A

SEQUENTIAL ESTIMATION WITH DIRICHLET PROCESS PRIORS

Thomas S. Ferguson[1]

Department of Mathematics
University of California at Los Angeles
Los Angeles, California, U.S.A.

I. INTRODUCTION AND SUMMARY

The purpose of this article is to investigate two simple sequential nonparametric problems from a Bayesian viewpoint using a Dirichlet process prior, the estimation of a distribution function on the real line, and the estimation of the mean of a distribution on the real line.

Let F denote a distribution function on the real line, $\mathbb{R}$, and let $X_1, X_2, \ldots$ represent a sample from F. In the problem of estimating F, the statistician is to choose a distribution function $\hat{F}$ with loss measured by the function

$$(1.1) \qquad L(F, \hat{F}) = \int (F(x) - \hat{F}(x))^2 dW(x)$$

where W is some finite measure (weight function) on $\mathbb{R}$. In the problem of estimating the mean of F, the statistician chooses a point $\hat{\mu} \in \mathbb{R}$ and loses the amount

$$(1.2) \qquad L(F, \hat{\mu}) = (\mu - \hat{\mu})^2$$

where $\mu = \int x dF(x)$ is assumed to exist for this problem. There is a positive cost $c > 0$ each time the statistician looks at a new observation. After each observation, the statistician must decide whether to take another observation or to stop sampling and choose an estimate.

[1]This research was supported by the National Science Foundation under Grant No. MCS 80-02732.

ISBN 0-12-307501-7

We attempt to find Bayes solutions to these problems when the prior distribution of F is the Dirichlet process, $\mathcal{D}(\alpha)$, where α is a given finite non-null measure on $\mathbb{R}$. See [3] for the elementary facts about this prior that are used below. In problem (1.1), it is assumed that α and W have no common atoms. It is useful to use the distribution function form of the measure α, in which α has the representation $\alpha = MF_0$, where $M = \alpha(\mathbb{R})$ is a positive number, and that F_0 is a distribution function. Thus, let $F \in \mathcal{D}(MF_0)$. Then $\mathcal{E}F = F_0$. Moreover, if $X_1, X_2, \ldots, X_n$ is a sample from F, then the posterior distribution of F given $X_1, \ldots, X_n$ is $\mathcal{D}((M+n)F_n)$ where

$$F_n = \frac{M}{M+n} F_0 + \frac{n}{M+n} \hat{F}_n, \tag{1.3}$$

where $\hat{F}_n$ is the sample distribution function,

$$\hat{F}_n(x) = \frac{1}{n} \sum_1^n I_{[X_i, \infty)}(x) \tag{1.4}$$

and $I_S(x)$ represents the indicator function of the set S.

We recommend the use of the 1- or 2-stage look-ahead rules for these problems. The theorems of this paper give a partial justification for this recommendation. In Theorems 1 and 3, it is seen that the k-stage look-ahead modified rules are easily computed for k as large as $[(M+n+2)/2]$, and the resulting rule is seen to be comparable to the 1-stage look-ahead rule. In Theorems 2 and 4, conditions are given under which the 1-stage look-ahead rule is optimal.

II. ESTIMATING A DISTRIBUTION FUNCTION

Let $F \in \mathcal{D}(MF_0)$ and let $X_1, \ldots, X_n$ be a sample of fixed size n from F. In estimating F with loss (1.1), the Bayes estimate is $\mathcal{E}(F|X_1, \ldots, X_n) = F_n$ of (1.3) and the minimum Bayes risk is

$$(2.1) \qquad \int \text{Var}\,(F(x)\,|\,X_1,\ldots,X_n)dW(x)$$

$$= \frac{1}{M+n+1} \int F_n(x)(1-F_n(x))dW(x).$$

In the extension of this problem to the sequential case, it is well-known (see, for example [2] Theorem 7.1) that once we decide to stop, the Bayes terminal estimate is the same as for the fixed sample size problem, so that if we stop after observing $X_1,\ldots,X_n$, the Bayes terminal estimate is F_n. Thus we are only concerned with finding the Bayes stopping rule.

It is only in very exceptional parametric cases that an optimal stopping rule can be found explicitly. Usually, an approximation to the optimal rule is sought, for example the rule optimal among those limited to a fixed bounded number N of observations. However, it is still only in special parametric cases -- when there exists a small dimensional sufficient statistic whatever the sample size -- that such a rule can be computed for moderate N in a reasonable length of time. In the problems considered here, sufficiency does not reduce the dimensionality of the observations, and the backward induction method necessary to compute such rules involves the approximation of functions of N variables.

As a practical matter, there are some very good suboptimal rules. These are the k-stage look-ahead rules.

The k-*stage look-ahead rule* (k-*sla*) *is the rule that at each stage stops or continues according to whether the rule optimal among those taking at most* k *more observations stops or continues.* Usually, and it is so for our problems, the 1-sla is trivial to compute. The 2-sla involves a 1-dimensional numerical integration to be performed at each stage, while the 3-sla (about as complex as one would like to consider) involves a 2-dimensional numerical integration at each stage.

The 1-sla is not only easy to compute; it is reasonably good for estimation problems, and the 2-sla or 3-sla is generally quite good. In fact, Bickel and Yahav [1] have shown that for sufficiently regular estimation problems with quadratic loss, the 1-sla is asymptotically pointwise optimal as $c \to 0$.

A simplification may be made in these rules. This is due to the fact that in our problems, if the k-stage look-ahead rule tells you to continue, it is optimal to continue, for there is at least one rule that continues (but stops in at most k stages) and gives a smaller expected cost plus loss than stopping at once. This property suggests a simplification of the 2-sla: use the 1-sla until it tells you to stop and then use the 2-sla. Similarly, the 3-sla is equivalent to: use the 1-sla until it tells you to stop, (then the 2-sla until it tells you to stop), and then use the 3-sla.

Let us evaluate the 1-stage look-ahead rule, sometimes called the "myopic" rule. First, we discover what it requires us to do at the very first stage before looking at any observations. If we stop and make a terminal decision without sampling, we lose (2.1) with n = 0, namely

$$\frac{1}{M+1} \int F_0(1-F_0)\,dW. \tag{2.2}$$

If we take one observation X_1 and then stop, we lose, conditional on X_1, the amount $c+(M+2)^{-1} \int F_1(1-F_1)\,dW$. On the average, we expect to lose

$$c+(M+2)^{-1}\mathcal{E}\int F_1(1-F_1)\,dW = c + \frac{M}{(M+1)^2} \int F_0(1-F_0)\,dW. \tag{2.3}$$

This computation is easily made using (1.3), (1.4) and $\mathcal{E}\hat{F}_1 = F_0$, which holds since the marginal distribution of X_1 is F_0. Therefore, the 1-stage look-ahead rule calls for stopping without taking any observations if (2.2) $\le$ (2.3), or equivalently, if

$$\int F_0(1-F_0)\,dW \le c(M+1)^2. \tag{2.4}$$

After observing $X_1,\ldots,X_n$, the 1-stage look-ahead rule calls for stopping if the updated posterior distribution satisfies an updated version of (2.4). That is, *the* 1*-stage look-ahead rule calls for stopping after the first* n *for which*

$$\int F_n(1-F_n)dW \leq c(M+n+1)^2. \tag{2.5}$$

Since the left side is bounded above by $W(\mathbb{R})/4$, and the right side increases to infinity as n tends to infinity, the 1-stage look-ahead rule eventually calls for stopping, and bounds on the maximum sample size can be found.

When the 1-sla finally tells us to stop, how can we tell if it is optimal to stop? Two partial answers to this question are given in Theorems 1 and 2 below. The first theorem involves the notion of the k-stage modified sequential decision problem, due originally to Magwire [4]. The only difference between this problem and the ordinary sequential decision problem is that if stage k is reached the terminal loss is set equal to zero, in which case it is certainly optimal to stop.

The k*-stage look-ahead modified rule* (k*-slam*) *is the rule that at each stage stops or continues according to whether the rule optimal* for the modified problem in which the terminal loss is set to zero if you take k more observations stops or continues.

In contrast to the k-sla which never takes too many observations, the k-slam never takes too few. *If the* k*-slam calls for stopping, it is optimal to stop,* for any rule that continues costs at least as much as the best rule for the modified problem which is the cost of stopping without taking any further observations. The optimal Bayes rule therefore lies somewhere between the k-sla and the k-slam. Bickel and Yahav [1] have shown that the 1-slam is asymptotically pointwise optimal as $c \to 0$ for sufficiently regular hypothesis testing problems.

If the 1-sla is myopic, the 1-slam is very myopic. It compares the expected terminal loss of taking no further observations with the cost of one more observation. It does not depend on the distribution of that observation. In our problem, the 1-slam calls for stopping without taking any observations if

$$\frac{1}{M+1} \int F_0(1-F_0)dW \leq c. \tag{2.6}$$

The general 1-*slam therefore calls for stopping after the first* n *for which*

$$\int F_n(1-F_n)dW \leq c(M+n+1). \tag{2.7}$$

This differs from the 1-sla only in the term (M+n+1) replacing $(M+n+1)^2$, but this is a very big difference if M or n is large. Since we expect the 1-sla to be reasonably good, this means we expect the 1-slam to be poor. However, the following theorem shows that we can completely describe the k-slam for $k \leq M/2 + 1$. This leads to a modified rule that is comparable to the 1-sla.

This theorem is based on two simple lemmas.

LEMMA 1.

$$\int F_n(1-F_n)dW = \frac{M}{M+n} \int F_0(1-F_0)dW + \frac{n}{M+n} \int \hat{F}_n(1-\hat{F}_n)dW + \frac{Mn}{(M+n)^2} \int (F_0-\hat{F}_n)^2 dW. \tag{2.8}$$

This identity is easily checked by straightforward calculation using (1.3).

LEMMA 2. *If* j *and* k *are non-negative integers and if* $M \geq 2(k-1)$ *then*

$$(k-j)(M+j)(M+j+1) \leq kM(M+1). \tag{2.9}$$

Proof. The result is obvious with equality if j = 0. So assume j > 0. From $(k-j)(M+j)(M+j+1)-kM(M+1) = kj(2M+1)+kj^2 - j(M+j)(M+j+1)$, we see that (2.9) holds if and only if

$$k(2M+j+1) \leq (M+j)(M+j+1). \tag{2.10}$$

Using the assumption $M \geq 2(k-1)$ or $k \leq M/2+1$, we will be finished if we show $(M/2+1)(2M+j+1) \leq (M+j)(M+j+1)$ or $(M+2)(M+(j+1)/2) \leq (M+j)(M+j+1)$. This is true with equality if $j = 1$, while for $j \geq 2$, $(M+2) \leq (M+j)$ and $M+(j+1)/2 \leq M+j+1$, completing the proof.

THEOREM 1. *If* $M \geq 2(k-1)$, *then the* k-*slam calls for stopping at stage* 0 *if and only if* $\int F_0(1-F_0)dW \leq k(M+1)c$.

Proof. The k-slam calls for stopping if and only if $\frac{1}{M+1} \int F_0(1-F_0)dW \leq c + \phi_1$ where, inductively, $\phi_k = 0$ and for $j = k-1,\ldots,1$.

$$\phi_j = \mathcal{E}[\min(\frac{1}{M+j+1} \int F_j(1-F_j)dW,\ c + \phi_{j+1})|X_1,\ldots,X_{j-1}].$$

Suppose that $\int F_0(1-F_0)dW = k(M+1)c$; we will show $\phi_j = (k-j)c$ for $j = 1,\ldots,k-1$. An application of (2.8) shows

$$\int F_j(1-F_j)dW \geq \frac{M}{M+j} \int F_0(1-F_0)dW \quad \text{a.s.}$$

$$= \frac{M(M+1)kc}{M+j}.$$

Inductively, using Lemma 2,

$$\frac{1}{M+k} \int F_{k-1}(1-F_{k-1})dW \geq \frac{M(M+1)kc}{(M+k)(M+k-1)} \geq c \quad \text{so} \quad \phi_{k-1} = c$$

$$\vdots$$

$$\frac{1}{M+2} \int F_1(1-F_1)dW \geq \frac{M(M+1)kc}{(M+2)(M+1)} \geq (k-1)c \quad \text{so } \phi_1 = (k-1)c.$$

Thus, if $\int F_0(1-F_0)dW = k(M+1)c$, the k-slam is indifferent. For larger values of c it is uniquely optimal to stop, and for smaller values of c, it is uniquely optimal to continue, completing the proof.

Since it is easy to use, we prefer the k-slam to the 1-slam where $k = \lfloor M/2+1 \rfloor$ ($\lfloor x \rfloor$ is the greatest integer less than or equal to x.) In fact, since M increases with n, we may let k depend on n and use the $\lfloor (M+n+2)/2 \rfloor$-slam: *stop after the first* n *for* *which*

$$\int F_n(1-F_n)dW \le c(M+n+1)\lfloor (M+n+2)/2 \rfloor . \tag{2.11}$$

This is much closer to the 1-sla (2.5).

This theorem is of use in the evaluation of the (k+1)-slam where $k \le (M+1)/2$ with the numerical computation of only one integral since after one observation M increases by 1 and the k-slam is then immediately computable.

We now consider the possible optimality of the 1-stage look-ahead rule. We have noted that if the 1-sla calls for continuing at a certain n, it is optimal to continue at that n. When the loss function is bounded as is the case here, there is a simple sufficient condition for the converse. *If the 1-stage look-ahead rule calls for stopping at a certain* n, *and if for almost all futures starting from that time the 1-stage look-ahead rule calls for stopping, then it is optimal to stop at that* n. This condition is well-known (see for example, [2] Theorems 7.5 and 7.6), but it is rather strong, and useful only in special cases. In our case, if (2.5) holds almost surely for $n = 0,1,2,\ldots$, then it is optimal to stop without taking any observations. The following theorem shows that this can occasionally occur.

THEOREM 2. *Conditions* (2.5) *hold almost surely for* $n = 0,1,2,\ldots$, *provided*

(i) $\int F_0(1-F_0)dW \le c(M+1)^2$

(ii) $M \max(\int F_0 dW, \int (1-F_0)dW) \le 4c(M+1)^3$, *and*

(iii) $M \ge (\sqrt{5} - 1)/2 = .618\ldots$.

Proof. Condition (2.4) is exactly condition (i). To evaluate the other inequalities (2.5) we need

$$\int F_n(1-F_n)dW = (M+n)^{-2}\int (MF_0+n\hat{F}_n)(M(1-F_0)+n(1-\hat{F}_n))dW \tag{2.12}$$

$$= (M+n)^{-2}\{M^2 \int F_0(1-F_0)dW + M\sum_{i=1}^{m}[\int_{-\infty}^{X_i} F_0dW + \int_{X_i}^{\infty}(1-F_0)dW]$$

$$+ n^2 \int \hat{F}_n(1-\hat{F}_n)dW\}.$$

It is easy to see that for all x,

$$\int_{-\infty}^{x} F_0dW + \int_{x}^{\infty}(1-F_0)dW \le \max(\int F_0dW, \int (1-F_0)dW) \tag{2.13}$$

so that using (i) and (ii),

$$\int F_1(1-F_1)dW \le (M+1)^{-2}\{M^2 \int F_0(1-F_0)dW$$

$$+ M\max(\int F_0dW, \int (1-F_0)dW)\}$$

$$\le (M+1)^{-2}\{M^2c(M+1)^2 + 4c(M+1)^3\} = c(M+2)^{-2} \quad \text{a.s.}$$

It remains to be shown that (2.5) holds a.s. for $n = 2,3,\ldots$. The summation in the last term on the right side of (2.12) may be bounded as follows.

$$n^2 \int \hat{F}_n(1-\hat{F}_n)dW \le \frac{n^2}{4}\int dW \le \frac{n^2}{2}\max(\int F_0dW, \int (1-F_0)dW).$$

Therefore, from (i), (ii), and (2.12),

$$\int F_n(1-F_n)dW \le \frac{c(M+1)^2}{(M+n)^2}\{M^2 + 4n(M+1) + 2n^2\frac{M+1}{M}\} \quad \text{a.s.}$$

Thus (2.5) holds a.s. if

$$(M+1)^2\{M^2+4n(M+1) + 2n^2\frac{M+1}{M}\} \le (M+n+1)^2(M+n)^2.$$

This reduces to the inequality

$$0 \le n^3+2n^2(2M+1) + n[4M^2-5-2M^{-1}]-2(M+1)(3M+2). \tag{2.14}$$

As a function of n, the right side of (2.14) has positive slope for $n \geq 2$ when (iii) is satisfied. Therefore, (2.5) holds a.s. for all $n \geq 2$ if (2.14) holds a.s. for $n = 2$; that is, if

$$0 \leq 2M^2 + 6M + 2 - 4M^{-1} = 2(M+2)(M^2 + M - 1)M^{-1}.$$

This is satisfied if $M^2 + M - 1 \geq 0$ which is exactly condition (iii), completing the proof.

This theorem gives conditions for the 1-stage look-ahead rule to be optimal at the initial stage. If (i) is satisfied, so that the 1-stage look-ahead rule calls for stopping, then it is optimal to stop without taking any observations provided (ii) and (iii) are satisfied. At subsequent stages, F_0 and M are updated to F_n and M+n so that condition (iii) becomes automatically satisfied.

COROLLARY. *If, after* $n \geq 1$ *observations have been taken, the 1-stage look-ahead rule calls for stopping, it is optimal to stop provided*

$$(2.15) \qquad (M+n)\max\{\int F_n dW, \int (1-F_n)dW\} \leq 4c(M+n+1)^3.$$

How likely is it, when (2.5) becomes satisfied for the first time, that (2.15) will be satisfied also? In a practical case where we expect to take many observations before stopping, we expect approximate equality in (2.5) when we do stop. In such a case, (2.15) becomes approximately

$$(2.16) \qquad \max\{\int F_n dW, \int (1-F_n)dW\} \leq 4 \int F_n(1-F_n)dW$$

(assuming M is small compared to n). Since F_n converges almost surely to the true F, this gives us an indication of how likely it will be that the 1-sla will stop only when it is optimal to stop.

As an example, let $F_0(x) = x$ on $[0,1]$, the uniform distribution, and suppose $dW = dx$ on $[0,1]$. Then

$\int F_0(1-F_0)dW = \int_0^1 x(1-x)dx = 1/6$, and $\int F_0 dW = \int (1-F_0)dW = 1/2$. The 1-sla calls for stopping without taking any observations if $1/6 \le c(M+1)^2$. This condition implies condition (ii) of Theorem 1, since $4c(M+1)^3 \ge 4(M+1)/6 \ge M/2$. Therefore, the 1-sla is optimal at the initial stage if $M \ge .618\ldots$. If the 1-sla calls for at least one observation, and if M is large, it is likely that condition (2.15) will be satisfied a.s. when we stop. This is because (2.16) is satisfied for $n = 0$ (i.e. $1/2 \le 2/3$), and, if M is large, it is likely that F will be close to the uniform. It is possible, though, that the 1-sla will call for stopping after the first observation and that it will not be optimal to stop. In cases when (2.15) is not satisfied when you stop, it is best to check the 2-sla.

We mention that the results of this section carry over to estimating a distribution function in d-dimensions. The statements of Theorems 1 and 2 are the same except that condition (ii) of Theorem 2 must be replaced by the more general statement

(ii') $$M \max_z \{\int [F_0(x)(1-I_{[z,\infty)}(x)) + (1-F_0(x))I_{[z,\infty)}(x)]dW(x)\} \le 4c(M+1)^3.$$

Here, $[z,\infty)$ represents the set of points $x \in \mathbb{R}^d$ such that the inequality $z \le x$ holds for each coordinate. In general for $d \ge 2$, (ii') does not reduce to (ii). The proofs were written to apply to this generalization unchanged.

Let us look briefly at the difficulties involved in computing the 2-sla and 3-sla. If we stop without sampling, we lose (2.2). If we take one observation X_1, and use the best 1-stage look-ahead procedure from there, we pay c plus the minimum of (2.2) and (2.3) updated by one observation, namely,

$$(2.17)\qquad c+\min\left(\frac{1}{M+2}\int F_1(1-F_1)dW, c+\frac{M+1}{(M+2)^2}\int F_1(1-F_1)dW\right)$$

$$= c+\frac{1}{M+2}\int F_1(1-F_1)dW + \min\left(0, c-\frac{1}{(M+2)^2}\int F_1(1-F_1)dW\right).$$

Thus, for the 2-sla, we stop without taking observations if (2.2) is smaller than the expectations of (2.17), that is if

$$(2.18)\qquad \frac{1}{(M+1)^2}\int F_0(1-F_0)dW \le c-\mathcal{E}\ \max\left(0, \frac{1}{(M+2)^2}\int F_1(1-F_1)dW-c\right).$$

This differs from the 1-sla of (2.4) by the subtraction of the nonnegative term on the right. The expectation in this term can be computed using numerical approximations without too much difficulty.

Similarly, for the 3-sla, if we take one observation and use the best 2-stage procedure from there, we pay c plus the minimum of (2.2) and the expectation of (2.17) updated by one observation. The rule corresponding to (2.18) is more completed and can be evaluated in general only by iterated numerical integration.

III. ESTIMATING THE MEAN OF A DISTRIBUTION

Let $F \in \mathcal{D}(\alpha)$ where $\alpha = MF_0$ as before, and consider estimating the mean μ of F with loss (1.2) based on a sample of fixed size n, $X_1,\ldots,X_n$ from F. It is assumed that the second moment of F_0 is finite. The Bayes estimate of μ is

$$(3.1)\qquad \mu_n = \frac{M}{M+n}\mu_0 + \frac{n}{M+n}\bar{X}_n$$

where $\mu_0 = \int x dF_0$ is the prior estimate of the mean and $\bar{X}_n = n^{-1}\Sigma_1^n \bar{X}_i$. The minimum conditional Bayes risk is

$$(3.2)\qquad \mathrm{Var}(\mu|X_1,\ldots,X_n) = \sigma_n^2/(M+n+1)$$

where σ_n^2 is the variance of the distribution F_n. It may be computed in a formula analogous to (2.8)

$$\sigma_n^2 = \int (x-\mu_n)^2 dF_n(x) = (M\sigma_0^2 + ns_n^2 + \frac{Mn}{M+n}(\bar{X}_n - \mu_0)^2)/(M+n) \tag{3.3}$$

where σ_0^2 is the variance of the distribution F_0 and $s_n^2 = n^{-1}\Sigma_1^n(X_i - \bar{X}_n)^2$.

Let us evaluate the 1-sla for the sequential problem. If we stop without sampling, our expected loss is $\sigma_0^2/(M+1)$. If we take one observation and stop, our conditional expected loss plus cost given X_1 is $c + \sigma_1^2(M+2)^{-1}$, so on the average we expect to lose

$$c + \frac{1}{M+2}\mathcal{E}\sigma_1^2 = c + \frac{1}{(M+1)(M+2)}(M\sigma_0^2 + \frac{M}{M+1}\sigma_0^2)$$
$$= c + \frac{M}{(M+1)^2}\sigma_0^2.$$

At the first stage, the 1-sla calls for stopping if

$$\sigma_0^2 \le c(M+1)^2. \tag{3.4}$$

Hence, the 1-stage look ahead rule is: *stop after the first* n *for which*

$$\sigma_n^2 \le c(M+n+1)^2. \tag{3.5}$$

On the other hand, the 1-slam calls for stopping without taking any observations if $\sigma_0^2 \le c(M+1)$, so that the general 1-slam is: *stop after the first* n *for which*

$$\sigma_n^2 \le c(M+n+1). \tag{3.6}$$

The difference between the stopping rules (3.5) and (3.6) is great if M or n is large. However, we can narrow the gap by considering the k-slam as in the following analogue of Theorem 1.

THEOREM 3. *If* $M \geq 2(k-1)$, *then the* k-*slam calls for stopping if and only if* $\sigma_0^2 \leq k(M+1)c$.

The proof of this theorem is exactly the same as Theorem 1 with $\int F_n(1-F_n)dW$ replaced by σ_n^2 and the use of Lemma 1 replaced by (3.3).

As before, we may use the k-slam and allow k to depend on n: *the* $\lfloor\frac{M+n+2}{2}\rfloor$-*slam calls for stopping after the first* n *for which*

$$\sigma_n^2 \leq (M+n+1)\lfloor\frac{M+n+2}{2}\rfloor c. \tag{3.7}$$

The difference between the stopping rules (3.5) and (3.7) is not great, and the optimal stopping rule lies between them.

As before, if the 1-sla calls for continuing at a certain n, it is optimal to continue at that n. If the 1-sla calls for stopping at a certain n, and if for almost all futures proceeding from that n the 1-stage look-ahead calls for stopping, it is optimal to stop at that n. A glance at (3.5) shows that if F_n were the distribution function of an unbounded distribution, (3.5) could not hold almost surely, since s_n^2 and $\bar{X}_n$ may take unboundedly large values. To obtain a result similar to Theorem 1, we must therefore assume that the distribution of X_1 is bounded. In the following theorem, we take the distribution of $|X_1|$ to be bounded by 1.

THEOREM 4. *Assume that* α *gives all of its mass to the interval* $[-1,1]$. *Then, conditions* (3.5) *are satisfied a.s. for* $n = 0,1,2,\ldots$, *provided*

(i) $\sigma_0^2 \leq c(M+1)^2$

(ii) $M(1+|\mu_0|)^2 \leq c(3M+4)(M+1)^2$, *and*

(iii) *if* $M|\mu_0| < 1$, *then* $1 + \frac{M}{2}\ {}^2_0 \leq c(3M^2 + 10M + 9)$.

Proof. Condition (3.5) with n = 0 is exactly condition (i). With n = 1, (3.5) is satisfied since

$$\sigma_1^2 = \frac{M}{M+1}\sigma_0^2 + \frac{M}{(M+1)^2}(X_1-\mu_0)^2$$

$$\leq \frac{M}{M+1}\sigma_0^2 + \frac{M}{(M+1)^2}(1+|\mu_0|)^2 \quad \text{a.s.}$$

$$\leq Mc(M+1) + c(3M+4) = c(M+2)^2$$

using only (i) and (ii). To show (3.5) for $n \geq 2$, it is sufficient to show

$$s_n^2 + \frac{M}{M+n}(\bar{X}_n-\mu_0)^2 \leq c\,\frac{M+n}{n}(M+n+1)^2 - c\,\frac{M(M+1)^2}{n} \quad \text{a.s.}$$

$$= c[n^2 + n(3M+2)+(M+1)(3M+1)]$$

using (i). Since all $|X_i| \leq 1$ a.s., we have

$$s_n^2 = \frac{1}{n}\sum_1^n X_i^2 - \bar{X}_n^2 \leq 1 - \bar{X}_n^2 \quad \text{a.s.}$$

so that it is sufficient to show

$$(3.8) \qquad 1-\bar{X}_n^2 + \frac{M}{M+n}(\bar{X}_n-\mu_0)^2 \leq c[n^2+n(3M+2)+(M+1)(3M+1)] \text{ a.s.}$$

We seek the maximum of the left side of this inequality over variation of $\bar{X}_n$ in [-1,1]. The function $1-x^2 + M(x-\mu_0)^2/(M+n)$ is quadratic in x with maximum at $x = -M\mu_0/n$.

Case 1. $M|\mu_0| < 1$. In this case, the maximum occurs at a point inside the interval [-1,1] for all n. Thus it is sufficient to show (3.8) with $\bar{X}_n$ replaced by $-M\mu_0/n$, namely

$$1 + \frac{M}{n}\mu_0^2 \leq c[n^2 + n(3M+2)+(M+1)(3M+1)]$$

for all $n \geq 2$. But the left side is decreasing in n and the right side is increasing in n, so it is sufficient to show this inequality for n = 2. For n = 2, this reduces to (iii).

Case 2. $M|\mu_0| \geq 1$. For $n \geq M|\mu_0|$, the left side of (3.8) is bounded above by $1 + \frac{M}{n}\mu_0^2$ as in Case 1. For $1 \leq n \leq M|\mu_0|$, the maximum of (3.8) as $\bar{X}_n$ varies in [-1,1] occurs at ± 1, and the maximum value is $M(1+|\mu_0|)^2/(M+n)$. These bounds for the left side of (3.8) are decreasing in n, and the right side is increasing in n, so it is sufficient to show this inequality for n = 1. But this is exactly (3.5) which we have already shown to be satisfied. This completes the proof.

The conditions of Theorem 4 are sufficient for the optimality of the decision to stop without taking any observations. Unfortunately, conditions (ii) and (iii) are somewhat stronger than their counterparts of Theorem 2 although (iii) could be weakened slightly. This makes it less likely that the theorem will be of use, partly because, no doubt, the 1-stage look-ahead rule is less likely to be optimal.

The uniform distribution on [-1,1] is a critical case. Let F_0 be the distribution function of $\mathcal{U}(-1,1)$, so that $\mu_0 = 0$ and $\sigma_0^2 = 1/3$. The 1-stage look-ahead rule calls for stopping without taking any observations if

$$\text{(3.9)} \qquad \frac{1}{3} \leq c(M+1)^2.$$

Condition (ii), that $M \leq c(3M+4)(M+1)^2$, and condition (iii), that $1 \leq c(3M^2 + 10M + 9)$, are then automatically satisfied. Thus, it is optimal to stop without taking observations if M and c are such that (3.9) is satisfied.

Thus, for the uniform distribution (ii) and (iii) are satisfied when (i) is, but barely. If the variance σ_0^2 were any smaller than 1/3, (ii) would not follow from (i) except for special values of M. Also, if the mean μ_0 were not exactly zero, (ii) would not follow in general from (i). The same goes for condition (iii).

If c is such that the sample size will be at least moderate before the 1-stage look-ahead calls for stopping, then we expect approximate equality in (i) with σ_0^2 and M updated. With M large, in order that (i) imply (ii) and (iii), it is necessary that $\sigma_0^2 \geq 1/3$ and μ_0 be close to zero. All in all, it is likely that the 1-stage look-ahead rule is optimal when stopping occurs only if the true distribution is U-shaped. More recourse to the 2-stage or 3-stage look-ahead rule will probably be required for this problem than for the problem of estimation of a distribution function.

The 2-stage look-ahead rule may be computed as in the earlier problem, and is found to have a similar form: stop without taking any observations if

$$\frac{\sigma_0^2}{(M+1)^2} \leq c - \mathcal{E}\max\left(0, \frac{\sigma_1^2}{(M+2)^2} - c\right).$$

REFERENCES

[1] Bickel, P. J., and Yahav, J. A. (1965). Asymptotically pointwise optimal procedures in sequential analysis. *Proc. Fifth Berkeley Symp. Math. Statist. Prob. I*, 401-413.

[2] Ferguson, T. S. (1967). *Mathematical Statistics: A decision theoretic approach.* Academic Press, New York and London.

[3] Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist. 1*, 209-230.

[4] Magwire, C. A. (1953). Sequential decisions involving the choice of experiments. Ph.D. thesis, Stanford University.