Rand
Research Memorandum
RM-1570

THE PREDICTION OF SEQUENCES

by

David Blackwell

I. INTRODUCTION

Suppose a predictor is allowed to observe the successive terms $\epsilon_1, \epsilon_2, \ldots$ of an infinite sequence of 0's and 1's until he decides to stop. If the first two unobserved terms are 1,0 in that order he loses; otherwise he wins. For a given number $p > 0$, does there exist a method of prediction (i.e., a randomized stop rule) for which the probability of winning is at least p for every infinite sequence $x = (\epsilon_1, \epsilon_2, \ldots)$ of 0's and 1's? If so, what is such a method? It is clear that for $p > 3/4$, no method exists, since if the sequence is chosen by coin tossing the average success probability over all sequences is exactly 3/4 for any method. For any $p < 3/4$, methods will be given below, and it will be shown that for $p = 3/4$ no method exists.

The problem described above is an instance of a class of problems, in which the predictor observes as many terms of an infinite sequence $\epsilon_1, \epsilon_2, \ldots$ of 0's and 1's as he pleases, after which he chooses an action from a finite set. He wins an amount depending only on the action chosen and on the first few unobserved terms of the sequence. One problem of this type, which furnished the stimulation for writing this paper, is the "two-move lag bomber-battleship game" solved by Dubins [1] and Isaacs and Karlin [2], in which the sequence $\epsilon_1, \epsilon_2, \cdots$ describes the motion of the battleship, with $\epsilon_1 + \cdots + \epsilon_N$ being its position at time N. The bomber watches the ship as long as he pleases, after which he drops a bomb, which lands two time units later at any position designated by the bomber. Thus the bomber wins if and only if he predicts correctly the sum of the first two unobserved terms.

As will be noted below, it follows from general theorems of Wald [4] and Karlin [3] that problems of this type, considered as games between the sequence chooser and the predictor, have a value, and that the sequence chooser has a good strategy. Some properties of these good strategies will be obtained, and it will be shown that there exists a stationary good strategy: one in which the probability that any specified finite sequence begins at time N is independent of N.

No systematic method exists for solving prediction games, even for the special case of predicting when a given sequence δ is not about to begin. A method of Milnor, which yields the value for certain δ, including all δ of length not exceeding 3, will be described below. It will be shown that, except in the trivial case in which all coordinates of δ are equal, no optimal prediction method exists.

II. DESCRIPTION OF PREDICTION GAMES

Denote by S the set of all finite sequences $s = (\epsilon_1, \cdots, \epsilon_N)$ of 0's and 1's, and by X the set of infinite sequences $x = (\epsilon_1, \epsilon_2, \cdots)$. We shall call a finite sequence $E = (e_1, \cdots, e_m)$ of elements of S a partition if every x has exactly one e_i as an initial segment. Let $E = (e_1, \cdots, e_m)$ be a partition, and let $A = ||A(i,j)||$ be an m x n matrix. Associated with (E,A) is a prediction game, in which the pure strategies for player I are sequences $x \in X$, and in which the pure strategies for player II are pairs $y = (F,g)$, where $F = (f_1, \cdots, f_r)$ is a partition, and g is a function associating with each f_k an integer $j = g(f_k)$, $1 \leq j \leq n$. The payoff to I is

$$M(x,y) = A(i, g(f_k)) ,$$

where i,k are the unique integers such that (f_k, e_i) is an initial segment of x.

A mixed strategy for I is determined by specifying for each $s \in S$ the probability P(s) of the set of all sequences x with s as initial segment. The function P satisfies $P(\theta) = 1$, $P(s) \geq 0$, and $P(s) = P(s,0) + P(s,1)$, where θ is the empty sequence, and any P satisfying these conditions determines a mixed strategy for I. Player II has only a countable set of pure strategies $y_1, y_2, \cdots$, so that a mixed strategy for him is specified by a sequence $Q = (\lambda_1, \lambda_2, \cdots)$, $\lambda_i \geq 0$, $\sum \lambda_i = 1$. Thus the payoff to I when he uses P against Q is

$$M(P,Q) = \sum_i \lambda_i M(P,y_i) ,$$

where

$$M(P,y) = \sum_{k,i} P(f_k, e_i) A(i, g(f_k)) .$$

This game satisfies the hypotheses of theorems of Wald [4] and Karlin [3]:

(a) <u>The strategy spaces are convex</u>: For every P_1, P_2 and λ, $0 < \lambda < 1$, there is a P with $M(P,Q) = \lambda M(P_1,Q) + (1 - \lambda)M(P_2,Q)$ for all Q, and similarly for II. (This is obvious.)

(b) <u>II's strategy space is separable</u>: There is a countable set Q_1, Q_2, $\cdots$ such that for any Q there is a subsequence Q_n' of Q_n such that $M(P,Q_n') \to M(P,Q)$ for all P. (This condition is always satisfied when the set of II's pure strategies is countable.)

(c) I's strategy space is compact: For any sequence P_n, there is a P and a subsequence P_n' of P_n such that $M(P_n',Q) \rightarrow M(P,Q)$ for all Q. (A subsequence of P_n can be selected which converges for each s to some P, and convergence of M follows.)

According to the Wald-Karlin theorem, the game has a value and I has a good strategy: there is a (unique) number v and a P* such that

$$M(P^*,Q) \geq v \text{ for all } Q$$

and

$$\inf_Q \sup_P M(P,Q) = v .$$

Theorem 1. If P is a good strategy and $P(s) > 0$, then P_s, defined by $P_s(t) = P(s,t)/P(s)$, is also a good strategy.

Proof: Let $(s_0,s_1,\cdots,s_t)$ be any partition with $s_0 = s$; let $y_0,\cdots,y_t$ be any strategies for II; and let y* be the strategy for II which consists of waiting until some s_k occurs and then using y_k on the sequence beginning immediately thereafter: If $y_k = (F_k,g_k)$, where $F_k = (f_{k1},\cdots,f_{kr_k})$, then $y^* = (F^*,g^*)$, where $F^* = \{(s_k,f_{kj})\}$ and $g^*(s_k,f_{kj}) = g_k(f_{kj})$.

Then for any optimal P we have

$$M(P,y^*) = \sum_k P(s_k) M(P_{s_k}, y_k) \geq v ,$$

so that

$$\sum_k P(s_k) m_k \geq v ,$$

where

$$m_k = \inf_y M(P_{s_k}, y) .$$

Since every m_k satisfies $m_k \leq v$, we have $m_k = v$ for $P(s_k) > 0$, and every P_{s_k} with $P(s_k) > 0$ is a good strategy.

<u>Theorem 2.</u> P <u>is a good strategy if and only if for every</u> s <u>with</u> P(s) > 0 <u>we have</u>

(1) $$\sum_i P_s(e_i) A(i,j) \geq v \text{ for all } j .$$

<u>Proof</u>: Clearly (1) holds if P is a good strategy since P_s, being also a good strategy, must yield at least v against taking action j immediately. Conversely, for any y and P,

(2) $$M(P,y) = \sum_k P(f_k) \left[\sum_i P_{f_k}(e_i) A(i, g(f_k)) \right] .$$

If (1) holds, each bracketed expression is at least v, so that P is a good strategy.

Corollary: For any P, $\inf_y M(P,y) \geq \inf_{s,j} \sum P_s(e_i)A(i,j)$.

Theorem 3. There is a good strategy P which is stationary, i.e., which satisfies $P(s) = P(1,s) + P(0,s)$ for every s.

Proof: For any P, let TP be defined by $TP(s) = P(1,s) + P(0,s) = P(1)P_1(s) + P(0)P_0(s)$, and let $Q_N = (P+TP+\cdots+T^{N-1}P)/N$. A subsequence of Q_N converges for each s to a limit, say, P^*. Since $TQ_N - Q_N = (T^N P - P)/N$, we have $TP^* = P^*$; thus P^* is stationary. If P is a good strategy, so is TP, being an average of the good strategies P_0 and P_1. Consequently each Q_N, being an average of good strategies, is good; and so is P^*, being a limit of good strategies.

III. THE GAME (1,0)

The game (1,0) described in the introduction is the particular prediction game with E = ((1,0), (0), (1,1)) and

$$A = \begin{Vmatrix} 1 \\ 0 \\ 0 \end{Vmatrix} .$$

We note that if the sequence is chosen by coin tossing, i.e., $P^*(s) = 2^{-N}$ for sequences of length N, equation (2) yields $M(P^*,y) = 1/4$ for all y, so that $v \geq 1/4$. A method of Milnor yields $v \leq 1/4$ as follows. If v is the value and P is a good strategy, then according to Theorem 2 we have $P_s(1,0) > v$ if $P(s) > 0$; i.e.,

$$(3) \qquad P_s(1)P_{s,1}(0) \geq v \quad \text{if} \quad P(s) > 0 .$$

Since $v > 0$, $P(s) > 0$ implies $P(s,1) > 0$. Let $g = \inf P_s(1)$, where the inf is over all s with $P(s) > 0$. From (3) we get $v \leq (1-g)P_s(1)$ and, choosing s_n for which $P_{s_n}(1) \to g$,

$$v \leq g(1-g) .$$

Since the maximum value of $g(1-g)$ is 1/4, we obtain $v \leq 1/4$.

Note that the foregoing method gives no clue as to optimal prediction methods. We show in the next section that no optimal prediction method exists. For any $\epsilon > 0$ we shall now describe an ϵ-optimal prediction method. Associated with every finite sequence $z = (n_1,\cdots,n_T)$ of integers with $n_i > 0$ for $i < T$, $n_T = 0$ is a stop rule as follows: Stop after any s which for some i (a) ends in exactly n_i zeros and (b) contains exactly $i-1$ other zeros. For fixed $\lambda_0 > 0,\cdots,\lambda_N > 0$, $\sum_0^N \lambda_i = 1$, we select a stop rule z by choosing $n_1,n_2,\cdots$ independently with $Pr(n_i = j) = \lambda_j$, with T as the smallest i for which $n_i = 0$. Thus

$$Q(z) = \lambda_{n_1}, \cdots, \lambda_{n_T} .$$

For any x, let us investigate $M(x,Q)$, the probability of stopping just before an occurrence of (1,0) in the sequence x when the stop rule z is selected according to Q. Let E_i be the event: Stopping occurs between the $(i-1)$-st and i-th zeros of x. Further, let F_i be the event E_i followed by the occurrence of (1,0) immediately after stopping. Now

$$\alpha_i = \Pr(F_i | E_i) = \frac{\Pr(n_i = a_i - 1)}{\Pr(n_i \le a_i)} ,$$

where a_i = No. of 1's between the $(i-1)$-st and i-th zeros of x. Thus $\alpha_i = 0$ if $a_i = 0$ or if $a_i > N+1$. For $1 \le a_i \le N$, $\alpha_i = \lambda_{i-1}/(\lambda_0 + \cdots + \lambda_i)$ and for $a_i = N+1$, $\alpha_i = \lambda_N$. Let

$$(4) \quad \max \left(\frac{\lambda_0}{\lambda_0+\lambda_1}, \cdots, \frac{\lambda_{i-1}}{\lambda_0+\cdots+\lambda_i}, \frac{\lambda_{N-1}}{\lambda_0+\cdots+\lambda_N}, \frac{\lambda_N}{\lambda_0+\cdots+\lambda_N}\right) = w .$$

Then

$$M(x,Q) = \sum_i \Pr(F_i) = \sum_i \Pr(E_i)\Pr(F_i|E_i) \le w \sum_i \Pr(E_i) = w .$$

Thus we have reduced the problem to that of finding the smallest w for which a solution of (4) with $\lambda_i > 0$ exists.

For any $w > 1/4$, a solution exists, for a solution of

$$z_i - z_{i-1} = wz_i$$

with initial conditions $z_{-1} = 0$ is

$$z_n = (2 \cos \theta)^n \sin (n + 1)\theta ,$$

where $4w \cos^2 \theta = 1$, $0 < \theta < \pi/2$. Choosing N so that $z_{-1} = 0 < z_1 < \cdots < z_N$, $z_N \geq z_{N+1}$, and defining $\lambda_i = z_i - z_{i-1}$, from $0 \leq i \leq N$ we obtain $\lambda_i = w(\lambda_0 + \cdots + \lambda_{i+1})$ for $0 \leq i < N$, $\lambda_N = wz_{N+1}/z_N \leq w = w(\lambda_0 + \cdots + \lambda_N)$, so that (4) is satisfied.

Thus, for any $w > 1/4$, the associated prediction scheme guarantees a probability of w or less that the stopping point is followed immediately by (1,0), against every sequence x. Operationally, the prediction scheme may be described as follows. For a given $w > 1/4$, define $\lambda_0, \cdots, \lambda_N$ as above. Select an integer n_1 according to the distribution $\{\lambda_i\}$. If $n_1 = 0$, stop initially, i.e., without observing anything. If $n_1 > 0$, wait until either n_1 1's or a 0 occurs. If n_1 1's occur initially, stop. If a 0 occurs before n_1 1's do, select a new n_2 according to $\{\lambda_i\}$ and wait until either n_2 1's or a second 0 occurs. If the former, stop immediately; if the latter, select n_3 according to $\{\lambda_i\}$, etc.

For $\theta = \pi/18$, corresponding to $w = .2577$, the distribution $\{\lambda_i\}$ has been computed ($N = 15$) and turns out to require an absurdly long expected waiting time against many sequences. For instance, we have $\lambda_0 = .000019$, so that, against the sequence $x = (0,0,\cdots)$, the expected number of observed digits is $1/\lambda_0 \sim 53000$. Again, we have $\lambda_1 = .000056$, so that, against the sequence $x = (1,0,1,0,\cdots)$, the expected number of digits is $2/(\lambda_0 + \lambda_1) \sim 27000$; and, of course, when we do stop, with probability w it will be after 0 so that (1,0) does occur immediately thereafter.

It would be desirable to exhibit near-optimal prediction schemes involving less observation. The best prediction scheme requiring at most k observations is near-optimal for large k, but it seems unlikely that any simple description of this scheme can be given.

IV. NONEXISTENCE OF OPTIMAL PREDICTION METHODS

For any sequence $\delta \in S$, we consider the game in which the predictor wins unless the unobserved sequence begins with δ.

Theorem 4. Except for the trivial case in which all digits of δ are alike, there is no optimum prediction method for δ.

Proof: Say v is the value and P is a good strategy for I, and let Q be any strategy for II. We may suppose the initial digit of δ to be 1. Except in the trivial case mentioned above, the following facts are easily verified:

(a) $0 < v < 1$;

(b) $P(s_N) > 0$ for all N, where $s_N = (1,1,\cdots,1)$ of length N ;

(c) $M(x^*,Q) = 0$, where $x^* = (1,1,\cdots)$.

If λ_n = Prob of stopping after n observations using Q against x^*, choose N so that $\sum_{0}^{N} \lambda_1 > 1 - v$. Then $M(x,Q) < v$ against any sequence x with initial segment s_{N+k}, where k = length of δ ; for II always wins against such an x when he stops after N or fewer observations, and N was chosen so that the probability of this already exceeds $1 - v$. However, $\int M(x,Q)dP(x) \geq v$, since P is an optimum strategy for I. Since $M(x,Q) < v$ on a set of sequences of positive P-probability (the set of sequences with s_{N+k} as initial segment), we must have $M(x,Q) > v$ for some x, so that Q cannot be optimal, and the theorem is proved.

REFERENCES

[1] Lester E. Dubins, "A Discrete Evasion Game," Tech. Note No. 2, Institute for Air Weapons Research, Univ. of Chicago, 9 December 1953.

[2] Rufus Isaacs and Samuel Karlin, A Game of Aiming and Evasion, The RAND Corporation, Research Memorandum RM-1316, 6 August 1954.

[3] Samuel Karlin, "Operator Treatment of Minmax Principle," Contributions to the Theory of Games, Annals of Math. Studies 24, Princeton, 1950.

[4] Abraham Wald, Statistical Decision Functions, John Wiley and Sons, New York, 1950.