CHAPTER III

BAYES SEQUENTIAL TEST PROCEDURES

3.0 Introduction

In all the problems that we have considered so far, the sample size was fixed in advance. Such a procedure does not take account the infomration being collected in the course of collecting observations. Also, the fixed sample size procedure ignores the fact that sampling is expensive and taking of each observation involves some cost. In this chapter an attempt is made to explain how a statistician can use the information collected in the course of experimentation and this information be used to take a decision. Basically the observations are being collected (may be to the extent of desired accuracy). However, once the sampling is terminated after, say n observations, the decision is taken as if being taken for a fixed sample size procedure.

In Bayesian sequential analysis, in order to decide when to stop sampling, one has to compare the posterior Bayes risk of an immediate decision with the expected Bayes risk of continuing sampling. At each stage one uses the current posterior as the basis of comparing the present with future.

In the following we have introduced sequential sample, components of sequential sampling procedure and gain due to sequential sampling procedure with the help of an example. A short discussion about the most commonly used sequential probability ratio test (SPRT) is given in (1.3.c). In Section (3.2) theorotical development of sequential decision procedure is discussed. Also, it is shown with the help of an example that a sequential Bayes procedure need not always exists. A technique of backward induction is stated as optimal sequential decision procedure and directly used for optimal bounded sequential decision procedure. Lastly an attempt is made to show that SPRT is Bayes procedure.

3.1. Preliminaries :

3.1.a : Sequential Sample :

Consider a statistical problem in which the statistician can take his observations $X_1, X_2, \ldots$ one at a time from some distribution involving a parameter W whose value is unknown. After each observation X_n he can evaluate the information he has obtained so far about W from the observations $X_1, X_2, \ldots, X_n$. And he can decide whether to terminate the sampling process or to take another observation X_{n+1}. A sample obtained in this way is called sequential sample.

Suppose a lot of large size of certain items is to be accepted or rejected based upon its quality. Let 'P' be the probability that an item is defective. A fixed sample size procedure would be : Take a random sample of 'n' items and accept the lot if the number of defective items in the sample is less than 'k' (k-specified) $0 \leq k < n$. If 'c' is the cost for each observation, then total sampling cost is n.c.

Instead of sampling all 'n' units at a time consider the following procedure. Take observations sequentially and stop for the first 'r' such that either $\sum_{i=1}^{r} d_i = k$ or $\sum_{i=1}^{r} [1-d_i] = n-k+1$ where

$$d_i = \begin{cases} 1 & \text{if i-th item is defective} \\ 0 & \text{otherwise.} \end{cases}$$

If 'N' denote the number of observations to stop the sampling

i.e. $$N = \left\{ \text{first } r : \sum_{i=1}^{r} d_i = k \text{ or } \sum_{1}^{r} (1-d_i) = n-k+1 \right\}$$

Here N is a discrete random variable taking values k, k+1,...,n. Observe that the decision reached by the fixed sample size procedure or by sequential sampling procedure are the same; the number of observations in the above sequential procedure never exceeds the fixed size 'n' ($N \leq n$).

In fact $E(N) < n$.

$$\therefore \quad C\,E(N) < n.c. \quad ; c > 0.$$

Thus by adopting the sequential procedure to reach the same decision we require less number of observations except on rare occassions.

<u>3.1.b. Components Sequential Decision Procedure</u> :

A sequential decision procedure has two components one is called sampling plan or a stopping rule. It specifies whether a decision in D should be chosen without any observation or whether at least one observation should be taken. If at least one observation is to be taken for every possible set of observed values $X_1 = x_1, \ldots, X_n = x_n$ $(n \geq 1)$ whether sampling should stop and a decision in D should be chosen without further observations or where another value X_{n+1} should be observed. a

The second component of sequential decision procedure is called a decision rule. It gives if no observations are to be taken the decision $d_o \in D$ is to be chosen. If at least one observation is to be taken the decision $\delta(x_1, \ldots, x_n) \in D$ should be chosen for each possible set of observed values $X_1 = x_1, \ldots, X_n = x_n$ after which sampling might be terminated.

With every stopping rule we associate a stopping random variable 'N' which takes the values 1,2,.... the total number of observations taken before sampling is terminated. If ultimate termination of sequential sampling procedure is guaranted; then it is called closed sampling procedures.

i.e. $P[N < \infty] = 1$.

or $P[N = \infty] = 1 - P[N < \infty] = 0$.

Gain that may result from sequential sampling :

Consider a statistical decision problem

Suppose $\Omega = \{w_1, w_2\}$ and $D = \{d_1, d_2\}$.

$$L(w_1, d_1) = L(w_2, d_2) = 0$$
$$L(w_1, d_2) = L(w_2, d_1) = b > 0.$$

Let X be a discrete random variable with $f_i(x) = P[X=x/W=w_i]$ as follows.

$$f_1(1) = 1 - \alpha \qquad f_1(2) = 0 \qquad f_1(3) = \alpha$$
$$f_2(1) = 0 \qquad f_2(2) = 1 - \alpha; f_2(3) = \alpha, \; 0 < \alpha < 1.$$

Suppose that the cost per observation is 'C'.

The prior distribution of W be specified as

$$\Pr[W = w_1] = \theta = 1 - \Pr[W = w_2], \; \theta \leq 1/2.$$

From the p.m.f of X; it follows that after an observation has been taken ;

$$P[W = w_1/X = x] = 1 \quad \text{if } x = 1$$
$$= 0 \quad \text{if } x = 2$$
$$= \theta \quad \text{if } x = 3$$

Similarly $P[W = w_2/X = x] = 0$ if $x = 1$
$$= 1 \quad \text{if } x = 2$$
$$= 1-\theta \quad \text{if } x = 3$$

Thus after an observation has been taken either the value of W becomes known or else the distribution of W remain just as it was before the observation was taken. Also, if the Bayes decision is taken either when $P(W = w_1) = 0$ or $P[W = w_1] = 1$ then expected loss will be zero. If Bayes decision is taken after 'n' observations $X_1, X_2, \ldots, X_n$ have been taken, the expected loss will be θb if $X_i = 3$ for every observation $i = 1,2,\ldots,n$ and will be zero if at least one of the observation is different from 3.

$$\Pr[\, X_i = 3,\ i = 1,2,\ldots,n] = \alpha^n \quad \text{for all } w \in W.$$

Therefore risk function $\rho(n)$ including the sampling cost, for the optimal procedure when exactly 'n' observations must be taken is,

$$\rho(n) = \theta\, b\, \alpha^n + cn$$

Assume that $\rho(1) < \rho(0)$; otherwise it is not necessary to go for an observation.

The optimum value n^* of n can be obtained by assuming n to be continious variable.

$$\frac{\partial \rho(n)}{\partial n} = 0 \Rightarrow \theta\, b\alpha^n \log \alpha + c = 0.$$

gives $$\alpha^n = \frac{c}{\theta\, b \log (1/\alpha)}$$

Therefore $$n^* = \frac{1}{\log(1/\alpha)} \left\{ \log \frac{\theta\, b \log(1/\alpha)}{c} \right\}$$

observe that $$\frac{\partial^2 \rho(n)}{\partial n} > 0.$$

Hence $\rho(n)$ by treating 'n' as continious is convex and hence the integer optimum value of 'n' must be $[n^*]$ or $[n^*]+1$ where $[x]$ is integer part of x.

Therefore risk corresponding to n* becomes

$$\rho(n^*) = \frac{c}{\log(1/\alpha)} \left[1 + \log \frac{e\, b \log(1/\alpha)}{c}\right]$$

$\rho(n^*)$ is the risk due to fixed sample size procedure with sample size n*.

Now we shall consider the following bounded sequential procedure: suppose we take n* observations sequentially and we suppose to stop sampling as soon as the value of the observation X_i is different from 3. In other words we will have to take all n* observations only when $X_i = 3$ for all $i = 1,2,\ldots,n^*-1$. Under such sequential procedure the posterior distribution when sampling terminates will be the same as it was for the fixed sample size procedure. However, the number N of observations that will be taken now is a random variable. Expected sample size is given by;

$$
\begin{aligned}
E(N) &= E(N/W = w_1) = E(N/W = w_2) \\
&= \sum_{j=1}^{n^*} j \Pr[N=j/W = w_i] \\
&= \sum_{j=1}^{n^*-1} j\, \alpha^{j-1}(1-\alpha) + n^*\, \alpha^{n^*-1} \\
&= \frac{1-\alpha^{n^*}}{1-\alpha} \\
&< n^*
\end{aligned}
$$

The total risk $\bar{\rho}$ from above sequential procedure

$$= \theta\, b\, \alpha^{n^*} + c\, E(N) < \rho(n^*)$$

$$= \frac{c}{1-\alpha} + \frac{c}{\log(1-\alpha)} \left[1 - \frac{c}{\theta b(1-\alpha)} \right]$$

Assumption $\rho(1) < \rho(0)$ is equivalent to the assumption that expression inside the bracket is positive. The procedure that we considered is nodoubt sequential but it is bounded and the number of observations that can be taken is at most n^*.

Finially we consider a sequential procedure without any upper bound on number of observations. Stopping rule is now stop as we get an observation different than 3. Decision rule is now given as; choose decision $W = w_1$ if $X_i = 1$ is observed or choose decision $W = w_2$ if $X_i = 2$ is observed. If this procedure is adopted we can always, when we terminate the sampling, choose a decision whose expected loss is zero. Now, a random variable N follows, waiting time distribution. i.e. $P[N=j] = \alpha^{j-1}(1-\alpha) \quad j=1,2,...$ Therefore, total risk for this pruely sequential procedure ρ^* is given by;

$$\rho^* = c\, E(N)$$

$$= c \sum_{j=1}^{\infty} j.\, \alpha^{j-1}(1-\alpha)$$

$$= \frac{c}{1-\alpha}$$

This shows that the sequential procedure we consider finally has a smaller total risk than other procedures that we considered before. However, it should be noted that this procedure may require more indeed many more than n* observations.

In sequential test procedure the expected number of observations required to reach a decision may be more than the size of the sample required in case fixed sample size procedure (with the same precision).

3.1.c. The Sequential Probability Ratio Test (SPRT) :

Let $X_1, X_2, \ldots,$ be a sequence of i.i.d. r.v. with common pmf(pdf) $f_w(x)$. Consider the hypothesis of testing $H_1: X \sim f_{w_1}(x)$ against a simple alternative $H_2: X \sim f_{w_2}(x)$ when observations are taken sequentially. Then

$$f_{jn}(x_1, x_2, \ldots, x_n) = \prod_{i=1}^{n} f_{w_j}(x_i), \quad j = 0,\ 1.$$

$$\text{Let } \lambda_n(x_1, x_2, \ldots, x_n) = \frac{f_{1n}(\underline{x})}{f_{2n}(\underline{x})}$$

where $\underline{x} = (x_1, x_2, \ldots, x_n)$

The sequential probability ratio test states : if at any stage of sampling

$$\lambda_n(\underline{x}) \geq A ;$$

stop and reject H_1

if $$\lambda_n(\underline{x}) \leq B ;$$

stop and accept H_1 ; and

if $\quad B < \lambda_n(\underline{x}) < A,$

continue sampling by taking another observation x_{n+1}.

Here A and B (A > B) are constants which are determined so that the test will have strength (α, β).

If N is stopping r.v.

$$\alpha = P_{w_1}\left\{\lambda_N(\underline{X}) \geq A\right\}, \quad \beta = P_{w_2}\left\{\lambda_N(\underline{X}) \leq B\right\}.$$

Take $Z_i = \log \dfrac{f_{w2}(X_i)}{f_{w1}(X_i)}$

Therefore $\log \lambda_n(\underline{x}) = \sum_{i=1}^{n} z_i = Sn$, say.

Now the SPRT may be written as follows :

At each stage compute $Sn = \sum_{i=1}^{n} Z_i$, If

$$b = \log B < Sn < \log A = a.$$

Continue by taking observation Z_{n+1}; if

$$a \leq Sn, \text{ reject } H_1$$

and if $\quad b \geq Sn$, reject H_2.

<u>Remark</u> :

With respect to any hypothesis H(not necessarily H_1 or H_2) for which $P\left\{|Z| > 0 \mid H\right\} > 0$ where $Z = \log \dfrac{f_2(X_1)}{f_1(X_1)}$

then $\quad P_H\left\{N < \infty\right\} = 1$ and

$$E_H\left\{e^{tN}\right\} < \infty \text{ for } -\infty < t < t_o,\ t_o > 0.$$

The SPRT terminates with probability 1 under both H_1 and H_2.

The most commonly used sequential procedure is SPRT introduced by Wald in the 1940S. The SPRT is designed for testing a simple null hypothesis against a simple alternative hypothesis, when sequential sample $X_1, X_2, \ldots$ is available. SPRT requires at least one observation. SPRTs are frequently used for testing problems which are more complicated than just testing simple against simple hypothesis. The most common such use is in testing H_1: $\theta \leq \theta_0$ v/s H_2: $\theta > \theta_0$ ($\theta_0 < \theta_1$). It is natural in this situation, that the problem is that of testing H_1: H_1: $\theta = \theta_0$ versus H_2: $\theta = \theta_1$ and to use the relevant SPRT. It can indeed, be shown that, if the X_i have a density with monotone likelihood ratio in θ, then the SPRT with error probabilities α_0 and α_1 gives error probabilities $\alpha_0(\theta) \leq \alpha_0$ for $\theta < \theta_0$ and $\alpha_1(\theta) \leq \alpha_1$ for $\theta > \theta_1$. In a classical sense, therefore, an SPRT is very reasonable in this situation (though probably not optimal).

Example (1.2.c) :

Let $X_1, X_2, \ldots$, be i.i.d $N(\mu, 1)$ r.v.'s. It is required to test H_1: $\mu = 0$ against H_2: $\mu = 1$.

Fixed numbers of observations required for the current test procedure of strength (α, β) is $n(\alpha, \beta)$. Test function ϕ can be given as

$$\phi(x) = \begin{cases} 1 ; & \text{if } \bar{x} > k \\ 0 ; & \text{otherwise.} \end{cases}$$

$$E_{H_1}[\phi(x)] = \alpha$$

$$E_{H_2}[\phi(x)] = \beta$$

gives

$$\int_k^\infty \frac{\sqrt{n}}{\sqrt{2\pi}} e^{-\frac{n(\bar{x})^2}{2}} d\bar{x} = \alpha \qquad (1)$$

and

$$\int_k^\infty \frac{\sqrt{n}}{\sqrt{2\pi}} e^{-\frac{n(\bar{x}-1)^2}{2}} d\bar{x} = 1-\beta \qquad (2)$$

Therefore

$$\Phi[\sqrt{n}\ (k - o)] = 1 - \alpha$$

and

$$\Phi[\sqrt{n}\ (k - 1)] = \beta$$

Let $\sqrt{n}(k) = \lambda_o$ and $\sqrt{n}\ (k - 1) = \lambda_1$

Therefore

$$\Phi(\lambda_o) = 1 - \alpha \quad \text{and } \lambda_o = Z_{1-\alpha}$$

$$\Phi(\lambda_1) = \beta \quad \text{and } \lambda_1 = Z_\beta$$

gives $\sqrt{n} = \lambda_o - \lambda_1$

$$n = (\lambda_o - \lambda_1)^2 = n(\alpha.\ \beta).$$

In particular for $\alpha = \beta = 0.05$

$\sqrt{n}\, k = \lambda_o = 1.65$ and $\sqrt{n}\,(k-1) = -1.65$ gives $n = 10.9$.

In sequential test procedure

$$A = \frac{0.95}{0.05} = 19 \quad , \quad B = \frac{0.05}{0.95} = 1/19 .$$

$Z = X - 1/2$.

So that $E_\mu(Z) = \mu - 1/2$

$$E_{H_1}(Z) = -\,1/2$$

$$E_{H_2}(Z) = 1/2.$$

$$E_\mu(Z)^2 = E_\mu\,(X - 1/2)^2 = \theta^2 - \theta + 5/4$$

$$E_{H_1} S_N \simeq \alpha \log A + (1-\alpha) \log B.$$

$$= -\,2.265.$$

$$E_{H_1}(N) \simeq 5.3.$$

$$E_{H_2} S_N \simeq 0.95 \log 19 - 0.05 \log 19$$

$$= 0.90 \log 19$$

so that $E_{H_2} N \simeq 5.3$.

In this case we see that there is some saving (on the average) if we use the sequential test over the fixed sample size test.

3.2 Sequential Decision Procedure :

Let 'S' denotes the sample space of any particular observation X_i, for $n = 1,2, \ldots$ we shall let $S^n = S \times S \times \ldots \times S$ (with n factors) be the sample space of the n observations $X_1, X_2, \ldots, X_n$ and we shall let S^∞ be the sample space of the infinite sequence of observations $X_1, X_2, \ldots$. A sampling plan in which at least one observation is to be taken can be characterised by a sequence of subsets $B_n \in S^n (n= 1,2,..]$ which have the following interpretation :

Sampling is terminated after the values $X_1= x_1,...,X_n= x_n$ have been observed if $(x_1,x_2,..,x_n) \in B_n$. Another value X_{n+1} is observed if $(x_1,...,x_n) \notin B_n$. If there is some value r for which $B_r = S^r$ or, more generally if

$P_r[(X_1,...,X_n) \notin B_n$ for $n= 1,2,...,r] = 0$ then sampling must stop after at most r observations have been taken.

The decision rule of a sequential decision procedure is characterised by a decision $d_o \in D$ and a sequence of functions $\partial_1, \partial_2,...$ with the following property : For any point $(x_1,x_2,...,x_n) \in D$, if the sampling plan specifies that an immediate decision in D is to be selected without any sampling then the decision $d_o \in D$ is chosen. If on the otherhand the sampling plan specifies that at least one observation is to be taken and if the observed values $X_1= x_1,...,X_n= x_n$ satisfy the condition that $(x_1,...,x_n) \in \{N=n\}$ then sampling is terminated and decision $\partial_n(x_1,...,x_n) \in D$ is chosen.

The values of the function ∂_n need only be specified on the subset $\{N=n\} \subset (S^n)$.

<u>Illustration</u> :

Suppose we wish to estimate the probability $\theta \in (0,1)$ of obtaining heads when a given coin is tossed. In a fixed procedure (say,n) we know that the proportion of heads is a 'good' estimator of θ.

Consider a sequential sampling procedure as follows:

Stopping rule is ; toss a coin until k-th head occurs.

Decision rule is ; after stopping the trails estimate θ ; by $\theta = \frac{K}{N}$, where 'N' the number of trails required to get 'k' heads first time; here N follows negative binomial distribution

$$P_r[N = n] = \binom{n-1}{k-1} \theta^k (1 - \theta)^{n-k} ; \; 0 \leq \theta \leq 1$$

$$n = k, k+1, \ldots$$

since $P_r[N < \infty] = \sum_{n=k}^{\infty} P_r[N = n] = 1$ above sampling procedure is closed one.

To illustrate idea of stopping regions consider example where

$$N = \left\{ r: \sum_{i=1}^{r} d_i = 3, \text{ or } \sum_{i=1}^{r} (1-d_i) = n-k+1 = 3 \right\}$$

In particular, $k = 3$ gives $B_1 = B_2 = \phi$; $B_3 = \{1,1,1\}$.

Since if S^n is the sample space corresponding to $X_1, X_2, \ldots, X_5$.

$$S^1 = \{1,0\}$$

$$S^2 = \{(1,1), (1,0), (0,1), (1,1)\}$$

$$S^3 = \{(1,1,1), (1,1,0), (1,0,1), (0,1,1), (1,0,0), (0,1,0), (0,0,1), (0,0,0)\} .$$

on similar line $B_4 = \{(0,1,1,1), (1,0,1,1), (1,1,0,1)\}$.

3.2.a A Bayes Sequential Decision Procedure :

The total risk $\rho(\theta, \delta)$ of a sequential decision procedure δ in which at least one observation is to be taken is

$$\rho(\theta, \delta) = E\left\{ L[W, \delta_N(X_1, \ldots, X_n)] + C_1 + \ldots + C_N \right\}$$

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

$$\theta(w \mid x_1, \ldots, x_n) d\ (w) d\ F_n(x_1, \ldots, x_n \mid \theta)$$

$$+ \sum_{n=1}^{\infty} (C_1 + \ldots + C_n)\ Pr\left\{N=n\right\} .$$

$$= \int_{\Omega} \left\{ \sum_{n=1}^{\infty} \int_{N=n} L[w, \delta_n(x_1, \ldots, x_n)] \prod_{i=1}^{n} f(x_i / w) \prod_{i=1}^{n} d_\mu(x_i)\ \ \theta(w)\ dv(w) \right\}$$

$$+ \sum_{i=1}^{\infty} (C_1 + \ldots + C_n)\ Pr\left\{N=n\right\} .$$

For further reference we shall denote the posterior g p.d.f. of W when $X_1 = x_1, \ldots, X_n = x_n$ where θ is the prior g p.d.f. of W is $\theta(x_1, \ldots, x_n)$.

For any g p.d.f. of W, let $\rho_o(\phi)$ be defined as,

$$\rho_o(\phi) = \inf_{d \in D} \int_{\Omega} L(w, d)\ \phi(w)\ dv\ (w) \ldots . \quad (3.2.a)$$

In this case $\rho_o(\phi)$ is the minimum risk from an immediate decision without any further observations when the g. p.d.f. of W is ϕ. If we continue to assume this against each possible g p.d.f. ϕ which arises during sampling process, there is Bayes decision in D which actually yields the minimum risk $\rho_0(\phi)$.

A Bayes sequential decision procedure, or an optimal sequential decision procedure, is a procedure ∂ for which the risk $\rho(\theta, \partial)$ is minimized. If sampling is to be terminated after the values $x_1, \ldots, x_n$ have been observed, then the decision $\partial_n(x_1, \ldots x_n)$ that is chosen should be Bayes against the posterior g p.d.f. $\theta(x_1, \ldots, x_n)$ of W. Suppose that a given sequential decision procedure ∂ will not always lead to a Bayes decision when sampling is terminated. Then a sampling procedure ∂^*, which has the same stopping rule as ∂ but which does always lead to a Bayes decision, will have the property that $\rho(\theta, \partial^*) \leq \rho(\theta, \partial)$.
Therefore, whenever a decision in D is chosen after sampling has been terminated, that decision is a Bayes decision against the posterior distribution of W. Hence it is not necessary to mention explicitly the decision rule in any procedure. Hence for any such procedure ∂ which specifies that at least one observation is to be taken satisfy following relation

$$\rho(\theta, \partial) = E\left\{\rho_0[\theta(X_1, \ldots, X_N)] + C_1 + \ldots + C_N\right\}.$$

And for the procedure ∂_0 which specifies that an immediate decision in D should be chosen without any observations satisfy,

$$\rho(\theta, \delta_o) = \rho_o(\theta).$$

<u>Example : (3.2.a)</u> :

Let $X_1, X_2, \ldots,$ is a sequential sample from a $N(\theta,1)$ density, and that it is desired to estimate θ under loss

$$L(\theta,a,n) = (\theta - a)^2 + \sum_{i=1}^{n} C_i$$

where $\quad C_i = [2\,i(i+1)]^{-1}$

It can be easily shown that

$$\sum_{i=1}^{n} C_i = \sum_{i=1}^{n} \frac{1}{2i(i+1)} = \frac{1}{2}\left(1 - \frac{1}{n+1}\right)$$

Hence

$$L(\theta,a,n) = (\theta - a)^2 + \frac{1}{2}\left(1 - \frac{1}{n+1}\right)$$

If now the prior is $N(0,1)$, so that posterior distribution is $N(\tilde{\theta}(x), (n+1)^{-1})$

where $\tilde{\theta}(x) = n\bar{x} = \mu_n$ (say).

The corresponding Bayes risk (Ref. example 1.2.e).

$$\rho^{*}_{\delta(x)}(n) = \frac{1}{n+1} + \frac{1}{2}\left(1 - \frac{1}{n+1}\right)$$

$$= \frac{1}{2}\left(1 + \frac{1}{n+1}\right),$$

which is independent of parameter.

This is decreasing in n, so that it never pays to stop sampling, another observation always lowers the posterior Bayes risk. Hence no proper Bayes procedure can therefore exists.

Remark :

A sequential Bayes procedure need not always exists as the above example shows.

3.2.b Backward Induction :

A sequential decision procedure ∂ is bounded if there is positive integer n such that $Pr(N \leq n) = 1$. In this section we shall consider problems in which there is fixed upper bound n on the number of observations.

When statistician employs the technique of backward induction, he begins by considering the final stage of observation and he then works backward to the first stage of observation. If X_n is observed, then the fixed limit on the number of observation makes it compulsory to terminate the sampling and to choose a decision in D. Hence, if sampling has not been stopped earlier, the statistician must determine at this final stage whether to choose a decision in D based on the values of $X_1,\ldots,X_{n-1}$ or to take exactly one more observation and then choose a decision in D. The optimal decision will usually depend on the values of $X_1,\ldots,X_{n-1}$ that have been observed. After determining this he can begin to work backward. Knowing the values of the observations $X_1,\ldots,X_{n-2}$ at the next to last stage, he can now decide whether it is worthwhile to take the next observation X_{n-1}. By working backward in this way to the first stage, the statistician determines, for each possible value of X_1, the optimal continuation throughout the remaining stages. He can evaluate the risk from observing X_1

and can compare this risk with the risk from the immediate choice of a decision in D without any observations. These comparision at the first and subsequent stages, determine the optimal sequential decision procedure.

Bellman (1957 a) has called the construction of optimal procedure by backward induction as principle of optimality. It is stated as : The optimal sequential decision procedure must satisfy the requirement that if, at any stage of the procedure, the values $X_1 = x_1, \ldots, X_j = x_j$ $(j < n)$ have been observed, then the continuation of the procedure must be the optimal sequential decision procedure for the problem where the prior distribution of W is $\theta(x_1, \ldots, x_j)$ and the maximum number of observations that can be taken is $n-j$.

3.2.c Optimal Bounded Sequential Decision Procedure :

Assume that there is fixed cost c per observation. If ϕ is any g p.d.f. of W and if the value $\rho_o(\phi)$ is defined by (3.2.a) then

$$E\left\{\rho_o[\phi(X)]\right\} = \int_{\mathcal{R}} \rho_o[\phi(x)]\, d\, F_1(x/\phi).$$

Suppose the values $X_1=x_1, \ldots, X_{n-1}=x_{n-1}$ have been observed and the statistician must decide whether to choose a decision in D without another observation or to observe X_n. The risk from choosing a decision in D without another observation is $\rho_o(\theta_{n-1})$. If the value $X_n=x$ is observed and

a decision in D is then chosen, the risk will be $\rho_o[\theta_{n-1}(x)]$.
Hence the expected total risk from observing X_n and then choosing a decision in D is $E\left\{\rho_o[\theta_{n-1}(x)]\right\} + C$.
The final choice in the optimal procedure will be as follows :

If $\rho_o(\theta_{n-1}) < E\left\{\rho_o[\theta_{n-1}(X)]\right\} + C$,

then sampling should be terminated and X_n should not be observed. If this inequality is reversed then X_n should not be observed. If the two sides of the relation are equal, then the risk is the same whether sampling is terminated or continued and we shall assume sampling is terminated in this case.

Let $\rho_1(\phi)$ denote the risk from the optimal procedure in which not more than one observation is taken and the g p.d.f. of W is ϕ.

$$\rho_1(\phi) = \min\left\{\rho_o(\phi),\ E[\phi(X)] + C\right\}.$$

In particular $\rho_1(\theta_{n-1})$ is the risk from the optimal continuation after the values $X_1=x_1,\ldots,X_{n-1}=x_{n-1}$ have been observed.

Now we shall move backward one stage, Suppose $X_1=x_1,\ldots,X_{n-2}=x_{n-2}$ have been observed. The risk from a decision in D without any further observations is $\rho_o(\theta_{n-2})$.
On the ohter hand, if the values of X_{n-1} is observed and

this value is x the posterior g p.d.f. becomes $\theta_{n-2}(x)$ and the risk from optimal continuation at the stage is $\rho_1[\theta_{n-2}(x)]$. The total risk from observing this value becomes $E\left\{\rho_1(\theta_{n-2}(X)]\right\}$ +C. Hence the optimal procedure can be described as follows :

If $\rho_o(\theta_{n-2}) \leq E\left\{\rho_1[\theta_{n-2}(X)]\right\} + C,$

then sampling should be terminated and X_{n-1} should not be observed.

Now let $\rho_2(\phi)$ the risk from optimal procedure when sampling is terminated after not more than two observations.

$$\rho_2(\phi) = \min\left\{\rho_o(\phi),\ E[\ \rho_1(\phi(X))] + C\right\}.$$

In particular, it follows from the above discussion that $\rho_2(\theta_{n-2})$ is the risk from the optimal continuation after the values $X_1=x_1,\ldots,X_{n-2}=x_{n-2}$ have been observed.
In general, for any g p.d.f. of W, let $\rho_o(\phi)$ be defined by (3.2.a) and let $\rho_1(\phi), \rho_2(\phi),\ldots, \rho_n(\phi)$ be defined recursively by :

$$\rho_{j+1}(\phi) = \min\left\{\rho_o(\phi),\ E[\rho_j(\phi(X))] + C\right\} \quad j = 0,1,\ldots,n-1.$$

It is assumed that each expectation in the above equation exists.

Despite the simple theoretical nature of the optimal procedure, the computation of the function ρ_n for a value of n larger than 4 or 5 is extremly difficult and time-consuming. The example presented here will provide further insight into the properties of optimal sequential decision procedure.

Example (3.2.c) :

Suppose that a sequential random sample $X_1, X_2, \ldots$ can be taken from the Bernoulli distribution for which the value of the parameter W is unknown. Suppose also that either $W = 1/3$ or $W = 2/3$ and that the statistician must decide which value of W is correct. Therefore, it is assumed that the decision space $D = \{d_1, d_2\}$ and that the loss function as specified in the table below. Suppose further that each observation costs 1 unit and that the prior distribution of W is specified by the probability $\theta = \Pr(W = 1/3) = 1 - \Pr(W = 2/3)$. We shall compute $\rho_0(\theta)$, $\rho_1(\theta)$ and $\rho_2(\theta)$.

loss table

	d_1	d_2
$W=1/3$	0	20
$W=2/3$	20	0

If $\theta \leq 1/2$, d_2 is Bayes decision.

If $\theta \geq 1/2$, d_1 is Bayes decision.

$$\rho_0(\theta) = \begin{cases} 20\,\theta & \text{for } 0 \leq \theta \leq 1/2 \\ 20(1-\theta) & \text{for } 1/2 \leq \theta \leq 1. \end{cases} \qquad (1)$$

Therefore $\rho_0(\theta) = \rho_0(1-\theta)$ for $0 \leq \theta \leq 1$.

Hence $\rho_j(\theta) = \rho_j(1-\theta)$ for $0 \leq \theta \leq 1$.

From the symmetry of the problem, we shall compute $\rho_1(\theta)$ and $\rho_2(\theta)$ only for the values of $0 \leq \theta \leq 1/2$. For $x = 0,1$ let $\theta(x)$ denote the posterior probability that $W = 1/3$ when the value of single observation X is x.

$$\theta(1) = \frac{\theta}{\theta+2(1-\theta)} \quad \text{and} \quad \theta(0) = \frac{2\theta}{2\theta+(1-\theta)} \tag{2}$$

Hence

$$\left.\begin{array}{ll} \rho_o[\theta(1)] = 20\ \theta(1) & \text{for} \quad 0 \leq \theta \leq 1/2. \\ \rho_o[\theta(0)] = \begin{cases} 20\ \theta(0) & \text{for} \quad 0 \leq \theta \leq 1/3. \\ 20\ [1-\theta(0)] & \text{for} \quad 1/3 < \theta \leq 1/2 \end{cases} \end{array}\right| \tag{3}$$

The marginal distribution of any observation X is

$$\Pr(X=1) = 1/3\ \theta + 2/3(1-\theta) = 1 - \Pr(X=0) \tag{4}$$

Therefore

$$E\left\{\rho_o[\theta(X)]\right\} = \rho_o[\theta(1)]\Pr(X=1) + \rho_o[\theta(0)]\Pr(X=0)$$

$$= \left\{\begin{array}{ll} 20\ \theta & \text{for} \quad 0 \leq \theta \leq 1/3 \\ 20/3 & \text{for}\ 1/2 < \theta \leq 1/2 \end{array}\right| \tag{5}$$

Since C = 1, it now follows that

$$\rho_1(\theta) = \min\left\{\rho_o(\theta),\ E[\ _o(\theta(X)))] + 1\right\}.$$

$$= \left\{\begin{array}{ll} 20\ \theta & \text{for}\ 0 \leq \theta \leq \frac{23}{60} \\ \frac{23}{3} & \text{for}\ \frac{23}{60} < \theta \leq 1/2 \end{array}\right| \tag{6}$$

From (2) it can be shown that

$$\theta(1) \leq \frac{23}{60} \quad \text{iff,} \quad \theta \leq \frac{46}{83},$$

$\theta(0) \leq \frac{23}{60}$ iff, $\theta \leq \frac{23}{97}$, and

$\theta(0) \geq \frac{37}{60}$ iff, $\theta \geq \frac{37}{83}$,

By making use of symmetry of the function $\rho(1)$ and equations (4) and (5) we get

$$E\left\{\rho_1[\theta(X)]\right\} = \begin{cases} 20\ \theta & \text{for } 0 \leq \theta \leq \frac{23}{97}, \\ \frac{83\ \theta + 23}{g} & \text{for } \frac{23}{97} < \theta \leq \frac{37}{83}, \\ \frac{20}{3} & \text{for } \frac{37}{83} < \theta \leq 1/2\ . \end{cases} \quad (7)$$

Hence $\rho_2(\theta) = \min\left\{\rho_o(\theta),\ E(\theta(X))] + 1\right\}$.

$$= \begin{cases} 20\ \theta & \text{for } 0 \leq \theta \leq \frac{32}{97}\ , \\ \frac{83.\ \theta + 32}{g} & \text{for } \frac{32}{97} < \theta \leq \frac{37}{83}, \\ \frac{20}{3} & \text{for } \frac{37}{83} < \theta \leq 1/2. \end{cases} \quad (8)$$

The functions ρ_o, ρ_1 and ρ_2 are sketched.

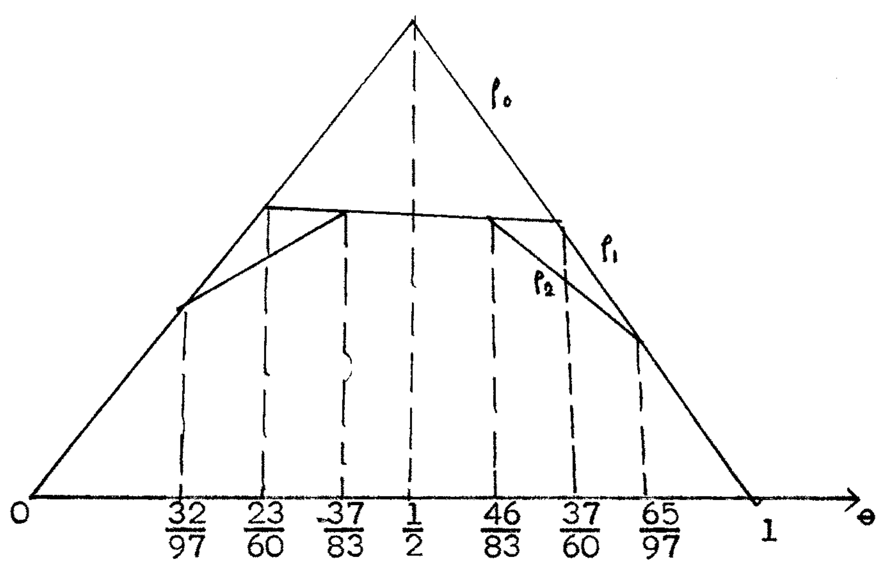

The computations of ρ_n for a large value of n is more difficult because the number of linear segments defining the function increases as n becomes larger. The computations will require more detail if symmetry that has been assumed in this example is deleted.

3.3 The Sequential Probability Ratio Test as Bayes Procedure :

Consider a sequential decision problem in which the parameter space is $\Omega = \{w_1, w_2\}$ and the decision space is $D = \{d_1, d_2\}$. Loss function is given by

	d_1	d_2
w_1	0	λ_1
w_2	λ_2	0

$\lambda_1, \lambda_2 > 0.$

Suppose that ~~ta~~ the sequential random sample $X_1, X_2, \ldots$ can be taken at a cost of C unit per observation. For $i = 1,2$ let f_i denote the conditional g.p.d.f. of any observation X when $W = w_i$. $\Pr(W = w_1) = \theta$. The risk $\rho_0(\theta)$ from choosing a decision immediately is

$$\rho_0(\theta) = \min\{\lambda_1\theta, \lambda_2(1 - \theta)\} .$$

Let Δ' denote the calss of all sequential decision procedure δ which requires that at least one observation should be taken

$$\rho'(\theta) = \inf_{\delta \in \Delta'} \rho(\theta,\delta) , \quad \text{for } 0 \leq \theta \leq 1.$$

The Bayes risk $\rho^*(\theta) = \min\left\{\rho_o(\theta), \rho'(\theta)\right\}$.

ρ' is concave continious function on the interval $0 \leq \theta \leq 1$ (Ref. 1.2.b), and every procedure $\partial \in \Delta'$ involves a sampling cost of at least C units. $\rho'(0) = \rho'(1) = C$ and $\rho'(\theta) \geq C$ for all values of θ. $(0 \leq \theta \leq 1)$. The functions ρ_o and ρ' are sketched below :

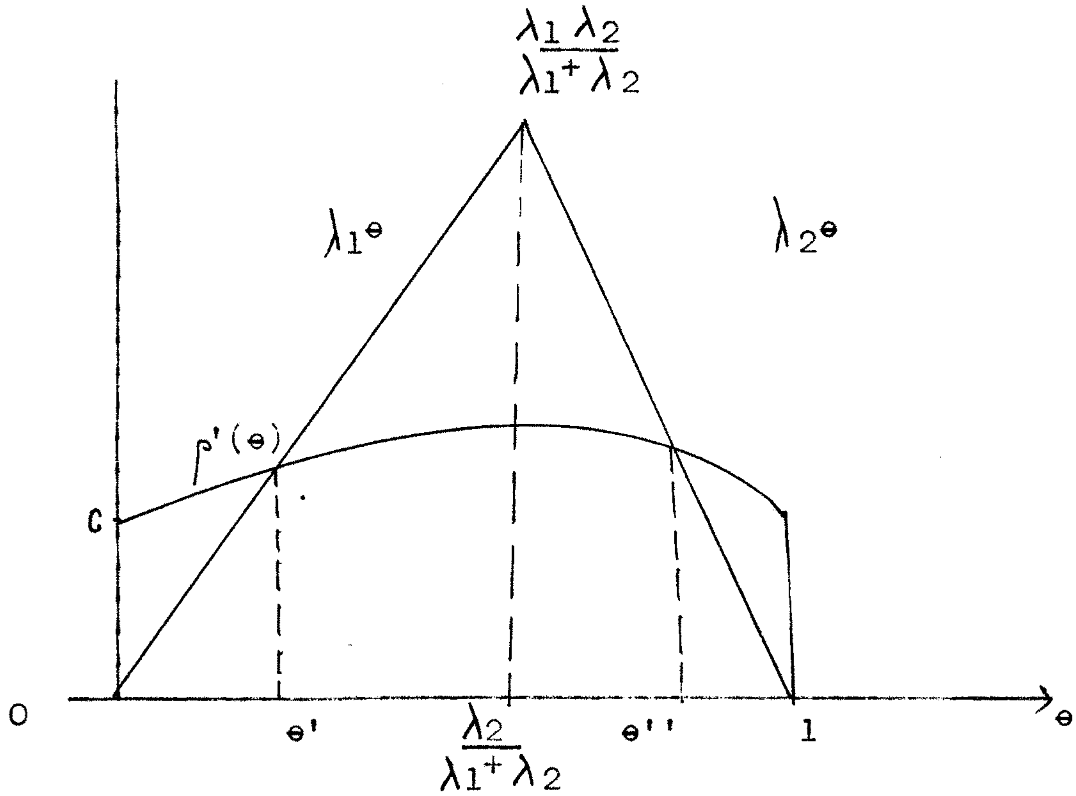

Let Σ^* be the set of all values of θ at which it is optimal to terminate sampling.

Therefore $\Sigma^* = \left\{\theta : \rho_o(\theta) \leq \rho'(\theta)\right\}$

suppose that

$$\rho'\left(\frac{\lambda_2}{\lambda_1 + \lambda_2}\right) = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2} \qquad (1)$$

This relation is satisfied in the figure.

Σ^* is the union of intervals $0 \leq \theta \leq \theta'$ and $\theta'' \leq \theta \leq 1$ where θ' and θ'' follows the following equations.

$$\lambda_1 \theta' = \rho'(\theta') \quad \text{and} \quad \lambda_2(1-\theta') = \rho'(\theta'').$$

If the e inequality (1) is not satisfied then it can be seen from figure that Σ^* is the entire interval $0 \leq \theta \leq 1$. In this case, regardless of the prior distribution of W, it is never worthwhile to take any observations. An interpretation of the above can be as follows:

Smallness of θ ($\theta < \theta'$) or largeness of θ ($\theta > \theta''$) indicates that the prior knowledge itself is so strong that it does not require additional observations.

The set Σ^* characterises the Bayes sequential decision procedure. Suppose that the prior probability θ lies in the interval $\theta' < \theta < \theta''$. It follows that the first observation should be taken. If $\theta(x_1,\ldots,x_n)$ is the posterior probability that $W = w_1$ after the values $X_i = x_i$ ($i = 1,\ldots,n$) have been observed, then optimal procedure is to continue taking observations whenever the following relation is satisfied.

$$\theta' < \theta(x_1,\ldots,x_n) < \theta'' \qquad (2)$$

Inequality (2) can be written as

$$\frac{1}{\theta''} < 1 + \frac{(1-\theta)\ f_2(x_1),\ldots\ldots,f_2(x_n)}{\theta f_1(x_1)\ \ldots\ldots\ f_1(x_n)} < \frac{1}{\theta'}$$

which reduces to

$$\frac{(1-\theta'')\ \theta}{\theta''(1-\theta)} < \frac{f_2(x_1),\ldots\ldots,f_2(x_n)}{f_1(x_1),\ldots\ldots,f_1(x_n)} < \frac{(1-\theta')\ \theta}{\theta'\ (1-\theta)}$$

Define $$A = \frac{(1-\theta')\theta}{\theta'(1-\theta)}$$

$$B = \frac{(1-\theta'')\theta}{\theta''(1-\theta)}$$

Therefore, continue sampling if

$$B < \frac{f_2(x_1),\ \ldots\ ,\ f_2(x_n)}{f_1(x_1),\ \ldots\ ,\ f_1(x_n)} < A.$$

If $$\frac{f_2(x_1),\ \ldots\ ,\ f_2(x_n)}{f_1(x_1),\ \ldots\ ,\ f_1(x_n)} < B$$

take decision d_1

If $$\frac{f_2(x_1),\ \ldots\ ,\ f_2(x_n)}{f_1(x_1),\ \ldots\ ,\ f_1(x_n)} > A$$

take decision d_2.

Since $\theta' < \theta < \theta''$ it follows that $A < 1$ and $B > 1$.

The problem of finding an optimal sequential decision procedure has now been reduced to the problem of finding an optimal choice of the constants A and B.

A sequential decision procedure of the type just described, is called a sequential probability-ratio test. In problem in which both the parameter space Ω and the decision space D, have exactly two points, the optimal sequential decision procedure is either to choose a decision immediately without any observations or to use a sequential probability-ratio test.