)$ and $g(a^1)$.

Thus, we can represent v as a convex combination of $g(a^1)$ and $g(p^*a^* + (1 - p^*)a^{**})$. Lemma 1 in the Appendix shows that a suitable sequence of a^1 's and randomizations between a^* and a^{**} has discounted payoffs v and is such that player 1's continuation payoff from any date onward is no greater than v_1 . We use this sequence to define the normal phase of our WRP equilibrium. It is important to note that not only are the

⁶ If $\Gamma(p^*)$ lies to the left of v on l^1 , then the corresponding construction using l^2 will work.

equilibrium payoffs just v , but also all normal-phase continuation payoffs are on the line segment between v and $g(a^1)$.

We next construct renegotiation-proof punishments to deter deviations. Intuitively, to punish player i we use the action pair a^i , repeated a suitable number of times before returning to the normal phase (with payoffs v). Observe that, for all large enough δ ,

$$(1 - \delta)v_i^{\max} + \delta c_i(a^i) < v_i, \tag{2}$$

where v_i^{\max} is player i 's maximax payoff:

$$v_i^{\max} = \max_{(a_1, a_2)} g_i(a_1, a_2).$$

Because (2) is strict, we can choose p_i to satisfy both

$$p_i > c_i(a^i) \tag{3}$$

and

$$(1 - \delta)v_i^{\max} + \delta p_i < v_i. \tag{4}$$

Choose a positive integer t_i to satisfy

$$t_i g_i(a^i) + v_i^{\max} < (t_i + 1)v_i.$$

Then, for δ near enough to 1,

$$p_i = (1 - \delta^{t_i})g_i(a^i) + \delta^{t_i}v_i \tag{5}$$

satisfies both (3) and (4).

With such a value of t_i , let player i 's punishment be "play actions a^i for t_i periods and then return to the normal phase; if i cheats during the t_i periods then restart the punishment." Then player i 's continuation payoff at the beginning of his punishment is p_i , and it rises toward v_i as the punishment proceeds. Hence, by (4), he will be deterred from cheating in the normal phase if the punishment is credible. Moreover, by cooperating during his punishment, player i gets at least p_i as continuation value, whereas if he cheats, he gets $c_i(a^i)$ for a period and restarts his punishment, with overall value $(1 - \delta)c_i(a^i) + \delta p_i$; thus, by (3), he will not cheat on his punishment. Finally, since $g_j(a^i) \geq v_j$, player j will not wish to deviate from punishing player i . This construction is illustrated in Fig. 2. Formally, define the equilibrium $\sigma(v)$ as follows.

Play begins in the *normal phase*, in which players are to play the actions we constructed whose payoffs average v . If player i cheats in the normal phase, the continuation equilibrium is "play a^i for t_i periods, then return to the normal phase." If he cheats during his punishment, the punishment begins again. If player j cheats during i 's punishment, then player j 's punishment begins immediately (player i is reprieved). If both players cheat simultaneously, player 2 is punished.

Because all continuation payoffs in the normal phase lie on the line segment between $g(a^1)$ and v , and all continuation payoffs in the punish- i phase lie on the line segment between $g(a^i)$ and v , none of the continuation payoffs of $\sigma(v)$ are Pareto-ranked. Hence, for δ near enough to 1, $\sigma(v)$ is WRP,⁷ establishing "sufficiency" in Theorem 1.

We now turn to necessity. For a given value of δ , consider a WRP equilibrium σ with payoffs v : $g^*(\sigma, \delta) = v$. We will show that an action pair a^1 satisfying Theorem 1 exists; the proof that there exists an action pair a^2 satisfying Theorem 1 is completely analogous.

First, if there exists an action pair $a \in A_1 \times A_2$ such that $g_2(a) \geq v_2$ and $c_1(a) = v_1 = g_1(a)$, then we are done. Suppose, therefore, that no such pair exists. By Lemma 2 (see the Appendix), we can assume without loss of generality that the equilibrium σ has a *worst* continuation equilibrium, σ^1 , for player 1: that is, $g_1^*(\sigma^1, \delta) \leq g_1^*(\sigma^c, \delta)$ for all continuation equilibria σ^c of σ . (If there are multiple such equilibria σ^1 , choose one that is best for player 2.) Let a^1 be the first-period actions and $\hat{\sigma}^1$ the continuation equilibrium after those first-period actions, specified by σ^1 : we will show that this action pair a^1 satisfies Theorem 1. For simplicity, we will now suppress the dependence of $g^*(\cdot, \delta)$ on δ and simply write $g^*(\cdot)$.

Clearly, σ^1 is no better for player 1 than is σ : that is, $g_1^*(\sigma^1) \leq v_1$. We claim also that σ^1 is no worse for player 2 than is σ : that is, $g_2^*(\sigma^1) \geq v_2$. For if not, and if $g_1^*(\sigma^1) < v_1$, then we must have $g_2^*(\sigma^1) \geq v_2$ because otherwise σ would strictly Pareto-dominate σ^1 , which cannot be since by assumption σ is WRP. And if $g_2^*(\sigma^1) < v_2$ and $g_1^*(\sigma^1) = v_1$, then $g_2^*(\sigma^1) \geq v_2$ because otherwise (by the parenthetic remark in the preceding paragraph) σ would contradict the choice of σ^1 . Thus indeed $g_2^*(\sigma^1) \geq v_2$.

Now we claim that $g_2(a^1) \geq g_2^*(\sigma^1)$. For if, to the contrary, $g_2(a^1) < g_2^*(\sigma^1)$, then by Eq. (1) we have $g_2^*(\hat{\sigma}^1) > g_2^*(\sigma^1)$. Thus, since $\hat{\sigma}^1$ cannot strictly Pareto-dominate σ^1 , we must have $g_1^*(\hat{\sigma}^1) \leq g_1^*(\sigma^1)$. But these last two inequalities imply that $\hat{\sigma}^1$ would contradict the definition of σ^1 . Thus, after all, $g_2(a^1) \geq g_2^*(\sigma^1)$. Since we showed in the preceding paragraph that $g_2^*(\sigma^1) \geq v_2$, it follows that $g_2(a^1) \geq v_2$, as claimed in Theorem 1.

We complete the proof by showing that $c_1(a^1) \leq g_1^*(\sigma^1) \leq v_1$ and that at least one of these inequalities is strict. As we noted above, $g_1^*(\sigma^1) \leq$

⁷ The strategy pair $\sigma(v)$ that we have constructed is, as we have shown, WRP for all δ near enough to 1. Of course, the payoff $g^*(\sigma(v), \delta)$ depends on δ and (in general) is equal to v only for the specific value of δ that we began with.

$g_1^*(\sigma) = v_1$ by construction. To see that $c_1(a^1) \leq g_1^*(\sigma^1)$, observe that otherwise player 1 could profitably cheat in the first period of σ^1 : he could get $c_1(a^1)$ in the first period and his continuation payoff after this cheating cannot (by definition) be worse for him than σ^1 , and hence by (1) he should cheat, a contradiction. Finally, to see that at least one of the inequalities is strict, note that if not, then $c_1(a^1) = g_1^*(\sigma^1) = v_1$, and this together with $g_2(a^1) \geq v_2$ contradicts our assumption above that no such action pair exists. We conclude that $c_1(a^1) < v_1$. ■

Equilibrium Paths of WRP Equilibrium

The proof of Theorem 1 not only characterizes WRP equilibrium payoffs for large δ , but also constructs corresponding equilibrium paths. An equilibrium path specifies an action pair $a(t)$ for each period t . The continuation values of the equilibrium path are the payoff pairs $v(t) = (1 - \delta)\sum_{\tau=t}^{\infty} \delta^{\tau-t}g(a(\tau))$. Then we have:

COROLLARY. *A sequence $\{a(t)\}$ of action pairs is the equilibrium path of a WRP equilibrium for all large enough $\delta < 1$ if (i) no two continuation values $v(t)$ are strictly Pareto-ranked, and (ii) there exist action pairs a^i ($i = 1, 2$) such that, for all t , $c_i(a^i) < v_i(t)$ and $g_j(a^i) \geq v_j(t)$ ($j \neq i$). Moreover, the converse holds for each δ if we replace $c_i(a^i) < v_i(t)$ by the weak inequality $c_i(a^i) \leq v_i(t)$.*

Proof. Follows directly from the proof of Theorem 1. After player i cheats, actions a^i are specified for some number of periods, followed by a return to the “normal phase.” ■

5. EXAMPLES

In this section we apply Theorem 1 to characterize WRP payoffs for four examples of economic interest that we will carry through the paper: the prisoner’s dilemma, Cournot and Bertrand duopolies, and a model of advertising.

Prisoner’s dilemma. Consider the version of the prisoner’s dilemma in Table I. If we take a^1 to be (cooperate, fink) and a^2 to be (fink, cooperate) then, for any payoff pair $v \in V^*$, $c_i(a^i) = 0 < v_i$ and $g_j(a^i) = 3 > v_j$, for $i = 1, 2$ and $j \neq i$. Hence, by Theorem 1, all elements of V^* are WRP payoffs for large enough δ .

Note that in this game, the same “universal” punishments a^i can be used to sustain all WRP payoffs. In standard repeated-game theory, “minimax the offender” is always a universal punishment in that sense.

TABLE I
PRISONER'S DILEMMA

	Player 2	
	Cooperate	Fink
Player 1		
Cooperate	2, 2	-1, 3
Fink	3, -1	0, 0

But when renegotiation is possible, different payoffs require different punishments in general, as we shall see in our other examples.

Cournot duopoly. Consider a Cournot duopoly in which marginal costs are zero and demand is given by $p = 2 - x$. Firm's pure strategies are quantities⁸ $a_i \in [0, 2]$, and payoffs are $g_i(a) = a_i(2 - a_1 - a_2)$. It is immediate that

$$V^* = \{v = (v_1, v_2) | v_i > 0 \text{ and } v_1 + v_2 \leq 1\}.$$

If a payoff pair v is WRP then, from Theorem 1, there exists an action pair a^1 (to punish firm 1 for deviations) such that

$$\max_{a_1} \{a_1(2 - a_1 - a_2^1)\} \leq v_1, \quad (6)$$

and

$$a_2^1(2 - a_1^1 - a_2^1) > v_2. \quad (7)$$

We seek conditions on v for such a punishment pair a^1 to exist. First, if a_2^1 were random, we could replace it by its mean: this would not affect the left hand side of (6) and would not lower the left hand side of (7). Thus we can assume that a_2^1 is a pure strategy. Now we can rewrite (6) as

$$\frac{1}{4}(2 - a_2^1)^2 \leq v_1. \quad (8)$$

To solve (7) and (8), we may as well take $a_1^1 = 0$; we must then satisfy (8) and

$$a_2^1(2 - a_2^1) \geq v_2. \quad (9)$$

⁸ Note that in this game and in the Bertrand duopoly example below, there are infinitely many pure strategies in the one-shot game g . Nonetheless, it is easy to show that Theorem 1 still applies.

It is easy to see that this task is possible if and only if

$$16v_1 \geq (v_2 + 4v_1)^2. \tag{10}$$

Similarly, of course, for a suitable punishment pair a^2 for player 2 to exist requires

$$16v_2 \geq (v_1 + 4v_2)^2. \tag{11}$$

From conditions (10) and (11) we see that the set of WRP payoffs, W , is that shown in Fig. 3. In particular, the *collusive* WRP payoffs (those with $v_1 + v_2 = 1$) are those on the line segment between $(\frac{1}{9}, \frac{8}{9})$ and $(\frac{8}{9}, \frac{1}{9})$. (And by the last sentence of Theorem 1, these endpoints are excluded.)

Bertrand duopoly. In Bertrand duopoly (with the same cost and demand assumptions), $c_i(a) \geq g_j(a)$ for all pure strategies a_j (player i can infinitesimally undercut player j 's price a_j). We claim that this implies that, if we considered only pure-strategy punishments a^i , no outcome other than the one-shot (zero-profit) Nash equilibrium would be WRP. For in any WRP equilibrium that gives strictly positive profits to one firm,

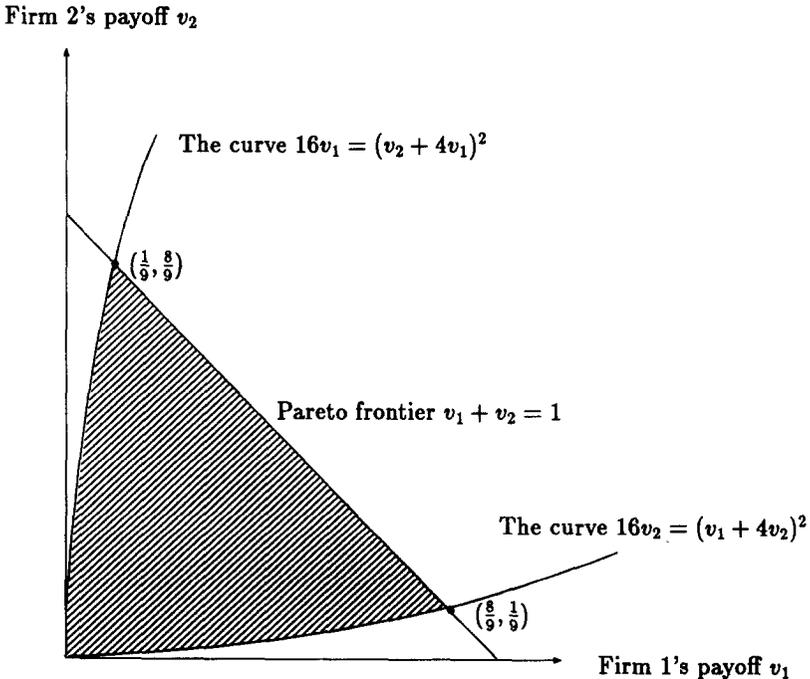


FIG. 3. Cournot duopoly. Shaded area represents WRP payoffs.

say j , the action pair a^i must punish player i while still giving player j a strictly positive profit. But, for pure strategies a^i , this implies that a^i_j is not a best response to a^i_j . Hence, we would have $v_i > c_i(a^i) \geq g_j(a^i) \geq v_j$. This implies that firm i also makes strictly positive profits in equilibrium, and so by symmetry $v_j > v_i$ also—a contradiction. Consequently, we must consider randomized punishments. To see what payoffs are WRP (with randomized punishments) for sufficiently large δ , we construct action pairs $a^i(\pi)$ that maximize $g_2(a^1)$ subject to $c_1(a^1) \leq \pi$, and use Theorem 1.

To maximize g_2 and minimize c_1 , we may as well take a^1_1 to be concentrated entirely on prices above 1 and a^1_2 to be concentrated entirely on prices no greater than 1. (To see the latter, observe that any weight on $p > 1$ in a^1_2 could be shifted to $2 - p$, with no reduction in g_2 and with no increase in c_1 . The former claim is then obvious.)

Moreover, in $a^1_2(\pi)$, there should be no weight on prices below $p^*(\pi)$, the smaller solution of $p(2 - p) = \pi$. (Firm 1 cannot achieve profits above π by pricing at or below p^* , so there is no danger of increasing c_1 above π by shifting weight from strictly below p^* to p^* itself; and such a shift increases g_2 , since the function $p(2 - p)$ is increasing in the interval $(0, 1)$.) Therefore, writing F for the distribution function of firm 2's price, we have $F(p) = 0$ for $p < p^*$, and $F(1) = 1$.

Within the range $[p^*, 1)$, g_2 is increasing in price, so we should make each $F(p)$ as small as possible subject to maintaining c_1 at π . Clearly this is achieved if we maintain $p(2 - p)(1 - F(p)) = \pi$ for $p \in [p^*, 1)$, so that

$$F(p) = \begin{cases} 0 & \text{for } p < p^*, \\ 1 - \frac{\pi}{p(2 - p)} & \text{for } p^* \leq p < 1, \\ 1 & \text{for } p \geq 1. \end{cases}$$

Now it is straightforward to calculate $g_2(a^1)$. With probability π , firm 2 sets price 1; otherwise, its price is continuously distributed on $(p^*, 1)$. Thus we have

$$g_2(a^1(\pi)) = \pi + \int_{p^*}^1 p(2 - p)f(p)dp,$$

where $f(p) = F'(p)$ is the density function. Changing variable to $y \equiv p(2 - p)$,⁹ we get

⁹ In the relevant range, y is a monotone function of p , so the change is legitimate.

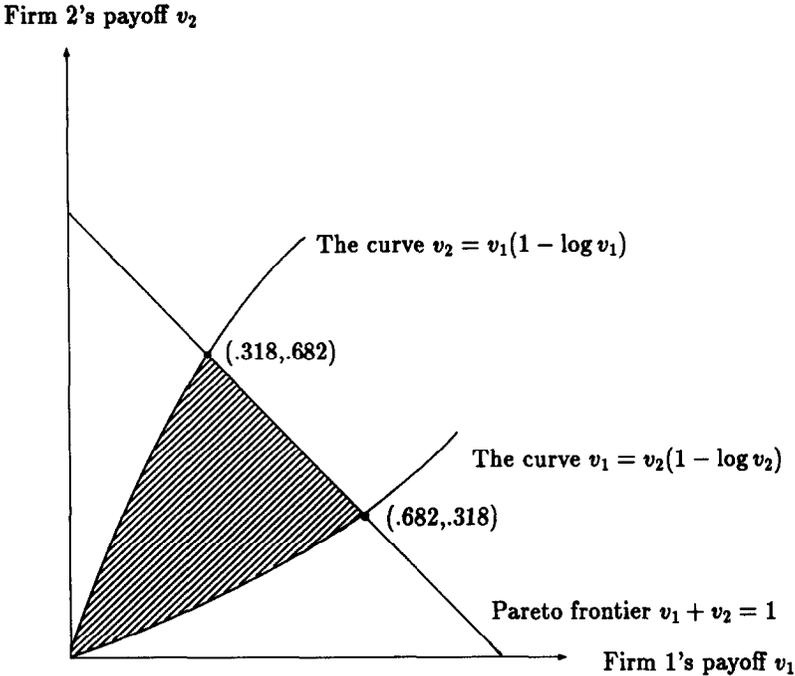


FIG. 4. Bertrand duopoly. Shaded area represents WRP payoffs.

$$g_2(a^1(\pi)) = \pi + \pi \int_{\pi}^1 \frac{dy}{y} = \pi(1 - \log \pi).$$

Consequently, by Theorem 1, a payoff pair $(v_1, v_2) \in V^*$ is WRP if $v_1(1 - \log v_1) > v_2$ and $v_2(1 - \log v_2) > v_1$, and is WRP only if those inequalities hold weakly: see Fig. 4. In particular, a Pareto-efficient (collusive) payoff pair (v_1, v_2) is WRP if and only if both v_1 and v_2 satisfy $v_i(2 - \log v_i) > 1$, or (approximately) $v_i > .318$. That is, each firm must be getting almost a third of the profits. This condition is considerably stronger than that for Cournot duopoly.

An advertising game. Suppose that each of two firms can advertise in three ways: "high," "low," and "dirty." If both choose low, they each get a payoff of zero. If one chooses high and the other chooses low, the high advertiser (who has created new demand for his product) earns 3 while his rival still earns zero. If they both choose high, their efforts somewhat cancel each other out, and each earns 1. Finally, "dirty" ad-

vertising is equivalent to low advertising unless the other firm chooses “high,” in which case the dirty advertiser can turn its rival’s high profile against it: the dirty advertiser earns 2 and the high advertiser earns -2 . Thus, we have:

		Firm 2		
		High	Low	Dirty
Firm 1	High	1, 1	3, 0	-2, 2
	Low	0, 3	0, 0	0, 0
	Dirty	2, -2	0, 0	0, 0

If we use punishment actions $a^1 = (H, D)$ against firm 1, and $a^2 = (D, H)$ against firm 2, Theorem 1 implies that any payoffs $v = (v_1, v_2) \in V^*$ such that $v_1 < 2$ and $v_2 < 2$ are WRP for δ near enough to 1. Moreover, we claim that all WRP payoffs satisfy $v_1 \leq 2$ and $v_2 \leq 2$. (See Fig. 5.) For

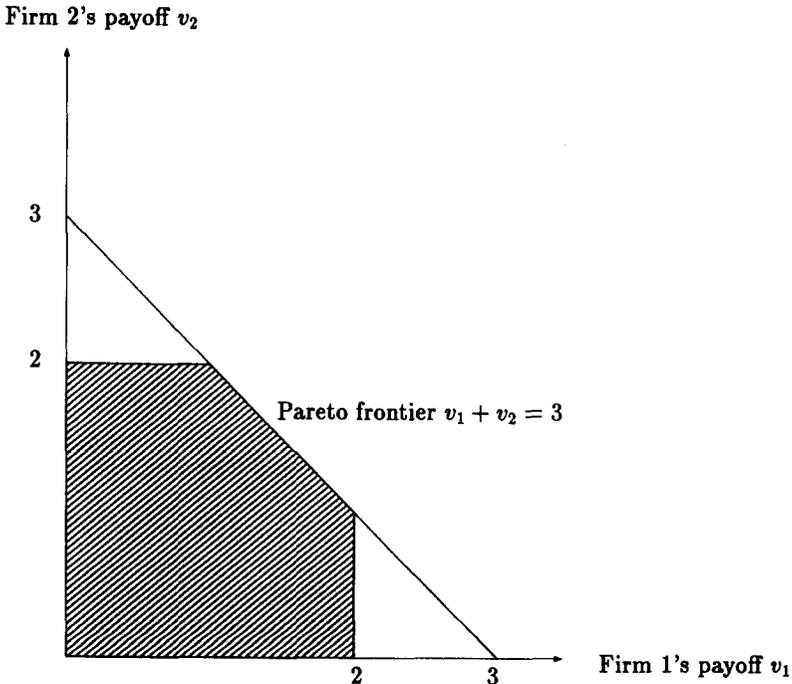


FIG. 5. Advertising game. Shaded area represents WRP payoffs.

suppose to the contrary that $v = (v_1, v_2)$ were WRP, where (say) $v_2 > 2$. Then, from feasibility, $v_1 < 1$, and hence, from Theorem 1, there must exist a^1 with $g_2(a^1) > 2$ and $c_1(a^1) < 1$. To try to construct such an a^1 , it is clear that we would want firm 1 to put no weight on “dirty” (since firm 2 prefers that 1 play either “high” or “low” rather than “dirty”), and firm 2 to put none on “low” (since firm 1 prefers that 2 play “low” rather than “dirty”). A routine calculation then proves that there is no such a^1 .

6. EFFICIENCY

The possibility of *Pareto-efficient* equilibria in g^* that are not Nash equilibria of g has been a central inspiration of the study of repeated games. In all the examples considered above, we have seen that at least part of $P(V^*)$, the (weak) Pareto frontier of V^* , is WRP for large δ . But this is not true for all games; to see this, consider the following game¹⁰:

		Player 2		
		Stone	Scissors	Paper
Player 1	Paper	1, 0	0, 1	0, 0
	Stone	0, 0	1, 0	0, 1
	Scissors	0, 1	0, 0	1, 0

In this game, no payoffs other than the one-shot Nash equilibrium payoffs are WRP. To see this, suppose that $v = (v_1, v_2)$ is the payoff vector $g^*(\sigma, \delta)$ for some $\delta < 1$ and for some WRP equilibrium σ . From Lemma 2, we can assume that σ has a continuation equilibrium σ^1 that is worst for player 1, and by Theorem 1, the first-period action pair a^1 in σ^1 must satisfy $c_1(a^1) \leq v_1$ and $g_2(a^1) \geq v_2$.

Now write a_2^1 as the vector $(p, q, 1 - p - q)$, where p is the probability of playing “stone” and q the probability of “scissors.” Then $c_1(a^1) = \max\{p, q, 1 - p - q\}$. Moreover, $g_2(a^1) \leq \max\{p, q, 1 - p - q\}$. Thus, $c_1(a^1) \geq g_2(a^1)$, and we have

$$v_1 \geq c_1(a^1) \geq g_2(a^1) \geq v_2.$$

Similarly $v_2 \geq v_1$. Hence, $v_1 = v_2 = u$, say.

¹⁰ Shapley (1962) used this example in a different context. It can be interpreted as a modified version of “stone–scissors–paper” where *both* players lose if they pick the same thing. Note also that, contrary to our standard normalization, the minimax is not zero.

Finally, we claim that (u, u) must correspond to a one-shot Nash equilibrium, i.e., $u = \frac{1}{3}$. For the argument above implies that, for all continuation equilibria σ' of σ , $g^*(\sigma', \delta) = (u', u')$ for some u' ; and since neither σ nor σ' may strictly Pareto-dominate the other, $u' = u$. Intuitively, this implies that no punishments are possible; hence, the conclusion. Formally, let a be the first-period action pair in σ , and let σ^c and $\hat{\sigma}$ be the continuation equilibria after a is played and after player 1 cheats on a , respectively. Both have continuation payoffs u . By Eq. (1), therefore, we have

$$u = (1 - \delta)g_1(a) + \delta u;$$

$$u \geq (1 - \delta)c_1(a) + \delta u.$$

Hence $g_1(a) = u = c_1(a)$, and likewise for player 2; thus a is the (unique) one-shot Nash equilibrium action pair, which yields payoffs $(\frac{1}{3}, \frac{1}{3})$ as claimed.

In this section, we give three results on WRP equilibrium and efficiency. Theorem 2 states that *generically* there does exist a Pareto-efficient WRP equilibrium. In Theorem 3, we characterize the set of Pareto-efficient WRP equilibrium payoffs. In Theorem 4, we give conditions under which there exists a WRP equilibrium *all* of whose continuation equilibria are Pareto-efficient.

THEOREM 2. *Given the players' action spaces A_1 and A_2 , for a generic¹¹ choice of payoff function g , and for all δ close enough to 1, there exists a WRP equilibrium that is Pareto-efficient.*

Proof. See Evans and Maskin (1989). ■

From now on, we will be concerned primarily with games in which there is a Pareto-efficient WRP equilibrium; Theorem 2 assures us that these games are generic. We write \overline{W} for the closure of the set W of payoff vectors that are WRP for large enough discount factors, and we write $w^1 = (w_1^1, w_2^1)$ for the payoff vector that is worst for player 1 in $\overline{W} \cap P(V^*)$. By Theorem 2, this intersection is generically nonempty, and it is certainly closed; therefore, this is well defined for generic games. (If there is more than one w_2^1 consistent with w_1^1 , take the largest one.) Similarly, define w^2 as the payoff vector in $\overline{W} \cap P(V^*)$ that is worst for player 2. We first show that the weak inequalities of Theorem 1 extend to the w^i 's—indeed, to any payoff vector in \overline{W} .

¹¹ Since we take the pure-strategy action sets to be finite, a payoff function g is just a finite-dimensional vector. Thus, "for a generic choice of payoff function" simply means "for an open and dense set of vectors in $2|A_1| \cdot |A_2|$ -dimensional space."

LEMMA 3. For $w \in \bar{W}$ and for $i = 1, 2$, there exists an action pair $a^i \in A_1 \times A_2$ such that $c_i(a^i) \leq w_i$ and $g_j(a^i) \geq w_j$ for $j \neq i$.

Proof. For definiteness, let us take $i = 1$. We first show that the function $c_1(\cdot)$ is continuous. To see this, observe that $c_1(a)$ is the maximum of $g_1(a'_1, a_2)$ over all (mixed) strategies a'_1 , but that this is the same as the maximum over pure strategies a'_1 . Hence, $c_1(a)$ is the maximum of finitely many functions of a_2 , of the form $g_1(a'_1, \cdot)$. Since $g_1(\cdot)$ is continuous in a_2 , so is $c_1(\cdot)$.

Now for $w \in \bar{W}$, w is the limit of WRP payoff vectors $v(n)$. By Theorem 1, for each n , there exists an action pair $a^1(n) \in A_1 \times A_2$ such that $c_1(a^1(n)) \leq v_1(n)$ and $g_2(a^1(n)) \geq v_2(n)$. Since $A_1 \times A_2$ is compact, there exists a convergent subsequence of the action pairs $a^1(n)$. Let a^1 be the limit of such a subsequence. Then, since the functions $g_2(\cdot)$ and $c_1(\cdot)$ are continuous, a^1 has the properties claimed. ■

Remark. Applying Lemma 3 to $w^i \in \bar{W}$ yields action pairs that we denote by α^i , such that $c_i(\alpha^i) \leq w^i_i$ and $g_j(\alpha^i) \geq w^i_j$ for $j \neq i$. We note, incidentally, that the α^i can be used as “universal” punishment pairs to enforce any Pareto-efficient WRP payoffs.

We now show that, in a generic game, any payoffs between w^1 and w^2 on the Pareto frontier of the (convex hull of the) feasible set can arise in a WRP equilibrium for δ near enough to 1. Moreover, the Pareto boundary of \bar{W} consists of this portion of the Pareto frontier of V^* , together with (perhaps) parts of the lines $v_1 = w^1_1$ and $v_2 = w^1_2$.

THEOREM 3. Consider a generic game (i.e., one in which $W \cap P(V^*)$ is nonempty). Any point on the Pareto frontier of V^* lying strictly between w^1 and w^2 is WRP for large enough δ . Furthermore, if $v = (v_1, v_2)$ lies on the Pareto frontier of \bar{W} , then either v lies on the Pareto frontier of V^* (nonstrictly) between w^1 and w^2 , or else $v_1 = w^1_1$ or $v_2 = w^1_2$.

Proof. The first claim follows immediately from Theorem 1, since any point in $P(V^*)$ lying strictly between w^1 and w^2 satisfies the (strict) conditions of Theorem 1, with the action pairs $a^1 = \alpha^1$ and $a^2 = \alpha^2$. Moreover, by the definition of the w^i , any Pareto-efficient point in \bar{W} must be (nonstrictly) between w^1 and w^2 .

Since the part of $P(V^*)$ strictly between w^1 and w^2 is in \bar{W} , and since no point in \bar{W} can strictly Pareto-dominate any other, it follows that any point $v \in P(\bar{W}) \setminus P(V^*)$ must lie (nonstrictly) outside the shaded area in Fig. 6: that is, either $v_1 \leq w^1_1$ and $v_2 \geq w^1_2$ or $v_1 \geq w^1_1$ and $v_2 \leq w^1_2$. Suppose without loss of generality that $v_1 \leq w^1_1$ and $v_2 \geq w^1_2$. The theorem claims that the second of these inequalities cannot hold strictly. If it did so, then (from the definition of w^1) so would the first: thus we would have $v_1 < w^1_1$ and $v_2 > w^1_2$, as illustrated in Fig. 6.

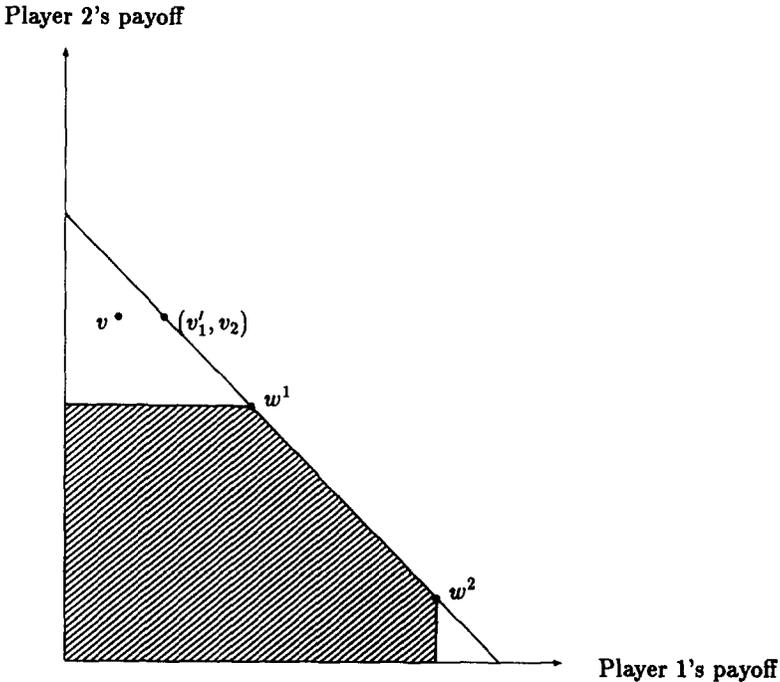


FIG. 6. Illustration of Theorem 3. \bar{W} lies (nonstrictly) within the shaded area.

Since v is Pareto-inefficient, there exists $v'_1 > v_1$ such that the point (v'_1, v_2) is Pareto-efficient. We will show, using Theorem 1, that this point is also WRP, contradicting the definition of w^1 . For punishing player 1, since v is WRP, there exists an action pair a^1 with $c_1(a^1) \leq v_1$ and $g_2(a^1) \geq v_2$; since $v'_1 > v_1$, this same action pair a^1 can be used to support (v'_1, v_2) against deviations by player 1. For punishing player 2, we can use the action pair α^2 : we have $c_2(\alpha^2) \leq w_2^2 \leq w_2^1 < v_2$, and $g_1(\alpha^2) \geq w_1^2 \geq w_1^1 \geq v'_1$. Thus (v'_1, v_2) is WRP, contradicting the definition of w^1 . This contradiction proves the last claim of the theorem. ■

Remark. Theorem 3 also implies that if $v = (v_1, v_2) \in \bar{W}$ then

$$v \in S \equiv \{v \in V^* | v_1 \leq w_1^2 \text{ and } v_2 \leq w_2^1\}.$$

The set S is depicted by the shaded area in Fig. 6. Observe that \bar{W} may fill S (as in the advertising game) or not (as in Cournot or Bertrand duopoly).

Although the equilibria of Theorems 2 and 3 are Pareto-efficient, their continuation equilibria need not be. As we shall argue more generally below, this may sometimes cast doubt on their credibility. As a step

toward resolving that doubt, and for independent interest, we next ask when, for δ near 1, there exists a WRP equilibrium all of whose continuation equilibria (including itself) are Pareto-efficient (with respect to V^*). Such an equilibrium is *strongly perfect* in the terminology of Rubinstein (1980).¹²

THEOREM 4. *If, for $i = 1, 2$, there exist (necessarily Pareto-efficient) action pairs a^i and also (necessarily Pareto-efficient) payoffs v^i , such that $c_i(a^i) < v_i^i < v_i^j$ for $j \neq i$, and such that all convex combinations of the payoff vectors $g(a^i)$ and v^i are Pareto-efficient, then there exists a strongly perfect equilibrium. Moreover, the nonstrict version of these strict conditions is necessary for the existence of such an equilibrium.*

Proof. First, sufficiency. Construct an equilibrium σ^1 as follows.

Begin by playing the actions a^1 for t_1 periods, and thereafter play actions leading to the payoffs v^1 . If, however, player 1 deviates at any point, begin σ^1 again. If player 2 ever cheats, begin σ^2 , which is defined as follows: "play the actions a^2 for t_2 periods, and thereafter play the actions leading to the payoffs v^2 ; if player 2 ever cheats, begin σ^2 again, and if player 1 cheats, begin σ^1 ."

This is a subgame-perfect equilibrium if t_i is chosen so that it does not pay player i to cheat, either on his punishment or during another phase of the equilibrium; as in Theorem 1, these conditions can be satisfied when δ is sufficiently large.

Moreover, for each continuation equilibrium, there exists i ($= 1, 2$) such that the equilibrium path involves only actions a^i and actions leading to the payoffs v^i . Because (by assumption) all convex combinations of $g(a^i)$ and v^i are Pareto-efficient, each continuation equilibrium is Pareto-efficient. Consequently, σ^1 is strongly perfect. This proves sufficiency.

Next, necessity. Suppose that σ is a strongly perfect equilibrium, i.e., all its continuation equilibria are Pareto-efficient, for discount factor δ . As usual we may assume that, for $i = 1, 2$, σ has a worst (for player i) continuation equilibrium, σ^i ; moreover, if there are multiple such equilibria, let σ^i be one that is best for player $j \neq i$.

Let a^i be the first-period action pair specified in σ^i , and \hat{v}^i the continuation payoffs after the first period in σ^i . Let v^i be the convex combination of \hat{v}^i and $g(a^i)$ such that $v_i^i = c_i(a^i)$.

Since, by assumption, σ^i is Pareto-efficient, the convex combination of $g(a^i)$ and \hat{v}^i , with weights $1 - \delta$ and δ , respectively, is Pareto-efficient. This implies that those two points are on the same line segment of the Pareto frontier of V^* , and therefore that *all* convex combinations are Pareto-efficient; hence, the same is true of $g(a^i)$ and v^i , as we wished to show.

¹² For two-player games, the concepts are equivalent. With more than two players, strong perfection requires more than perfection and Pareto-efficiency of all continuation equilibria.

Finally, we must show that $v_i^j \leq v_i^i$. Since σ^i is the worst continuation equilibrium for player i , we have $c_i(a^i) \leq g_i^*(\sigma^i)$. By definition of σ^i , we have $g_i^*(\sigma^i) \leq g_i^*(\sigma^j)$. Hence, $v_i^i = c_i(a^i) \leq g_i^*(\sigma^i) \leq g_i^*(\sigma^j)$. But likewise $v_j^j \leq g_j^*(\sigma^j)$. Now suppose first that $v_j^j < g_j^*(\sigma^j)$. Then since v^j is Pareto-efficient, $v_i^j \geq g_i^*(\sigma^j)$. Suppose next that $v_j^j = g_j^*(\sigma^j)$. If $v_i^j < g_i^*(\sigma^j)$, then $g_j(a^j) = v_j^j = \hat{v}_j^j$ and $v_i^j < g_i^*(\sigma^j) < \hat{v}_i^j$, and so \hat{v}^j contradicts the fact that σ^j is a worst continuation equilibrium for j that is best for i . Hence, again, $v_i^j \geq g_i^*(\sigma^j)$. We conclude that $v_i^j \leq v_i^i$ as claimed. ■

Examples

Prisoner's dilemma. As we saw above, in this game there are Pareto-efficient punishments that drive an offender down to his minimax payoff. Consequently, demanding efficient punishments imposes no extra constraints on WRP equilibrium paths. In Theorem 4, we can let a^1 be the action pair (*cooperate, fink*) and a^2 the analogous pair for player 2; we can take v^i to be (2, 2) for $i = 1, 2$.

Cournot duopoly. To characterize the WRP payoffs that can be sustained in strongly perfect equilibrium, we must solve inequalities (7) and (6) with $a_1^1 + a_2^1 = 1$. As before, we may as well take $a_1^1 = 0$, so this means solving them with $a_2^1 = 1$. But then (7) is trivially satisfied and the constraint is (6) with $a_2^1 = 1$, which yields $v_1 > \frac{1}{4}$. Symmetrically, of course, we require $v_2 > \frac{1}{4}$. Thus, WRP collusion with Pareto-efficient punishments requires $\frac{1}{4} < v_1 < \frac{3}{4}$, a strictly smaller range than when we allowed inefficient punishments.

Bertrand duopoly. Randomized punishments are Pareto-inefficient, so there exists no strongly perfect equilibrium.

Advertising. The only Pareto-efficient action pairs are (H, L) and (L, H), with payoffs (3, 0) and (0, 3), respectively. But $c_2(H, L) = 2$ and $c_1(L, H) = 2$, so there can exist no Pareto-efficient payoffs v^1, v^2 satisfying the necessary conditions of Theorem 4.

7. STRONGLY RENEGOTIATION-PROOF EQUILIBRIUM

We noted at the outset that the requirement of weak renegotiation-proofness is, as the name suggests, too weak a condition to guarantee credibility when renegotiation is possible. Even if an equilibrium is WRP, there may be another equilibrium that Pareto-dominates it. Of course, any strongly perfect equilibrium is free from this problem; but, as we have seen, such equilibria need not exist. And, clearly, demanding that all continuation equilibria be Pareto-efficient is unduly restrictive: it should not be considered an objection to a proposed WRP equilibrium to point out that it is Pareto-dominated by another subgame-perfect equilibrium

that is itself *not* WRP. We therefore call a WRP equilibrium *strongly renegotiation-proof* if none of its continuation equilibria is strictly Pareto-dominated by another WRP equilibrium. In this section we characterize the SRP equilibria for δ near enough to 1.

Trivially, for sufficiently *small* δ , the only subgame-perfect equilibria consist of sequences of one-shot Nash equilibria. In this case, it will be SRP as well as WRP to repeat infinitely often an undominated one-shot Nash equilibrium. For larger values of δ , even for δ close to 1, SRP equilibrium may fail to exist, as we shall see below. But SRP equilibria do exist in many games of economic interest, including three of our four examples.

We first provide a simple sufficient condition for existence of SRP for large δ . Recall that w^i is the payoff vector that is worst for player i in $\bar{W} \cap P(V^*)$, and that α^i is an action pair such that $c_i(\alpha^i) \leq w_i^i$ and $g_j(\alpha^i) \geq w_j^j$. (If there is more than one such action pair, let α^i be one that minimizes $c_i(\alpha^i)$.)

THEOREM 5. *Consider a generic game (i.e., one for which $W \cap P(V^*)$ is nonempty). If, for $i = 1, 2$ and for $j \neq i$, $c_i(\alpha^i) < w_i^i < w_i^j$, then every payoff vector $v \in \bar{W} \cap P(V^*)$ is SRP for all δ sufficiently close to 1.*

Proof. From Lemma 1 in the Appendix, for every sufficiently large δ , there is a sequence of Pareto-efficient action pairs $\{a^i(v)\}$ whose discounted-average payoffs are v and all of whose continuation payoffs lie in $\bar{W} \cap P(V^*)$. Define an equilibrium $\sigma(v, \delta)$ as follows: In the normal phase, follow the actions $\{a^i(v)\}$. But should player i ever deviate, punish him by playing α^i for a finite number of periods and then move to the sequence of action pairs $a^i(w^i)$ (beginning with $t = 1$). The number of periods is chosen as in the proof of Theorem 1, so that the punishment continuation payoff is sufficiently small and yet not so small that cheating on the punishment actions becomes attractive. The same punishment is used if a player deviates during a punishment phase.

Clearly, no equilibrium Pareto-dominates the continuation payoffs in the normal phase or, in a punishment, once the players reach the stage of playing $a^i(w^i)$. But we must show that no WRP equilibrium can Pareto-dominate the continuation payoffs while players are supposed to be taking the actions α^i at (or near) the beginning of i 's punishment: in general, *these* continuation payoffs v^i are *not* Pareto-efficient.

Suppose, then, that some WRP payoff vector \hat{v} strictly Pareto-dominates v^1 . Since $\hat{v} \in W$, there exists $v^* \in P(\bar{W})$ that weakly Pareto-dominates \hat{v} , and consequently strictly Pareto-dominates v^1 . Now v^1_2 is a convex combination of $g_2(\alpha^1)$, which is at least equal to w^1_2 , and w^1_2 ; thus $v^1_2 \geq w^1_2$, and hence $v^*_2 > w^1_2$. Since w^1 is Pareto-efficient, this implies that

$v_1^* \leq w_1^1$. But this and the fact that $v_2^* > w_2^1$ together contradict the definition of w^1 .¹³ ■

Theorem 5 stated that, under certain conditions, every Pareto-efficient WRP payoff vector is also SRP for large enough δ . Theorem 6 is a converse of a sort. Part (ii) asserts that if σ is SRP for all discount factors near 1, then, for any given such δ , the corresponding equilibrium payoffs v lie on the Pareto boundary of S (the shaded area in Fig. 6); that is, v is either Pareto-efficient or is on the horizontal or vertical parts of the boundary of S . The hypothesis that the same strategies σ be SRP for a range of discount factors is strong (although it is satisfied by those constructed in the proof of Theorem 5; see footnote 13). Accordingly, part (i) of Theorem 6 considers a sequence of SRP equilibria as the discount factor converges to 1 and shows that the corresponding payoffs converge to the Pareto boundary of S .

THEOREM 6. *Consider a generic game. (i) Suppose that, for some $\hat{\delta} < 1$ and for all $\delta > \hat{\delta}$, there exists a strategy profile $\sigma(\delta)$ that is an SRP equilibrium for δ . If $v = (v_1, v_2) \in V^*$ is such that there exists a sequence $\{\delta_n\}$ with $\delta_n \rightarrow 1$ and with $g^*(\sigma(\delta_n), \delta_n) \rightarrow v$, then either (a) $v \in \bar{W} \cap P(V^*)$ or (b) for some $i \neq j$, $v_j = w_j^i$ and $c_i(\alpha^i) < w_i^i < w_i^j$. Moreover, (ii) if there exists a (fixed) strategy profile σ that is SRP for all $\delta > \hat{\delta}$, then for all such δ (not only in the limit), the payoff vector $g^*(\sigma, \delta)$ satisfies either (a) or (b).*

Proof. Every WRP payoff vector lies (nonstrictly) in the shaded set S of Fig. 6. The main claim of part (i) is that every payoff vector v that is the limit of SRP payoff vectors for discount factors converging to 1 actually lies on the boundary of S .

The reason for this is that, if v were in the interior of S , then v , and hence the points $g^*(\sigma(\delta_n), \delta_n)$ for large enough n ,¹⁴ would be strictly Pareto-dominated by some point v' on the Pareto frontier of V^* lying strictly between w^1 and w^2 . But (by Theorem 3) v' is WRP for large enough δ , which contradicts the assumption that $\sigma(\delta_n)$ is SRP for all n .

This proves part (i) of Theorem 6; part (ii) is proved in the Appendix. ■

¹³ Note that the equilibrium $\sigma(v, \delta)$ that we construct in this proof is SRP not only for δ but for all discount factors sufficiently near 1. Of course, for these other discount factors, the corresponding average payoffs need not be v .

¹⁴ Formally, since v' strictly Pareto-dominates v , there exists a neighborhood N of v such that v' strictly Pareto-dominates all points in N . Since $g^*(\sigma(\delta_n), \delta_n) \rightarrow v$, $g^*(\sigma(\delta_n), \delta_n) \in N$ for large enough n .

Examples

Prisoner's dilemma. We saw above that for δ near 1, any payoffs in $P(V^*)$ can be sustained in a strongly perfect equilibrium. Such an equilibrium is certainly SRP. Hence, for δ near 1, the SRP payoffs consist of those in $P(V^*)$.

Cournot duopoly. We have $w^1 = (\frac{1}{3}, \frac{2}{3})$, $w^2 = (\frac{2}{3}, \frac{1}{3})$, and $c_i(\alpha^i) = \frac{1}{3}$ for $i = 1, 2$. Hence, by Theorem 6, every equilibrium that is SRP for large enough δ must consist entirely of Pareto-efficient points. But, as we saw earlier, the payoffs that can be sustained in strongly perfect equilibrium are precisely those on the Pareto frontier between $(\frac{1}{4}, \frac{3}{4})$ and $(\frac{3}{4}, \frac{1}{4})$.

Bertrand duopoly. It can be shown that, for all δ sufficiently close to 1, no SRP equilibria exists.¹⁵ Because that demonstration is complicated, however, we content ourselves here with showing that there exists no equilibrium that is SRP for all δ near 1.

From an earlier argument, we know that there exists no strongly perfect equilibrium. Thus, if there were an equilibrium that was SRP for all sufficiently large δ , it would have a Pareto-inefficient continuation equilibrium, and case (b) of Theorem 6 would apply: without loss of generality, we can write this as $c_1(\alpha^1) < w_1^1$. But from our previous analysis of this example, we know that

$$w_1^1(2 - \log w_1^1) = 1,$$

and also that the cumulative distribution function of the (random) punishment action α_2^1 is given by

$$F(p) = \begin{cases} 0, & p < p^*, \\ 1 - \frac{w_1^1}{p(2 - p)}, & p^* \leq p < 1; \\ 1, & p \geq 1, \end{cases}$$

where p^* is the smaller solution of $p(2 - p) = w_1^1$. From this it is easy to check that $c_1(\alpha^1) = w_1^1$. Thus, there can be no such equilibrium.

Advertising. We know from above that there exists no strongly perfect equilibrium. But

$$c_1(\alpha^1) = c_1(H, D) = 0 < w_1^1 = 1 < 2 = w_1^2.$$

¹⁵ Bernheim and Ray (1989) and Farrell and Maskin (1987) give examples of games where SRP equilibrium fails to exist for a given $\delta < 1$. In those examples, however, existence is restored for discount factors near enough to 1.

Hence, from Theorem 5, all payoffs on the line segment between (1, 2) and (2, 1) are SRP for δ sufficiently close to 1. This game therefore illustrates that the class of SRP equilibria is, in general, strictly larger than the class of strongly perfect equilibria.

In the definition of SRP equilibrium, we have ruled out any WRP equilibrium that is strictly Pareto-dominated by any other WRP equilibrium. But of course the very motivation for defining SRP equilibrium is that not all WRP equilibria are truly credible. Thus, although (when renegotiation is possible) all truly credible equilibria are WRP, the converse does not hold. Since we should rule out only equilibria strictly Pareto-dominated by truly credible equilibria, it follows that SRP equilibrium is perhaps too demanding a solution concept: every SRP equilibrium is truly credible, but not vice versa.

One attempt at devising a suitable concept that is intermediate between WRP and SRP is that of *relatively strong renegotiation-proofness* (RSRP).¹⁶ An equilibrium σ is RSRP for δ if there exists a subset $\Sigma_\delta \subseteq \Sigma_\delta$ such that, with discount factor δ , all the continuation payoffs of σ lie on the Pareto frontier of $g^*(\Sigma_\delta, \delta)$, and there exists no strictly larger subset of Σ_δ for which there exists such an equilibrium. One advantage of the RSRP equilibrium concept is that such an equilibrium exists for any discount factor; see Farrell and Maskin (1987).

8. RELATED WORK

The last few years have seen an explosion of work on renegotiation and renegotiation-proofness. In this section we give a brief survey of work on renegotiation in dynamic games.

Bernheim and Ray (1989), in simultaneous and independent work, defined "consistent" equilibria. Their definition of "internal consistency" (IC) coincides with our definition of WRP equilibrium. They, like us, strengthen that concept, but do not concentrate on SRP equilibrium. They say that one IC set, A , "directly dominates" another, B , if some equilibrium in A strictly Pareto-dominates some equilibrium in B . Thus, an IC set which no other IC set directly dominated would (in our terminology) be strongly renegotiation-proof. Since no such set need exist, they proceed as follows. Say that one IC set, A , "dominates" another, B , if there is a finite chain $A = C_1, C_2, \dots, C_n = B$ such that, for all i , C_i directly dominates C_{i+1} . Then an IC set P^* is "consistent" if P^* dominates every IC set Q that dominates it. They show that such a set P^* always exists.

¹⁶ For another attempt, see Bernheim and Ray (1989) and our discussion of their work below.

Intuitively, the idea seems to be that players will resist renegotiating from P^* to Q (when they are meant to play an equilibrium in P^* that is dominated by one in Q) by reflecting that perhaps Q is not really plausible, since some continuation equilibrium in Q is dominated by some equilibrium in P^* . Of course, this reflection might be unconvincing, especially if Q were itself consistent, but as we have seen above, strengthening WRP (i.e., IC) is hard and we must take what we can. They calculate the consistent equilibria for the repeated prisoner's dilemma, and find (for a range of parameter values) that only cyclic behavior is consistent for some "intermediate" values of δ .

Benoit and Krishna (1988) consider renegotiation in finitely repeated games. Using the natural recursive definition of renegotiation-proofness, they find that in the limit as the horizon becomes long, the set of renegotiation-proof payoffs is connected, and is either a singleton or Pareto-efficient, with the latter case being generic.

Blume (1987) relates the finite-horizon and infinite-horizon (stationary) concepts of renegotiation-proofness, using the Fudenberg–Levine (1983) characterization of subgame-perfect equilibria of infinite-horizon games as limits of ε -perfect equilibria of finite-horizon approximations to those games.

Pearce (1987) takes a different view of renegotiation. He too looks for a self-generating set of equilibria as a solution concept for a repeated game, but unlike us and Bernheim–Ray, he does not forbid Pareto-ranked equilibria *within* a solution set. Instead, he requires that no equilibrium within the set be Pareto-dominated by *the worst* equilibrium in any other self-generating set of equilibria. He thus emphasizes a form of "external" consistency to the exclusion of what we have called "internal" consistency. He shows that, in the limit as $\delta \rightarrow 1$, the welfare loss imposed by (his form of) renegotiation-proofness vanishes.¹⁷

We have assumed that, if players defect to a new continuation equilibrium, the old (putative) equilibrium has no force. Thus any kind of "social norm" underlying an *ex ante* equilibrium cannot undo a renegotiation. More technically, we assumed that renegotiation—whether from a renegotiated equilibrium or from the original equilibrium—requires the consent of all players. By contrast, de Marzo (1988) assumes that the (*ex ante*) original equilibrium is backed by an influential social norm, and that a renegotiation can work only if all players involved will benefit from the proposed renegotiation, not only at the current node but at all future nodes: equivalently, any player at any time can insist on a *return* to the originally specified equilibrium! De Marzo finds that, for a class of two-player games, his "strong sequential equilibrium" (SSE) can achieve all

¹⁷ See also Abreu *et al.* (1989).

strictly individually rational payoffs for large enough δ . And for all games, every strictly individually rational Pareto-efficient payoff is SSE.¹⁸

Asheim (1989) uses a somewhat similar definition of renegotiation-proofness, involving a "social norm" whose focal power survives renegotiation, in the sense that any player can at any time insist on reverting to that norm. A particularly interesting feature of this definition is that, in certain circumstances, renegotiation (while still feasible for the players) is *deterred!* Roughly speaking, this works as follows. After a defection by (say) player 1, it is intended to punish player 1 with a continuation equilibrium that is unfortunately worse for *both* players than was the original equilibrium. Why do they not renegotiate back to that equilibrium? Asheim's answer is that, if they did, then either player (say, player 2) could now cheat, hurting player 1 even more than his intended punishment, *and could then insist on restoring the norm*. Once the norm is reestablished, the players see that *neither* player was doing what he was meant to do in order to carry out the specified punishment of player 1; according to the social norm, this "simultaneous defection" is not to be punished. Thus, by *pretending* to agree to the proposed renegotiation of the punishment, player 2 gets to cheat and not be punished himself!

Cave (1987) studies what he calls renegotiation-proof equilibria in a dynamic game of exploitation of a fishery. His concept of renegotiation-proofness is that (i) equilibrium payoffs should be undominated, relative to all subgame-perfect equilibria; and (ii) punishment continuation equilibria are undominated by other feasible sufficient punishments. It seems, therefore, that he seeks relatively efficient punishments to sustain those outcomes that are sustainable using arbitrary (subgame-perfect) punishments.

Maskin and Tirole (1988) investigate strong renegotiation-proofness in a dynamic price-setting duopoly. They show that for discount factors near 1 there is a unique Markov perfect equilibrium that is SRP. In this equilibrium firms share the monopoly level of profit.

Van Damme (1989) shows directly that mutual cooperation is WRP in the repeated prisoner's dilemma. He emphasizes the fact that cooperation is WRP not only for "sufficiently large" δ , but indeed for the same set of values of δ as cooperation is sustainable in subgame-perfect (or, indeed, Nash) equilibrium. As we saw above, this is atypical, and results from the special feature of the prisoner's dilemma that player 1 can be held down to his minimax payoff while nevertheless providing player 2 with a payoff greater than his maximum available in V^* .

Finally, in very recent work, Bergin and MacLeod (1989) take an axio-

¹⁸ De Marzo proves this for "competitive games," in which there is not an outcome unanimously preferred by all players. But if there is such an outcome, it is the whole Pareto frontier, and it is clearly an SSE.

matic view of the various "renegotiation-proofness" concepts, and formalize some of the logical relationships among them.

9. CONCLUSION

While renegotiation subverts some subgame-perfect equilibria in repeated games, this need not bar cooperation enforced by threats, since some punishments would not hurt the aggrieved player. In many games, such renegotiation-proof threats can enforce cooperation, although in general less effectively than when renegotiation is impossible and more severe threats are credible. We introduced two equilibrium concepts that formalize this fact: *weakly* renegotiation-proof equilibrium, which always exists, and whose payoffs we characterized; and *strongly* renegotiation-proof equilibrium, which may fail to exist but which does exist and seems compelling in many games of economic interest.

We have confined our analysis to two-player games. In fact, all our results have natural and immediate generalizations to games with three or more players. We are not convinced, however, that WRP and SRP equilibria are appropriate concepts with more than two players. Renegotiation-proofness is to some extent a *cooperative* game theoretic requirement: players agree to follow a given equilibrium if there is no better one available. But with more than two players, renegotiation could presumably involve coalitions of players smaller than the set of all players, as in Bernheim *et al.* (1987). Thus, renegotiation-proofness needs to be strengthened to allow for such possibilities.

APPENDIX

Proof of Theorem 1

We complete the proof of Theorem 1. First, we show (Lemma 1) that if v is a convex combination of two other payoff vectors, v' and v'' , then for all large enough δ we can judiciously choose $v(t)$ equal to either v' or v'' for $t = 1, 2, \dots$, so that the discounted average payoff is v and so that at each continuation date player 1's continuation payoff is close to, but no greater than, v_1 . Second, we demonstrate (Lemma 2) that we can assume without loss of generality that a WRP equilibrium has a continuation equilibrium that is worst for player 1 (and likewise, one that is worst for player 2).

LEMMA 1. *Suppose that v is a convex combination of v' and v'' . For any $\varepsilon > 0$, there exists $\hat{\delta} < 1$ such that, for every $\delta > \hat{\delta}$, there is a sequence $\{v(t)\}$, where (i) each $v(t)$ is either v' or v'' , (ii) the average payoffs are v , and (iii) for all t , players 1's continuation payoff is between $v_1 - \varepsilon$ and v_1 .*

Proof. Assume without loss of generality that $v'_1 < v_1 < v''_1$. Take $v(1) = v''$, and define $v(t)$ inductively by

$$v(t) = \begin{cases} v' & \text{if } \sum_{\tau=1}^{t-1} \delta^{\tau-1} v_1(\tau) + \delta^{t-1} v'_1 > \frac{1-\delta^t}{1-\delta} v_1 \\ v'' & \text{otherwise.} \end{cases}$$

We claim that

$$0 \leq \sum_{\tau=1}^t \delta^{\tau-1} v_1(\tau) - \frac{1-\delta^t}{1-\delta} v_1 \leq \delta^{t-1} (v''_1 - v'_1). \quad (12)$$

The first inequality in (12) is true by construction. The second clearly holds for $t = 1$, and we now prove it by induction on t . Suppose it holds for $t - 1$. Then

$$\sum_{\tau=1}^t \delta^{\tau-1} v_1(\tau) - \frac{1-\delta^t}{1-\delta} v_1 = \delta^{t-2} \kappa + \delta^{t-1} (v_1(t) - v_1), \quad (13)$$

where $\kappa \leq v''_1 - v'_1$. If $\kappa < \delta(v_1 - v'_1)$, then $v_1(t) = v''_1$, and so the right hand side of (13) is less than $\delta^{t-1} (v''_1 - v'_1)$. If $\kappa \geq \delta(v_1 - v'_1)$, then $v_1(t) = v'_1$, and so an upper bound on the right hand side of (13) is $\delta^{t-1} (v''_1 - v'_1)$ for δ near enough 1. Hence, the second inequality in (12) holds for t . Now, letting $t \rightarrow \infty$, we obtain

$$\sum_{\tau=1}^{\infty} \delta^{\tau-1} v_1(\tau) = \frac{v_1}{1-\delta}. \quad (14)$$

From (12) and (14),

$$(1-\delta) \sum_{\tau=t+1}^{\infty} \delta^{\tau-t-1} v_1(\tau) - v_1 \geq -(1-\delta) \delta^{-1} (v''_1 - v'_1). \quad (15)$$

But for δ near enough 1, the right hand side of (15) exceeds $-\varepsilon$. ■

LEMMA 2. *Let σ be a WRP equilibrium for a discount factor $\delta < 1$. For $i = 1, 2$, let v_i be the infimum of the numbers $g_i^*(\sigma^c, \delta)$ as σ^c ranges over all continuation equilibria of σ . Then there exists a WRP equilibrium σ' (for the same discount factor δ) with the same equilibrium path as σ ,*

and such that, for $i = 1, 2$, σ' has a worst continuation equilibrium σ^i for player i , in which player i 's payoff is v_i .

Proof. For $i = 1, 2$, choose a sequence $\{\sigma_n^i\}$ of continuation equilibria of σ , such that

$$\lim_{n \rightarrow \infty} g_i^*(\sigma_n^i, \delta) = v_i.$$

Because $A_1 \times A_2$ is compact, we can choose a subsequence $\{\sigma_{n,1}^i\}$ of $\{\sigma_n^i\}$ such that (i) the first-period actions converge to a pair $a^i(1)$ and (ii) for all n , the payoffs for the first-period actions of $\sigma_{n,1}^i$ are within $1/n$ of $g(a^i(1))$. Continuing iteratively, given $\{\sigma_{n,t}^i\}$, choose a subsequence $\{\sigma_{n,t+1}^i\}$ such that (i) the $(t + 1)$ -period actions converge to a pair $a^i(t + 1)$, and (ii) for all n , the payoffs for the $(t + 1)$ -period actions of $\sigma_{n,t+1}^i$ are within $1/n$ of $g(a^i(t + 1))$.

Define a strategy pair σ^i as follows. Provided nobody deviates, play the sequence of actions $a^i(1), a^i(2), \dots$. But (a) should player i deviate, start the sequence again; and (b) should player j ($j \neq i$) deviate, switch to σ^j . Now define a strategy pair σ' as follows. Play the equilibrium path of σ ; but should player i deviate, switch to σ^i . We will show that this σ' satisfies the claims of Lemma 2: this is immediate once we show that σ' is a WRP equilibrium.

First, we show that player i 's payoff in σ^i is equal to v_i . Take any $\varepsilon > 0$. Choose n so large that $\delta^n v_i^{\max} < \varepsilon/2$ and $1/n < \varepsilon/2$. Then

$$|g_i^*(\sigma^i, \delta) - g_i^*(\sigma_{n,n}^i, \delta)| < (1 - \delta^n) \frac{1}{n} + \delta^n v_i^{\max} < \varepsilon.$$

Thus, since $g_i^*(\sigma_{n,n}^i, \delta) \rightarrow v_i$, we have $g_i^*(\sigma^i, \delta) = v_i$.

Similarly, we can show that no continuation value in σ' strictly Pareto-dominates any other. It is enough to show that this is true of σ^i . Given any continuation equilibrium $\hat{\sigma}^i$ of σ^i and any $\varepsilon > 0$, there exists a continuation equilibrium $\hat{\sigma}$ of σ such that $|g_j^*(\hat{\sigma}^i, \delta) - g_j^*(\hat{\sigma}, \delta)| < \varepsilon$ for $j = 1, 2$. Therefore, since no continuation equilibrium of σ strictly Pareto-dominates any other, the same is true of the continuation equilibria of σ^i .

Finally, we show that σ^i is a subgame-perfect equilibrium. Suppose that player i deviates at time t . For any $\varepsilon > 0$ we can find an equilibrium $\hat{\sigma}$ whose t th period actions and continuation payoffs are within ε of those of σ^i . Hence, player i 's one-period gain from deviating is close to that from deviating from $\hat{\sigma}$ in the t th period. Moreover, i 's punishment for deviating— v_i —is at least as severe as that for $\hat{\sigma}$. Hence, because $\hat{\sigma}$ is an equilibrium, so is σ^i . ■

Proof of Theorem 6, part (ii)

Let σ be a pair of repeated-game strategies, such that for every $\delta > \delta^*$, σ is SRP for δ . We will show that for every such δ , $g^*(\sigma, \delta)$ lies on the outer boundary of S .

First, we argue that if $g^*(\sigma, \delta)$ lies in the interior of S , then so does $g^*(\sigma, \delta')$ for all but countably many δ' . For $g^*(\sigma, \delta')$, considered as a function of δ' , is analytic for $\delta' < 1$. Consequently, the set of values δ' for which $g^*(\sigma, \delta')$ satisfies one or more of the finite set of equations that comprise the boundary of S —the equations $g_i^*(\sigma, \delta') = w_i^j$ ($i = 1, 2; j \neq i$) and the finitely many linear equations in g_1^* and g_2^* that comprise the relevant part of $P(V^*)$ —is either countable or the whole interval; and we know (from the assumption about the case $\delta' = \delta$) that it is not the latter.

Consequently, we can pick a sequence $\{\delta_n\}$, with $\delta_n \rightarrow 1$, such that for all n , $g^*(\sigma, \delta_n)$ is in the interior of S , and such that the sequence $g^*(\sigma, \delta_n)$ converges to some point, say v , which must lie in S . We will now show that this cannot happen, by showing (i) that v cannot lie in the interior of S ; (ii) that v cannot lie in $\bar{W} \cap P(V^*) \setminus \{w^1, w^2\}$; and (iii) that v cannot satisfy $v_i = w_i^j$.

First, then, we cannot have $v \in \text{int}(S)$, by part (i) of the theorem (which we proved in the text).

Second, can we have $v \in \bar{W} \cap P(V^*) \setminus \{w^1, w^2\}$? For any point v in that set, there exist $\varepsilon > 0$ and $\delta(\varepsilon) < 1$ such that all points $v' \in \bar{W} \cap P(V^*)$ within ε of v are WRP for all discount factors greater than $\delta(\varepsilon)$. But, for n large enough, not only is $\delta_n > \delta(\varepsilon)$ but also $g^*(\sigma, \delta_n)$ is strictly Pareto-dominated by some such v' . But this contradicts the assumption that σ is SRP for such δ_n .

Third, can we have $v_i = w_i^j$? Without loss of generality, suppose that $v_1 = w_1^2$. We will show that our assumptions imply that w^2 is itself WRP for large enough discount factors; and hence that there exists $\varepsilon > 0$ and $\delta(\varepsilon) < 1$ such that for all $\delta' > \delta(\varepsilon)$, all points $v' \in \bar{W} \cap P(V^*)$ within ε of w^2 are WRP for δ' . But, for large enough n , not only is $\delta_n > \delta(\varepsilon)$ but also $g^*(\sigma, \delta_n)$ is strictly Pareto-dominated by some such v' . But this contradicts the assumption that σ is SRP for such δ_n .

To show that w^2 is WRP, we begin by showing that some payoff vector (w_1^2, v_2) is WRP for δ . Then, if $v_2 = w_2^2$ we are of course done; if $v_2 < w_2^2$ we can use the punishments for (w_1^2, v_2) to support w^2 (by Theorem 1). To show that some (w_1^2, v_2) is WRP for δ , Lemma 2 states that it is enough to show that the upper bound, B_1 , of player 1's payoffs in continuation equilibria σ' of σ is equal to w_1^2 . This is what we now show.

By assumption, $g_1^*(\sigma, \delta_n) \rightarrow w_1^2$. This means that every sufficiently long undiscounted time average of player 1's payoffs on the equilibrium path of σ is equal to w_1^2 . Intuitively, this is inconsistent with $B_1 < w_1^2$. Formally, we proceed as follows.

Let player 1's period- t payoff on the equilibrium path of σ be $v_1(t)$. Since we have normalized the minimax payoffs to zero, all continuation payoffs must be nonnegative, and therefore for any m ,

$$(1 - \delta)v_1(2) + \delta v_1(3) + \dots + \delta^{m-2}v_1(m) \leq B_1 \tag{16}$$

Similarly, of course,

$$\begin{aligned} (1 - \delta)v_1(2) + \delta v_1(3) + \dots + \delta^{m-2}v_1(m) &\leq B_1 \\ (1 - \delta)v_1(3) + \delta v_1(4) + \dots + \delta^{m-3}v_1(m) &\leq B_1 \\ &\vdots \\ (1 - \delta)v_1(m) &\leq B_1. \end{aligned} \tag{17}$$

Multiplying the inequalities in (17) by $(1 - \delta)$ and adding the inequality (16), we get

$$(1 - \delta)(v_1(1) + \dots + v_1(m)) \leq (1 + (m - 1)(1 - \delta))B_1.$$

Similarly, for any integer t , we have

$$(1 - \delta)(v_1(tm + 1) + \dots + v_1(tm + m)) \leq (1 + (m - 1)(1 - \delta))B_1 \tag{18}$$

But since

$$\lim_{n \rightarrow \infty} (1 - \delta_n) \Big/ \sum_{t=1}^{\infty} \delta_n^{t-1} v_1(t) = w_1^2,$$

the fact that inequality (18) holds for all t and m implies that $B_1 \geq w_1^2$. Since no WRP equilibrium can give player 1 more than w_1^2 , we conclude that $B_1 = w_1^2$.

Now, as described above, Lemma 2 implies that there exists a payoff pair (w_1^2, v_2) that is WRP for δ . Hence, the payoff pair w^2 itself is WRP for large enough discount factors, and so are all points $v' \in \overline{W} \cap P(V^*)$. If the limit v of the interior points $g^*(\sigma, \delta_n)$ is on the vertical segment of the boundary of S and is strictly below w^2 , the fact that w^2 itself is WRP for large discount factors already contradicts the assumption that σ is SRP for such large discount factors. If $v = w^2$, the convergence might involve no points strictly Pareto-dominated by w^2 , but must involve points strictly Pareto-dominated by some points v' near w^2 on the Pareto frontier.

Thus we have shown that, if σ is SRP for all large enough discount factors, then its payoffs lie on the boundary of S for all such discount factors, not only in the limit. ■

REFERENCES

- ABREU, D., PEARCE, D., AND STACCHETTI, E. (1989). "Renegotiation and Symmetry in Repeated Games," mimeo, Yale University, May.
- ASHEIM, G. (1989). "Extending Renegotiation-Proofness to Infinite Horizon Games," mimeo, revised, Norwegian School of Economics and Business.
- BENOIT, J.-P., AND KRISHNA, V. (1988). "Renegotiation in Finitely Repeated Games," Harvard Business School Working Paper 89-004.
- BERGIN, J., AND MCLEOD, W. B. (1989). "Efficiency and Renegotiation in Repeated Games," mimeo, Queen's University.
- BERNHEIM, B. D., PELEG, B., AND WHINSTON, M. (1987). "Coalition-Proof Nash Equilibrium. I. Concepts," *J. Econ. Theory* **42**, 1-12.
- BERNHEIM, B. D., AND RAY, D. (1989). "Collective Dynamic Consistency in Repeated Games," *Games Econ. Behav.* **1**, 295-326.
- BLUME, A. (1987). "Renegotiation-Proof Theories in Finite and Infinite Games," mimeo, U.C. San Diego, November.
- CAVE, J. (1987). "Long-Term Competition in a Dynamic Game: The Cold Fish War," *Rand J. Econ.* **18**, 596-610.
- VAN DAMME, E. (1989). "Renegotiation-Proof Equilibria in Repeated Prisoners' Dilemma," *J. Econ. Theory* **46**, 206-217.
- EVANS, R., AND MASKIN, E. (1989). "Efficient Renegotiation-Proof Equilibria in Repeated Games," *Games Econ. Behav.* **1**, 361-369.
- FARRELL, J., AND MASKIN, E. (1987). "Renegotiation in Repeated Games," Working Paper, Harvard Institute of Economic Research, June.
- FUDENBERG, D., AND LEVINE, D. (1983). "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *J. Econ. Theory* **31**, 251-268.
- FUDENBERG, D., AND MASKIN, E. (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica* **54**, 533-554.
- FUDENBERG, D., AND MASKIN, E. (1988). "On the Dispensability of Public Randomization in Discounted Repeated Games," mimeo, Harvard and MIT.
- DE MARZO, P. (1988). "Coalitions and Sustainable Social Norms in Repeated Games," IMSSS Technical Report 529, Stanford.
- MASKIN, E., AND TIROLE, J. (1988). "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica* **56**, 571-599.
- MASKIN, E., AND TIROLE, J. (1989). "Markov Equilibrium," mimeo, MIT.
- PEARCE, D. (1987). "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," mimeo, Yale.
- RUBINSTEIN, A. (1980). "Strong Perfect Equilibrium in Supergames," *Int. J. Game Theory* **9**, 1-12.
- SHAPLEY, L. (1962). "On the Nonconvergence of Fictitious Play," RAND memorandum RM-3026-PR.