EX ANTE EQUILIBRIUM:

STRATEGIC BEHAVIOUR AND THE PROBLEM OF PERFECTION

David G. Pearce

Econometric Research Program
Research Memorandum No. 301

June 1982

I am grateful to Dilip Abreu, Bob Anderson, Vijay Krishna and Hugo Sonnenschein for their invaluable assistance. All errors are the responsibility of the author.

Please send comments to David Pearce c/o Department of Economics, Princeton University, Princeton, NJ 08544.

Econometric Research Program
PRINCETON UNIVERSITY
207 Dickinson Hall
Princeton, New Jersey

I. INTRODUCTION

"What constitutes rational behaviour in a non-cooperative strategic situation?" This paper explores the issue in the context of a wide class of finite non-cooperative games in extensive form. The traditional answer relies heavily upon the idea of Nash equilibrium (Nash, 1951). The position developed here, however, is that as a criterion for judging a profile of strategies to be "reasonable" choices for players in a game, the Nash equilibrium property is neither necessary nor sufficient. Some Nash equilibria are intuitively unreasonable, and not all reasonable strategy profiles are Nash equilibria.

The fact that a Nash equilibrium can be intuitively unattractive is well-known: the equilibrium may be "imperfect." Introduced into the literature by Selten (1965), the idea of imperfect equilibria has prompted game theorists to search for a narrower definition of equilibrium. While this research, some of which will be discussed here, has been extremely instructive, it remains inconclusive. Theorists often agree about what should happen in particular games, but to capture this intuition in a general solution concept has proved to be very difficult. If this paper is successful it should make some progress in that direction.

The other side of the coin has received less scrutiny. Can all non-Nash profiles really be excluded on logical grounds? I believe not. The standard justifications for considering only Nash profiles are circular in nature, or make gratuitous assumptions about players' decision criteria or beliefs. This will be discussed in Section II.

Most of the paper is devoted to the development and evaluation of an alternative solution concept, which I call "ex ante equilibrium." It is offered as an answer to my opening question: "What constitutes rational behaviour in a non-cooperative strategic situation?" No attempt is made to single out a unique strategy profile for each game; instead, a profile is considered to be in ex ante equilibrium if each player has selected any strategy that is "reasonable" in a sense to be made precise. A single player might have many such strategies.

While allowing for more flexibility than the Nash solution concept permits, the alternative proposal attempts to eliminate the problem of imperfection. This is complicated by the fact that there are actually two types of behaviour that have been labelled "imperfect" in the literature. The first involves "implausible behaviour at unreached information sets," and arises only in games having some sequential nature. The second is intimately related to the first, but can occur even in perfectly simultaneous games. It concerns the taking of risks that seem "likely" to be costly, when there are no offsetting advantages for a player to consider. The first type of imperfection can be ruled out on the basis of rather innocuous rationality postulates. Elimination of the second type, however, requires an additional assumption, amounting to the assertion that players will exercise prudence when it is costless to do so. Accordingly, I define two solution concepts. The first, ex ante equilibrium, relies upon little more than logical deduction, and ignores the second type of imperfect behaviour. A narrower solution concept, which I call perfect ex ante equilibrium, makes the additional assumption needed to eliminate imperfections of the second type.

For expositional purposes the early sections of the paper deal only with normal form representations of games. Because I believe that the additional structure provided by the extensive form is often important in determining how players will act, I interpret a normal form game as a convenient representation of a perfectly simultaneous game, in which no one can observe any move of any other player before moving himself. (See Note 1.) Such games can be analyzed without the encumbrance of the extensive form structure. The analysis of Sections III and IV should be understood as an investigation of a special class of extensive form games. Indeed, the general solution concepts ultimately proposed in Sections VI and VII reduce to those of Sections III and IV for games in which everyone moves simultaneously. Many of the central themes of the paper come across more clearly in these special games.

The strong influence that a number of papers on imperfect equilibria have had on the work reported here will be evident to the reader. Emphasis is given to those ideas in the literature that I consider crucial, as well as those with which I must take issue. A less obvious, but major intellectual debt should be recorded here, one that I owe to my colleague Dilip Abreu. Our countless discussions on game theory have played a central role in shaping my ideas about strategic behaviour. Of course only I can be held responsible for the statements made in this paper.

II. A BRIEF CRITIQUE OF NASH EQUILIBRIUM

In the literature one can find many alternative justifications for concentrating upon Nash equilibria, and conceivably others could be concocted. To attempt to prove that all such justifications must be inadequate might be overly ambitious. Instead, this section examines a few of the best-known arguments in favour of the Nash solution concept, and concludes that they are less than compelling.

The discussion here concerns finite N-person non-cooperative normal form games of complete information. Such a game

$$G = (S^1, \ldots, S^N;\ U^1, \ldots, U^N)$$

is completely characterized by the finite nonempty sets S^i of pure strategies, and real-valued utility functions U^i having domain $\prod_{i=1}^{N} S^i$. For each i, there are $k(i)$ pure strategies $s^i_1, \ldots, s^i_{k(i)}$. A mixed strategy for i is a probability distribution over S^i represented by a vector $p^i = (p^i_1, \ldots, p^i_{k(i)})$ in $k(i)$-dimensional Euclidean space. The components of p^i satisfy $0 \leq p^i_j \leq 1$ and sum to unity. Denote the i^{th} player's set of mixed strategies by M^i. The utility functions are extended to mixed strategy profiles by a straightforward expected utility calculation: for any profile $p^1, \ldots, p^N$

$$U^i(p^1, \ldots, p^N) = \sum_{j=1}^{k(1)} \cdots \sum_{\ell=1}^{k(N)} p^1_j \cdots p^N_\ell \, U^i(s^1_j, \ldots, s^N_\ell) .$$

A strategy profile $(\sigma^1, \ldots, \sigma^N)$ is a <u>Nash equilibrium</u> of G if for every i and all $\bar{\sigma}^i \in M^i$

$$U^i(\sigma^1, \ldots, \sigma^N) \geq U^i(\sigma^1, \ldots, \sigma^{i-1}, \bar{\sigma}^i, \sigma^{i+1}, \ldots, \sigma^N) .$$

Thus no player can improve his position by deviating unilaterally from a Nash equilibrium.

On what grounds does a theory of non-cooperative strategic behaviour single out Nash equilibrium profiles? In the context of two-person zero-sum games, Luce and Raiffa (1957) specify ". . . a demand to be met by any theory of strictly competitive games . . . the mere knowledge of the theory should not cause either of the players to change his choice . . ." (page 63). Similarly, in the case of N-person games, Kreps and Wilson (forthcoming) remark that although there have been many motivations of Nash equilibrium, ". . . a thread common to all of them is that *if* players are to arrive at some 'agreed-upon' mode of behavior, then it is *necessary* that this behavior constitutes a Nash equilibrium. Otherwise, some player would in his own self-interest defect from the agreement." (Section 3). Since communication amongst players is forbidden, the mode of behaviour cannot be agreed upon explicitly, but rather must be obvious to all players because of some compelling features of the game.

A tacit assumption underlying Luce and Raiffa's requirement is that a theory of games must associate with each game, a *single* strategy profile σ that is supposed to describe how players behave. Such a theory enables the game theorist to predict the outcome, and each player can presumably do the same. Naturally the prediction is then sensible only if each player is responding optimally to the predicted strategies of others. This is also the thrust of Kreps and Wilson's remarks. But why *should* we require a theory of games to predict unique strategic choices for all players, with certainty? To assume a priori that all strategy profiles but one are excludable, rather than deriving this by some logical process, makes the argument circular and inconclusive.

Defenders of the Nash solution concept will insist that it must be possible for players to predict the strategic choices of others. After all, each player knows "all there is to know" about his opponents; everyone knows the rules of the game. To predict player j's choice of strategy, player i simply puts himself in j's position, and imagines what i himself would do in that situation. Whatever i would do (if he were given j's strategy space and utility function) i assumes j

will actually do. I claim that this argument too is circular. It assumes that there is a unique strategy that i would prefer if he were put in j's position. Suppose instead that there are possibly many choices that each player could reasonably make. In an attempt to resolve his own uncertainty, i imagines himself in j's position. To his dismay, he realizes that if he were j, he would be torn in several directions, having a number of reasonable-looking strategies from which to choose. This mental exercise simply pushes the problem back one step, where it is encountered afresh. None of this depends upon whether or not the game in question has a unique Nash equilibrium; for those who wish to make single-valued predictions regarding strategy profiles, the occurrence of multiple Nash equilibria is just an added embarrassment.

Sometimes one hears the following opinion expressed: although Nash equilibrium may not be achieved in a single play of a game, behaviour will converge eventually to a Nash profile. Using various "rules of thumb" for dynamic adjustments by players, Bernheim (1981) shows that convergence occurs only under very restrictive conditions. (See Note 2.) In any case, convergence arguments cannot justify the use of the Nash equilibrium concept in "one-shot" games. On the other hand, if a game repeated fifty times, for example, is really the object of interest, this defines one overall game that can be decomposed into its fifty natural component games. When this large game is analyzed, typically it will still be necessary to make arbitrary assumptions about behaviour in order to ensure Nash behaviour in each component game. This cannot be a satisfactory defense of Nash equilibrium.

In the case of two-person zero-sum games, it is supposed that the maximization of each player's "security level" (the least expected utility a player can receive when he plays a certain strategy, regardless of what his opponent does) provides an independent reason for Nash equilibrium to arise. I believe that it is only because "maximin" behaviour turns out to have the "best response" property, that it receives any attention at all. In N-person games maximin behaviour does not have this property, and no one mentions maximizing security levels in such games.

Only a pathologically pessimistic person would care about nothing except his security level; our players, expected utility maximizers who accept risk willingly, do not fit this description. One of our players would maximize his security level only if he were convinced that his opponent could guess his move, and then inflict maximal damage on him. In a two-person zero-sum game, the opponent at least has an incentive to do so. But one still needs the assumption that each player can predict with certainty the strategy employed by the other. There is no independent argument here after all. For an early protest against the orthodox theory of two-person zero-sum games, the reader should consult the provocative paper by Ellsberg (1956).

Perhaps the central position that Nash equilibrium occupies in the theory of games is partly explained by the notion that in "equilibrium," everyone must be satisfied with his own choice. It is an error to draw from this requirement the conclusion that the outcomes of games should have a "no regret" property, that is, no one wishes, after the game is played, that he had chosen differently. The error creeps in because in the statement "everyone must be satisfied with his own choice," no time is specified. As long as no one wants to change his strategy at the time when he is required to commit himself, the relevant requirement is met. The possibility that someone may later be disappointed upon discovering the strategies actually chosen by others, does not make his choice less reasonable- he must act with limited information. Ex ante, before knowing the strategies of other players, a player makes a decision; ex post, he may regret it. No game theorist requires that a player have no regrets about the action he chose, given the actions (particular realizations of mixed strategies) of others—such an equilibrium need not even exist. The actions are not known to the player in question when he must make his decision. Similarly, if other players' strategies are not known when decisions are made, the possibility of ex post regret must be considered natural.

The discussion above deals exclusively with the notion that the attainment of Nash equilibrium must not be considered a necessary condition for judging players' behaviour "reasonable." Less controversial is the reciprocal contention that some Nash equilibria can themselves represent unreasonable behaviour, because they are imperfect. There is no need to belabour the latter point; many illustrative examples will arise in subsequent sections dealing specifically with perfection.

III. EX ANTE EQUILIBRIUM IN THE NORMAL FORM

The doubts expressed above regarding the Nash equilibrium concept suggest an alternative approach to the analysis of a game. Rather than require that a theory predict a unique strategic choice for each player, this approach isolates a set of strategies for each player that cannot be excluded on the basis of a few underlying assumptions, interpreted as basic rationality postulates. Such an exercise is carried out in this section for simultaneous games represented by their normal forms. In Section IV a slightly more restrictive theory is proposed to deal with imperfect behaviour in simultaneous settings. The analysis of these two sections is extended to extensive form games in Sections VI and VII.

In this view, a theory of N-person non-cooperative games associates with each such game

$$G = (S^1,\ldots,S^N;\ U^1,\ldots,U^N)$$

a vector of sets $(E^1,\ldots,E^N)$, where $E^i \subseteq M^i\ \forall\ i$, in the notation of the previous section. The vector $(E^1,\ldots,E^N)$ is the <u>solution</u> of the game G, and each profile $(e^1,\ldots,e^N)$ with $e^i \in E^i\ \forall\ i$ is an <u>equilibrium</u> of G. The interpretation given to E^i is that it contains all those strategies (pure or mixed) available to player i that do not contradict the underlying assumptions upon which the theory is based. There is no presumption that these sets will always be singletons; they may not even be proper subsets of the original sets M^i. Isolation of a particular strategy for each individual may occur for certain games, but in those cases it is an implication of the analysis, not a restriction imposed arbitrarily. The formulation of the problem does, however, automatically force the solution concept to exhibit the property of <u>interchangeability</u>: if $(e^1,\ldots,e^N)$ and

$(f^1,\ldots,f^N)$ are equilibria of G, and $(g^1,\ldots,g^N)$ satisfies $g^i \in \{e^i,f^i\} \; \forall \; i$, then $(g^1,\ldots,g^N)$ is also an equilibrium of G. Given that players have no way of communicating with one another, it does not seem reasonable to allow a theory to violate this property.

The particular solution concept I wish to propose retains an implicit assumption present in the traditional theory of games: each player's goals and strategic possibilities are "common knowledge" (see Note 3). Essentially this means that for any piece of information θ involving the structure of the game, any statement of the form "i knows that j knows that . . . k knows θ" is true, for all players i, j,...,k .

Another feature of standard non-cooperative theory is the assumption that each player maximizes expected utility subject to some "point expectation" regarding all other players' strategies. While preserving the expected utility-maximizing behaviour of the participants, the theory forwarded here does not, of course, assume that players are certain about the strategies of others. Instead, they form subjective probability distributions over their opponents' strategies, and maximize expected utility given those distributions. This manner of dealing with uncertainty is also taken to be common knowledge.

Before using the basic notions outlined above to motivate the particular requirements to be imposed upon the solution sets, I need to introduce some notation and terminology. For any set C which is a subset of a finite-dimensional Euclidean space $\mathbb{R}^n$, define the convex hull of C by

$$\overline{C} = \{c \in \mathbb{R}^n : c = \sum_{i=1}^{r} a_i c_i \;,\; \sum_{i=1}^{r} a_i = 1,\; c_i \in C \text{ and } a_i \geq 0 \;\; \forall \; i\} .$$

In a game $G = (S^1,\ldots,S^N;\, U^1,\ldots,U^N)$ a strategy $r \in M^i$ is <u>strongly dominated</u> if $\exists t \in M^i$ such that for every strategy profile $(m^1,\ldots,m^N)$, it is the case that

$$U^i(m^1,\ldots,m^{i-1},t,m^{i+1},\ldots,m^N) > U^i(m^1,\ldots,m^{i-1},r,m^{i+1},\ldots,m^N) .$$

If the above inequality is replaced by a weak inequality, with strict inequality holding for at least one profile $\bar{m}^1,\ldots,\bar{m}^N$, then r is weakly dominated by t . A strategy $b \in M^i$ is a best response for i to a profile $(m^1,\ldots,m^N)$ if $\forall\ d \in M^i$

$$U^i(m^1,\ldots,m^{i-1},b,m^{i+1},\ldots,m^N) \geq U^i(m^1,\ldots,m^{i-1},d,m^{i+1},\ldots,m^N).$$

Similarly, given any subset $B^i \subseteq M^i$, $b^i \in B^i$ is a best response in B^i to $(m^1,\ldots,m^N)$ if $\forall\ d \in B^i$

$$U^i(m^1,\ldots,m^{i-1},b,m^{i+1},\ldots,m^N) \geq U^i(m^1,\ldots,m^{i-1},d,m^{i+1},\ldots,m^N) .$$

One wishes to exclude from the sets $E^1,\ldots,E^N$ precisely those strategies that "rational" players could never choose, given that the structure of the game, and the way uncertainty is dealt with, are common knowledge amongst all players. Consider player 1's problem. While he may be unsure about exactly which mixed strategies his opponents will choose, those that can be ruled out on logical grounds are not in $E^2,\ldots,E^N$, and can be ignored. For example, some strategies of player 2 might be strongly dominated, so that 1 could eliminate them from consideration. According to the behavioural assumptions explained above, player 1 will form some subjective probability distribution over each set E^i, $i=2,\ldots,N$, and then maximize expected utility subject to this "conjecture" about other players' strategies. Formally, a conjecture on a set of strategies $L^i \subseteq M^i$ is a probability measure μ on the Borel sets of L^i . But since each element of L^i is a (mixed) strategy, μ induces a probability distribution over pure strategies. Thus for the purposes of expected utility maximization, μ can be identified with a mixed strategy which is a "weighted average" of points in L^i , the weights being determined by the measure chosen. This will not necessarily lie in L^i unless L^i is convex; but the "weighted average" will be some element of $\bar{L}^i$.

Thus, a strategy m^1 can reasonably be chosen by 1, and hence be included in E^1, only if it is a best response to some point in $\prod_{i=1}^{N} \bar{E}^i$. The same argument applies to each of the other players, so it is appropriate to impose the following condition upon the solution sets $E^1,\ldots,E^N$:

if $e^i \in E^i$, then e^i is a best response to some profile $(m^1,\ldots,m^N) \in \prod_{r=1}^{N} \bar{E}^r$, $\forall$ i.

Call this property the <u>best response property</u> of the solution sets.

Player 1's choice of some $m^1 \in E^1$ is "justified" by the existence of a profile in $\prod_{i=1}^{N} \bar{E}^i$ to which m^1 is a best response. A natural suggestion at this point might be that a higher order "best response property" is also appropriate: if 1's conjecture assigns positive probability to some mixed strategy m^2 for 2, say, then m^2 should in turn be justifiable for 2, or else 1's conjecture would not be sensible. But 1's conjecture is over strategies in E^2, all elements of which can be justified, because of the requirement that $E^1,\ldots,E^N$ have the best response property. Similarly, those strategy profiles that justify the strategies that 1 thinks 2 might choose, are in $\prod_{i=1}^{N} \bar{E}^i$, and can be justified themselves. This process of successive justifications can continue indefinitely; in fact, for any $m^i \in E^i$, one can find an infinite succession of conjectures "supporting" m^i.

For a given game G, there will typically exist many sets $X^1,\ldots,X^N$ satisfying the best response property. For example, suppose that G has three Nash equilibria: $r = (r^1,\ldots,r^N)$, $s = (s^1,\ldots,s^N)$ and $t = (t^1,\ldots,t^N)$. Then the singleton sets $\{r^1\},\ldots,\{r^N\}$ have the best response property, by the definition of a Nash equilibrium. So do the sets $\{s^1\},\ldots,\{s^N\}$ and the sets $\{t^1\},\ldots,\{t^N\}$. I wish to argue that the solution sets $E^1,\ldots,E^N$ should be "at least as large" as any sets $X^1,\ldots,X^N$ having the best response property, that is $X^1 \subseteq E^1,\ldots,X^N \subseteq E^N$. If not, the theory is asserting that for one of the X^i, some $x \in X^i$ is an irrational choice for i. But i can insist that there is an infinite succession of conjectures (of the type outlined in the previous paragraph) supporting the choice x . At each stage, these conjectures are in accordance with what player i knows about the game: they involve players' forming probability distributions over other players' strategies, and maximizing expected utility on the basis of those distributions.

This observation suggests that the solution sets $E^1,\ldots,E^N$ be defined, for a given G , as follows: $E^i = \{x \in M^i : \exists X^1,\ldots,X^N$ with the best response property, and $x \in X^i\} \forall i$. That is, E^i should be the union of all sets X^i which are part of a vector of sets having the best response property. It is necessary to show that if defined in this manner, the sets $E^1,\ldots,E^N$ themselves have the best response property. $x \in E^i \Rightarrow \exists X^1,\ldots,X^N$ with the best response property, and $x \in X^i$. Then x is a best response to some element of $(\bar{X}^1,\ldots,\bar{X}^N)$, so x is a best response to something in $(\bar{E}^1,\ldots,\bar{E}^N)$, because $\bar{X}^i \subseteq \bar{E}^i \forall i$. This establishes the result. Consequently it is meaningful to describe the sets $E^1,\ldots,E^N$ as the "largest" sets having the best response property. With the E^i defined in this way, the vector $(E^1,\ldots,E^N)$ is the <u>ex ante solution</u> of G , and each $(e^1,\ldots,e^N)$ with $e^i \in E^i \forall i$ is an <u>ex ante equilibrium</u>.

Before remarking upon some of the properties of the ex ante solution, including the nonemptiness of the sets E^i , I shall present an example, in the hope of clarifying the reasoning upon which the solution is based. Below is the payoff matrix of a two-person game G_1 . Each entry in the matrix is a pair, with the first and second coordinates giving the payoffs (utilities) to players 1 and 2 respectively.

2

		β_1	β_2
1	α_1	(0,5)	(-1,3)
	α_2	(0,0)	(-1,3)

G_1

1 has pure strategies α_1 and α_2 , and 2 has pure strategies β_1 and β_2 . Let (p_1,p_2) denote a mixed strategy chosen by 1, and (q_1,q_2) a mixed strategy

for 2. As is customary, I identify a pure strategy such as α_1 with the mixed strategy (1,0) in M^1. It is easily verified that there are three types of Nash equilibrium strategy profiles of G_1. First, any profile (m^1,m^2) of the form $m^1 = (p_1,1-p_1)$ with $p_1 \geq \frac{3}{5}$, and $m^2 = (1,0)$ is a Nash equilibrium. Player II is content to choose β_1 as long as 1 plays α_1 with fairly high probability. In the second type of equilibrium, 1 chooses $p_1 = \frac{3}{5}$, leaving 2 indifferent between his pure strategies: $(m^1,m^2) = ((\frac{3}{5},\frac{2}{5}), (q_1,1-q_1))$ with $q_1 \in [0,1]$. Finally, when 1 sets $p_1 < \frac{3}{5}$, 2 prefers β_2, resulting in Nash equilibria of the third type: $(m^1,m^2) = ((p_1,1-p_1), (0,1))$ with $p_1 < \frac{3}{5}$.

Which of these equilibria are "reasonable"? It is unclear that there is any one equilibrium that should be singled out as "the" solution of the game. Noting that ((1,0),(1,0)) Pareto dominates all other equilibria (gives each player at least as much utility as any other equilibrium, and for each alternative equilibrium, gives some player strictly greater utility than the alternative), some game theorists might wish to single out ((1,0),(1,0)) as the solution to G_1. Opposition is bound to come from others who will insist that in the face of player 1's indifference between α_1 and α_2 (regardless of 2's strategic choice), 2 should consider it equally likely that α_1 or α_2 will be played. 2 would then choose β_2, which is not his strategy in the Pareto efficient equilibrium.

There seems to be no compelling reason for 1 not to choose (1,0) or (0,1) or any mixture whatsoever. Similarly, 2 could sensibly choose any strategy. Think of each player handing to a "gamesman" an instruction, which is a choice of a pure strategy, or an instruction to randomize in a certain way. Given that α_1 is reasonable for 1, and β_2 is among the sensible choices for 2, should we be astonished to observe, in some actual play of the game, that 1 and 2 have given the gamesman the instructions (1,0) and (0,1) respectively? This is not a Nash equilibrium, but not in the least an unlikely occurrence.

The ex ante solution of G_1 is simply the pair (M^1,M^2) of the original mixed strategy sets, since these have the best response property. Every $m^1 \in M^1$ is a best response to anything in M^2, and $(q_1,1-q_1) \in M^2$ is a best response to $(\frac{3}{5},\frac{2}{5}) \in M^1$, for any probability q_1. This reflects the fact that each player can give a sensible justification for any of the choices open to him.

The preceding example shows how large the ex ante solution sets can be, and illustrates that seemingly attractive guidelines such as favouring Pareto efficient profiles or assigning equal probability to two strategies for i that always give i the same utility, may be in conflict with one another. A limitation of the example, however, is that every pure strategy that was given positive weight in the ex ante solution sets, occurs with positive probability in some Nash equilibrium. The next example demonstrates that this need not be true.

Let G_2 be defined by the following matrix of payoff pairs:

$$\begin{array}{cc|c|c|c|c}
 & & & 2 & & \\
 & & \beta_1 & \beta_2 & \beta_3 & \\
\hline
 & \alpha_1 & (1,-1) & (-1,1) & (0,\frac{1}{2}) & \\
\hline
1 & \alpha_2 & (-1,1) & (1,-1) & (0,\frac{1}{2}) & G_2 \\
\hline
 & \alpha_3 & (\frac{1}{2},0) & (\frac{1}{2},0) & (\frac{1}{2},-1) & \\
\hline
\end{array}$$

Let (p_1,p_2,p_3) and (q_1,q_2,q_3) represent mixed strategies of 1 and 2 respectively. Since the payoff to 1 is $p_1(q_1-q_2) + p_2(q_2-q_1) + \frac{1}{2}p_3$ it is never the case that positive weight is given simultaneously to α_1 and α_2; this would require that $q_1-q_2 \geq \frac{1}{2}$, and $q_2-q_1 \geq \frac{1}{2}$, a contradiction. Hence in a Nash equilibrium, 1's strategies must be of the form $(p_1, 0, 1-p_1)$ or $(0, p_2, 1-p_2)$.

Without loss of generality, suppose $p_2 = 0$. If $p_1 > 0$, 2's best response would be $(0, 1, 0)$ but this is not a Nash equilibrium; $p_2 = 1$ would be 1's best response. Thus player 1's Nash equilibrium strategy is the pure strategy α_3. 2's best response to this is anything of the form $(q_1, 1-q_1, 0)$. Conversely, α_3 is a best response to $(q_1, 1-q_1, 0)$ whenever $\frac{1}{4} \leq q_1 \leq \frac{3}{4}$. This establishes that the Nash equilibria of G_2 are all pairs of the form $((0,0,1), (q_1, 1-q_1, 0))$, with $\frac{1}{4} \leq q_1 \leq \frac{3}{4}$.

By contrast, it is easy to show that α_1 and α_2, for example, are in 1's ex ante solution set. Consider the sets $T^1 = \{\alpha_1, \alpha_2\}$, and $T^2 = \{\beta_1, \beta_2\}$. T^1, T^2 have the best response property because α_i is a best response to β_i, $i=1,2$, and β_2 and β_1 are best responses to α_1 and α_2 respectively. By definition, the ex ante solution sets E^1 and E^2 contain T^1 and T^2. This demonstrates that a strategy such as α_1 that is never given positive weight in any Nash equilibrium can nonetheless be an element of an ex ante solution set. Further calculations show that for G_2, $M^i = E^i$, $i=1,2$. This does not mean that considerations of a more speculative nature could not narrow down the set of likely outcomes; it just says that no configuration of mixed strategies, if chosen by the players, would contradict their rationality, or throw into question their knowledge of the structure of the game.

Nash (1951) proved that any finite N-person non-cooperative game (of complete information) has at least one equilibrium, say $(e^1, \ldots, e^N)$. Since the sets $\{e^1\}, \ldots, \{e^N\}$ have the best response property, $e^i \in E^i \; \forall \; i$. This establishes nonemptiness of these sets, and the existence of an ex ante equilibrium in every finite game. An alternative elementary proof of existence will be given later.

For any game, the solution sets $E^1, \ldots, E^N$ have a particularly simple structure. If some pure strategies $s_1, \ldots, s_\ell$ are all given positive weight by some mixed strategy $m^i \in E^i$, then every strategy that is a mixture (including degenerate mixtures such as s_j and so on) over strategies $s_1, \ldots, s_\ell$ also appears in E^i. Since $m^i \in E^i$ is a best response to some profile $\gamma \in \prod_{r=1}^{N} E^{-r}$, each pure strategy given positive weight by m^i must also be a best response to γ and any convex combination of these strategies is a best response to γ. (These statements are immediate consequences of the fact that for any fixed profile of opponents' strategies, U^i is linear in the probabilities with which i's strategies are employed).

A set A of mixed strategies has the <u>pure strategy property</u> if $\alpha \in A$ implies that all pure strategies given positive weight by α are in A. Each of the solution sets E^i satisfies this property and is nonempty, and hence an ex ante equilibrium in pure strategies always exists. The need for players to randomize in many Nash equilibria has long been considered somewhat puzzling (see for example the discussion in Luce and Raiffa (1957), pages 74-76). The incentive for randomization seems to be the need to "evade" one's opponents. But in the present context, the opponents are not always able to figure out a player's strategic choice; he can hide without randomizing, camouflaged by the uncertainty of the other players.

A particularly useful characterization of the ex ante solution sets is available. Consider arbitrary nonempty closed sets $H^i \subseteq M^i$, $i=1,\ldots,N$, each satisfying the pure strategy property. Let $H^i(0) = H^i \ \forall \ i$, and for each positive integer t define $H^i(t)$ recursively by $H^i(t)$ = the set of all strategies in $H^i(t-1)$ that are best responses among strategies in $H^i(t-1)$ to some element of $\prod_{r=1}^{N} \overline{H}^r(t-1)$. If at stage t-1 the sets $H^i(t-1)$ are nonempty, closed, and have the pure strategy property, those properties are satisfied by the sets $H^i(t)$ also. To establish nonemptiness, choose any element δ of $\prod_{r=1}^{N} \overline{H}^r(t-1)$; a best response to δ can be found because the continuous function U^i attains a maximum on the nonempty compact set $H^i(t-1)$. Next, notice that if $\alpha^\ell \in H^i(t)$, $\ell=1,2,\ldots$ and α^ℓ converges to α, then there exists an integer V such that $W \geq V$ implies that α^W gives positive weight to (at least) all those pure strategies given positive weight by α. Then since α^V is a best response to some $\gamma \in \prod_{r=1}^{N} \overline{H}^r(t-1)$, so is α. Furthermore $\alpha \in H^i(t-1)$, because $H^i(t-1)$ is closed by hypothesis. Therefore $\alpha \in H^i(t)$ and $H^i(t)$ is closed. Finally, suppose $\beta \in H^i(t)$ gives positive weight to pure strategies $\beta_1,\ldots,\beta_e$. Then β is a best response to some $\rho \in \prod_{r=1}^{N} \overline{H}^r(t-1)$, and $\beta_1,\ldots,\beta_e$ are best responses also.

But $\beta \in H^i(t-1)$ implies $\beta_1,\ldots,\beta_e$ are elements of $H^i(t-1)$, and as they are best responses to ρ, $\beta_1,\ldots,\beta_e$ are in $H^i(t)$. This establishes the pure strategy property. Thus the fact that the original sets $H^1,\ldots,H^N$ are nonempty, closed and have the pure strategy property ensures that at every stage t, $H^1(t),\ldots,H^N(t)$ have these three properties.

$H^i(t+1)$ can differ from $H^i(t)$ only if for some j, $\overline{H}^j(t) \neq \overline{H}^j(t-1)$. But since $H^j(t)$ and $H^j(t-1)$ both satisfy the pure strategy property, their convex hulls differ only if some pure strategy in $H^j(t-1)$ is absent from $H^j(t)$. In other words, the iterative procedure "stops" unless some pure strategies are eliminated at each stage. Since there are only $\sum_{r=1}^{N} k(r)$ pure strategies in total, of which at most $\sum_{r=1}^{N} k(r) - N$ can be removed in this manner, $H^i(t) = H^i(k)$ $\forall$ $t \geq k = \sum_{r=1}^{N} k(r) - N$. Let $H = (H^1,\ldots,H^N)$ and define $D^i(H) = H^i(k)$ $\forall$ i. In particular, consider the sets $D^1(M),\ldots,D^N(M)$, where $M = (M^1,\ldots,M^N)$. These are of special interest because they coincide with the ex ante solution sets $E^1,\ldots,E^N$.

By construction the sets $D^1(M),\ldots,D^N(M)$ have the best response property (see Note4). Therefore $D^i(M) \subseteq E^i$ $\forall$ i. Conversely, if $\alpha \in E^i$, then α is a best response to some element of $\prod_{r=1}^{N} \overline{E}^r$ and hence is a best response to an element of $\prod_{r=1}^{N} \overline{M}^r(1)$. Therefore $\alpha \in M^i(2)$, and $E^i \subseteq M^i(2)$ $\forall$ i. k-fold repetition of this argument. establishes that

$$E^i \subseteq M^i(k) = D^i(M) .$$

In the special case of two-person games, this iterative procedure is identical to the iterative removal of strongly dominated strategies. Iterative dominance procedures have long been a part of the game theoretic literature (see Gale (1953), Farquharson (1957/1969), and Luce and Raiffa (1957), as well as the more recent work of Moulin (1979)). In two-person games, a strategy is strongly dominated if and only if it fails to be a best response to any strategy of the opponent (see Appendix). This equivalence does <u>not</u> hold for $N \geq 3$.

In games with more than two players, the ex ante solution sets may be smaller than (and are always contained by) those resulting from the iterative removal of strongly dominated strategies. This is due to the fact that even in the first round, strategies may be undominated and yet fail to be best responses to any prior over opponents' strategies, as the game G_3 shows. There are three players, the first of whom has pure strategies α_1, α_2, and α_3. 2 and 3 each play H or T. Only the payoffs to 1 matter here; those are shown below:

	HH	HT	TH	TT
α_1	6	6	6	6
α_2	10	10	10	0
α_3	0	10	10	10

G_3

In this game there is no conjecture over 2's and 3's strategies, for which α_1 would be a best response. But α_1 is neither strongly nor weakly dominated (the numerical calculations are omitted in the interests of brevity). It is the inability of 2 and 3 to coordinate their random choices that makes this example work.

The principal drawback of the ex ante solution is clear: it typically does not allow a specific prediction to be made about strategic choice. But this indeterminacy is an accurate reflection of the difficult situation faced by players in a game. The rules of a game and its numerical data are seldom sufficient for logical deduction alone to single out a unique choice of strategy for each player. To do so one requires either richer information (such as institutional detail or perhaps historical precedent for a certain type of behaviour) or bolder assumptions about how players choose strategies. Putting further restrictions on strategic choice is a complex and treacherous task. But one's intuition frequently points to patterns of behaviour that cannot be isolated on the grounds of consistency alone. Formalizing this intuition in specific solution concepts would seem to be a matter of high priority; I interpret papers such as Harsanyi (1976) to be in this spirit.

IV. PERFECT EX ANTE EQUILIBRIUM IN THE NORMAL FORM

The notion of an imperfect equilibrium was originally conceived (see Selten (1965)), and is still most commonly perceived, as a problem arising because of "implausible behaviour at unreached information sets." This is obviously applicable only to extensive form games, which are treated in later sections. But a related phenomenon appears in normal form games, and has received some attention. In particular the paper by Myerson (1978) on perfect and proper equilibria concerns exactly this issue.

Myerson's opening example is perhaps the simplest illustration of the problem at hand.

2

		β_1	β_2
1	α_1	(1,1)	(0,0)
	α_2	(0,0)	(0,0)

G_4

G_4 has two Nash equilibria. In the first, I and II select the pure strategies α_1 and β_1 respectively. In the second, they choose α_2 and β_2 respectively. The latter equilibrium is, as Myerson indicates, counterintuitive: "it would be unreasonable to predict (α_2,β_2) as the outcome of the game. If player 1 thought that there was any chance of player 2 using β_1, then 1 would certainly prefer α_1." (Myerson (1978), page 74). It is clear that 1 is taking an unnecessary risk by choosing α_2. He has nothing to gain by doing so, and possibly something to lose. The same applies to player 2, who would be foolish to choose β_2.

Explanations of why a certain equilibrium is to be considered "imperfect" usually involve stories about players making mistakes with small positive probabilities. This is a departure from tradition in the theory of games, and one senses a certain reluctance in Selten's remarks: "There cannot be any mistakes if the players are absolutely rational. Nevertheless, a satisfactory interpretation of equilibrium points in extensive games seems to require that the possibility of mistakes is not completely excluded. This can be achieved by a point of view which looks at complete rationality as a limiting case of incomplete rationality." (Selten, (1975), Section 7). The same reasoning is employed in normal form games, and Myerson concludes his commentary on the game G_4 by saying that ". . . there is always a small chance that any strategy might be chosen, if only by mistake. So in our example, α_1 and β_1 must always get at least an infinitesimal probability weight, which will eliminate (α_2, β_2) from the class of perfect (and proper) equilibria." (Myerson (1978), page 74).

I do not believe that the "slight mistakes" story does justice to our intuition about how players make their decisions. In game G_4 , if 1 prefers α_1 to α_2 , it is <u>not</u> because he believes that 2 might "make a mistake" and play β_1. On the contrary, β_1 would be an eminently reasonable choice for 2 (regardless of 1's choice). 1's reluctance to choose α_2 reflects 1's belief that 2 is likely to choose β_1 deliberately, not as a result of incomplete rationality. Similarly, 2 is likely to use β_1 because he expects that 1 will probably select α_1; no errors enter the picture.

If one really believed that players entertain the notion that their opponents may commit errors with some actual positive probability, then for large enough x , one would predict that the first player in G_5 would choose α_2 :

2

		β_1	β_2
1	α_1	(10,10)	(0,0)
	α_2	(0,10)	(x,0)

G_5

If there is a positive probability (at least in 1's mind) that β_2 will be played, 1 will choose α_2 if x is sufficiently large. But β_2 is strongly dominated by β_1 , and the latter is obviously going to be 2's choice; neither the solution concept of Selten nor that of Myerson predicts that 1 would ever choose α_2 .

I will argue that there is no need to base an analysis of imperfect behaviour on incomplete rationality; an alternative is available which conforms more closely to intuition. First, an extremely brief sketch of the solution concepts proposed by Selten and Myerson is given. This is not meant to be a substitute for reading the original definitions.

In a game $G = (S^1,\dots,S^N; U^1,\dots,U^N)$, a totally mixed strategy for player i is a mixed strategy giving positive weight to each pure strategy in S^i. For any small positive number ε, an ε-equilibrium of G is a profile of totally mixed strategies $(t^1,\dots,t^N)$ such that for each i, player i gives weight greater than ε to a given element s of S^i only if s is a best response to $(t^1,\dots,t^N)$. If $(z^1,\dots,z^N)$ is the limit of ε-equilibria as $\varepsilon \to 0$, $(z^1,\dots,z^N)$ is said to be a perfect equilibrium of G. (Each component of $(t^1,\dots,t^N)$ is an element of Euclidean space; convergence is with respect to the usual Euclidean metric). This is Myerson's formulation (Myerson (1978)), of what is often called "trembling hand perfect equilibrium," originally defined by Selten (1975) on the extensive form.

Roughly speaking, an ε-proper equilibrium is a "combination of totally mixed strategies in which every player is giving his better responses much more probability weight than his worse responses (by a factor $\frac{1}{\varepsilon}$), whether or not those 'better' responses are 'best' . . . We now define a proper equilibrium to be any limit of ε-proper equilibria." (Myerson (1978), page 78).

Requiring, as proper equilibrium does, that when contemplating an opponent's "trembles," a player should give much higher weight to relatively innocuous mistakes than to those which would cause the opponent serious damage, suggests that one is interested in "sensible trembles." In other words, the idea behind proper equilibrium seems to be that a player should be open-minded about various reasonable alternative strategies his opponents might use; the random component attributed to an opponent's action must not be arbitrary. While it is important to insist that doubts entertained by a player regarding his opponents' strategies should be concentrated upon reasonable possibilities, proper equilibrium attempts to enforce this without reference to any theory specifying what possibilities are realistic. This explains the failure of proper equilibrium to rule out unreasonable choices in many games. One well-known example is presented later in this section.

I believe that the ex ante equilibrium concept provides the kind of theory that is required to determine what "reasonable doubts" players can rationally entertain regarding the choices of their opponents. For each game, the ex ante theory distinguishes those strategies that players could employ without violating the implications of the common knowledge they possess, from those that are patently unreasonable. If the condition that players do not take unnecessary risks is to be imposed by requiring that their conjectures give positive weight to all "likely" alternatives, those strategies not in the ex ante solution sets should still be given zero weight.

Ideally one would like to proceed in the same way as in the previous section: call the vector of sets $(F^1,\dots,F^N)$ the perfect ex ante solution of $G = (S^1,\dots,S^N;\ U^1,\dots,U^N)$ if $F^1,\dots,F^N$ are the "largest" sets satisfying a suitable modification of the best response property. $X^1,\dots,X^N$ have the <u>cautious response property</u> if for each $\alpha \in X^i$, α is a best response to some $(\sigma^1,\dots,\sigma^N) \in \prod_{j=1}^{N} \bar{X}^j$, where for each j, α gives positive weight to every pure strategy in X^j. (Call α a <u>cautious response</u> to $X^1,\dots,X^N$). The requirement that the perfect ex ante solution sets have this property ensures that each player takes into account all of the alternative strategies admitted by the solution concept.

Unfortunately, this procedure is not well-defined, because in general there are no "largest" sets $F^1,\dots,F^N$ with the desired property. Consider the following example:

2

		β_1	β_2
1	α_1	(1,1)	(1,1)
	α_2	(1,1)	(0,0)

G_6

Let M^i be the mixed strategy set of player i , $i=1,2$. If A^1 is the singleton containing α_1 only, and A^2 contains β_1 only, then A^1, M^2 have the cautious response property, and so do M^1, A^2 . But the property is not preserved under unions: $A^1 \cup M^1$, $A^2 \cup M^2$ do not have the cautious response property. Underlying this fact is a paradox that presents a dilemma in the modelling of games with "cautious" players. If it is clear to the theorist that 1 and 2 will restrict themselves to α_1 and β_1 respectively, then this should be equally clear to 1, who then has no strict incentive not to choose α_2 . On the other hand, if there is the slightest doubt about what the players might choose, they have a strict incentive to stick religiously to their first strategies; but this removes all doubt about what might happen. In summary, there is a limit to the logical consistency of any solution concept for cautious strategic behaviour.

When faced with games where at least one player should "play it safe" and players are in essentially symmetric situations, as in G_6 , I am definitely willing to assume that careful players will all select their safe strategies. However, there are less symmetric situations requiring separate consideration. In G_7, rationality implies that 2's choice of β_1 is a certainty. 1 should be indifferent, then, between α_1 and α_2 . Notice that for 1, any strategy is a best response to β_1, the sole element of 2's ex ante solution set, and hence a cautious response to that set.

2

		β_1	β_2
1	α_1	(10,10)	(10,0)
	α_2	(10,10)	(0,0)

G_7

Since the most direct approach to formulating a solution concept for games with cautious players is unavailable, consider the following iterative method. Given the ex ante solution sets $E^1,\ldots,E^N$ of a game G , let $C^i(1) = \{\alpha \in E^i:$

α is a cautious response to $E^1,\ldots,E^N\}$, $\forall$ i . For $t > 1$, define $C^i(t)$ recursively for each i by $C^i(t) = \{\alpha \in D^i(C(t-1))$: α is a cautious response to $D^1(C(t-1),\ldots,D^N(C(t-1))\}$, where $C(t-1) = (C^1(t-1),\ldots,C^N(t-1))$, and D is the operation defined in the previous section (page 16). At each "round", strategies that are not best response strategies are eliminated first, and then those that are not cautious responses are removed.

Arguments analogous to those used for the previous section's iterative technique establish that the nonemptiness, closedness, and pure strategy property of the sets $E^1,\ldots,E^N$ are inherited at each stage t by the sets $C^1(t),\ldots,C^N(t)$. Hence the "finite-stopping" proof given in Section III is easily adapted to the present procedure. For $k = \sum_{r=1}^{N} k(r) - N$, $C^i(t) = C^i(k)$ for all $t \geq k$.

Let $Q^i = C^i(k)$ $\forall$ i . The vector $(Q^1,\ldots,Q^N)$ is the perfect ex ante solution of G , and each $(q^1,\ldots,q^N)$ with $q^i \in Q^i$ $\forall$ i is a perfect ex ante equilibrium of G . Since Q^i equals the nonempty set $C^i(k)$, there is an α in Q^i that is a cautious response to the sets $C^1(k),\ldots,C^N(k)$. Every pure strategy given positive weight by α is also in Q^i , so pure strategy perfect ex ante equilibria always exist.

The solution concept performs as desired on Myerson's example G_4 , and the reader can easily verify that the perfect ex ante solution concept is equally appropriate when applied to another example (not given here) constructed in Myerson (1978), for which proper equilibrium also does well. But consider G_8 , the normal form of a well-known extensive form game (to be called Γ_2) that is discussed in the next section.

2

		β_1	β_2
	α_1	(1,1)	(1,1)
1	α_2	(2,-1)	(-10,-2)
	α_3	(0,-2)	(0,-1)

G_8

Notice that (α_1,β_2) is one of the Nash equilibrium profiles of this game; in fact, one can show (α_1,β_2) is both a trembling hand perfect, and a proper equilibrium. Why would 2 ever select β_2? β_2 is preferable to β_1 only if 1 gives considerable weight to α_3. But 2 knows that α_3 is strongly dominated for 1 by α_1, and will never be played. Thus, there is no risk to playing β_1, and a superior return for playing β_1 rather than β_2 if α_2 is played. If 2 were a "cautious" player, it would be ridiculous for him to play β_2; knowing this, 1 plays α_2. In the notation developed above, $C^1(1)$ contains all strategies giving zero weight to α_3, while $C^2(1) = \{\beta_1\}$. Then $D^1(C(1)) = \{\alpha_2\}$, and $D^2(C(1)) = \{\beta_1\}$. No further reduction can take place; the unique perfect ex ante equilibrium isolates the only reasonable Nash equilibrium of G_8, namely (α_2,β_1).

On the other hand, the perfect ex ante solution concept was also designed with games such as G_7 in mind, where it singles out β_1 for 2, but respects 1's legitimate indifference between α_1 and α_2. Critics may object that a "nongeneric" phenomenon such as the weak dominance of α_2 by α_1 in G_7 should not be a consideration in the formulation of a solution concept. This position is contested in the next section.

V. A NOTE ON NONGENERIC PROPERTIES

Results involving the "generic equivalence" of solution concepts or the "nongeneric nature" of some event are encountered with increasing frequency in the study of game theory. While these results are often enlightening, there is some danger of their being misinterpreted. Specifically, occurrences that arise only on a set of measure zero (according to a particular measure on some relevant space) may be ignored, and the poor performance of a solution concept on that set excused, because such occurrences are "infinitely unlikely." The purpose of this section is to question the appropriateness of the measure most commonly applied, and to explain thereby the concern expressed in the last section about a solution concept's performance characteristics even in "nongeneric" examples.

Consider a vector $S = (S^1,\ldots,S^N)$ where each S^i is comprised of a finite number $k(i)$ of pure strategies. Let W be the set of all N-vectors of utility functions having domain S. There are $s = k(1)\times\ldots\times k(N)$ profiles of the form $(s^1,\ldots,s^N)$, $s^i \in S^i \; \forall \; i$. Label these profiles $q_1,\ldots,q_s$, the order chosen being immaterial. Then for a given $U = (U^1,\ldots,U^N) \in W$, $V(U) = (U^1(q_1),\ldots,U^1(q_s);\ldots;U^N(q_1),\ldots,U^N(q_s))$ is an element of Ns-dimensional Euclidean space. A property P on the elements of W is said to be <u>nongeneric</u> if the closure of the set $\{V(U)\colon \; U \text{ satisfies } P\}$ has Lebesgue measure zero. Intuitively, this suggests that if a vector of payoff functions, represented by an Ns-element vector, were chosen "at random" from a subset of Euclidean space having positive Lebesgue measure, it is "infinitely unlikely" that the chosen vector would satisfy the nongeneric property P. In particular, the property that "ties" occur in the payoffs (i.e. $\exists$ pure strategy profiles p and q, $p \neq q$, such that for some i, $U^i(p) = U^i(q)$) is nongeneric.

Whatever theory of strategic behaviour one develops should be suitable for analyzing a wide variety of real-world situations that are likely to be of interest. Since the payoff functions in actual examples are not drawn at random, and often are partly determined by man-made institutions, can we confidently assert that phenomena such as ties will never arise? A few examples suffice to provide the answer.

The theory of voting is one field in which the application of game theory has proved fruitful. Consider the simplest case in which each of the N players must vote for one of two candidates. While there are 2^N strategy profiles, there are only two possible outcomes; ties must occur for $N \geq 2$, for any nonstochastic rule that selects a victor as a function of the voting. With more candidates, say M of them, ties are even more ubiquitous: M^N is usually vastly larger than M. The problem is aggravated further if more complicated strategies are allowed, such as rank-ordering of alternatives.

Another major application of game theory is to the study of auctions. Consider a sealed bid auction at which N persons privately and independently submit bids to the auctioneer. In order to retain the finite character of the game, suppose that the bids must be integer dollar values between \$0 and \$D, where D is some number huge enough to be unrestrictive. The object to be auctioned, let us say a painting, goes to the highest bidder for the price he bids. If the high bid is made by several people, some rule known to all players dictates which player gets the painting. The standard assumption is that each player cares only about whether or not he gets the painting, and at what price. If so, a player can realize at most D+2 different payoffs. But there are $(D+1)^N$ strategy profiles; ties will occur in each person's utility function whenever $N \geq 2$ and $D \geq 1$. If players actually care about who gets the painting, other than themselves, there are many more utility levels possible. But given that 1 has the high bid of \$10, say, person 2 does not care whether 3 bids \$4 or \$5; this means that a tie occurs in 2's utility function.

Whether or not actual games are ever precisely zero-sum in practice, it is often convenient to model certain situations (such as dividing a cake between two persons) as zero-sum games. But the zero-sum property is nongeneric, so two solution concepts could differ drastically on every zero-sum game, and still be regarded as generically equivalent to one another.

These examples leave little doubt that the measure typically used in declaring an event or example nongeneric is unsuitable if game theory is to be a general tool for studying what will happen in varied strategic settings. Therefore efforts to ensure that irrational behaviour is ruled out by a theory should not be dismissed merely because the problem is nongeneric according to the traditional measure.

VI. EX ANTE EQUILIBRIUM IN THE EXTENSIVE FORM

This section generalizes the analysis of Section III to games having some sequential nature. In this context it is possible to study the best-known type of imperfect behaviour, namely unreasonable behaviour at unreached information sets. The problem is attacked using the idea of consistent conjectures, without the additional assumptions needed to ensure cautious behaviour. Those assumptions are invoked in Section VII, because what I have called imperfections of the second type may still arise in the extensive form.

A complete formal description of an extensive form game would be too lengthy to be appropriate here. Some knowledge of extensive form games and their normal forms is taken for granted, but many initial definitions are unavoidable. My presentation follows the more detailed treatment to be found in Selten (1975), with some changes in notation.

A finite extensive form game has the structure $\Gamma = (K, P, I, S, p, h)$ and is interpreted as follows. K is a topological tree defined by a set of _vertices_ or _nodes_ X and a set of _edges_ or _alternatives_ A connecting certain pairs of vertices. Z denotes the set of non-terminal vertices. $P = (P^0, P^1, \ldots, P^N)$ partitions the non-terminal vertices of K into player sets. Without loss of generality suppose that only the origin of K is in the player set P^0. The 0th player is "nature," who makes a random move at the beginning of the game. $P^1, \ldots, P^N$ are associated with _personal players_. Let I partition the non-terminal nodes into information sets $I^{ij} \subseteq P^i$, where I^{ij} is the j^{th} information set of the i^{th} player. $A^{ij} \subseteq A$ is the set of all alternatives at vertices $x \in I^{ij}$. A _choice_ at I^{ij} is a subset of A^{ij} that contains exactly one alternative for every $x \in I^{ij}$; each $a \in A^{ij}$ is part of exactly one choice at I^{ij}. S^{ij} is the set of all choices at I^{ij}. A positive probability $p(a)$ is assigned to each $a \in S^{01} = A^{01}$, the random player's set of alternatives. The payoff vector h assigns payoffs $h^1(x), \ldots, h^N(x)$ to each terminal node x.

A vertex x is said to come before y if $x \neq y$, and the path (set of edges) from the origin to y contains the path from the origin to x. An information set I^{ij} is a _predecessor_ of $I^{k\ell}$ if there are vertices $x \in I^{ij}$ and

$y \in I^{k\ell}$ such that x comes before y; $I^{k\ell}$ is called a successor of I^{ij}. A vertex x comes after a choice c if one of the edges in c is on the path to x. I restrict myself to games of perfect recall (Kuhn (1953)): for each $i=1,2,\ldots,N$, and any j, if $y \in I^{ij}$ comes after a choice c at I^{ik}, then every $x \in I^{ij}$ comes after c. This condition would be violated only if some player were "forgetting" information as the game proceeded.

A pure strategy f for player i is a function associating with each information set I^{ij} of i one of the choices in S^{ij}; denote this choice by $f(i,j)$. If f and g are pure strategies for i, g is an ij-replacement for f if for all $\ell \neq j$ such that $I^{i\ell}$ is not a successor to I^{ij},

$$g(i,\ell) = f(i,\ell).$$

This says that f and g agree everywhere except on I^{ij} and its successors.

A mixed strategy for i is a probability distribution over player i's pure strategies. If $n(i)$ is the number of information sets of i, and the ij^{th} information set has $k(i,j)$ choices, then i has $\prod_{j=1}^{n(i)} k(i,j)$ pure strategies. Hence, a mixed strategy is represented as a point m in Euclidean space of dimension $\prod_{j=1}^{n(i)} k(i,j)$. The components of m lie between zero and one and sum to unity. As usual the mixed strategy assigning probability one to some pure strategy α is considered identical to that pure strategy; these two representations of α are used interchangeably. M^i denotes the set of mixed strategies of player i.

A profile $(m^1,\ldots,m^N)$ of mixed strategies together with the probability assignment function p (for the random player) induces a probability distribution on the set of terminal nodes of the game. Number these nodes $x_1,\ldots,x_T$ and let $\delta_r(m^1,\ldots,m^N)$ be the probability of reaching the terminal node x_r when players select $(m^1,\ldots,m^N)$. Then utility functions are defined over the set of profiles of mixed strategies by the expected payoff calculation:

$$U^i(m^1,\ldots,m^N) = \sum_{r=1}^{T} \delta_r(m^1,\ldots,m^N)h^i(x_r),\ i=1,\ldots,N .$$

The mixed strategy profile $(\sigma^1,\ldots,\sigma^N)$ is a Nash equilibrium of Γ if for each i and for all $\bar{\sigma}^i \in M^i$,

$$U^i(\sigma^1,\ldots,\sigma^N) \geq U^i(\sigma^1,\ldots,\sigma^{i-1}, \bar{\sigma}^i,\sigma^{i+1},\ldots,\sigma^N).$$

Consider a particular information set I^{ij} and a profile $m = (m^1,\ldots,m^N)$ of (mixed) strategies. If for each terminal node x_r reached with positive probability when m is played, and each $y \in I^{ij}$, y does not lie on the path from the origin to x_r, then I^{ij} is not reached by m. If the condition is violated, I^{ij} is reached by m .

Consider the game Γ_1 having perfect information (all information sets are singletons) and no randomness. (When representing games where the random move is restricted to one choice, I simply omit the random player's information set.)

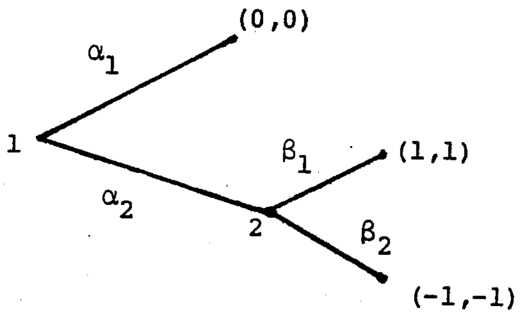

Γ_1

Although the outcome yielding (0,0) is absurd, it is among the Nash equilibrium outcomes of Γ_1 . If 1 specifies the choice α_1 (with probability 1) and 2 chooses β_2 , neither has an incentive to deviate. But everyone must agree that if 1 were to play α_2 2 would, upon being reached, respond by playing β_1 . Knowing this, 1 should play α_2 . The imperfect behaviour arises because in the dubious equilibrium, 2's information set is not reached with positive probability. Consequently 2 can specify any choice with impunity.

Subgame perfect equilibrium (Selten (1965), (1975)) deals nicely with examples of this variety. A Nash equilibrium is subgame perfect if the strategies it induces on any proper subgame of Γ (see Selten (1975)) constitute a Nash equilibrium of that subgame. In Γ_1 , 2's choice of α_2 is not Nash on the subgame starting at 2's information set.

Unfortunately there are often too few proper subgames to allow subgame perfection to enforce intuitively reasonable behaviour in a game. This prompted Selten (1975) to introduce a further notion, perfect equilibrium, or trembling hand perfect equilibrium. The set of perfect equilibria is a subset of the set of subgame perfect equilibria. Loosely speaking, a perfect equilibrium is a limit of a sequence of ε-equilibria, each of which assigns at least some small probability to each choice at every information set. As a result, all information sets are reached in an ε-equilibrium, and strategies such as α_2 in Γ_1 are not played.

As was noted in Section IV, the indiscriminate nature of the "trembles" allowed causes problems for the perfect equilibrium concept. The attempt by Myerson (1978) to correct this by limiting the class of admissible trembles was only partially successful; proper equilibrium remains too deeply rooted in the stochastic "small mistakes" framework to escape all the difficulties created by that approach. A major alternative has been suggested by Kreps and Wilson (forthcoming). Their solution concept, sequential equilibrium, is based upon an examination of rational beliefs rather than the possibilities for error. As it is not practical to present the complicated definitions here, the reader should consult the original paper for a full account. He will find there an excellent discussion of many of the issues involved in the perfection debate, as well as a rich supply of examples.

While all of the solution concepts mentioned above have features that are extremely attractive, examples abound in which none of the equilibrium notions is satisfactory (one well-known example is presented later in this section). Equally important is the fact that they all admit Nash profiles only; this paper attempts to escape that restriction. Let us try to apply the idea of consistent conjectures to examples such as Γ_1 .

The possibility of collapsing series of choices into timeless contingent strategies must not obscure the fact that the phenomenon actually being modelled is some sequential game, in which conjectures may actually be contradicted in the course of play. In Γ_1 , it is ludicrous to maintain that if 2 is called upon to move, having been reached, he might choose β_2, thinking that α_1 was played by 1. By the time he must commit himself to a course of action, 2 knows that it is a fact that 1 played α_2. The observation that a conjecture must not be maintained in the face of evidence that refutes it is a central element of the sequential equilibrium concept; it is combined here with a further principle and the iterative techniques of previous sections to construct a new solution concept for extensive form games.

Since a player's beliefs about others' strategies may be refuted as a play of the game progresses, he might need to formulate new conjectures as the old ones are disproven. Consequently I associate a <u>conjecture</u> c^{ij} with each information set in Γ; c^{ij} represents what an "agent" ij for player i believes about what everyone's mixed strategies are, once I^{ij} is reached. Such a conjecture involves a probability measure on each set M^r. But since M^r is a set of probability distributions over the pure strategies of r, a measure on M^r induces a probability distribution over pure strategies, and can be regarded for all strategic purposes as a point in M^r. More generally, for subsets $A^r \subseteq M^r$, $r=1,\ldots,N$, a conjecture c^{ij} over the sets $A^1,\ldots,A^N$ is an element $(c^{ij}(1),\ldots,c^{ij}(N))$ of $\prod_{r=1}^{N} \bar{A}^r$, i.e. each component $c^{ij}(r)$ is a "weighted average" of points in A^r, where the weights depend on the measure that "agent ij" has in mind. The A^r were not assumed convex; the weighted averages $c^{ij}(r)$ need only be in the convex hulls $\bar{A}(r)$.

I have noted that an agent ij, upon being reached, should not entertain a conjecture that does not reach I^{ij}. A further restriction, not invoked in other solution concepts, is appropriate: if the information set can be reached without violating the rationality of any player, then the agent's conjecture must not attribute an irrational strategy to any player. In other words, he should seek a <u>reasonable</u> explanation for what he has observed. This principle is applied within an iterative procedure similar to that of Section III, suitably elaborated to exploit the additional information in the extensive form.

For later reference, the iterative procedure is defined for arbitrary closed, nonempty sets $H^1,\ldots,H^N$ having the pure strategy property; our immediate interest is in the technique applied to $M^1,\ldots,M^N$. Let $H = (H^1,\ldots,H^N)$ and let $H^i(0) = H^i$, $i=1,\ldots,N$. For any $t \geq 1$, define the sets $H^1(t),\ldots,H^N(t)$ recursively as follows. For each pure strategy $\beta \in H^i(t-1)$, let $J^i(\beta,H,t)$ contain all those j such that I^{ij} can be reached by some profile of the form $(m^1,\ldots,m^{i-1},\beta,m^{i+1},\ldots,m^N)$, where $m^r \in H^r(t-1)$, $r=1,\ldots,N$. (The eventual interpretation will be that at stage t of the logical deduction process, i knows that if he plays β, no information set I^{ij} will be reached unless $j \in J^i(\beta,H,t)$.) A strategy $\alpha \in H^i(t-1)$ giving positive weight to pure strategies $\alpha_1,\ldots,\alpha_h$ is an element of $H^i(t)$ if there exist conjectures c_z^{ij}, $z=1,\ldots,h$ such that for all z, and all $j \in J^i(\alpha_z,H,t)$,

(i) $c_z^{ij}(i) = \alpha_z$

(ii) $c_z^{ij}(\ell) = c_1^{ij}(\ell)$, $\ell \neq i$

(iii) for $r, s \in J^i(\alpha_z,H,t)$, if I^{ir} is a predecessor of I^{is} and c_z^{ir} reaches I^{is}, then $c_z^{is} = c_z^{ir}$.

(iv) c_z^{ij} reaches I^{ij}.

(v) $c_z^{ij} \in \prod_{r=1}^{N} \bar{H}^r(t-1)$

and (vi) α_z is a best response to c_z^{ij} among all ij-replacements for α_z in $H^i(t-1)$.

At each stage, additional restrictions are placed on conjectures and actions only at information sets that can be reached by profiles of strategies not previously eliminated. In a particular play of the game, player i uses some pure strategy α_z which is a realization of the mixed strategy α. Condition (i) says that i's "conjecture" about his own strategy is correct. The next requirement stipulates that conjectures about others' strategies do not depend upon which of the $\alpha_1,\ldots,\alpha_h$ player i ends up using. According to (iii), a conjecture should not be discarded unless it is contradicted (by arrival at an information set unreachable by the conjecture in question). Condition (iv) ensures that a conjecture at I^{ij} explains how that information set could have been reached. The principle that the explanation should be "reasonable" is embodied in (v), which restricts conjectures to strategies that have not been eliminated at a previous stage. Finally, the strategy chosen by i should at all times be an optimal response to the conjectures he holds. The most convenient way to express this condition is to consider ij-replacements for α_z ; these represent the options still open to i at I^{ij} . Among these, α_z must constitute an optimal contingent plan, given that beliefs about others' mixed strategies are described by c_z^{ij} .

The sets $H^i(t)$, $i=1,\ldots,N$ inherit the pure strategy property, nonemptiness, and closedness from the original sets H^i . This is easy to see in the case of the pure strategy property, because if the pure strategies of which a mixed strategy α is comprised can collectively satisfy (i) to (vi), each of the pure strategies satisfies the conditions individually. To show nonemptiness, assume $H^1(t-1),\ldots,H^N(t-1)$ are nonempty and closed, and choose any conjecture $\tilde{c} = (\tilde{c}(1),\ldots,\tilde{c}(N))$ such that $\tilde{c}(r) \in \bar{H}^r(t-1)$ gives positive weight to every pure strategy in $H^r(t-1)$. Since U^i is continuous and $H^i(t-1)$ is nonempty and compact, there exists an α that is a best response in $H^i(t-1)$ to $\tilde{c}$. α may be chosen to be a pure strategy, because $H^i(t-1)$ has the pure strategy property. For every $j \in J^i(\alpha,H,t)$, define

$$c_1^{ij} = (\tilde{c}(1),\ldots,\tilde{c}(i-1),\alpha,\tilde{c}(i+1),\ldots,\tilde{c}(N)).$$

α and c_1^{ij} satisfy (i) to (vi). (i) holds by definition. (ii) is trivially satisfied because there is only one pure strategy involved. (iii) is equally clear since c_1^{ij} is not a function of j as defined. In all components except i , c_1^{ij} gives positive weight to all pure strategies not eliminated in previous rounds; hence c_1^{ij} reaches I^{ij} for all $j \in J^i(\alpha,H,t)$, and (iv) is satisfied. (v) holds by the definition of $\tilde{c}$. Since α is a best response to c_1^{ij} in $H^i(t-1)$, α is certainly a best response to c_1^{ij} in the set of all ij-replacements for α in $H^i(t-1)$, therefore $\alpha \in H^i(t)$. To establish that $H^i(t)$ is closed, consider a sequence $\beta_1,\beta_2,\ldots$ in $H^i(t)$ converging to a strategy β . $H^i(t-1)$ is closed by hypothesis, so $\beta \in H^i(t-1)$. For some integer V , it must be the case that for all $W \geq V$, β_W gives positive weight to (at least) all the pure strategies given positive weight by β. But there exists a set of conjectures c_z^{ij} (where z indexes the pure strategies comprising β_V) such that β_V and the c_z^{ij} satisfy (i) to (vi). Then β and the c_z^{ij} (omitting any conjectures corresponding to pure strategies not given positive weight by β) satisfy (i) to (vi). Thus $\beta \in H^i(t)$, and the set is closed.

As in the normal form, the fact that the sets $H^i(t)$ have the pure strategy property means that the sets differ at successive stages only if pure strategies have been eliminated in the preceding stage. As pure strategies are in finite supply, the process stops after at most $k = (\sum_{r=1}^{N} \prod_{j=1}^{n(r)} k(r,j)) - N$ steps: for all $t \geq k$ and $i=1,\ldots,N$,

$$H^i(t) = H^i(k).$$

Let $D^i(H) = H^i(k) \; \forall \; i$, where $H = (H^1,\ldots,H^N)$.

In particular, the objects of interest in this section are the sets $D^1(M),\ldots,D^N(M)$, where $M = (M^1,\ldots,M^N)$ and the M^i are the full sets of mixed strategies for players $i=1,\ldots,N$. Let $D^i(M) = E^i \; \forall \; i$. The vector of sets $(E^1,\ldots,E^N)$ is defined to be the ex ante solution of Γ, and any $(e^1,\ldots,e^N)$ with $e^i \in E^i \; \forall \; i$ is an ex ante equilibrium. The nonemptiness of the sets $H^i(t)$ for all i and t guarantees the existence of an ex ante equilibrium for every finite extensive form game Γ. The pure strategy property ensures that there is an equilibrium in pure strategies.

To get some feeling for how this solution concept operates, consider two examples, starting with the familiar Γ_1. In that game, 1 is unable to eliminate any strategy in the first round. Since strategies of 1 that reach 2's information set must give positive weight to α_2, 2 must remove all strategies that are not best responses to some such strategy. This eliminates all strategies of 2 except β_1, so in the next round, 1 retains the only strategy that is a best response to β_1, namely α_2.

A more challenging test for the theory is an example that Kreps and Wilson (forthcoming) attribute to E. Kohlberg. (The example is "generic": small perturbations in the payoffs will not alter any of the statements made below.)

In the game Γ_2 , player 2 has only one information set, which is indicated in the game tree by enclosing the two nodes in that information set by an oblong figure.

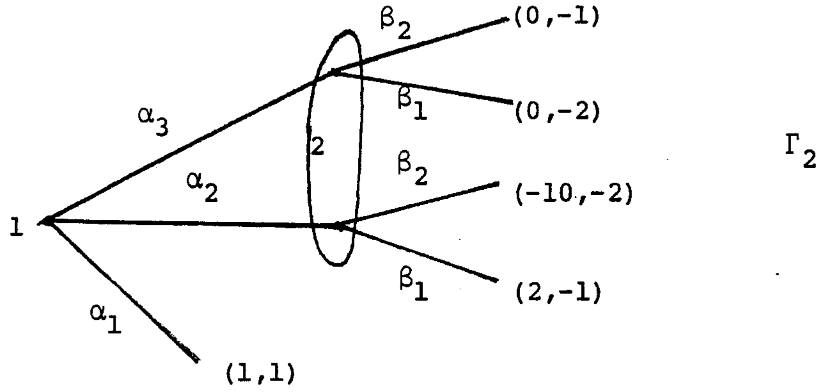

Notice that α_1 strongly dominates α_3; the latter will never be played with positive probability by a rational player. If reached, 2 should conclude that α_2 was played and respond optimally by playing β_1 . Knowing that this would be 2's response, 1 should play α_2 . Despite this simple argument, another Nash equilibrium (which can actually be shown to be a trembling hand perfect, proper, and sequential equilibrium) has 1 playing α_1 with certainty and 2 playing α_3 . This is not an ex ante equilibrium. In the first "round," all strategies giving α_3 positive weight are removed. In the second round, since these strategies are absent from $M^1(1)$, 2 eliminates every strategy except β_1 , because elements of $M^1(1)$ reaching 2's information set are those giving some positive weight to α_2 . In the third round, 1 has a unique best response α_2 to the single element β_1 in $M^2(2)$. The only ex ante equilibrium of Γ_2 is what Kreps and Wilson agree is the only reasonable profile. Their general remarks on what beliefs should be admissible are interesting:

> "Some sequential equilibria are supported by beliefs that the analyst can reject because they are supported by beliefs that are implausible. We will not propose any formal criteria for 'plausible beliefs' here. In certain cases, such as Myerson's concept of properness, some formalization is possible. In other cases, it is not clear that any formal criteria can be devised--it may be that arguments must be tailored to the particular game." (Kreps and Wilson, forthcoming, Section 8)

The ex ante solution formalizes the notion that beliefs may be implausible at an information set because

(i) the set could not have been reached had those beliefs been true,

or (ii) they are inconsistent with the results of logical deductions based on what players know about one another and the rules of the game.

If the ex ante solution fails to narrow down the possible outcomes significantly in a given game, one might then consider applying criteria of a more ad hoc description, and perhaps make predictions on a game-by-game basis as Kreps and Wilson suggest.

VII. PERFECT EX ANTE EQUILIBRIUM IN THE EXTENSIVE FORM

It is straightforward to verify that in a perfectly simultaneous game, the ex ante solution coincides with the normal form definition given in Section III, applied to the normal form of the game in question. But in such games, ex ante equilibrium behaviour is not always "cautious": the solution concept does not prevent imperfection of the second type. A simple demonstration that this applies equally to the extensive form is given by Γ_3, whose normal form is G_4, Myerson's example.

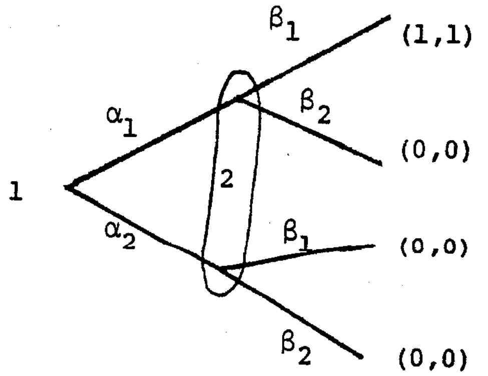

Γ_3

If both players make prudent choices, (α_1,β_1) will result. But (α_2,β_2) is also an ex ante equilibrium. Such behaviour can be avoided by the same technique as that employed in Section IV. A natural generalization of the normal form analysis is accomplished here as briefly as possible.

Consider a game Γ with notation as defined in the previous section. For arbitrary nonempty sets $A^i \subseteq M^i$, $i = 1,\ldots,N$, a strategy $\alpha \in A^i$ is a <u>cautious response</u> to the sets $A^1,\ldots,A^N$ if α is a best response among the strategies in A^i, to some $(\sigma^1,\ldots,\sigma^N) \in \prod_{r=1}^{N} \bar{A}(r)$, and $\forall\ r \neq i$, σ^r gives positive weight to every pure strategy in A^r .

Given the (extensive form) ex ante solution sets $E^1,\ldots,E^N$, for each i let $C^i(1) = \{\alpha \in E^i\colon\ \alpha$ is a cautious response to $E^1,\ldots,E^N\}$. Let $C(1) = (C^1(1),\ldots,C^N(1))$. For $t > 1$, define $C^i(t)$ recursively by $C^i(t) = \{\alpha \in D^i(\mathbb{C}(t-1))\colon\ \alpha$ is a cautious response to $D(C(t-1))\}$ where $C(t-1) = (C^1(t-1),\ldots,C^N(t-1))$, and $D = (D^1,\ldots,D^N)$ is the operation defined in the previous section (page 34). The sets $C^i(t)$ are closed, nonempty, and satisfy the pure strategy property. At each round, strategies that are not best responses are discarded first, and then those that are not cautious responses are eliminated. As usual, the iterations produce no change unless pure strategies are eliminated, and so for $k = \sum_{i=1}^{N} \prod_{j=1}^{n(i)} k(i,j) - N$ and all $t \geq k$,

$$C^i(t) = C^i(k)\ \forall\ i\ .$$

Let $Q^i = C^i(k)\ \forall\ i$. The vector of sets $(Q^1,\ldots,Q^N)$ is the <u>perfect ex ante solution</u> of Γ , and any $(q^1,\ldots,q^N) \in \prod_{i=1}^{N} Q^i$ is a <u>perfect ex ante equilibrium</u>. The nonemptiness of the $C^i(t)$ ensure the existence of at least one perfect ex ante equilibrium in each game.

This solution concept has the attractive feature that in the play of a game, no one's conjectures are ever contradicted. Since each person's conjecture gives positive weight to every perfect ex ante strategy of each other player, nothing that is believed by any player to have zero probability ever occurs, so long as others choose cautiously.

It might appear at first glance that in a game such as Γ_4 in which 1 should be indifferent between α_1 and α_2 (according to subgame perfection or backward induction), perfect ex ante equilibrium forces 1 to choose α_1 , by eliminating α_2 in the first round, before β_2 has been removed.

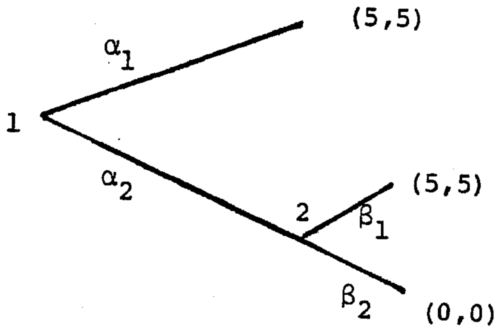

Γ_4

In fact this does not happen. Recall that before the cautious response criterion comes into play, the ordinary ex ante solution is calculated. For 2, this eliminates all strategies except β_1 ; in "cautious response" to <u>this</u>, 1 plays either α_1 or α_2. It is in order to preserve this sort of performance characteristic in more complicated examples, that the iterative procedure generating perfect ex ante equilibrium removes strategies in the particular order specified.

VIII. CONCLUSION

In response to the opening question: "What constitutes rational behaviour in a non-cooperative strategic situation?", an extremely conservative theory of strategic behaviour, ex ante equilibrium, has been developed. Without pretending to predict behaviour uniquely in all games, the solution concept rules out strategic choices on the basis of rather fundamental principles such as maximization of expected utility, and the common knowledge assumption. Ex ante equilibrium is well-suited to dealing with implausible behaviour at "unreached" information sets, but an additional assumption that players are in some sense cautious is needed to deal with a second kind of imperfection. If cautious behaviour seems plausible, then the perfect ex ante solution is an interesting concept.

While the discussion has been restricted to games of complete information, Harsanyi (1967) has shown that games having various sorts of incomplete information, such as incomplete knowledge of others' utility functions, can be handled as standard games of complete information by an ingenious use of the random move at the beginning of the game. The ex ante theory should apply to these cases as well. I feel that one of the most promising applications of this concept lies in the resolution of various puzzles that could not be solved using Nash equilibrium. For example, a notorious problem in labour economics is to rationalize the occurrence of strikes. This is extremely hard to do using full rationality and any perfect Nash solution concept. The phenomenon can be explained very simply, however, if one allows for mismatching of expectations about players' future behaviour. Similarly in industrial organization, price wars between two rivals, one of whom eventually leaves the market, are best explained by the position that both expected to "win"; these expectations need not be irrational, although they could not be held simultaneously in a Nash equilibrium. I hope to elaborate on these ideas in subsequent papers.

Notes

1. This interpretation of the normal form is not uncommon. Selten (1973) notes that ". . . a simultaneity game is adequately described by its normal form . . . every normal form is isomorphic to the normal form of some simultaneity game." (page 160).

2. Bernheim's paper is a rarity in the current literature, in that it contains a serious critical analysis of the Nash theory. My conclusions are very similar to his. In fact, in the special case of perfectly simultaneous finite games of complete information, ex ante equilibrium and one form of Bernheim's solution concept reduce to the same definition, with the exception that Bernheim requires players to employ pure strategies. Our work was done entirely independently; the differences in mine largely reflect the fact that among my main interests have been the problem of imperfect equilibria, and the importance of information contained in the extensive form.

3. The concept of "common knowledge" was formalized by Aumann (1976). It is not helpful to give his definition here, as it implicitly depends upon various elements of the definition (such as statistical priors over a certain set) being common knowledge in an intuitive sense that is not defined.

4. $\alpha \in D^i(M) \Rightarrow \alpha$ is a best response $\underline{\text{in } D^i(M)}$ to some $\gamma \in \prod_{r=1}^{N} \overline{D}^r(M)$. Let α^* be a best response in M^i to γ. Since $\gamma \in \prod_{r=1}^{N} \overline{D}^r(M)$, $\gamma \in \prod_{r=1}^{N} \overline{M}^r(t)\ \forall t$, so α^* cannot be removed at any stage. Thus $\alpha^* \in D^i(M)$, so α is in fact a best response (in M^i) to γ.

Appendix

This appendix furnishes a proof of a claim made in the body of the paper. The result has probably been established in various contexts in the literature, but is included here for completeness. Dilip Abreu suggested the method of proof followed here. Note that the argument is <u>not</u> restricted to zero-sum games, but cannot be generalized to N-person games, where the proposition is false.

Proposition. Let $G = (S^1, S^2; U^1, U^2)$ be a finite non-cooperative game, with associated mixed strategy sets M^1 and M^2. $\alpha \in M^1$ is strongly dominated if and only if $\nexists m \in M^2$ such that α is a best response to m.

<u>Proof</u>: If some $\beta \in M^1$ strongly dominates α, then $\forall \gamma \in M^2$, $U^1(\beta,\gamma) > U^1(\alpha,\gamma)$, so α is never a best response.

To establish the converse, suppose α is not a best response to any element of M^2. Then there exists a function $b: M^2 \rightarrow M^1$ with $U^1(b(m),m) > U^1(\alpha,m) \; \forall \, m$. Consider the zero-sum game $\bar{G} = (S^1, S^2; \bar{U}^1, \bar{U}^2)$ where $\bar{U}^1(x,y) = U^1(x,y) - U^1(\alpha,y)$ and $\bar{U}^2(x,y) = -\bar{U}^1(x,y)$.

Let (x^*,y^*) be a Nash equilibrum of $\bar{G}$, and hence a pair of maximin strategies.

For any $m \in M^2$,

$$\begin{aligned} \bar{U}^1(x^*,m) &\geq \bar{U}^1(x^*,y^*) \\ &\geq \bar{U}^1(b(y^*),y^*) \\ &> \bar{U}^1(\alpha,y^*) \\ &= 0 . \end{aligned}$$

But $\bar{U}^1(x^*,m) > 0 \; \forall \, m$

$\Rightarrow U^1(x^*,m) > U^1(\alpha,m) \; \forall \, m$

$\therefore$ α is strongly dominated by x^*. Q.E.D.

Bibliography

Aumann, R. (1976), "Agreeing to Disagree," The Annals of Statistics, 4.

Bernheim, D. (1981), "On the Choice of an Equilibrium Concept for Economic Games," unpublished manuscript, M.I.T.

Ellsberg, D. (1956), "Theory of the Reluctant Duelist," American Economic Review, 46.

Farquharson, R. (1957; 1969), Theory of Voting, New Haven, Yale University Press.

Gale, D. (1953), "A Theory of N-Person Games with Imperfect Information," Proceedings of the National Academy of Sciences, 39.

Harsanyi, J. C. (1967-1968), "Games with Incomplete Information Played by 'Bayesian' Players, I-III," Management Science, 14, nos. 3, 5, 7.

Harsanyi, J. C. (1976), "A Solution Concept for n-Person Noncooperative Games," International Journal of Game Theory, 5.

Kreps, D. and R. Wilson (1982), "Sequential Equilibria," Econometrica, forthcoming.

Kuhn, H. W. (1953), "Extensive Games and the Problem of Information," H. Kuhn and A. Tucker (eds.), Contributions to the Theory of Games, Vol. 2, Princeton University Press.

Luce, R. D. and H. Raiffa (1957), Games and Decisions, New York, John Wiley and Sons.

Moulin, H. (1979), "Dominance Solvable Voting Schemes," Econometrica, 47, no. 6.

Myerson, R. B. (1978), "Refinements of the Nash Equilibrium Concept," International Journal of Game Theory, 7, no. 2.

Nash, J. F. (1951), "Non-cooperative Games," Annals of Mathematics, 54.

Selten, R. (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," Zeitschrift für die Gesamte Staatswissenschaft, 121.

Selten, R. (1973), "A Simple Model of Imperfect Competition, where 4 are Few and 6 are Many," International Journal of Game Theory, 2, no. 3.

Selten, R. (1975), "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory, 4.