COLLECTED PAPERS OF KENNETH J. ARROW

The Economics of Information

COLLECTED PAPERS OF KENNETH J. ARROW

The Economics of Information

Basil Blackwell

Preface

The work of Jerzy Neyman and Egon Pearson ("On the Problem of the Most Efficient Tests of Statistical Hypotheses," *Philosophical Transactions of the Royal Society of London*, series A, 231:289–337, 1933) gave an economic cast to the foundations of statistical method. They presented criteria of performance and sought to optimize them. Abraham Wald (*Statistical Decision Functions*, New York: Wiley, 1950) gave a more explicitly economic account. Even earlier, statistical methodology in acceptance inspection and quality control had been based on explicitly economic concepts (producers' and consumers' risks). Statistical method was an example for the acquisition of information. In a world of uncertainty it was no great leap to realize that information is valuable in an economic sense. Nevertheless, it has proved difficult to frame a general theory of information as an economic commodity, because different kinds of information have no common unit that has yet been identified. In different ways the papers in this volume have sought to set out the dimensions of the problem or problems and have proposed approaches in certain specific cases. But a general approach is still elusive.

I should like to thank Mary Ellen Geer for her careful and thorough editing and Michael Barclay and Robert Wood for preparation of the index.

Contents

1 Bayes and Minimax Solutions of Sequential Decision Problems

Abraham Wald developed the idea and methods of testing statistical hypotheses by sequential analysis at the Statistical Research Group, formed to develop statistical methods for use in the national defense in World War II. No doubt, as in other such efforts, many of the fruits were not available for use until after the emergency that called them forth was over. The memorandum (Statistical Research Group, 1945), originally marked "Confidential," was circulated, and some of us at the Weather Division of Air Force headquarters were using it within a few months to test whether or not the long-range weather forecasts produced there were significantly bet-

This chapter was written with D. Blackwell and M. A. Girshick. Reprinted from *Econometrica,* 17 (1949): 213–244. The research for this paper was carried on at the RAND Corporation, a nonprofit research organization under contract with the United States Air Force. It was presented at a joint meeting of the Econometric Society and the Institute of Mathematical Statistics at Madison, Wisconsin, September 9, 1948, under the title "Statistics and the Theory of Games." Many of the results in this paper overlap with those obtained previously by Wald and Wolfowitz (1948), and also with some prior unpublished results of Wald and Wolfowitz, announced by Wald at the meeting of the Institute of Mathematical Statistics at Berkeley, California, June 22, 1948. Sections 3 and 6 of the present paper contain analogues of lemmas 1–4 of Wald and Wolfowitz (1948), though both the statements and the proofs differ because of the generally different approach. The proof that the sequential probability ratio test of a dichotomy minimizes the expected number of observations under either hypothesis, in Section 5 of the present paper, follows from Section 3 in the same way that the proof of the same theorem follows from lemmas 1–8 in Wald and Wolfowitz. The previously mentioned unpublished results of Wald and Wolfowitz include the main result of Section 2 (structure of the optimum sequential procedure for the finite multidecision problem) in the special case of linear cost functions.

ter than chance. (They were not.) I became especially interested in Wald's more general formulations of statistical decision theory, both nonsequential and sequential (Wald, 1947b).

While at the RAND Corporation in the summer of 1948, I worked especially with Meyer A. Girshick. About that time, Girshick had attended a meeting of the Institute of Mathematical Statistics at which Wald and Jacob Wolfowitz presented some new results about the structure of sequential analysis when there were more than two alternative hypotheses. He returned with great excitement and stimulated David Blackwell and myself to join him in attempting to reconstruct the results in a more transparent form; the original presentation was certainly hard to understand, and its underlying logic unclear. The three of us grasped that the essential idea was the repetition of the decision situation at each step, though with varying values of the parameters. Hence the decision rule consisted in specifying regions in the parameter space, the same for all time. This point of view was of course implicit in the studies of Wald (1947a) and of Wald and Wolfowitz (1948) but had not been made central.

An unpleasant episode was connected with this paper: a version was circulated that had inadequate acknowledgment to the work of Wald and Wolfowitz, and they felt that there was a challenge to their priority. The published version presents the relation between the papers fairly.

The paper sets forth explicitly the notion of recursive optimization. It provided me with a model argument to be applied to the determination of optimal inventories in work done jointly with Theodore E. Harris and Jacob Marschak. More important, it helped to suggest to Richard Bellman (1957) the general principle of dynamic programming, which has found so many applications.

The problem of statistical decisions has been formulated by Wald (1947b) as follows. The statistician is required to choose some action a from a class A of possible actions. He incurs a loss $L(u,a)$, a known bounded function of his action a and an unknown state u of Nature. What is the best action for the statistician to take?

If u is a chance variable, not necessarily numerical, with a known a priori distribution, then $\mathscr{E}L(u,a) = R(a)$ is the expected loss from action a, and any action, or randomized mixture of actions, which minimizes $R(a)$ has been called by Wald a *Bayes solution* of the decision problem, corresponding to the given a priori distribution of u.

Now suppose there is a sequence x of chance variables $x_1, x_2, \ldots,$ whose joint distribution is determined by u. Instead of choosing an action immediately, the statistician may decide to select a sample of x's since this will yield partial information about u, enabling him to make a wiser selection of a. There will be a cost $c_N(x)$ of obtaining the sample $x_1, \ldots,$ x_N and, in choosing a sampling procedure, the statistician must balance the expected cost against the expected amount of information to be obtained.

Formally, the possibility of making observations leaves the situation unchanged, except that the class A of possible actions for the statistician has been extended. His action now consists of choosing a sampling procedure T and a decision function D specifying what action a will be taken for each possible result of the experiment. The expected loss is now $R(T,D) = l(T,D) + c(T)$, where $l(T,D)$ is the expected value of $L(u,a)$ for the specified sampling procedure and decision rule, and $c(T)$ is the expected cost of the sampling procedure. A Bayes solution is now a pair (T,D), or randomized mixture of pairs (T,D), for which $R(T,D)$ assumes its minimum value.

The minimizing $T = T^*$ has been implicitly characterized by Wald and may be described by the following rule: At each stage, take another observation if and only if there is some sequential continuation which reduces the expected risk below its present level. The main difficulty here is that various quantities which arise are not obviously measurable: for instance, if the first observation is x_1, we must compare our present risk level, say $w_1(x_1)$, with $z(x_1) = \inf w(x_1,T,D)$, where $w(x_1,T,D)$ is the expected risk for any possible continuation (T,D); we take another observation if and only if $w_1 > z$. It is not a priori clear that z will be a measurable function of x_1, so that the set of points x_1 for which we stop may not be measurable. Actually, z always is measurable, as will be shown.[1]

A characterization of the minimizing $T = T^*$ is obtained for hypotheses involving a finite number of alternatives under the condition of random sampling. It consists of the following. We are given k hypotheses H_i $(i = 1,2, \ldots, k)$ which have an a priori probability g_i of occurring, a risk matrix $W = (w_{ij})$ where w_{ij} represents the loss incurred in choosing H_j when H_i is true, and a function $c(n)$ which represents the cost of taking n observations. It is shown that for each sample size N, there exist k convex regions S_j^* in the

1. The possibility of nonmeasurability is not considered in Wald (1947a) or Wald and Wolfowitz (1948).

$(k-1)$-dimensional simplex spanned by the unit vectors in Euclidean k-space whose boundaries depend on the hypotheses H_i, the risk matrix W, and the cost function $c_N(n) = c(N+n) - c(n)$. These regions have the property that if the vector $\mathbf{g}(N)$ whose components represent the a posteriori probability distribution of the k hypotheses lies in S_j^*, the best procedure is to accept H_j without further experimentation. However, if $\mathbf{g}(N)$ lies in the complement of $\bigcup_{j=1}^{k} S_j^*$, the best procedure is to continue taking observations. At any stage, the decision whether to continue or terminate sampling is uniquely determined by this sequence of k regions, and moreover, this sequence of regions completely characterizes T^*.

A method for determining the boundaries of these convex regions is given for $k = 2$ (dichotomy) when the cost function is linear. It is shown that in this special case, T^* coincides with Wald's sequential probability ratio test.

The minimax solution to multivalued decision problems is considered, and methods are given for obtaining them for dichotomies. It is shown that in general, the minimax strategy for the statistician is pure, except when the hypotheses involve discrete variates. In the latter case, mixed strategies will be the rule rather than the exception. Examples of double dichotomies, binomial dichotomies, and trichotomies are given to illustrate the construction of T^* and the notion of minimax solutions.

It may be remarked that the problem of optimum sequential choice among several actions is closely allied to the economic problem of the rational behavior of an entrepreneur under conditions of uncertainty. At each point in time, the entrepreneur has the choice between entering into some imperfectly liquid commitment and holding part or all of his funds in cash pending the acquisition of additional information, the latter being costly because of the foregone profits.

1. Construction of Bayes Solutions

The Decision Function

We have seen that the statistician must choose a pair (T,D). It turns out that the choice of D is independent of that of T.

LEMMA. *There is a fixed sequence of decision functions D_m such that*

(1-1) $$R(T,D_m) \to \inf R(T,D) = w(T) \quad \text{for all } T.$$

This will be the main result of this section. It follows that the expected loss from a procedure T may be taken as $w(T)$, since this loss may be approximated to arbitrary accuracy by appropriate choice of D_m, and a best sequential procedure T^* of a given class will be one for which $w(T^*) = \inf w(T)$, where the infimum is taken over all procedures T of the class under consideration.

We are considering, then, a chance variable u and a sequence x of chance variables $x_1, x_2, \ldots$. A *sequential procedure* T is a sequence of disjunct sets $S_0, S_1, \ldots, S_N, \ldots$, where S_N depends only on $x_1, \ldots, x_N$ and is the event that the sampling procedure terminates with the sample $x_1, \ldots, x_N$; we require that $\Sigma_{N=0}^{\infty} P(S_N) = 1$. S_0 is the event that we do not sample at all, but take some action immediately; it will have probability either 0 or 1.

A *decision function* D is a sequence of functions d_0, $d_1(x_1), \ldots, d_N(x_1, \ldots, x_N), \ldots$, where each d_N assumes values in A and specifies the action taken when sampling terminates with $x_1, \ldots, x_N$. We admit only decision functions D such that $L[u,d_N(x)]$ is for each N a measurable function.

Proof of Lemma. The loss from (T,D) is $G(u,x,T,D) = L[u,d_N(x)] + c_N(x)$ for $x \in S_N$, and $\mathcal{E}G = R(T,D)$. Here, $c_N(x)$ depends only on $x_1, \ldots, x_N$. Then, denoting by $\mathcal{E}_N$ the conditional expectation given $x_1, \ldots, x_N$, we have $\mathcal{E}_N G = \mathcal{E}_N L[u,d_N(x)] + c_N(x)$ for $x \in S_N$, and

$$R(T,D) = \sum_{N=0}^{\infty} \int_{S_N} \mathcal{E}_N L(u,d_N)\, dP + c(T). \tag{1-2}$$

Now fix N; we shall show that we can choose a sequence of functions $d_{Nm}(x)$, $m = 1, 2, \ldots$, such that

(a) $\mathcal{E}_N L(u,d_{Nm}) \geqq \mathcal{E}_N L(u,d_{N,m+1})$ for all x,

(b) $\mathcal{E}_N L(u,d_N) \geqq r_N$ for all d_N and all x, where

$$r_N(x) = \lim_{m\to\infty} \mathcal{E}_N L(u,d_{Nm}).$$

(c) $r_N \geqq \mathcal{E}_N r_n$ if $n \geqq N$.

First choose a sequence d'_{Nm} such that

$$\mathcal{E}L(u,d'_{Nm}) \to \inf_{d_N} \mathcal{E}L(u,d_N) = r.$$

Now define d_{Nm} inductively as follows: $d_{N1} = d'_{N1}$; $d_{Nm} = d'_{Nm}$ for those values of x such that $\mathscr{E}_N L(u,d'_{Nm}) \leqq \mathscr{E}_N L(u,d_{N,m-1})$, otherwise $d_{Nm} = d_{N,m-1}$. Then certainly (a) holds, so that $\lim_{m\to\infty} \mathscr{E}_N L(u,d_{Nm}) = r_N(x)$ exists. Also $\mathscr{E}_N L(u,d_{Nm}) \leqq \mathscr{E}_N L(u,d'_{Nm})$, so that $\mathscr{E} r_N = r$. Choose any d_N and any $\delta > 0$, and let S be the event $\{\mathscr{E}_N L(u,d_N) < r_N(x) - \delta\}$. Then, defining $d^*_{Nm} = d_N$ on S, $d^*_{Nm} = d_{Nm}$ elsewhere, we have

$$\mathscr{E} L(u,d^*_{Nm}) \leqq \int_S r_N(x)\, dP + \int_{CS} \mathscr{E}_N L(u,d_{Nm})\, dP - \delta P(S),$$

so that $\lim_{m\to\infty} \mathscr{E} L(u,d^*_{Nm}) \leqq r - \delta\, P(S)$, and $P(S) = 0$. This establishes (b). Finally, (c) follows from the fact that every $d_N(x)$ is also a possible $d_n(x)$ if $n > N$. This means that, defining $d^*_n = d_N$, we have $\mathscr{E}_N L(u,d_N) = \mathscr{E}_N[\mathscr{E}_n L(u,d^*_n)] \geqq \mathscr{E}_N r_n$ for all d_N, and consequently (c) holds.

Now define $D_m = \{d_{Nm}\}$. Since $\mathscr{E}_N L(u,d_{Nm})$ decreases with m, (1-2) yields that

$$R(T,D_m) \to \sum_{N=0}^{\infty} \int_{S_N} r_N(x)\, dP + c(T) = w(T),$$

and, using (b), that $R(T,D) \geqq w(T)$ for all D. Thus we have reduced the problem of finding Bayes solutions to the following. We are given a sequence x of chance variables $x_1, x_2, \ldots$, and a sequence of nonnegative expected loss functions $w_0, \ldots$, where $w_N = r_N(x_1, \ldots, x_N) + c_N(x_1, \ldots, x_N)$. c_N is the cost of the first N observations, and r_N is the loss due to incomplete information. With each sequential procedure $T = \{S_N\}$ there is associated a risk $w(T) = \Sigma_N \int_{S_N} w_N(x)\, dP$. How can T be chosen so that $w(T)$ is as small as possible?

The Best Truncated Procedure

Among all sequential procedures not requiring more than N observations, there turns out to be a best, that is, one whose expected risk does not exceed that of any other. Moreover, the procedure can be explicitly described, by induction backwards, in such a way that its measurability is clear. After $N-1$ observations $x_1, \ldots, x_{N-1}$, we compare the present risk w_{N-1} with the conditional expected risk $\mathscr{E}_{N-1} w_N$ if we take the final observation. Thus, by choosing the better course, we can limit our loss to $\alpha_{N-1} = \min(w_{N-1}, \mathscr{E}_{N-1} w_N)$, which may be considered as the attainable risk with the observations $x_1, \ldots, x_{N-1}$. We can then decide, on the basis of $N-2$ observations, whether the $(N-1)$st is worth taking by comparing the present risk, w_{N-2}, with $\mathscr{E}_{N-2}\alpha_{N-1}$, the attainable risk if x_{N-1} is observed. Continuing backwards, we obtain at each stage an expected attainable risk α_k for the observations $x_1, \ldots, x_k$, and a description of how to attain this risk, that is, of when to take another observation. This is formalized in the following theorem.

THEOREM 1. *Let $x_1, \ldots, x_N; w_0, \ldots, w_N$ be any chance variables, $w_i = w_i(x_1, \ldots, x_i)$. Define $\alpha_N = w_N$, $\alpha_j = \min(w_j, \mathscr{E}_j\alpha_{j+1})$ for $j < N$, $S_j = \{w_i > \alpha_i$ for $i < j$, $w_j = \alpha_j\}$. Then for any disjoint events $B_0, \ldots, B_N$, B_i depending only on $x_1, \ldots, x_i$, $\sum_{i=0}^{N} P(B_i) = 1$, we have*

$$\sum_{j=0}^{N} \int_{S_j} w_j \, dP \leq \sum_{i=0}^{N} \int_{B_i} w_i \, dP.$$

Proof. We shall show that, for fixed i and any $(x_1, \ldots, x_i)$-set A,

$$\sum_{j \geq i} \int_{AS_j} \alpha_j \, dP = \sum_{j \geq i} \int_{AS_j} \alpha_i \, dP, \tag{1-3}$$

and that, for fixed j, and any disjoint sets $A_j, \ldots, A_N$ with A_i depending only on $x_1, \ldots, x_i$ and $\bigcup_{i=j+1}^{N} A_i$ depending only on $x_1, \ldots, x_j$,

$$\sum_{i \geq j} \int_{A_i} \alpha_j \, dP \leq \sum_{i \geq j} \int_{A_i} \alpha_i \, dP. \tag{1-4}$$

Choosing $A = B_i$ in (1-3) and summing over i, choosing $A_i = B_i S_j$ in (1-4) and summing over j, and adding the results yields

$$\sum_{i,j=0}^{N} \int_{B_i S_j} \alpha_j \, dP \leq \sum_{i,j=0}^{N} \int_{B_i S_j} \alpha_i \, dP. \tag{1-5}$$

Now on S_j, $\alpha_j = w_j$, and always $\alpha_i \leq w_i$. Making these replacements in (1-5) yields the theorem.

We now prove (1-3) and (1-4). The relationship (1-3) is clear for $i = N$; for $i < N$,

$$\sum_{j \geq i} \int_{AS_j} \alpha_i \, dP = \int_{AS_i} \alpha_i \, dP + \int_{A(S_{i+1} \cup \ldots \cup S_N)} \alpha_i \, dP.$$

But on $S_{i+1} \cup \ldots \cup S_N$, $\alpha_i = \mathscr{E}_i \alpha_{i+1}$; making this replacement in the final integral and using induction backwards on i completes the proof. The relationship (1-4) is clear for $j = N$; for $j < N$,

$$\sum_{i>j} \int_{A_i} \alpha_j \, dP = \int_{A_{j+1} \cup \ldots \cup A_N} \alpha_j \, dP \leqq \int_{A_{j+1} \cup \ldots \cup A_N} \alpha_{j+1} \, dP$$
$$= \int_{A_{j+1}} \alpha_{j+1} \, dP + \sum_{i>j+1} \int_{A_i} \alpha_{j+1} \, dP,$$

where the inequality is obtained from the fact that always $\alpha_j \leqq \mathcal{E}_j \alpha_{j+1}$. An induction backwards on j now completes the proof of (1-4).

The Best Sequential Procedure

We are given now a sequence of functions $w_0, w_1, \ldots, w_N, \ldots$, where $w_N = r_N(x_1, \ldots, x_N) + c_N(x_1, \ldots, x_N)$. The sequence $r_N(x)$ is uniformly bounded, since we supposed the original loss function $L(u,a)$ to be bounded, and we have shown that $r_N \geqq \mathcal{E}_N r_n$ for $n > N$. We shall suppose that $c_N(x)$ is a nondecreasing sequence, $c_N(x) \to \infty$ as $N \to \infty$ for all x. We now construct a best sequential procedure.[2]

The best sequential procedure is obtained as a limit of the best truncated procedures given in the preceding section.

We first define $\alpha_{NN} = w_N$, $\alpha_{jN} = \min(w_j, \mathcal{E}\alpha_{j+1,N})$, $S_{jN} = \{w_i > \alpha_{iN}$ for $i < j$, $w_j = \alpha_{jN}\}$. For fixed j, α_{jN} is a decreasing sequence of functions; say $\alpha_{jN} \to \alpha_j$ as $N \to \infty$. Then $\alpha_j = \min(w_j, \mathcal{E}_j \alpha_{j+1})$. Define $S_j = \{w_i > \alpha_i$ for $i < j$, $w_j = \alpha_j\}$. We shall prove that $T^* = \{S_j\}$ is a best sequential procedure, that is, T^* is a sequential procedure, and for any sequential procedure $T = \{B_j\}$,

$$w(T^*) = \sum_{j=0}^{\infty} \int_{S_j} w_j \, dP \leqq \sum_{i=0}^{\infty} \int_{B_i} w_i \, dP = w(T).$$

Now

$$\sum_{i=N+1}^{\infty} \int_{B_i} w_i \, dP \geqq \int_{\bigcup_{i>N} B_i} c_N \, dP \geqq \int_{\bigcup_{i>N} B_i} w_N \, dP - M \sum_{i>N} P(B_i),$$

2. The assumption made here is somewhat weaker than condition 6 in Wald (1947a, p. 297). The only other assumption made, that $L(u,a)$ is bounded, is condition 1 in Wald (1947a, p. 297).

where M is the uniform upper bound of $r_1(x)$, $r_2(x)$, . . . Thus

$$\sum_{i=0}^{N-1} \int_{B_i} w_i\, dP + \int_{B_N} w_N\, dP + \int_{\bigcup_{i>N} B_i} w_N\, dP$$

$$\leqq w(T) + MP(B_{N+1} \cup \ldots),$$

so that $w(T_N) \leqq w(T) + MP(B_{N+1} \cup \ldots)$, where T_N is the truncated test $B_0, \ldots, B_{N-1}, B_N \cup B_{N+1} \cup \ldots$. From the preceding section,

$$\sum_{j=0}^{N} \int_{S_{jN}} w_j\, dP \leqq w(T_N)$$

for all N. Then

$$\sum_{j=0}^{\infty} \int_{S_j} w_j\, dP \leqq w(T),$$

letting $N \to \infty$, and using Lebesgue's convergence theorem and the easily verified fact that the characteristic function of S_{jN} approaches that of S_j, and therefore $w(T^*) \leqq w(T)$.

It remains to prove that T^* really is a sequential test, that is,

$$\sum_{N=0}^{\infty} P(S_N) = 1.$$

Write $A_N = C(S_0 + \ldots + S_N)$, $\Pi_{N=1}^{\infty} A_N = A$; we show that $P(A) = 0$. It is easily verified by induction that $\alpha_{jN} \geqq c_j$ for all N. The relation (1-3), with $i = 0$, $A =$ sample space, yields that

$$\alpha_{0N} \geqq \sum_{j=m+1}^{N} \int_{S_{jN}} c_m\, dP = \int_{C(S_{0N} + \ldots + S_{mN})} c_m\, dP$$

for all $m < N$. Then

$$\alpha_0 \geqq \int_{A_m} c_m\, dP \geqq \int_A c_m\, dP$$

for all m. Since $c_m \to \infty$, $P(A) = 0$.

We now prove that $w(T^*) = \alpha_0$. If T_N^* denotes the truncated test $S_0, S_1, \ldots, S_N + S_{N+1} + \ldots$, the proof that $w(T_N) \to w(T)$ shows that

$w(T^*_N) \to w(T^*)$. Also (1-3), with $i = 0$, A = sample space, shows that

$$\alpha_{0N} = \sum_{j=0}^{N} \int_{S_{jN}} w_j \, dP.$$

Since $\{S_{0N}, \ldots, S_{NN}\}$ is the best of all procedures truncated at N, and T^* is the best of all sequential procedures, $w(T^*) \leq \alpha_{0N} \leq w(T^*_N)$. Letting $N \to \infty$ yields $w(T^*) = \alpha_0$.

Now $S_0 = \{w_0 \leq \alpha_0\}$; that is, the best procedure T^* is to take no observations if and only if there is no sequential procedure which reduces the risk below its present level. This remark, which identifies our procedure with that characterized by Wald, at least at the initial stage, will be useful in the next section.

2. Bayes Solutions for Finite Multivalued Decision Problems

In this section we shall seek a characterization of the optimum sequential procedure developed in Section 1 in cases where the number of alternative hypotheses is finite. It will be shown that the optimum sequential test for a k-valued decision problem is completely defined by k (or a sequence of k) convex regions in a $(k-1)$-dimensional simplex spanned by the unit vectors. No procedure has yet been developed for determining the boundaries of these regions in the general case. However, for $k = 2$ (dichotomy) and for a linear cost function, a method for determining the two boundaries has been found and the optimum test is shown to be the sequential probability ratio test developed by Wald (1947b).

Statement of the Problem

We are given k hypotheses $H_1, H_2, \ldots, H_k$, where each H_i is characterized by a probability measure u_i defined over an R-dimensional sample space E_R and has an a priori probability g_i of occurring. We are also given a risk matrix $W = (w_{ij})$, $(i,j = 1,2, \ldots, k)$, where w_{ij} is a nonnegative real number and represents the loss incurred in accepting the hypothesis H_j when in fact H_i is true. (We shall assume that $w_{ii} = 0$ for all i. This is based on the supposition, which appears reasonable, that the decision maker is not to be penalized for selecting the correct alternative, no matter how unpleasant its consequences may be.) In addition to the risk matrix (w_{ij}) we shall assume that the cost of experimentation depends only on the number (n) of observations taken and is given by a function $c(n)$ which approaches infinity as n approaches infinity. The problem is to characterize the procedure for deciding on one of the k alternative hypotheses which results in a minimum average risk. This risk is defined as the average cost of taking observations plus the average loss resulting from erroneous decisions.

Structure of the Optimum Sequential Procedure

Let G_k stand for the convex set in the k-dimensional space defined by the vectors $\mathbf{g} = (g_1, g_2, \ldots, g_k)$ with components $g_i \geqq 0$ and $\Sigma_{i=1}^k g_i = 1$; and let $H = (H_1, H_2, \ldots, H_k)$ represent the k hypotheses under consideration. Then every vector $\mathbf{g}$ in G_k may be considered as a possible a priori probability distribution of H.

For any $\mathbf{g}$ in G_k and for any sequential procedure T (see the definition in Section 1), let $R(\mathbf{g}|T)$ represent the average risk entailed in using the test T when the a priori distribution of H is $\mathbf{g}$. Then

$$R(\mathbf{g}|T) = \sum_{i=1}^{k} g_i \mathscr{E}_i[c(n)|T] + \sum_{i=1}^{k} \sum_{j=1}^{k} g_i w_{ij} P_{ij}(T), \tag{1-6}$$

where $\mathscr{E}_i[c(n)|T]$ is the average cost of observations when the sequential test T is used and H_i is true, and $P_{ij}(T)$ is the probability that the sequential test T will result in the acceptance of H_j when H_i is true. The risk involved in accepting the hypothesis H_j prior to taking any observations will be designated by $R_j (j = 1,2, \ldots, k)$ and is given by

$$R_j = \sum_{i=1}^{k} g_i w_{ij}. \tag{1-7}$$

We now define k subsets $S_1^*, S_2^*, \ldots, S_k^*$ of G_k as follows. A vector $\mathbf{g}$ of G_k will be said to belong to S_j^* if (a) min $(R_1, R_2, \ldots, R_k) = R_j$ and (b) $R(\mathbf{g}|T) \geqq R_j$ for all T. We observe that since the unit vector with 1 in the jth component belongs to S_j^*, the subsets S_j^* are nonempty. We now prove the following theorem.

THEOREM 2. *The sets S_j^* are convex. That is, if $\mathbf{g}_1$ and $\mathbf{g}_2$ belong to S_j^* so does $\mathbf{g} = a\mathbf{g}_1 + (1 - a)\mathbf{g}_2$ for all a, $0 \leqq a \leqq 1$.*

Proof. Assume the contrary. Then there exists a sequential procedure T such that

$$R(\mathbf{g}|T) = \sum_{i=1}^{k} g_i \mathscr{E}_i[c(n)|T] + \sum_{i=1}^{k} \sum_{j=1}^{k} g_i w_{ij} P_{ij} < \sum_{i=1}^{k} g_i w_{ij}. \tag{1-8}$$

But by definition, if either $\mathbf{g}_1$ or $\mathbf{g}_2$ represents the a priori distribution of the hypotheses H, we must have for all sequential procedures and hence for T,

$$R(\mathbf{g}_1|T) = \sum_{i=1}^{k} g_{1i}\mathscr{E}_i[c(n)|T] + \sum_{i=1}^{k}\sum_{j=1}^{k} g_{1i}w_{ij}P_{ij} \geqq \sum_{i=1}^{k} g_{1i}w_{ij}, \tag{1-9}$$

and

$$R(\mathbf{g}_2|T) = \sum_{i=1}^{k} g_{2i}\mathscr{E}_i[c(n)|T] + \sum_{i=1}^{k}\sum_{j=1}^{k} g_{2i}w_{ij}P_{ij} \geqq \sum_{i=1}^{k} g_{2i}w_{ij}. \tag{1-10}$$

If we now multiply (1-9) by a and (1-10) by $(1 - a)$ and add, we see that the resulting expression contradicts (1-8). This proves the theorem.[3]

It is easily seen that for given hypotheses H the shape of the convex regions S_j^* will depend on the cost function $c(n)$ and the risk matrix W. Thus if the cost of taking a single observation were prohibitive, the region S_j^* in G_k would simply consist of all vectors $\mathbf{g}$ for which min $(R_1,R_2, \ldots, R_k) = R_j$. On the other hand, if the cost of taking observations were negligible and the risk of making an erroneous decision large, the regions S_j^* would shrink to the vertices of the polyhedron G_k. To exhibit the dependence of the regions S_j^* on H, $c(n)$, and W, we shall use the symbol $S_j^*[H,c(n),W]$. We shall also use the symbol $S^*[H,c(n),W]$ to represent the region consisting of all vectors $\mathbf{g}$ in G_k which belong to the complement of $\bigcup_{j=1}^{k} S_j^*[H,c(n),W]$.

We now define $c_N(n) = c(N + n) - c(N)$ for all $N = 0, 1, 2, \ldots$. Thus $c_N(n)$ represents the cost of taking n observations when N observations have already been taken. It will now be shown that, for random sampling, the problem of characterizing the optimum sequential procedure T^* for a given H, $c(n)$, and W reduces itself to the problem of finding the boundaries of the regions $S_j^*[H,c_N(n),W]$ for all N. The truth of this can be seen from the following considerations.

We are initially given a vector $\mathbf{g}$ in G_k as the a priori distribution of the hypotheses H. Initially we are also given a matrix W and a function $c(n) = c_0(n)$. Now assume we have taken N independent observations ($N = 0,1,2, \ldots,$). These N observations transform the initial state into one in which (a) the vector $\mathbf{g}$ goes into a vector $\mathbf{g}^{(N)}$ in G_k where each component $g_i^{(N)}$ of $\mathbf{g}^{(N)}$ represents the new a priori probability of the hypothesis H_i (that

3. A similar proof shows the convexity of the corresponding regions in cases where the number of alternatives is infinite.

is, the a posteriori probability of H_i given the values of the N observations), (b) the risk matrix W remains unchanged, and (c) the cost function $c(n)$ goes into the function $c_N(n)$.

Assume now that the boundaries of the regions $S_j^*[H,c_N(n),W]$ are known for each j and N. Then, if we take the observations in sequence, we can determine at each stage $N(N = 0,1,2, \ldots ,)$ in which of the $k+1$ regions the vector $\mathbf{g}^{(N)}$ lies. If $\mathbf{g}^{(N)}$ lies in $S^*[H,c_N(n),W]$, then, by definition of this region, there exists a sequential test T which, if performed from this stage on, would result in a smaller average risk than the risk of stopping at this stage and accepting the hypothesis corresponding to the smallest of the quantities $R_j^{(N)} = \Sigma_{i=1}^k g_i^{(N)} w_{ij} (j = 1,2, \ldots , k)$. But it has been shown in Section 1 that if any sequential test T is worth performing, the optimum test T^* is also worth performing. Now T will coincide with T^* for at least one additional observation. But when that observation is taken $\mathbf{g}^{(N)}$ will become $\mathbf{g}^{(N+1)}$ and $c_N(n)$ will become $c_{N+1}(n)$. Again if $\mathbf{g}^{(N+1)}$ lies in $S^*[H,c_{N+1}(n),W]$, the same argument will show that it is worth taking another observation. However, if $\mathbf{g}^{(N+1)}$ lies in $S_j^*[H,c_{N+1}(n),W]$ for some j, it implies that there exists no sequential test T which is worth performing, and hence the optimum procedure is to stop sampling and accept H_j.

Thus we see that the optimum sequential test T^* is identical with the following procedure. Let $N = 0, 1, 2, \ldots ,$ represent the number of observations taken in sequence. For each value of N we compute the vector $\mathbf{g}^{(N)}$ representing the a posteriori probabilities of the hypotheses H. As long as $\mathbf{g}^{(N)}$ lies in $S^*[H,c_N(n),W]$ we take another observation. We stop sampling and accept $H_j (j = 1,2, \ldots , k)$ as soon as $\mathbf{g}^{(N)}$ falls in the region $S_j^*[H,c_N(n),W]$.

We have as yet no general method for determining the boundaries of $S_j^*[H,c_N(n),W]$ for arbitrary H, $c_N(n)$, and W. However, in the case of a dichotomy ($k = 2$) and a linear cost function, such a method has been found and will be discussed here in detail. Some illustrative examples of the optimum sequential test for trichotomies ($k = 3$) will also be given.

3. Optimum Sequential Procedure for a Dichotomy when the Cost Function is Linear

We are given two alternative hypotheses H_1 and H_2, which, for the sake of simplicity, we assume are characterized respectively by two probability densities $f_1(x)$ and $f_2(x)$ of a random vector X in an R-dimensional Euclidean space. (If X is discrete, $f_1(x)$ and $f_2(x)$ will represent the probability under the