CHAPTER IV

GAMES AGAINST NATURE*

by

John Milnor

PRINCETON UNIVERSITY

1. INTRODUCTION

The object of this paper will be to study games of the following type. A matrix (a_{ij}) is given in which a player must choose a row. A column will be chosen by "Nature", a fictitious player having no known objective and no known strategy. The payoff to the player will then be given by the entry in that particular row and column. This entry should represent a numerical utility in the sense of von Neumann and Morgenstern. (See [3] or [1].)

It will be shown that several known criteria for playing such games can be characterized by simple axioms. An axiomatic procedure will also be used to criticise these criteria, and to study the possibilities for other criteria.

(Our basic assumption that the player has absolutely no information about Nature may seem too restrictive. However such no-information games may be used as a normal form for a wider class of games in which certain types of partial information are allowed. For example if the information consists of bounds for the probabilities of the various states of Nature, then by considering only those mixed strategies for Nature which satisfy these bounds, we construct a new game having no information. Unfortunately in practice partial information often occurs in vague, non-mathematical forms which are difficult to handle.)

The following criteria have been suggested for such games against Nature.

Laplace. If the probabilities of the different possible states of Nature are unknown, we should assume that they are all equal.

*The preparation of this paper was sponsored in part by the RAND Corporation. The author was a National Science Foundation fellow during 1952-53.

Thus if the player chooses the i-th row his expectation is given by the average $(a_{i1}+\cdots+a_{in})/n$, and he should choose a row for which this average is maximized.

Wald [4] (Minimax principle). If the player chooses the i-th row then his payoff will certainly be at least $\operatorname*{Min}_j a_{ij}$. The safest possible course of action is therefore to choose a row for which $\operatorname*{Min}_j a_{ij}$ is maximized. This corresponds to the pessimistic hypothesis of expecting the worst.

If mixed strategies for the player are also allowed, then this criterion should be formulated as follows. Choose a probability mixture $(\xi_1,\cdots,\xi_m)$ of the rows so that the quantity $\operatorname*{Min}_j (\xi_1 a_{1j}+\cdots+\xi_m a_{mj})$ is maximized. In other words play as if Nature were the opposing player in a zero sum game.

Hurwicz[1]. Select a constant $0 \leqq \alpha \leqq 1$ which measures the player's optimism. For each row [or probability mixture of rows] let a denote the smallest component and A the largest. Choose a row [or probability mixture of rows] for which $\alpha A + (1-\alpha)a$ is maximized. For $\alpha = 0$ this reduces to the Wald criterion.

Savage [2] (Minimax Regret). Define the (negative) regret matrix (r_{ij}) by $r_{ij} = a_{ij} - \operatorname*{Max}_k a_{kj}$. Thus r_{ij} measures the difference between the payoff which actually is obtained and the payoff which could have been obtained if the true state of Nature had been known. Now apply the Wald criterion to the matrix (r_{ij}). That is choose a row [or mixture of rows] for which $\operatorname*{Min}_j r_{ij}$ [or $\operatorname*{Min}_j (\xi_1 r_{1j}+\cdots+\xi_m r_{mj})$] is maximized.

These four criteria are certainly different. This is illustrated by the following example, where the preferred row under each criterion is indicated.

$$\begin{pmatrix} 2 & 2 & 0 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 4 & 0 & 0 \\ 1 & 3 & 0 & 0 \end{pmatrix} \begin{matrix} \text{Laplace} \\ \text{Wald} \\ \text{Hurwicz (for } \alpha > 1/4) \\ \text{Savage} \end{matrix}$$

2. AXIOMATIC CHARACTERIZATION OF CRITERIA

In this section we will consider criteria which assign to each matrix (a_{ij}) a preference relation $\gtrsim$ between pairs of rows[2] of the matrix. It will be shown that each of the four criteria of 1 is characterized by certain of the following axioms. The first five axioms are compatible with all four criteria.

1. Ordering. The relation $\gtrsim$ is a complete ordering of the rows. That is it is a transitive relation, such that for any two rows r, r' either $r \gtrsim r'$ or $r' \gtrsim r$.

2. Symmetry. This ordering is independent of the numbering of the rows and columns.
(Thus we are not considering situations where there is any reason to expect one state of Nature more than another.)

3. Strong domination. If each component of r is greater than the corresponding component of r', then $r > r'$ (shorthand for: $r \gtrsim r'$ but not $r' \gtrsim r$).

4. Continuity. If the matrices $a_{ij}^{(k)}$ converge to a_{ij}, and if $r^{(k)} > r_1^{(k)}$ for each k, then the limit rows r and r_1 satisfy $r \gtrsim r_1$.

5. Linearity. The ordering relation is not changed if the matrix (a_{ij}) is replaced by (a'_{ij}) where $a'_{ij} = \lambda a_{ij} + \mu$, $\lambda > 0$.

The following four axioms serve to distinguish between the four criteria.

6. Row adjunction. The ordering between the old rows is not changed by the adjunction of a new row.

7. Column linearity. The ordering is not changed if a constant is added to a column.

(This can be interpreted as an assertion that Nature has no prejudices for or against the player. It also asserts that the utility is linear, not only with respect to known probabilities, but also with respect to unknown probabilities of the type under consideration.)

8. Column duplication. The ordering is not changed if a new column, identical with some old column, is adjoined to the matrix. (Thus we are only interested in what states of Nature are possible, and not in how often each state may have been counted in the formation of the matrix.)

9. Convexity. If row r is equal to the average $\frac{1}{2}(r' + r'')$ of two equivalent rows, then $r \gtrsim r'$.
(Two rows are equivalent, $r' \sim r''$, if $r' \gtrsim r''$ and $r'' \gtrsim r'$. This axiom asserts that the player is not prejudiced against randomizing. If two rows are equally favorable, then he does not mind tossing a coin to decide between them.)

Finally we will need a modified form of axiom 6 which is compatible with all four criteria.

10. Special row adjunction. The ordering between the old rows is not changed by the adjunction of a new row, providing that no component of this new row is greater than the corresponding components of all old rows.

The principal results of this section are all incorporated in the following diagram, which describes the relations between the ten axioms and the four criteria. The symbol "X" indicates that the corresponding axiom and criterion are compatible. Each criterion is characterized by those axioms which are marked "⊠".

	Laplace	Wald	Hurwicz	Savage
1. Ordering	⊠	⊠	⊠	⊠
2. Symmetry	⊠	⊠	⊠	⊠
3. Str. Domination	⊠	⊠	⊠	⊠
4. Continuity	X	⊠	⊠	⊠
5. Linearity	X	X	⊠	X
6. Row adjunction	⊠	⊠	⊠	
7. Col. linearity	⊠			⊠
8. Col. duplication		⊠	⊠	⊠
9. Convexity	X	⊠		⊠
10. Special Row adj.	X	X	X	⊠

Diagram 1. X = compatibility.
Each criterion is characterized by axioms marked ⊠

Theorem 1. The Laplace criterion is compatible with all of these axioms other than axiom 8; the Wald criterion with all but axiom 7; the Hurwicz criterion with all but 7 and 9; the Savage criterion with all but 6.

The proofs are all completely trivial. Perhaps the following two examples are of interest. In the first matrix the Hurwicz criterion (for $\alpha > 0$) is not compatible with axiom 9 (convexity). In the second pair the Savage criterion is not compatible with axiom 6 (row adjunction).

$$\begin{pmatrix} 2 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 2 & 0 \end{pmatrix} \qquad \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix} \longrightarrow \begin{pmatrix} 2 & 1 \\ 0 & 2 \\ -7 & 4 \end{pmatrix}$$

Theorem 2. The Laplace criterion is characterized by axioms 1,2,3,6,7.

It is first necessary to prove the following.

Lemma 1. Assuming axioms 1,2,6 (ordering, symmetry, row adjunction) two rows which differ only in the order of their components are equivalent.

Adjoin a sequence of intermediate rows so that two consecutive rows differ only by a permutation of two components. The result now follows by an application of the symmetry axiom to each pair of consecutive rows.

The proof of theorem 2 follows. Suppose that the average of the components of r equals the average of the components of r'. Alternately perform the following two operations on the matrix:

a) Permute the elements of r and r' so that they are in order of increasing size. (Permissible by lemma 1, and axiom 6).

b) Subtract from each column the component in r or the component in r', whichever is smaller. (Permissible by axiom 7). After a finite number of steps, all of the components of r and r' will be zero. It follows that $r \sim r'$.

Now using axioms 3 and 6 it follows that $r > r'$ whenever the average of the elements of r is greater than the average of the elements of r'. Thus the criterion is that of Laplace.

Theorem 3. The Wald criterion is characterized by axioms 1,2,3,4,6,8,9.

Two lemmas are first necessary.

Lemma 2. Assuming axioms 3 and 4 (domination and continuity), if each component of r is greater than or equal to the corresponding component of r', then $r \gtrsim r'$.

The proof is clear.[3]

Lemma 3. Assuming axioms 1,2,3,4,6,8, two rows which have the same minimum element and the same maximum element are equivalent.

Let $(a_1, \cdots, a_n)$ be any row having the minimum component a and the maximum component A. From lemmas 1 and 2 it follows that

$$(a, \cdots, a, A) \lesssim (a_1, \cdots, a_n) \lesssim (A, \cdots, A, a).$$

But $(a, \cdots, a, A)$ is equivalent to $(A, \cdots, A, a)$ since the matrix $\begin{pmatrix} a \cdots a & A \\ A \cdots A & a \end{pmatrix}$ can be obtained from the symmetrical matrix $\begin{pmatrix} a & A \\ A & a \end{pmatrix}$ by column duplication. Therefore any two rows having minimum element a and maximum element A are equivalent.

Proof of theorem 3. By lemma 3 it is sufficient to consider pairs (a,A) with $a \leqq A$ in place of rows. Applying the convexity axiom (9) to the matrix

$$\begin{pmatrix} a & \frac{1}{2}(a+A) & \frac{1}{2}(a+A) \\ a & a & A \\ a & A & a \end{pmatrix}$$

we have $(a,A) \lesssim (a, \frac{1}{2}(a+A))$. By repeated application of this rule, together with the continuity axiom, we have $(a,A) \lesssim (a,a)$, hence $(a,A) \sim (a,a)$. It follows easily that the criterion is that of Wald.

Theorem 4. The Hurwicz criteria are characterized by axioms 1,2,3,4,5,6,8.

Again it suffices to consider pairs (a,A) with $a \leqq A$. Let α be the supremum of all numbers α' such that

$$(\alpha', \alpha') \lesssim (0,1).$$

By the domination axiom it follows that $0 \leqq \alpha \leqq 1$. By continuity it follows that $(\alpha, \alpha) \sim (0,1)$. By linearity

$$(\alpha A + (1-\alpha)a, \alpha A + (1-\alpha)a) \sim (a,A),$$

whenever $a < A$. It follows easily that the given criterion is just that criterion of Hurwicz which corresponds to the parameter value α.

Theorem 5. The Savage criterion is characterized by axioms 1,2,3,4,7,8,9,10.

A matrix will be called normalized if it contains a row r_0 consisting entirely of zeros, and if it contains no positive components. Any given matrix can be normalized by first subtracting the maximum element from each column, and then adjoining the row r_0. By axioms 7 and 10 these operations do not change the ordering relation between the old rows. In a normalized matrix we are free, by axiom 10, to adjoin any row which contains no positive elements and to delete any row other than r_0. The proof is now completely parallel to the proof of theorem 3. It is only necessary to require that all matrices considered be normalized.

3. CRITICISM OF THE CRITERIA

There is one fundamental principle which has not yet been mentioned: that of domination (or admissibility). One strategy is said to dominate another if it is just as good in all states of Nature and definitely better in at least one. It is natural to require that the following axiom be satisfied.

3'. If r dominates r' then $r > r'$.

This axiom is not compatible with the criteria of Wald, Hurwicz, and Savage. Each of these criteria could be modified in a trivial way[4] so as to satisfy 3', but the result would violate the equally fundamental axiom of continuity. This difficulty is illustrated by the following two examples.

Example 1. Consider the family of matrices

$$\begin{pmatrix} 2 & 1 & 1 \\ 2 & 2 & k \end{pmatrix}$$

where $0 \leqq k \leqq 1$. Mixed strategies are to be allowed. In the case $k = 1$ the second row dominates the first. It is therefore natural to expect that the second row should be chosen exclusively for $k = 1$, and should be chosen with high probability for k close to 1. But according to the Wald and Hurwicz criteria ($\alpha < 1$) the first row should be chosen whenever $k < 1$. (Compare diagram 2). In this example the Savage criterion has the expected behavior, but in the following, more complicated, example the Savage criterion is also unsatisfactory.

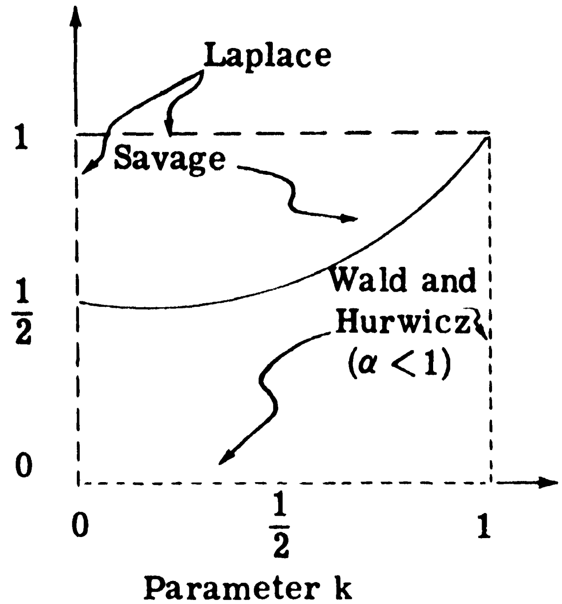

Diagram 2. Probability of choosing second row (Example 1.)

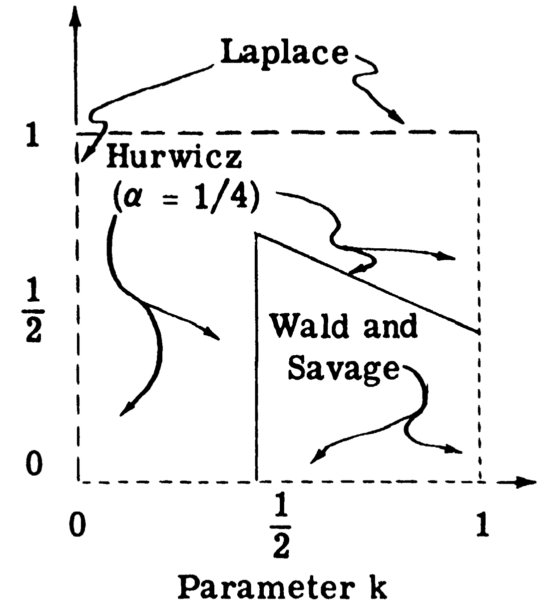

Diagram 3. Probability of choosing last two rows (Example 2.)

Example 2. Consider the matrices

$$\begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & k & 0 \\ 1 & 1 & 0 & k \end{pmatrix}$$

where $0 \leq k \leq 1$. For $k = 1$ the first two rows are dominated, yet according to the Wald and Savage criteria these two rows should be chosen exclusively whenever $k < 1$. (Compare diagram 3). Only the Laplace criterion gives a satisfactory solution in this example.

The Laplace criterion has been successful under all of the tests which have been made of it, with the single exception of axiom 8 (column duplication). It appears that, if we are willing to sacrifice this axiom, then the Laplace criterion is definitely the best. However in many applications it is desirable to preserve axiom 8. This is particularly true in cases where there is no clear and natural separation of the possible states of Nature into a finite number of distinct alternatives.

Thus all of the criteria under consideration seem unsatisfactory in that they fail to satisfy certain rather basic axioms.

S, corresponding to the n possible states of Nature. The negative regret is defined by $r_j(s,S) = p_j(s) - \underset{s' \in S}{\text{Max}}\, p_j(s')$. The Savage choice set C(S) consists of the set of all strategies $s \in S$ for which $\underset{j}{\text{Min}}\, r_j(s,S)$ attains its maximum M. Instead we will consider the set $C_{\mathcal{E}}(S)$ consisting of all $s \in S$ such that

$$\underset{j}{\text{Min}}\, r_j(s,S) \geqq M - \mathcal{E} .$$

The required criteria are now constructed as follows. Choose as parameters an infinite sequence of positive numbers $\mathcal{E}_1, \mathcal{E}_2, \cdots$ which converge to zero. Define the sets $S_0 \supset S_1 \supset \cdots$ by $S_0 = S$, $S_i = C_{\mathcal{E}_i}(S_{i-1})$. As choice set $C_{(\mathcal{E}_1, \mathcal{E}_2, \ldots)}(S)$ we take the intersection of the S_i.

The axioms I through VIII may now be verified. The proofs will not be carried out, since they are rather involved (at least for domination and continuity). In any case these criteria are probably too difficult computationally to be of practical interest.

A further interesting property which is possessed by these criteria is the following. The n payoff functions are all constant on the choice set. Thus any two elements in the choice set are completely equivalent.

It is interesting to ask if there exist any simple, computable criteria which satisfy all of the preceding conditions.

FOOTNOTES

1. Suggested by L. Hurwicz in an unpublished paper.

2. For simplicity, only pure strategies for the player are considered in this section. However the results can easily be generalized to the (more natural) case where mixed strategies are allowed.

3. Lemma 2 suggests the following criterion. Define $r \gtrsim r'$ if and only if each component of r is $\geqq$ the corresponding component of r'. It may be shown that this criterion satisfies all axioms except 1, and is characterized by 2,3,4,6,7,8, together with the transitivity portion of 1.

4. Let r be preferred to r' (in the modified sense) if either $r > r'$ in the old sense (of Wald, Hurwicz, or Savage) or r dominates r'.

BIBLIOGRAPHY

1. Herstein, I. N. and Milnor, J., An Axiomatic Approach to Measurable Utility, Econometrica, 21 (1953), 291-297.

2. Savage, L. J., The Theory of Statistical Decision, Journal of the American Statistical Association, 48 (1947), 238-248.

3. von Neumann, J., and Morgenstern, O., Theory of Games and Economic Behavior, Princeton, 1944.

4. Wald, A., Statistical Decision Functions, New York, 1950.

CHAPTER V

NOTE ON SOME PROPOSED DECISION CRITERIA*

by

Roy Radner and Jacob Marschak

COWLES COMMISSION FOR RESEARCH IN ECONOMICS

1. SUMMARY

The purpose of this paper is to apply two currently advocated statistical decision procedures to a simple problem and show that they result in solutions that have certain undesirable properties. Each of the two procedures is a generalization or interpretation of the minimax principle. The problem consists of a game in which an individual observes and bets on the outcomes of tosses of a coin with constant but unknown probability of falling heads.

2. INTRODUCTION

2.1. The Rational Decision-Maker. In this discussion we shall consider an individual decision-maker who is rational in the following sense: if he can specify a set of "states of nature" such that for a given state n and a given strategy he knows the probability distribution of outcomes, then he will always

(1) choose some admissible strategy (when possible),[1]

(2) choose the strategy so as to maximize his expected utility, if he knows the true state of nature.

Let U(s,n) be the expected utility when n is the true state of nature and the individual uses strategy s. A strategy s_o is admissible if there is no other strategy s such that

$$U(s,n) \geqq U(s_o,n), \quad \text{for all } n, \quad \text{and}$$

*This research was carried out under contract with the Office of Naval Research. We are indebted to E. Lehmann for several valuable suggestions and discussions.

$$U(s,n_0) > U(s_0,n_0)\,, \quad \text{for some } n_0\,.$$

2.2. The Minimax Principle. If the individual does not know the true state of nature then, in general, the criterion of admissibility will not be sufficient to enable him to choose a strategy; thus some further criteria are needed. Such a criterion is Wald's "minimax principle." (Cf. [8], p. 18) One interpretation of this principle is: "Minimax the negative expected utility," i.e., choose s to achieve

$$\min_s \max_n [-U(s,n)]\,.$$

This interpretation has been attacked by many as too pessimistic (cf. for example [7], p. 63), and it is this undesirable property which has, in part, led to the proposal of two alternative criteria which we will now consider.

2.3. The Hurwicz Criterion. The first of these, which might be considered as comprising a whole class of criteria, including the minimax principle as just stated, is a generalized form of a criterion proposed by L. Hurwicz [3]. In this generalized form it requires that a strategy be chosen which maximizes:

$$H(s) = \phi\,(\operatorname{Sup}_n U(s,n),\ \operatorname{Inf}_n U(s,n)) \tag{1}$$

where ϕ is some fixed monotone increasing function of each of its two arguments. ϕ itself is chosen by the decision-maker and in some sense characterizes his attitude towards uncertainty. A special case (the one actually suggested by Hurwicz) is

$$\phi = \alpha \operatorname{Sup}_n U(s,n) + (1-\alpha) \operatorname{Inf}_n U(s,n) \tag{2}$$

where α is some fixed number between 0 and 1. Here α might be regarded as a degree of optimism (cf. [4], p. 344).

We will present an example in which application of the Hurwicz criterion leads to the conclusion that at most one observation should be taken in a situation in which common sense demands that a large number of observations be taken.

2.4. The Minimax Regret Criterion. A different direction is taken by L. J. Savage, [7], who gives good reason why Wald could not have considered negative expected utility as the appropriate thing to minimax (cf. Wald [8], p. 8). Instead, Savage says the proper interpretation of the "minimax rule" is: "Choose that strategy which minimizes $\operatorname{Sup}_n R(s,n)$ where

$$R(s,n) = \operatorname*{Sup}_{s'} U(s',n) - U(s,n) . \tag{3}$$

We will call R(s,n) the regret function.[2] Chernoff, [2], has criticized this principle because there are cases in which, if the domain S of the player's available strategies is enlarged, a new minimax regret solution is obtained which differs from the old one, yet is contained in the original S. (This is not surprising, since the value of R(s,n) for any pair (s,n) depends upon the domain S. Note that this is not true of H(s).)

It is interesting to note that the idea behind Chernoff's objection has analogues in Nash's treatment of the bargaining problem (postulate 7, p. 159 of [5]) and Arrow's discussion of social welfare functions (Condition 3, p. 27 of [1]). Borrowing Arrow's terminology we shall say that in the kind of cases described above the minimax regret solution is "dependent upon irrelevant alternatives."

3. A NON-SEQUENTIAL GAME

3.1. General Description. Consider the following game: The player observes an odd number (2k + 1) of tosses of a coin with a constant but unknown probability p of falling heads (and $q = 1 - p$ of falling tails), whereupon he makes a bet on the outcome of the next toss, wins one dollar if his prediction is correct and loses one dollar if incorrect. Each toss costs the player c dollars, and he must decide in advance the (odd) number of tosses he will observe before betting. The player is also free not to enter the game at all. This last possible decision we will call the null strategy. Aside from it, any pure strategy of the player consists of a number k, which determines that he will bet after 2k + 1 tosses, and a rule r, which determines for every set of observations (sample) which way he will bet. A mixed strategy is a probability distribution on the set of pure strategies (k,r).

We shall see that for both types of solutions there is an optimal rule r_m which requires that the player bet with the majority of previous tosses. This will be called the maximum likelihood rule. It will be shown that the Hurwicz solution has the property that for any (positive) cost c, and any ϕ, no more than one observation should be taken. In the special case of a linear ϕ [equation (2)], the solution is: if $\alpha \geqq 2c$, bet after one observation, if $\alpha \leqq 2c$ do not play.

The minimax regret solution is of the form: randomize between two adjacent values of k, these values being certain non-increasing functions of c; if c is less than a certain quantity, randomize between one observation and not playing. However, if we modify the game by compelling the player always to use the maximum likelihood rule, the optimal number of observations will be seen to differ from that in the solution of the more general game.

3.2. Hurwicz Solution. For any non-null strategy the expected gain cannot be more than $1 - c(N + 1)$ where N is the expected number of observations, since this is the gain if the prediction is correct with certainty. On the other hand, for any non-null strategy the expected gain for $p = \frac{1}{2}$ is $-c(N + 1)$, hence the minimum cannot be more than $-c(N + 1)$.

If the player uses the maximum likelihood rule, then the expected gain is exactly $1 - c(N + 1)$ for $p = 0$ or 1, while it is never less than $-c(N + 1)$. (If the reader is not immediately convinced of this latter statement he can examine the expected gain function in more detail in the next section.) Moreover, given that a non-null strategy is used, the smallest possible value of N is 1. Hence among non-null strategies both $\operatorname{Sup}_n U(n,s)$ and $\operatorname{Inf}_n U(n,s)$ are maximized by using the maximum likelihood rule and taking one observation, and no matter what the ϕ and c the optimal procedure will have the property that no more than one observation is taken. If ϕ has the linear form of (2) and if s_0 is the strategy which consists of using the maximum likelihood rule after one observation with probability v and not playing with probability $1 - v$, then

$$\begin{aligned} H(s_0) &= \alpha\, v(1 - 2c) + (1 - \alpha)\, v(-2c) \\ &= v(\alpha - 2c)\,. \end{aligned}$$

Thus $H(s_0)$ is maximized by taking v equal to 1 or 0 according as $\alpha \geqq 2c$ or $\alpha \leqq 2c$.

3.3 Minimax Regret Solution. This solution is not so easily obtained as the one imposed by the Hurwicz criterion, and we will only sketch the method of arriving at it.

Let $d(k,r)$ denote a joint probability distribution of k and the rule r and let $d(r|k)$ be the conditional distribution of r given k.

Denote the player's expected money gain using $d(k,r)$, given p, by $U(d(k,r),p)$; the expected money gain using $d(r|k)$, given k and p, by $U_k(d(r|k),p)$; and the null strategy by $k = -1$. Then:

$$(4) \qquad U(d(k,r),p) = \sum_{k=-1}^{\infty} f(k)\, U_k(d(r|k),p)\,.$$

Let $\bar{h}_i$ and $\bar{t}_i$ be the events of getting i heads and i tails respectively with respective probabilities $h_i(p)$ and $t_i(p)$ $(i = 0, \cdots, k+1)$. Any $d(r|k)$ is a rule of the form:

"For given k, if $\bar{h}_i$, bet on heads with the probability η_i and if $\bar{t}_i$ bet on tails with probability τ_i."

The maximum likelihood rule r_m is defined by $\eta_i = \tau_i = 1$. If $d(r|k)$ differs from r_m, it will do so exactly on certain events $\bar{h}_j$ (j in J) and $\bar{t}_\ell$ (ℓ in L). It is easily verified that:

$$(5) \qquad U_k(r_m,p) = (p-q)\sum_i [h_i(p) - t_i(p)] - 2(k+1)c$$

$$(6) \qquad U_k(d(r|k),p) = U_k(r_m,p) + 2(p-q)\Big[\sum_L t_\ell(p)(1-\tau_\ell) - \sum_J h_j(p)(1-\eta_j)\Big].$$

It is not hard to show that we can reject as inadmissible all strategies such that there is some k (with $f(k) \neq 0$), for which J and L are not disjoint. The set of remaining strategies we will call S.

We want that $d(k,r)$ which minimizes the supremum, with respect to p, of the regret:

$$R[d(k,r),p] = \hat{U}(p) - U[d(k,r),p], \text{ where}$$

$$\hat{U}(p) = \sup_{d(k,r)} U[d(k,r),p]\,.$$

$\hat{U}(p)$ is attained, for every p, if the player bets on heads when $p \geqq \frac{1}{2}$, on tails when $p \leqq \frac{1}{2}$ and pays as small a cost as possible (i.e., $k = 0$) provided the resulting expected gain is positive; otherwise it is attained by not playing. Thus:

$$\hat{U}(p) = \max\{|p-q| - 2c,\, 0\}\,.$$

Let $h(k,p)$ and $t(k,p)$ be the probabilities of majorities of heads and tails respectively.

Then for a strategy using the maximum likelihood rule, the regret is:

$$\rho(f,p) = \sum_{k=-1}^{\infty} f(k)\,\rho_k(p) \quad \text{where, for } k \geq 0,$$

$$\rho_k(p) = \begin{cases} 2(q - p)\,h(k,p) + 2kc, & (q - p) \geq 2c \\ 2(k + 1)c - (p - q)(h(k,p) - t(k,p)), & p - q \leq 2c \\ 2(p - q)\,t(k,p) + 2kc, & (p - q) \geq 2c \end{cases}$$

and

$$\rho_{-1}(p) = U(p).$$

The function $\rho(f,p)$ is symmetric in p, for all f, and has a maximum at two points, say p_1 and $q_1 = 1 - p_1$, if $c < c_0$; or at $\frac{1}{2}$, if $c \geqq c_0$, where

$$\frac{1}{2} + c < p_1 < 1$$

and c_0 is defined by

$$c_0 = \sum_{k=0}^{\infty} f(k)(p_1 - q_1)\,t(k,p_1)$$

i.e., c_0 is the cost per observation for which the three relative maxima of $\rho(f,p)$ are equal.

Next, it can be shown that one of the minimax regret strategies uses the maximum likelihood rule. The important step in the proof of this point is the fact that (when $c < c_0$), if the regret for some strategy s at $p = p_1$ is less than $\rho(f,p_1)$ then at $p = q_1$ it is greater than $\rho(f,q_1)$, and vice versa.

It remains now to find the optimal distribution f of k, when the maximum likelihood rule is used.

Although $\rho_k(p)$ is defined only for integral values of k, it is, for every fixed p, analogous to a convex function of k, in that for every integer k (and fixed p):

$$\rho_{k+1}(p) - \rho_k(p) \leqq \rho_{k+2}(p) - \rho_{k+1}(p).$$

It is shown in [6] that in such a case the only admissible strategies (using the maximum likelihood rule) are those such that f(k) is concentrated on at most two consecutive integers. Since such a distribution is determined by its mean we can express the solution by a single number $\hat{k}$, which will be a function of c. The approximate value of this function $\hat{k}(c)$ has been determined numerically for several values of c, and the results are given in Table 1 below.[3]

3.4. Dependence on "Irrelevant Alternatives." We shall now show that in this game the minimax regret solution "depends upon irrelevant alternatives."

Suppose we modify the above game by requiring the player to use the maximum likelihood rule. We proceed to obtain the minimax regret solution for this case.

The expected gain using f(k) is given by (4) and (5). Again the negative of $U_k(r_m,p)$ is convex in k for every p, in the sense described above, and the only admissible f's are those which are zero at all but at most two consecutive values of k. As in Section 3.3, we have obtained the value of the function $\hat{k}(c)$, describing the optimal strategy, for various values of c. The results are given in Table 1.

Cost c	.001	.002	.005	.010	.020	.050
$\hat{k}$ for General Strategy Domain	13.8	8.9	4.3	2.4	1.2	0.2
$\hat{k}$ for maximum likelihood Strategy Domain	9.1	5.4	2.6	1.8	.7	

Table 1. Minimax Regret Solutions for 2 Strategy Domains

We recall that the two games considered differ only in that in the first game the player is free to use strategies which do not incorporate the maximum likelihood rule, while in the second he must use that rule. Nevertheless, in the first game the optimal strategy is shown to use the maximum likelihood rule, but with a different number of trials $2k + 1$.

4. A SEQUENTIAL GAME—THE HURWICZ SOLUTION

It is worthwhile pointing out[4] that the essential feature of the Hurwicz solution in Section 3.2 carries over to a sequential generalization of the first game. That is, if we allow the player to decide when he will make his bet after having seen any number of observations, it remains true that any optimal strategy will not involve taking more than one observation. The proof of this for general ϕ is practically the same as that for the non-sequential game.

FOOTNOTES

1. In our examples there will always be admissible strategies.

2. Savage calls this the "loss function" but economists and others are liable to confuse this with negative income.

3. Computations for this and the following section were made under the direction of J. Templeton and W. Parrish.

4. We are indebted to E. L. Lehmann for doing so to us.

BIBLIOGRAPHY

1. Arrow, K. J., Social Choice and Individual Values, New York, 1951.

2. Chernoff, H., Remarks on Rational Selection of a Decision Function, Cowles Commission Discussion Papers, Statistics Numbers 326 and 326A (mimeographed).

3. Hurwicz, L., A Class of Criteria for Decision Making under Ignorance, Cowles Commission Discussion Paper, Statistics No. 356 (mimeographed).

4. Hurwicz, L., Some Specification Problems and Applications to Econometric Methods (Abstract), Econometrica, 19 (1951), 343-344.

5. Nash, J. F. Jr., The Bargaining Problem, Econometrica, 18 (1950), 155-162.

6. Radner, R., Note on Generalized Convex Functions and the Decision Problem, Cowles Commission Discussion Paper, Statistics No. 362 (mimeographed).

7. Savage, L. J., The Theory of Statistical Decisions, Journal of the American Statistical Association, 48 (1947), 238-248.

8. Wald, A., Statistical Decision Functions, New York, 1950.