

MIXED AND BEHAVIOR STRATEGIES
IN INFINITE EXTENSIVE GAMES

Robert J. Aumann

Econometric Research Program
Research Memorandum No. 32
October 6, 1961

Princeton University
Econometric Research Program
92-A Nassau Street
Princeton, N. J.

1. Introduction.

We are concerned with infinite extensive games — not necessarily of perfect information — in which there may be a continuum of alternatives at some or all the moves; the games may also have unbounded or infinite play length. Our object is to define the notion of mixed strategy for such games, and to use this definition to prove the appropriate generalization of Kuhn's theorem on optimal behavior strategies in games of perfect recall [K]. Also, our methods give a solution to the conceptual problem raised by McKinsey under the heading "games played over function space" [Mc, pp. 355-357].

By-products are that our proof of Kuhn's theorem makes no use of the rather cumbersome "tree" model for extensive games, that it explicitly uses conditional probabilities (which are implicitly used by Kuhn), and that it explicitly proves that in a game which is of perfect recall for one player, that player can restrict himself to behavior strategies (this also is implicit in Kuhn's proof). Our proof is considerably longer and more complicated than Kuhn's proof, but only because of the problems introduced by the non-denumerably infinite character of the game.

2. Examples.

We give three examples to motivate this study and to illustrate some of the difficulties.

In our first example, there are two players, the "attacker" and the "defender"; for concreteness, one may think of the attacker as a bomber. The attacker starts the play by choosing a course of action (such as a flight course). The defender has some mechanism (such as radar) for determining the course chosen by the attacker, and he decides on his

course of action on the basis of the information he gets from this mechanism. But the mechanism is not perfect; it only gives an apparent attacker course x , the true attacker course being distributed around x according to a known probability distribution (which may vary with x). Thus the defender gets some information about the attacker's course, but not perfect information.

Denote by X the set of all possible apparent attacker courses, i.e. the set of possible information states of the defender. Denote by Y the set of courses of action available to the defender. Clearly a pure defender strategy is a function from X into Y . What about mixed strategies? If X and Y are finite, then there are only finitely many pure strategies, so there is no difficulty about defining mixed strategies. But in many cases the most appropriate model would be one in which X and Y are, say, copies of the unit interval. It is then still possible to define some kinds of mixed strategies; for example, we can mix finitely or denumerably many pure strategies, or we can adopt a fixed continuous distribution over Y regardless of what information we have — i.e., we can mix a continuum of pure strategies, each of which is a constant function from X into Y . But is this the best we can do? Can't we mix a continuum of pure strategies that are not constants?

A mixed strategy can be thought of as a probability distribution, i.e. a measure, on the set of all pure strategies. But before one defines a measure on a non-denumerable space, one must define a measurable structure on the space, i.e. one must define which subsets are measurable. It is by no means clear how this should be done in our case, or even what kind of measurable structure on the pure strategy space should be considered "appropriate" for this purpose.

For our second example we can do no better than quote McKinsey [Mc, p.356]:

"A game has four moves: in the first move P_1 (player 1) chooses a real number x_1 ; in the second move, P_2 , knowing x_1 , chooses a real number y_1 ; in the third move, P_1 , knowing y_1 , but having forgotten x_1 , chooses a real number x_2 ; and in the last move, P_2 , knowing y_1 and x_2 , but not knowing x_1 , chooses a real number y_2 . (The payoff is then some function of the four variables x_1, x_2, y_1 , and y_2 .) A pure strategy for P_1 is now an ordered couple $\{a, f\}$, where a is a real number and f is a function of one real variable (it depends on y_1); and a pure strategy for P_2 is an ordered couple $\{g, h\}$, where g is a function of one real variable (it depends on x_1) and h is a function of two real variables (it depends on y_1 and x_2) ...

"It is clear that the payoff function for a game of the type just described need not necessarily have a saddle point, and hence it is natural to suppose that the players will make use of mixed strategies ... " The difficulties that McKinsey goes on to describe correspond precisely to those we discussed in connection with the first example.

Our third example involves the notion of the supergame of a given game G . This is a game each play of which consists of a number of repeated plays of G ; the payoff to the "superplay" usually is defined as some kind of average of the payoffs to the individual plays. The supergame and related notions¹ have received considerable attention in the literature; this is partly because supergames occur naturally in the applications, and partly because an analysis of a supergame sometimes yields clues as to rational behavior for a single play.²

Supergames are usually analyzed on a step-by-step basis; that is,

¹such as that of stochastic game.

²cf. [Mc, the discussion at top of p. 134]; also [A₁] and [A₂, §10].

it is assumed that each player decides on a strategy for each of the component plays separately. These "local" strategies may or may not depend on the outcomes of the previous component plays, and may be pure or mixed; but the possibility of mixing a number of pure "grand strategies" for the whole supergame is usually ignored. Of course this makes analysis of the supergame much easier.

The supergame may be considered a game in extensive form, a move being a choice of a pure strategy for a component play. Obviously it is a game of perfect recall — at each component play each player remembers what he knew at previous component plays. What we are doing when we limit analysis of the supergame to consideration of mixed strategies for the component plays is that we are considering only behavior strategies in the supergame. Now we lose no generality by this restriction if Kuhn's theorem on behavior strategies in games of perfect recall³ applies, which is the case when the originally given game is finite and is only repeated finitely often. Wolfe [W, p.15] has pointed out that Kuhn's theorem may be extended to games with infinite play length, and it is easily seen that we can also allow a denumerable infinity of alternatives at some (or all) of the moves. The difficulties enter when there may be a continuum of alternatives at some of the moves; in our case this corresponds to a G with a continuum of strategies.

What is the importance of supergames of games with a continuum of strategies? Suppose we wish to consider the supergame of a cooperative game. To analyze this supergame properly, we must formalize the pre-play bargaining for each component play. Such a formalization must involve a

³Kuhn's theorem asserts that in a game of perfect recall each mixed strategy m has an equivalent behavior strategy, i.e., a behavior strategy which yields the same payoff as m (to all players) no matter what the other players do.

continuum of pure strategies for the bargaining session — for example we already have a continuum in the set of correlated strategies that can be offered by a player for the consideration of a coalition that he wishes to form.⁴ Thus a satisfactory analysis of a cooperative supergame cannot proceed without first proving an analogue of Kuhn's theorem for the continuous case. Indeed it was this problem that originally motivated this study.

3. Mixed Strategies.

Let us take a closer look at the first example in the previous section; take X and Y to be copies of the unit interval. We shall need to consider probability distributions on X and Y , and as we remarked in the previous section, this involves defining measurable structures on them. Any such measurable structure should be rich enough to enable us to define the probability of an interval; this means it would have to contain all Borel sets. Let us denote by I the unit interval on which has been imposed a measurable structure consisting of all the Borel sets, and let us once and for all⁵ take X and Y to be copies of I .

Henceforth we will write "m-" for "measurable."

⁴cf. [A₁, §6] or [A₂, §10].

⁵We have adopted the smallest structure that fills our needs. An overly rich structure is self-defeating. For example, if the structure on X consists of all subsets, then the only measures on X are purely atomic (under the continuum hypothesis [S, p.107]); if it consists of all Lebesgue measurable sets, then the only measures are sums of absolutely continuous and purely atomic ones (thus excluding all those with a singular non-atomic component). We therefore see that increasing the set of measurable sets beyond a certain point actually reduces the set of available measures. If we want all intervals to be measurable, the largest set of measures is obtained if we let the structure consist of the Borel sets. (In this connection we remark that there is a confusing misprint in [Mc, p.357, line 7]; here "Lebesgue measurable" should read "Borel.")

Suppose the defender has adopted a strategy f , i.e., a function from X into Y , and that the action of chance and the strategy of the attacker have induced a probability distribution on X . The strategy f , acting on this X -distribution, should induce a distribution on Y . Does it? Suppose $B \subset Y$ is a Borel set. The probability that a member of B is chosen by the defender

$$\begin{aligned}
 &= \text{prob} \{x : f(x) \in B\} \\
 &= \text{prob} \{f^{-1}(B)\} .
 \end{aligned}$$

This expression is meaningless unless $f^{-1}(B)$ is measurable in X ! The same holds for all m -subsets B of Y . In order to have an induced distribution on Y , we want the inverse image under f of a measurable set in Y to be measurable in X . In other words, we want f to be a measurable transformation. So we redefine a pure defender strategy; it is not just any function from X into Y , but an m -transformation.⁶ We denote by Y^X the set of all m -transformations from X into Y .

A mixed strategy, then, is a probability measure on Y^X , the latter having been endowed with an "appropriate" measurable structure. Let us define a function $\phi : Y^X \times X \rightarrow Y$ by $\phi(f, x) = f(x)$. Suppose we again start out with a distribution on X , and suppose that the defender has chosen a mixed strategy; we wish to calculate the induced distribution on Y . For m -sets $B \subset Y$, the probability that the defender chooses a member of B

$$\begin{aligned}
 &= \text{prob} \{(f, x) : f(x) \in B\} \\
 &= \text{prob} \{(f, x) : \phi(f, x) \in B\} \\
 &= \text{prob} \{\phi^{-1}(B)\} .
 \end{aligned}$$

⁶This redefinition of pure strategy is a consequence of the demand that distributions on X induce distributions on Y . Besides being intuitively desirable, this is absolutely necessary for the formal analysis, as the reader will see later. Perhaps the most compelling intuitive argument, though, is that this is needed so that a pair of pure attacker and defender strategies should induce a payoff distribution, e.g., that we should be able to assign a probability to the attacker's payoff being positive.

As before, we conclude that the structure R must be chosen so that ϕ is an m-transformation. But as we have shown elsewhere $[A_3, A_4]$, there is no structure R for which this is so; no structure on Y^X is "appropriate"!

There are two ways out of the blind alley into which we have been led by our apparently sound reasoning. The first is as follows: Recall from the previous section that it is possible to randomize over certain subsets of Y^X , e.g., over denumerable subsets of Y^X and over the set of all constants in Y^X ; on the other hand we have just seen that it is not possible to randomize over all of Y^X . This suggests that instead of trying to define mixed strategies on all of Y^X , we try to characterize those subsets of Y^X which will not lead us to a "blind alley" of the type we have encountered above, and limit ourselves to defining mixed strategies over such subsets.

For each $F \subset Y^X$, define $\phi_F : F \times X \longrightarrow Y$ by $\phi_F(f, x) = f(x)$. A structure R on F is called "admissible" if ϕ_F , considered as a mapping from $(F, R) \times X$ into Y, is an m-transformation. If F has at least one admissible structure, then F itself (as an abstract set without structure) is also called admissible. The admissible subsets of Y^X are precisely those sets whose members can be "mixed," i.e., those sets over which mixed strategies can be defined. Admissible sets have been studied in considerable detail elsewhere $[A_3, A_4]$. Here let us quote only the chief result of those studies as it applies to our situation: A subset of Y^X is admissible if and only if it is contained in some Baire class (which may have arbitrarily high denumerable order).

We are now ready for the definition of mixed defender strategy under the first "way out": A mixed defender strategy is a triple (F, R, μ) , where F is an admissible subset of Y^X , R is an admissible

structure on F , and μ is a measure on the m -space (F, R) . Thus in choosing a mixed strategy, the defender must choose not only a probability measure as usual, but also the m -space on which it is to be defined.

The second "way out" uses a completely different approach. Let us recall the intuitive meaning of a mixed strategy: It is a method for choosing a pure strategy by the use of a random device. Physically, one tosses a coin, and according as to which side comes up chooses a corresponding pure strategy; or, if one wants to randomize over a continuum of pure strategies, one uses a continuous roulette wheel. Mathematically, the random device — the set of sides of the coin or of points on the edge of the roulette wheel — constitutes a probability measure space, sometimes called the sample space; a mixed strategy is a function from this sample space to the set of all pure strategies. In other words what we have here is precisely a random variable whose values are pure strategies. Up to now we have been working with the distribution of this random variable; we now suggest that the use of the random variable itself will simplify matters considerably.

Let us denote by Ω the measure space that results when we impose Lebesgue measure on I . All of our sample spaces will be copies of Ω . The intuitive justification for this is that every "real-life" random device is either "discrete," "continuous," or a combination of the two; that is, the sample space involved must either be finite or denumerable, or it must be a copy of I (with a measure that is not necessarily Lebesgue measure).⁷ All such random devices can be represented by random variables whose sample space is actually a copy of Ω .

⁷Physically, of course, all sample spaces are discrete and even finite; but it is often convenient to use a continuous or a denumerable model.

In our example, therefore, we should define a mixed strategy to be a function from Ω to the space Y^X of all pure strategies. We can expect that not all such functions will be "eligible" as mixed strategies, because of the by now familiar condition that a mixed strategy and a distribution on X must induce a distribution on Y . Fortunately, the appropriate condition is not that the mixed strategy as defined above be a measurable transformation, because this would again involve defining a measurable structure on Y^X . To state the correct condition, we recall that to every function from Ω to Y^X there is a corresponding function from $\Omega \times X \rightarrow Y$; to $f : \Omega \rightarrow Y^X$ there corresponds the function $g : \Omega \times X \rightarrow Y$ defined by $g(\omega, x) = f(\omega)(x)$. The correct condition on a mixed strategy is that this corresponding function be an m -transformation. Thus it is most convenient to redefine a mixed strategy to be an m -transformation from $\Omega \times X$ into Y , and this is the definition we adopt.

As they now stand, both of the above definitions of mixed strategy apply only to the highly simplified situation treated in the first example of the introduction. However, both definitions can be extended without difficulty to more complicated, many-move games. We prefer the second definition, both because it is inherently simpler and more intuitive, and because it bypasses very great difficulties which are encountered in the proof of Kuhn's theorem if the first definition is used.

The relation between the two definitions — similar to that between a random variable and its distribution — will be explored in detail in section 11.

4. Extensive Games.

We first give the formal definitions, then discuss their intuitive meaning and their relation to other definitions in the literature.

An m -space is called standard⁸ either if it is finite or denumerable with the discrete structure (i.e. all subsets are measurable), or if it is isomorphic⁹ with I . Most m -spaces that one "encounters in practice" are standard; for example, any Borel subset of any Euclidean space or of Hilbert space is standard.

Definition:

A game from an individual player's viewpoint, or simply a game, consists of

(i) A (finite or infinite) sequence Y_1, Y_2, \dots of standard m -spaces called action spaces;

(ii) A corresponding sequence X_1, X_2, \dots of standard m -spaces called information spaces;

(iii) A set Z called the set of strategies of the opponents;

(iv) A sequence of functions $g_i : Z \times Y_1 \times \dots \times Y_{i-1} \rightarrow X_i$, called information functions, which for each fixed $z \in Z$, are m -transformations on $Y_1 \times \dots \times Y_{i-1}$ into X_i ;

(v) A standard m -space H called the payoff space;

(vi) A function

$$h : Z \times Y_1 \times Y_2 \times \dots \rightarrow H$$

called the payoff function. The payoff function is assumed to be an m -transformation for each fixed $z \in Z$.

Intuitively, the game is played as follows: First the "opponents,"

⁸This use of the word is due to Mackey [M].

⁹An isomorphism is a one - one correspondence that is measurable in both directions.

including chance, each pick a strategy; the composite of these strategies is a member z of Z . Next, our player is informed of the value of $g_1(z)$; this is a member of X_1 , and represents our player's state of information for his first move. Our player then chooses a member y_1 of Y_1 . Next, he is informed of the value of $g_2(z, y_1)$; this is a member of X_2 , and on the basis of this he must choose a member y_2 of Y_2 . Next, he is informed of the value of $g_3(z, y_1, y_2)$; the game continues in this way. The payoff is determined as a function of the strategy z chosen by the opponents, and the actions y_1, y_2, \dots taken by our player. Usually it will be most convenient to take the payoff space H to be a Euclidean space of dimension equal to the number of players. However this need not always be so,¹⁰ and since we do not use any particular form for H in the sequel, we have left H as general as possible.

Note that up to the present we have not assumed that our player remembers anything on the occasion of a given choice except what he is told by the value of the function g_i . This can be made plausible if we think of the choices of y_1, y_2, \dots as being made by distinct agents of our player, who are not allowed to communicate with each other.

The mappings g and h have been assumed to be m -transformations in the variables y_i for the familiar reason, namely to ensure that distributions on the domain spaces induce a distribution on the range space. This has not been required for the variable z in order to avoid the necessity of defining a measurable structure on the strategy space Z , which leads to difficulties, as we have seen. The results should thus be conceived as holding for each z separately. In a particular case it might be possible to integrate over some components of z (e.g. that

¹⁰For instance, for some purposes it is convenient to consider the payoff to a supergame as being simply the sequence of payoffs to the individual plays, rather than the average (in some sense) of these payoffs.

belonging to chance); this can be done without difficulty after the results have been established for fixed z .

The above definition is a compromise between the normal and extensive forms of a game. The game has been retained in extensive form for our player, but has been normalized for the other players. Even for finite games, most of the important theorems in extensive game theory, such as the theorems on games of perfect information and on games of perfect recall, are best stated for one player at a time; the process of normalizing the game for the other players enables us to focus attention on the single player and thus simplifies the proofs.

Not all finite extensive games in the sense of Kuhn [K] are included in the above definitions; however all games of perfect recall are included, as are all finite extensive games in the sense of von Neumann and Morgenstern [N-M]. The condition for a Kuhn game to be included is that the game can be "serialized," time-wise, for the player in question. For example the game in Figure 1 does not come under our




Figure 1

definition if the information sets A and B belong to the same player, but does come under our definition if they belong to distinct players. Of course the possibility of serialization is not at all equivalent with perfect recall (but the latter implies the former).

Next, we define games of perfect recall.

Definition: A game is said to be of perfect recall if there are sequences of m -transformations

$$u_j^i : X_i \longrightarrow Y_j , \quad j < i$$

and

$$t_j^i : X_i \longrightarrow X_j , \quad j < i$$

such that

$$u_j^i g_i(z, y_1, \dots, y_{i-1}) = y_j$$

and

$$t_j^i g_i(z, y_1, \dots, y_{i-1}) = g_j(z, y_1, \dots, y_{j-1}) .$$

Intuitively, u is the function by which a player remembers what he previously did, and t is the function by which he remembers what he previously knew.

Note that we have given an analytic definition of games of perfect recall which, while retaining complete generality, avoids the cumbersome geometric tree model. This has been made possible by the device of normalizing the game for all but one player.

5. Mixed and Behavior Strategies: The Formal Definitions.

We may assume without loss of generality that the X_i and Y_i are all copies¹¹ of I ; for if one of them is only finite or denumerable, we can always add a continuum of identical copies. The cartesian products $X_i X_i$ and $X_i Y_i$ will be denoted \underline{X} and \underline{Y} respectively, and their members will be denoted $\underline{x} = (x_1, x_2, \dots)$ and $\underline{y} = (y_1, y_2, \dots)$. We remind the reader that the phrase "sample space" means a copy of Ω . Sample spaces will be denoted by Ω , $\underline{\Omega}$, Ω_i , etc.; the measures on them by λ , $\underline{\lambda}$, λ_i , etc., respectively.

¹¹For the definitions of I and Ω see section 3.

In this and the following four sections, the word "subset," when applied to an m -space, will always mean "measurable subset," and $B \subset Y$, when Y is an m -space, will mean "B is an m -subset of Y ."

Definition:

A mixed strategy is a sequence $\underline{m} = (m_1, m_2, \dots)$ of m -transformations $m_i : \Omega \times X_i \longrightarrow Y_i$, where Ω is a fixed sample space.

A behavior strategy is a sequence $\underline{b} = (b_1, b_2, \dots)$ of m -transformations $b_i : \Omega_i \times X_i \longrightarrow Y_i$, where $\{\Omega_1, \Omega_2, \dots\}$ is a sequence of sample spaces.¹²

With a behavior strategy \underline{b} we may associate a mixed strategy $\underline{m} = \underline{m}^b$, with sample space $\underline{\Omega} = \Omega_1 \times \Omega_2 \times \dots$, defined by $m_i^b(\underline{\omega}, x) = b_i(\omega_i, x)$, where $\underline{\omega} = (\omega_1, \omega_2, \dots)$. Intuitively, \underline{m}^b may be thought of as having the same "effect" as \underline{b} .

Every triple $(\omega, \underline{m}, z)$ consisting of a member of the sample space, a mixed strategy, and a strategy of the opponents uniquely determines a member $\underline{v}(\omega; \underline{m}, z)$ of \underline{Y} ; $\underline{v} = (v_1, v_2, \dots)$ is defined recursively by

$$v_i = m_i(\omega, g_i(z, v_1, \dots, v_{i-1})).$$

Intuitively, \underline{v} is the sequence of choices that actually occur when the game is played. Furthermore every pair (\underline{m}, z) uniquely determines a distribution (i.e., measure) μ on \underline{Y} ; this is defined for $\underline{B} \subset \underline{Y}$ by

$$\mu(\underline{B}) = \mu(\underline{B}; \underline{m}, z) = \lambda \{ \omega : \underline{v}(\omega; \underline{m}, z) \in \underline{B} \}.$$

Intuitively, μ is the distribution of the random variable $v(\cdot; \underline{m}, z)$. Two mixed strategies are said to be equivalent if they determine the same distribution on \underline{Y} .

^{12:}

See section 10 for a discussion of this definition.

6. Kuhn's Theorem.

A behavior strategy \underline{b} and a mixed strategy \underline{m} are said to be equivalent if \underline{m} and $\underline{m}^{\underline{b}}$ are equivalent. By definition, for every behavior strategy there is an equivalent mixed strategy. The converse is

Kuhn's Theorem. In a game of perfect recall, every mixed strategy has an equivalent behavior strategy.

7. Lemmas for the Proof of Kuhn's Theorem.

Our first lemma deals with a property of conditional probabilities.

Let Y and Y' be copies of I , let Ω be a sample space with measure λ , let $g : \Omega \times Y \rightarrow Y'$, and $v : \Omega \rightarrow Y$ be m -transformations, and let $B' \subset Y'$ and $B \subset Y$. According to a known theorem [H, p. 210, example 5], the conditional probability

$$\text{cond } \lambda(\Gamma \mid v(\omega) = y)$$

can be defined¹³ as a probability measure on the measurable subsets Γ of Ω , for each fixed $y \in Y$. When we use the symbol $\text{cond } \lambda$ and similar symbols, we shall henceforth mean this probability measure.¹⁴

Lemma A. Under the above conditions

$$\begin{aligned} \int_B \text{cond } \lambda(\{\omega : g(\omega, y) \in B'\} \mid v(\omega) = y) d\lambda v^{-1}(y) \\ = \lambda\{\omega : g(\omega, v(\omega)) \in B' \text{ and } v(\omega) \in B\} \dots \end{aligned}$$

Remark The unusual feature of the integral on the left-hand side is that

¹³We trust that our notation for conditional probabilities, though not standard, is sufficiently transparent as to cause no confusion. There are good reasons for using it rather than one of the standard notations.

¹⁴Actually $\text{cond } \lambda$ is defined uniquely only up to a set of y which is of (λv^{-1}) -measure 0. But all our statements will hold for any particular version of $\text{cond } \lambda$, so the particular choice can be made arbitrarily.

the subset of Ω of which the conditional probability is being taken — the set $\{\omega : g(\omega, y) \in B'\}$ — varies with the condition y . If it were not that $\text{cond } \lambda$ is defined essentially uniquely as a probability measure (for example if Ω were not standard), the integral would have no meaning, because $\text{cond } \lambda$ could be assigned an arbitrary value for each y . What the lemma says is that since the condition asserts $v(\omega) = y$, we may substitute $v(\omega)$ for y on the left side of the $|$ sign, and then obtain the correct answer by using the usual theorem about integrals of conditional probabilities.

Proof. Let $C \subset \Omega \times X$ be defined by $C = g^{-1}(B')$. Denoting by C^y the section $\{\omega : (\omega, y) \in C\}$, we obtain that Lemma A is equivalent to

$$(A1) \quad \int_B \text{cond } \lambda (C^y \mid v(\omega) = y) d\lambda v^{-1}(y) \\ = \lambda\{\omega : (\omega, v(\omega)) \in C \text{ and } v(\omega) \in B\}.$$

Both sides of (A1), as functions of C , are measures on $\Omega \times Y$ (since $\text{cond } \lambda$ is a measure on Ω for each y). Hence it is sufficient to prove (A1) when C is a rectangle $\Gamma \times A$ in $\Omega \times Y$. In this case the left side of (A1)

$$= \int_{A \cap B} \text{cond } \lambda(\Gamma \mid v(\omega) = y) d\lambda v^{-1}(y) \\ = \lambda\{\Gamma \cap v^{-1}(A \cap B)\} \\ = \lambda\{\omega : \omega \in \Gamma \text{ and } v(\omega) \in A \text{ and } v(\omega) \in B\} \\ = \lambda\{\omega : (\omega, v(\omega)) \in \Gamma \times A \text{ and } v(\omega) \in B\} \\ = \lambda\{\omega : (\omega, v(\omega)) \in C \text{ and } v(\omega) \in B\}.$$

This demonstrates (A1) when C is a rectangle, and (A1) and therefore also Lemma A follows in the general case.

Now let us return to our game. First we introduce some further notation. We write $\underline{Y}_i = Y_1 \times \dots \times Y_i$. Similarly, for $\underline{y} \in \underline{Y}$, we write $\underline{y}_i = (y_1, \dots, y_i)$. If $B_1 \subset Y_1, B_2 \subset Y_2, \dots$, then we write

$\underline{B}_i = B_1 \times \dots \times B_i$, and $\underline{B} = B_1 \times B_2 \times \dots$. The symbol \underline{B} will always be reserved for a rectangle of this kind.

Let us consider a mixed strategy \underline{m} with sample space Ω , a strategy z of the opponents, and a sequence $\underline{y} \in \underline{Y}$. Then for each $i = 1, 2, \dots$ we may define a sequence $\underline{v}^i = (v_1^i, v_2^i, \dots) = \underline{v}^i(\omega, \underline{y}; \underline{m}, z)$ inductively as follows:

$$v_j^i = \begin{cases} y_j & , \text{ for } j < i \\ m_j(\omega, g_j(z, \underline{v}_{j-1}^i)) & , \text{ for } j \geq i . \end{cases}$$

We have $\underline{v}^1 = \underline{v}$, and $v_i^i = m_i(\omega, g_i(z, \underline{v}_{i-1}^i))$, which is the decision on the i^{th} play if \underline{v}_{i-1}^i has been chosen on the previous plays. Denote (v_1^i, \dots, v_k^i) by \hat{v}_k^i (for $k \geq i$). Note that v_j^i depends only on \underline{v}_{i-1}^i rather than on all of \underline{y} , so we may write $v_j^i(\omega, \underline{v}_{i-1}^i; \underline{m}, z)$ rather than $v_j^i(\omega, \underline{y}; \underline{m}, z)$. As \underline{m} and z will be fixed throughout most of this discussion, we will usually write $v_j^i(\omega, \underline{v}_{i-1}^i)$, and omit explicit mention of \underline{m} and z . The expression $v_j^i(\cdot, \underline{v}_{i-1}^i)^{-1}(B_j)$ means $\{\omega : v_j^i(\omega, \underline{v}_{i-1}^i) \in B_j\}$. For future reference note that

$$(B1) \quad \hat{v}_j^{i+1}(\omega, (\underline{v}_{i-1}^i, v_i^i(\omega, \underline{v}_{i-1}^i))) = v_j^i(\omega, \underline{v}_{i-1}^i) .$$

Next, define a sequence $\lambda, \lambda_{\underline{y}_1}, \lambda_{\underline{y}_2}, \dots$ of measures on Ω as follows:

$$\lambda_{\underline{y}_i}(\Gamma) = \text{cond } \lambda_{\underline{y}_{i-1}}(\Gamma) \mid v_i^i(\omega, \underline{v}_{i-1}^i) = y_i$$

(where of course $\lambda_{\underline{y}_0}$ stands for λ).

Lemma B. Let $B_i \subset Y_i, \dots, B_k \subset Y_k$. Then

$$\begin{aligned} & \int_{B_i} \dots \int_{B_k} d\lambda_{\underline{y}_{k-1}} v_k^k(\cdot, \underline{y}_{k-1})^{-1}(y_k) \dots d\lambda_{\underline{y}_{i-1}} v_i^i(\cdot, \underline{y}_{i-1})^{-1}(y_i) \\ & = \lambda_{\underline{y}_{i-1}} \hat{v}_k^i(\cdot, \underline{y}_{i-1})^{-1}(\hat{B}_k^i) . \end{aligned}$$

Proof. We use reverse induction on i . The start, at $i = k$, is immediate. For the inductive step ($i + 1$ implies i) we have

$$\begin{aligned} & \int_{B_i} \int_{B_{i+1}} \dots \int_{B_k} \\ &= \int_{B_i} \lambda_{\underline{y}_i} \hat{v}_k^{i+1}(\cdot, \underline{y}_i)^{-1}(\hat{B}_k^{i+1}) d\lambda_{\underline{y}_{i-1}} v_i^i(\cdot, \underline{y}_{i-1})^{-1}(y_i) \\ &= \int_{B_i} \text{cond } \lambda_{\underline{y}_{i-1}} \{ (\omega : \hat{v}_k^{i+1}(\omega, (\underline{y}_{i-1}, y_i)) \in \hat{B}_k^{i+1} \} | \\ & \quad v_i^i(\omega, \underline{y}_{i-1}) = y_i \} d\lambda_{\underline{y}_{i-1}} v_i^i(\cdot, \underline{y}_{i-1})^{-1}(y_i). \end{aligned}$$

Applying Lemma A with $\lambda_{\underline{y}_{i-1}}$ instead of λ , B_i instead of B , \underline{y}_i instead of \underline{y} , \hat{B}_k^{i+1} instead of B , $\hat{v}_k^{i+1}(\cdot, (\underline{y}_{i-1}, \cdot))$ instead of g , \hat{v}_k^{i+1} instead of \underline{y} , and $v_i^i(\omega, \underline{y}_{i-1})$ instead of v , we obtain that the last expression above

$$\begin{aligned} &= \lambda_{\underline{y}_{i-1}} \{ \omega : \hat{v}_k^{i+1}(\omega, (\underline{y}_{i-1}, v_i^i(\omega, \underline{y}_{i-1}))) \in \hat{B}_k^{i+1} \\ & \quad \text{and } v_i^i(\omega, \underline{y}_{i-1}) \in B_i \}, \end{aligned}$$

and from (B1) we deduce that this

$$= \lambda_{\underline{y}_{i-1}} \{ \omega : \hat{v}_k^i(\omega, \underline{y}_{i-1}) \in \hat{B}_k^i \}.$$

This completes the induction.

Corollary C. Let $B_1 \subset Y_1, \dots, B_k \subset Y_k$. Then

$$\begin{aligned} & \int_{B_1} \dots \int_{B_k} d\lambda_{\underline{y}_{k-1}} v_k^k(\cdot, \underline{y}_{k-1})^{-1}(y_k) \dots d\lambda_{\underline{y}_1}^{-1}(y_1) \\ &= \lambda_{\underline{y}_k}^{-1}(B_k). \end{aligned}$$

Corollary D. Let f be an m -transformation from \underline{Y}_k to the real numbers. Then

$$\begin{aligned} & \int_{B_1} \dots \int_{B_k} f(\underline{y}_k) d\lambda_{\underline{y}_{k-1}} v_k^k(\cdot, \underline{y}_{k-1})^{-1}(y_k) \dots d\lambda_{\underline{y}_1}^{-1}(y_1) \\ &= \int_{B_k} f(\underline{y}_k) d\lambda_{\underline{y}_k}^{-1}(y_k). \end{aligned}$$

Proof. If f is the characteristic function of a rectangular parallel-
opiped in \mathbb{Y}_k , this follows from Corollary C. The general case follows by
the usual methods.

8. Further Lemmas.

The object of this section is to prove that a family of distribu-
tions can be "inverted" to yield a family of random variables.

Lemma E. Let f be a non-decreasing upper semi-continuous function¹⁵

on I such that $f(0) \geq 0$ and $f(1) = 1$. For $0 \leq y \leq 1$ define

$$f^{-1}(y) = \begin{cases} \sup \{x : f(x) \leq y\}, & \text{if the set in curly brackets} \\ & \text{is non-empty} \\ 0 & , \text{ if it is empty.} \end{cases}$$

Then

- (1) f^{-1} is non-decreasing
- (2) f^{-1} is upper-semi-continuous
- (3) $f^{-1}(0) \geq 0$, $f^{-1}(1) = 1$
- (4) $(f^{-1})^{-1} = f$.

Proof. For the proof, note that we can restate the theorem as follows:

Let f be a non-decreasing upper-semi-continuous function on $[-1, 1]$ such
that $f(x) = x$ for $x < 0$ and $f(1) = 1$. For $-1 \leq y \leq 1$ define

$$f^{-1}(y) = \sup \{x : f(x) \leq y\}.$$

Then

- (1) f^{-1} is non-decreasing
- (2) f^{-1} is upper-semi-continuous
- (3) $f^{-1}(1) = 1$, $f^{-1}(x) = x$ for $x < 0$
- (4) $(f^{-1})^{-1} = f$.

We now proceed to prove the lemma in this restated form.

¹⁵i.e. $f(x) = \lim_{y \rightarrow x} \sup f(y) = \lim_{y \rightarrow x^{++}} f(y)$.

(1) is obvious.

(2) First we have, by (1), that

$$\lim_{z \rightarrow y^+} f^{-1}(z) \geq \lim_{z \rightarrow y^+} f^{-1}(y) = f^{-1}(y).$$

To prove the opposite inequality, let $f^{-1}(y) = \eta$. Suppose

$$\lim_{z \rightarrow y^+} f^{-1}(z) > \eta, \text{ say}$$

$$\lim_{z \rightarrow y^+} f^{-1}(z) = \eta + \epsilon.$$

Then for each $z > y$, $f^{-1}(z) \geq \eta + \epsilon$, and hence $\sup \{x : f(x) \leq z\} \geq \eta + \epsilon$ for all such z . Hence it follows that for all $z > y$, there is an $x \geq \eta + \epsilon$ such that $f(x) \leq z$; in particular, since f is non-decreasing, it follows that for all $z > y$, $f(\eta + \epsilon) \leq z$. But from this it follows at once that $f(\eta + \epsilon) \leq y$, which contradicts $\sup \{x : f(x) \leq y\} = \eta < \eta + \epsilon$. Hence

$$\lim_{z \rightarrow y^+} f^{-1}(z) \leq f^{-1}(y),$$

and the proof of (2) is complete.

(3) is obvious.

(4) Let $f^{-1} = g$. Then

$$g^{-1}(x) = \sup \{y : g(y) \leq x\}.$$

Suppose $g^{-1}(x) > f(x)$. Then

$$\sup \{y : g(y) \leq x\} > f(x).$$

Hence $\exists y > f(x)$ such that $g(y) \leq x$.

Hence $\exists y > f(x)$ such that $\sup \{z : f(z) \leq y\} \leq x$.

Now

$$\begin{aligned} \sup \{z : f(z) \leq y\} \leq x &\iff (f(z) \leq y \implies z \leq x) \\ &\iff (z > x \implies f(z) > y) \\ &\iff (\forall z > x)(f(z) > y). \end{aligned}$$

Hence

$$(\exists y > f(x))(\forall z > x)(f(z) > y).$$

But since $f(x) > y$ for all $z > x$, it follows from upper semi-continuity that $f(x) \geq y$; this contradicts $y > f(x)$. So we have proved $g^{-1}(x) \leq f(x)$.

To prove the opposite inequality, suppose

$$g^{-1}(x) < f(x).$$

Then $\sup \{y : g(y) \leq x\} < f(x)$, say $\sup \{y : g(y) \leq x\} = f(x) - \epsilon$.

Hence $\forall y (g(y) \leq x \implies y \leq f(x) - \epsilon)$.

Hence $\forall y (\sup \{z : f(z) \leq y\} \leq x \implies y \leq f(x) - \epsilon)$.

Hence $\forall y (\forall z (f(z) \leq y \implies z \leq x) \implies y \leq f(x) - \epsilon)$.

Hence

$$(E1) \quad \forall y (\forall z (z > x \implies f(z) > y) \implies y \leq f(x) - \epsilon).$$

Since f is non-decreasing, we have $z > x \implies f(z) \geq f(x) > f(x) - \frac{\epsilon}{2}$.

So if we set $y = f(x) - \frac{\epsilon}{2}$, then $(\forall z)(z > x \implies f(z) > y)$, but clearly not $y \leq f(x) - \epsilon$. Hence we have a y that contradicts (E1). This completes the proof of (4) and hence of Lemma E.

Let X and Y be copies of I , and let \mathcal{B} be the σ -ring of m -sets in Y . Let $\beta : X \times \mathcal{B} \rightarrow \Omega$ be a function which is measurable in X for each fixed $B \in \mathcal{B}$ and a probability in \mathcal{B} for each fixed $x \in X$.

Lemma F. Under the above conditions, there is a family of random variables whose distributions are given by β ; more precisely, there is an m -transformation $b : X \times \Omega \rightarrow Y$ such that

$$\lambda\{\omega : b(x, \omega) \in B\} = \beta(x, B)$$

for each $x \in X$ and $B \in \mathcal{B}$.

Proof. For $y \in Y$, define

$$\pi(x, y) = \beta(x, [0, y]).$$

Write $\pi_x = \pi(x, \cdot)$; π_x is a non-decreasing upper-semi-continuous function of y , so by Lemma E it has a well-defined inverse, which we denote b_x ;

set $b(x, \omega) = b_x(\omega)$.

Lemma Fl. $b(\cdot, \omega)$ is Borel measurable in x for each fixed ω .

Proof. For $B \in \mathcal{B}$ we must show that $\{x : b(x, \omega) \in B\}$ is measurable in X . It is sufficient to show this when B is of the form $[0, y_0)$.

Now

$$\begin{aligned} & \{x : b(x, \omega) \in [0, y_0)\} \\ &= \{x : \sup \{y : \pi(x, y) \leq \omega\} < y_0\} \\ &= \{x : \exists \text{ rational } r < y_0 \text{ such that } \pi(x, r) > \omega\} \\ &= \bigcup_{r < y_0} \{x : \pi(x, r) > \omega\} \\ &= \bigcup_{r < y_0} \{x : P(x, [0, r)) > \omega\} \\ &= \text{union of Borel sets} = \text{a Borel set.} \end{aligned}$$

This completes the proof of Lemma Fl.

Next we show that b is measurable in the two variables simultaneously. It is sufficient to prove that sets of the form $[y_0, 1]$ have measurable inverse images. Indeed,

$$\begin{aligned} b^{-1}[y_0, 1] &= \{(x, \omega) : b(x, \omega) \geq y_0\} \\ &= \{(x, \omega) : (\forall \text{ rational } s), (s > \omega \implies b(x, s) \geq y_0)\} \\ &\quad (\text{because of upper semi-continuity of } b) \\ &= \bigcap_s \{(x, \omega) : (b(x, s) \geq y_0) \text{ or } (s \leq \omega)\} \\ &= \bigcap_s (\{(x, \omega) : b(x, s) \geq y_0\} \cup \{(x, \omega) : (s \leq \omega)\}) \\ &= \bigcap_s (\{x : b(x, s) \geq y_0\} \times \Omega \cup X \times [s, 1]) \end{aligned}$$

and this is Borel measurable in $X \times \Omega$ (by Lemma Fl).

Finally, we show that

$$\lambda\{\omega : b(x, \omega) \in B\} = \beta(x, B).$$

It is sufficient to demonstrate this when B is of the form $[0, y]$. Then

$$\begin{aligned} \lambda\{\omega : b(x, \omega) \in [0, y]\} &= \lambda\{\omega : b(x, \omega) \leq y\} \\ &= \sup \{\omega : b_x(\omega) \leq y\} = b_x^{-1}(y). \end{aligned}$$

But $b_x = \pi_x^{-1}$; so $b_x^{-1} = \pi_x$ (by Lemma E).

Hence

$$\begin{aligned} \lambda\{\omega : b(x, \omega) \in [0, y]\} &= b_x^{-1}(y) = \pi_x(y) \\ &= \pi(x, y) = \beta(x, [0, y]) , \end{aligned}$$

and the proof of Lemma F is complete.

9. Proof of Kuhn's Theorem.

Fix \underline{m} ; we wish to find an equivalent behavior strategy, which we will call \underline{b} . We first define the distributions β_i of the random variables $b_i(\cdot, x)$, and then only the random variables themselves.

For $B \subset Y_i$ and $x \in X_i$, define

$$\begin{aligned} \beta_i(B, x) &= \text{cond } \lambda (\{\omega : m_i(\omega, x) \in B\} | m_{i-1}(\omega, t_{i-1}^i(x)) = \dots \\ &= u_{i-1}^i(x) | \dots | m_1(\omega, t_1^i(x)) = u_1^i(x)) . \end{aligned}$$

The expression on the right is to be interpreted as an iterated conditional probability, similar to the definition of $\lambda_{\underline{y}_i}$. To underscore the similarity, note that

$$(K1) \quad \beta_i(B, g_i(z, \underline{y}_{i-1})) = \lambda_{\underline{y}_{i-1}} v_i^i(\cdot, \underline{y}_{i-1}; \underline{m}, z)^{-1}(B) .$$

According to Lemma F we can find b_i so that the $\beta_i(\cdot, x)$ are the distributions of the $b_i(\cdot, x)$, i.e. such that

$$(K2) \quad \lambda_i\{\omega_i : b_i(\omega_i, x) \in B\} = \beta_i(B, x) .$$

Let $B_1 \subset Y_1, B_2 \subset Y_2, \dots$; for each n write $B_n^* = B_n \times Y_{n+1} \times \dots$. To show that \underline{m} and \underline{b} are equivalent, it is only necessary to show that

$$(K3) \quad \mu(B_n^*; \underline{m}, z) = \mu(B_n^*; \underline{m}^b, z)$$

for every $z \in Z$ and every n and arbitrary choice of the B_i .

Let us write $\underline{\omega}_n$ for $(\omega_1, \dots, \omega_n)$. We first note that $\mu(B_n^*; \underline{m}^b, z)$ depends only on $\underline{\omega}_n$ rather than on all of $\underline{\omega}$. In fact, if we define $w_n(\underline{\omega}_n)$ recursively by

$$w_n(\underline{\omega}_n) = b_n(\omega_n, g_n(z, w_{n-1}(\underline{\omega}_{n-1}))) ,$$

then

$$(K4) \quad w_n(\underline{\omega}_n) = v_n(\underline{\omega}; \underline{m}^b, z) .$$

Henceforth we will use v for \underline{m} exclusively (unless we explicitly indicate otherwise); thus $\underline{v}(\omega)$ will mean $\underline{v}(\omega; \underline{m}, z)$, and similarly for v_j^i , etc.

The proof of (K3) is by induction on n ; the induction is easily started (at $n = 1$). For the inductive step (n implies $n + 1$) note that

$$\lambda_{\underline{n-n}} w_{\underline{n-n}}^{-1}(\underline{B}_n) = \mu(\underline{B}_n^*; \underline{m}^b, z)$$

because of (K4), and

$$\lambda \underline{v}_n^{-1}(\underline{B}_n) = \mu(\underline{B}_n^*; \underline{m}, z) .$$

Since by induction hypothesis the two right sides are equal, it follows that the left sides also are; but since this holds for all $\underline{B}_n \subset \underline{Y}_n$, it follows that

$$(K5) \quad \lambda_{\underline{n-n}} w_{\underline{n-n}}^{-1} = \lambda \underline{v}_n^{-1}$$

as measures on \underline{Y}_n . Next, we have

$$\begin{aligned} & \mu(\underline{B}_{n+1}^*; \underline{m}^b, z) \\ &= \lambda \{ \underline{\omega} : w_{n+1}(\underline{\omega}_{n+1}) \in \underline{B}_{n+1} \} \\ &= \lambda_{n+1} \{ \underline{\omega}_{n+1} : w_{n+1}(\underline{\omega}_{n+1}) \in \underline{B}_{n+1}, w_n(\underline{\omega}_n) \in \underline{B}_n \} \\ &= \lambda_{n+1} \{ \underline{\omega}_{n+1} : b_{n+1}(\underline{\omega}_{n+1}, g_{n+1}(z, w_n(\underline{\omega}_n))) \in \underline{B}_{n+1}, w_n(\underline{\omega}_n) \in \underline{B}_n \} \\ &= \int_{\underline{w}_n^{-1}(\underline{B}_n)} \lambda_{n+1} \{ \omega_{n+1} : b_{n+1}(\omega_{n+1}, g_{n+1}(z, w_n(\omega_n))) \in \underline{B}_{n+1} \} d\lambda_{\underline{n-n}}(\underline{\omega}_n) \\ &= \int_{\underline{B}_n} \beta_{n+1}(\underline{B}_{n+1}, g_{n+1}(z, \underline{y}_n)) d\lambda_{\underline{n-n}}^{-1}(\underline{y}_n) , \end{aligned}$$

(because of (K2) and the change of variables $\underline{y}_n = w_n(\underline{\omega}_n)$)

$$= \int_{\underline{B}_n} \lambda_{\underline{y}_n} v_{n+1}^{n+1}(\cdot, \underline{y}_n)^{-1}(\underline{B}_{n+1}) d\lambda_{\underline{n-n}}^{-1}(\underline{y}_n)$$

(because of (K1) and (K5))

$$= \int_{B_1} \dots \int_{B_n} \lambda_{\underline{y}_n} v_{n+1}^{n+1}(\cdot, \underline{y}_n)^{-1}(B_{n+1}) d\lambda_{\underline{y}_{n-1}} v_n^n(\cdot, \underline{y}_{n-1})^{-1}(y_n) \dots d\lambda_1^{-1}(y_1)$$

(because of Corollary D)

$$= \int_{B_1} \dots \int_{B_n} \int_{B_{n+1}} d\lambda_{\underline{y}_n} v_{n+1}^{n+1}(\cdot, \underline{y}_n)^{-1}(y_{n+1}) \dots d\lambda_1^{-1}(y_1)$$

$$= \lambda_{\underline{y}_{n+1}}^{-1}(B_{n+1})$$

(because of Corollary C)

$$= \mu(B_{n+1}^* ; \underline{m}, z) .$$

This completes the proof of Kuhn's theorem.

10. Remarks on the Definition of Behavior Strategy.

Regarding the definition of behavior strategy given in Section 4, it might be objected that the intuitive idea of behavior strategy demands that the functions $b_i(\cdot, x)$ be mutually independent random variables for distinct x , even when we are dealing with a single i . To demand this, though, would mean that we must have a non-denumerable number of mutually independent random variables on the same sample space, and this is in fact impossible (when the phrase "sample space" is used in our restricted sense, which corresponds to the intuitive idea of a random device). For example, suppose we wish to construct a non-denumerable number of random variables on the same sample space Ω , representing independent Bernoulli trials with probability $\frac{1}{2}$ of success. If we associate 1 with "success" and 0 with "failure," this means that we must have non-denumerably many characteristic functions of Borel subsets of Ω of measure $\frac{1}{2}$, and these characteristic functions must be mutually independent. Let the subsets of Ω that are involved be denoted S_x , where x runs through the non-denumerable index set X . We have $\lambda(S_x) = \frac{1}{2}$, and from the

independence assumption we obtain $\lambda(S_x \cap S_y) = \frac{1}{4}$ for distinct x and y . Now let us call two Borel subsets of Ω equivalent, if they differ only by a set of measure 0; denote the equivalence class of S by $\{S\}$. Next, let us construct the metric space whose points are equivalence classes of subsets of Ω , and where distance is defined by

$$\rho(\{S\}, \{T\}) = \lambda(S - T) + \lambda(T - S).$$

Then for distinct $x, y \in X$, we have $\rho(\{S_x\}, \{S_y\}) = \frac{1}{2}$. This means that there is a non-denumerable set of points in our metric space such that the distance between any pair of points in the set is $\frac{1}{2}$. It follows that the metric space cannot be separable; but it is known that it is separable [H, p.168]. So our assumption about the existence of non-denumerably many random variables on Ω representing independent Bernoulli trials is false. Similarly (but in a somewhat more complicated way) it can be proved that there cannot be any non-denumerable set of independent random variables on the same sample space.

It may seem that this makes any genuine analogue of Kuhn's theorem in the continuous case impossible. What we have shown is that in a game of perfect recall, one may restrict oneself to deciding on strategies for each of the stages separately, rather than deciding on a grand strategy at the beginning of the game. However, one decides at each stage before receiving the information for that stage, rather than afterwards, and in this our theorem apparently is weaker than Kuhn's.

The difference, however, is illusory, and there is no real loss of strength in our theorem. We have seen that the $b_i(\cdot, x)$ must necessarily be correlated as x ranges over X_i , simply because of the cardinality of X_i . However, this correlation is entirely irrelevant to the game, and cannot affect the payoff in any way. In fact, the payoff distribution depends only on the distributions of the individual $b_i(\cdot, x)$,

and not on any of the joint distributions (this follows from section 9).

In other words, the $b_i(\cdot, x)$ are correlated (for fixed i and varying x) not because this correlation is necessary to mimic the effect of the given mixed strategy \underline{m} ; in fact, if this were necessary (as it may be when the game is not of perfect recall), this kind of in-stage correlation could not accomplish it,¹⁶ and we would have to resort to interstage correlation. The correlation is rather in the way of being an irrelevant mathematical accident.

Yet another way of saying this is that as long as they have the proper distributions, the $b_i(\cdot, x)$ can be chosen in any way we please, without any regard to each other, except that in the end b_i must be simultaneously measurable in both ω and x . Though the $b_i(\cdot, x)$ must be correlated, what form the correlation takes is of no concern to us.

Finally, we venture to say that the main point of Kuhn's theorem lies in the possibility of removing interstage rather than in-stage correlation from mixed strategies in games of perfect recall; we say this because in finite games, in-stage correlation can always be removed from strategies from which inter-stage correlation has been removed, even when a game is not of perfect recall (as we remarked in the next to the last paragraph). Of course in the set-up of [K] there is no distinction between stages; the above remark only makes sense when one introduces stages as we have done.

11. The Range of a Mixed Strategy.

How large a set of pure strategies can be mixed by a mixed strategy?

¹⁶cf. the italicized statement above.

In Section 3 we gave two definitions of "mixed strategy": The "distribution" definition, which we subsequently abandoned, and which we will henceforth call a mixed₁ strategy; and the "random variable" definition, which we subsequently adopted, and which we henceforth call a mixed₂ strategy. For mixed₁ strategies, the above question was answered in Section 3, in the case in which the game has only one stage. In that case the pure strategies coincide with the m -transformations from the (single) information space X to the (single) action space Y , both X and Y being copies of I . The answer that we gave is that a set F of pure strategies constitutes the set of pure "ingredients" of some mixed strategy if and only if F is a subset of some Baire class. Here we wish to examine the same question for mixed₂ strategies. For simplicity we again assume that there is only one stage, and use the same notations and assumptions as above. The results we will obtain may be easily extended by the reader to the general case.

In keeping with the one-stage restriction, we may denote a mixed₂ strategy by $m : \Omega \times X \longrightarrow Y$. A mixed₂ strategy m and a member ω of the sample space determine a pure strategy m_ω , defined by $m_\omega(x) = m(\omega, x)$; m_ω is a "section" of m , and therefore indeed a pure strategy (i.e., an m -transformation). We define the range of m to be the set of all m_ω as ω ranges over Ω ; it is the set of pure strategies that are mixed by the mixed₂ strategy m , or alternatively the set of pure strategies that can occur when m is played. By a range we mean a subset of Y^X that is the range of some mixed₂ strategy.

Range Theorem. Every range is a subset of some Baire class; every Baire class is a subset of some range.

This theorem does not give a complete characterization of ranges, similar to the complete characterization of admissible sets that was

obtained for mixed₁ strategies. For example, I do not even know whether every Baire class is a range, though I suspect that it is; on the other hand, it is highly likely that there exist subsets of Baire classes that are not ranges. What the range theorem does do is give an "order of magnitude" characterization, answering the question with which we began this section: a range can be as large as a Baire class of arbitrarily high order, but no larger.

Proof of the Range Theorem. For the proof of the first part, it suffices to show that every range F is admissible, because it then follows that it must be a subset of a Baire class [A_3 or A_4 , Theorem D], as we have remarked above. Let m be a mixed₂ strategy, and F its range. For every $f \in F$ choose one member ω of Ω , such that $m_\omega = f$; let Ω' be the subspace of Ω obtained in this way, with the subspace structure (a set is measurable if and only if it is the intersection of Ω' with an m -set in Ω). Let m' be the restriction of m to $\Omega' \times X$. Now the restriction of an m -transformation to a subspace is still an m -transformation; hence if we give $\Omega' \times X$ the subspace structure (i.e. as a subspace of $\Omega \times X$), then m' will be an m -transformation. But it is easily verified that the subspace structure on $\Omega' \times X$ is the same as the product structure; hence m' is an m -transformation also when $\Omega' \times X$ has the product structure.

Ω' and F are in one-one correspondence under the correspondence $\omega \longleftrightarrow m_\omega$. Let us impose on F the structure corresponding to that of Ω' ; then Ω' and F are isomorphic. Hence $\Omega' \times X$ and $F \times X$ are also isomorphic. Let us denote the isomorphism by $\zeta : F \times X \longrightarrow \Omega' \times X$; we have $\zeta(m_\omega, x) = (\omega, x)$, where on the right side ω is uniquely defined because of the definition of Ω' .

Now $\phi_F(m_\omega, x) = m_\omega(x) = m(\omega, x) = m'(\omega, x) = m'\zeta(m_\omega, x)$; thus $\phi_F = m'\zeta$.
 But both m' and ζ are m -transformations, and therefore ϕ_F also is an m -transformation. Therefore F is admissible, and the proof of the first part is complete.

To prove the second part of the range theorem, let us define a transfinite sequence $\{F_\alpha\}$, where α ranges over all denumerable ordinals, inductively as follows: F_0 is the set of all continuous functions from X to Y ; F_α is the set of all functions that are pointwise upper limits of sequences of functions in $\bigcup_{\beta < \alpha} F_\beta$. We will prove by induction on α that F_α is a range; since the α' -th Baire class is clearly a subset of F_α , this will complete our proof. The induction is started by

Lemma G . F_0 is a range.

Proof. This is accomplished most easily if we use a result stated in $[A_3]$ and proved in $[A_4]$. Let us consider the uniform convergence topology on F_0 , and let R_0 be the σ -ring of m -sets generated by this topology; our result says that R_0 is admissible.¹⁷ Now F_0 , in the uniform convergence topology, is compact and separable (in the topological sense); hence it is a compact subspace of the Hilbert cube, and therefore in particular (F_0, R_0) is a Borel subspace of the Hilbert cube. But the Hilbert cube is a copy of I ; hence by a known theorem $[M, \text{Corollary 1, p.139}]$, (F_0, R_0) is also a copy of I . To show that F_0 is a range it is now only necessary to let $\psi : \Omega \longrightarrow F_0$ be a Borel isomorphism and to define $m_\omega = \psi(\omega)$. This completes the proof of the lemma.

For the inductive step, let α be a finite or denumerable ordinal, and suppose it has been shown that F_β is a range for all $\beta < \alpha$.

¹⁷In fact, it says more, namely that it is a so-called "natural" admissible structure; cf. $[A_3, A_4]$. We do not need this concept here.

Let $\Omega_1, \Omega_2, \dots$ be a sequence of copies of Ω , and let m_1, m_2, \dots be a sequence of mixed₂ strategies such that each F_α with $\alpha < \beta$ is the range of infinitely many of the m_i . Now let $\underline{\Omega} = \Omega_1 \times \Omega_2 \times \dots$. Define $\underline{m} : \underline{\Omega} \times X \rightarrow Y$ by $\underline{m}(\underline{\omega}, x) = \limsup_{i \rightarrow \infty} m_i(\omega_i, x)$. From the fact that the m_i are m -transformations, it follows easily that \underline{m} is an m -transformation; furthermore it may be seen that the range of \underline{m} is exactly F_α . This completes the proof of the range theorem.

We conclude with a discussion of the relation between the two definitions of mixed strategy. Suppose we are given a mixed₂ strategy m . We may try to define a mixed₁ strategy (F_m, R_m, μ_m) as follows: F_m is the range of m . To define the structure R_m of F_m , we use the function $\psi : \Omega \rightarrow F_m$ given by $\psi(\omega) = m_\omega$; we define R_m to be the identification structure on F_m , i.e. $G \subset F_m$ is measurable if and only if $\psi^{-1}(G)$ is measurable in Ω . Finally, μ is defined by $\mu(G) = \lambda\psi^{-1}(G)$. The idea is that if m is thought of as a "random variable," then (F_m, R_m, μ_m) is the natural candidate for the "distribution" of m .

The trouble is that I cannot establish any correspondence between the random variables and the distributions. On the one hand, if m is a mixed₂ strategy, then I cannot prove that R_m is an admissible structure¹⁸ on F_m ; this means that (F_m, R_m, μ_m) may not be a mixed₁ strategy. On the other hand if a mixed₁ strategy (F, R, μ) is given, it may not be possible¹⁹ to find a mixed₂ strategy m such that $(F, R, \mu) = (F_m, R_m, \mu_m)$. In other words, for all I know, some random variables may have no distributions, and some distributions no corresponding random variables.

¹⁸Though according to the range theorem F_m is admissible, i.e. has admissible structures.

¹⁹I have no example for this, but am fairly convinced that one exists.

The second of these two "paradoxes" could have been expected: while allowing (F, R) to be an arbitrary m -space, we restricted the m -structure of Ω to be a copy of I . The first of the two "paradoxes" is more startling, and we would like to discuss some of the reasons for it. To show that R_m is admissible, we would have to show that ϕ_{F_m} is an m -transformation. Let $\rho : X \rightarrow X$ be the identity, and define $\psi \times \rho : \Omega \times X \rightarrow F_m \times X$ by $(\psi \times \rho)(\omega, x) = (\psi(\omega), x)$. It is easily verified that $\phi_{F_m}(\psi \times \rho) = m$, and hence $(\psi \times \rho)^{-1} \phi_{F_m}^{-1} = m^{-1}$. Now let B be an m -subset of Y . Then $m^{-1}(B)$ is an m -subset of $\Omega \times X$, hence $(\psi \times \rho)^{-1} \phi_{F_m}^{-1}(B)$ also is. We know that (F_m, R_m) is an identification space of Ω under the identification map ψ ; if we only knew that $(F_m, R_m) \times X$ is an identification space of $\Omega \times X$ under $\psi \times \rho$, we would be finished, for it would then follow that $\phi_{F_m}^{-1}(B)$ is measurable, which is exactly what we need. The proposition that "if ψ is an identification map and ρ an identity map, then $\psi \times \rho$ is also an identification map" is intuitively very compelling, but unfortunately we have not succeeded in proving it. Let us call this proposition "the identification space hypothesis"; only the following special cases are known to me:

Mackey's Result²⁰ The identification space hypothesis holds if the domains and images of both ψ and ρ are analytic²¹ m -spaces.

Ernest's Result²² The identification space hypothesis holds if ψ carries m -sets onto m -sets.

²⁰Private correspondence with Professor G. W. Mackey.

²¹i.e., isomorphic with analytic subspaces of I ; cf. [M].

²²Private correspondence with Professor J. Ernest.

We summarize some of the problems left open by this section:

- 1) Characterize ranges of mixed₂ strategies.
- 2) In particular, is every Baire class a range?
- 3) Characterize the mixed₁ strategies that correspond to mixed₂ strategies.
- 4) Prove or disprove the identification space hypothesis.
- 5) Prove or disprove the proposition that for every mixed₂ strategy m , (F_m, R_m, μ_m) is a mixed₁ strategy.

BIBLIOGRAPHY

- [A₁] R. J. Aumann, Acceptable points in general cooperative n-person games, Contributions to the Theory of Games IV, Annals of Math. Study 40, Princeton University Press, 1959, pp. 287-324.
- [A₂] ———, The core of a cooperative game without side payments, Trans. Amer. Math. Soc. 98 (1961), pp. 539-552.
- [A₃] ———, Spaces of measurable transformations, Bul. Amer. Math. Soc. 66 (1960), pp. 301-304.
- [A₄] ———, Borel structures for function spaces, to appear in Ill. Jour. Math., December 1961.
- [H] P. R. Halmos, Measure Theory, Van Nostrand, 1950.
- [K] H. W. Kuhn, Extensive games and the problem of information, Contributions to the Theory of Games II, Annals of Math. Study 28, Princeton University Press, 1953, pp. 245-266.
- [N-M] J. von Neumann and O. Morgenstern, Theory of games and economic behavior, Princeton University Press, 1944, Third Edition, 1953.
- [M] G. W. Mackey, Borel structures in groups and their duals, Trans. Amer. Math. Soc. 85 (1957), pp. 283-295.
- [Mc] J. C. C. McKinsey, Introduction to the theory of games, McGraw-Hill, 1952.
- [S] W. Sierpinski, Hypothèse du continu, Chelsea, 1956.
- [W] P. Wolfe, editor, Report of an informal conference on recent developments in the theory of games, Princeton University, 1955 (mimeographed).