Universität Bielefeld/IMW

Working Papers
Institute of Mathematical Economics

Arbeiten aus dem
Institut für Mathematische Wirtschaftsforschung

Nr.15

John C. Harsanyi

The Tracing Procedure: A Bayesian Approach to Defining a Solution for n-Person Noncooperative Games. PART I

May 1974

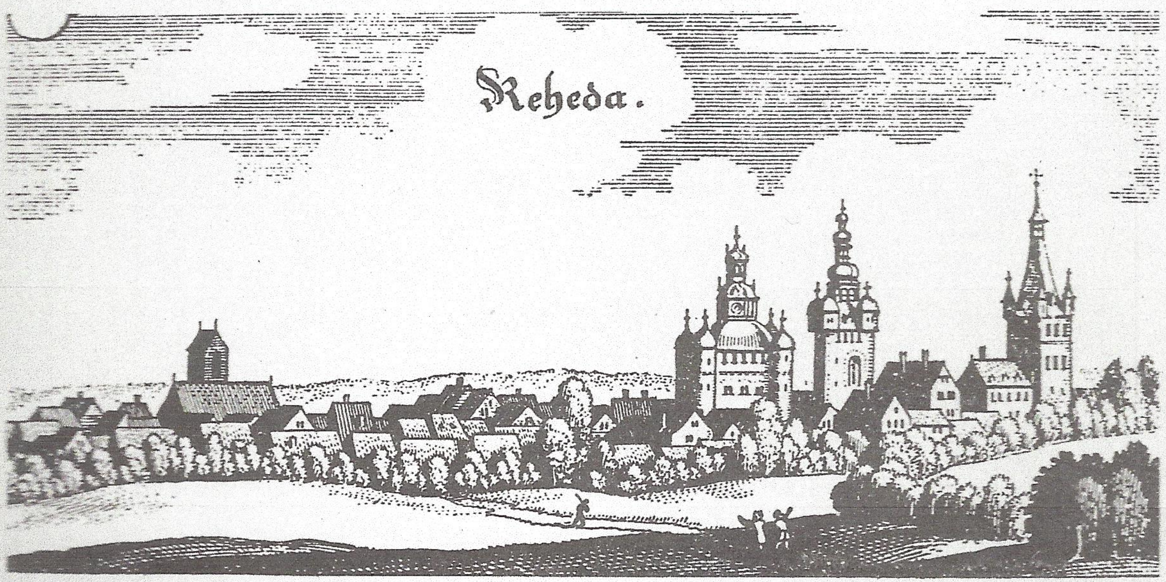

Anno 1600 nach Merian

Institut für Mathematische Wirtschaftsforschung
an der
Universität Bielefeld
Adresse/Address:
Schloß Rheda
484 Rheda
Bundesrepublik Deutschland
Federal Republic of Germany

The Tracing Procedure: A Bayesian Approach to Defining a Solution for n-Person Noncooperative Games.

PART I

By John C. Harsanyi
University of California, Berkeley,
and
University of Bielefeld

ABSTRACT

(For Parts I and II)

The paper proposes a Bayesian approach to selecting a particular equilibrium point s^* of any given finite n-person noncooperative game Γ as solution for Γ. It is assumed that each player i starts his analysis of the game situation by assigning a subjective prior probability distribution p_j to the set of all pure strategies available to each other player j. (The prior distributions p_j used by all other players i in assessing the likely strategy choice of any given player j will be identical, because all these players i will compute this prior distribution p_j from the basic parameters of game Γ in the same way.) Then, the players are assumed to modify their subjective probability distributions p_j over each other's pure strategies systematically in a continuous manner until all of these probability distributions will converge, in an appropriate sense, to a specific equilibrium point s^* of Γ, which, then, will be accepted as solution.

A mathematical procedure, to be called the <u>tracing procedure,</u> is proposed to provide a mathematical representation for this intellectual process of convergent expectations. Two variants of this procedure are described. One, to be called the <u>linear</u> tracing procedure, is shown to define a unique solution in "almost all" cases but not quite in all cases. The other variant, to be called the <u>logarithmic</u> tracing procedure, always defines a unique solution in all possible cases. Moreover, in all cases where the linear procedure yields a unique solution at all, both procedures always yield the same solution. For any given game Γ, the solution obtained in this way heavily depends on the prior probability distributions $p_1, \ldots, p_n$ used as a starting point for the tracing procedure. In the last section, the results of the tracing procedure are given for a simple class of two-person variable-sum games, in numerical detail.

SYNOPSIS

Part I

Part II

1. Introduction

In this paper I shall propose a mathematical procedure, to be called the tracing procedure, for choosing one particular equilibrium point s^* of any finite n-person noncooperative game Γ as solution for Γ.[1)] Choice of s^* as solution for Γ will be based both on game-theoretical considerations in a narrower sense, and on the Bayesian theory of decision making under uncertainty. Use of the latter is motivated by the fact that selecting a strategy in a game, without knowing the other players' strategies in advance, always amounts to decision making under uncertainty, even if this uncertainty is of a very special kind.

I shall assume that, before the players have adopted a more specific theory for predicting the outcome (or solution) of the game, the likely strategy choice of each player i will be assessed by any other player $j \neq i$ in terms of a subjective prior probability distribution p_i over all possible pure strategies that player i may choose. Moreover, all other players j will assign the same prior probability distribution p_i to player i's strategies. This will be the case because each player j will use the same mathematical model and, therefore, will use the same mathematical function F_i, in computing this prior probability distribution $p_i = F_i(\Gamma)$ from the basic parameters of game Γ. Since p_i is a probability distribution over player i's pure strategies, mathematically it will have the nature of a mixed strategy for player i. But, of course, its substantive game-theoretical interpretation will be quite different from that of an ordinary mixed strategy. When another player j ascribes a specific prior distribution p_i to player i's pure strategies, typically he will not mean to assume that player i will intentionally randomize among his alternative pure strategies (as he would do if he used a mixed strategy in the usual sense). Even if player j were fully convinced that player i would always choose a pure strategy, and would never use a mixed strategy at all, it would still make good sense for him to assign subjective probabilities to player i's various pure strategies, so long as j felt uncertain about which particular pure strategy i would actually choose.

More specifically, the prior probability distribution p_i that any other player j assigns to player i's pure strategies $a_i^1,\ldots,a_i^k,\ldots,a_i^{K_i}$ will try to assess, for each pure strategy a_i^k, the theoretical probability p_i^k that this strategy a_i^k would be used by a rational individual playing the role of player i in game Γ. In other words, it will try to assess the probability p_i^k that any given strategy a_i^k should be player i's <u>best reply</u> to the strategies he might expect the other players to use. This means that, other things being equal, the prior probability p_i^k assigned to any given pure strategy a_i^k will be greater the greater the range of possible situations in which a_i^k will be a best reply in the game.

Let $p=(p_1,\ldots,p_n)$ denote the n-tuple of prior probability distributions that the players assign to one another's pure strategies. Then, the tracing procedure I shall propose will be of the form $s^* = T(\Gamma,p)$, in the sense that, for any given game Γ, and for any given n-tuple p of prior distributions, it will always select one particular equilibrium point s^* of Γ as solution for Γ.

In this paper, I shall simply assume that the prior probability distributions p_i used by the players are <u>given</u>. But in a forthcoming joint paper with Reinhard Selten [2), a mathematical procedure will be proposed for generating these prior distributions from the basic parameters of game Γ, i.e., for defining the functions F_i. Even though the problem of choosing appropriate priors p_i for any given game Γ will not be discussed in this paper, I wish to emphasize that resolution of this problem will be a very essential part of our whole theory of solutions for n-person noncooperative games, because in general the tracing procedure to be described will select very different equilibrium points s^* as solutions for any given game Γ, depending on the n-tuple p of prior probability distributions used as a starting point. [3)]

2. Definitions and notations

The kth pure strategy of player i (i=1,...,n) will be called a_i^k, whereas the set of all his K_i pure strategies will be called A_i. Let

$$K = \prod_{i=1}^{n} K_i, \tag{1}$$

and

$$\overline{K_i} = \prod_{j \neq i} K_j = K/K_i. \tag{2}$$

We shall assume that the K possible n-tuples of pure strategies are numbered consecutively as $b^1,\ldots,b^k,\ldots,b^K$. Let

$$b^k=(a_1^{k_1},\ldots,a_i^{k_i},\ldots,a_n^{k_n}). \tag{3}$$

Then we shall write

$$b^k(i) = a_i^{k_i}, \tag{4}$$

to denote the pure strategy used by player i in the strategy n-tuple b^k. The set of all K possible pure-strategy n-tuples will be called B. We have $B = A_1 x\ldots x A_n$.

We shall also assume that the $\overline{K_i}$ possible (n-1)-tuples which can be formed of the pure strategies of the (n-1)players other than a given player i are also numbered consecutively as

$c_i^1,\ldots,c_i^k,\ldots,c_i^{\overline{K_i}}$. The set of all $\overline{K_i}$ (n-1)-tuples of form c_i^k will be called C_i. We have $C_i = A_1 x\ldots x\ A_{i-1}\ x\ A_{i+1}\ x\ldots x\ A_n$.

Any mixed strategy of a given player i (i = 1,...,n) can be identified with a probability vector s_i of the form

$$s_i = (s_i^1,\ldots,s_i^k,\ldots,s_i^{K_i}), \tag{5}$$

where $s_i^1,\ldots,s_i^k,\ldots,s_i^{K_i}$ are the probabilities that this mixed strateg

assigns to i's pure strategies $a_i^1,\ldots,a_i^k,\ldots,a_i^{K_i}$,respectively.The set S of all mixed strategies available to player i will be called his <u>strategy space</u> . S_i is a simplex of (K_i-1) dimensions, consisting of all K_i-vectors satisfying the conditions

$$s_i^k \geqq 0, \quad \text{for } k = 1,\ldots,K_i, \tag{6}$$

and

$$\sum_{k=1}^{K_i} s_i^k = 1. \tag{7}$$

The set $S = S_1 \times \ldots \times S_n$ of all possible n-tuples $s = (s_1,\ldots,s_n)$ of mixed strategies is a convex and compact polyhedron of (K^*-n) dimensions, where

$$K^* = \sum_{n=1}^{n} K_i. \tag{8}$$

S will be called the <u>(collective) strategy space</u> of the n players, or simply the strategy space of game Γ.

We shall write $s = (s_i, \overline{s_i})$, where $\overline{s_i} = (s_1,\ldots,s_{i-1}, s_{i+1},\ldots,s_n)$ is the strategy (n-1)-tuple used by the (n-1) players other than player i. The set $\overline{S_i}$ of all possible strategy (n-1)-tuples of form $\overline{s_i}$ is again a convex and compact polyhedron. Its dimensionality will be $(K^*-n) - (K_i-1) = K^*-K_i-n+1$. We can write $\overline{S_i} = S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$. $\overline{S_i}$ will be called the <u>(collective) strategy space</u> of the (n-1) players opposing player i.

Besides the mixed strategies s_i, which are probability distributions over a given set A_i, we shall also have to consider probability distributions σ_i over a given set C_i, i.e., over all possible (n-1)-tuples c_i^k of pure strategies representing strategy combinations of the (n-1) players opposing a given player i. In a cooperative game, such a probability distribution σ_i could be interpreted as a jointly randomized mixed strategy of these (n-1) players. But, since we shall be dealing with noncooperative games, any such distribution σ_i will always be interpreted as player i's <u>subjective</u> probability distribution over all possible strategy (n-1)-tuples c_i^k that the other (n-1) players could possibly use against him.[4)] Any probability distribution σ_i will be a probability vector of the form

$$\sigma_i = (\sigma_i^1, \ldots, \sigma_i^k, \ldots, \sigma_i^{\overline{K_i}}), \tag{9}$$

where $\sigma_i^1, \ldots, \sigma_i^k, \ldots, \sigma_i^{\overline{K}}$ are the probabilities assigned to the strategy (n-1)-tuples $c_i^1, \ldots, c_i^k, \ldots, c_i^{\overline{K_i}}$, respectively. Of course, these probabilities σ_i^k must satisfy restrictions analogous to (6) and (7).

Let $\sigma_i = \sigma_i(\overline{s_i})$ be that particular probability distribution over the set C_i which obtains when the (n-1) players other than players i use the (n-1)-tuple $\overline{s_i}$ of mixed strategies. Of course, the individual components $\sigma_i^k = \sigma_i^k(\overline{s_i})$ of the probability vector $\sigma_i = \sigma_i(\overline{s_i})$ can be computed by means of the probability-multiplication law in the usual way. For convenience, we shall identify the probability distribution $\sigma_i(\overline{s_i})$ with $\overline{s_i}$ itself and shall write

$$\sigma_i(\overline{s_i}) = \overline{s_i}. \tag{10}$$

Let $\overline{s_i}$ and $\overline{p_i}$ be two strategy (n-1)-tuples of the (n-1)players other than i. Let t be a number with $0 \leq t \leq 1$. Suppose that the probability of $\overline{s_i}$ being used is t, whereas the probability of $\overline{p_i}$ being used is (1-t). Then, this will give rise to a probability distribution π_i of the following form over set C_i:

$$\pi_i = t\sigma_i(\overline{s_i}) + (1-t)\sigma_i(\overline{p_i}). \tag{11}$$

In other words, each component π_i^k of the probability vector $\pi_i = (\pi_i^1, \ldots, \pi_i^{\overline{K_i}})$ will have the form

$$\pi_i^k = t\sigma_i^k(\overline{s_i}) + (1-t)\,\sigma_i^k(\overline{p_i}), \quad k=1,\ldots,\overline{K_i}. \tag{12}$$

In view of (12) we shall sometimes write π_i more concisely as

$$\pi_i = t\ \overline{s_i} + (1-t)\ \overline{p_i}. \tag{13}$$

The set $Q_i(s_i)$ of all strategies a_i^k to which a given mixed strategy s_i assigns positive probabilities $s_i^k > 0$ will be called the carrier of s_i. If the carrier $Q_i(s_i)$ contains only one pure strategy a_i^k, then s_i will be identified with this pure strategy a_i^k, so that we shall write $s_i = a_i^k$. On the other hand, if $Q_i(s_i)$ contains two or more pure strategies, then s_i will be called a proper mixed strategy.

If $Q_i(s_i)$ contains all K_i pure strategies of player i, then s_i will be called a complete(ly) mixed strategy. Finally, if s_i is a proper mixed strategy but is not complete, then it will be called an incomplete(ly) mixed strategy.

For any strategy n-tuple $s = (s_1,\ldots,s_n)$, its carrier $Q(s)$ will be defined as the union of the carriers of its component strategies, that is, as

$$Q(s) = \bigcup_{i=1}^{n} Q_i(s_i). \tag{14}$$

For any strategy n-tuple s having a given set $Q(s) = Q^*$ as its carrier set, we shall say that it belongs to the carrier class Q^*. Since any carrier set Q^* is a subset of the set $A^* = \bigcup_i A_i$ of all K^* pure strategies in game Γ, there can be only a finite number of different carrier sets Q^*, i.e., all strategy n-tuples s of Γ will fall into a finite number of possible carrier classes Q^*.

Suppose that the ith component of a pure-strategy n-tuple b^k is $b^k(i) = a_i^m$, and that a given mixed strategy s_i of player i assigns the probability s_i^m to a_i^m. Then, we shall write

$$q_i^k\ (s_i) = s_i^m\ . \tag{15}$$

Of course, if $s_i = a_i^m$ is a pure strategy, then we have

$$q_i^k\ (a_i^m) = 1 \qquad \text{when } b^k(i) = a_i^m, \tag{16}$$

but

$$q_i^k\ (a_i^m) = 0 \qquad \text{when } b^k(i) \neq a_i^m\ . \tag{17}$$

When the n players use an n-tuple b^k of pure strategies, then player i (i = 1,...,n) will receive the payoff

$$U_i\ (b^k) = u_i^k, \tag{18}$$

whereas if they use an n-tuple $s = (s_1,\ldots,s_n)$ of mixed strategies then i's payoff will be

$$U_i(s) = \sum_{k=1}^{K}\left[\prod_{i=1}^{n} q_i^k\ (\overline{s_i})\right] u_i^k\ . \tag{19}$$

Let $\mathcal{G} = \mathcal{G}(n;\ K_1,\ldots,K_n)$ be the set of all n-person games in which players 1,...,n have exactly $K_1,\ldots,K_n$ pure strategies, respectively. Thus, $\mathcal{G}$ is a set of all games of a given size. Each specific game Γ in $\mathcal{G}$ can be characterized by an (nk)-vector $u = (u_1^1,\ \ldots,u_1^K;\ldots;u_i^1,\ldots,u_i^K;\ldots;u_n^1,\ldots,u_n^K)$, whose components $u_i^k = U_i(b^k)$ are the payoffs to various players i for different pure-strategy n-tuples b^k. This vector $u=u(\Gamma)$ will be called the <u>vector of possible payoffs</u> in game Γ. Each game Γ can be identified with its vector $u=u(\Gamma)$. Accordingly, any set $\mathcal{G} = \{u\}$ can be considered to be an (nK)-dimensional Enclidean space.

Within any given set $\mathcal{G}$, let $\bar{\mathcal{S}}(\mathcal{G})$ be the set of all games Γ about which a certain mathematical statement $\mathcal{S}$ is false. We shall say that $\mathcal{S}$ is true about <u>almost all</u> games Γ if, for every possible set $\mathcal{G}$ of games of a given size, this set $\bar{\mathcal{S}}(\mathcal{G})$ is a closed set of measure zero (or is a subset of a set of this kind) within the (nK)-dimensional Enclidean space $\mathcal{G}$. (Regarding the requirement that $\bar{\mathcal{S}}(\mathcal{G})$ should be <u>closed</u> (or should be a subset of a closed set of an appropriate type), see Debreu [1970, p.387]).

More generally, we shall say that a certain statement $\mathcal{S}$ is true about <u>almost all</u> elements ω of a given finite-dimensional set Ω, if those elements ω about which $\mathcal{S}$ is false form a closed set of measure zero (or form a subset of a set of this kind) within Ω.

3. Best replies and equilibrium points

Any pure or mixed strategy s_i^* of player i is called a best reply to a given strategy (n-1)-tuple $\overline{s_i}$ of the other (n-1) players if

$$U_i(s_i^*,\overline{s_i}) \geqq U_i\ (r_i,\overline{s_i}), \quad \text{for all } r_i \varepsilon S_i. \tag{20}$$

It is easy to verify that:

Lemma 1. Any mixed strategy s_i^* will be a best reply to a given strategy (n-1)-tuple $\overline{s_i}$ if and only if every pure strategy a_i^k in the carrier $Q_i\ (s_i^*)$ of s_i^* is itself a best reply to $\overline{s_i}$.

Let $\overline{S_i}\ (s_i^*)$ be the set of all strategy (n-1)-tuples $\overline{s_i}$ to which a given strategy s_i^* is a best reply. Then, $\overline{S_i}\ (s_i^*)$ will be called the stability set of s_i^* .

Lemma 2. Every stability set $\overline{S_i}\ (s_i^*)$ is a closed subset (possibly empty) of the relevant strategy space $\overline{S_i}$.

The lemma follows from the fact that any stability set is always defined by weak inequalities [of form (20)].

We shall use the shorter notation $\overline{S}_i^k = \overline{S_i}\ (a_i^k)$ to denote the stability set of a pure strategy a_i^k .

Lemma 3. The stability set $\overline{S_i}\ (s_i)$ of a mixed strategy s_i is simply the intersection of the stability sets $\overline{S}_i^k$, $\overline{S}_i^{k'}$,... belonging to the pure strategies a_i^k, $a_i^{k'}$,... in the carrier $Q_i(s_i)$ of s_i.

This lemma is a direct consequence of Lemma 1.

In what follows, unless it is stated otherwise, by "stability set" we shall always mean a stability set $\overline{S}_i^k$ of a pure strategy a_i^k.

Lemma 4. The stability sets $\bar{S}_i^1, \ldots, \bar{S}_i^{K_i}$ cover the whole strategy space $\bar{S}_i$ so that

$$\bigcup_{k=1}^{K_i} \bar{S}_i^k = \bar{S}_i \ . \qquad (21)$$

Proof. The expression $U_i \ (r_i, \bar{s}_i)$ must reach its maximum value at least at one point $r_i = s_i^* \in S_i$, because U_i is a continuous function while S_i is a compact set. Hence, for any $\bar{s}_i$, there exists a strategy s_i^* that is a best reply to $\bar{s}_i$. By Lemma 1, any pure strategy a_i^k in the carrier set $Q_i(s_i^*)$ will have the same property. Therefore, every point $\bar{s}_i$ in $\bar{S}_i$ will lie at least in one stability set $\bar{S}_i^k$. This establishes the lemma.

Lemma 5. Lemma 2 will remain true even if in (21) we use only stability sets $\bar{S}_i^k$ with nonempty interiors.

Proof. Let X be the union of all stability sets $\bar{S}_i^m$ which have an empty interior. Since X is a finite union of sets with empty interiors, it will itself also have this property. Let Y be the union of all stability sets $\bar{S}_i^k$ which have a nonempty interior. Let $\bar{Y} = \bar{S}_i - Y$. By Lemma 2, Y is closed, while $\bar{Y}$ is open. But, by Lemma 4, $X \cup Y = \bar{S}_i$. Therefore, $\bar{Y} \subseteq X$. Thus, $\bar{Y}$ is an open subset of a set with an empty interior. Hence, $\bar{Y}$ must be empty, so that $Y = \bar{S}_i$. This implies the lemma.

Lemma 6. In almost all games Γ, the intersection of two stability sets $\bar{S}_i^k$ and $\bar{S}_i^{k'}$, $k \neq k'$, is either empty or lies within the hypersurface (hyperboloid) $H_i^{kk'}$, that forms the common boundary of $\bar{S}_i^k$ and of $\bar{S}_i^{k'}$.

Proof. In view of (21), any point $\bar{s}_i$ lying in the intersection of $\bar{S}_i^k$ and of $\bar{S}_i^{k'}$, $k \neq k'$, must satisfy

$$U_i \ (a_i^k, \bar{s}_i) - U_i \ (a_i^{k'}, \bar{s}_i) = 0. \qquad (22)$$

The left-hand side of this equation is a multilinear form in the probabilities s_j^m $(j \neq i)$ characterizing the various component strategies s_j of the (n-1)-tuple $\bar{s}_i$; and its coefficients are payoff differences of the form $(u_i^k - u_i^{k'})$, $k \neq k'$. If this equation is not an identity, then all points $\overline{s_i}$ satisfying it will lie on a hypersurface (hyperboloid) $H_i^{kk'}$, and so the lemma will be true.

The lemma can fail only if equation (22) is in fact an identity, which can happen only if a sufficient number of the payoff differences $(u_i^k - u_i^{k'})$ vanish, i.e., if the vector $u = u\,(\Gamma)$ of pure-strategy payoffs has a sufficient number of pairwise equal components. But the set $\mathcal{G}^*$ of all games having this property is a closed set of measure zero in $\mathcal{G}$ (because it will be the intersection of a certain number of hyperplanes in $\mathcal{G}$). This completes the proof of Lemma 6.

Lemma 7. For almost all games Γ the following statement is true: For almost all strategy (n-1)-tuples $\overline{s_i}$, the set of all best replies to $\overline{s_i}$ in Γ will consist of a unique pure strategy a_i^k.

Proof. By Lemmas 3 and 7, any strategy (n-1)-tuple $\overline{s_i}$ to which there exist two or more different best replies must lie on the common boundary of two or more stability sets $\bar{S}_i^k$, $\bar{S}_i^{k'}$, ... But, by Lemma 6, in almost all games Γ, the union of these boundary sets is a closed set of measure zero in the strategy space $\overline{S_i}$. This establishes the lemma.

Any pure strategy a_i^k whose stability set $\bar{S}_i^k$ has an empty interior as a subset of $\overline{S_i}$ will be called an inferior strategy. A special case of an inferior strategy is a strictly inferior strategy whose stability set $\bar{S}_i^k$ as a whole is empty.[5)] In view of Lemma 5, to any strategy (n-1)-tuple $\overline{s_i}$, player i will always have a best reply that is not an inferior strategy.

I shall now argue that, this being the case, it is always advantageous for player i to use a noninferior best-reply strategy when he expects the other players to use the strategy combination $\overline{s_i}$. This is so because it is always preferable to use a best-reply strategy with a larger stability set than one with a smaller stability set since this way one increases the chance that one's strategy will remain a best reply even if the other players somewhat deviate from their expected strategies. Yet, the stability set of a noninferior strategy is a set of full dimensionality in the space $\bar{S}_i$, and therefore is incomparably "larger" than the stability set of an inferior strategy which is always a set of less-than-full dimensionality.

Of course, there are even stronger reasons for any player to avoid using strictly inferior strategies, since these are never best replies in any conceivable situation.

Accordingly, we feel that, before any further analysis is applied to any given game Γ, all inferior pure strategies should be eliminated, since they are strategies that will never be used by rational players. Of course, this very elimination of these strategies may make inferior strategies out of some other strategies that did not use to be inferior. Therefore, elimination of all inferior strategies should be repeated as many times as is necessary in order to obtain a game completely without inferior strategies. Since the set $A^* = \bigcup_i A_i$ of all pure strategies in game Γ is a finite set, this procedure will always come to an end after a finite number of repetitions.[6)] In what follows, it will be assumed that game Γ does not any longer contain inferior strategies.

A given strategy n-tuple $s = (s_1,\ldots,s_n)$ is an equilibrium point [Nash, 1951] if every component s_i in s is a best reply to the (n-1) - tuple $\bar{s}_i$ formed by the remaining (n-1) components.

In view of Lemma 1 and condition (20), a necessary and sufficient condition for any given strategy n-tuple $s = (s_1,\ldots,s_n)$ to be an equilibrium point is that the following equations and inequalities should be satisfied:

$$U_i(a_i^k,\bar{s}_i) = U_i(a_i^{k'},\bar{s}_i), \quad \text{if } a_i^k, a_i^{k'} \in Q_i(s_i), \tag{23}$$

and

$$U_i(a_i^k, \overline{s_i}) \geqq U_i\ (a_i^{k'}, \overline{s_i}), \quad (24)$$

$$\text{if } a_i^k \varepsilon Q_i(s_i) \text{ whereas } a_i^{k'} \notin Q_i(s_i),$$

$$\text{for } i = 1,\ldots,n.$$

In the special case where all n equilibrium strategies $s_1,\ldots,s_n$ are pure strategies, we shall have only inequalities of form (24); while in the special case where all n equilibrium strategies are completely mixed, we shall have only equations of form (23).

An equilibrium point s is called strong if every player's equilibrium strategy s_i is his only best reply to the other players' strategy combination $\overline{s_i}$. (In view of Lemma 1, any strong equilibrium point s must be in pure strategies, so that we can write $s = b^k$.) An equilibrium point that is not strong is called weak.

An equilibrium point s is called quasi-strong if no player i has pure-strategy best replies to $\overline{s_i}$ other than the pure strategies belonging to the carrier $Q_i(s_i)$ of his equilibrium strategy s_i. An equilibrium point that is not even quasi-strong is called extra-weak.

4. The linear tracing procedure

Most game theorists agree that the solution s^* of any noncooperative game Γ must be an equilibrium point. The reason is that any solution concept yielding a non-equilibrium-point strategy combination s as solution would be self-defeating: the very expectation that all other players j would use the strategies s_j prescribed by this solution concept would give an incentive for at least one player i not to use the strategy s_i prescribed by it. (This is so because, by assumption, s would not be an equilibrium point and, therefore, at least one player i would find that his prescribed strategy s_i was not a best reply

to the other players' prescribed strategy combination $\overline{s_i}$.)
The requirement that the solution s^* of a noncooperative game should always be an equilibrium point will be called the requirement of <u>internal optimality</u>. (It is called "internal" because it requires that the <u>components</u> s_i^* of the solution s^* should be related to one another in a certain particular way.)

On the other hand, Bayesian decision theory suggests that the solution strategy s_i^* of each player i should be a best reply to the subjective probability distribution σ_i he is entertaining over all possible strategy combinations c_i^k that the other (n-1) players may use against him. Moreover, this probability distribution σ_i should be derived, in a suitable manner, from the prior distributions $p_1, \ldots, p_{i-1}, p_{i+1}, \ldots p_n$ that the originally assigned to these players' strategies. This will be called the <u>external-optimality requirement</u>. (It is called "external" because it requires that the solution s^* should be related to a mathematical entity external to s^*, viz. to the n-tuple p of prior distributions, in a certain particular way.)

Of course, the simplest interpretation of this external-optimality requirement would be to demand that the solution strategy s_i^* of each player i should be a best reply to the prior distributions $p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_n$ <u>themselves</u>, with $\sigma_i = \overline{p_i} = (p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_n)$. But this simple-minded interpretation of the external-optimality requirement would clearly violate the internal-optimality requirement, because normally the n-tuple $s^* = (s_1^*, \ldots, s_n^*)$ of these best-reply strategies would not be an equilibrium point.

For this reason, we shall interpret the external-optimality requirement as follows: The n players will find the solution s^* of a given game Γ through an intellectual process of <u>convergent expectations</u>, to be called the <u>solution process</u>. At the beginning of this process, the players' expectations about each other's strategies will take the form of the prior distributions p_i. During this process, they will continually and systematically modify these expectations - - until, at the end of this process, their expectations will come to converge to one particular equilibrium point s^* of game Γ, which, then, will be accepted as the solution for Γ. The tracing procedure I shall describe is meant to be a mathematical representation of this solution process through which the players choose a particular equilibrium point $s^* = T(\Gamma, p)$ as solution.

In order to define this tracing procedure for a given game Γ, and for a given n-tuple p of prior distributions, I shall first define a one-parameter family of n-person games, $\{\Gamma^t\}$, $0 \leqq t \leqq 1$, as follows. In every game Γ^t, each player i will have the same strategy space S_i as he has in the original game Γ. For any specific value of t, let $\Gamma^t = G(\Gamma, p, t)$ be a game in which the payoff function V_i of each player i ($i = 1,\ldots,n$) is of the form

$$V_i(s_i, \overline{s_i}; p, t) = t\, U_i(s_i, \overline{s_i}) + (1-t) U_i(s_i, \overline{p_i}), \tag{25}$$

where U_i is player i's payoff function in game Γ, s_i is i's own strategy, $\overline{s_i}$ is the strategy (n-1)-tuple used by the other (n-1) players, whereas p is the n-tuple of prior distributions, and $\overline{p}_i$ is an (n-1)-tuple of these priors.

Clearly, $\Gamma^1 = \Gamma$, while Γ^0 is a game with the payoff functions

$$V_i(s_i, \overline{s_i}; p, 0) = U_i\ (s_i, \overline{p_i}). \tag{26}$$

Thus, Γ^0 is a game of a rather special form, in which the payoff of each player i will depend only on his own strategy s_i, but will not depend on the other players' strategy combination $\overline{s_i}$. Any strategy n-tuple $s = (s_1, \ldots, s_n)$ will be an equilibrium point in Γ^0 if and only if each component s_i of s is a best reply, in the original game Γ, to the (n-1)-tuple $\overline{p_i}$ of prior distributions. Consequently, in view of Lemma 7, we can state:

Lemma 8. For almost all games Γ, the following statement is true: For almost all choices of the n-tuple p of prior distributions, game $\Gamma^0 = G(\Gamma, p, 0)$ will have exactly one equilibrium point $s^0 = (s^0, \ldots, s^0_n)$. Each component s^0_i of s^0 will be a pure strategy $s^0_i = a_i^{k_i}$, and will be characterized by being the only best reply, in the original game Γ, to the (n-1)-tuple $\overline{p_i}$ of prior distributions. (Hence, s^O will be a strong equilibrium point of Γ^O.)

Lemma 8 can also be stated in a slightly different form:

Lemma 8*. For a given choice of p, almost all games Γ will give rise to a game $\Gamma^O = G(\Gamma,p,O)$ having only one strong equilibrium point $s^O = b^k$ in pure strategies.

Proof of Lemma 8* is similar to the proof of Lemma 7.

Let E^t be the set of all equilibrium points in game Γ^t. As each game Γ^t is a finite game, by Nash's [1951] theorem, E^t will always be nonempty. Let $P = P(\Gamma,p)$ be the graph of the correspondence $t \rightarrow E^t$, $0 \leqq t \leqq 1$. P will be typically a collection of one-dimensional piecewise algebraic curves,[7] though in degenerate cases it may also include isolated points and/or subsets of more than one dimension.

P is always a subset of the polyhedral set $R = I \times S$, where $I = [0,1]$ is the closed unit interval, while S is the strategy space of game Γ. Each point of R will have the form $x = (t,s)$. (The t coordinate will always be written first.) The strategy n-tuple s occurring in $x = (t,s)$ will be called the strategy component of this point x.

For any specific value of t, $0 \leqq t \leqq 1$, let R^t be the set of all points x in R having this specific t value as their first coordinate. Clearly, any set R^t can be considered to be the representation, within set R, of the strategy space S of game Γ^t.

Suppose the graph P contains a path L connecting a point of form (O,s^O) (where s^O is of course an equilibrium point of game Γ^O) with a point of form $(1,s^*)$ (where s^* is of course an equilibrium point of game $\Gamma^1 = \Gamma$). Then L will be called a feasible path.

Clearly, the strategy part s of this point x will be an equilibrium point of game Γ^t. We shall call s a distinguished equilibrium point of Γ^t. In most cases, any given game Γ^t will have only one distinguished equilibrium point s, because normally L will intersect each set R^t in a unique point $x = (t,s)$. But it can happen that a whole one-dimensional segment Λ of L will lie in a given set R^t. In this case, of course, game Γ^t will have infinitely many distinguished equilibrium points. However - - if a distinguished path L exists at all - - game Γ^O [illegible] distinguished equilibrium point s^O.

point $x^O = (O,s^O)$ of the distinguished path L; then he will move continuously along this path L; until, at the end of the solution process, he will come to base his analysis on the solution s^* itself, defined by the end point $x^* = (1,s^*)$ of L.

At the <u>beginning</u> of the solution process, each player i will be in a state of complete <u>predictive uncertainty</u>, in the sense that he will lack any specific theory about what the outcome of the game will be. Therefore, he will feel unable to make specific predictions about the other players' strategy (n-1)-tuple $\overline{s_i}$. Accordingly, he will entertain a subjective probability distribution of the form

$$\pi_i\ (O,s^O) = \sigma_i(\overline{p_i}) = \overline{p_i} \tag{27}$$

over all possible pure-strategy combinations c_i^k that the other (n-1) players may use. Lacking any alternative theory, this probability distribution will be completely based on the prior probability distributions $p_1,\ldots,p_{i-1}, p_{i+1},\ldots,p_n$ he assigns to the other players' pure strategies.

In contrast, at the <u>end</u> of the solution process, each player i will be in the state of <u>predictive certainty</u> because he will feel able to predict that the other (n-1) players will use the strategy (n-1)-tuple $\overline{s_i^*} = (s_1^*,\ldots,s_{i-1}^*, s_{i+1}^*,\ldots,s_n^*)$ prescribed by the solution s^*. Therefore, his subjective probability distribution over all possible pure-strategy combinations c_i^k of the other (n-1) players will be of the form

$$\pi_i\ (1,s^*) = \sigma_i(\overline{s_i^*}) = \overline{s_i^*}\ . \tag{28}$$

At any given moment <u>during</u> the solution process, when his analysis of the game situation is based on a given distinguished equilibrium point s of a particular game Γ^t, this equilibrium point s will give him the prediction that the other (n-1) players will use the strategy (n-1)-tuple $\overline{s_i}=(s_1,\ldots,s_{i-1}, s_{i+1},\ldots,s_n)$ corresponding to s. But he will know that he <u>cannot have full confidence in this prediction</u> because in general a distinguished equilibrium point s (or, for that matter, any other equilibrium point s') of any game Γ^t with $t \neq 1$ will not be an equilibrium point of the original game Γ, and therefore will not be a possible solution for Γ. Hence, player i will give only the probability weight t to this prediction, and will retain the remaining probability weight (1-t) for his initial probability distribution $\pi_i(0,s^0) = \overline{p_i}$. Accordingly, his subjective probability distribution $\pi_i(t,s)$ over all possible (n-1)-tuples c_i^k will now take the form

$$\pi_i(t,s) = t\,\sigma_i\,(\overline{s_i}) + (1-t)\,\sigma_i\,(\overline{p_i}) = t\overline{s_i} + (1-t)\,\overline{p_i} \tag{29}$$

(Of course, equations (27) and (28) are special cases of equation (29).)

This, however, means that, if player i <u>had to act at this very moment</u>, then he would choose a strategy that would be a best reply [8], in the original game Γ, to this probability distribution π_i (t,s). More specifically, I shall assume that, if he had to act at this moment, then he would always choose his equilibrium strategy s_i, corresponding to the distinguished equilibrium point s. (In view of (25), this strategy s_i will always have the required property: it will always be a best reply, in the original game Γ, to $\pi_i\,(t,s) = t\,\overline{s_i} + (1-t)\,\overline{p_i}$.)

Thus, if player i had to act at the beginning of the solution process, then he would choose strategy s_i^0 (which is his best reply, in game Γ, to $\pi_i\,(0,s^0) = \overline{p_i}$, i.e., to the (n-1)-tuple $\overline{p_i}$ of prior distributions). On the other hand, when he has to act at the end of the solution process or at any later time, then he will always choose strategy s_i^*, prescribed for him by the solution s^*.

In other words, any distinguished equilibrium point s of a given game Γ^t with $t < 1$, will furnish a correct <u>conditional</u> prediction about each player i's behavior: if he had to act at that stage of the solution process, then he would in fact choose strategy s_i. But it will not necessarily furnish a correct prediction about what i's behavior will be at the end of the solution process, when he actually has to make his final choice of strategy in the game - - since at that time he will choose his solution strategy s_i^*, which may be different from s_i. By the same token, player i will assign only the probability weight $t < 1$ to the prediction that the other (n-1) players will use the strategy (n-1)-tuple $\overline{s_i}$.

Accordingly, under this model, as any given player is moving along the distinguished path L during the solution process, he will place more and more confidence in the predictions that the distinguished equilibrium points s of the various games Γ^t provide for him, and will rely less and less on the prior distributions p_i as predictive devices - - until, at the end of the solution process, he will place full confidence in the prediction that the solution s^* itself yields him about the other players' behavior. Hence, the solution process is a process of continually increasing predictive certainty, and of continually decreasing predictive uncertainty.

6. The logarithmic tracing procedure

The purpose of the logarithmic tracing procedure is to approximate the piecewise algebraic graph P of the linear tracing procedure by a fully algebraic graph P^*, which has certain desirable mathematical properties that P itself lacks.

Once more, for any given game Γ, for any given n-tuple p of prior distributions - - and, this time, also for any given small positive constant ε - - I shall define a one-parameter family of n-person games, $\{\Gamma^t_*\}$, $0 \leq t \leq 1$, as follows. In every game Γ^t_*, each player i will have the same strategy space S_i as he has in the original game Γ. For any specific value of t, let $\Gamma^t_* = G^*(\Gamma, p, \varepsilon, t)$ be a game in which the payoff function V_i^* of each player i (i = 1,...,n) is of the form

$$V_i^*(s_i,\overline{s_i};p,\varepsilon,t) = t\,U_i(s_i,\overline{s_i}) + (1-t)\,U_i(s_i,\overline{p_i}) + \varepsilon(1-t)\sum_{k=1}^{K_i} \log s_i^k. \qquad (30)$$

Obviously, $\Gamma_*^1 = \Gamma$, while Γ_*^0 is a game with payoff functions of the form

$$V_i^* (s_i,\overline{s_i};\ p,\varepsilon,0) =$$

$$= U_i \ (s_i,\overline{p_i}) + \varepsilon \sum_{k=1}^{K_i} \log s_i^k = \sum_{k-1}^{K_i} s_i^k \, U_i(a_i^k,\overline{p_i}) + \varepsilon \sum_{k=1}^{K_i} \log s_i^k \ . \tag{31}$$

Thus, game Γ_*^0 is again a game of a very special form, in which the payoff of each player i will depend only on his own strategy s_i, but will not depend on the other players' strategy (n-1)-tuple $\overline{s_i}$.

<u>Lemma 9</u>. In any game Γ_*^t with $0 \leqq t < 1$, to any strategy (n-1)-tuple $\overline{s_i}$ of the other players, each player i will have exactly one best reply s_i^*, which will always be a completely mixed strategy.

<u>Proof.</u> The uniqueness of the best reply $s_i = s_i^*$ follows from the fact that V_i^* is a strictly concave function in the probabilities s_i^k. On the other hand, s_i^* must be completely mixed, because any pure or incompletely mixed strategy would yield player i an infinite negative payoff $V_i^* = -\infty$ whereas any completely mixed strategy will yield him a finite payoff $V_i^* > -\infty$.

<u>Lemma 10</u>. Game Γ_*^0 will always have exactly one equilibrium point $s^0 = (s_1^0,\ldots,s_n^0)$.

<u>Proof.</u>Lemma 10 directly follows from Lemma 9 in view of the fact that, by (31), the payoff of each player i in Γ_*^0 will depend on his own strategy s_i.

<u>Lemma 11</u>. In any game Γ_*^t with $0 \leqq t < 1$, any strategy n-tuple $s = (s_1,\ldots,s_n)$ will be an equilibrium point if and only if its component strategies s_i satisfy the equations:

$$\sum_{k=1}^{K_i} s_i^k - 1 = 0, \quad \text{for } i = 1,\ldots,n; \tag{32}$$

and

$$t\left[U_i(a_i^k,\overline{s_i}) - U_i(a_i^1,\overline{s_i})\right] + (1-t)\left[U_i\;(a_i^k,\overline{s_i}) - U_i(a_i^1,\overline{s_i})\right] +$$

$$+\ \varepsilon(1-t)(s_i^1-s_i^k)\Big/(s_i^1\ s_i^k) = 0\ , \qquad (33)$$

for $i=1,\ldots,n$; and, given any specific value of i, for $k=2,\ldots,K_i$.

Equation (33) can also be written in the form

$$t\ s_i^1\ s_i^k\left[U_i(a_i^k,\overline{s_i}) - U_i(a_i^1,\overline{s_i})\right] + (1-t)s_i^1\ s_i^k\left[U_i\;(a_i^k,\overline{p_i}) - U_i(a_i^1,\overline{s_i})\right] +$$

$$+\ \varepsilon(1-t)(s_i^1-s_i^k) = 0 \qquad (34)$$

Proof. Equation (32) is a restatement of (7). The equations of form (33) are the first-order conditions for maximizing the payoff functions V_i^*. Since the latter are strictly concave functions in the probabilities s_i^k , these first-order conditions are not only necessary conditions of maximization but are also sufficient conditions. The admissibility of writing (33) in the form (34) follows from the fact that, in view of 9, all variables s_i^k will be nonzero at any equilibrium point.

Lemma 12. Any equilibrium point of game $\Gamma_*^1 = \Gamma$ will satisfy conditions (32) and (34). But, in general, not every strategy n-tuple s satisfying these conditions will be an equilibrium point of Γ.

Proof. The first sentence of the lemma follows from (23). The second sentence follows from the fact that e.g., every n-tuple $s = b^k$ of pure strategies in Γ will satisfy (32) and (34) if we set $t = 1$, whether $s = b^k$ is an equilibrium point or not.

Let E_*^t be the set of all equilibrium points in a given game Γ_*^t. Let $P^* = P^*(\Gamma,p,\varepsilon)$ be the graph of the correspondence $t \rightarrow E_*^t$, $0 \leqq t \leqq 1$.

Let Π be the set of all points (t,s) satisfying conditions (32) and (34). We have n equations of form (32) and $(K^* - n)$ equations of form (34), where K^* is the number defined by (8). Thus, we have all together K^* independent equations in K^* variables of form s_i^k and in the one additional variable t, i.e., in $(K^* + 1)$ variables. Consequently, Π will be typically a one-dimensional algebraic variety, i.e., an algebraic curve. Let $\Pi' = \Pi \cap R$, where $R = I \times S$ is the polyhedral set defined in Section 3. Clearly, Π' will be typically a subset of a one-dimensional algebraic curve, though in degenerate cases it may include isolated points and/or subsets of more than one dimension.

Lemma 13. $P^* \subseteq \Pi'$. More particularly, in the region $0 \leq t < 1$, P^* will coincide with Π' but, in the region $t = 1$, Π' may include points not belonging to P^*.

Lemma 13 follows from Lemma 12.

By Lemma 13, graph P^*, unlike the graph P used in the linear tracing procedure is (a subset of) a fully algebraic curve, instead of being merely a piecewise algebraic curve.

Let $\hat{R} = I^0 \times \bar{S}$, where $I^0 = (0,1)$ is the open unit interval whereas $\bar{S}$ is the boundary of the strategy space S, consisting of all strategy n-tuples s having at least one component s_i that is a pure or an incompletely mixed strategy. Let R^0 and R^1 be the sets of all those points in R which have the form $(0,s)$ or the form $(1,s)$, respectively. Clearly, the boundary $\bar{R}$ of set R consists of the three disjoint sets, R^0, R^1, and $\hat{R}$.

Lemma 14. Let s^0 be the unique equilibrium point of game Γ_*^0. Then, graph P^* will have a branch L^* starting at point $x^0 = (0,s^0)$, and locally unique in a finite neighborhood of this point. If we continue this branch L^* analytically long enough, then it will eventually intersect set R^1 at a point $x^* = (1,s^*)$. The strategy part s^* of this point x^* will be an equilibrium point of game $\Gamma_*^1 = \Gamma$.

Proof. Let us define K^* functions φ_i^k as follows. For $i = 1,\ldots,n$, we define φ_i^1 (t,s) as the left-hand side of equation (32). For $i = 1,\ldots,n$, and for $k = 2,\ldots,K_i$, we define $\varphi_i^k(t,s)$ as the left-hand side of equation (34). Let J be the Jacobian

$$J = J(t,s) = \frac{\partial(\varphi_1^1,\ldots,\varphi_1^{K_1};\ldots;\varphi_i^1,\ldots,\varphi_i^{K_i};\ldots;\ \varphi_n^1,\ldots,\varphi_n^{K_n})}{\partial(s_i^1,\ldots,s_1^{K_1};\ldots;s_i^1,\ldots,s_i^{K_i};\ldots;\ s_n^1,\ldots,s_n^{K_n})}\,. \quad (35)$$

It is easy to verify that, at the point $x^O = (O,s^O)$, this Jacobian J will never vanish. Consequently, by the Implicit Function Theorem, graph P^* will have a branch L^* starting at x^O, and this branch will be locally unique in a finite neighborhood of x^O. On the other hand, since L^* is an algebraic curve, if it is analytically continued long enough, then it will once more intersect the boundary $\bar{R}$ of set R (cf. Harsanyi [1973], Lemmas 2 and 3, on pp. 241-242). However, this second intersection point x cannot again be the point x^O with $x^O = x$, because then L^* would not be locally unique near x^O. Nor can x be another point $x \neq x^O$ of set R^O because, by Lemma 10, s^O is the only equilibrium point of game Γ_*^O. Finally, x cannot be a point of set $\hat{R}$, either, because by Lemma 9, none of the games Γ_*^t with $O<t<1$ can have an equilibrium point lying on the boundary $\bar{S}$ of the strategy space S. Consequently, this intersection point x must lie in set R^1 and, therefore, must have the form $x = (1,s)$.

I now wish to show that the strategy part s of this point x is an equilibrium point of game Γ. Let $\bar{L} = L^* - \{x\}$. Since $L^* \subseteq \Pi'$, all points (t,s) of $\bar{L}$ must satisfy condition (34). But $\varepsilon>0$, and in view of Lemma 9, $s_i^1>0$ and $s_i^k>0$. Therefore, the expression

$$\Delta = t\left[U_i(a_i^k,\overline{s_i}) - U_i(a_i^1,\overline{s_i})\right] + (1-t)\left[U_i(a_i^k,\overline{p_i}) - U_i(a_i^1,\overline{p_i})\right] \quad (36)$$

must have the same sign as the expression $(s_i^k-s_i^1)$ has.

Consequently, we can write

$$\Delta(s_i^k - s_i^1) \geqq 0. \tag{37}$$

On the other hand, point x = (1,s) is a limit point of $\bar{L}$ and, therefore, its strategy part s must likewise satisfy (34) and (37) if we set t = 1. When we do this, then (34) will take the form

$$s_i^1 \; s_i^k \left[U_i(a_i^k,\overline{s_i}) - U_i(a_i^1,\overline{s_i})\right] = 0, \tag{38}$$

whereas, in view of (36), (37) will take the form

$$(s_i^k - s_i^1) \left[U_i(a_i^k,\overline{s_i}) - U_i(a_i^1,\overline{s_i})\right] \geqq 0. \tag{39}$$

Without loss of generality, we can assume that player i's strategies are numbered in such a way that

$$s_i^1 \geqq s_i^k, \qquad \text{for } k = 2,\ldots,K_i. \tag{40}$$

in view of (7), this implies that

$$s_i^1 > 0. \tag{41}$$

However, in view of (38 and (41), we must have

$$U_i(a_i^k,\overline{s_i}) = U_i(a_i^1,\overline{s_i}), \quad \text{whenever } s_i^k > 0. \tag{42}$$

On the other hand, in view of (39), (40) and (41), we must have

$$U_i(a_i^k,\overline{s_i}) \leqq U_i(a_i^1,\overline{s_i}), \quad \text{whenever } s_i^k = 0. \tag{43}$$

Yet, (42) and (43) imply (23) and (24). Consequently, the strategy part $s = s^*$ of $x = (1,s) = x^* = (1,s^*)$ is an equilibrium point of game Γ. This complets the proof of Lemma 14.

Since, in general, the curve L^* defined by Lemma 14 will depend on Γ, p and ε, we shall write $L^* = L^*(\Gamma,p,\varepsilon)$. Let $s^* = s^*(\Gamma,p,\varepsilon)$ be the strategy part of the end point $(1,s^*)$ of L^*. Then, we define

$$s^{**} = s^{**}(\Gamma,p) = \lim_{\varepsilon \to 0} s^*(\Gamma,p,\varepsilon). \tag{44}$$

<u>Lemma 15</u>. The limit indicated by (44) always exists. Moreover, the limit point s^{**} is always an equilibrium point of game Γ.

<u>Proof.</u> The existence of this limit follows from the fact that L^* is an algebraic curve in ε, in t and in the probabilities s_i^k. On the other hand, by Lemma 14, any point $s^*(\Gamma,p,\varepsilon)$ with $\varepsilon > 0$ is an equilibrium point of Γ. Therefore, s^{**} is a limit point of the set E^1 of all equilibrium points in Γ. As E^1 is a closed set, s^{**} itself must also be an equilibrium point of Γ.

We are now in a position to give our <u>first definition</u> for the logarithmic tracing procedure and for the solution s^{**} specified by it. (Shortly, we shall give a second definition, which will, however, always yield the same solution.) For any given game Γ, for any n-tuple p of prior distributions, and for any small positive constant ε, the <u>logarithmic tracing procedure</u> consists in following the curve $L^* = L^*(\Gamma,p,\varepsilon)$ from its starting point $x^0 = (0,s^0)$ to its end point $x^* = (1,s^*)$, and then in finding the <u>solution</u> s^{**} by means of the limit operation indicated by (43).

For any point $x = (t,s)$ of the curve $L^* = L^*(\Gamma,p,\varepsilon)$, let $\lambda(t,s)$ denote its distance from the starting point $x^0 = (0,s^0)$, as measured along the curve L^*. We define

$$\lambda^*(t,s) = \lambda(t,s) / \lambda(1,s^*) \ , \tag{45}$$

where $x^* = (1,s^*)$ is the end point of L^*. The variable λ^* can be used to parametrize the curve L^*, by means of one equation of the form

$$t = \psi_0(\lambda^*,\varepsilon) \tag{46}$$

and by means of K^* equations of the form

$$s_i^k = \psi_i^k (\lambda^*,\varepsilon), \tag{47}$$

for $i=1,\ldots,n$; for $k=1,\ldots,K_i$; and for $0 \leqq \lambda^* \leqq 1$.

The K^* equations of form (47) can also be written more concisely in vector notation, in the form

$$s = \psi(\lambda^*,\varepsilon). \tag{48}$$

We now define a new curve $L^{**} = L^{**}(\Gamma,p)$, by means of the two parametric equations

$$t = \bar{\psi}_o(\lambda^*) = \lim_{\varepsilon \to 0} \psi_o(\lambda^*,\varepsilon), \tag{49}$$

and

$$s = \bar{\psi}(\lambda^*) = \lim_{\varepsilon \to 0} \psi(\lambda^*,\varepsilon), \tag{50}$$

for $0 \leqq \lambda^* \leqq 1$.

Since L^* is an algebraic curve in ε, in t and in the probabilities s_i^k, the limits indicated by (49) and by (50) will always exist. We shall call the curve L^{**} defined by (49) and (50) the <u>limit curve</u> since it represents the limit position of curve L^* when ε goes to zero.

Now, we shall state the <u>second definition</u> for the logarithmic tracing procedure and for the solution specified by it. We shall say that, for a given game Γ, and for a given n-tuple p of prior distributions, the <u>logarithmic tracing procedure</u> consists in following the limit curve $L^{**} = L^{**}(\Gamma,p)$ from its starting point $x^{oo} = (0,s^{oo})$, corresponding to $\lambda^* = 0$, to its end point $x^{**} = (1,s^{**})$, corresponding to $\lambda^* = 1$. The <u>solution</u> $s^{**} = T^*(\Gamma,p)$ specified by this logarithmic tracing procedure is the strategy part $s^{**} = \bar{\psi}(1)$ of this end point x^{**}.

Lemma 16. The solutions s^{**} defined by the first and by the second definitions are identical.

Proof. For any given choice of Γ, of p, and of ε, we have $s^*(\Gamma,p,\varepsilon) = \psi(1,\varepsilon)$ since both expressions denote the strategy part s^* of the end point $(1,s^*)$ of the curve $L^* = L^*(\Gamma,p,\varepsilon)$. Consequently,

$$\lim_{\varepsilon \to 0} s^*(\Gamma,p,\varepsilon) = \lim_{\varepsilon \to 0} \psi(1,\varepsilon). \tag{51}$$

This establishes the lemma.

Theorem 1. The solution $s^{**} = T^*(\Gamma,p)$ defined by the logarithmic tracing procedure always exists and is always unique. Accordingly, the logarithmic tracing procedure is always well defined.

Proof. The theorem follows from Lemmas 14,15 and 16.

Lemma 17. The limit curve L^{**} is always a subset of the graph P, used in the linear tracing procedure.

Proof. For any given value of λ^* with $0 \leqq \lambda^* \leqq 1$, let $\hat{t} = \bar{\psi}_0(\lambda^*)$ and $\hat{s} = \bar{\psi}(\lambda^*)$. We have to show that, for all choices of λ^*, $\hat{s}$ is an equilibrium point of game $\Gamma^{\hat{t}}$. For $\lambda^* = 1$, this follows from Lemma 15. Thus, we have to consider only the case where $0 \leqq \lambda^* < 1$. For any positive ε, the variables $t = \psi_0(\lambda^*, \varepsilon)$ and $s_i^k = \psi_i^k(\lambda^*,\varepsilon)$ will satisfy condition (34), because (t,s) will be a point of curve $L^* = L^*(\Gamma,p,\varepsilon)$. In view of (25), this condition can also be written in the form

$$s_i^1 s_i^k \left[V_i(a_i^k,\bar{s}_i;p,t) - V_i(a_i^1,\bar{s}_i;p,t) \right] + \varepsilon(1-t)(s_i^1-s_i^k) \tag{52}$$

However, as $\lambda^* < 1$, we have $t < 1$, and $1-t > 0$. Also $\varepsilon > 0$. Moreover, by Lemma 9, $s_i^1 > 0$ and $s_i^k > 0$. Consequently, the expression in square brackets in (52) must have the same sign as the expression $(s_i^k - s_i^1)$. Therefore, we can write

$$(s_i^k - s_i^1)\left[V_i(a_i^k,\overline{s_i};\ p,t) - V_i(a_i^1,\overline{s_i};\ p,t)\right] \gtreqless 0. \qquad (53)$$

Since the point $(\hat{t},\hat{s})$ is the limit of such points (t,s) when ε goes to zero, the variables $\hat{t}$ and $\hat{s}_i^k$ must likewise satisfy (52) and (53) if we set $\varepsilon = 0$. Moreover, we can assume without loss of generality that player i's pure strategies a_i^k are numbered in such a way that

$$\hat{s}_i^1 \geqq \hat{s}_i^k, \qquad \text{for } k = 2,\ldots,K_i \qquad (54)$$

so that

$$\hat{s}_i^1 > 0. \qquad (55)$$

But then, by (52) and (55), we must have

$$V_i\ (a_i^k,\overline{\hat{s}_i}\ ;\ p,t) = V_i\ (a_i^1,\overline{\hat{s}_i}\ ;\ p,t), \qquad (56)$$
$$\text{whenever } \hat{s}_i^k > 0;$$

and

$$V_i(a_i^k,\overline{\hat{s}_i}\ ;\ p,t) \leqq V_i(a_i^1,\overline{\hat{s}_i}\ ;\ p,t), \qquad (57)$$
$$\text{whenever } \hat{s}_i^k = 0.$$

Yet, (56) and (57) imply conditons (23) and (24) as applied to game $\Gamma^{\hat{t}}$. Consequently, the strategy n-tuple $\hat{s}$ is an equilibrium point of game $\Gamma^{\hat{t}}$. This completes the proof.

<u>Theorem 2</u>. For any possible choice of Γ and of p, the linear tracing procedure is always feasible.

<u>Proof</u>. We have to show that graph $P = P(\Gamma,p)$ always contains a feasible path. But this follows from Lemma 17, which implies that at least the limit curve L^{**} itself is such a feasible path.

<u>Theorem 3</u>. For some choices of Γ and p, the linear tracing procedure is not well defined.

<u>Proof.</u> We shall adduce a numerical example for which the linear tracing procedure is not well defined. Consider the following two-person nonzero-sum game Γ:

	a_2^1	a_2^2
a_1^1	2,1	0,0
a_1^2	0,0	1,2

This game has three equilibrium points. Two of them, viz. $b^1 = (a_1^1, a_2^1)$ and $b^2 = (a_1^2, a_2^2)$, are in pure strategies. The third is in mixed strategies, and has the form $s^* = (s_1^*, s_2^*)$,

where $s_1^* = \left(\frac{2}{3}, \frac{1}{3}\right)$ while $s_2^* = \left(\frac{1}{3}, \frac{2}{3}\right)$.

Now, suppose the players choose the prior probability vectors $p_1 = p_2 = \left(\frac{1}{2}, \frac{1}{2}\right)$. Then, the graph P will have the following form. In the region $0 \leqq t < \frac{1}{3}$, P will have only one branch α, whose points will all have the form $(t, s_1, s_2) = (t, a_1^1, a_2^2)$. In the region $t = \frac{1}{3}$, P will have two branches. One of them, β, will have points of the form $(\frac{1}{3}, a_1^1, s_2)$, with s_2 ranging over all of player 2's mixed strategies. The other, γ, will have points of the form $(\frac{1}{3}, s_1, a_2^2)$, this time with s_1 ranging over all of players 1's mixed strategies. Finally, in the region $\frac{1}{3} < t \leqq 1$, P will have three branches. One of them, δ, will have points of the form (t, a_1^1, a_2^1). Another, ε, will have points of the form (t, a_1^2, a_2^2). The third branch, ζ, will have points of the form $(t, s_1(t), s_2(t))$, where

$$s_1(t) = \left(\frac{3t+1}{6t}, \frac{3t-1}{6t}\right), \qquad s_2(t) = \left(\frac{3t-1}{6t}, \frac{3t+1}{6t}\right).$$

Accordingly, this graph P contains three essentially different feasible paths. Path L_1 consists of the segments α, β, and δ, and leads to the end point $(1, a_1^1, a_2^1)$, suggesting $b^1 = (a_1^1, a_2^1)$ as solution. Path L_2 consists of the segments α, γ, and ε, and leads to the end point $(1, a_1^2, a_2^2)$, suggesting

$b^2 = (a_1^2, a_2^2)$ as solution. Finally, path L_3 consists of the segments α and ζ, and leads to the end point $(1, s_1^*, s_2^*)$, suggesting $s^* = (s_1^*, s_2^*)$ as solution. Thus, there is a feasible path leading to each of the three equilibrium points: the linear tracing procedure is clearly not well defined.

Lemma 18. Whenever the linear tracing procedure is well defined, the distinguished path L of the linear tracing procedure and the limit curve L^{**} of the logarithmic tracing procedure will coincide.

Proof. When the linear procedure is well defined, graph P contains only one feasible path L. But, by Lemma 17, the limit curve L^{**} is always a feasible path contained by P. Therefore, in this case, we must have $L = L^{**}$.

Theorem 4. Whenever the linear tracing procedure is well defined, the solution $s^* = T(\Gamma, p)$ specified by the latter and the solution $s^{**} = T^*(\Gamma, p)$ specified by the logarithmic tracing procedure will coincide.

Theorem 4 directly follows from Lemma 18.

Finally, we state:

Theorem 5. For any given n-tuple p of prior distributions, almost all games Γ will give rise to a well-defined linear tracing procedure.

Since the proof of Theorem 5 is rather long, it will be presented in a separate section (Section 7), in Part II of this paper.

Note. In the light of these results, we can easily extend our model for the solution process (described in Section 4), to the case where the linear tracing procedure is not well defined. All we have to do in this case is to replace the non-existent distinguished path L by the limit curve L^{**}, which always exists (and coincides with the distinguished path L when the latter also exists). Thus, we can assume that, at any given moment of the solution process, every player will base his analysis of the game situation on one particular point $x = (t,s)$ of the limit curve L^{**}, and that this point x will continuously move from the starting point $x^{oo} = (0,s^{oo})$ of L^{**} to its end point $x^{**} = (1,s^{**})$ as the solution process progresses, with full acceptance, by every player, of the equilibrium point s^{**} as solution of the game in the end.

REFERENCES

Rudolf Carnap and Richard C. Jeffrey,
Inductive Logic and Probability. University of California Press, 1971.

Gerard Debreu,
"Economies with a Finite Number of Equilibria ", Econometrica, 38 (1970), 387-392.

John C. Harsanyi,
"Oddness of the Number of Equilibrium Points", International Journal of Game Theory, 2 (1973), 235-250.

John F. Nash,
"Noncooperative Games", Annals of Mathematics, 54 (1951), 286-295.

Leonard J.Savage.
The Foundations of Statistics, John Wiley & Sons, 1954.

Footnotes

1). The present author has been engaged for several years in a joint research project with Professor Reinhard Selten (now of the University of Bielefeld) on the problem of defining a suitable solution concept for n-person noncooperative games. The results of this work will now be published partly as joint papers and partly as individual papers by each of us. This paper is the first in this series. I wish to express my indebtedness to Reinhard Selten for this very stimulating and still ongoing cooperation and, more specifically, for many helpful discussions on the results to be described in this paper.
Thanks are due also to the National Science Foundation for its ongoing support of this research through Grant GS-3222 to the University of California, administered through its Center for Research in Management Science, Berkeley.

I am also grateful to the Institute of Mathematical Economics, the University of Bielefeld in Rheda, Westphalia, for its support during my Visiting Professorship at the Institute.

2). See footnote 1).

3). As is well known, Bayesian decision theorists fall into two groups. Some regard the decision maker's prior probability distribution simply as a datum, and consider it to be beyond the scope of formal decision theory to provide definite rules for the choice of rational priors. Though, as a practical matter, most decision theorists in this group will readily admit that in many specific situations the choice of a rational prior is uniquely determined by symmetry and other similar criteria, they prefer not to incorporate these criteria into their formal theories [e.g., Savage, 1954]. Others feel that formal decision theory should provide definite criteria also for the construction of rational prior distributions, at least in some specified classes of situations [e.g., Carnap and Jeffrey, 1971]. Our own theory of solutions for noncooperative games sides with this second group of Bayesian decision theorists, at least with respect to game situations, because it tries to specify a unique prior probability distribution $p_i = F_i(\Gamma)$ over the strategies of each player i in any given game Γ.

4). In the situation to be described by equation (10), $\sigma_i = \sigma_i(\overline{s_i})$ may also have the nature of an <u>objective</u> probability distribution. Yet, our interest in σ_i will result from the fact that, under our assumptions (see below), σ_i will also express player i's expectations (i.e., his subjective probability distribution).

5). The concepts of inferior and of strictly inferior strategies, as well as the requirement that they should be eliminated from the game, are due to Reinhard Selten.

6). It can never happen that, after repeated elimination of all inferior strategies, any given player should be left without any pure strategies altogether. (This follows from Lemma 5.) But it can happen that he will be left with <u>one</u> pure strategy only. Of course, if this occurs, then this player himself can be formally eliminated from the game, and the remaining players can play the game with one another on the assumption that the player in question will automatically always use the one noninferior strategy available to him.

7). The points of P will be of form x = (t,s), where s is an equilibrium point of game Γ^t. As s is an ewuilibrium point, it will have to satisfy a finite number of inequalities of form (6) and (24), and of equations of form (7) and (23). A given arc P' of P will be an algebraic curve if at every point x of P' the <u>same</u> inequalities and equations are binding; but P will be nonalgebraic in the neighborhood of any point x where the set of binding inequalities and equations changes.

8). To save space, I have not formally defined the concept of a best reply to a probability distribution π_i over all possible pure-strategy (n-1)-tuples c_i^k of the other (n-1) players (which is mathematically the same thing as a best reply to a jointly randomized mixed strategy of these latter players). But the reader can no doubt readily supply the missing definition.