The Annals of Statistics
1985, Vol. 13, No. 3, 1129–1139

A BAYESIAN NONPARAMETRIC SEQUENTIAL TEST FOR THE MEAN OF A POPULATION[1]

BY MURRAY K. CLAYTON

University of Guelph

We may take observations sequentially from a population with unknown mean θ. After this sampling stage, we are to decide whether θ is greater or less than a known constant ν. The net worth upon stopping is either θ or ν, respectively, minus sampling costs. The objective is to maximize the expected net worth when the probability measure of the observations is a Dirichlet process with parameter α. The stopping problem is shown to be truncated when α has bounded support. The main theorem of the paper leads to bounds on the exact stage of truncation and shows that sampling continues longest on a generalized form of neutral boundary.

1. Introduction and summary. We are permitted to sample sequentially from a population with unknown mean θ, paying a positive cost c per observation. We can choose to stop sampling at any time. Upon stopping we are to decide whether θ is greater or less then a known constant ν. The net worth is θ or ν, respectively, minus sampling costs. We wish to determine sequential strategies which maximize the expected net worth.

Let X_i be the ith observation taken and let X be a generic observation. We assume that the sequence $X, X_1, X_2, \cdots$, is nonterminating, even though we might only take a finite number of observations. Given the unknown probability measure P, we assume $X, X_1, X_2, \cdots$ are independent and identically distributed with probability measure P. Given P, the unknown mean θ is the expectation of X: $\theta = \int X\, dP$.

As in many sequential problems, we are faced with a conflict between the desire to gain information about the unknown θ through sampling and the desire to avoid excessive sampling costs. The issue then is to decide when to stop sampling and make the terminal decision, and so we are faced with a typical "stopping problem." (See Chow et al., 1971, for related problems.)

To model the unknown P, we take the nonparametric Bayesian approach of Ferguson (1973) and assume that P is a Dirichlet process with parameter α. We make some comments regarding notation and refer the reader to Ferguson (1973) for further details regarding the Dirichlet process.

We take α to be a bounded nonnull measure on the reals, $\mathbb{R}$, and the Borel sets. To ensure that the problem is well defined, assume that α has a finite first moment. Following the notation of Ferguson (1973), let $\mathscr{P}$ denote the probability measure on $[0, 1]^{\mathbb{R}}$ and its associated σ-field generated by the cylinder sets which

Received January 1984; revised August 1984.

[1] Research supported by a Natural Sciences and Engineering Research Council of Canada Postgraduate Scholarship and NSERC Grant A8725.

AMS 1980 *subject classifications.* Primary 62L15; secondary 62C10.

Key words and phrases. Dirichlet process, sequential decisions, optimal stopping.

determines the distribution of P, and let $\mathscr{E}$ denote expectation with respect to $\mathscr{P}$. The phrase "almost sure" is with respect to $\mathscr{P}$. Let $M = \alpha(\mathbb{R})$ and define F by $F(x) = \alpha(-\infty, x]/M$, $x \in \mathbb{R}$. F is the normalized distribution form of α and is the mean of P, in the sense that $F(x) = \mathscr{E}P[y{:}y \le x]$. M may be interpreted as the "prior weight." (See Ferguson, 1973, page 223.) The prior mean, μ, for an observation is $\mathscr{E}\theta$, and can be computed as the mean of F. We shall often denote the parameter α by the equivalent quantity MF, referring to either as the Dirichlet parameter. When necessary, we shall denote the dependence of $\mathscr{P}$, $\mathscr{E}$, M, F, and μ on α by a subscript: $\mathscr{P}_\alpha$, $\mathscr{E}_\alpha$, etc.

The Dirichlet parameter α summarizes the prior "information" about P. Conditional on observations $X_1, X_2, \cdots, X_k$, the measure P is a Dirichlet process with parameter $\alpha + \sum_1^k \delta_{X_i}$, where δ_x gives mass one to x (Ferguson, 1973, Theorem 1). The conditional expectation of a function g of X given $X_1, \cdots, X_k$ can be computed from $\alpha + \sum_1^k \delta_{X_k}$ by integrating against the normalized form of

$$\alpha + \textstyle\sum_1^k \delta_{X_k}\text{: } \mathscr{E}_\alpha(g(X) \mid X_1, \cdots, X_n) = M(M+n)^{-1} \int g(x)\, dF(x) + (M+n)^{-1} \sum_1^n g(X_i).$$

We shall sometimes denote this conditional expectation by $\mathscr{E}_\gamma g(X)$ where $\gamma = \alpha + \sum_1^k \delta_{X_i}$. In particular, we have $\mathscr{E}_\alpha X = \mu$.

Let N be a stopping rule; that is, N is a map from the set of all sequences of observations $(X_1, X_2, \cdots)$ into $\{0, 1, 2, \cdots, \infty\}$. The event $\{N = n\}$ is interpreted to mean that sampling stops at the nth stage, and is measurable with respect to the σ-field generated by $X_1, \cdots, X_n$. If $X_1, X_2, \cdots, X_n$ have been observed and $N(X_1, \cdots, X_n) = n$, then conditional on $X_1, \cdots, X_n$ the expected net worth is

$$r(n) = [\mathscr{E}_\alpha(X \mid X_1, \cdots, X_n) \vee \nu] - nc.$$

N is *optimal* if $\mathscr{E}_\alpha r(N)$ is maximized; such rules are shown to exist by DeGroot (1970, Section 12.9). Let N^* be such an optimal rule. In Clayton (1983), we show $N^* < \infty$ almost surely. Of greater interest is the setting in which the decision problem is truncated, that is, when there exists an n' such that $N^* \le n'$ almost surely. The smallest such n', denoted n^*, is the "exact stage of truncation" (Ray, 1965, page 860). The existence of n', or equivalently, of n^* allows us to determine N^* using backward induction (DeGroot, 1970, page 277).

In this paper we show that n^* exists when α has bounded support. The proof of this is given in Section 3; Section 2 contains some preliminary results. An example is given in Section 4.

It is an open question whether n^* exists for arbitrary α with unbounded support. In Clayton (1983, Proposition 6.2.8), we show that $n^* = 0$ when c is sufficiently large, namely, when $c \ge \mathscr{E}_\alpha(X \vee \nu) - (\mu_\alpha \vee \nu)$. We also show that $n^* \le 1$ when $\mu = \nu$ and $c \ge \frac{1}{2} M(M+1)^{-1}\mathscr{E}_\alpha[(X-\nu) \vee 0]$ (Clayton, 1983, Proposition 6.2.9). This latter result may be interpreted to say that $n^* \le 1$ for sufficiently small M. We conjecture that, in fact, n^* exists for all $M > 0$ and all F.

The stopping problem of this paper has been discussed in various forms by Moriguti and Robbins (1962), Lindley and Barnett (1965), Ray (1965) and others.

The first two assume that the X_i's are conditionally Bernoulli with parameter p. A priori, p has a beta distribution with parameters t and u. In the sequel, we refer to this as the "Bernoulli-beta(t, u)" model. It is a special case of that based on a Dirichlet process and is obtained by setting $\alpha = t\delta_0 + u\delta_1$. Lindley and Barnett (1965) show $n^* \le \nu(1 - \nu)/c - 1$, and use backward induction to determine optimal strategies for various ν, c, t, and u. Ray (1965, Section 4) considers a more general model than the Bernoulli, in which the X_i's have a density in the one-dimensional exponential family. He shows for such models that sampling continues longest on the "neutral boundary": the set of points $(x_1, \cdots, x_n) \in \mathbb{R}^n$ such that $E(X \mid X_1 = x_1, \cdots, X_n = x_n) = \nu$. In Theorem 3.1 we derive a result similar to Ray's for the model based on the Dirichlet process.

The appeal of the Dirichlet process model follows from a result of Proposition 3 of Ferguson (1973): with respect to the topology of convergence in distribution, the support of P is the set of all distributions whose support is contained in the support of α. We thus have an essentially nonparametric approach allowing us to model those situations in which the responses can take on values in a specified set, the support of α.

There are three other papers related to this problem which deserve attention. Chernoff (1968) and Lechner (1962) both examine continuous-time versions of the stopping problem, and Pratt (1966) describes the asymptotic behavior of the continuation region for $c \to 0$ when the Bernoulli-beta(t, u) model is used.

Lindley and Barnett (1965) and Pratt (1966) discuss the stopping problem in terms of its risk. We prefer to follow Moriguti and Robbins (1962) and speak instead of its worth (or utility). Of course, by using a simple transformation the two problems can be seen to be equivalent. Moreover, both formulations are equivalent to the setting of Chernoff (1968) who discusses the problem as one in which we are to decide whether or not θ is greater than zero, with absolute error loss in the terminal decision.

2. Preliminary results.

In this section we give some preliminary results to be used in proving Theorem 3.1. To begin, let $V_0(\alpha, \nu)$ denote the expected worth when no observations are taken: $V_0(\alpha, \nu) = \mu_\alpha \vee \nu$. Define

$$\Delta(\alpha, \nu) = \mathscr{E}_\alpha V_0(\alpha + \delta_X, \nu) - V_0(\alpha, \nu). \tag{2.1}$$

Δ is the expected gain in taking one free observation and stopping, over stopping immediately. The next lemma lists two intuitive and useful properties of Δ. A formal proof is straightforward and is omitted.

LEMMA 2.1. *For given α and ν,*

(i) $\Delta(\alpha, \nu) \ge 0$.

(ii) *If α' is defined by $\alpha'(x) = \alpha(x + a)$, and if $\nu' = \nu - a$, then $\Delta(\alpha', \nu') = \Delta(\alpha, \nu)$, for any constant a.*

Suppose $\Delta(\alpha, \nu) \le c$. If we are limited to taking at most one observation, we should not sample at all. This notion is extended by the following theorem.

THEOREM 2.1. *There is an optimal stopping rule N^* which is truncated if there exists an n' such that*

$$n \geq n' \text{ implies } \Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, \nu) \leq c \text{ almost surely.} \tag{2.2}$$

If n_0' is the smallest such n', then $n^* \leq n_0'$.

PROOF. The theorem follows from Ray (1965, Theorem 3.2), a straightforward modification of DeGroot (1970, Theorem 12.13.1), and the fact that $\lim_{n\to\infty}\mathscr{E}_\alpha V_0(\alpha + \sum_1^n \delta_{X_n}, \nu) = \mathscr{E}_\alpha(X \vee \nu)$. This limit is a consequence of the strong law of large numbers. □

It is quite possible to have $n^* < n_0'$. For example, let $\alpha = \frac{1}{2} M(\delta_0 + \delta_1)$ and $\nu = \frac{1}{2}$. Using the equivalent Bernoulli-beta($M/2$, $M/2$) model, Ray shows n^* to be approximately $n_0'/3$. (Equation (5.27) of Ray (1965) should read $a^1 \sim a^0/3$, not $a^1 \sim a^0/9$.)

If we are to use Theorem 2.1 to show that the stopping problem is truncated, then it is necessary to impose the restriction that α have bounded support, as the following example shows.

EXAMPLE 2.1. Let α be the standard normal distribution function (so $M = 1$), let $\nu = 0$, and let c be given. For n even, suppose $x_1, \cdots, x_n$ are as follows: $x_i = -8nc$ if i is odd and $x_i = 8nc$ if i is even. Then $\Delta(\alpha + \sum_1^n \delta_{x_i}, 0) \geq 4n^2c/[(n+1)(n+2)] > c$ if $n \geq 2$. □

Restricting α to have bounded support is sufficient to ensure that the stopping problem is truncated. One proof of this is given in Clayton (1983, Theorem 6.3.2). In the next section we give a theorem which results in another proof and which may also be used to derive a bound for n^*.

We end this section by defining some notation. Let F be as in Section 1. Following DeGroot (1970, Sections 11.8–9) define the transform T_F by

$$T_F(s) = \mathscr{E}_F(X \vee s) - s. \tag{2.3}$$

We list some properties of T_F:

LEMMA 2.2. *Under the model of Section* 1,

(i) *$T_F(s)$ is continuous and nonincreasing in s.*

(ii) *$T_F(s) = \mu - s$ iff $\mathscr{P}_F(X < s) = 0$ and $T_F(s) = 0$ iff $\mathscr{P}_F(X > s) = 0$.*

(iii) *$T_{\delta_y}(s) = (y - s) \vee 0$.*

(iv) *If $G = \eta F_1 + (1-\eta)F_2$, $0 \leq \eta \leq 1$, then $T_G = \eta T_{F_1} + (1-\eta)T_{F_2}$.*

The first two properties listed above are stated for a more general setting than ours in DeGroot (1970, Sections 11.8–9). The last two are easily proved.

We can use (2.3) to reexpress $\Delta(\alpha, \nu)$ in a form which we use in the next section:

$$\Delta(\alpha, \nu) = (M+1)^{-1}T_F[(M+1)(\nu - \mu) + \mu] + \nu - (\mu \vee \nu). \tag{2.4}$$

3. Bounds on Δ and n*. The principle result of this section is a theorem which is analogous to Theorem 4.1 of Ray (1965) and which coincides with his result for the Bernoulli-beta(t, u) model. Before we state the theorem, we define some notation. If $n \geq 1$, let

$$\tilde{k}(n, \nu) = [M(\mu - \nu) + n(b - \nu)]/(b - a) \tag{3.1}$$

and define $k^*(n, \nu)$ as follows: $k^* = 0$ if $\tilde{k} < 0$, $k^* = n$ if $\tilde{k} > n$, and $k^* = \tilde{k}$ otherwise. When appropriate we shall delete the arguments of $\tilde{k}$ and k^* from the notation.

THEOREM 3.1. *Suppose there exist* a *and* b *such that* $\mathscr{P}_\alpha(a \leq X \leq b) = 1$. *For all* x_i *such that* $a \leq x_i \leq b$, $i = 1, \cdots, n$,

$$\Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, \nu) \leq \Delta(\alpha + k^*(n, \nu)\delta_a + (n - k^*(n, \nu))\delta_b, \nu).$$

We make a series of remarks before proving the theorem.

Let $\beta(k) = \alpha + k\delta_a + (n - k)\delta_b$. For $\tilde{k}$ defined by (3.1), $\mathscr{E}_{\beta(\tilde{k})}(X) = \nu$, and if $a < \nu < b$, (2.4) and Lemma 2.2(iv) give

$$\begin{aligned}\Delta(\beta(\tilde{k}), \nu) &= \frac{M}{(M+n)(M+n+1)} T_F(\nu) + \frac{\tilde{k}}{(M+n)(M+n+1)} T_{\delta_a}(\nu) \\ &\quad + \frac{n-\tilde{k}}{(M+n)(M+n+1)} T_{\delta_b}(\nu) \\ &= \frac{M}{(M+n)(M+n+1)} T_F(\nu) + \frac{M(\nu-\mu)+n(\nu-a)}{(M+n)(M+n+1)} \cdot \frac{b-\nu}{b-a}.\end{aligned} \tag{3.2}$$

The second equality in (3.2) holds by Lemma 2.2 (iii) and the definition of $\tilde{k}$. Let $f(n) = \Delta(\beta(\tilde{k}(n, \nu)), \nu)$. If n is sufficiently large, $k^* = \tilde{k}$, and in that case Theorem 3.1 and the definition of $\beta(\tilde{k})$ imply $\Delta(\alpha + \sum_1^n \delta_{x_i}, \nu) \leq f(n)$. But for sufficiently large n, $f(n)$ decreases to 0 in n, so there exists an n' such that $n \geq n'$ implies $f(n) \leq c$. Combining this with Theorem 2.1 proves that the problem is truncated when the support of α is bounded. Moreover, if n_0 is the smallest n' such that $n \geq n'$ implies $\Delta(\beta(k^*(n, \nu)), \nu) \leq c$, then $n^* \leq n_0' \leq n_0$ by Theorem 2.1 and the fact that $\Delta(\beta(k^*(n, \nu)), \nu)$ is an upper bound for $\Delta(\alpha + \sum_1^n \delta_{x_i}, \nu)$.

Note that if $\tilde{k} \in [0, n]$, then $\mathscr{E}_{\beta(k^*)}(X) = \nu$, while if $\tilde{k} \notin [0, n]$ then $\mathscr{E}_{\beta(k^*)}(X)$ comes as close to ν as possible under the restriction $k^* \in [0, n]$. In a general sense, then, the *measure* $\beta(k^*)$ lies on a "neutral boundary" and Theorem 3.1 says that $\Delta(\beta, \nu)$ is maximized for a measure β lying on such a neutral boundary. A similar result was proved by Ray (1965, Theorem 4.1) for data with a density in the one-parameter exponential family.

EXAMPLE 3.1. Let $\alpha = M(q\delta_0 + p\delta_1)$, $p + q = 1$, $0 < p < 1$, $0 < \nu < 1$. Here $\tilde{k} = M(p - \nu) + n(1 - \nu)$ and $T_F(\nu) = p(1 - \nu)$. For sufficiently large n, $k^* = \tilde{k}$, so we may apply (3.2) and Theorem 3.1 to deduce that, for large n

and $x_i \in \{0, 1\}$, $i = 1, \cdots, n$,

$$\Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, \nu) \le \nu(1 - \nu)/(M + n + 1). \tag{3.3}$$

If n_0 is the smallest n such that $\nu(1 - \nu)/(M + n + 1) \le c$, then

$$n^* \le n_0 \doteq \nu(1 - \nu)c^{-1} - M - 1. \tag{3.4}$$

As previously noted, the model of this example is equivalent to the Bernoulli-beta(Mp, Mq) model. The bound (3.4) coincides with the bound derived by Ray (1965, Equation 5.16) and Lindley and Barnett (1965, Equation 10) for this latter form of the model. □

We now proceed with the proof of Theorem 3.1. The first step is to show that it holds when $n = 1$. This will follow from Lemma 3.1.

LEMMA 3.1. *Let α be given, let $M = \alpha(\mathbb{R})$, let F be the normalized form of α and let μ be the mean of F. Considered as a function of x, $\Delta(\alpha + \delta_x, 0)$ is nonincreasing for $x \ge -M\mu$ and nondecreasing for $x \le -M\mu$.*

PROOF. Let $d(x) = \Delta(\alpha + \delta_x, 0)$. From (2.4) and Lemma 2.2(iii and iv),

$$\begin{aligned} d(x) = {} & \frac{M}{(M+1)(M+2)} T_F(-M\mu - x) \\ & + \frac{2}{(M+1)(M+2)}\left[\left(x + \frac{1}{2}M\mu\right) \vee 0\right] - \left(\frac{x + M\mu}{M+1} \vee 0\right). \end{aligned} \tag{3.5}$$

The proof now follows from the examination of a number of cases. We distinguish two:

CASE 1. $\mu \ge 0$, $x \le -M\mu$. Here $x \le -M\mu/2$, so (3.5) gives

$$d(x) = \frac{M}{(M + 1)(M + 2)} T_F(-M\mu - x)$$

which is nondecreasing in x by Lemma 2.2(i).

CASE 2. $-M\mu < x \le -M\mu/2$. Here (3.5) and (2.3) give

$$d(x) = \frac{M}{(M + 1)(M + 2)} \mathscr{E}_F[X \vee (-M\mu - x)] - \frac{2(M\mu + x)}{(M + 1)(M + 2)}$$

which is nonincreasing in x.

Other cases follow similarly. □

We now prove Theorem 3.1 for the case $n = 1$ and $\nu = 0$.

PROPOSITION 3.1. *Let α be given, let F be the normalized form of α, let μ be the mean of F, and let $\nu = 0$. If there exist a and b, $a < 0 < b$, such that*

$\mathscr{P}_\alpha(a \le X \le b) = 1$, *then*

$$\Delta(\alpha + \delta_x, 0) \le \Delta(\alpha + k^*(1, 0)\delta_a + (1 - k^*(1, 0))\delta_b, 0)$$

where k^* *is defined following* (3.1).

PROOF. Let $\tilde{k} = \tilde{k}(1, 0)$ be as in (3.1). Suppose $x \in [a, b]$. If $\tilde{k} < 0$, then $b < -M\mu$, and so $x < -M\mu$. It follows from Lemma 3.1 that

$$\Delta(\alpha + \delta_x, 0) \le \Delta(\alpha + \delta_b, 0) = \Delta(\alpha + k^*\delta_a + (1 - k^*)\delta_b, 0).$$

Similarly, if $\tilde{k} > 1$,

$$\Delta(\alpha + \delta_x, 0) \le \Delta(\alpha + \delta_a, 0) = \Delta(\alpha + k^*\delta_a + (1 - k^*)\delta_b, 0).$$

Suppose now that $0 \le \tilde{k} \le 1$. By Lemma 3.1,

$$\begin{aligned}\Delta(\alpha + \delta_x, 0) &\le \Delta(\alpha + \delta_{-M\mu}, 0)\\ &= \frac{M}{(M+1)(M+2)} T_F(0) + \frac{(-M\mu) \vee 0}{(M+1)(M+2)}.\end{aligned} \tag{3.6}$$

The equality in (3.6) follows from (2.5) and Lemma 2.2 (iii and iv). By (3.2),

$$\begin{aligned}\Delta(\alpha + \tilde{k}\delta_a &+ (1 - \tilde{k})\delta_b, 0)\\ &= \frac{M}{(M+1)(M+2)} T_F(0) - \frac{M\mu + a}{(M+1)(M+2)} \cdot \frac{b}{b-a}.\end{aligned} \tag{3.7}$$

Since $0 \le \tilde{k} \le 1$, $k^* = \tilde{k}$, and, by (3.6) and (3.7),

$$\begin{aligned}\Delta(\alpha + k^*\delta_a &+ (1 - k^*)\delta_b, 0) - \Delta(\alpha + \delta_{-M\mu}, 0)\\ &= \frac{1}{(M+1)(M+2)}\left\{-\frac{(a + M\mu)b}{b-a} - [(-M\mu) \vee 0]\right\}.\end{aligned} \tag{3.8}$$

It is easy to verify that the right side of (3.8) is nonnegative under the condition $0 \le \tilde{k} \le 1$, in which case we have

$$\Delta(\alpha + \delta_x, 0) \le \Delta(\alpha + \delta_{-M\mu}, 0) \le \Delta(\alpha + k^*\delta_a + (1 - k^*)\delta_b, 0). \quad \square$$

We can restate Proposition 3.1 to say that, given α, there exists a $k \in [0, 1]$ such that $\Delta(\alpha + \delta_x, 0) \le \Delta(\alpha + k\delta_a + (1 - k)\delta_b, 0)$ for all $x \in [a, b]$. We use this restatement in the proof of the next lemma.

LEMMA 3.2. *For* α *given, suppose there exist* a *and* b *such that* $\mathscr{P}_\alpha(a \le X \le b) = 1$. *Let* $x_1, \cdots, x_n$ *lie in* $[a, b]$. *There exists a constant* $k = k(x_1, \cdots, x_n)$ *in* $[0, n]$ *such that*

$$\Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, 0) \le \Delta(\alpha + k\delta_a + (n - k)\delta_b, 0).$$

PROOF. Apply Proposition 3.1 successively. Take $\alpha + \sum_2^n \delta_{x_i}$ for α and x_1 for

x in the proposition; there exists a $k_1 \in [0, 1]$ such that

$$\Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, 0) \le \Delta(\alpha + \sum_2^n \delta_{x_i} + k_1\delta_a + (1 - k_1)\delta_b, 0).$$

Now take $\alpha + \sum_3^n \delta_{x_i} + k_1\delta_a + (1 - k_1)\delta_b$ for α and x_2 for x in Proposition 3.1; there exists a $k_2 \in [0, 1]$ such that

$$\begin{aligned}\Delta(\alpha + \textstyle\sum_2^n \delta_{x_i} + k_1\delta_a + (1 - k_1)\delta_b, 0)\\ \le \Delta(\alpha + \textstyle\sum_3^n \delta_{x_i} + (k_1 + k_2)\delta_a + (2 - k_1 - k_2)\delta_b, 0).\end{aligned}$$

By continuing in this fashion, we see that there exist $k_1, k_2, \cdots, k_n \in [0, 1]$ such that

$$\Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, 0) \le \Delta(\alpha + (\sum_1^n k_i)\delta_a + (n - \sum_1^n k_i)\delta_b, 0).$$

The lemma follows upon taking $k = \sum_1^n k_i$. □

Lemma 3.2 is used with the next lemma in proving Theorem 3.1.

LEMMA 3.3. *For α given, let F denote the normalized form of α, and let μ be the mean of F. Suppose there exist a and b such that $\mathscr{P}_\alpha(a \le X \le b) = 1$, and suppose $a < 0 < b$. Considered as a function of k, $\Delta(\alpha + k\delta_a + (n - k)\delta_b, 0)$ is maximized for $k \in [0, n]$ at $k = k^*(n, 0)$.*

PROOF. Let $\beta(k) = \alpha + k\delta_a + (n - k)\delta_b$ and let $\Delta(k) = \Delta(\beta(k), 0)$. Since $\mathscr{E}_{\beta(k)}(X) = [M\mu + ka + (n - k)b]/(M + n)$, (2.4) gives

$$\Delta(k) = \frac{1}{M+n+1} T_G(-M\mu - ka - (n-k)b) - \left(\frac{M\mu + ka + (n-k)b}{M+n} \vee 0\right) \tag{3.9}$$

where $G = \beta(k)/(M + n)$, the normalized form of $\beta(k)$. Let $g(k) = -M\mu - nb + k(b - a)$, and note that $g(k)$ is increasing in k. From (3.9), Lemma 2.2 and some manipulation give

$$\begin{aligned}\Delta(k) = {} & \frac{M}{(M+n)(M+n+1)} T_F(g(k))\\ & + \frac{k}{(M+n)(M+n+1)}[(a - g(k)) \vee 0]\\ & + \frac{n-k}{(M+n)(M+n+1)}[(b - g(k)) \vee 0] + \frac{g(k) \wedge 0}{M+n}.\end{aligned} \tag{3.10}$$

There are three cases to consider: (1) $g(k) < a$, (2) $g(k) > b$, and (3) $a \le g(k) \le b$. In Case 1, $g(k) < a$, so $g(k) < b$ also. By Lemma 2.2(ii), $T_F(g(k)) = \mu - g(k)$. Combining this with the definition of $g(k)$ and (3.10) gives $\Delta(k) = 0$. But since $\Delta(k^*) \ge 0$ by Lemma 2.1, we have $\Delta(k^*) \ge \Delta(k)$ for $g(k) < a$. A similar argument shows $\Delta(k^*) \ge \Delta(k)$ in Case 2: $g(k) > b$.

In Case 3: $a \le g(k) \le b$, we consider one subcase, $g(k) \ge 0$. The subcase

$g(k) < 0$ is handled similarly. When $0 \le g(k) \le b$, (3.10) reduces to

$$(3.11) \qquad (M + n)(M + n + 1)\Delta(k) = MT_F(g(k)) + h(k),$$

where $h(k) = (b - a)(k - n)(k - r)$ is a parabola opening upwards, with roots n and $r = [M\mu + (n + 1)b]/(b - a)$. Note that the restriction $0 \le g(k) \le b$ is equivalent to $(M\mu + nb)/(b - a) \le k \le r$, and so we must have $k \le n \wedge r$. However, $h(k)$ is decreasing in k for these values of k. In addition, since $g(k)$ is increasing in k, by Lemma 2.2(i) $T_F(g(k))$ is nonincreasing in k. From (3.11) then, $\Delta(k)$ is nonincreasing in k for $k \le n \wedge r$. Therefore, in this subcase, $\Delta(k)$ is maximized by taking k as small as possible, namely, by setting $k = \max\{0, (M\mu + nb)/(b - a)\} = k^*$. □

PROOF OF THEOREM 3.1. First assume that $\nu = 0$ and $a < 0 < b$. Let $x_1, \cdots, x_n \in [a, b]$. By Lemma 3.2, there exists a $k(x_1, \cdots, x_n) \in [0, n]$ such that

$$\Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, 0) \le \Delta(\alpha + k(x_1, \cdots, x_n)\delta_a + (n - k(x_1, \cdots, x_n))\delta_b, 0).$$

However, for all $x_1, \cdots, x_n \in [a, b]$, and $k(x_1, \cdots, x_n) \in [0, n]$,

$$\begin{aligned}\Delta(\alpha + k(x_1, \cdots, x_n)\delta_a + (n - k(x_1, \cdots, x_n))\delta_b, 0)& \\ \le \Delta(\alpha + k^*(n, 0)\delta_a + (n - k^*(n, 0))\delta_b, 0)&\end{aligned}$$

by Lemma 3.3. This proves Theorem 3.1 when $\nu = 0$.

If $\nu \ne 0$, an easy argument combining the immediately preceding discussion and Lemma 2.1(ii) gives

$$\Delta(\alpha + \textstyle\sum_1^n \delta_{x_i}, \nu) \le \Delta(\alpha + k^*(n, \nu)\delta_a + (n - k^*(n, \nu))\delta_b, \nu) \quad \text{for all} \quad x_i \in [a, b]. \ □$$

4. An example. In Example 3.1 we have applied the theory of this paper to essentially all α supported by two points. In this section we apply the theory to a more complex measure and, for such a measure, we suggest a means for approximating optimal rules and their worths.

Let H be the distribution function for a uniform random variable on [0, 1]: $H(x) = x$, $x \in [0, 1]$. We are interested in the case $\alpha = H$, in which case $M = 1$ and $F = H$. We fix $\nu = 0.5$ and consider the costs $c = 0.03$, 0.02, and 0.01.

From (3.1) $\tilde{k} = n/2$ and so $k^* = \tilde{k}$. $T_H(s) = (1 - s)^2/2$ for $s \in [0, 1]$, and so Theorem 3.1 and (3.2) give

$$\Delta(H + \textstyle\sum_1^n \delta_{x_i}, 0.5) = (2n + 1)/[8(n + 1)(n + 2)].$$

For $c = 0.03$, $n_0 = 6$; for $c = 0.02$, $n_0 = 10$; and for $c = 0.01$, $n_0 = 23$. So, for example, when $c = 0.03$, it follows from Theorems 2.1 and 3.1 that $n^* \le 6$. Since we must use backward induction to actually determine n^*, in the case $c = 0.03$ we must evaluate a six-dimensional integral. For smaller c, the dimension of the integral is even larger. The evaluation of such integrals on most computers is difficult and costly and so we suggest a means for finding suitable approximations.

Let H_m be the distribution function giving mass $1/m$ to the points $(i - 0.5)/m$; $i = 1, 2, \cdots, m$. For any α, let $V(\alpha, \nu) = \mathscr{E}_\alpha r(N^*)$, the worth of an optimal rule N^*. A result of Christensen (1983) may be modified to show that $V(H_m, \nu)$ converges in m to $V(H, \nu)$. We may thus use $V(H_m, \nu)$ to approximate $V(H, \nu)$. Moreover, this approximation can be found relatively easily: first find n_0 using Theorem 3.1 and then use backward induction to find $V(H_m, \nu)$. Since H_m is discrete, the backward induction involves the computation of sums rather than integrals. The results of such calculations are given in Table 4.1 for $c = 0.03$, 0.02, 0.01 and $m = 3, 4, 5, 6$.

While the convergence of $V(H_m, 0.5)$ is not monotone, it seems reasonable to use $V(H_6, 0.5)$ to approximate $V(H, 0.5)$. Also, on the basis of Table 4.1, we might conjecture that, for $\alpha = H$, n^* is 1 if $c = 0.03$ or 0.02 and 2 or 3 when $c = 0.01$. We can evaluate this conjecture by referring to Table 4.2. That table lists values of $V_n(H, 0.5)$, where $V_n(H, 0.5)$ is the optimal worth maximizing over all stopping rules which never take more than n observations. (These values were determined by using backward induction and multiple numerical integration.) A straightforward modification of DeGroot (1970, Section 12.13) shows that $V_n(H, 0.5)$ converges in n to $V(H, 0.5)$. From Table 4.2 it seems safe then to assume $n^* = 1$ when $c = 0.03$; $n^* = 3$ when $c = 0.02$; and $n^* = 4$ when $c = 0.01$. The conjectured values of n^* mentioned above are therefore too low, except in the case $c = 0.03$. However, the conjectured n^* values based on Table 4.1 are "nearly" optimal: when $c = 0.02$, $V_3(H, 0.5)$ exceeds $V_1(H, 0.5)$ by only 0.002; and when $c = 0.01$, $V_4(H, 0.5)$ exceeds $V_3(H, 0.5)$ by only 0.001. In addition, from Tables 4.1 and 4.2 it appears that $V(H_6, 0.5)$ closely approximates $V(H, 0.5)$.

The above discussion suggests that we can give good approximations to $V(H, 0.5)$ and n^*, but it does not give much practical help in actually carrying out a sequential design. We end this section with some suggestions in this direction and leave their evaluation to future study.

TABLE 4.1
n_0, n^*, and $V(H_m, 0.5)$ for given c and m.

m	$c = 0.03$			$c = 0.02$			$c = 0.01$		
	n_0	n^*	$V(H_m, 0.5)$	n_0	n^*	$V(H_m, 0.5)$	n_0	n^*	$V(H_m, 0.5)$
3	6	1	0.526	10	1	0.536	23	2	0.548
4	6	1	0.533	10	1	0.543	23	3	0.553
5	5	1	0.530	10	1	0.540	22	2	0.552
6	6	1	0.533	10	1	0.543	23	3	0.554

TABLE 4.2
Values of $V_n(H, 0.5)$ for given c and n.

n	$c = 0.03$	$c = 0.02$	$c = 0.01$
1	0.533	0.543	0.553
2	0.533	0.543	0.554
3	0.533	0.545	0.558
4	0.533	0.545	0.559

For the model $\alpha = MH$ and ν arbitrary, divide the range [0, 1] into m intervals of equal length, centered at the points $(i - 0.5)/m$, $i = 1, \cdots, m$. For any observation x, treat it as if it were equal to i_0, where i_0 is such that $x \in ((i_0 - 1)/m, i_0/m]$. (For example, if $m = 10$, we round each observation to the nearest tenth. We may do whatever we please with the outcome $x = 0$, since it occurs with probability zero.) Proceed now using the optimal strategy derived for the discrete measure MH_m. The idea underlying such an approach is that the strategy can be feasibly determined since H_m is discrete. We conjecture that for any $\varepsilon > 0$ there is an m such that the worth of the procedure described above exceeds $V(MH, \nu) - \varepsilon$.

Next, to generalize this approach, suppose that $\alpha = MF$ is an arbitrary measure. We now take F_m to be a discrete measure with mass $1/m$ assigned to each of the $(i - 0.5)/m$ quantiles of F. If an observed x lies between the $(i - 1)/m$ and i/m quantiles, we take x to be equal to the $(i - 0.5)/m$ quantile and follow the strategy which would be optimal if the Dirichlet parameter were MF_m. We conjecture that this will give nearly optimal procedures for sufficiently large m. Moreover, for any m, F_m will have bounded support even if F does not. We thus have a means for handling α with unbounded support in a fashion which we conjecture will be nearly optimal.

Acknowledgements. This paper is adapted from the author's Ph.D. dissertation presented to the University of Minnesota. I thank Professor Donald A. Berry for his support and guidance throughout the preparation of the dissertation and this paper. I am grateful to Mr. Neil Price-Jones for his computer programming assistance, and I thank the referees and associate editor for their helpful comments.

REFERENCES

CHERNOFF, H. (1968). Optimal stochastic control. *Sankhyā Ser. A* **30** 221–252.

CHOW, Y. S., ROBBINS, H. and SIEGMUND, D. (1971). *Great Expectations: The Theory of Optimal Stopping.* Houghton Mifflin, Boston.

CHRISTENSEN, R. (1983). Searching for the lowest price when the unknown distribution of prices is modeled with a Dirichlet process. Ph.D. Thesis. University of Minnesota.

CLAYTON, M. K. (1983). Bayes sequential sampling for choosing the better of two populations. Ph.D. Thesis. University of Minnesota.

DEGROOT, M. H. (1970). *Optimal Statistical Decisions.* McGraw-Hill, New York.

FERGUSON, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.* **1** 209–230.

LECHNER, J. A. (1962). Optimum decision procedures for a Poisson process parameter. *Ann. Math. Statist.* **33** 1384–1402.

LINDLEY, D. V. and BARNETT, B. N. (1965). Sequential sampling: two decision problems with linear losses for binomial and normal random variables. *Biometrika* **52** 507–532.

MORIGUTI, S. and ROBBINS, H. (1962). A Bayes test of $p \leq 1/2$ versus $p > 1/2$. *Rep. Statist. Appl. Res. Un. Japan. Sci. Engs.* **9** 39–60.

PRATT, J. W. (1966). The outer needle of some Bayes sequential continuation regions. *Biometrika* **53** 455–467.

RAY, S. N. (1965). Bounds on the maximal sample size of a Bayes sequential procedure. *Ann. Math. Statist.* **36** 859–878.

DEPARTMENT OF STATISTICS
UNIVERSITY OF WISCONSIN
MADISON, WISCONSIN 53706