Reprinted from JOURNAL OF ECONOMIC THEORY
Vol. 50, No. 1, February 1990
Printed in Belgium

Signaling, Forward Induction, and Stability in Finitely Repeated Games*

MARTIN J. OSBORNE

Department of Economics, McMaster University, Hamilton, Ontario, Canada L8S 4M4

Received December 11, 1987; revised January 2, 1989

In a finitely repeated two-person game, suppose that after a deviation by player i from the path P in period t there is only one continuation path Q in which player i's payoff from period t on is higher than it is in P. Suppose also that player j cannot benefit from deviating from Q, whatever outcomes ensue. Then it is shown that the path P is not stable in the sense of Kohlberg and Mertens (*Econometrica* **54** (1986), 1003–1037). It follows that, in a repeated game of coordination, among the set of pure outcome paths which consist of sequences of one-shot Nash equilibria, only those with payoffs very nearly Pareto efficient are stable. *Journal of Economic Literature* Classification Number: 026.

1. INTRODUCTION

Many economic phenomena can be captured effectively only in models in which agents interact repeatedly. It is frequently appealing to formulate such models as repeated games. Yet the simplest and most widely used solution—Nash equilibrium—fails to predict a definite outcome in a large class of repeated games: the "folk theorem" asserts that a wide range of outcomes is consistent with equilibrium. Beyond showing that fully cooperative (efficient) outcomes can be supported by "implicit agreements," this result tells us little about the outcome of repeated interaction. In particular, it fails to rule out continual noncooperation (a sequence of one-shot equilibria), or any degree of partial cooperation, as outcomes.

* The work on which this paper is based was supported at Columbia University by the National Science Foundation under Grant SES-8510800 and at McMaster University by the Social Sciences and Humanities Research Council of Canada and the Natural Sciences and Engineering Research Council of Canada. The paper was completed during a visit to the Kyoto Institute of Economic Research at Kyoto University, the generous hospitality of which it is a pleasure to acknowledge. I am grateful to Jean-Pierre Benoît, John Hillas, Elon Kohlberg, Vijay Krishna, Jean-François Mertens, Michael Peters, Carolyn Pitchik, Phil Reny, Ariel Rubinstein, and an anonymous referee for helpful comments and encouragement at various stages of this work.

Thus many interesting questions are left unanswered: Under what conditions is the outcome efficient (or "collusive")? What role do threats play in determining the outcome? What effect does the presence or absence of "signals" have?

The notion of Nash equilibrium attempts to capture the idea of "strategic stability." There are many examples which show that it does so imperfectly. In response to these examples, alternative solutions have been proposed. One of the most recent is due to Kohlberg and Mertens [8], who have delved deeply into the character of a satisfactory solution. Here I investigate the implications of Kohlberg and Mertens' solution—the set of "stable outcomes"—for finitely repeated games.

I show that an outcome path which fails to satisfy an intuitive notion of strategic stability is not stable in the sense of Kohlberg and Mertens. The intuitive notion of strategic stability is illustrated in the following example. Consider the (subgame perfect equilibrium) outcome path $((T, L), (T, L))$ in the two-fold play of the game in Fig. 1. This path yields each player a payoff of 2. Suppose that player 1 deviates from the path in period 1, by playing B rather than T. This action can lead to a path in which player 1's payoff exceeds 2 only if the outcome in the second period is (B, R). Thus if player 1 deviates to B in period 1, player 2 can deduce that player 1 will use B in period 2, in which case it is better for player 2 to play R than to play L in period 2. Hence it is in player 2's interest to adopt a strategy in which she responds to a first-period deviation of player 1 to B by using R in the second period. But if player 2 adopts such a strategy, player 1 can obtain a payoff of 3 by deviating from the path; thus he has an incentive to deviate, upsetting the equilibrium.

There are two elements in this argument. First, the deviation by player 1 is an unambiguous signal of the path, say Q, he intends to follow in the future. Second, it is in player 2's interest to play consistently with this path Q. Since the game has only two periods, player 2's decision of whether or not to follow Q is straightforward: her action in the second period has no repercussions for future periods. When there are two or more periods left

	L	R
T	1,1	0,0
B	0,0	3,3

FIG. 1. The game G_1.

after a deviation by player 1, player 2 needs to consider what outcomes will occur if she deviates from Q. It may be that player 2 is made worse off by deviating from Q *whatever* outcomes follow her deviation. If so, it is certainly in her interest to play consistently with the path Q. Consequently player 1 should carry out the deviation, upsetting the equilibrium.

Thus, suppose that the following condition holds for the pure Nash equilibrium outcome path $P=(a^1, ..., a^T)$ in the T-fold repetition of an arbitrary two-player strategic game.

> There is a deviation by player i in some period τ which generates the outcome $d^\tau \neq a^\tau$ in period τ, with the property that there is precisely one sequence of outcomes $(d^{\tau+1}, ..., d^T)$ in the remaining periods for which player i is at least as well off in $(d^\tau, ..., d^T)$ as he is in $(a^\tau, ..., a^T)$, and player i is in fact better off in $(d^\tau, ..., d^T)$ than he is in $(a^\tau, ..., a^T)$. Further, player j's payoff is higher when she adheres to the path $(d^{\tau+1}, ..., d^T)$ than when she deviates from this path, *whatever* sequence of outcomes her deviation induces.

In this case I say that the path P can be *upset by a convincing deviation.*

I show (Proposition 1) that every pure outcome path which can be upset by a convincing deviation is not stable in the sense of Kohlberg and Mertens [8]. I proceed to investigate how this result restricts the set of stable outcome paths in finitely repeated games. A path consisting of a sequence of equilibria of the one-shot game may well be upset by a convincing deviation (and hence not be stable), even though every such path is subgame perfect (and sequential). From the work of Benoît and Krishna [3] we know that a limiting folk theorem applies to the set of subgame perfect equilibria of a large class of games. In Section 4, I give an example of a game in this class for which, by contrast, the set of average payoffs to pure stable outcome paths converges to a singleton.[1] However, the example, though generic, is artificial, and it is not clear that there is a wide class of games for which the set of equilibria which cannot be upset by a convincing deviation is so small.

Nevertheless, if attention is restricted to outcome paths which consist of strings of pure one-shot Nash equilibria, the condition does have substantial power in the finite repetition of (almost) any coordination game $G(\alpha, \beta)$ (see Fig. 2) in which each player has two strategies and $\beta > \alpha > 0$. Precisely, I show (Proposition 2) that in the T-fold repetition $G^T(\alpha, \beta)$ of

[1] Benoît and Krishna [3] restrict attention to pure *strategies*. My result concerns pure *outcome paths*, which are supported by strategies which are pure along the equilibrium path, but may be mixed off this path. (It is essential for the analysis of stability that mixed strategies be considered.)

	a	b
a	α, α	0,0
b	0,0	β, β

FIG. 2. The game $G(\alpha, \beta)$: a two-player game of coordination in which each player has two pure strategies. It is assumed that $\beta > \alpha \geq 0$.

such a game the only strings of pure one-shot equilibria which cannot be upset by a convincing deviation—and hence are possibly stable—contain at most k occurrences of (a, a), where k depends on α and β, but is independent of T. Thus the average payoff in all stable strings of pure one-shot equilibria converges to β, the highest payoff in the game.

In Sections 6 and 7, I consider the extent to which this result can be generalized. The signaling argument appears to lose much—though not all—of its force both when the players have more than two actions in the one-shot game and when equilibrium paths containing outcomes which are either not Nash equilibria of the one-shot game or involve randomization are considered. In these cases, there may be no deviation which is an unambiguous signal of a player's intentions. When, in the one-shot game, there is conflict over what is the best outcome, the power of the signaling argument appears also to be limited.

The condition of being upset by a convincing deviation bears a family resemblance to the "intuitive criterion" of Cho and Kreps [6], which is concerned with (two-stage) signaling games. The criteria differ because the structure of a repeated game is different from that of a signaling game, and a repeated game involves, in general, more than two periods. Further, Cho and Kreps' criterion focuses on the reasonableness of beliefs at unreached information sets; the information sets at which my criterion has force are singletons, so no issue of beliefs arises. Cho [5] extends the ideas of Cho and Kreps to arbitrary extensive games; like Cho and Kreps, he focuses on beliefs at unreached information sets.

The independent work of van Damme [9] is complementary to mine. He explores the implications of stability in a number of interesting examples, including some in which a game is played twice. The papers of Kalai and Samet [7] and Aumann and Sorin [1] address the same basic problem that I address. Kalai and Samet study "unanimity games," which are related to the repeated games I consider, and show that persistent equilibria in which the outcome is the same in every subgame of the same length are efficient. Aumann and Sorin study the pure strategy equilibria of infinite repetitions of two-person games in which there is a unique efficient outcome. Rather than selecting an outcome which is robust to all perturbations of a certain sort (as stability does), Aumann and Sorin select an outcome which is robust to a single perturbation in which every strategy with finite memory is used with positive probability.

2. Stable Equilibria

The notion of Nash equilibrium, and the stronger notions of perfect equilibrium and subgame perfect equilibrium, assesses the strategic stability of single strategy profiles. Kohlberg and Mertens [8] argue that one must instead consider *sets* of equilibria. They define a set of equilibria to be stable if it is robust to all perturbations of a certain type in the structure of the game. Precisely, let G be an n-player strategic game, let $\sigma = (\sigma_1, ..., \sigma_n)$ be a mixed strategy profile, and let $\delta = (\delta_1, ..., \delta_n) > 0$. The perturbed game $G[\sigma, \delta]$ has the same set of players and strategy sets as G; the payoff of player i in $G[\sigma, \delta]$ when he uses the strategy τ_i is equal to his payoff in G when he uses the strategy $(1-\delta_i)\tau_i + \delta_i\sigma_i$.

Definition. A closed set E of Nash equilibria of G is *stable* if it is minimal with respect to the following condition:

> for any $\varepsilon > 0$ there exists $\delta_0 > 0$ such that for any completely mixed strategy profile σ of G and any $\delta = (\delta_1, ..., \delta_n)$ with $0 < \delta_i < \delta_0$ for all i, the perturbed game $G[\sigma, \delta]$ has a Nash equilibrium within ε of E.

Suppose that G is the strategic form of an extensive game Γ. If the members of a stable set of equilibria of G generate the same (pure) outcome path in Γ then I call this path a *stable outcome path.*[2] The results of Kohlberg and Mertens imply that almost every extensive game has a stable outcome path, and that any such path satisfies the following properties.

(D) A stable outcome path is stable in any game obtained by deletion of a dominated strategy.

(I) A stable outcome path is stable in any game obtained by deletion of a strategy which is an inferior response to all the equilibria in the stable set which generates the path.

Kohlberg and Mertens refer to property (I) as "forward induction."

[2] Note that this concept depends on the extensive form.

	LLL	LLR	LRL	LRR	RLL	RLR	RRL	RRR
TTT	2	2	1	1	1	1	0	0
TTB	2	2	1	1	0	0	3	3
TBT	1	1	4	4	1	1	0	0
TBB	1	1	4	4	0	0	3	3
BTT	1	0	1	0	4	3	4	3
BTB	1	0	1	0	3	6	3	6
BBT	0	3	0	3	4	3	4	3
BBB	0	3	0	3	3	6	3	6

FIG. 3. The game G_1^2: the reduced strategic form of the two-fold play of G_1. The single number in each box is the payoff to each player. A strategy for player 1 of the form xyz means play x in the first period, and in the second period play y if player 2 chose L in the first period and z if player 2 chose R in the first period. Similarly, player 2's strategies give an action for period 1, and then actions for period 2 contingent on whether player 1 chose T or B in period 1.

As an immediate application, I argue that the outcome path $((T, L), (T, L))$ in the two-fold play of the game G_1 discussed above is not stable. (The reduced strategic form G_1^2 of this game is given in Fig. 3.) In order for $((T, L), (T, L))$ to be stable in G_1^2, by (I) it must be stable in the game obtained by deleting the strategies TBT, TBB, BTT, and BTB of player 1 (which are all inferior responses to *all* strategies of player 2 consistent with the path). But now in the reduced game we can use (D) to eliminate the strategy LLL of player 2, which is dominated by LLR. This having be done, there is no equilibrium which generates the path $((T, L), (T, L))$, so that this path is certainly not stable in the reduced game. Thus it is not stable in G_1^2.

3. FINITELY REPEATED TWO-PLAYER GAMES

Let G be a two-player strategic game in which the set of pure strategies of each player i is finite, denoted A_i. Let $u_i: A_1 \times A_2 \to \mathbb{R}$ be the payoff function of player i in G. A pair of actions $(a_1, a_2) \in A_1 \times A_2$ is an *outcome* of G. If a is an outcome of G and $b_i \in A_i$ then $a \backslash b_i$ is the outcome in which player i uses b_i and player j uses a_j.

Suppose that G is played T times in succession. A pure strategy of player i in this game is a sequence of functions, indexed by $t = 1, \ldots, T$, the tth of which prescribes, for each history of pairs of actions through period $t-1$, the action i takes in period t. Some of these strategies are duplicates: for any strategy of j, any two strategies of i which react in the same way to histories of actions by j, but differ in their reactions to histories of play by i himself, generate the same outcome path. Thus, since Kohlberg and Mertens' stable sets depend only on the reduced strategic form of a game, we can work with the strategic game G^T in which a strategy for each player i is a sequence of functions $(s_i^1, ..., s_i^T)$, where each s_i^t associates with each history $(a_j^1, ..., a_j^{t-1})$ of actions by j in G an action $a_i \in A_i$ of i in period t.

A sequence of T outcomes in G is an *outcome path* of G^T. Let P be an outcome path of G^T. A strategy of player i which follows P so long as player j does is *consistent with* P. If there is a Nash equilibrium of G^T which results in the outcome path P, then P is a *Nash equilibrium outcome path* of G^T.

Let b^i and c^i be the first- and second-ranked outcomes of G for player i $(=1, 2)$ (order tied outcomes arbitrarily). The following result formalizes the condition under which an outcome path is upset by a convincing deviation, and asserts that such an outcome path is not stable.

PROPOSITION 1. *Let $P = (\bar{a}^1, ..., \bar{a}^T)$ be a pure Nash equilibrium outcome path of G^T. Suppose that there exist $\tau \in \{1, ..., T-1\}$, $i \in \{1, 2\}$, and $\tilde{a}_i \in A_i$ such that*

$$u_i(\bar{a}^\tau \backslash \tilde{a}_i) + u_i(c^i) + (T-\tau-1)\, u_i(b^i) < \sum_{t=\tau}^{T} u_i(\bar{a}^\tau) < u_i(\bar{a}^\tau \backslash \tilde{a}_i) + (T-\tau)\, u_i(b^i) \tag{1}$$

and

$$(T-\tau)\, u_j(b^i) > \max_{a_j} \{u_j(b^i \backslash a_j) \colon a_j \in A_j \text{ and } a_j \neq b_j^i\} + (T-\tau-1)\, u_j(b^j), \tag{2}$$

where $j \neq i$. Then the outcome path P is not stable.

Proof. The argument consists of three steps.

(a) Any strategy s_i of player i in G^T with $s_i^\tau(\bar{a}_j^1, ..., \bar{a}_j^{\tau-1}) = \tilde{a}_i$ is an inferior response to every mixture of strategies of j consistent with P unless

$$s_i^t(\bar{a}_j^i, ..., \bar{a}_j^\tau, b_j^i, ..., b_j^i) = b_i^i \qquad \text{for} \quad t = \tau+1, ..., T. \tag{3}$$

Proof. If (3) is violated, and s_j is a pure strategy of j consistent with P, then (s_i, s_j) generates an outcome path in which the outcome in at least one of the periods after τ is not b^i. But then by the left-hand inequality in (1) player i's payoff is less than his payoff in P.

Now consider the game $\bar{G}^T$ obtained from G^T by eliminating the strategies of player i shown to be inferior in (a). By property (I) of stable sets the outcome path P must be stable in $\bar{G}^T$ in order to be stable in G^T.

(b) Let s_j be a strategy of player j which is consistent with the path P, and for which

$$s_j^{t^*}(\bar{a}_i^1, ..., \bar{a}_i^{\tau-1}, \tilde{a}_i, b_i^i, ..., b_i^i) \neq b_j^i \qquad \text{for some} \quad t^* \in \{\tau+1, ..., T\}.$$

Then s_j is dominated in $\bar{G}^T$ by the strategy $\hat{s}_j = (\hat{s}_j^1, ..., \hat{s}_j^T)$ defined by

$$\hat{s}_j^t(h^{t-1}) = \begin{cases} b_j^i & \text{if } t \geq \tau+1 \text{ and } h^{t-1} = (\bar{a}_i^1, ..., \bar{a}_i^{\tau-1}, \tilde{a}_i, b_i^i, ..., b_i^i) \\ s_j^t(h^{t-1}) & \text{otherwise.} \end{cases}$$

Proof. If s_i is a strategy of i in $\bar{G}^T$ which first deviates from P by playing $\tilde{a}_i$ in period τ then player j's total payoff at $(\hat{s}_j, s_i)$ in periods $\tau+1, ..., T$ is $(T-\tau)\, u_j(b^i)$. By contrast, her payoff at (s_j, s_i) cannot exceed the right-hand side of (2). If s_i is any other strategy of i in $\bar{G}^T$ then (s_j, s_i) and $(\hat{s}_j, s_i)$ generate the same outcome path, so that player j obtains the same payoff in each. Hence (2) implies that $\hat{s}_j$ dominates s_j in $\bar{G}^T$.

Now consider the game $\hat{G}^T$ obtained by eliminating from $\bar{G}^T$ all the strategies of j shown to be dominated in (b). By property (D) of stable sets, P must be stable in $\hat{G}^T$ in order to be stable in G^T.

(c) The game $\hat{G}^T$ has no Nash equilibrium which generates the path P.

Proof. Let s_j be a strategy of j in $\hat{G}^T$. Then (by (b)) we have $s_j^t(\bar{a}_i^1, ..., \bar{a}_i^{\tau-1}, \tilde{a}_i, b_i^i, ..., b_i^i) = b_j^i$ for all $t = \tau+1, ..., T$. Suppose that player i uses a strategy s_i which follows the path P through period $\tau-1$, then deviates to $\tilde{a}_i$, then plays b_i^i so long as player j uses b_j^i. Then i's payoff at (s_i, s_j) is $\sum_{t=1}^{\tau-1} u_i(\bar{a}^t) + u_i(\bar{a}^\tau \backslash \tilde{a}_i) + (T-\tau)\, u_i(b^i)$, which, by the right-hand-side inequality of (1), exceeds his payoff on the path P.

Thus P is not stable in G^T, completing the proof. ∎

The proof shows that the first part of the signaling argument—that i's deviation is an unambiguous signal that i intends to follow some path Q in the future—is equivalent to the fact that any strategy of player i which deviates and then does not follow Q is an inferior response to every strategy of player j consistent with P. The second part of the argument—that it is in player j's interest to follow Q—is equivalent to the fact that, after player i's inferior responses are eliminated, a strategy which follows Q dominates all others.

A number of variations on the result are possible. For example, suppose that after a deviation by player i we eliminate those future paths in which he is worse off than in the equilibrium. Even if more than one path remains—so that player i's deviation is not an unambiguous signal of the path he intends to follow—some of player j's strategies may be dominated. If i has a deviant strategy in which his payoff is higher than in the equilibrium *whatever* mixture of the undominated strategies player j uses, then the equilibrium is not stable. However, none of the stronger versions of the result seem to allow the results in the following sections to be strengthened.

Note that the proposition gives only a sufficient condition for a path to be unstable, not a necessary condition. The full implications of stability go significantly beyond properties (D) and (I), as examples of Cho and Kreps [6, pp. 216–219] and Banks and Sobel [2, Fig. 3, p. 655] show. Further, the argument in Proposition 1 uses (I) and (D) in only a very limited way. Nevertheless, the result provides a simple condition, with a clear interpretation, that an outcome path must satisfy in order to be stable.

4. An Example

The game G_2 in Fig. 4, in which $0 < \varepsilon < 2$, satisfies the conditions of the subgame perfect folk theorem of Benoît and Krishna [3, Theorem 3.7], so that any feasible payoff above $(0, 0)$ can be approximately achieved in a subgame perfect equilibrium of G_2^T for sufficiently large T. However, we can use Proposition 1 to show that among pure outcome paths, only those in which the average payoffs are close to 3 can be stable.

(a) The path $(a, a), (b, b), ..., (b, b), (a, a), (b, b), ..., (b, b)$ is unstable: A deviation to z in the first period by player 1 satisfies the conditions of the proposition.

(b) Any path beginning with an outcome with payoff $(-\varepsilon, 0)$ or $(0, -\varepsilon)$ which consists subsequently solely of occurrences of (b, b) is unstable: If the outcome in period 1 has payoff $(-\varepsilon, 0)$ then there is an action of player 1 in period 1 which gives him a payoff of 0. Such an action is a deviation which satisfies the conditions of the proposition. If the out-

	z	a	b
z	$-\epsilon, 0$	$0, -\epsilon$	$-\epsilon, 0$
a	$-\epsilon, 0$	$1, 1$	$0, -\epsilon$
b	$0, -\epsilon$	$-\epsilon, 0$	$3, 3$

Fig. 4. The game G_2.

come in period 1 has payoff $(0, -\varepsilon)$, then there is an action of player 2 with a similar property.

(c) Any path beginning with an outcome with payoff $(0, -\varepsilon)$ or $(-\varepsilon, 0)$ which consists subsequently solely of occurrences of (a, a) and (b, b), with precisely one occurrence of (a, a), is unstable: If the outcome in period 1 has payoff $(0, -\varepsilon)$ then there is an action of player 1 in period 1 which gives him a payoff of $-\varepsilon$. Such an action is a deviation which satisfies the conditions of the proposition. If the outcome in period 1 has payoff $(-\varepsilon, 0)$, then there is an action of player 2 with a similar property.

Given that the outcome in the last period must be a Nash equilibrium of G_2, it follows that any pure outcome path which ends with any sequence of the form in (a), (b), or (c) is unstable. Thus the only pure outcome paths which can be stable are those which contain at most one occurrence of (a, a), all the remaining outcomes being (b, b). Hence the average payoff of each player in every pure stable outcome path of G_2^T converges to 3 as $T \to \infty$. Further, by the arguments of Proposition 4 below, this statement is not vacuous: at least one of these paths—the one consisting solely of occurrences of (b, b)—is in fact stable. Thus stability, in contrast to subgame perfection, gives a very sharp prediction among pure outcome paths in G_2^T.

The argument is restricted to pure outcome paths, since these are the only ones covered by the result in Proposition 1. It is not clear what can be said about the payoffs supported by stable mixed outcome paths in this game.

Note that in order for Proposition 1 to be applied it is essential that $\varepsilon > 0$. At any outcome of G_2 other than (a, a) or (b, b) we require that there be an action of one of the players which reduces his payoff (as in (c)), and also an action by one of the players which increases his payoff (as in (b)). In terms of the signaling argument, a positive value of ε allows a player to signal his future intentions: by losing a payoff of ε, a player ensures that a future string of (b, b)'s is the only path which makes him better off.

5. The Stability of Outcome Paths Consisting of Pure One-Shot Nash Equilibria in 2×2 Games of Coordination

In this section I apply Proposition 1 to repetitions of the game $G(\alpha, \beta)$ in Fig. 2. I restrict attention to those pure outcome paths of $G^T(\alpha, \beta)$ which consist of strings of the pure Nash equilibria (a, a) and (b, b) of $G(\alpha, \beta)$. I show that given α and β, there is a positive integer k such that, for any value of T, every such outcome path of $G^T(\alpha, \beta)$ which is stable contains fewer than k occurrences of (a, a).

PROPOSITION 2. *Let T be a positive integer. Suppose that there exists a positive integer k with $2 \leq k \leq T$ such that $[k/(k-1)]\,\alpha < \beta < [(k-1)/(k-2)]\,\alpha$. Let P be a pure outcome path of $G^T(\alpha, \beta)$ in which the outcome at every stage is either (a, a) or (b, b), and in which there are k or more occurrences of (a, a). Then P is not stable.*

Proof. Apply Proposition 1 as follows. Let τ be the last period at which precisely k plays of (a, a) remain (including period τ), let i be either 1 or 2, and let $\tilde{a}_i = b$. Then $u_i(\bar{a}^\tau \backslash \tilde{a}_i) = 0$, $u_i(c^i) = \alpha$, and $u_i(b^i) = \beta$, so that, given the condition on k, condition (1) in Proposition 1 is satisfied. Condition (2) is satisfied also, since the highest one-shot payoff j can obtain by deviating from b is zero. ∎

COROLLARY 3. *For any given values of $\beta > \alpha > 0$ with $(k-1)\,\beta \neq k\alpha$ for every positive integer k, the average payoff of each player in any stable pure outcome path of $G^T(\alpha, \beta)$ which consists of a string of pure Nash equilibria of $G(\alpha, \beta)$ converges to β as $T \to \infty$.*

The following result shows that this corollary is not vacuous.

PROPOSITION 4. *For any given $\beta > \alpha \geq 0$ and any positive integer T, the outcome path P^* in $G^T(\alpha, \beta)$ consisting solely of occurrences of (b, b) is stable.*

This follows from the fact that any strategy which deviates from the path P^* is an inferior response to every strategy consistent with this path.

Proposition 2 and Corollary 3 do not deal with the case in which $(k-1)\,\beta = k\alpha$ for some integer k. When $k = 1$ (i.e., $\alpha = 0$), iterative application of property (D) shows that, among *all* outcome paths of $G^T(\alpha, \beta)$ (not just those that are pure, or consist of strings of equilibria of $G(\alpha, \beta)$), there is only one that is stable—namely that in which the outcome is (b, b) in every period.

If $(k-1)\,\beta = k\alpha$ for some integer $k \geq 2$ then we can no longer use Proposition 1 to argue that a path of length T in which there are k occurrences of (a, a) and $T-k$ occurrences of (b, b) is unstable. A player who deviates to b in the period of the first occurrence of (a, a) can do no better than obtain β in every subsequent period, which results in a payoff equal to (not better than) that along the path. In fact, by the same argument as for Proposition 4, such a path *is* stable. In paths which contain $k+1$ or more occurrences of (a, a) (the remainder being (b, b)), no deviation unambiguously signals a player's future intentions. Suppose, for example, that the path P contains precisely $k+1$ occurrences of (a, a), and a player deviates from the first occurrence of (a, a). Then this player is at least as well off as in P if the outcome is (b, b) in every period subsequent to his deviation, *or* if the outcome is (b, b) in every subsequent period but one, and (a, a) in the remaining period. Whether the full implications of stability can be used to show that a path with $k+1$ or more occurrences of (a, a) is stable in this game is an open question.

6. Extensions: Coordination Games

(a) *More Than Two Strategies per Player*

Proposition 1 has strong implications for only a relatively small class of coordination games in which each player has more than two strategies. Consider the coordination game $G(\alpha_1, ..., \alpha_m, \beta)$ in which each player has $m+1$ pure strategies $a_1, ..., a_m, b$, and the diagonal payoffs are $\alpha_1, ..., \alpha_m, \beta$ with $\beta > \alpha_m \geq \cdots \geq \alpha_1 \geq 0$. If α_1 and α_m are sufficiently close then Proposition 1 can be applied to conclude that a path consisting of a string of pure one-shot equilibria is stable only if the average payoff is close to β. Otherwise, a deviation may not be an unambiguous signal of a player's future intentions. Consider, for example, the game $G(1, 3, 7)$. If player 1 deviates from the path $((a_1, a_1), (a_1, a_1), (b, b))$ in the first period, then he could be anticipating the outcomes $((a_2, a_2), (b, b))$, $((b, b), (a_2, a_2))$, or $((b, b), (b, b))$ in the last two periods, since all of these yield more than the path. Player 2 thus does not know how to act in periods 2 and 3—she has no strategy which dominates all others. Consequently the path $((a_1, a_1), ..., (a_1, a_1), (b, b))$ in $G^T(1, 3, 7)$, which yields an average payoff close to 1, is not ruled out as unstable by the arguments of Proposition 1.

(b) *Outcome Paths Containing Mixed Nash Equilibria*

Every game $G(\alpha_1, ..., \alpha_m, \beta)$ has mixed Nash equilibria in addition to the $m+1$ pure equilibria, including a completely mixed equilibrium. Consider the strategy profile of $G^T(\alpha_1, ..., \alpha_m, \beta)$ in which each player uses his strategy in the completely mixed equilibrium of $G(\alpha_1, ..., \alpha_m, \beta)$ in each period, independent of past events. This is a completely mixed equilibrium of $G^T(\alpha_1, ..., \alpha_m, \beta)$, and hence is stable (it is an equilibrium of any sufficiently close perturbed game). Further, the payoff in this equilibrium is lower than the payoff in any pure equilibrium. Thus there is no hope of extending the result in Corollary 3 to all paths consisting of strings of Nash equilibria (rather than just strings of *pure* Nash equilibria). The extent to which the arguments I have made apply to outcome paths which include mixed equilibria of $G(\alpha_1, ..., \alpha_m, \beta)$ other than the completely mixed equilibrium is an open issue.

(c) *Outcome Paths Containing Outcomes Which Are Not Nash Equilibria of the One-Shot Game*

The difficulties which arise in trying to apply Proposition 1 to (Nash equilibrium) paths containing outcomes which are not Nash equilibria of $G(\alpha_1, ..., \alpha_m, \beta)$ are similar to those discussed in (a) above. In some cases, there appears to be no deviation which unambiguously signals a player's future intentions. For example, consider the (subgame perfect) equilibrium outcome path $((b, a), (b, b), (b, b), (a, a))$ in the game $G^4(1, 3)$. The payoff on this path is 7. Suppose that player 1 deviates in the first period. He then obtains a payoff of 1, rather than 0, in that period. Consequently he must intend to receive a payoff of at least 6 in the following three periods. There are several ways he can do so, so that his deviation does not unambiguously signal his future intentions; consequently player 2 does not have a dominant strategy when player 1 is restricted to best responses. Now augment the path by putting a string of outcomes (b, a) at the start. At no point can a player deviate in such a way that his future intentions are clear, so that the arguments of Proposition 1 cannot be invoked to demonstrate that this path is unstable.

It appears that allowing players to publicly discard payoff before each stage of the game sufficiently enriches the possible signals that a player's intention may be indicated clearly by a deviation. Thus, if player 1 can discard two units of payoff when he deviates in the first period from the path $((b, a), (b, b), (b, b), (a, a))$ in the game $G^4(1, 3)$, the only future path in which he is better off is $((b, b), (b, b), (b, b))$. Consequently even paths containing outcomes which are not one-shot equilibria are not stable unless the payoff is close to 3. (The effect on the set of stable outcomes of allowing players to discard payoff is considered by van Damme [9], who studies repetitions of the game G_3 in Fig. 5 below, and by Ben-Porath and Dekel [4], who consider one-shot games (including coordination games) in which there is a unique efficient outcome.)

	L	R
T	3,1	0,0
B	0,0	1,3

FIG. 5. The game G_3.

7. Extensions: Other Games

Proposition 1 can be applied to some games in which the players' payoffs differ, so long as the players agree on the best outcomes. For example, the off-diagonal payoffs do not have to be zero, as they are in the games I have studied so far. Once the players' disagreement is more substantial, however, Proposition 1 is less powerful. Consider, for example, the game G_3 in Fig. 5. Suppose this game is played twice, and let P be the outcome path $((B, R), (B, R))$. We have $u_i(b^i) = 3$ and $u_i(c^i) = 1$ for $i = 1, 2$, so that (1) in Proposition 1 is satisfied with $i = 1$, $\tau = 1$, and $\tilde{a}_i = T$; (2) is also satisfied. Thus the path $((B, R), (B, R))$ is not stable. Similarly the path $((T, L), (T, L))$ is not stable. By a separate argument, the paths $((B, R), (T, L))$ and $((T, L), (B, R))$ *are* stable. (Van Damme [9] has independently studied this example. He shows also that in a related game, no pure outcome path is stable.)

Once the game G_3 is played more than twice, however, for most paths condition (2) is not satisfied for any τ which satisfies (1). Intuitively, a deviation by a player may signal unambiguously the path he intends to follow in the future (condition (1) is satisfied), but it may no longer be in the interest of the other player to act so that this path is realized. Consider, for example, the path $((B, R), (B, R), (T, L))$ in G_3^3. If player 1 deviates in the first period, he obtains 0, so in order to get more than his payoff of 5 on the path he must "anticipate" the outcome (T, L) in each of the next two periods. But it may not be in player 2's interest to play in such a way that these anticipated outcomes are realized. If she does so, then she obtains of payoff of 2; if she deviates from what player 1 expects (by playing R in period 2) then the outcome (B, R) might occur in the last period, yielding her a payoff of 3. Indeed, by applying the forward induction argument to player 2's deviation, player 1 could deduce that player 2 intends to play R in the last period, so that the outcome then is in fact (B, R), and player 2 is encouraged to deviate from the "expected" path. Thus player 1 should not deviate, for he cannot reasonably expect player 2 to cooperate. A detailed analysis of the path $((B, R), (B, R), (T, L))$ shows that it is in fact stable; the question of precisely what paths are stable in G_3^T is open.

This example raises the question of precisely how deviations should be interpreted. Perhaps player 2 should deduce from a deviation that player 1 intends to play along with some *stable* path in the future—so that the fact that unstable payoff-improving paths exist should be ignored. However, even then it seems likely that a deviation by player 1 cannot unambiguously signal his intentions, since some permutations of a stable path are likely to be stable. Indeed, if any pure outcome path is strategically stable in G_3^3 then it seems that both $((T, L), (T, L), (B, R))$ and $((T, L),$

(B, R), $(T, L))$ should be, in which case a deviation in the first period from the outcome path $((B, R), (B, R), (B, R), (T, L))$ in G_3^4 does not signal unambiguously.

8. Conclusion

This paper is an attempt to understand the implications of Kohlberg and Mertens' notion of stability, and, in particular, the property of "forward induction," in finitely repeated games. The results suggest that an analysis of a repeated game which ignores the logic of forward induction incorrectly identifies many outcomes as "strategically stable." I have isolated a condition which any stable outcome must satisfy; this condition highlights the role of deviations as signals of future behavior. In some games the condition can be used to show that many Nash equilibrium paths are not stable. However, my analysis suggests that without a language in which to communicate, it is hard for players to signal clearly. If the payoffs are arranged appropriately (as in the game G_2 in Fig. 4), or if the players can publicly discard payoff (see Section 6(c)), then actions may be as good as words; otherwise ambiguities may exist.

References

1. R. J. Aumann and S. Sorin, Cooperation and bounded recall, *Games and Economic Behavior* **1** (1989), 5–39.
2. J. S. Banks and J. Sobel, Equilibrium selection in signaling games, *Econometrica* **55** (1987), 647–661.
3. J.-P. Benoît and V. Krishna, Finitely repeated games, *Econometrica* **53** (1985), 905–922.
4. E. Ben-Porath and E. Dekel, Coordination and the potential for self sacrifice, draft, 1988.
5. I.-K. Cho, A refinement of sequential equilibrium, *Econometrica* **55** (1987), 1367–1389.
6. I.-K. Cho and D. M. Kreps, Signaling games and stable equilibria, *Quart. J. Econ.* **102** (1987), 179–221.
7. E. Kalai and D. Samet, Unanimity games and Pareto optimality, *Int. J. Game Theory* **14** (1985), 41–50.
8. E. Kohlberg and J.-F. Mertens, On the strategic stability of equilibria, *Econometrica* **54** (1986), 1003–1037.
9. E. van Damme, Stable equilibria and forward induction, *J. Econ. Theory* **48** (1989), 476–496.

Printed by Catherine Press, Ltd., Tempelhof 41, B-8000 Brugge, Belgium