Universität Bielefeld/IMW

Working Papers
Institute of Mathematical Economics

Arbeiten aus dem
Institut für Mathematische Wirtschaftsforschung

Nr. 31

Hans W. Gottinger

Sequential Analysis and Optimal Stopping

May 1975

Anno 1600 nach Merian

Institut für Mathematische Wirtschaftsforschung
an der
Universität Bielefeld
Adresse/Address:
Schloß Rheda
484 Rheda
Bundesrepublik Deutschland
Federal Republic of Germany

Abstract

SEQUENTIAL ANALYSIS AND OPTIMAL STOPPING

Hans W. Gottinger

This exposition reviews main ideas in the theory of sequential-decision-making and optimal stopping. It attempts to show the wide range of possible applications in operations research, management science, control and system engineering, economics and statistics.
The bibliography provides a (not neccessarily representative) sample on increasing activities of sequential decision theory and methods in many areas of interest.

Sequential Analysis and Optimal Stopping[+]

1. General Considerations

We are here going to analyze and discuss a very important class of decision problems which involve time explicitly as an irreversible resource. These problems are known as dynamic or sequential decision problems. They have a most natural formulation since every real-life decision has to take care of 'time-induced' changes to which the decision maker has to adjust or to adapt. These types of problems may be extremely complex: they may involve changes in preferences, technology and resources, the environment. Complexities may be added by uncertainty or lack of information and multidimensionality. A general, very useful technique of resolving dynamic decision problems has been introduced by R. Bellman's dynamic programming [1]. The original class of decision problems treated by dynamic programming were restricted to deterministic problems. Later dynamic programming in conjunction with the theory of Markov chains and general stochastic processes have covered uncertainty, and the case of conflict among many decision-makers acting sequentially in time has been treated by differential game theory. In statistics, sequential analysis was developed by A. Wald [16] in the forties as a consequent extension of his statistical decision theory.

All these problems, although originating in various subjects, have common elements and also involve similar methods.

We first describe some of the problems and methods and then turn toward statistical problems in which Bayesian methods play a crucial role. Bayes' theorem obtains new importance in view of obtaining new information by sequential experimentation.

[+]Manuscript completed: May 13,1975.

2. Sequential Decision Problems

In every conventional decision problem one is faced with the situation to act in one or another way. Then if the decision is taken and a particular event occurs, a certain result will be obtained. In a sequential decision problem one has another option which could be summarized as 'wait and see', or 'go on and take another observation'. This choice problem constitutes one stage, if the choice problem is carried over several stages then the obvious question arises when should the decision maker stop in collecting further information, this depends upon his expected utility of taking one more observation. But it is difficult to compute the expected utility of taking one more observation. In order to find the best decision now, i.e. whether to stop and make a decision or to go on and take another observation, it is necessary to know the best decision in the future. Consequently, the search for an optimal decision should not proceed according to chronological time but in reverse order to work backwards in time since the present optimum involves the future optimum. This fundamental fact is incorporated in the principle of dynamic programming.

Let $U(w,t)$ be the expected utility of the best sequential scheme when starting from a situation in which the parameters describing the distribution of θ are w and t, w changes randomly and t deterministically. (w could be the mean of the distribution of θ, and t the inverse of the variance, i.e. the precision.) Let $\bar{U}(w,t)$ be the expected utility of taking the best decision now, without further observations. Then $U(w,t)$ is either $\bar{U}(w,t)$ (and it is not worth taking further observations) or it is worth taking further observations. In the latter case we start with $U(w,t)$ and look at the change of the situation, i.e. w and t change to w' (randomly) and $t+h$ (deterministically). Let then $P(w'|w,t)$ be the distribution of w' given w and t, let $c(t,h)$ be the cost of 'taking one more observation' from t to $t+h$.

Then in general, by taking further observations the expected utility functional is

$$U(w,t) = \int U(w',t+h)P(w'|w,t)dw'-c(t,h).$$

The optimality principle, according to dynamic programming, requires that

$$U(w,t) = \max\{\bar{U}(w,t) : \int U(w',t-h)P(w'|w,t)dw' - c(t,h)\}$$

This optimality principle yields a unique criterion on optimal stopping depending on whether or not

$$U(w,t) = \max\{\bar{U}(w,t)\}$$

3. The Marriage Problem

This type of sequential decision problem is representative for a very general class of decision problems that can be solved via dynamic programming, [14] or with other tools [8].

A known number of ladies , n, are going to be inspected in a random order. You are able to rank them according to some fixed criterion catalogue as to which lady will best meet your standards. Let r, an integer, be that number indicating the rank among n ladies, $1 \leq r \leq n$. At any stage of this procedure you may either propose to one lady (by which the procedure stops) or continue inspecting. Whenever you inspected a girl and you didn't propose she will never come back, i.e. she will never get inspected again.

If you propose to a lady she will always accept. What is the optimal stopping rule? The desirability of every lady to be inspected is represented by a utility index U_i, the utility of being married to the $_i$-th lady with the $_i$-th rank with

$$U_{i-1} \geq U_i \geq U_{i+1}$$

We denote by r the number of n ladies, and by s the apparent rank after some ladies have been inspected, hence r changes deterministically and s changes randomly.

Correspondingly, we denote the expected utility by U(s,r) and $\bar{U}(s,r)$, respectively. Now, the probability that the r-th lady of apparent rank s will have true rank S is easily calculated by the binomial equation

$$\binom{S-1}{s-1}\binom{n-S}{r-s}\Big/\binom{n}{r} = P_{S:s,r}$$

Hence we have

(2) $\bar{U}(w,t) = \bar{U}(s,r) = \sum_{i=s}^{s+n-r} P_{S:s,r}$

as the expected utility, and $\bar{U}(s,r)$ is considered to be a known function. Given the situation to have chosen the s-th rank out of r inspections, the probability that the next lady will have apparent rank s' is clearly 1/(r+1) for all s' so that $P(s'|s,r)$ equals 1/(r+1).

Hence, the expected utility functional becomes

(3) $U(s,r) = \max \{\bar{U}(s,r) : \sum_{i=1}^{r+1} U(i,r+1)/r+1)\}$

Consider two cases.

1) Set $U_1 = 1$ and $U_i = 0$ for $i > 1$, i.e. follow the instruction 'always take the best'. Then the optimality criterion - to search for - is according to (3).

(4) $U(1,r) = \max \{r/n : \sum_{i=1}^{r+1} U(i,r+1)/(r+1)\}$

and

(5) $U(s,r) = \sum_{i=1}^{r+1} U(i,r+1)/(r+1)$ for $s > 1$

U(s,r) must be a function only of r since with increasing r it is more likely to find the true top rank S which coincides with s. (4) and (5) may be written in terms of recursive functions.

(6) $U(1,r) = \max\{r/n, n_r\}$

(7) $n_r = \frac{1}{r+1}\{U(1,r+1) + rn_{r+1}\}$.

2) Suppose $U(1,r) > \bar{U}(1,r) = r/n$, i.e. the utility of continuing exceeds that of proposing. It follows from (6) that $U(1,r) = n_r$ and from (7), by reducing the value of r by one , $u_{r-1} = u_r$.

Therefore, $u_r > r/n$, $u_{r-1} > (r-1)/n$ and from (6) $U(1,r) > \bar{U}(1,r) > U(1,r-1) > (r-1)/n$. If it is not worth proposing to a lady who is best out of r it is not worth proposing to a lady who is best out of (r-1). The best strategy must be to propose to a lady who is best out of r, provided r is large enough. How large should r be?

Suppose that $U(1,r) = \bar{U}(1,r) = r/n$ and $U(1,r') = r'/n$ for all $r' \geq r$. From (7) we derive

(8) $u_r = \frac{1}{r+1} \{\frac{r+1}{n} + ru_{r+1}\}$ or if $v_r = u_r/r$: $v_r = \frac{1}{nr} + v_{r+1}$,

for all r. Adding together the r.h.s. of these equations we get

(9) $$v_r = \frac{1}{n}\left\{\frac{1}{r} + \frac{1}{r+1} + \ldots + \frac{1}{n-1}\right\}$$ and

(10) $$\frac{1}{r} + \frac{1}{r+1} + \ldots + \frac{1}{n-1} < 1$$ (bounded by 1).

As long as (10) exists it is worth proposing to the best lady out of r.

Let r = R be the least such value, that is

(11) $$\frac{1}{R} + \frac{1}{R+1} + \ldots + \frac{1}{n-1} < 1 < \frac{1}{R-1} + \frac{1}{R} + \frac{1}{R+1} + \ldots + \frac{1}{n-1}$$

Then from (7) with $r = R-1$:

(12) $$u_{R-1} = \frac{1}{R}\left\{\frac{R}{n} + (R-1)u_R\right\} = \frac{R-1}{n}\left\{\frac{1}{R-1} + \frac{1}{R} + \ldots + \frac{1}{n-1}\right\}$$

from (9).

If n is large the value of R is given by $\int_R^n dx/x = 1$ where the series (10) is approximated by an integral and hence $n/R = e$ is the base of natural logarithm. Hence for large n the optimum rule is to inspect until a proportion e^{-1}(0.368) have been inspected and then to propose to any subsequent lady of apparent rank one. The expected utility, given by (12), is calculated by e^{-1}. If someone looks for a marriage partner at 18 through 40 (i.e. 22 years) one should never propose until age 18 + 0.368 = 26.

4. Stopping Rule Problems

The essential features of a stopping rule problem can be split into two parts, consisting of:

1. A probabilistic mechanism, that is, a random device that moves from state to state under a known, partially known, or unknown probability law.

2. A payoff and decision structure such that, after observing the current state, we have a choice of at most two decisions.

(a) Take your accumulated payoff to date and quit.

(b) Pay an entrance fee for the privilege of watching one more observation.

This procedure is very natural for casino or gambling problems.

(9) $$v_r = \frac{1}{n}\left\{\frac{1}{r} + \frac{1}{r+1} + \dots + \frac{1}{n-1}\right\}$$ and

(10) $$\frac{1}{r} + \frac{1}{r+1} + \dots + \frac{1}{n-1} < 1$$ (bounded by 1).

As long as (10) exists it is worth proposing to the best lady out of r.

Let $r = R$ be the least such value, that is

(11) $$\frac{1}{R} + \frac{1}{R+1} + \dots + \frac{1}{n-1} < 1 < \frac{1}{R-1} + \frac{1}{R} + \frac{1}{R+1} + \dots + \frac{1}{n-1}$$

Then from (7) with $r = R-1$:

(12) $$u_{R-1} = \frac{1}{R}\left\{\frac{R}{n} + (R-1)u_R\right\} = \frac{R-1}{n}\left\{\frac{1}{R-1} + \frac{1}{R} + \dots + \frac{1}{n-1}\right\}$$

from (9).

If n is large the value of R is given by $\int_R^n dx/x = 1$ where the series (10) is approximated by an integral and hence $n/R = e$ is the base of natural logarithm. Hence for large n the optimum rule is to inspect until a proportion e^{-1}(0.368) have been inspected and then to propose to any subsequent lady of apparent rank one. The expected utility, given by (12), is calculated by e^{-1}. If someone looks for a marriage partner at 18 through 40 (i.e. 22 years) one should never propose until age 18 + 0.368 = 26.

4. Stopping Rule Problems

The essential features of a stopping rule problem can be split into two parts, consisting of:

1. A probabilistic mechanism, that is, a random device that moves from state to state under a known, partially known, or unknown probability law.

2. A payoff and decision structure such that, after observing the current state, we have a choice of at most two decisions.

(a) Take your accumulated payoff to date and quit.

(b) Pay an entrance fee for the privilege of watching one more observation.

This procedure is very natural for casino or gambling problems.

Some of the following examples have been lucidly described by L. Breiman [6].

Unrestricted coin-tossing, restricted coin-tossing, house-hunting problem, purchasing a used car, parking place problem, the stock market problem, the job searching problem, the dynamic inventory problem.

Other classes or problems contain the product design problem, medical diagnosis and sequential control processes.

The gambling problem often involves the following scheme:

An urn contains N red balls and M blue balls. You are allowed to draw with or without replacement as long as there are any balls in the urn. Each time you draw a red ball you receive one dollar and each time you draw a blue ball you lose one dollar. This situation can be formulated as follows.

Let $\{X_k, k = 1,2,...\}$ be a sequence of random variables, and let $X_k = +1$ provided the k-th toss results in heads and -1 provided it results in tails. Then $S_n = \Sigma_1^n X_k$ is the accumulated profit at the n-th toss. If you and your opponent are infinitely wealthy you could adopt two decision rules: Stop and collect your profits (or pay your losses) or continue for at least one more toss. (In case of a restricted coin-tossing situation, where your initial fortune S_o is finite, there may be at some stage a forced stopping rule, e.g. $S_n = -S_o$, you have lost all your money and must quit.)

An attractive strategy would be 'to stop when you are ahead'. Such coin-tossing game can be treated fully by Markov devices. Examples of this type will be encountered in a subsequent section on gambling problems. As a probabilistic device we can understand a Markov chain with a countable number of states and specified stationary transition probabilities. Under the conditions specified, Markov devices work well far gambling problems. Other devices are possible, those which require more or less restrictions such as simple random walks, Brownian motion, martingales, or other stochastic processes (see L. Breiman [7], L. Doob [9], Dubins and Savage [10], Griffeath and J.L. Snell [13].) For practical applications and computational work they rely in most cases on the functional equation approach of dynamic programming. The initial state S_o covers all relevant information of the past. The Markov property is defined by the assumption that if S_n (e.g. accumulated winnings) is the state at time n, then $P(S_{n+1} = s_j$ given $S_n = s_j$ and all past history up to time n) $= P(s_j|s_i)$, where $P(s_j|s_i)$ is the specified transition probability of moving to state s_j from state s_i, satisfying

$$P(s_j|s_i) \geq 0, \sum_{s_j} P(s_j|i) = 1.$$

One immediate empirical outcome of the Markov property is to be seen in the fact that the past is incorporated in the present so that the transition probabilities are only conditioned on the present state.

5. Payoff and decision structure

Any stopping rule problem may involve a specification of the initial conditions and a payment of the entrance fee (as a compensation for taking part in a game). After some duration of the play you may either collect your winnings or losses to date or continue playing. The collection of your winnings or losses my be referred to as your (sequential) payoff. The terminal payoff $F(s_i)$ is the integral payoff over time. Payoff and entrance fees are virtually dual notions in this context. To the decision rule 'stop and collect the payoff $F(s_i)$' we will associate a stopping set T_s containing all forced stopping states at which we must stop and collect $F(s_i)$. For example, you may end in a state s_i where the game is forced to stop, for whatever reason. Likewise, consider a set T_c, disjoint from T_s, containing all forced continuation states associated to the decision rule 'continue and pay the fee $f(s_j)$', i.e. being in a state s_j we may be forced to continue playing, the set of these states is the set of forced continuation states. So payoffs and costs (of observation) are dually related.

Example: Consider a simple coin-tossing game. The probabilistic device has the Markov property with transition probabilities:

$$P(j|i) = \frac{1}{2} \text{ for } j = i+1,$$

$$= \frac{1}{2} \text{ for } j = i-1$$

$$= 0 \text{ for other values of } j.$$

A well defined stopping rule must tell us when to stop along each possible sequence of states, since otherwise it would be possible for the device to produce a sequence of states along which our rule would not hold. Therefore, a stopping rule involves a matching process between a criterion catalogue represented by expected utility and computational costs in terms of 'costs of observation'. The best stopping rule, the solution we are seeking, makes the expected utility or expected monetary value as large as possible. In general, expected total payoff (rather than utility) is defined by

$$EZ = \Sigma Z(s_o, s_1, \ldots, s_n) P(s_o, s_1, \ldots, s_n)$$

where the sum is taken over all stopping sequences ont the list weighted by the probability of the sequence. $Z(s_o, s_1, \ldots, s_n)$ is the payoff function.

A stopping rule then involves a binary choice 'either stop or continue observing, sampling etc.'. Let T be the set of these rules, then a stopping rule T^* is <u>optimal</u> with respect to the total expected payoff if $E_{T^*}Z \geq E_T Z$ for all other stopping rules T.

Stability of a stopping rule means that it can be approximated, in terms of payoff, by rules in which we decide to quit after a large but fixed number of plays. This corresponds to a forced stopping rule.

6. Decision Trees

A heuristic device for representing sequential decisions is given by decision trees. Also it presents an intuitive meaning to the idea of structuring and organizing complex decisions in a time context, where a decision problem can be broken down into a sequence of problems which follow one another in a natural time order. Standard examples of that sort are the product decision problem or other R & D decision problems, the medical diagnosis and treatment problem and the investment decision problem, but the list can be extended almost indefinitely. Example of a product decision problem can be traced as follows: The decision grows from the left to the right, and it reflects the structure of decisions (decision nodes) and uncertainty (random nodes) in a sequential framework. Although it proceeds in a chronological order, it has been demonstrated by the optimality principle of dynamic programming that to obtain an optimal decision in terms of maximizing expected utility (or minimizing expected loss) it is necessary to proceed in the reverse direction since an optimal sequential decision can only be maintained if each of the next steps of the decision have shown to be optimal. A decision tree consists of a series of branches (corresponding to the complexity of the problem). Summarizing, the decision tree method proceeds in the following stages:

(1) the tree is written out in chronological order, the decisions and events being described by branches in the order in which they occur,

(2) probabilities are attached to branches emerging from random nodes in any coherent way,

(3) utilities are attached to the terminal branches,

(4) proceeding back from the terminals to the base, by taking expectations at random nodes and maximizing at decision nodes, the best decisions and their expected utilities are determined.

This procedure is well adapted to computational work and to the structure of dynamic programming, see D.V. Lindley [15]. Chapt. 8.

7. Adaptive Processes and Optimal Stopping

Another type of application of optimal stopping rule is motivated by adaptive processes occuring in sequential sampling. Suppose that there are two drugs available for treatment of a certain disease. It is not known which one is more efficient. How are the drugs to be used to save as many patients as possible? If two large samples are used initially we can determine which is better with high confidence. But this implies that a high proportion of patients would have been treated with the inferior drug. A more successful procedure would consist in trying the drugs initially on small sample groups, observing the outcomes, weighting the next test in favor of which seems to be the better drug on the basis of current knowledge, and continuing in this fashion. This procedure involves optimal stopping at a stage when sufficient information is collected allowing to choose the superior drug. Learning and acting simultaneously are already involved. For adaptive control processes a decision-making device (or controller) is called upon to perform under various conditions of uncertainty, conditions which may range from complete knowledge to total ignorance. In fact, optimal stopping rules can be viewed as control devices governing a deterministic or stochastic dynamic system. The evolution of the system will be the result of the interaction between the laws of notion of the system and the sequence of actions taken over time.

To show the connection between adaptive system theory and optimal stopping rules we exhibit the simple transformation of a deterministic process.

We assume that when a system is in the current state s and a decision a is made, the new state of the system s_1 is $s_1=T(s,a)$. Then if the system is observed to be in state s_1 and the decision a_1 is made, the system is transformed into $s_2 = T(s_1,a_1)$.

Next with the system observed to be in state s_2 and the decision a_2 to be chosen, the new state is $s_3 = T(s_2,a_2)$, etc.

Then after N stages the system will be transformed into state $s_N = T(s_{N-1}, a_{N-1})$. Therefore the pair (s,a) generates a semigroup of transformations over the state space S. The optimality of process is to obtain a final state at which the(expected) net gain or loss is maximal or minimal. The expected gain or loss consists of the difference of the gain or loss associated to the final state and the costs involved to obtain that state, hence it is a function of the final state $\phi(s_N)$.

Bellman's principle of optimality is very natural in pursuing this goal, and the functional equation technique is very appropriate in reaching it. It states: An optimal sequence of decisions has the property that whatever the initial state and decision are, i.e. given the pair (s_o, a_o) the remaining decisions must be optimal with regard to the state resulting from the initial decision.

A more interesting and less restrictive case is that of partially known transformations. Those involve processes in which the outcome of a decision is not precisely known. This means instead of considering deterministically known states we are concerned with random states. Let it be assumed that the controller or decision-maker does not know the exact distribution of possible resulting random states, rather he has an initial estimate of this distribution that may be justified by adopting the Bayesian philosophy, and in the process of decsion-making he is able to modify this estimate in the light of the actually observed history of the process which can be transformed via Bayes formula to obtain new information. This idea is basic to learning system theory in a random environment (K.S. Fu [12]). An interesting application of this procedure for the treatment of patients has been given by J. Cornfield [18].

We first give a general treatment of this situation before we turn to specific sequential estimation techniques in this framework.

Let the decision maker's knowledge be specified by an information pattern I. This information pattern contains all the information about the past on which the future actions are to be based. It may be represented by an a priori probability density function. The overall state of the system plus the decision maker's knowledge is specified by a point in a new state (s,I). Then, if the state of the system is (s,I) and the decision a is made, the new state is $s_1 : I_1 = T_1(s,I,a,r)$, where r is a random variable having an a priori probability distribution function G(s,I,a,r), knowledge of which in itself is part of the information pattern. The new information pattern I_2 is specified by a transformation

$$I_2 = T_2(s,I_1;a,r).$$

The goal is to determine a sequence of decision $\{a_1,a_2,\ldots,a_N\}$ that will minimize the expected value of a preassigned function of the final state $\phi(s_N,I_N)$. Since the exact distribution functions are not known, the expected value is taken regarding the a priori distribution functions as the true ones.

Introduce a sequence of functions $\{f_k(s,I)\}$

$$f_k(s,I) = \min_{\{a_1,\ldots,a_k\}} E\left[\phi(s_k,I_k)\right].$$

Then the principle of optimality yields the relationships

$$(1) \qquad f_{N+1}(s,I) = \min_a \int_r f_N(T_1(s,I;a,r),T_2(s,I_1;a,r))\,dG(s,T;a,r), \quad N = 1,2,\ldots$$

and for N = 1 we have

$$(2) \qquad f_1(s,I) = \min_a \int_r \phi(T_1(s,I;a,r),T_2(s,I_1,a,r))\,dG(s,I;a,r).$$

The relationship (1) and (2) can be used for establishing the existence of optimal policies.

8. Sequential Estimation

Sequential estimation and related sequential detection processes of this type to be discussed occur in radar and communication technology where the receiver uses variable rather than fixed sample sizes. In such cases the principle of optimality provides a natural mathematical formulation and a numerical solution. Let us assume that the task of a controller is to estimate the value of an unknown probability p. Take a binary sequential experiment where p is the unknown probability that a certain random variable takes the value unity, and 1-p is the probability that it takes the value zero. The controller is to conduct a series of experiments, record the outcomes, and make an estimate of p on the basis of this experience plus any a priori information available. There are also costs associated with performing each experiment and possibly making wrong estimates of p. The problem is to determine when the experiment should be stopped and what estimate should be made by the controller. Let us specify the situation in detail. Suppose that at the beginning of the process the controller is in possession of the prior information that n ones out of s trials have been observed. Regarding the observation of the outcome of the process itself, we assume that n of r trials have resulted in a one, but here we disregard information concerning the order in which the events occurred. Since p is unknown we regard it as a random variable, its distribution function changes during the course of the process. First consider only the prior information, the change of the distribution function is given by

$$(1) \qquad dG(p) = \frac{p^{n-1}(1-p)^{s-n-1}}{B(n,s-n)}\, dp,$$

where B is the beta function.

Second after m ones have been observed in r additional trials, we consider it to be

$$(2) \qquad dG_{r,m} = \frac{p^{m}(1-p)^{r-m}dG(p)}{\int_0^1 p^{m}(1-p)^{r-m}dG(p)},$$

a Bayes' approach.

Let $c_{r,m}$ denote the expected cost of incorrect estimation after r additional trials have resulted in ones and set

$$(3) \qquad c_{r,m} = \alpha\int_0^1 (p_{r,m}-p)^2 dG_{r,m}(p),$$

where $p_{r,m}$ is the estimate which minimizes $c_{r,m}$. The value of $p_{r,m}$ is given by the formula

$$(4) \qquad p_{r,m} = \int_0^1 p\,dG_{r,m}(p) = \frac{m+n}{r+s}$$

which yields an intuitively reasonable estimate for p. A calculation then shows that

$$(5) \qquad c_{r,m} = \alpha \frac{m+n}{r+s}\left[\frac{m+n+1}{r+s+1} - \frac{m+n}{r+s}\right] .$$

Now suppose that if m experiments have been performed the cost of the next experiment is k(m), allowing for the cost of the experiment to vary during the process, a feature that entails interesting possibilities. We shall assume that in the absence of additional information estimated probabilities are to be regarded as true probabilities. Also we wish to require that no more than R experiments be performed, thus we introduce a <u>forced termination rule</u> of the sequential process. This forced termination rule makes sense in many practical situations, particular those which are alike the marriage problem. If the termination rule comes into effect the process must be truncated at this point. By dynamic programming one can determine the optimal control policy. In doing this the cost function $f_r(m)$ is defined by

(6) $f_r(m)$ = expected cost of a process beginning with m ones in r experiments having been observed, and using an optimal sequence of decisions.

Then the principle of optimality yields the functional equation

$$(7) \qquad f_r(m) = \text{Min} \begin{cases} T_c : k(r) + p_{r,m} f_{r+1}(m+1) + (1-p_{r,m}) f_{r+1}(m) \\ T_s : \alpha \frac{m+n}{r+s}\left[\frac{m+n+1}{r+s+1} - \frac{m+n}{r+s}\right] \end{cases}$$

which holds for m = 0,1,2,...,r and r = R-1,R-2,...,0. The sets T_c and T_s denote the continuation and stopping rule sets respectively. In view of the termination rule we also have

$$(8) \qquad f_R(m) = \alpha \frac{m+n}{R+s}\left[\frac{m+n+1}{R+s+1} - \frac{m+n}{R+s}\right].$$

These relations quickly enable us to calculate the sequence of functions $f_R(m)$, $f_{R-1}(m), \dots, f_0(m)$.

At the same time we can determine whether to stop or continue and what estimate to make of p in the event the process is terminated.

The functional equations have been investigated computationally by R. Bellman and others [2], [3], [4] for a wide range of values for the parameters α and R, and for several cost functions k(m). When the cost of experimentation was constant from experiment to experiment, or when it increased and when one out of two ones had been observed a priori it was found that the optimal policy essentially consisted in:

1. Continuing the experiments if r was small (not enough information present on which to base an estimate).
2. Stopping the experiments if r was sufficiently large.
3. Continuing the experiments for intermediate values of r, unless extreme runs of either zeros or ones occurred and stopping otherwise.

On the other hand, in the case of a decreasing cost of experiment the optimal control policy is more complex. This is intuitively plausible since the cost of experimentation may have dropped to such a low level that it might be profitable to do at least one more experiment before making the estimate.

9. Gambling Problems

Consider the classical ruin problems as formulated by Feller [11, Ch. 14] Suppose a gambler with an initial capital S_o plays against an infinitely rich adversary but the gambler always has the option to stop playing whenever he likes to. The gambler then adopts the strategy (policy) of playing until he loses his capital or obtains a net gain $S_n - S_o$ at the n-th play. Then p is his probability of losing and 1-p the probability of winning. In other words, the gambler's net gain G is a random variable with the values $S_n - S_o$ at probability 1-p and $-S_o$ at probability p. The expected gain is $E(G) = S_n(1-p) + p \cdot -S_o$. The treatment of this problem can be facilitated by interpreting the gambler's process as a random walk with absorption barrier $(0, S_n)$. Such kind of problem immediatley leads to a problem of sequential sampling. Let a particle start from a position S_o such that $0 < S_o < S_n$, we seek the probability p_{S_o} that the particle will attain some position ≤ 0 before reaching any position $\geq S_n$. Then the position of the particle at time n is the point $S_o+X_1+X_2+\ldots+X_n$ where the $\{X_k\}$ are mutually independent random variables with the common distribution $\{p_r\}$. The process stops when for the first time either $X_1+\ldots+X_n \leq -S_o$ or $X_1+\ldots+X_n \geq S_n-S_o$. In sequential sampling the X_k represent certain characteristics of samples or observations. Measurements are taken until a sum $X_1+\ldots+X_k$ falls outside or inside the preassigned limits. In the first case it leads to rejection, in the second case to acceptance. The main ideas of sequential sampling are due to A. Wald (1947). The whole problem can be formulated in terms of a Markov chain. The idea of finding optimal gambling strategies for favorable and unfavorable games has been pursued rigorously in the literature on stochastic processes. In particular, martingales have been found very useful for studying optimal stopping times and stopping rules. However, their investigation require more advanced methods than developed here. The more advanced reader is advised to consult L. Breiman [7], Ch. 5, Doob [9] Ch. II,VII and Dubins and Savage [10]; the standard reader is referred to Feller [11].

10. Sequential Statistical Problems

Consider a statistical problem where the statistician can take observations $X_1, X_2, \ldots, X_n$ at different time from some population involving a parameter W whose value is unknown.

After each observation he can evaluate the information having accumulated up to that time and then make a decision whether to terminate the sampling process or to take another observation (to continue) sampling. This is called sequential sampling. In general, there are some costs of observation (the costs of an experiment), at some stage of sampling the incremental benefits of taking one more observation are offset by the incremental costs of observation. The criterion for the statistician is to minimize the total risk, therefore, in many situations he will compare fixed sample size sampling with sequential sampling with respect to this criterion. Although the benefits of sequential sampling may be determined in advance, the costs of sampling may assume different forms, one particular reasonable form is that costs of sampling may be sharply increasing in the process of taking more and more observations. The risk of the sequential decision rule d in which at least one observation is to taken is

$$\rho(\xi,\delta)= E\{L[W,\delta_N(X_1,\ldots,X_n)] + c_1+\ldots+c_N\}$$

$$=\sum_{n=1}^{\infty}\int_{\{N=n\}}\int_{\Omega}L(w,\delta_n(x_1,\ldots,x_n))\cdot$$

$$\xi(w|x_1,\ldots,x_n)\,d\nu(w)\,dF_n(X_1,\ldots,X_n|\xi)$$

$$+\sum_{n=1}^{\infty}(c_1+\ldots+c_n)\rho\{N=n\}$$

with $\xi(\cdot|X_1,\ldots,X_n)$ being the posterior generalized probability density function of W after the values $X_1 = x_1,\ldots,X_n = x_n$ have been observed. A Bayes sequential decision procedure is a procedure δ for which the risk $\rho(\xi,\delta)$ is minimized, hence it is optimal.

It is said a sequential decision process δ is bounded if there is a positive integer n such that $P_r(N\leq n) = 1$. The existence of bounded sequential procedures reflects the existence of a termination rule of the game. For practical applications there are many reasons for introducing a forced termination rule since there are situations which force ourselves to make a decision not to continue.

11. Existence of Optimal Stopping Rules

For a g.p.d.f. ϕ of W let $\rho_o(w)$ be defined as follows:

$$\rho_o(\phi) = \inf_{d\ D}\int_{\Omega}L(w,d)\phi(w)\,d\nu(w)$$

Then $\rho_o(\phi)$ is the minimum risk from an immediate decision without any further observation when the p.d.f. of W is ϕ.

Let $X_1,X_2,\ldots,$ be a sequence of observations which have a specified joint distribution, and for $n = 1,2,\ldots,$ let $Y_n = Y_n(X_1,\ldots,X_n)$

be a random variable whose value depends on the first n-observations $X_1,\ldots,X_n$. Suppose that the statistician terminates the sampling process after having observed the values of $X_1,\ldots,X_n$, his gain is Y_n. The question is does there exist a stopping rule which maximizes the expected gain $E(Y_n)$. For a given stopping rule the expection $E(Y_n)$ exists iff the following relation is satisfied.

$$E(|Y_N|) = \sum_{n=1}^{\infty} E(|Y|N=n)P\ (N=n) < \infty.$$

We are interested in determining whether there exists a stopping rule which maximizes the expected gain $E(Y_N)$. For a given stopping rule, the expectation $E(Y_N)$ exists iff the following relation is satisfied:

$$E(|Y_N|) = \sum_{n=1}^{\infty} E(|Y_N|N=n)P(N=n) < \infty.$$

12. Sequential Statistical Analysis

After having shown that the idea of sequential decision making pertains to many real-life decision processes we are going to demonstrate now that they are also particularly useful for the theory of statistical decision.

Suppose that there is a stream of potential observations $X_1,X_2,\ldots,$ generally infinite, but sometimes finite, as in the case of sampling from a finite population. In the simplest case the variables $X_1,\ldots,X_k$, could be considered as independent observations from a fixed population with probability function

$$f_k(x_1,\ldots,x_k|\theta) = f(x_1|\theta)f(x_2|\theta)\ \ldots\ f(x_k|\theta),$$

where $f(x|\theta)$ is the probability function of the population, if we restrict to the discrete case only.

An optimum decision procedure is one that would minimize the overall expected loss (or, equivalently, maximize expected utility). One special problem that one encounters in the loss structure, and which is not considered in samples of fixed size, is the cost of obtaining observations or cost of sampling: In the sequential case the cost of sampling must be added to the loss ordinarily associated with the consequences of taking a certain action. In general, it is reasonable to let depend the cost of sampling on the state of nature, the number of observations and sometimes even upon the values of the observations. Hence, define $C(\theta,k,X_1,\ldots,X_k)$ as the cost of observations.

A very simple special kind of assumption is that all costs being proportional and independent of the state of nature that obtains, e.g.

$$C(X_1,\ldots,X_k) = kC.$$

For simplicity we will work with the latter function in what follows. The sequential nature of sampling is generally exhibited in two ways. First, the sequential nature of the experiment has to be defined, and second a termination rule has to provide a criterion at which step of the sequential process one has to stop taking further observations. Therefore, the experimental design involves two key notions: a stopping rule and a terminal decision rule. A sequential decision rule is specified for each number of observations $X_1,\ldots,X_k$ by a function d_k, so that $d_k(X_1,\ldots,X_k)$ represents a certain action after $X_1,\ldots,X_k$ observations are at hand. A class of such sequential decision functions $d_o,d_1,d_2,\ldots$, defined by $d = \{d_o,d_1(X_1),d_2(X_1,X_2),\ldots\}$ is called a terminal decision rule for a given sequential process under consideration, where d_o is one of the action when no data are at hand.

A stopping rule is associated to a terminal decision, characterized by a family of functions.

$$s = \{s_o,s_1(X_1),s_2(X_1,X_2),\ldots\}, \qquad \text{where}$$

$$s_k(X_1,\ldots,X_k) = \begin{cases} 0, & \text{if at least one further observation should be taken, given } X_1,\ldots,X_k \text{ have been observed.} \\ 1, & \text{if no more observations should be taken given that } X_1,\ldots,X_k \text{ are at hand.} \end{cases}$$

Now given a sequence of observations $X_1, X_2, \ldots$ the function s_k should be uniquely determined by the conditional probability that precisely k observations will be taken, i.e.

$$P(N = k | X_1, X_2, \ldots) = s_k(X_1, \ldots, X_k).$$

Then the probability of stopping, to be computed before the observations will be known, and given a particular stopping rule is

$$P_s(N = k) = E[s_N(X_1, \ldots, X_N)].$$

An interesting example is provided by a hypothesis testing problem to be discussed later more extensively.

Example:

Let $a<m<b$ and let $L_n = L_n(X_1, \ldots, X_n)$ denote the likelihood ratio

$$L_n = \frac{f_o(X_1) f_o(X_2) \ldots f_o(X_n)}{f_1(X_1) f_1(X_2) \ldots f_1(X_n)}$$

where f_o and f_1 are probability (or density) functions that characterize two states of nature (the hypothesis H_o and H_1, respectively). Then define the sequential decision rule

$$d_n(X_1, \ldots, X_n) = \begin{cases} \text{reject } H_o & \text{if } L_n < m \\ \text{accept } H_o & \text{if } L_n \geq m. \end{cases}$$

The stopping rule will be defined as follows: let $s_o = 0$, and for $n > 0$,

$$s_n(X_1, \ldots, X_n) = \begin{cases} 0, & \text{if } a < L_n < b \\ 1, & \text{if } L_n \leq a \text{ or } L_n \geq b. \end{cases}$$

Then the pair (s,d) defines a sequential decision procedure for the problem of testing a simple hypothesis against a simple alternative.

13. Bayesian Procedures

Suppose a sequence of observations $X_1, X_2, \ldots$ is available at cost kC for k observations. For a given sequential procedure (s,d) the total loss, including the costs of observation will be

$$\ell(\theta, d_N(X_1, \ldots, X_N)) + NC$$

with N being a random variable (its distribution determined by θ and the stopping rule s), whose value is given by the number of observations actually used in reaching a decision.

The expected loss or risk is then

$$R(\theta,(s,d)) = E\left[\ell(\theta,d_N(X_1,\ldots,X_N)) + NC\right]$$

and the Bayes risk for a prior $P(\theta)$ is obtained by further averaging with respect to that prior

$$B(s,d) = E\left[R(\theta,(s,d))\right].$$

The problem then is to find a pair (s,d) which minimizes the Bayes risk. This involves a two-stage procedure: first determine the minimizing d for each stopping rule, and then choose the stopping rule that produces the overall minimum. We can state the following result whose various technicalities do not permit a proof here for which the reader is referred to Blackwell & Girschick [5], Chap.9.

For a given stopping rule s and a given prior $P(\theta)$ the Bayes risk $B(s,d)$ is minimized by the decision rule $d^* = (d_o^*, d_1^*, \ldots, d_i^*, \ldots)$ where d_i^* is the Bayes rule applied to any fixed-sample size problem with i observations $X_1, X_2, \ldots, X_i, \ldots$

Sometimes a sequential statistical problem has only a finite number of stages, this demonstrates the similarity to problems such as restricted coin tossing games, the marriage problems and other problems discussed in the previous section where, in general, termination is enforced. In view of statistical sequential analysis the introduction of forced termination is motivated by the consideration that one may run out of data, for instance, if one takes samples without replacement from a finite population.

When the number of stages is finite, the above result can be obtained in a process of backward induction, on the basis of computational procedures as developed by dynamic programming, to determine the optimal stopping rule s^* such that the Bayes procedure (s^*,d^*) is 'best' for a given prior distribution. In order to outline the approach suppose that the stages of observation are restricted to n(and not more), corresponding to observations $X_1,\ldots,X_n$. If it happens that the Bayes procedure requires taking all n observations, the terminal decision is made according to the Bayes criterion, i.e. the posterior distribution obtained on the basis of n observed values is applied to the given loss function to obtain averages in terms of which the available actions are ordered and the optimal one is chosen. If the stopping rule s^*, on the other hand, requires at least n-k observations, the problems of whether to stop (and use the Bayes terminal rule for those observations) or to obtain more sampling data is resolved by comparing two conditional expected losses;

(i) the expected loss conditional on the previous n-k observations, and

(ii) the expected conditional loss if one takes more observations (including costs for future observations).

Having determined the optimum procedure for k = o, the computations and comparisons can be made for k = 1,2,... revealing the optimum among rules calling for at least n-1,n-2,... observations. Proceedings recursively this way the Bayes rule is completely determined via backward induction.

14. Hypothesis Testing

This section exhibits an example of testing a statistical hypothesis with two states of nature and two actions, e.g. testing a simple null against a simple alternative hypothesis. Consider the problem of testing the hypothesis that a population is normal with mean $\mu_1 = 0$ against the alternative that it is normal with mean $\mu_2 = 2$ the variance being the same in both cases: $\sigma_1^2 = \sigma_2^2 = 1$. The actions available are a_1, to accept H_o, and a_2, to reject H_o with losses as given in the following table (in which c_1 and c_2 are positive constants).

	H_o	H_1
a_1	0	c_1
a_2	c_2	0

Furthermore $P(\theta_1) = p$, $P(\theta_2) = 1-p$ be the prior distributions for states of nature θ_1 and θ_2 respectively.

If there are no observations available (no data case) the Bayes action is chosen on the basis of Bayes risks which yields in this particular case:

$$B_p(a_1) = c_1(1-p)$$

$$B_p(a_2) = c_2p.$$

The minimum of this determines the proper action. Now suppose that at least one observation from the population is being taken, at a certain cost. The cost of observation being fixed at each observation, the Bayes decision rule for a given observation is the one that minimizes the Bayes risk, e.g. the likelihood ratio rule:

$$\begin{cases} \text{Reject } H_o \text{ if } \dfrac{f_o}{f_1} < \dfrac{c_1(1-p)}{c_2p} \\ \text{Accept } H_o \text{ otherwise} \end{cases}$$

where f_i denotes the density of the observation under H_i.

Since $$\frac{f_o}{f_1} = \frac{e^{-x^2/2}}{e^{-(x-2)^2/2}} = e^{2-2x}$$

the inequality for rejecting H_o can be expressed in terms of x as

$$x > 1 - \frac{1}{2}\log\frac{c_1}{c_2} + \frac{1}{2}\log\frac{p}{1-p} = F(p).$$

For example, if $c_1(1-p) = c_2p$, then do act as though the mean is 0 if the observation is closer to 0 than to 2, otherwise act as though the mean is 2. The minimum Bayes risk for given $p = P(H_o)$ implying the rejection limit F(p) is then

$$\rho(p) = c_2p\alpha + c_1\beta(1-p) = c_2p\,[1-\Phi(F)] + c_1(1-p)\Phi(F-2)$$

where $\Phi(\cdot)$ is the probability distribution function of the standard normal distribution. The total cost of taking an observation and using the Bayes rule for that observation is the value of $\rho(p)$ plus the cost of the observation. The minimum Bayes risk, over all rules that use either no observations or one observation, can be found graphically (see Figure 1). The figure shows $\rho(p)$ with zero cost of observation, and also with two other constant costs of observation (moderate and high) and the rejection limit F(p) is shown for the two hypotheses.

The Bayes rule for the situation pictured in the figure is as follows

(i) if $p \leq c$, reject H_o with no observation,

(ii) if $c \leq p \leq d$ take the observation and use the appropriate Bayes rule,

(iii) if $p \geq d$, accept H_o with no observation.

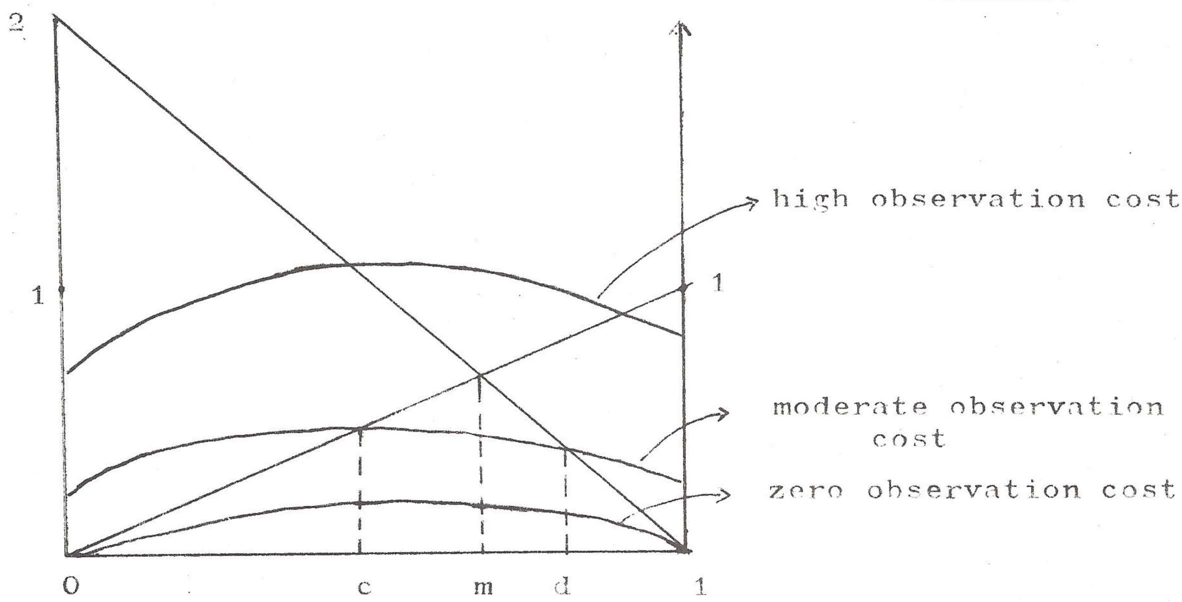

Figure 1

The argument above can also be used for the study of a sequential problem in which at most n+1 steps are permitted in order to determine the Bayes procedures in the class of procedures requiring either n or n+1 steps.

Thus suppose one considers only stopping rules that require at least n observations. If such a stopping rule is to be a part of a Bayes procedure, the associated terminal decision rule, wherever the sampling stops, is an ordinary fixed sample size Bayes decision rule. For stopping rules calling for exactly n observations, with a terminal decision at that point, a Bayesian decision rule is obtained by applying the posterior probabilities for H_o and H_1 directly to the original loss structure.

Thus if at least n and at most n+1 observations are assumed one is in almost the same position as in the example above - with the notable difference that in place of the prior p in this example one now uses the posterior probabilities based on the first n observations. Given the random vector $Z = (X_1,\ldots,X_n)$ of the first n observations and the original prior probabilities $p = P(H_o)$ and $1-p = P(H_1)$ the posterior probability can be calculated as follows:

$$\hat{p}_o = P(H_o|Z = z) = \frac{pf_o(z)}{w(z)}$$

$$\hat{p}_1 = P(H_1|Z = z) = \frac{(1-p)f_1(z)}{w(z)}$$

where

$$f_o(z) = f(z|H_o) = P(Z = z|H_o)$$

$$f_1(z) = f(z|H_1) = P(Z = z|H_1)$$

denote the probability (or density) functions of the data Z, and

$$w(z) = pf_o(z) + (1-p)f_1(z).$$

The expected posterior loss is computed for each action.

$$\text{Accept} \qquad H_o: 0 + c_1\hat{p}_1 = \frac{c_1(1-p)f_1(z)}{w(z)}$$

$$\text{Reject} \qquad H_o: c_2\hat{p}_o + 0 = \frac{c_2pf_o(z)}{w(z)}.$$

If the first is larger one should reject H_o, if the second is larger one should accept H_o.

Therefore, the critical region is the set of Z-values defined by the rule

$$\text{Reject} \qquad H_o \text{ if } \frac{f_o(Z)}{f_1(Z)} < \frac{c_1(1-p)}{c_2p}.$$

The constant $c_1(1-p)/c_2p$ cannot be calculated without knowing the losses c_1 and c_2 and the prior probabilities, however, it can be shown that the family of tests satisfying this rule includes all Bayes tests and only Bayes tests. The ration $L = f_o/f_1$ is called the likelihood ratio and tests with critical region are called likelihood ratio tests. Analogously, we can deal with the case of sequential tests.

The inequality that determines whether or not the observation is to be taken, e.g. $c<\hat{p}_o<d$ where c and d are determined as in the example above. They can be expressed in terms of the likelihood ratio, by substitution for the posterior probability $\hat{p}_o$;

$$c < \frac{wf_o(z)}{wf_o(z)+(1-w)f_1(z)} < d .$$

By rearranging this formula

$$c < 1+ \frac{1-wf_1(z)}{wf_o(z)} < \frac{1}{d}$$

or in terms of

$$L_n = f_o(z)/f_1(z)$$

with

$$a = \frac{1-w}{w} \quad \frac{c}{1-c} , \quad b = \frac{1-w}{w} \quad \frac{d}{1-d} .$$

Then the Bayes procedure is as follows:

<u>Calculate the likelihood ratio based an n observations, if it falls in the interval from a to b, take one more observation and use the Bayes test for n+1 observations. If the likelihood ratio for the n observations is less than a, reject H_o, if it exceeds b accept H_o without taking the (n+1)-st observation.</u>

15. The Sequential Likelihood (Probability) Ratio Test

This test has been first developed by A. Wald [16] and is designed for testing a simple hypothesis H_o against a simple alternative H_1. For a test to achieve error sizes α and β, define constants $a = \alpha/(1-\beta)$ and $b = (1-\alpha)/\beta$ and use these as limits for the likelihood ratio L_n computed after each observation is taken. If $L_n \leq a$, the sampling stops and the null hypothesis is rejected, if $L_n \geq b$ the sampling stops and the null hypothesis is accepted. If $a<L_n<b$, another observation is taken. The test assumes the availability of observations $X_1, X_2, \ldots$ and that at least one observation is taken. It can be shown that although the error sizes actually achieved with the test are not exactly those

specified, they are close enough for practical purposes. It can be shown that the sequential likelihood ratio test terminates with probability 1 both under H_o and H_1. Wald and Wolfowitz [17] have shown that for assigned error sizes the sequential likelihood ratio test minimizes the expected number of observations n. Define the expected number by

$$E(n) = \frac{E(\log L_n)}{EZ}$$

where Z is the logarithm of the likelhood for a single observation, and the numerator is approximately given by

$$E(\log L_n^{\cdot}) = (\log a)P(\text{rej}.H_o)+(\log b)P(\text{acc}.H_o)$$

For the case of Bernoulli population, given the observation z,

$$\log \frac{\ell(\theta_o | z)}{\ell(\theta_1 | z)} = \log \left(\frac{\theta_o}{\theta_1}\right)^z \left(\frac{1-\theta_o}{1-\theta_1}\right)^{1-z}$$

$$= z \log \frac{\theta_o(1-\theta_1)}{\theta_1(1-\theta_o)} + \log \frac{1-\theta_o}{1-\theta_1}$$

where $\ell(\cdot|z)$ denotes the likelihood function of the Bernoulli parameter to be tested for $z = 0, 1$.

Thus we have

$$EZ = z \log \frac{\theta_o(1-\theta_1)}{\theta_1(1-\theta_o)} + \log \frac{1-\theta_o}{1-\theta_1} .$$

The sequential likelihood ratio test looks very much like the Bayes sequential procedure defined in the preceeding section. The essential difference is that there is no zero stage in the sequential likelihood ratio test, with the probability of making a decision with no data at all because a specific loss structure is not assumed. It can be shown, however, that given any sequential likelihood ratio test there exist losses c_1 and c_2 and a sampling cost per observation such that the Bayes sequential test for some prior distribution is exactly the given sequential likelihood ratio test.

Various further results, extensions and generalizations on optimal stopping rules and sequential analysis can be found in recent issues of Annals of Math. Statistics (Annals of Prcbability and Annals of Statistics) and the Proc. Berkeley Symposia on Math. Statistics and Probability.

Suggested Readings

[1] R. Bellman, Dynamic Programming, Princeton Univ. Press, Princeton, N.Y.,1957.

[2] R. Bellman, Adaptive Control Processes: A Guided Tour, Princeton Univ. Press, Princeton, N.Y. 1961.

[3] R. Bellman and R. Kalaba, 'A Mathematical Theory of Adaptive Control Processes', Proc. Nat. Acad. Science, U.S.A. 45, 1959, 1288-1290.

[4] R. Bellman, R. Kalaba and D. Middleton, 'Dynamic Programming, Sequential Estimation and Sequential Detection Processes', Proc. Nat. Acad. Sci., U.S.A. 47, 1961,338-341.

[5] D. Blackwell and M.A. Girshick, Theory of Games and Statistical Decisions, New York: Wiley 1954.

[6] L. Breiman, 'Stopping Rule Problems', in Applied Combinatorial Mathematics, New York: Wiley 1964, 284-319.

[7] L. Breiman, Probability, Reading (Mass): Addison-Wesley 1968.

[8] Y.S. Chow et.al., 'Optimal Selection based on Relative Rank (The Secretary Problem)', Israel Jour. of Math. 2, 1964, 81-90.

[9] L. Doob, Stochastic Processes, New York: Wiley 1963.

[10] L.E. Dubins and L.J. Savage, How to gamble if you must: Inequalities for Stochastic Processes, New York:McGraw Hill 1965.

[11] W. Feller, An Introduction to Probability Theory and its Applications, Vol. 1, New York: Wiley.

[12] K.S. Fu, Sequential Methods in Pattern Recognition and Machine Learning, New York: Academic Press 1968.

[13] D. Griffeath and J.L. Snell, 'Optimal Stopping in the Stock Market', Annals of Probability 2, 1974, 1-13.

[14] D.V. Lindley, 'Dynamic Programming and Decision Theory', Applied Statistics 10, 1961, 39-51.

[15] D.V. Lindley, Making Decisions, London: Wiley 1970.

[16] A. Wald, Sequential Analysis, New York, Wiley 1947.

[17] A. Wald and J. Wolfowitz, 'Optimum Character of the Sequential Probability Ratio Test', Ann. Math. Statist. 19, 1948,326-339.

[18] J. Cornfield et. al., 'An Adaptive Procedure for Sequential Clinical Trials', Jour. American Statist. Assoc. 64,1969, 795-770.