The Annals of Statistics
1988, Vol. 16, No. 1, 236–253

A DYNAMIC SAMPLING APPROACH FOR DETECTING A CHANGE IN DISTRIBUTION

BY DAVID ASSAF

The Hebrew University

The problem of detecting a change in drift of Brownian motion is considered in the Bayesian framework with the time of change having a (prior) exponential distribution. To the well known problem of finding an optimal stopping rule for "declaring a change," we add the option of continuously controlling the sampling rates—resulting in controlling the variance coefficient of the process. The combined problem of finding an optimal rate function (dynamic sampling) together with an optimal stopping rule is solved and explicit expressions for the quantities of interest are derived.

The dynamic sampling procedure is shown to be significantly superior to constant rate sampling. The comparison is most favorable when the expected time until change tends to infinity, where the relative efficiency between the two procedures tends to infinity.

1. Introduction and summary. Let $\{x(t);\ t \geq 0\}$, $x(0) = 0$, be Brownian motion with drift coefficient μ and variance coefficient $\sigma^2 > 0$. The process starts with drift coefficient $\mu = \mu_0$, but at some time T the drift coefficient changes to $\mu = \mu_1$. The time of change T is unknown, but is assumed to have a (prior) exponential distribution with mean λ^{-1} ($\lambda > 0$). The parameters μ_0, μ_1, σ^2 and λ are all known. Occasionally, in accordance with some stopping rule, we "raise an alarm" indicating that we have evidence that the change has already occurred. The process is assumed to reset itself every time such an alarm is raised. The object is to find a stopping rule which, in some sense, minimizes both the probability of a false alarm and the expected delay (i.e., the time elapsed between the actual change and the sounding of the alarm).

To fix ideas, think of a machine with continuous production (ice cream, for example). At some point in time, the quality of production suddenly deteriorates and we want to find that point as quickly as possible after it occurs. Once the change is declared, the machine is checked, fixed if needed, and production with good quality starts again until the next change. As another example, consider testing for air or water pollution using a test tube.

Information regarding the quality of production is obtained by continuous sampling. We may, for example, decide to continuously sample 2% of the total production. The process $\{x(t);\ t \geq 0\}$ is a result of this sampling.

Let τ be a stopping time. Denote by $\alpha = P(\tau < T)$ the probability of a false alarm and by $\beta = E(\tau - T)^+$ the expected delay. The optimal stopping problem may formally be set up as that of finding a stopping time τ which minimizes β for some fixed $0 < \alpha < 1$.

Received December 1985; revised February 1987.

AMS 1980 *subject classifications*. Primary 62N10, 62K05; secondary 60J60, 60J65, 62L15, 93E20.

Key words and phrases. Detecting a change, expected delay, probability of false alarm, Brownian motion, diffusion process, stochastic control, dynamic sampling, optimal stopping.

OPTIMAL STOPPING PROBLEM.

Find $$\inf_{\tau} \beta$$

subject to $$P(\tau < T) = \alpha.$$

The solution to the optimal stopping problem may be found in the last chapter of the book by Shiryayev (1978). Different variations of the problem, including discrete formulations, cases in which some of the parameters are unknown, non-Bayesian formulations and others, have been studied by many authors. See for example Ashir and Muni (1975), Bather (1967) and Pollak and Siegmund (1985). For an up to date survey and a more complete list of references, see Pollak (1985, 1987).

In this paper the following problem is studied: Suppose that in addition to the optimal stopping time, we are allowed, at no additional cost, to vary the sample size as we like as long as on the average we do not sample more than before. In the machine production example, we may, as one possibility, sample 1% half the time and 3% the other half, rather than sampling 2% all of the time. (In the air pollution example, we may increase or decrease the amount of air flowing into the tube.) Can such a "dynamic sampling" procedure improve the performance, i.e., give a lower value of β for fixed α?

Think of σ^2 as the variance due to sampling error when sampling at some standard rate $a = 1$. Then sampling with rate $a(t)$ results in variance coefficient $\sigma^2/(a(t))$ (see Section 6, Comment 3). The rate at time t, $a(t)$, is naturally nonnegative (the case $a = 0$ is discussed later) and may depend on the history of the process up to time t [i.e., on $\{x(s); 0 \le s \le t\}$ and on $\{a(s); 0 \le s < t\}$]. The total sample size up to time t is then $\int_0^t a(s)\,ds$ and the average sample size per unit time is given by

$$\frac{1}{t}\int_0^t a(s)\,ds.$$

For given stopping time τ, denote by $C(\tau)$ the total sample size during the time from 0 to τ, that is,

$$C(\tau) = \int_0^\tau a(t)\,dt.$$

Throughout this paper, we restrict the discussion to a and τ satisfying $E\tau < \infty$ and $EC(\tau) < \infty$ (see Section 6, Comment 4). Assuming naturally that the decision process $a(t)$ is probabilistically reset whenever $x(t)$ is and applying a standard renewal-reward argument, the long-run average sample size per unit time is given by $EC(\tau)/E\tau$. We keep this value equal to some level of sampling $\gamma \ge 0$ (typically, $\gamma = 1$) and thus set up the dynamic sampling problem:

DYNAMIC SAMPLING PROBLEM.

Find $$\inf_{a,\tau} \beta$$

subject to $$P(\tau < T) = \alpha$$

and
$$\frac{EC(\tau)}{E\tau} = \gamma.$$

A heuristic derivation of the optimal a and τ for the dynamic sampling problem is given in Section 2. The performance of the resulting policy is approximated by a family of suboptimal policies in Section 3, where expressions for $EC(\tau)$, $E\tau$ and β are also derived. A rigorous proof of optimality is then given in Section 4.

Some numerical aspects are discussed in Section 5. These include comparisons between the dynamic sampling procedure and previous results obtained for fixed sampling rate $a(t) \equiv \gamma$. The comparisons show the dynamic sampling procedure to be significantly superior to the fixed rate one. The comparison is most striking when $\lambda \to 0$, where the optimal delay tends to infinity in the fixed rate case while tending to a finite constant in the dynamic sampling case. A slightly different comparison shows the relative efficiency between the two procedures to tend to infinity as $\lambda \to 0$.

Some additional comments and indications of possible extensions of these results are discussed in Section 6.

Models allowing for sampling at different rates have been studied, mainly in problems of quality control where several levels of sampling are permitted. The possibility of continuously controlling the variance in a similar context is discussed by Bather (1976). It seems that the only paper in which this possibility is mentioned regarding the detection problem is by Girshick and Rubin (1952), who discuss a discrete case version and allow a sample/do not sample scheme.

Theorems and techniques in stochastic processes used in this paper may be found in texts such as Itô and McKean (1965), Karlin and Taylor (1975, 1981) and Ross (1983). For the results used in stochastic control and dynamic programming, see Dynkin (1963), Fleming and Rishel (1975) and Strauch (1966).

2. Preliminary results and heuristic arguments. Denote by F_t the σ-field generated by the history of the process up to time t. It is well known [see Arrow, Blackwell and Girshick (1949)] that the process $\{y(t);\ t \geq 0\}$ defined by $y(t) = P(T \leq t | F_t)$ contains all relevant information, and we, thus, consider the problem in terms of this process. Shiryayev (1978) proves that $y(t)$ is a time homogeneous diffusion process with state space $[0,1]$ and derives its drift and variance coefficients as

$$\mu(y) = \lambda(1-y), \tag{1}$$

$$\sigma^2(y) = 2a(y)\rho y^2(1-y)^2, \tag{2}$$

where

$$\rho = \frac{1}{2}\frac{(\mu_1-\mu_0)^2}{\sigma^2} \tag{3}$$

and, naturally, $y(0) = 0$.

REMARK 1. Shiryayev (1978) computes the coefficients (1) and (2) for the constant sampling rate $a(y) \equiv 1$. This derivation, as well as others to follow, easily extends to arbitrary sampling rates and the proofs are omitted.

REMARK 2. The sub-σ-fields generated during time intervals with $a = 0$ are trivial and the process $y(t)$ is then a deterministic one. This corresponds to the situation in which no random fluctuations due to sampling are present, and the change is a result of the prior exponential only.

LEMMA 1. *For any stopping time* τ

$$\alpha = 1 - Ey(\tau), \tag{4}$$

$$\beta = E\int_0^\tau y(t)\,dt. \tag{5}$$

PROOF. A rigorous proof for the case $a \equiv 1$ is given in Shiryayev (1978). Intuitively (4) follows, since by definition $y(t) = P(T \le t|F_t)$, so that $Ey(\tau) = P(T < \tau) = 1 - \alpha$. The representation of β by (5) also follows intuitively from the definition of $y(t)$ because if the process is at y for Δt units of time, then the resulting expected delay is $y\Delta t$. □

The following lemma is used in this paper and may be of independent interest as well.

LEMMA 2. *For any stopping time* τ

$$E\tau = \beta + \frac{Ey(\tau)}{\lambda}. \tag{6}$$

PROOF. Let $\{z(t);\, t \ge 0\}$ be a diffusion process with drift coefficient $\mu(z)$ and let τ be a stopping time (as always assumed, $E\tau < \infty$). Dynkin's formula with the identity function $u(z) = z$ states that

$$E\int_0^\tau \mu(z(t))\,dt = Ez(\tau) - z(0). \tag{7}$$

Substitute the process $y(t)$ in (7) and use (1) to obtain

$$E\int_0^\tau \lambda(1 - y(t))\,dt = Ey(\tau) - y(0). \tag{8}$$

Since $y(0) = 0$, (8) becomes

$$\lambda E\tau - \lambda E\int_0^\tau y(t)\,dt = Ey(\tau), \tag{9}$$

from which (6) follows directly, using Lemma 1. □

For a heuristic derivation of the optimal policy, we make the two conjectures:

1. There exists an optimal stopping time which is stationary in y.
2. There exists an optimal sampling rate function which is monotone nondecreasing in y.

Conjecture 1 is reasonable since the problem is basically time homogeneous. Conjecture 2 seems reasonable since the value of $y(t)$ measures our current suspicion that the change has already occurred. It is natural, then, to try and obtain information quicker as our suspicion becomes higher.

Due to continuity of paths and (4), it is evident that the only stopping time which is both stationary and satisfies the constraint $P(\tau < T) = \alpha$ is the stopping time $\tau(\alpha)$ given by

$$\tau(\alpha) = \inf\{t;\ y(t) = 1 - \alpha\}. \tag{10}$$

(6) may now be written as

$$\beta = E\tau(\alpha) - \frac{1-\alpha}{\lambda}. \tag{11}$$

Since $(1 - \alpha)/\lambda$ is constant, it follows that minimizing the expected delay β is equivalent to minimizing the expected cycle time $E\tau(\alpha)$. The dynamic sampling problem may thus be equivalently represented as: Find

$$\inf_a E\tau(\alpha) \tag{12}$$

subject to

$$\frac{EC(\tau(\alpha))}{E\tau(\alpha)} = \gamma. \tag{13}$$

Using standard constrained optimization arguments, the problem may equivalently be written in an unconstrained form as: Find

$$\inf_a \{E\tau(\alpha) + kEC(\tau(\alpha))\} \tag{14}$$

or

$$\inf_a \left\{E\int_0^{\tau(\alpha)} (1 + ka(t))\, dt\right\} \tag{15}$$

for some $k > 0$.

For $0 \le y < 1 - \alpha$, let

$$H(y) = \inf_a \left\{E\int_0^{\tau(\alpha)} (1 + ka(t))\, dt | y(0) = y\right\} \tag{16}$$

denote the optimal value function for the problem. Formal substitution into the optimality equation for this case [see Fleming and Rishel (1975), Chapter 6] yields the quantity to be minimized as a function of a as

$$\inf_{a \ge 0} \{1 + ka + \mu(y)H'(y) + \tfrac{1}{2}\sigma^2(y)H''(y)\}. \tag{17}$$

Substituting the values of $\mu(y)$ and $\sigma^2(y)$ from (1) and (2) yields, after some simple rearrangements,

$$1 + \lambda(1 - y)H'(y) + \inf_{a \ge 0} \left\{a\left[k + \rho y^2(1 - y)^2 H''(y)\right]\right\}. \tag{18}$$

Since (18) is linear in a, it follows that the optimal value of a is either zero or infinity, depending on the sign of the quantity $k + \rho y^2(1 - y)^2 H''(y)$.

The function $H(y)$ is unknown at this stage, but is explicitly derived later in this paper. The proper regions, however, may readily be deduced from conjecture 2, and the optimal policy may, thus, be summarized as: There exists a value $0 < y_0 \le 1 - \alpha$ such that no sampling is performed in the interval $[0, y_0)$, while sampling with an infinite rate is performed in the interval $(y_0, 1-\alpha)$. A change is declared as soon as the process reaches $1-\alpha$ for the first time.

The correct value of y_0 naturally depends on γ, including the case $y_0 = 1-\alpha$, when $\gamma = 0$. The explicit form of this dependence, as well as the dependence of y_0 on k in the alternative representation, is discussed later.

In order for the preceding "policy" to have a manageable form, two basic questions still need to be answered. The first question regards the interpretation of sampling at an infinite rate. This may be solved by setting some high rate L and letting $L \to \infty$. It turns out to be computationally more convenient to set the rates as $a(y) = \lambda M/[\rho y(1-y)]$, with $M \to \infty$. The second question regards the sampling procedure around y_0, which may be resolved by the approximation: Set $0 < \varepsilon < y_0$ and use $a = 0$ in $[0, y_0 - \varepsilon]$ and $a = \infty$ in $[y_0, 1-\alpha)$. The sampling rates in the interval $(y_0 - \varepsilon, y_0)$ depend on the "direction" of the process, and we set $a = 0$ when the process is "moving to the right" (from $y_0 - \varepsilon$ to y_0) and as $a = \infty$ when it is "moving to the left." A rigorous formulation of this procedure using stopping times is given in Section 3.

Rather than proving optimality for the 0–∞ policy directly, its performance is approximated and the limiting value function is derived. This value function is then shown to, indeed, be the optimal one in Section 4.

3. The $A(M, \varepsilon, y_0)$ policies. Throughout this section, we fix α, $0 < \alpha < 1$. Consider a family of policies $\{A(M, \varepsilon, y_0);\ M > 1,\ \varepsilon < y_0 \le 1-\alpha\}$:

1. The stopping time τ $[= \tau(\alpha)]$ is defined by

$$\tau = \inf\{t\colon y(t) = 1-\alpha\}. \tag{19}$$

2. The sampling rates using $A(M, \varepsilon, y_0)$ are given for $y \notin (y_0 - \varepsilon, y_0)$, by

$$a(y) = \begin{cases} 0, & 0 \le y \le y_0 - \varepsilon, \\ \dfrac{\lambda M}{\rho y(1-y)}, & y_0 \le y < 1-\alpha. \end{cases} \tag{20}$$

For $y = y(t) \in (y_0 - \varepsilon, y_0)$, define the Markov time $u = u(t)$ by

$$u = \sup\{0 \le s \le t;\ y(s) = y_0 - \varepsilon \text{ or } y(s) = y_0\}$$

and let $a(y) = a(y(t), u, y(u))$ be defined as

$$a(y) = \begin{cases} 0, & y(u) = y_0 - \varepsilon, \\ \dfrac{\lambda M}{\rho y(1-y)}, & y(u) = y_0. \end{cases} \tag{21}$$

REMARK 3. The policies $A(M, \varepsilon, y_0)$ have an obvious "physical" interpretation, namely, do not sample at all until reaching the level y_0 for the first time. Whenever reaching y_0, sample at rate $\lambda M/[\rho y(1-y)]$ until the level reaches $1-\alpha$, in which case a change is declared, or drops a bit to $y_0 - \varepsilon$, in which case sampling is stopped until y_0 is reached again.

Expressions and bounds for $EC(\tau)$, $E\tau$ and β are next derived. The derivations are fairly standard and some technical details are omitted. A convenient reference to the concepts and techniques used is Karlin and Taylor [(1981), Chapter 15].

The basic behavior of $y(t)$ when applying the $A(M, \varepsilon, y_0)$ policy is: It starts at 0 and moves deterministically to y_0. Whenever at y_0, the process behaves like a diffusion process until hitting either $y_0 - \varepsilon$ or $1 - \alpha$. The drift coefficient is $\mu(y) = \lambda(1 - y)$, while the variance coefficient is given by (2) and (21) as

$$\sigma^2(y) = 2\lambda My(1 - y). \tag{22}$$

Upon reaching $y_0 - \varepsilon$, the process moves deterministically back to y_0 and a new "inner" cycle begins. Once $1 - \alpha$ is reached, a change is declared and an "external" cycle begins with the process reset at 0.

Consider one such inner cycle. Since

$$\frac{2\mu(y)}{\sigma^2(y)} = \frac{1}{My}, \tag{23}$$

it follows that the scale density is given by

$$s(y) = y^{-\delta}, \tag{24}$$

while the scale function is given by

$$S(y) = \frac{y^{1-\delta}}{1 - \delta} \tag{25}$$

with

$$\delta = \frac{1}{M}. \tag{26}$$

Note that $0 < \delta < 1$.

The probability of hitting $1 - \alpha$ before $y_0 - \varepsilon$, starting at y_0, is given by

$$\frac{S(y_0) - S(y_0 - \varepsilon)}{S(1 - \alpha) - S(y_0 - \varepsilon)}. \tag{27}$$

Thus the expected number of inner cycles is given by

$$N = \frac{S(1 - \alpha) - S(y_0 - \varepsilon)}{S(y_0) - S(y_0 - \varepsilon)}. \tag{28}$$

The Green function for an inner cycle is given by

$$G(y) = \begin{cases} \dfrac{2(S(y) - S(y_0 - \varepsilon))(S(1 - \alpha) - S(y_0))}{(S(1 - \alpha) - S(y_0 - \varepsilon))\sigma^2(y)s(y)}, & y_0 - \varepsilon \le y \le y_0, \\[2ex] \dfrac{2(S(y_0) - S(y_0 - \varepsilon))(S(1 - \alpha) - S(y))}{(S(1 - \alpha) - S(y_0 - \varepsilon))\sigma^2(y)s(y)}, & y_0 \le y \le 1 - \alpha. \end{cases} \tag{29}$$

Derivation of $EC(\tau)$. By the definition of the $A(M, \varepsilon, y_0)$ policies, $a = 0$ except for the inner cycles. Applying Wald's equation, we have

$$EC(\tau) = \int_{y_0-\varepsilon}^{1-\alpha} Na(y)G(y)\,dy, \tag{30}$$

with $a(y) = \lambda M/[\rho y(1-y)]$. Note that by (2),

$$[2a(y)]/[\sigma^2(y)] = 1/\left[\rho y^2(1-y)^2\right].$$

The portion of the integral from $y_0 - \varepsilon$ to y_0 is given by

$$I_1 = \int_{y_0-\varepsilon}^{y_0} \frac{(S(1-\alpha) - S(y_0))(S(y) - S(y_0-\varepsilon))}{(S(y_0) - S(y_0-\varepsilon))\rho s(y) y^2(1-y)^2}\,dy. \tag{31}$$

Obviously, $I_1 \geq 0$. Since $S(y) \leq S(y_0)$ for $y \leq y_0$ and $1/[s(y)] = y^\delta \leq 1$, we have

$$\frac{S(y) - S(y_0-\varepsilon)}{(S(y_0) - S(y_0-\varepsilon))s(y)} \leq 1, \tag{32}$$

so that, from (25),

$$0 \leq I_1 \leq \frac{1}{\rho}\frac{(1-\alpha)^{1-\delta} - y_0^{1-\delta}}{1-\delta}\int_{y_0-\varepsilon}^{y_0} \frac{dy}{y^2(1-y)^2}. \tag{33}$$

The upper bound on I_1 may thus be computed explicitly [see (38) for a similar computation]. An obvious approximation for the integral part is $\varepsilon/[y_0^2(1-y_0)^2]$. In any case, the upper bound is an $O(\varepsilon)$ term.

The portion from y_0 to $1 - \alpha$ is given by

$$I_2 = \int_{y_0}^{1-\alpha} \frac{S(1-\alpha) - S(y)}{s(y)}\frac{1}{\rho}\frac{1}{y^2(1-y)^2}\,dy. \tag{34}$$

For $0 < \delta < 1$ and $y_0 \leq y \leq 1 - \alpha$, the following inequalities are straightforward:

$$(1-\delta)y_0^\delta(1-\alpha-y) \leq \frac{(1-\alpha)^{1-\delta} - y^{1-\delta}}{y^{-\delta}} \leq 1 - \alpha - y. \tag{35}$$

Using (35), it follows that

$$\frac{y_0^\delta}{\rho}\int_{y_0}^{1-\alpha} \frac{1-\alpha-y}{y^2(1-y)^2}\,dy \leq I_2 \leq \frac{1}{\rho(1-\delta)}\int_{y_0}^{1-\alpha} \frac{1-\alpha-y}{y^2(1-y)^2}\,dy. \tag{36}$$

From (33) and (36), we have

$$EC(\tau) = \frac{1}{\rho}\int_{y_0}^{1-\alpha} \frac{1-\alpha-y}{y^2(1-y)^2}\,dy + O(\varepsilon) + O(\delta). \tag{37}$$

Taking the limits $\varepsilon \to 0$, $\delta \to 0$ and computing the integral explicitly yields the limiting value of $EC(\tau)$ as

$$(38)\qquad EC(\tau) = \frac{1}{\rho}\left[(1-\alpha-y_0)\frac{1-2y_0}{y_0(1-y_0)} + (1-2\alpha)\log\left(\frac{(1-\alpha)(1-y_0)}{\alpha y_0}\right)\right].$$

Derivation of $E\tau$. $E\tau$ is composed of the first hitting time of y_0 (starting at 0), given by

$$(39)\qquad t_0 = \frac{1}{\lambda}\log\left(\frac{1}{1-y_0}\right),$$

plus the expected accumulated times during the inner cycles, given by

$$(40)\qquad t_1 = N\left[\frac{1}{\lambda}\log\left(\frac{1-y_0+\varepsilon}{1-y_0}\right) + \int_{y_0-\varepsilon}^{1-\alpha} G(y)\,dy\right].$$

For $0 < \delta < 1$, $0 < y_0 \le 1-\alpha$ and $0 < \varepsilon < y_0$, let $g = g(\varepsilon, \delta) = (1-\delta)(y_0 - \varepsilon)^\delta$. Apply the inequalities

$$(41)\qquad g\frac{1-\alpha-y_0+\varepsilon}{\varepsilon} \le \frac{(1-\alpha)^{1-\delta}-(y_0-\varepsilon)^{1-\delta}}{y_0^{1-\delta}-(y_0-\varepsilon)^{1-\delta}} \le \frac{1}{g}\frac{1-\alpha-y_0+\varepsilon}{\varepsilon}$$

and

$$(42)\qquad x - \frac{x^2}{2} \le \log(1+x) \le x, \qquad 0 < x < 1,$$

to obtain

$$(43)\qquad N\frac{1}{\lambda}\log\left(\frac{1-y_0+\varepsilon}{1-y_0}\right) = \frac{1}{\lambda}\frac{1-\alpha-y_0}{1-y_0} + O(\varepsilon) + O(\delta).$$

For the second term in (40), apply (32) and (35) to obtain

$$(44)\qquad \int_{y_0-\varepsilon}^{1-\alpha} NG(y)\,dy = O(\varepsilon\delta) + O(\delta).$$

Upper and lower bounds for t_1 may easily be obtained from the corresponding inequalities applied. Taking the limits $\varepsilon \to 0$ and $\delta \to 0$ yields the limiting value

$$(45)\qquad E\tau = \frac{1}{\lambda}\left[\log\left(\frac{1}{1-y_0}\right) + \frac{1-\alpha-y_0}{1-y_0}\right].$$

REMARK 4. A person observing the $y(t)$ process when the $A(M, \varepsilon, y_0)$ policy is used in the limiting sense, will see the process travelling smoothly from 0 to y_0, reaching y_0 after time t_0. The process will then appear to be "standing at y_0" for some time (with expected value $(1/\lambda)[(1-\alpha-y_0)/(1-y_0)]$) and then suddenly "shooting in no time" to $1-\alpha$.

REMARK 5. For large M, the behavior of the $y(t)$ process during the inner cycles may be approximated by Brownian motion with zero drift. This may be better motivated by taking $a(y) = M/[2\rho y^2(1-y)^2]$, which yields a constant variance $\sigma^2(y) \equiv M$ and a uniformly bounded drift. Applying standard convergence theorems, one may then use the simpler forms of hitting probabilities and the Green function to derive the limiting quantities (38) and (45) more easily.

Derivation of β. Applying Lemma 2, the limiting value of β is easily computed as $\beta = E\tau - (1-\alpha)/\lambda$, that is,

$$\beta = \frac{1}{\lambda}\left[\log\left(\frac{1}{1-y_0}\right) - \frac{\alpha y_0}{1-y_0}\right]. \tag{46}$$

An appropriate $O(\varepsilon) + O(\delta)$ should be added when using an $A(M, \varepsilon, y_0)$ policy. The exact formula for the added term is, of course, the same as for $E\tau$.

REMARK 6. An alternative derivation of the limiting value of β is offered by Lemma 1. As the process travels from 0 to y_0, we have $y(t) = 1 - e^{-\lambda t}$. Upon reaching y_0, we have $y(t) \equiv y_0$ for an expected time of $(1/\lambda)[(1-\alpha-y_0)/(1-y_0)]$. Thus, by (5),

$$\beta = \int_0^{t_0}(1 - e^{-\lambda t})\,dt + \frac{y_0}{\lambda}\frac{1-\alpha-y_0}{1-y_0}, \tag{47}$$

which may be checked to equal (46).

4. Proving optimality. Taking Lemmas 1 and 2 into account and applying standard arguments as in Section 2, the dynamic sampling problem may equivalently be set up as: Find

$$\inf_{a,\tau}\{E\tau + k_1 EC(\tau) - k_2 Ey(\tau)\} \tag{48}$$

or

$$\inf_{a,\tau} E\left[\int_0^\tau (1 + k_1 a(t))\,dt - k_2 y(\tau)\right] \tag{49}$$

for some positive k_1 and k_2.

For $0 \le y \le 1$, let

$$F(y) = E\left[\int_0^\tau (1 + k_1 a(t))\,dt - k_2 y(\tau) | y(0) = y\right] \tag{50}$$

be the value function when using the $A(M, \varepsilon, y_0)$ policy in the limiting $\varepsilon \to 0$, $M \to \infty$ sense, and with $1 - \alpha = y_1$, that is, with τ defined as

$$\tau = \inf\{t;\ y(t) \ge y_1\}. \tag{51}$$

The value of $F(0)$ is essentially computed in (38) and (45). Using similar methods, it may readily be shown that for $0 < y_0 < y_1 < 1$,

$$F(y) = \begin{cases} F_0(y) - k_2 y_1, & 0 \le y \le y_0, \\ F_1(y) - k_2 y_1, & y_0 \le y \le y_1, \\ -k_2 y, & y_1 \le y \le 1, \end{cases} \tag{52}$$

where

$$F_0(y) = \frac{1}{\lambda}\log\left(\frac{1-y}{1-y_0}\right) + L(y_0, y_1), \tag{53}$$

$$\begin{aligned} F_1(y) &= \frac{y_1 - y}{y_1 - y_0} L(y_0, y_1) \\ &+ \frac{k_1}{\rho}\left[(1-2y)\log\left(\frac{y}{1-y}\right) - \frac{(y_1 - y)(1-2y_0)}{y_1 - y_0}\log\left(\frac{y_0}{1-y_0}\right)\right. \\ &\left. + (2y_1 - 1)\frac{y - y_0}{y_1 - y_0}\log\left(\frac{y_1}{1-y_1}\right)\right] \end{aligned} \tag{54}$$

and

$$\begin{aligned} L(y_0, y_1) &= \frac{1}{\lambda}\frac{y_1 - y_0}{1 - y_0} \\ &+ \frac{k_1}{\rho}\left[(y_1 - y_0)\frac{1 - 2y_0}{y_0(1-y_0)} + (2y_1 - 1)\log\left(\frac{y_1(1-y_0)}{(1-y_1)y_0}\right)\right]. \end{aligned} \tag{55}$$

For $y_0 = y_1 > 0$ (the no sampling case), the middle $F_1(y) - k_2 y_1$ term is omitted in $F(y)$. An additional case may occur when $y_0 = y_1 = 0$. This corresponds to immediate stopping (and no sampling) for all $0 \le y \le 1$, and the appropriate value function is $F(y) = -k_2 y$.

THEOREM. *For proper choice of $y_0 \le y_1$, $F(y)$ is the value function for the dynamic sampling problem as set up in* (49), *i.e.*,

$$F(y) = \inf_{a,\tau} E\left[\int_0^\tau (1 + k_1 a(t))\,dt - k_2 y(\tau) | y(0) = y\right]. \tag{56}$$

The choice of y_0 and y_1 as functions of k_1 and k_2 is as follows.

CASE 1. *If*

$$\left(\frac{\lambda k_1}{\rho}\right)^{1/2} < 1 - \frac{1}{\lambda k_2}, \tag{57}$$

then

$$y_0 = \left(\frac{\lambda k_1}{\rho}\right)^{1/2} \tag{58}$$

and $y_0 < y_1 < 1$ is the unique solution to the equation

$$F_1'(y_1) = -k_2. \tag{59}$$

CASE 2. *If*

$$\left(\frac{\lambda k_1}{\rho}\right)^{1/2} \ge 1 - \frac{1}{\lambda k_2} \quad \textit{and} \quad \lambda k_2 > 1, \tag{60}$$

then

$$y_0 = y_1 = 1 - \frac{1}{\lambda k_2}. \tag{61}$$

CASE 3. *If*

$$\lambda k_2 \le 1, \tag{62}$$

then

$$y_0 = y_1 = 0. \tag{63}$$

PROOF. We prove the theorem for Case 1 only. The arguments for the two remaining cases are similar and somewhat easier.

Note first that $F(y)$ is continuous. Next differentiate to obtain

$$F_0'(y) = -\frac{1}{\lambda(1-y)}, \tag{64}$$

$$F_1'(y) = -\frac{1}{\lambda(1-y_0)} + \frac{k_1}{\rho}\left[2\log\left(\frac{y_0(1-y)}{(1-y_0)y}\right) + \frac{1-2y}{y(1-y)} - \frac{1-2y_0}{y_0(1-y_0)}\right]. \tag{65}$$

Also,

$$F_0''(y) = -\frac{1}{\lambda(1-y)^2}, \tag{66}$$

$$F_1''(y) = -\frac{k_1}{\rho}\frac{1}{y^2(1-y)^2}. \tag{67}$$

Since $F_1'(y_0) = -1/[\lambda(1-y_0)]$, $F_1''(y) < 0$ and $F_1'(y) \to -\infty$ as $y \to 1$, it readily follows that the solution of (59) for y_1 is indeed unique and satisfies $y_0 < y_1 < 1$. The choice of y_0 and y_1 also results in continuity of F' as well as continuity of F'' at y_0 (F'' is evidently not continuous at y_1).

Since $F(y)$ is attained as a limit of given policies, it suffices to prove that it satisfies the optimality conditions as inequalities [this follows from general theory; see Dynkin (1963), Fleming and Rishel [(1975), Chapter 6] and Strauch (1966)]. We begin by showing that the stopping region defined by τ is optimal. For this, we need to prove that for all $a \ge 0$, $y_1 < y \le 1$,

$$1 + k_1 a + D^a F(y) \ge 0 \tag{68}$$

and that for all $0 \le y \le y_1$,

$$F(y) \le -k_2 y, \tag{69}$$

where

$$D^a = \mu(y)\frac{d}{dy} + \rho a y^2(1-y)^2\frac{d^2}{dy^2} \tag{70}$$

is the stochastic operator corresponding to sampling at rate a.

To prove (68), note that $F(y) = -k_2 y$ for $y_1 < y \le 1$ and (68) thus becomes

$$1 + k_1 a - k_2\lambda(1-y) \ge 0. \tag{71}$$

Since the left-hand side is minimal for $a = 0$ and $y = y_1$, it suffices to show that

(72) $$\lambda k_2(1 - y_1) \le 1,$$

that is,

(73) $$F_0'(y_1) = -\frac{1}{\lambda(1 - y_1)} \le -k_2.$$

By the definition of y_1, we have $F_1'(y_1) = -k_2$. Also, by continuity of F' at y_0,

(74) $$F_0'(y_0) = F_1'(y_0).$$

To prove (73), it thus suffices to show

(75) $$F_0''(y) \le F_1''(y), \qquad y_0 \le y \le y_1.$$

Inequality (75), however, follows easily from (58), (66), and (67).

The proof of inequality (69) follows along similar lines. Since $F(y_1) = -k_2 y_1$, it suffices to prove $F'(y) \ge -k_2$ for all $0 \le y \le y_1$. However, $F'(y_1) = F_1'(y_1) = -k_2$, hence, it suffices to prove $F''(y) < 0$ for all $0 \le y < y_1$, which in turn follows directly from (66) and (67).

To complete the proof of the theorem it remains to show that $F(y)$ "cannot be improved" in the continuation region $[0, y_1)$. Formally, this amounts to proving that for all $a \ge 0$, $0 \le y < y_1$,

(76) $$1 + k_1 a + D^a F(y) \ge 0.$$

[Note that this is simply (68) for the continuation region.] Inequality (76) is implied by the pair of inequalities

(77) $$1 + \mu(y)F'(y) \ge 0,$$

(78) $$k_1 + \rho y^2(1 - y)^2 F''(y) \ge 0,$$

which we next show to hold for all $0 \le y < y_1$.

(i) Take $0 \le y \le y_0$. Then $\mu(y)F'(y) \equiv -1$ and (77) holds as equality. Also,

(79) $$\rho y^2(1 - y)^2 F''(y) = -\frac{\rho}{\lambda} y^2 \ge \frac{-\rho}{\lambda} y_0^2 = \frac{-\rho}{\lambda}\frac{\lambda k_1}{\rho} = -k_1,$$

so that (78) holds.

(ii) For $y_0 \le y < y_1$, we have $\rho y^2(1 - y)^2 F''(y) \equiv -k_1$ and (78) holds as equality. It remains to verify that (77) holds. In this case, we have

(80) $$\mu(y)F'(y) = \frac{\lambda k_1}{\rho}\left[2(1 - y)\log\left(\frac{y_0(1 - y)}{(1 - y_0)y}\right) + \frac{1 - 2y}{y} - \frac{(1 - 2y_0)(1 - y)}{y_0(1 - y_0)}\right] - \frac{1 - y}{1 - y_0}.$$

Thus, $\mu(y_0)F'(y_0) = -1$, so (77) holds at $y = y_0$. To complete the proof, it is sufficient to show that $\mu(y)F_1'(y)$ is nondecreasing for $y > y_0$. Now

$$(\mu(y)F_1'(y))' \tag{81}$$
$$= \frac{\lambda k_1}{\rho}\left[2\log\left(\frac{y(1-y_0)}{y_0(1-y)}\right) - \frac{2}{y} - \frac{1}{y^2} + \frac{1-2y_0}{y_0(1-y_0)}\right] + \frac{1}{1-y_0}.$$

Omitting the (nonnegative) logarithmic term, it suffices to show

$$\frac{\lambda k_1}{\rho}\frac{1-2y_0}{y_0(1-y_0)} + \frac{1}{1-y_0} \geq \frac{\lambda k_1}{\rho}\left(\frac{2}{y} + \frac{1}{y^2}\right), \tag{82}$$

but (82) is obvious since it holds (as equality) for $y = y_0$ and its right-hand side decreases with y while its left-hand side is constant. □

5. Numerical aspects and results. In this section, we check some properties of the dynamic sampling procedure via the limiting expressions (38), (45) and (46) for $EC(\tau)$, $E\tau$ and β. In practice, this may mean that our choice of M and ε is such that the limiting expressions are sufficiently close to the corresponding actual ones. The values of M (or $\delta = M^{-1}$) and ε needed to obtain a given approximation may be calculated using the exact expressions or, more easily, using the appropriate bounds derived in Section 3. It is perhaps worth noting that for given M and ε, the performance of the bounds depends on the values of the parameters λ and ρ as well as on the values of α and γ (though y_0). Examination of this behavior reveals that special care is needed around the extreme values, most notably around $\lambda = 0$ and $y_0 = 0$.

Consider first given values of $0 < \alpha < 1$, $\lambda > 0$ and $\rho > 0$. Straightforward examination of (38), (45) and (46) reveals that as functions of y_0, $EC(\tau)$ decreases while $E\tau$ and β increase. In fact, $EC(\tau) \to 0$ as $y_0 \uparrow 1 - \alpha$ and $EC(\tau) \to \infty$ as $y_0 \downarrow 0$. $E\tau$, on the other hand, remains bounded (and bounded away from zero) for all $0 \leq y_0 \leq 1 - \alpha$. Thus, as a function of y_0, $EC(\tau)/E\tau$ is monotone decreasing and tends to infinity as $y_0 \downarrow 0$ and to zero as $y_0 \uparrow 1 - \alpha$. As a result, the equation $EC(\tau)/E\tau = \gamma$ has a unique solution $y_0 = y_0(\gamma)$ for each $\gamma \geq 0$.

Furthermore, for fixed $0 < \alpha < 1$, the relation may be written as

$$\frac{EC(\tau)}{E\tau} = \frac{\lambda}{\rho}g(y_0), \tag{83}$$

where g is decreasing as discussed previously. The solution of y_0 as a function of λ, γ and ρ is thus given by

$$y_0 = g^{-1}\left(\frac{\gamma\rho}{\lambda}\right). \tag{84}$$

In particular, it follows that y_0 decreases as a function of γ and ρ and increases as a function of λ. Once y_0 is determined for given level γ, the expected delay β is readily computed using (46).

A second type of computation, needed for comparing efficiencies, requires the average sampling rate γ for obtaining a given expected delay. This is done by first solving (46) for y_0 and then substituting in (83) to compute the correct γ. Since all formulas are explicit, all computations are easily programmed on a calculator.

Asymptotic behavior. (i) as $\lambda \downarrow 0$. Note first that $EC(\tau)$ does not depend on λ while $E\tau$ is a multiple of $1/\lambda$. Thus, $EC(\tau)/E\tau$ tends to zero as $\lambda \to 0$ as long as all other quantities remain fixed. To obtain interesting results, take $0 < \alpha < 1$, $\rho > 0$ and $\gamma > 0$ fixed and change y_0 as $\lambda \to 0$ so as to keep $EC(\tau)/E\tau$ equal to γ. From the previous discussion, it follows that y_0 must then tend to zero with λ. For small values of y_0, however, the following approximations may be used:

$$EC(\tau) \cong \frac{1}{\rho}\frac{1-\alpha}{y_0}, \tag{85}$$

$$E\tau \cong \frac{1-\alpha}{\lambda}, \tag{86}$$

$$\beta \cong \frac{1-\alpha}{\lambda} y_0. \tag{87}$$

To keep $EC(\tau)/E\tau = \gamma$, we thus need

$$y_0 \cong \lambda/\gamma\rho, \tag{88}$$

which results in an approximate value for β,

$$\beta \cong \frac{1-\alpha}{\gamma\rho}. \tag{89}$$

Thus as $\lambda \to 0$, the expected delay *remains bounded* in the dynamic sampling procedure. This is in contrast to the results for the fixed rate procedure, where the delay tends to infinity as $\lambda \to 0$.

(ii) As $\lambda \to \infty$. To keep $EC(\tau)/E\tau = \gamma$, we need $y_0 \uparrow 1 - \alpha$. The expected delay is then approximately

$$\beta \cong \frac{1}{\lambda}\left[\log\left(\frac{1}{\alpha}\right) - (1-\alpha)\right], \tag{90}$$

which is of hyperbolic form $\beta \cong c/\lambda$ ($c \cong 1.4$ for $\alpha = 0.1$). The approximation in (90) also holds when no sampling at all is done [check that in this latter case $E\tau = (1/\lambda)\log(1/\alpha)$ and use Lemma 2]. As a result, (90) holds for *any* sampling procedure, in particular, for the fixed rate one. The improvement obtained by the dynamic sampling procedure thus loses its effect as $\lambda \to \infty$.

Numerical results. The expected delay when sampling at constant rate $a \equiv 1$ is given in Shiryayev (1978) as

$$\beta = \frac{1}{\rho}\left[\int_{1/(1-\alpha)}^{\infty} \frac{e^{\Lambda x}(x-1)^{\Lambda}}{x^2}\left(\int_x^{\infty} \frac{ue^{-\Lambda u}}{(u-1)^{2+\Lambda}}\,du\right)dx\right], \qquad \Lambda = \frac{\lambda}{\rho}. \tag{91}$$

The integral tends (slowly) to infinity as $\lambda \to 0$ (for any fixed $0 < \alpha < 1$, $\rho > 0$). The values in the second row of Table 1 were computed by numerical

TABLE 1

λ	100	10	1	0.1	0.01	0.001	0.0001	...	0
β (fixed rate)	0.0138	0.131	0.869	2.63	4.70	6.78	8.85	...	∞
β (dynamic sampling)	0.0138	0.125	0.649	1.01	0.931	0.905	0.901	...	0.9
γ needed (efficiency)	0.999	0.577	0.521	0.364	0.210	0.137	0.102	...	0

integration of the integral in (91) using the DBLIN subroutine. The upper limits of integration were set up as finite ones (very large though for small values of λ), which indicates that the actual values are even slightly higher than those given in the table.

The table makes the comparison between fixed rate and dynamic sampling procedures, with parameters set as $\alpha = 0.1$ and $\rho = 1$. In the first row, different values of λ are considered. The second row lists the results of the numerical integration of (91) for the expected delay with $a \equiv 1$. The third row is the expected delay using the dynamic sampling procedure with the same average rate $\gamma = 1$. The last row is the computation of the γ needed in the dynamic sampling procedure for obtaining the same expected delay as the $a \equiv 1$ procedure yields (for example, the dynamic sampling needs 0.21 on the average to obtain an expected delay of 4.70 for $\lambda = 0.01$). The relative efficiency of the dynamic sampling procedure is thus ~ 2 for $\lambda = 1$ and ~ 7 for $\lambda = 0.001$. As $\lambda \to 0$, the relative efficiency tends to infinity, while as $\lambda \to \infty$, the relative efficiency tends to 1.

All numbers are rounded off to three significant digits. Values in the second row are approximate lower bounds as mentioned previously.

6. Additional comments.

1. The dynamic sampling model is suitable for problems in which the variance is due to sampling error, and large samples which lead to smaller variance are possible.

2. In many problems a $0-\infty$ sampling procedure will not be a practical one and some additional constraints relating to (frequent and large) changes in sampling rates may be present. The solution of the unconstrained problem described in this paper may then serve as a bound on performance and as a reference value.

3. The optimal solution in other models may, in some cases, be an obvious generalization of the present one. For example, if bounds on sampling rates are present [say $a_1 \le a(t) \le a_2$], then the optimal solution is an a_1–a_2 one [see (18)]. Similarly, if the "production" rate is some number R, then a more appropriate model may be set up by taking the variance coefficient of $\{x(t);\ t \ge 0\}$ as $[\sigma^2/a(t)][1 - a(t)/R]$. It may readily be checked that a 0–R policy is optimal in this latter case as well.

4. Throughout the paper, only finite-expectation stopping times were considered. In view of Lemma 2, with a suitable truncation if needed, it is clear that a stopping time τ with $E\tau = \infty$ cannot be optimal for minimizing β in any

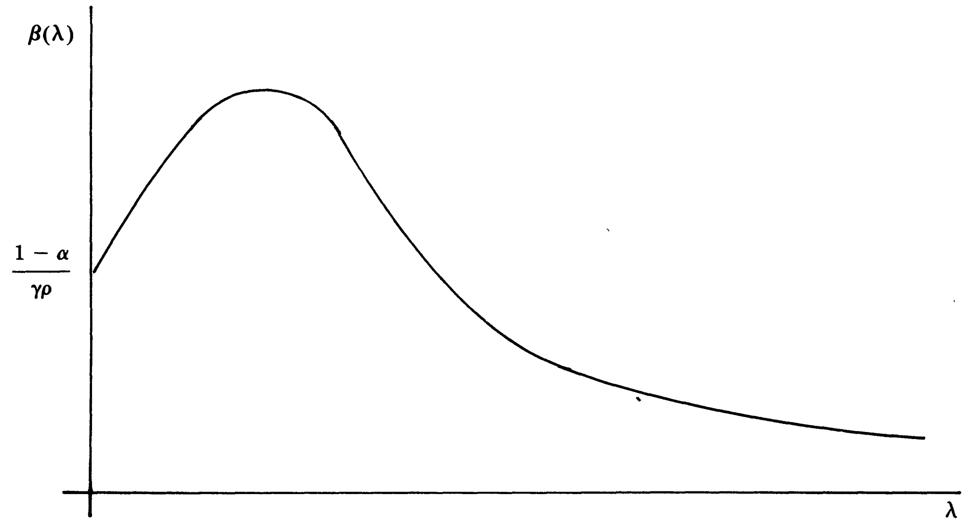

FIG. 1.

reasonable sense. Similarly, a sampling rate procedure with $EC(\tau) = \infty$ will not be optimal except for the trivial case $\gamma = \infty$, which leads to $y_0 = 0$ and $\beta = 0$.

5. Consider fixed $0 < \alpha < 1$, $\rho > 0$ and $\gamma > 0$. For any λ, let $\beta(\lambda)$ be the expected delay in the optimal dynamic sampling procedure with $EC(\tau)/E\tau = \gamma$. Then, $\beta(\lambda) \to 0$ as $\lambda \to \infty$ and $\beta(\lambda) \to (1 - \alpha)/\gamma\rho$ as $\lambda \to 0$ (see Section 5). It is perhaps surprising that $\beta(\lambda)$ is not monotone (as is also apparent from the third row of Table 1). The schematic form of $\beta(\lambda)$ is sketched in Figure 1.

Acknowledgments. I am indebted to Ya'akov Ritov for several observations which were most valuable in reaching the optimal solution. I have benefited from several discussions with Larry Brown, Bob Keener, Moshe Pollak and David Siegmund regarding different aspects of the problem. Many thanks are due an Associate Editor and two referees for comments on an earlier version of this paper. Thanks are also due Israel Einot for the computer results of Section 5.

REFERENCES

ARROW, K. J., BLACKWELL, D. and GIRSHICK, M. A. (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica* **17** 213–244.

ASHIR, S. and MUNI, S. S. (1975). On tests for detecting change in mean. *Ann. Statist.* **3** 98–108.

BATHER, J. A. (1967). On a quickest detection problem. *Ann. Math. Statist.* **38** 711–724.

BATHER, J. A. (1976). A control chart model and a generalized stopping problem for Brownian motion. *Math. Oper. Res.* **1** 209–224.

DYNKIN, E. B. (1963). Optimal selection of stopping time for a Markov process. *Dokl. Akad. Nauk SSSR* **150** 238–240.

FLEMING, W. H. and RISHEL, R. W. (1975). *Deterministic and Stochastic Optimal Control.* Springer, New York.

GIRSHICK, M. A. and RUBIN, H. (1952). A Bayes approach to a quality control model. *Ann. Math. Statist.* **23** 114–125.

ITÔ, K. and MCKEAN, H. P., JR. (1965). *Diffusion Processes and Their Sample Paths.* Springer, New York.

KARLIN, S. and TAYLOR, H. M. (1975). *A First Course in Stochastic Processes.* Academic, New York.

KARLIN, S. and TAYLOR, H. M. (1981). *A Second Course in Stochastic Processes.* Academic, New York.

POLLAK, M. (1985). Optimal detection of a change in distribution. *Ann. Statist.* **13** 206–227.

POLLAK, M. (1987). Average run lengths of an optimal method of detecting a change in distribution. *Ann. Statist.* **15** 749–779.

POLLAK, M. and SIEGMUND, D. (1985). A diffusion process and its application to detecting a change in the drift of a Brownian motion process. *Biometrika* **72** 267–280.

ROSS, S. M. (1983). *Stochastic Processes.* Wiley, New York.

SHIRYAYEV, A. N. (1978). *Optimal Stopping Rules.* Springer, New York.

STRAUCH, R. E. (1966). Negative dynamic programming. *Ann. Math. Statist.* **37** 871–890.

DEPARTMENT OF STATISTICS
THE HEBREW UNIVERSITY
JERUSALEM 91905
ISRAEL