International Journal of Game Theory, Vol. 13, Issue 3, page 129–144.

Persistent Equilibria in Strategic Games[1])

By *E. Kalai*, and *D. Samet*, Evanston[2])

Abstract: A perfect equilibrium [*Selten*] can be viewed as a Nash equilibrium with certain properties of local stability. Simple examples show that a stronger notion of local stability is needed to eliminate unreasonable Nash equilibria. The persistent equilibrium is such a notion. Properties of this solution are studied. In particular, it is shown that in each strategic game there exists a pesistent equilibrium which is perfect and proper.

1. Introduction

The notion of Nash equilibrium is used to predict or prescribe the strategies that will be used by the players of a noncooperative game. Nash's rational is to impose necessary conditions on such strategies and to show that the imposition of these conditions on the potential strategy combinations rules out many of them and leaves a significantly smaller set to consider. At the same time Nash's conditions are not too restrictive in the sense that noncooperative games in which each of the players has finitely many (pure) strategies, have always at least one Nash equilibrium.

Selten [1975], as well as others [see *Harsanyi; Myerson; Kreps/Wilson; Kohlberg* for additional references], pointed out that Nash's conditions are not restrictive enough and that in some cases not all the Nash equilibria of a game are reasonable as outcomes of the game.

To deal with this difficulty Selten introduced the notion of perfect equilibrium. In order for a strategy combination of the players to be perfect equilibrium it has to satisfy some properties of local stability in addition to Nash's pointwise stability. Invoking his conditions Selten restricts the set of outcomes of the game further than Nash but still without losing existence.

Myerson [1978] pushes Selten's idea further and defines a set of proper equilibria, which is contained in the set of perfect equilibria, and shows existence for his solution concept.

---

[1]) Parts of this research were supported by the Center for Advanced Studies in Managerial Economics, J.L. Kellogg Graduate School of Management. In addition, Samet's work was supported by an N.S.F. grant #SOC-79-05900.

[2]) Professors *Ehud Kalai*, and *Dov Samet*, Department of Managerial Economics and Decision Sciences, Northwestern University, 2001 Sheridon, Evanston, IL 60201, U.S.A.

0020-7276/84/030129-144$2.50

A typical example to demonstrate Selten's idea of perfectness is the following two-person game.

Γ_1:

		Player 2	
		L	*H*
Player 1	*L*	1,1	0,0
	H	0,0	0,0 .

In this game if both players play their low strategies (L, L) they each receive a payoff of one. If they play any other combination of strategies both receive zero.

There are two Nash equilibria in this game. These are the strategy combinations (L, L) and (H, H). To check whether a strategy combination is a Nash equilibrium we have to check whether every player is maximizing his payoff assuming that his opponent, or opponents, behave as prescribed by the given strategy combination. This is certainly the case for the pair of combinations (L, L) and (H, H). However while we feel that the combination (L, L) is a good one, we are sceptical of the combination (H, H). It is true that if player 2 plays H, player 1 cannot improve by defecting to his L strategy. However he cannot lose by defecting to L and if moreover there is even an infinitesimally small probability that player 2 will play L then he is strictly better off playing L. The strategy combination (L, L) is much more stable. Each player loses by shifting from L and moreover, even if player i expects his opponent to use with some probability his strategy H, player i is still willing to play L.

A generalization of these ideas leads to the definition of Selten's perfect equilibrium. A Nash equilibrium is perfect if there exist certain infinitesimal trembles away from the equilibrium strategies (by using all the strategies of all the players), such that each player still wishes to play his strategy, even though the others tremble. The proper equilibrium of Myerson narrows down the set of perfect equilibrium by imposing rationality restrictions on the permissible trembles. Both the perfect and the proper equilibrium have certain properties of stability in the neighborhood of the equilibrium.

Motivated by some examples of games in which perfect and proper equilibrium yield unsatisfactory results, we study in this paper a notion of neighborhood stability.

Consider first the following 3-person game.

Γ_2:

		Player 3: L			*Player 3: H*	
		Player 2			Player 2	
		L	*H*		*L*	*H*
Player 1	*L*	1,1,1	0,0,0	*L*	0,0,0	0,0,0
	H	0,0,0	0,0,0	*H*	0,0,0	1,1,0

In this game if all three players play their low strategies they are paid each one unit. If all three players play their high strategy then player 1 and 2 are paid 1 and player 3 is paid 0. Any other strategy combination yields a payoff 0 to all three players. Player 3 only chance of being paid is if he plays his low strategy. The strategy combination (L,L,L) is a perfect and proper equilibrium of this game. Not only that there exist certain trembles under which each player still plays this combination, but for every possible tremble each player will still do it. [*Kohlberg*, calls such an equilibrium, truly perfect].

But somewhat surprisingly the strategy combination (H,H,L) is also a Nash equilibrium which is perfect and proper. The reason is, that if players 1 and 2 except that player 3 trembles into his H strategy much more then they tremble into their L strategy, they will be paid by the combination (H,H,H) (with payoffs (1,1,0)) more than they are paid by the combination (L,L,L).

But a very specific set of trembles is needed to justify the perfectness of (H, H,L). What happens if the players tremble in another way, one that might be considered even more logical? If for example player 3 trembles much less into his H strategy (which does not yield him any payoff), than player 1 and 2 tremble into their L strategy then they will all end up in the combination (L,L,L). From there, as we have said, no further trembles will get them out. It seems that the weak stability of (H,H,L) is stressed and amplified by the strong stability of (L,L,L).

Our next example deals with the neighborhood stability of strategy combinations in which all the players use all their strategies. We call such strategy combination, inner strategy combination. By the definition of perfect and proper equilibrium, trembles are applied only to those players who do not use all their strategies. As a result, inner equilibria are always perfect and proper. But trembles can be applied also to inner strategies. Consider for example the following game.

Γ_3: *The Battle of the Sexes*

		Player 2	
		L	H
Player 1	L	1,1	0,0
	H	0,0	1,1

In this game there are three Nash equilibria, (L,L), (H,H) and the mixed strategy ((1/2, 1/2), (1/2, 1/2)) which yields the expected payoff (1/2, 1/2). The equilibria (L,L) and (H,H) have strong neigborhood stability; they are immune against any tremble. The equilibrium ((1/2, 1/2)), (1/2, 1/2)) is not. If the trembles $((1/2 + \epsilon, 1/2 - \epsilon), (1/2 + \epsilon, 1/2 - \epsilon))$ are considered, then the players will shift to (L,L) which is immune against trembles.

Although neighborhood instability is inherent to inner strategy combinations in general, one cannot rule them out as a solution. Consider for example the next game.

Γ_4: *The Matching Pennies*

		Player 2	
		L	H
Player 1	L	1,−1	−1,1
	H	−1,1	1,−1

There is a unique Nash equilibrium in this game, namely ((1/2, 1/2) (1/2, 1/2)), but any tremble from this inner equilibrium will lead the players out of it. The point in favor of this equilibrium is the fact that there is no strategy combination that threatens it as do the strongly stable equilibria in the games Γ_2 and Γ_3 with respect to the weaker ones. Trembling out of the inner equilibrium in Γ_4 leads nowhere.

In this paper we develop a notion of neighborhood stability. This notion helps us to choose in each game a neighborhood-stable equilibria set which we call persistent. In particular the only persistent equilibrium in the game Γ_2 is (L,L,L). The inner equilibrium in Γ_3 (The Battle of the Sexes) is not persistent while the other two are persistent. In Γ_4 (Matching Pennies) the inner equilibrium is persistent.

We start by defining the notion of absorbing subsets of strategies. We restrict our attention to subsets of mixed strategies which are products of compact and convex subsets of individual mixed strategies. We refer to such products as retracts. A retract is self absorbing, or Nash, if every strategy τ in it has another strategy σ in it which is a best reply to τ. Thus this is a generalization of the notion of Nash equilibria to sets. Indeed we show that the set of minimal Nash retracts coincides with the set of Nash equilibria.

To strengthen the notion of stability to a neighborhood stability we consider absorbing Nash retracts which we define to be retracts that absorb some neighborhood of themselves in addition to absorbing themselves. We then look at minimal absorbing Nash retracts which we call persistent retracts. We call a strategy persistent if it belongs to some persistent retract.

It turns out that there may exist persistent strategies which are not Nash equilibrium. We show however that every game has at least one strategy which is persistent, Nash, perfect, and proper equilibrium. We study the structure of persistent retracts and persistent strategies. We show for example that every persistent strategy does not use any dominated pure strategy. We give the basic structure of persistent retracts which make their computations easier. We finally show how by applying the notion of persistency, the difficulties described by the previous examples disappear.

2. Definitions and Notations

An n-person strategic game Γ consists of:

1. A set N of n players; denote $N = \{1, 2, \ldots, n\}$,
2. for each $i \in N$ a finite set S_i (the *pure strategies* of i),
3. for each $i \in N$ a real valued function u_i defined on $S = \times_{i=1}^{n} S_i$ (the *utility payoff* to i).

An element $s = (s_1, \ldots, s_n)$ in S is called a *pure strategy combination*.

For every finite set F we let $\Delta(F)$ denote the set of probability distributions on F. Thus $\sigma \in \Delta(F)$ if and only if for every $x \in F$ $\sigma(x) \geqslant 0$ and $\sum_{x \in F} \sigma(x) = 1$.

We call $\Delta(S_i)$ the set of (mixed) *strategies of player i* and we let $M = \overset{n}{\underset{i=1}{\times}} \Delta(S_i)$ denote the set of (mixed) *strategy combinations*. We will often denote by s_i the mixed strategy that puts its entire mass on the pure strategy s_i.

The functions U_i are naturally extended to M by:

$$U_i(\sigma) = U_i(\sigma_1, \sigma_2, \ldots, \sigma_n) = \sum_{s \in S} U_i(s) \overset{n}{\underset{i=1}{\times}} \sigma_i(s_i), \text{ for each } \sigma \in M.$$

The set of mixed strategies of player i, $\Delta(S_i)$ can be viewed as a simplex in $E^{|S_i|}$, the Euclidian space for dimension $|S_i|$. Our topological notions with respect to $\Delta(S_i)$ refer to the topology induced on $\Delta(S_i)$ as a subset of $E^{|S_i|}$. The topology on M is the one induced on M as a subset of $\overset{n}{\underset{i=1}{\times}} E^{|S_i|}$. The interior of a set $A \subseteq M$ will be denoted by $\mathring{A}$.

For $\sigma \in M$ and $\tau_i \in \Delta(S_i)$ we define $(\sigma \mid \tau_i)$ by $(\sigma \mid \tau_i)_j = \sigma_j$ if $j \neq i$ and $(\sigma \mid \tau_i)_i = \tau_i$. Thus $(\sigma \mid \tau_i)$ is the same as σ except for player i who replaces σ_i by τ_i.

For player i and a strategy combination $\sigma \in M$ we define the set of *maximizers of i at* σ by

$$M_i(\sigma) = \{s_i \in S_i : U_i(\sigma \mid s_i) \geqslant U_i(\sigma \mid \tau_i) \text{ for each } \tau_i \in \Delta(S_i)\}.$$

The set of *best replies of player i at* σ is defined by

$$BR_i(\sigma) = \Delta(M_i(\sigma)).$$

We define the set of *best replies at* σ by

$$BR(\sigma) = \overset{n}{\underset{i=1}{\times}} BR_i(\sigma).$$

The following conditions are well known easy consequences of the definitions above.

1. $\tau_i \in BR_i(\sigma)$ if and only if $U_i(\sigma \mid \tau_i) \geqslant U_i(\sigma \mid \bar{\tau}_i)$ for every $\bar{\tau}_i \in \Delta(S_i)$.
2. BR_i is an upper semi-continuous correspondence, with $BR_i(\sigma)$ being nonempty and convex for every $\sigma \in M$.
3. BR is an upper semi-continuous correspondence with $BR(\sigma)$ being nonempty and convex for every $\sigma \in M$.

A strategy combination $\sigma \in M$ is a *Nash equilibrium* if and only if $\sigma \in BR(\sigma)$.

3. Persistent Equilibria

A retract of the game Γ is a restriction of each of the sets of mixed strategies of the players. Formally a subset R of M, is a *retract* if $R = \overset{n}{\underset{i=1}{\times}} R_i$, with each R_i being a non-empty closed convex subset of $\Delta(S_i)$.

For a retract R and a set of mixed strategies $A \subseteq M$, we say that R *absorbs* A if for every $\sigma \in A$, $BR(\sigma) \cap R \neq 0$. That is, for every $\sigma \in A$ and for every player i there is a $\tau_i \in R_i$ such that τ_i is a best reply of player i to σ.

A retract R will be called a *Nash retract* if R, absorbs itself. The notion of a Nash retract is a generalization of the notion of Nash equilibrum. Once it is proposed that the players play strategies of R no player i has individual incentive to use strategies outside of his R_i; in other words this proposal is self enforcing.

Notice that the entire set of strategy combinations M is trivially a Nash retract. Also for a strategy combination σ, $\{\sigma\}$ is a Nash retract if and only if σ is a Nash equilibrium.

Nash retracts may be viewed as a set valued solution concept. They capture Nash's idea of being self enforcing. But in general, a Nash retract may be very large (e.g. the entire game). To make theory more useful we are interested in narrowing Nash retracts down as much as possible to minimal ones. A Nash retract is *minimal* if it does not properly contain another Nash retract. This approach leads us again to the Nash equilibrium concept.

Proposition 1: A Nash retract R is minimal if and only if $R = \{\sigma\}$ where σ is a Nash equilibrium.

Proof: It is trivial that every Nash equilibrium constitutes a minimal Nash retract. The other direction of the proposition follows immediately from:

Lemma 1: Every Nash retract R contains a Nash equilibrium.

Proof: This lemma is proved by the standard method of demonstrating the existence of a Nash equilibrium using Kakutani Fixed Point theorem.

We define the restriction (domain and range) of the best reply corresponence to R by

$$BR^R : R \to R,\ BR^R(\sigma) = BR(\sigma) \cap R.$$

For every $\sigma \in R$, $BR^R(\sigma)$ is nonempty (R is Nash) and convex. Also the upper semi-continuity of BR implies that BR^R is upper semi-continuous. Since R is compact and convex Kakutani's theorem applies and we have a $\sigma \in R$ with $\sigma \in BR^R(\sigma)$. Hence $\sigma \in BR(\sigma)$ and it is a Nash equilibrium. Q.E.D.

Our goal is to strengthen Nash's notion of stability which is pointwise stability to a notion of neighborhood stability. To that end we follow the same analysis as before but instead of starting with Nash retracts which are self absorbing we start with retracts which absorb a neighborhood of themselves. We call such retracts absorbing. Formally, a retract R is called an *absorbing retract* if for some neighborhood T of R, R absorbs T, i.e., for every $\sigma \in T$ there is a $\tau_i \in R_i$ with τ_i being a best reply of player i to σ.

Using our previous notations the last condition is just: $BR(\sigma) \cap R \neq \emptyset$ for every $\sigma \in T$. Notice that every absorbing retract is in particular a Nash retract. A somewhat weaker notion of neighborhood stability is the following one. A rectract R is called *sequentially absorbing Retract* if there exists a nonincreasing sequence of Nash retracts $(R^k)_{k=1}^{\infty}$ such that R is contained in the interior of each R^k and $R = \bigcap_{k=1}^{\infty} R^k$. Instead of requiring that R absorbs a neighborhood of itself, sequential absorbing requires that close enough neighborhoods of R absorb themselves. For strategic games with finitely many strategies, which are the subject of this paper, we can state the following equivalence.

Lemma 2: A retract R is an absorbing retract if and only if it is a sequentially absorbing retract.

Proof: Assuming that R is an absorbing retract let T be a neighborhood of R which is absorbed by R. Since R absorbs T, R absorbs any retract A satisfying $R \subseteq A \subseteq T$. Thus any such retract A is Nash. So it suffices to construct any nonincreasing sequence of retracts $(R^k)_{k=1}^{\infty}$ with R being in the interior of each R^k and $R = \bigcap_{k=1}^{\infty} R^k$. Since R is a compact set contained in the interior of T this is obviously possible.

Conversely, assume that $R = \bigcap_{k=1}^{\infty} R^k$ with $(R^k)_{k=1}^{\infty}$ satisfying the conditions required for sequential absorbing. It suffices to show that for some k, R absorbs R^k. If this is not the case we can construct a sequence $(\sigma^k)_{k=1}^{\infty}$, $\sigma^k \in R^k$ such that $BR(\sigma^k) \cap R \neq \emptyset$. For every σ^k and every player i consider the maximizers of i at σ^k, $M_i(\sigma^k)$. Recall that $M_i(\sigma^k)$ is a subset of the finite set S_i and therefore there are finitely many sets of the form $M_i(\cdot)$. Since $BR_i(\cdot) = \Delta(M_i(\cdot))$ there are also finitely many sets of the form $BR_i(\cdot)$. Since there are finitely many players we may assume without loss of generality, that for each player i the sets $BR_i(\sigma^k)$ are the same for all k. Let B_i denote this common set for player i. For every k and every i we have $B_i \cap R_i^k \neq \emptyset$. By the closedness of B_i and R_i^k is must be that $B_i \cap R_i \neq \emptyset$ which is a contradiction. Q.E.D.

Following the same philosophy as in the Nash retracts development, it is desirable to reduce the absorbing retracts as much as possible. We call a retract R *persistent* if it is a minimal absorbing retract, i.e., if it does not properly contain an absorbing retract.

Theorem 1: Every game has a persistent retract.

Proof: The set of absorbing retracts of the game is partially ordered by the containment relation. By Zorn's lemma it suffices to show that each set of absorbing retracts $\{R^{\alpha}\}_{\alpha \in A}$ which is completely ordered has a lower bound. We will show that $R = \bigcap_{\alpha \in A} R^{\alpha}$ is such a lower bound; by showing that R is an absorbing retract. We first observe that there is a countable nonincreasing sequence $(R^t)_{t=1}^{\infty}$ with $R^t \in \{R^\alpha\}_{\alpha \in A}$ for which $R = \bigcap_{t=1}^{\infty} R^t$. The fact that R is an absorbing retract follows now by applying the characterization given by Lemma 2 to each of the R^t's and then to R (for each R^t we choose a Nash retract containing R^t in its interior in a way that these choices converge down to R). Q.E.D.

Remark: Zorn's lemma is a too heavy tool to prove Theorem 1, and indeed a more constructive proof is given in the next section. The reasons for using Zorn's lemma are as follows. Consider strategic games in which for each player the set of strategies is a nonempty convex and compact set Π_i in some separable metric space. The payoff functions U_i are assumed to be continuous on $M = \times_{i=1}^{n} \Pi_i$. The notions of absorbing retract and sequentially absorbing retract are applied to these games without change. But for this class of games the two concepts are not equivalent. (It is still true that absorbing implies sequential absorbing). Moreover, the existence of a minimal absorbing retract is not guaranteed for every game. Yet the arguments of Theorem 1 are applicable to these games and they prove the existence of sequentially absorbing retract in every game.

With the above existence theorem in mind we define a *strategy combination* $\sigma \in M$ to be *persistent* if it belongs to some persistent retract.

A strategy combination $\sigma \in M$ will be called a *persistent equilibrium* if σ is a Nash equilibrium and it is persistent. The existence of a persistent equilibrium is now settled because Lemma 1 directly implies the correctness of:

Theorem 2: Every game has a persistent equilibrium.

We will shortly see that unlike in the Nash development the minimal absorbing retracts contain also strategies which are not equilibrium strategies. In other words there are persistent strategies which are not equilibrium strategies. However by replacing 'Nash retracts' by 'absorbing retracts' in Lemma 1, we can significantly strengthen this lemma.

Absorbability guarantees not only the existence of a Nash equilibrium but also the existence of a perfect and even a proper equilibrium. Thus in every game we will be able to select a persistent proper equilibrium. We leave these issues to section 6 after we study some additional properties of persistent retracts.

4. The Structure of Persistent Retracts

We call two strategies of player i, η_i and τ_i *equivalent* for *player i* if for every $\sigma \in M$, $U_i(\sigma \mid \eta_i) = U_i(\sigma \mid \tau_i)$. Obviously the equivalence of two strategies of i can be checked only against pure strategies of the others.

Lemma 3: Two strategies τ_i and σ_i of player i are equivalent for i if and only if for every $s \in S$, $U_i(s \mid \eta_i) = U_i(s \mid \tau_i)$.

The next lemma gives another test for the equivalence of strategies.

Lemma 4: Let O be a nonempty open subset of M. Then two strategies of player i, η_i and τ_i are equivalent if and only if for every $\sigma \in O$

$$U_i(\sigma \mid \eta_i) = U_i(\sigma \mid \tau_i).$$

Proof: We need to prove only the "if" direction of the lemma. Let $\delta \in M$ be any strategy combination. We will that $U_i(\delta \mid \eta_i) - U_i(\delta \mid \tau_i) = 0$. Choose any $\sigma \in O$ and let

$$f(\lambda) = U_i(\lambda\sigma + (1-\lambda)\delta \mid \eta_i) - U_i(\lambda\sigma + (1-\lambda)\delta \mid \tau_i).$$

Now we want to show that $f(0) = 0$ while we know that $f(\lambda) = 0$ for all $\lambda \in [1-t, 1]$ for some $t > 0$. But observe that $f(\lambda)$ is a polynomial function in λ. Therefore, since it is zero on an open interval, it must be zero for every λ. Q.E.D.

For every player i let $Q_i^1, Q_i^2, \ldots, Q_i^{k(i)}$ denote the partition of S_i into the equivalence classes generated by equivalence of strategies. We let $\Delta(Q_i^j)$ denote the convex hull of Q_i^j.

Our goal is to show that every presistent retract consists of a small selection of strategies for every player. We will show that in a persistent retract R each R_i is the convex hull of a finite selection of strategies with no more than one strategy being selected from each Q_i^j.

We show first that in every neighborhood in M one can find a point at which best replies are unique up to equivalence of strategies, i.e., all the best replies of player i at this point belong to the same $\Delta(Q_i^j)$.

Lemma 5: Let O be any nonempty open subset of M. For every i there exists a $\sigma \in O$ such that if $s_i \in M_i(\sigma)$ and $\tau_i \in M_i(\sigma)$ then s_i and τ_i are equivalent.

Proof: Since $|M_i(\cdot)|$ is a function with positive integer values we can choose a $\sigma \in O$ with minimal $|M_i(\sigma)|$ among all the strategies in O. By the continuity of the U_i's we can find a neighborhood W of σ, such that $W \subseteq O$ and for every $\tau \in W$, $M_i(\tau) \subseteq M_i(\sigma)$. By the minimality of $|M_i(\sigma)|$ it follows that $M_i(\tau) = M_i(\sigma)$ for every $\tau \in W$. Now Lemma 5 follows as an immediate consequence of Lemma 4.

Q.E.D.

A strategy selection of player i is a finite selection of strategies containing at most one strategy from the convex hull of every nonempty equivalence class of strategies

Q_i^j. Formally, F_i is a *strategy selection* of player i if:

1. $\emptyset \neq F_i \subseteq \Delta(S_i)$,
2. if $\sigma_i \in F_i$ then $\sigma_i \in \Delta(Q_i^j)$ for some $1 \leqslant j \leqslant K(i)$, and
3. if $\sigma_i \in F_i$ and $\tau_i \in F_i$ then σ_i is not equivalent to τ_i.

Observe that if player i has no equivalent pure strategies then F_i is just a subset of the pure strategies of player i.

A retract R is a *selection retract* if for every player i, $R_i = \Delta(F_i)$ = convex hull (F_i), for some strategy selection F_i of player i.

Lemma 6: Every absorbing retract R contains an absorbing selection retract.

Proof: For each player i let $J_i = \{j : R_i \cap \Delta(Q_i^j) \neq \emptyset\}$. By Lemma 5 and the fact that R is an absorbing retract it follows that for each i, $J_i \neq \emptyset$. For every $j \in J_i$ let σ_i^j be a choice of one strategy from $R_i \cap \Delta(Q_i^j)$ and let F_i be the set of these σ_i^j. Clearly F_i is a nonempty selection of player i. We will show now that $\hat{R} = \bigtimes_{i=1}^{n} \Delta(F_i)$ absorbs a neighborhood T of itself.

Let T be a neighborhood of R which is absorbed by R. We will show that it is absorbed by $\hat{R}$. Let i be fixed. It suffices to show that for every $\eta \in T$ there is a $\tau_i \in \Delta(F_i) \cap BR_i(\eta)$. By Lemma 5 we can find a sequence $(\eta^l)_{l=1}^{\infty}$ with $\eta^l \to \eta$ as $l \to \infty$, $\eta^l \in T$, and such that for some j, $1 \leqslant j \leqslant K(i)$, $M_i(\eta^l) = Q_i^j$ for every l. It follows that $R_i \cap \Delta(Q_i^j) \neq \emptyset$ and hence there exists a $\sigma_i^j \in F_i$ with $\sigma_i^j \in BR_i(\eta^l)$ for every l. By the upper semi-continuity of BR_i it follows that $\sigma_i^j \in BR_i(\eta)$.

Lemma 6 gives us immediately the structure of persistent retracts. Q.E.D.

Theorem 3: Every persistent retract is a selection retract.

We also re-obtain the existence of persistent retracts very easily and this time without the use of Zorn's lemma. For a selection retract R let $(F_1, F_2, \ldots, F_n)$ be the players' selections in R. Let the *cardinality* of R be defined by $|R| = \sum_{i=1}^{n} |F_i|$ then we have a new easy proof of:

Theorem 1: Every game has a persistent retract.

Proof: By theorem 3 it is enough to demonstrate the existence of a minimal absorbing selection retract. Lemma 6 guarantees the existence of an absorbing selection retract, the cardinality of every such retract is a positive integer ($\geqslant n$). Therefore there exists one with minimal cardinality which must be persistent.

5. Persistent Retracts of Games with Dominated Strategies

Our next goal is to show that persistent strategies do not use pure dominated strategies.

A pure strategy of player i, s_i, is *dominated* by $\sigma_i \in \Delta(S_i)$ if for every $\tau \in M$

$$U_i(\tau \mid s_i) \leqslant U_i(\tau \mid \sigma_i),$$

with the inequality being strict for some $\tau \in M$. A pure strategy s_i is *dominated* if it is dominated by some $\sigma_i \in \Delta(S_i)$.

The following two lemmas are immediate consequences of Lemma 3.

Lemma 7: τ_i dominates s_i if and only if for every pure strategy combination $r \in S$, $U_i(r \mid \tau_i) \geqslant U_i(r \mid s_i)$ with a strict inequality for some $r \in S$.

Lemma 8: τ_i dominates s_i if and only if for every $\sigma \in \mathring{M}$, $U_i(\sigma \mid \tau_i) > U_i(\sigma \mid s_i)$.

Theorem 4: If σ is a persistent strategy and $s_i \in S_i$ is dominated then $\sigma_i(s_i) = 0$.

Proof: Let R be an absorbing retract containing σ. Let $\hat{R}$ be the retract defined by $\hat{R}_j = R_j$ for $j \neq i$ and $\hat{R}_i = R_i \cap \{\tau_i \in \Delta(S_i): \tau_i(s_i) = 0\}$. It suffices to show that $\hat{R}_i$ is an absorbing retract. Let T be a neighborhood of R which is absorbed by R. We will show that $\hat{R}$ absorbs T. For every $\eta \in T$ there exists a sequence of $\{\eta^j\}_{j=1}^{\infty}$ with $\eta^j \to \eta$, and $\eta^j \in \mathring{M} \cap T$ for every j. For every $\delta_i^j \in BR_i(\eta^j) \cap R_i$ we have by Lemma 8, $\delta_i^j(s_i) = 0$. By the upper semi-continuity of BR_i we have $\delta_i \in BR_i(\eta) \cap R_i$ with $\delta_i(s_i) = 0$. Q.E.D.

Theorems 3 and 4 can be used to find persistent retracts in the following way.

Corollary 1: In every game there exists a persistent retract R of the form $\mathop{\times}\limits_{i=1}^{n} \Delta(\hat{S}_i)$ where $\hat{S}_i$ is a subset of undominated pure strategies of i, which contains no equivalent strategies.

Proof: Let F_i be a strategy selection for player i which contains one pure strategy from each of the equivalence classes in S_i. The retract $\mathop{\times}\limits_{i=1}^{n} \Delta(F_i)$ absorbs all of M and therefore is absorbing. Clearly it contains a persistent retract which by theorems 3 and 4 should be of the form described in the corollary. Q.E.D.

By Corollary 1 there exists a finite family of retracts of very simple form, from which one can pick up a persistent retract. Under the conditions specified in the following corollary, all persistent retracts are of that simple form and therefore there are only finitely many persistent retracts.

Corollary 2: If no player has equivalent strategies among his undominated strategies then every persistent retract is of the form $\bigtimes_{i=1}^{n} \Delta(\hat{S}_i)$ where $\hat{S}_i$ is a subset of undominated pure strategies of i.

We can further simplify the problem of finding all the persistent retracts by observing that under the condition of the last corollary every two persistent retracts must be disjoint.

Corollary 3: Suppose no player has equivalent strategies among his pure undominated strategies and R^1 and R^2 are two distinct persistent retracts then $R^1 \cap R^2 = \emptyset$.

Proof: For $q = 1, 2$ let $F^q = (F_1^q, F_2^q, \ldots, F_n^q)$ be the selection describing R^q. Then $F_i^q \subseteq S_i$ and F_i^q does not contain dominated strategies. Suppose contrary to the statement of the corollary that $\sigma \in R^1 \cap R^2$. It follows that for every player j, $F_j^1 \cap F_j^2 \neq \emptyset$. Consider the retract R defined by the selection

$$F = (F_1^1 \cap F_1^2, F_2^1 \cap F_2^2, \ldots, F_n^1 \cap F_n^2).$$

For $q = 1, 2$ let T^q be a neighborhood absorbed by R^q and let $T = T^1 \cap T^2$. T is a neighborhood of R.

Since R is a proper subset of the retracts R^1 and R^2 there is a $\tau \in R$ which is not absorbed by R. Hence there is a player i with $BR_i(\tau) \cap R_i = \emptyset$. On the other hand for $q = 1, 2, BR_i(\tau) \cap R_i^q \neq \emptyset$. So $M_i(\tau) \cap F_i^1 \cap F_i^2 = \emptyset$ whereas for $q = 1, 2, M_i(\tau) \cap F_i^q \neq \emptyset$. Using arguments similar to those in Lemma 5, we can find an open set $O \subseteq T$ such that for some $s_i^1 \in F_i^1$ and some $s_i^2 \in F_i^2$ with $s_i^1 \neq s_i^2$ we have and $M_i(\delta) \cap F_i^2 = \{s_i^2\}$ and $M_i(\delta) \cap F_i^1 = \{s_i^2\}$ for every $\delta \in 0$. But then s_i^1 and s_i^2 must be equivalent on the open set O and hence by Lemma 4 they are equivalent. This contradicts the assumption of the corollary. Q.E.D.

6. Perfect, Proper and Persistent Equilibria

We show in this section that every game has a proper equilibrium which is persistent. Since every proper equilibrium is perfect, every game has also a perfect persistent equilibrium. We follow the definitions given in *Myerson* [1978].

For a given $\epsilon > 0$ we say that a strategy $\sigma \in M$ is an *ϵ-proper equilibrium* if the following two conditions hold for every player i:

1. $\sigma_i(s_i) > 0$ for every $s_i \in S_i$, and
2. for every pair of pure strategies $s_i, \bar{s}_i \in S_i$ if

$$U_i(\sigma \mid s_i) > U_i(\sigma \mid \bar{s}_i) \text{ then } \epsilon\, \sigma_i(s_i) \geqslant \sigma_i(\bar{s}_i).$$

A strategy $\sigma \in M$ is a ***proper equilibrium*** if there exists a sequences ϵ_k and σ_k with $\epsilon_k \to 0$, σ_k being an ϵ_k proper equilibrium and $\sigma_k \to \sigma$.

Lemma 9: Let R be an absorbing retract. Then for sufficiently small $\epsilon > 0$ there exists an ϵ-proper equilibrium σ such that for some $\tau \in R$ $| \sigma_i(s_i) - \tau_i(s_i) | \leqslant \epsilon$, for each i and $s_i \in S_i$.

Proof: For $\epsilon > 0$ we denote

$$\delta = \min_{i \in N} \epsilon^{|S_i|+1} \Big/ \sum_{j=1}^{|S_i|} \epsilon^j,$$

and

$$B_i = B_i(\delta) = \{\sigma \in \Delta(S_i) : \min_{s_i \in S_i} \sigma(s) \geqslant \delta\},$$

and

$$B = B(\delta) = \mathop{\times}_{i=1}^{n} B_i(\delta).$$

We define a retract W by

$$W_i = \{\sigma_i \in \Delta(S_i) : \text{for some } \tau_i \in R_i, \max_{s_i \in S_i} | \sigma_i(s_i) - \tau_i(s_i) | \leqslant \epsilon\}.$$

We have $R \subseteq W$ and if ϵ is sufficiently small then R absorbs W.

We define a correspondence $F_i : W \to W_i$ by

$$F_i(\sigma) = \{\sigma_i^* \in W_i : \text{for every } s_i, \bar{s}_i \in S_i \text{ if } U_i(\sigma \mid s_i) > U_i(\sigma \mid \bar{s}_i) \mid \text{then } \epsilon\, \sigma_i^*(s_i) \geqslant \geqslant \sigma_i^*(\bar{s}_i)\}.$$

The correspondence $F : W \to W$ is defined by $F(\sigma) = \mathop{\times}_{i=1}^{n} F_i(\sigma)$. It is clear that a fixed point of F which also belongs to B would constitute the desired ϵ-proper equilibrium.

It is easy to see that $F(\sigma) \cap B$ is convex for every $\sigma \in W$. Also $F(\cdot) \cap B$ is upper semi-continuous. We need only to show that $F(\sigma) \cap B \neq \emptyset$ for each $\sigma \in W$.

Given $\sigma \in W$ we let $\tau \in BR(\sigma) \cap R$. We will perturb τ to generate a point η in $F(\sigma) \cap B$.

For every player i we order his pure strategies by a function $O : S_i \to \to \{1, 2, \ldots, |S_i|\}$ satisfying

1. O is one to one, and
2. for every $s_i, \bar{s}_i \in S_i$ if $U_i(\sigma \mid s_i) > U_i(\sigma \mid \bar{s}_i)$ then

$$O(s_i) < O(\bar{s}_i).$$

It is clear that if $s_i \in M_i(\sigma)$ and $\bar{s}_i \notin M_i(\sigma)$ then $O(s_i) < O(\bar{s}_i)$.
We define η_i by

$$\eta_i(s_i) = (1-\epsilon)\,\tau_i(s_i) + \epsilon\,\epsilon^{O(s_i)} \Big/ \sum_{j=1}^{|S_i|} \epsilon^j.$$

Clearly $\eta_i \in B_i$. Also by the construction of η_i it must satisfy the inequalities of $F(\sigma)$. Finally $|\eta_i(s_i) - \tau_i(s_i)| \leq \epsilon$ for every i and $s_i \in S_i$ so $\eta \in W$. Q.E.D.

Theorem 5: Every game contains a proper (hence perfect) Nash equilibrium which is also persistent.

Proof: By the previous lemma every persistent retract R has a sequence σ_k of ϵ_k-proper equilibrium which converges to R as $\epsilon_k \to 0$. We can find a subsequence of σ_k which converges to a point in R which must be a proper equilibrium. Q.E.D.

It is not the case however that every persistent equilibrium is proper or even perfect. Consider the following 3-person game.

		Player 3: L			*Player 3: H*	
		Player 2			Player 2	
		L	*H*		*L*	*H*
Player 1	*L*	0,0,0	0,0,1	*L*	0,1,0	1,0,0
	H	0,1,0	1,0,1	*H*	1,0,1	0,1,0

We claim that every strategy in this game is persistent. Yet the strategy (L,L,L) constitutes a Nash equilibrium which is not perfect.

To see that every strategy is persistent observe that this game has no equivalent strategies for any one of the players. Therefore by Corollary 2 every persistent retract consists of a finite number of pure strategies (and their convex combinations) for any one of the players. Let R be such a retract and consider first the case that $H \in R_3$. It follows immediately by inspection of the right hand side table of the game (the one corresponding to player 3 playing H) that for R_1 and R_2 to be Nash closed we must have $R_1 = R_2 = \{L, H\}$. But then (since Player 3 prefers (H,H,L) to (H,H,H)) $L \in R_3$ and $R_3 = \{L,H\}$. So every persistent retract R with $H \in R_3$ coincides with the entire game, and every strategy must be persistent.

So we are left to consider the case when $H \notin R_3$. It follows by the fact that R is a Nash retract that $(H,L,L) \notin R$, and therefore $(H,H,L) \notin R$ and therefore $(L,H,L) \notin R$.

It must be then that $R = \{(L,L,L)\}$. Now consider a mixed strategy σ described by

$$\sigma_1(L) = 1 - \epsilon_1 \qquad \sigma_1(H) = \epsilon_1$$
$$\sigma_2(L) = 1 - \epsilon_2 \qquad \sigma_2(H) = \epsilon_2$$
$$\sigma_3(L) = 1 - \epsilon_3 \qquad \sigma_3(H) = \epsilon_3.$$

Where each ϵ_i satisfy $0 < \epsilon_i < 1/3$. Clearly R must absorb such strategies for some sufficiently small ϵ_1, ϵ_2 and ϵ_3.

But consider player 1's payoff at σ

$$U_1(\sigma) = \epsilon_1 \epsilon_2 (1 - \epsilon_3) + \epsilon_1 (1 - \epsilon_2) \epsilon_3 + (1 - \epsilon_1) \epsilon_2 \epsilon_3$$
$$= \epsilon_1 (\epsilon_2 + \epsilon_3 - 3\epsilon_2 \epsilon_3) + c \quad \text{Where } c \text{ is independent of } \epsilon_1,$$
$$= \epsilon_1 [\epsilon_2 (1 - 3\epsilon_3) + \epsilon_3] + c.$$

The term inside the brackets is positive when $\epsilon_3 < 1/3$. So every best reply in a neighborhood of (L,L,L) must contain the strategy H for player 1 and $(H,L,L) \in R$, a contradiction. So we have seen that the only persistent retract is the entire game and every strategy is persistent.

The arithemtic done in the previous paragraph also shows that (L,L,L) could not be a perfect equilibrium and this completes the analysis of this example.

7. Unanimity Games

All of the examples discussed in the introduction are of the family of unanimity games (or diagonal games). These game were discussed by *Harsanyi* [1981]. A game is a unanimity game if it has the following special structure. All the players have the same set of strategies which we denote by C, (i.e. for each i, $S_i = C$). A combination of pure strategies in C^n is called *diagonal* if it is of the form $(c, c, \ldots, c)$ for some $c \in C$. We denote such a strategy by $\bar{c}$. For each combination of pure strategies s which is not diagonal and for each player i, $U_i(s) = 0$. (Our analysis would apply in the same manner if $U_i(s) = \alpha_i \neq 0$ without this utility normalization.) We call a combination of pure strategies s, *positive* if $U_i(s) > 0$ for each player i. Clearly if s is positive than s is diagonal.

Theorem 6: If a unanimity game has a positive combination of pure strategies then a strategy combination is persistent if and only if it is positive.

Proof: Suppose $\bar{c} = (c, \ldots, c)$ is a positive combination of pure strategies. We will show that the retract $\{\bar{c}\}$ absorbs a neighborhood of itself, for a given $1 > \epsilon > 0$ consider the set $T = \{\sigma \in \Delta^n(C) \mid \sigma_i(c) > 1 - \epsilon \text{ for } i = 1, \ldots, n\}$. For every $\sigma \in T$, $U_i(\sigma \mid c) = U_i(\bar{c}) \prod_{j \neq i} \sigma_j(c) \geqslant U_i(\bar{c})(1 - \epsilon)^{n-1}$.

For every $c' \neq c$, $U_i(\sigma \mid c') = U_i(\bar{c}') \prod_{j \neq i} \sigma_i(c') \leqslant U_i(\bar{c}') \epsilon^{n-1}$. If ϵ is small enough then c is the unique maximizer for each player at each point of T, which shows that $\{\bar{c}\}$ absorbs T. It follows then that $\{\bar{c}\}$ is a persistent retract.

Conversely, suppose that R is a persistent retract and $\sigma \in R$. Observe that since there is a positive combination of strategies no player has equivalent strategies among his undominated strategies. Thus we can evoke Corollary 2 to conclude that $R = \mathop{\times}\limits_{i=1}^{n} \Delta (C_i)$ where $C_i \subseteq C$ is a set of undominated strategies for the player i. Notice also, that if $\mathop{\times}\limits_{i=1}^{n} C_i$ contains a diagonal combination of strategies $\bar{c}_0$, then necessarily $\bar{c}_0$ is positive. We will show that R contains a diagonal combination of strategies $\bar{c}_0$ and therefore by the "if" direction of the theorem and because of the minimality of R, $R = \{\bar{c}_0\}$. Let $\bar{c}_0$ be a positive combination of strategies and assume that $s = (s_1, \ldots, s_n)$ is a combination of strategies in $\mathop{\times}\limits_{i=1}^{n} C_i$ which is not diagonal. For $0 < \epsilon < 1$ consider the strategy σ^ϵ defined by $\sigma_i^\epsilon = \epsilon c_0 + (1 - \epsilon) s_i$ Clearly $U_i (\sigma^\epsilon \mid c_0) = U_i (\bar{c}_0) \mathop{\times}\limits_{j \neq i} \sigma_j^\epsilon (c_0) > 0$. We assume first that $n > 2$. Suppose now that there are no $n - 1$ coordinates in s which are identical. Then for each i and for each $c' \neq c_0$, $U_i (\sigma^\epsilon \mid c') = U_i (\bar{c}') \mathop{\times}\limits_{j=i} \sigma_j^\epsilon (c') = 0$. It follows then that the only maximizer, for each i, at σ^ϵ is c_0, i.e. $\bar{c}_0 \in R$. Assume now that there are exactly $n - 1$ coordinates which are identical. Without loss of generality we may assume $s_2 = s_3 = \cdots = s_n = c'$ and $s_1 \neq c'$. The same argument as above shows that $c_0 \in C_i$ for $i \geqslant 2$. As for the first player, $U_1 (\sigma^\epsilon \mid c) = 0$ for each c for which $c \neq c_0$ and $c \neq c'$. Therefore either $c_0 \in C_1$ in which case $\bar{c}_0 \in \mathop{\times}\limits_{i=1}^{n} C_i$, or $c' \in C_1$ in which case $\bar{c}' \in \mathop{\times}\limits_{i=1}^{n} C_i$. The argument for $n = 2$ is similar. Q.E.D.

Applying theorem 6 to the games Γ_1, Γ_2, and Γ_4 discussed in the introduction we see that the persistent strategies coincide exactly with the solutions suggested there. In the matching pennies game, Γ_3, it is easy to see that the set of persistent strategies coincide with the entire set of mixed strategies of the game.

However, ruling out bad strategies will occur also in non-unanimity game. For example we could perturb the battle of the sexes game to make it nondiagonal while keeping the same best reply functions and the mixed strategy would still not be persistent.

References

Harsanyi, J.C.: Solutions for Some Bargaining Games Under the Harsanyi-Selten Solution Theory, Part II: Analysis of Specific Bargaining Games. CP-432 Center for Research in Management Science, University of California, Berkeley 1981.

Kreps, D.M., and *R. Wilson*: Sequential Equilibria. Graduate School of Business, Stanford University, Stanford 1980.

Kohlberg, E.: Some Problems with the Concept of Perfect Equilibrium. N.B.E.R., Conference on the Theory of General Economic Equilibrium. University of California, Berkeley 1981.

Myerson, R.: Refinements of the Nash Equilibrium Concept. Int. J. of Game Theory 7, 1978, 73–80.

Selten, R.: Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. Int. J. of Game Theory 4, 1975, 25–55.

Received May 1983.