Cowles Foundation Paper 96
Reprinted from Thrall, Davis and Coombs , eds.,
Decision Processes, John Wiley, 1954

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_0 is admissible if there is no other strategy s such that

$$U(s,n) \geqq U(s_0,n), \quad \text{for all } n, \quad \text{and}$$

*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.

$$U(s,n_O) > U(s_O,n_O) \ , \quad \text{for some } n_O \ .$$

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:

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

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

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

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

(3) $$R(s,n) = \sup_{s'} U(s',n) - U(s,n) .$$

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:

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

$\hat{U}(p)$ 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 $c < c_o$; or at $\frac{1}{2}$, if $c \geqq c_o$, where

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

and c_o is defined by

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

i.e., c_o 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 $c < c_o$), 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.