

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 34	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Numerical Methods For Bayes Sequential Decision Problems		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Herman Chernoff A. John Petkau		8. CONTRACT OR GRANT NUMBER(s) N00014-75-C-0555
9. PERFORMING ORGANIZATION NAME AND ADDRESS Statistics Center Massachusetts Institute of Technology Cambridge, Massachusetts 02139		10. PROGRAM ELEMENT PROJECT, TASK AREA & WORK UNIT NUMBERS (NR-609-001)
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Statistics and Probability Code 436 Arlington, Virginia 22217		12. REPORT DATE January 1984
		13. NUMBER OF PAGES 127
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES This report also appears as Technical Report No. 83-26 (November, 1983) in the Institute of Applied Mathematics and Statistics series at the University of British Columbia, Vancouver, B.C., Canada.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Backward induction; Bayes risk; Decision theory; Dynamic programming; Free boundary problem; Numerical methods; Optimal stopping; Random walks; Sequential analysis; Weiner process		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) See reverse side.		

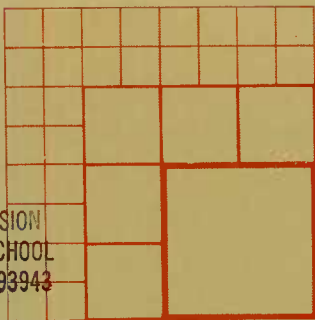
Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

ABSTRACT

The general approach to sequential decision-theoretic problems where sums of successive observations are approximated by a continuous time Wiener process has a number of fundamental advantages. Simple numerical techniques which can be employed to obtain explicit descriptions of the solutions of the resulting continuous time optimal stopping problems are described. The techniques are not well adapted for very precise results, but are surprisingly effective for reasonably accurate approximations. Special features of particular problems can be exploited to reduce the necessary computational effort. The techniques are illustrated in a number of problems thereby clearly indicating their properties.

LIBRARY
RESEARCH REPORTS DIVISION
NAVAL POSTGRADUATE SCHOOL
MONTEREY, CALIFORNIA 93943



STATISTICS CENTER

Massachusetts Institute of Technology

77 Massachusetts Avenue, Rm. E40-111, Cambridge, Massachusetts 02139 (617) 253-8722

NUMERICAL METHODS FOR BAYES SEQUENTIAL DECISION PROBLEMS

BY

HERMAN CHERNOFF

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

AND

A. JOHN PETKAU

UNIVERSITY OF BRITISH COLUMBIA

TECHNICAL REPORT NO. ONR 34

JANUARY 1984

PREPARED UNDER CONTRACT

N00014-74-C-0555 (NR-609-001)

FOR THE OFFICE OF NAVAL RESEARCH

Reproduction in whole or in part is permitted for
any purpose of the United States Government

This document has been approved for public release
and sale; its distribution is unlimited

NUMERICAL METHODS FOR BAYES SEQUENTIAL DECISION PROBLEMS

by

Herman Chernoff and A. John Petkau

ABSTRACT

The general approach to sequential decision-theoretic problems where sums of successive observations are approximated by a continuous time Weiner process has a number of fundamental advantages. Simple numerical techniques which can be employed to obtain explicit descriptions of the solutions of the resulting continuous time optimal stopping problems are described. The techniques are not well adapted for very precise results, but are surprisingly effective for reasonably accurate approximations. Special features of particular problems can be exploited to reduce the necessary computational effort. The techniques are illustrated in a number of problems thereby clearly indicating their properties.

Some key words: Backward induction; Bayes risk; Decision theory;
Dynamic programming; Free boundary problem; Numerical methods;
Optimal stopping; Random walks; Sequential analysis; Weiner process.

AMS 1970 subject classifications: Primary 62L10, 62L04;

Secondary 62C10, 62L15, 65D99.

NUMERICAL METHODS FOR BAYES SEQUENTIAL DECISION PROBLEMS

by

Herman Chernoff

Statistics Center
Massachusetts Institute of Technology

and

A. John Petkau

Department of Statistics
University of British Columbia

DEDICATION

The work reported here was initiated when the second author was a Ph.D. student in the Department of Statistics at Stanford working under the supervision of the first author, Herman Chernoff. The techniques were developed further over the intervening years in conjunction with our other joint research efforts. Throughout this period, Herman Chernoff has been a constant source of stimulation and inspiration to the second author. In recognition of this fact, the second author would like to dedicate his work on this project to Herman Chernoff on the recent occasion of his sixtieth birthday (July 1, 1983).

TABLE OF CONTENTS

	<u>Page</u>
§1. Introduction	1
§2. Problems under investigation	4
§3. Numerical techniques	13
§4. Illustration of techniques	24
§5. Practical implementation in statistical problems	38
§6. The Anscombe problem with ethical cost	48
§7. Summary and comments	62
REFERENCES	67
TABLES	70
FIGURES	103

1. INTRODUCTION

A natural formulation for many statistical problems is one combining Bayesian, sequential and decision-theoretic aspects. For the problem of deciding the sign of a normal mean, Chernoff (1961, 1965a, 1965b), Breakwell & Chernoff (1964) and Bather (1962) develop an approach to such a formulation where sums of successive observations are replaced by a continuous time Wiener process. Subsequently, this approach has been employed by Chernoff & Ray (1965), Chernoff (1967), Bather & Chernoff (1967a, 1967b), Feder & Stroud (1971), Petkau (1978), and Chernoff & Petkau (1981) in a wide variety of problems.

The continuous time problem has a number of fundamental advantages over the discrete time problem for which it is an approximation. First, the continuous time problem can be normalized so that many of the parameters which appear in the original (discrete time) problem are eliminated; thus, a single continuous time problem corresponds to an entire class of discrete time problems. Second, the continuous time problem is equivalent to an optimal stopping problem for a Wiener process where the cost associated with stopping depends only on the stopping point; any such problem is related to a problem in analysis, a free boundary problem (FBP) involving the heat equation. This relationship facilitates obtaining bounds and asymptotic approximations for the solution of the continuous time problem.

While these bounds and asymptotic approximations provide valuable insight, in most problems they do not provide an adequate approximation to the solution. Techniques are required which will provide numerical descriptions of the solution of the continuous time problem; this solution will then provide approximations to the solutions of an entire class of discrete time problems.

In this paper we will describe simple numerical techniques which can be easily employed to obtain explicit descriptions of the solutions of such continuous time problems. The basic idea is straightforward: the Wiener process is approximated by a discrete time process and backward induction is employed to solve the optimal stopping problem for this new process. The techniques will be illustrated in a number of problems thereby clearly indicating their properties. Some of these problems have and some have not been previously investigated in the literature.

The reader may feel that the path that has just been traced is somewhat circular. We begin with a discrete time Bayes sequential decision (optimal stopping) problem which can be solved by backward induction and approximate it by a continuous time optimal stopping problem for a Wiener process which we propose to solve by applying backward induction to a discrete time version of the Wiener process. However, the situation is not quite so empty. First, as already indicated, the continuous time problem allows one to derive valuable characteristics of the solution including asymptotic approximations. Second, there are available excellent approximations to the difference between the solution of the continuous time problem and those of its various discrete time versions. Thus, we can solve the optimal stopping problem for a particular discrete time version and use the solution, properly adjusted, to approximate that of the continuous time problem. This approximation can then itself be adjusted further to approximate the solution of any of the original discrete time problems. Thus, the single backward induction applied to the discrete time version of the Wiener process provides solutions for the normalized continuous time problem and all of its discrete time versions.

In this paper we will focus on obtaining numerical solutions for continuous time problems. The question of whether these continuous time solutions, when

properly adjusted, yield accurate approximations to the solutions of the original discrete time problems must be considered on a problem-by-problem basis. We merely mention here that this question has been considered in detail for a problem involving Bernoulli data by Petkau (1978) and for a different problem involving normal data by Chernoff and Petkau (1981); in both cases, the adjusted continuous time solutions provided accurate approximations to the solutions of the original discrete time problems.

2. PROBLEMS UNDER INVESTIGATION

The techniques to be described can be applied to obtain a numerical description of the solution of our special class of optimal stopping problems involving zero drift Wiener processes. With one exception, the normalized forms of the continuous time Bayes sequential decision problems to be considered in this paper are special cases of the following optimal stopping problem: Given a Wiener process $\{Y(s): s \geq s_1\}$ in the $-s$ scale, described by $E\{dY(s)\} = 0$ and $\text{Var}\{dY(s)\} = -ds$ and starting at $Y(s_0) = y_0$ ($s_0 > s_1$), find a stopping time S to minimize the risk, $E\{d(Y(S), S)\}$; here $d(y, s)$ is the cost associated with stopping at the point (y, s) and stopping is enforced at the end of the problem, namely, when $s = s_1$.

Some characteristics of the solution of such an optimal stopping problem can now be described. If we define $\hat{d}(y_0, s_0) = \inf b(y_0, s_0)$, where $b(y_0, s_0)$ is the risk associated with a particular stopping time (procedure) and the infimum is taken over all procedures, then since $Y(s)$ is a process of independent increments, $\hat{d}(y, s)$ represents the best that can be achieved once (y, s) has been reached, irrespective of how it was reached. Thus, the rule "Stop as soon as $\hat{d}(Y(s), s) = d(Y(s), s)$," which yields an optimal procedure if one exists, can be described by the continuation set $C = \{(y, s): \hat{d}(y, s) < d(y, s)\}$ or by its complement, the stopping set $S = C^c = \{(y, s): \hat{d}(y, s) = d(y, s)\}$; attention can therefore be restricted to procedures which can be so described. Note that this characterization does not depend upon the initial point (y_0, s_0) and thus yields the solution for all initial points simultaneously. Chernoff (1968, 1972) has demonstrated that one should expect the solution (\hat{d}, C) of the stopping problem to be a solution of the following free boundary problem (FBP) involving the heat equation (∂C denotes the boundary of the set C):

$$(2.1) \quad \begin{aligned} \frac{1}{2} \hat{d}_{yy}(y, s) &= \hat{d}_s(y, s) & \text{for } (y, s) \in C, \\ \hat{d}(y, s) &= d(y, s) & \text{for } (y, s) \in S, \\ \hat{d}_y(y, s) &= d_y(y, s) & \text{for } (y, s) \in \partial C; \end{aligned}$$

this relationship enables one to obtain bounds and asymptotic approximations for the solution. One particular result is that for any such stopping problem one should never stop at points (y, s) where $H(y, s) \equiv \frac{1}{2} \hat{d}_{yy}(y, s) - \hat{d}_s(y, s) < 0$; if one thinks of the optimal stopping problem as a gambling problem, then $H(y, s)$ can be heuristically thought of as the "rate of losing" at the point (y, s) . Further, it is obvious from (2.1) that changing the stopping cost function $d(y, s)$ by adding to it any solution of the heat equation leaves the solution S of the FBP unchanged.

Some special cases of this general optimal stopping problem which have already been investigated in the literature are now described.

Example 2.1. Testing for the sign of a normal mean. Chernoff (1961). X_1, X_2, \dots are independent $N(\mu, \sigma^2)$ random variables (σ^2 known). It is desired to test $H_1: \mu \geq 0$ versus $H_2: \mu < 0$, where the cost of a wrong decision is $k|\mu|$ and the cost of observing n X 's is cn . If the parameter μ is assumed to have a normal prior, what is the Bayes sequential strategy? A normalized form of the continuous time version of this problem is a special case of the general stopping problem formulated above with

$$(2.2) \quad d(y, s) = s^{-1} + s^{\frac{1}{2}} \psi(y/s^{\frac{1}{2}}) \quad \text{for } s > 0;$$

here

$$(2.3) \quad \begin{aligned} \gamma(x) &= \phi(x) - x[1 - \phi(x)] & \text{for } x \geq 0, \\ &= \gamma(-x) & \text{for } x < 0, \end{aligned}$$

while ϕ and Φ are the standard normal density and cumulative respectively. For further detail, the reader is referred to Chernoff (1961, 1965a, 1965b, 1972), Breakwell & Chernoff (1964) and Bather (1962). Closely related work appears in Lindley (1961), Moriguti & Robbins (1962) and Lindley & Barnett (1965). We will refer to this problem as the sequential analysis problem.

Example 2.2. One-armed bandit problem, Chernoff & Ray (1965). X_1, X_2, \dots, X_n are independent $N(\mu, \sigma^2)$ random variables (σ^2 known). The payoff for stopping at $n \leq N$ is $X_1 + X_2 + \dots + X_n$. When μ has a normal prior, the normalized continuous time version leads to the special case

$$(2.4) \quad d(y, s) = -y/s \quad \text{for } s \geq 1.$$

The variation where X_i is either a or b with unknown probabilities p and $1 - p$ and p has a beta prior is relevant to (a) a one-armed bandit problem with a limited number of pulls available, (b) the rectified sampling inspection problem in which context this problem first appeared, and (c) clinical trials comparing a new treatment against a known one with a finite horizon of patients to be treated, Chernoff (1967). For discussion of the continuous time version, see Chernoff (1967, 1972).

Example 2.3. Sequential medical trials involving paired data, Anscombe (1963).

There is a horizon of N patients to be treated with one of two available treatments. In the initial (experimental) phase, n pairs of patients are treated sequentially with different treatments randomly assigned to the patients in each pair; the remaining $N - 2n$ patients are all assigned to the treatment which is inferred to be superior. The differences in the values of the outcomes for each pair are

independent $N(\mu, \sigma^2)$ random variables (σ^2 known) and the cost of treating any patient with the inferior treatment is proportional to $|\mu|$. If the parameter μ is assumed to have a normal prior, what is the Bayes sequential strategy? The continuous time version, recently studied in detail by Chernoff & Petkau (1981), leads to the special case

$$(2.5) \quad d(y, s) = -(1 - 1/s) |y| \quad \text{for } s \geq 1.$$

Related work appears in Begg & Mehta (1979), Petkau (1980), Lai, Levin, Robbins & Siegmund (1980) and Lai, Robbins & Siegmund (1983). We will refer to this problem as the Anscombe problem.

Example 2.4. Sequential medical trials for comparing an experimental with a standard treatment, Petkau (1978). There is a horizon of N patients to be treated with either the standard treatment, characterized by a known probability of success p_0 , or the experimental treatment, characterized by an unknown probability of success p . Sampling is to be initiated with the experimental treatment and continued with this treatment during an experimental period until a decision is made in favor of one of the treatments; the remaining patients are then treated with the favored treatment. There is a cost incurred for each unsuccessful application of either treatment as well as a cost of experimentation which is incurred for each patient treated during the experimental period. If a beta prior is assumed for the parameter p , what is the optimal design? A continuous time version of this problem leads to the special case (here γ is a normalized cost of experimentation parameter)

$$(2.6) \quad \begin{aligned} d(y, s) &= \gamma/s - y & \text{for } y \geq 0, s \geq 1, \\ &= \gamma/s - y/s & \text{for } y < 0, s \geq 1. \end{aligned}$$

The above examples arise naturally in the statistical problems described. In each case, closed form solutions are unavailable; complete descriptions of optimal procedures are available only through numerical techniques such as those to be described. In order to fully illustrate the properties of these numerical techniques, a problem of the same general form as Examples 2.1 - 2.4 but for which the solution is available in closed form will be useful. The following modification of Example 2.3 will serve our purpose.

Example 2.5. Modified Anscombe problem. This artificial problem corresponds to the special case

$$(2.7) \quad d(y,s) = -(1 - 1/s) |y| - 2s^{1/2} \psi(\tilde{y}(s)/s^{1/2}) \quad \text{for } s \geq 1,$$

where $\tilde{y}(s)$ is defined by

$$(2.8) \quad 1 - \psi(\tilde{y}(s)/s^{1/2}) = s^{-1/2}.$$

It is easily verified (see, for example, Chernoff, 1968, 1972) that the optimal solution (\hat{d}, C) for this problem is given by

$$(2.9) \quad \begin{aligned} C &= \{(y,s): |y| \leq \tilde{y}(s), s \geq 1\}, \\ \hat{d} &= -|y| - 2s^{1/2} \psi(y/s^{1/2}) \quad \text{for } (y,s) \in C, \end{aligned}$$

where ψ is defined in (2.3); of course, $\hat{d} \equiv d$ for $(y,s) \in S = C^c$.

The statistical problems described above all lead to special cases of the general optimal stopping problem for a zero drift Wiener process in the (y,s) scale which was described in the first paragraph of this section. While these

statistical problems will be the main interest in this paper, the techniques to be described apply equally well to a class of gambling problems, the general case of which can be described as follows: Given a Wiener process $\{X(t): t \leq t_1\}$ in the t scale, described by $E\{dX(t)\} = 0$ and $\text{Var}\{dX(t)\} = dt$ and starting at $X(t_0) = x_0$ ($t_0 < t_1$), find a stopping time T to maximize the expected reward $E(g(X(T), T))$; here $g(x,t)$ is the reward associated with stopping at the point (x,t) and stopping is enforced at the end of the problem, namely, when $t = t_1$. The solution of any such problem will not depend upon the initial point (x_0, t_0) and we will denote the optimal reward at (x,t) by $\hat{g}(x,t)$. For a detailed study of this general problem, the interested reader is referred to Van Moerbeke (1974a, 1974b, 1975) and Shiryaev (1978). Two results of particular interest are that the solution of this stopping problem can be represented as the solution of a free boundary problem for the backward heat equation, $\frac{1}{2}u_{xx} + u_t = 0$, and that one should never stop at points (x,t) where $H(x,t) \equiv \frac{1}{2}g_{xx}(x,t) + g_t(x,t) > 0$; $H(x,t)$ can be thought of as the payoff rate or "rate of winning" at the point (x,t) . Again, changing the stopping reward function $g(x,t)$ by adding to it any solution of the backward heat equation leaves the solution S of the FBP unchanged.

It should not be surprising that there is a close relation between this general optimal stopping problem for a zero drift Wiener process in the (x,t) scale and the general problem in the (y,s) scale which was defined earlier; a simple change of variables transforms one into the other (see, for example, Van Moerbeke, 1974a, p.547). In spite of this close relation, we will prefer to work with the statistical problems in the (y,s) scale and the gambling problems in the (x,t) scale since these are the natural scales.

Two special cases of this general optimal stopping problem in the (x,t) scale will be considered.

Example 2.6. Van Moerbeke's gambling problem, Van Moerbeke (1974a, 1974b).

A gambler loses an amount of money equal to the amount of time the process spends in the region $x \leq 0$ and wins an amount of money equal to the amount of time spent in the region $x > 0$. If the gambler is permitted to stop the process at any time t , $0 < t \leq 1$, what is the optimal strategy?

For this problem, the payoff rate $H(x, t) = \pm 1$ depending on whether the process is in the positive or negative x half-plane; clearly the gambler should never stop in the win region, $x > 0$. A naive gambler would stop as soon as the lose region, $x \leq 0$, was entered, but it may pay to lose a bit now in the hope of winning in the more remote future. The reward function described above differs from

$$(2.10) \quad \begin{aligned} g(x, t) &= 1 - t + 2x^2 & \text{for } x > 0, \\ &= 1 - t & \text{for } x \leq 0, \end{aligned}$$

by a solution of the backward heat equation. Hence this problem is equivalent to the special case with reward function $g(x, t)$. Van Moerbeke has established that the optimal solution (\hat{g}, C) for this problem is given by

$$(2.11) \quad \begin{aligned} C &= \{(x, t): x \geq -\alpha(1-t)^{1/2}, t \leq 1\}, \\ \hat{g}(x, t) &= 2(1-t)(1+w^2) + \alpha(1-t)(w\phi(w) \\ &\quad - (1+w^2)[1-\phi(w)])/\phi(\alpha) \quad \text{for } (x, t) \in C, \end{aligned}$$

where $w = x/(1-t)^{1/2}$ and $\alpha \approx 0.5061$ is the solution of $\alpha\phi(\alpha) = \phi(\alpha)$.

Example 2.7. The S_n/n problem with finite horizon. In the infinite horizon version of this problem a gambler is allowed to view successively as many terms as he pleases of a sequence X_1, X_2, \dots of independent random variables with common distribution F . If upon stopping at time n the gambler receives a payoff of S_n/n , where $S_n = X_1 + \dots + X_n$, what is the optimal strategy?

This problem was first studied by Chow & Robbins (1965), who proved the existence of an optimal stopping rule when F is a two point distribution. They also proved the intuitively obvious but nontrivial fact that an optimal rule is to stop at the first n at which $S_n \geq \beta_n$, and provided a method of calculating the sequence of numbers β_n in principle. Dvoretzky (1965) and Teicher & Wolfowitz (1966) proved that the same result holds for any F with finite second moment (the β 's depend upon F , of course). Dvoretzky also showed that if F has zero mean and unit variance then $0.32 \dots \leq \beta_n/n^{1/2} \leq 4.06 \dots$ for n sufficiently large, and conjectured that $\lim \beta_n/n^{1/2}$ existed. This conjecture was proved independently by Taylor (1968), Walker (1969), and Shepp (1969), who found $0.8399 \dots$ as the limiting value. They pointed out that considered for large values of n , this problem would be equivalent to its Wiener process analogue, the special case where for $0 < t < \infty$

$$g(x, t) = x/t;$$

the optimal solution (\hat{g}, C) for this problem is given by

$$C = \{(x, t): x \leq \alpha t^{1/2}, 0 < t < \infty\}$$

$$\hat{g}(x, t) = (1 - \alpha^2)t^{-1/2}\phi(w)/\phi(\alpha) \quad \text{for } (x, t) \in C,$$

where $w = x/t^{1/2}$ and $\alpha \approx 0.8399$ is the solution of $\alpha\phi(\alpha) = (1 - \alpha^2)\phi(\alpha)$.

The finite horizon variation of this problem, in which the gambler is permitted to observe at most N terms of the sequence X_1, X_2, \dots , has been considered by J. L. Snell & H. Tisdale (1978). A normalized form of the

continuous time version of this problem leads to the special case where for $0 < t \leq 1$,

$$(2.12) \quad g(x,t) = x/t ;$$

it is this particular version of the S_n/n problem which will be considered here.

In the remainder of this section we briefly preview the rest of this paper. In Section 3, the backward induction methods for the normal and binomial discrete time versions are described together with alternative versions of continuity corrections. In Section 4, the examples we have presented are discussed to illustrate and evaluate the continuity corrections. The modified Anscombe problem, Example 2.5, for which the solution is known illustrates the case of a symmetric region where the boundary is monotone. Example 2.6, Van Moerbeke's gambling problem, illustrates the case where truncation may be used to capitalize on one-sided stopping regions. Example 2.7 illustrates the case where the boundary is not monotone.

In Section 5, the problem of computing solutions over large ranges of s values is addressed by a technique of changing increments in s . This method is used to present numerical results for the important classical sequential analysis problem, Example 2.1, and one-armed bandit problem, Example 2.2.

Finally in Section 6, a new example is introduced. This is the Anscombe problem with ethical cost considerations. It is new in two senses. It has not been treated before in the literature. It is different from Examples 2.1 to 2.7 in that the posterior risk on stopping depends not only on the current position of the Wiener process but also on the past history. This problem can still be solved numerically by backward induction or it can be transformed into a stopping problem of the same form as the others.

3. NUMERICAL TECHNIQUES

In this section we describe the techniques to be employed in obtaining numerical descriptions of the solution of the general optimal stopping problem for a zero drift Wiener process in the (y,s) scale defined at the beginning of the previous section. As already indicated, the basic idea is straightforward: the process $Y(s)$ which is a Wiener process in the $-s$ scale, is approximated by a discrete time process, and backward induction is used to solve the optimal stopping problem for this new process. Using asymptotic results concerning the relation of the solution of the discrete time problem to the solution of the Wiener process problem, the discrete time solution can then be adjusted to provide an approximation to the solution of the continuous time problem.

A natural approximation to the continuous time problem results if one considers the discrete time problem where one is permitted to stop only on the discrete set of possible values of s , $(s_1 + i\Delta, i = 0, 1, \dots)$. While the value of s decreases by Δ between these successive possible stopping times, the process $Y(s)$ changes by a normal deviate with mean 0 and variance Δ ; in effect, the Wiener process is being approximated by an appropriate sum of independent normal random variables. At any point (y,s) where s corresponds to a permissible stopping time, the choice between either stopping at this point or continuing on to the next permissible stopping time and proceeding optimally thereafter is made on the basis of which of $d(y,s)$ or $E\{\hat{d}(Y(s-\Delta), s-\Delta) | Y(s)=y\}$ is smaller. Thus, the backward induction algorithm which yields the optimal solution to the stopping problem for this discrete time process is specified by

$$(3.1) \quad \begin{aligned} \hat{u}(y,s) &= d(y,s) && \text{for } s = s_1, \\ &= \min[d(y,s), E\{\hat{d}(y + Z\Delta^{1/2}, s - \Delta)\}] && \text{for } s > s_1. \end{aligned}$$

where Z represents a standard normal deviate.

This approximation is a natural one since the discrete time problem is embedded within the continuous time problem; the former corresponds to the special case of the latter where one is permitted to stop only on the discrete set of values of s given by $\{s_i + i\Delta, i = 0, 1, \dots\}$. From this point of view it is obvious that the continuous time problem is more favorable. For a sequence of values of Δ approaching 0, the solution of the discrete time problem (both the continuation region and the risk) would approach that of the continuous time problem in a monotonic fashion.

Note that the evaluation of the expectation appearing in (3.1) would require a numerical integration for which purpose the y axis would also be discretized. Thus, in practice, the backward induction is carried out on a grid of (y, s) points each of which is classified as either a stopping or continuation point during the course of the computation.

How would one use the results of the backward induction algorithm (3.1) to obtain an approximation to the boundary $\bar{y}(s)$ of the continuation region for the continuous time problem? Chernoff (1965b) presents a detailed investigation of the relation of this discrete time problem to the continuous time problem; the results lead to two distinct methods of approximating the continuous time boundary $\bar{y}(s)$.

The first method consists of simply adjusting the boundary of the optimal continuation region for the discrete time problem; this boundary is determined by the "break-even" points $\bar{y}_\Delta(s)$ at which $d(y, s) = E\{d(y + Z\Delta^{1/2}, s - \Delta)\}$. Chernoff (1965b) has established that

$$(3.2) \quad \bar{y}(s) = \bar{y}_\Delta(s) + k\Delta^{1/2} + o(\Delta^{1/2}),$$

where the sign is determined so as to make the continuation region for the Wiener process problem larger and $k = -\zeta(1/2)/\sqrt{2\pi} \approx 0.5826$, where ζ is the Riemann zeta function. The first method should then be clear: For a (reasonably small) value Δ , carry out the backward induction algorithm and obtain the break-even points $\bar{y}_\Delta(s)$ at each fixed value of s . Then use the correction implied by (3.2) to approximate $\bar{y}(s)$. Note that since the entire backward induction is carried out on a grid of (y, s) points, the break-even points $\bar{y}_\Delta(s)$ would only be obtained approximately, presumably by some interpolation or extrapolation scheme (Day, 1969, provides the details of a scheme for carrying out the backward induction together with an interpolation scheme for approximating the break-even points for the discrete time version of Example 2.3). We shall call this the adjustment method and label it A.

For the second method, the break-even points need not be approximated. Chernoff (1965b) has established that, in the neighbourhood of the boundary of the optimal continuation region for the continuous time problem, the difference between the optimal risk for the discrete time problem and the cost for stopping behaves asymptotically (as $\Delta \rightarrow 0$), at every fixed value of s , like a simple function which depends upon the (unknown) location of the continuous time boundary at this value of s and whose form he provides; indeed, it is this result which leads to the relationship (3.2). This result forms the basis of the second method: For a (reasonably small) value of Δ , carry out the backward induction algorithm to obtain the optimal risk for the discrete time problem at each grid point. Then, at each fixed value of s , fit the known values of the difference between the optimal risk for the discrete time problem and the cost for stopping at those grid points in the interior of the continuation region (but adjacent to the boundary) to the relationship provided by Chernoff

(1965b) in order to approximate (or, more precisely, to extrapolate to) the location of the continuous time boundary (further details for a closely related scheme will be provided below). We shall call this the extrapolation method and label it E.

While the discrete time process with normal increments is the most natural approximation to the Wiener process, we propose to use the simpler approximation in which the Wiener process Y is replaced by the simple random walk process where $Y(s - \Delta) = Y(s) \pm \Delta^{1/2}$, each with probability 1/2. This approximation to the Wiener process results in a very simple corresponding backward induction algorithm; the standard normal deviate Z in (3.1) is replaced by a random variable which is ± 1 , each with probability 1/2, leading to the algorithm

$$(3.3) \quad \begin{aligned} \hat{d}(y, s) &= d(y, s) && \text{for } s = s_1, \\ &= \min[d(y, s), (\hat{d}(y + \Delta^{1/2}, s - \Delta) + \hat{d}(y - \Delta^{1/2}, s - \Delta))/2] && \text{for } s > s_1. \end{aligned}$$

Obviously, this algorithm is considerably simpler to implement than that specified by (3.1) which requires a numerical integration to evaluate the risk at each grid point (y, s) . As was the case with the previous approximation, the backward induction is carried out on a grid of (y, s) points; in the present approximation, however, the discretization of the y -axis is necessarily related to the discretization of the s -axis. Whereas the Wiener process was previously being approximated by the sum of its increments, in this simpler approximation the increment of the Wiener process is itself replaced by a Bernoulli random variable. While the second moment of the Bernoulli variable is chosen to match that of the increment it is replacing, the higher even moments do not match. In general, therefore, it is not clear whether this discrete problem is more

or less favorable than the Wiener process problem. Further, while the solution of this discrete time problem (both the continuation region and the risk) would also approach that of the continuous problem as Δ approached 0, one would not necessarily expect the behavior to be monotone.

Chernoff & Petkau (1976) have investigated the relation of this discrete time simple random walk problem to the original continuous time problem. They establish that the appropriate analogue of (3.2) for the present case is

$$(3.4) \quad \tilde{y}(s) = \tilde{y}_\Delta(s) \pm 0.5\Delta^{1/2} + o(\Delta^{1/2}),$$

and also provide the form of the simple function which describes, for each fixed value of s , the asymptotic (as $\Delta \rightarrow 0$) behavior of the difference between the optimal risk for this discrete time problem and the cost for stopping in the neighbourhood of the boundary of the optimal continuation region for the continuous time problem. These results enable the same two general approaches described above of approximating the continuous time boundary to be used in connection with the backward induction algorithm (3.3). Further details of these methods in the context of this discrete time simple random walk approximation will now be provided. For simplicity of discussion, we will suppose throughout the remainder of this section that we are in the case of a one-sided problem where the optimal continuation region for the continuous time problem is given by $C = \{(y, s): s > s_1 \text{ and } y < \tilde{y}(s)\}$, where $\tilde{y}(s)$ is monotonically increasing in s .

To implement the adjustment approach, the break-even points $\tilde{y}_\Delta(s)$ at which $d(y, s) = d^*(y, s)$ where

$$(3.5) \quad d^*(y, s) = (\hat{d}(y + \Delta^{1/2}, s - \Delta) + \hat{d}(y - \Delta^{1/2}, s - \Delta))/2$$

must be approximated at each fixed value of $s \in \{s_1 + i\Delta: i = 1, 2, \dots\}$. In carrying out the algorithm (3.3) at s , one would discover the grid level $y_0 = y_0(s)$ on the y axis determined by the condition that the grid points $(y_0 - j\Delta^{\frac{1}{2}}, s) = (y_j, s)$ say, be classified as continuation points ($d^* \leq d$) for $j = 0, 1, \dots$ and as stopping points ($d^* > d$) for $j = -1, -2, \dots$. The grid level $y_0(s)$ might be called the highest, or last, continuation level at s ; the sequence of highest continuation levels would be nondecreasing and a naive approximation to $\bar{y}_\Delta(s)$, the break-even point at s , would be provided by $y_0(s)$. Note, however, that in the case of this discrete time random walk approximation, the grid being employed has a vertical spacing of $\Delta^{\frac{1}{2}}$ which is coarse compared to the horizontal spacing of Δ ; for reasonably small values of Δ , therefore, as s increases successive values of $y_0(s)$ would often be identical and this naive approach would produce a series of steps as an approximation to the gradually increasing sequence of break-even points. One might attempt to smooth the sequence of levels $y_0(s)$ to form an improved (hopefully) approximation to the sequence of break-even points $\bar{y}_\Delta(s)$. This could be accomplished by the crude approach in which $\bar{y}_\Delta(s)$ is approximated by $y_0(s)$ only at those values on the s grid where the last continuation level changes, that is, $y_0(s) > y_0(s - \Delta)$; line segments connecting these approximation points could then be used to approximate $\bar{y}_\Delta(s)$ at any intermediate value of s . The approximation can then be adjusted by $0.5\Delta^{\frac{1}{2}}$ to provide the crude adjusted estimate of $\bar{y}(s)$; this method is labelled CA.

The above method of approximating the break-even points may seem crude since the computed values of the risk at the grid points are completely ignored, except that they are employed to classify the grid points as either

stopping or continuation points. In carrying out the algorithm (3.3), the value of $d-d^*$ is determined at each grid point; at each fixed value of s then, the values of $d-d^*$ at the grid points could be used in an interpolation to approximate the break-even point $\bar{y}_\Delta(s)$. The simplest such scheme would be a linear one based on the values of $d-d^*$ at $y_0(s)$ and $y_{-1}(s) = y_0(s) + \Delta^{\frac{1}{2}}$, the grid levels between which it is known that the break-even point $\bar{y}_\Delta(s)$ lies. Alternately, one might employ a quadratic interpolation scheme based on the values of $d-d^*$ at either $y_1(s)$, $y_0(s)$ and $y_{-1}(s)$ or $y_0(s)$, $y_{-1}(s)$ and $y_{-2}(s)$. Day (1969, p.306) points out that for two-sided problems with normal increments where $d-d^*$ is symmetric and convex (in y at each s) and has a monotone decreasing second derivative, these two quadratic interpolations will actually yield an underestimate and an overestimate respectively of the break-even point $\bar{y}_\Delta(s)$. This suggests, and we shall use, the average of the two interpolated values as the approximation. The estimates of $\bar{y}_\Delta(s)$ described here can be adjusted by $0.5\Delta^{\frac{1}{2}}$ to give variations of A which may be called LA and QA for linear adjusted and quadratic adjusted.

Each of the above adjustment methods involves adjusting an estimate of $\bar{y}_\Delta(s)$. It is possible, at considerable computational expense, to approximate the points $\bar{y}_\Delta(s)$ more precisely by repeating the discrete backward induction with each of a series of related grids. By using the grid

$$(3.6) \quad ((y, s): s = s_1 + i\Delta, y = c + k\Delta^{\frac{1}{2}}; i = 0, 1, \dots, k = 0, \pm 1, \dots)$$

with many fractional values of $c/\Delta^{\frac{1}{2}}$ (without loss of generality, we assume $0 \leq c < \Delta^{\frac{1}{2}}$), one can estimate the break-even points $\bar{y}_\Delta(s)$ arbitrarily well.

We now consider EX, an analogue of the extrapolation method E, which bypasses the explicit calculation of $\tilde{y}_\Delta(s)$. Defining

$$D(y,s) = \hat{d}(y,s) - d(y,s) ,$$

where \hat{d} is the optimal risk in the discrete time problem (the function evaluated by the algorithm (3.3)), the results of Chernoff & Petkau (1976) indicate that for the one-sided problem under discussion, at each fixed value of s , one should expect

$$(3.7) \quad D(\tilde{y}(s) + v\Delta^{1/2}, s) \approx -H(\tilde{y}(s), s)r(v)\Delta$$

where the function $r(v)$ is given by

$$(3.8) \quad \begin{aligned} r(v) &= 0 & \text{for } v > -1/2 , \\ &= v^2 - \inf_j (v + j)^2 & \text{for } v \leq -1/2 , \end{aligned}$$

and

$$(3.9) \quad H(y,s) = \frac{1}{2} d_{yy}(y,s) - d_s(y,s)$$

is the "rate of losing". Suppose then that the algorithm (3.3) has been carried out and we wish to approximate $\tilde{y}(s)$. At those values on the s grid, the value of $D(y,s)$ is available at $y_j(s)$ for $j = 0, 1, \dots$ (of course, $D \equiv 0$ at $y_j(s)$ for $j = -1, -2, \dots$). If we represent

$$(3.10) \quad y_0(s) = \tilde{y}(s) + v\Delta^{1/2} ,$$

we only require an approximation for v . Fitting the known values of D at $y_0(s)$ and $y_1(s)$ to the relation (3.7) leads to

$$(3.11) \quad \begin{aligned} D_0 &\equiv D(y_0(s), s) = ar(v) , \\ D_1 &\equiv D(y_1(s), s) = ar(-1 + v) , \end{aligned}$$

where the unknown constant $a = -H(\tilde{y}(s), s)$. Assuming as suggested by (3.4) that $-1/2 < v \leq -1/4$, (3.11) becomes

$$(3.12) \quad \begin{aligned} D_0 &= a(v^2 - (v+1)^2) = -a(2v+1) , \\ D_1 &= a((v-1)^2 - (v+1)^2) = -a(4v) . \end{aligned}$$

Solving the system (3.12) leads to the approximations

$$(3.13) \quad \begin{aligned} a &= -D_1/4v , \\ v &= D_1/2(2D_0 - D_1) ; \end{aligned}$$

this value for v is then substituted into (3.10) to yield the extrapolation estimate of $\tilde{y}(s)$; this method is labelled EX. Note that in the special case where $y_0(s)$ is itself a break-even point, $D_0 = 0$ and this extrapolation scheme calls for estimating $\tilde{y}(s)$ as $y_0(s) + 0.5\Delta^{1/2}$, while in the case $D_0 < 0$ the scheme

calls for a correction which is larger than $0.5\delta^{\frac{1}{2}}$; these properties agree with what is suggested by (3.4).

In summary, the technique which we propose to employ to solve the general optimal stopping problem for a zero drift Wiener process in the (y,s) scale defined earlier is as follows: The Wiener process $Y(s)$ is approximated by a discrete time simple random walk process and backward induction is employed to solve the optimal stopping problem for this discrete time problem. The solution of the discrete time problem is then adjusted by one of the methods CA, LA, QA or EX to approximate the solution of the Wiener process problem. In the above, details of the methods of adjusting the discrete time solution have been discussed in the context of a one-sided problem with a monotone increasing boundary. It should be clear that the same methods can be used in problems with more complicated types of optimal continuation regions. Further, it should not be surprising that exactly the same techniques can be employed to solve the general optimal stopping problem for a zero drift Wiener process in the (x,t) scale defined earlier.

The reader will already have noticed that while we have dwelt at some length on adjusting the boundary of the optimal continuation region for the discrete time problem to provide an improved approximation to the boundary of the optimal continuation region for the continuous time problem, nothing has been said about how one might similarly adjust the optimal risk. In order to do so, a relationship between the discrete and continuous time risks analogous to the relationship between the boundaries given by (3.4) would be necessary; unfortunately, no such relationship is known at present.

In the next section, these techniques will be illustrated on some of the examples described in Section 2; the behavior of the optimal discrete time risk as an approximation (unadjusted) to the optimal continuous time risk will also

be considered. We remark that for certain problems there are a number of ways of reducing the labor involved in carrying out the backward induction algorithm (3.3); these typically depend upon the particular problem under consideration and will be discussed as the opportunity arises in the next section.

4. ILLUSTRATION OF TECHNIQUES

In this section we illustrate the behavior of the general technique described in the previous section in the context of some of the examples presented in Section 2; each example has its own particular features, but the basic algorithm is in every case the same. While application of this technique to derive refined estimates of the optimal boundary and risk for the continuous time problem would require an exorbitant amount of computation, nevertheless, it is extremely easy to program and relatively coarse grids on the s axis yield surprisingly accurate estimates.

The (y,s) problems which have been described all have the property that the interval of possible values of s is infinite. For these statistical problems, the region of large values of s is of particular interest since it corresponds in each case to the "beginning" of the problem where little information is yet available. The question of how one obtains estimates in a practical manner for large values of s will be discussed in the next section; in this section we restrict attention in each case to the interval $100 \geq s \geq s_1$.

We begin with the examples for which exact solutions are known; these permit a careful examination of the convergence of the estimates as the grid spacing is refined. We then discuss the implementation of the techniques for the other examples and present a few results.

Example 2.5. Modified Anscombe problem. This problem is symmetric in y with optimal continuation region $C = \{(y,s): |y| \leq \bar{y}(s), s \geq 1\}$, where the monotonically increasing boundary $\bar{y}(s)$ is specified by $1 - \phi(\bar{y}(s)/s^{1/2}) = s^{-1/2}$. Note that $\bar{y}(s) \sim (\pi/2)^{1/2}(s-1)$ as $s \rightarrow 1$ and $\bar{y}(s) \sim (2s \log s)^{1/2}$ as $s \rightarrow \infty$.

Consider carrying out the algorithm (3.3) to determine the solution of the corresponding random walk problem, using a grid of the form (3.6) for some value of c . The computation proceeds in stages: At the initial (zero-th) stage, the values of \hat{d} are assigned at all points of the grid corresponding to the final value of s , namely $s = s_1 = 1$. At the k th stage, \hat{d} has already been evaluated at all points of the grid corresponding to the values of $s = 1 + j\Delta$, for $j = 0, 1, \dots, k-1$; \hat{d} is then evaluated at all points of the grid corresponding to $s = 1 + k\Delta$. In the course of this computation which yields the optimal risk for the random walk problem, each of the individual grid points is classified as either a stopping point or a continuation point for the random walk problem. Thus, the sequence of highest continuation levels corresponding to the particular grid being employed are determined and any of the methods described in the previous section can be employed to obtain an approximation to the continuous time boundary $\bar{y}(s)$.

While this computation is straightforward, there are a number of fairly obvious modifications which reduce the amount of computation involved in carrying out the algorithm (3.3) for this particular problem. First, due to the symmetry, we have $\hat{d}(-y,s) = \hat{d}(y,s)$ at each s . Using a grid which is symmetric about $y = 0$ (use of $c = 0$ in (3.6)) then allows attention to be restricted to the positive y half-plane. Second, it is intuitively obvious and easy to show that the sequence of break-even points for the random walk problem inherits the monotonicity property of the continuous time boundary $\bar{y}(s)$. Thus, at stage k where $s = 1 + k\Delta$, the grid levels $0, \Delta^{1/2}, 2\Delta^{1/2}, \dots, y_0(1 + (k-1)\Delta)$, where $y_0(1 + (k-1)\Delta)$ is the highest continuation level corresponding to $s = 1 + (k-1)\Delta$, are known to be in the continuation region. At stage k then,

\hat{d} can simply be assigned the value of d^* (see (3.5)) at these grid levels. The minimization indicated by the algorithm (3.3) need be carried out only at successively higher grid levels until the first stopping point is encountered; all higher grid levels will also be within the stopping region for the random walk problem. In fact, for reasonable values of Δ , the minimization need be carried out only at the single grid level $y_{-1}(1 + (k - 1)\Delta) = y_0(1 + (k - 1)\Delta) + \Delta^{\frac{1}{2}}$ since if the highest continuation level does change at stage k , it will change from $y_0(1 + (k - 1)\Delta)$ to $y_{-1}(1 + (k - 1)\Delta)$.

These computations have been carried out for a sequence of grids specified by decreasing values of the grid spacing Δ . Since the only apparent pattern in the size of the errors, $e = \hat{y} - \bar{y}$ of estimation of the continuous time boundary $\bar{y}(s)$ was a very slight tendency for the errors to decrease as s increased, an overall summary should be an adequate description. Such an overall summary for each of the methods CA, LA, QA and EX is provided in Table 1.

Examination of Table 1 reveals that while methods CA and QA underestimate the correction required to approximate the continuous time boundary for coarse grid spacings and overestimate it for the (more reasonable) finer grid spacings, method LA overestimates the correction for all spacings considered. Method EX underestimates the correction for coarse grid spacings, but this bias begins to disappear as the spacing is refined. Perhaps the most important observation to be made about Table 1, however, is the apparent relationship between the size of the errors made and the grid spacing for method EX: refining the grid spacing in s by a factor of 4 appears to reduce the size of the errors, as measured by either $\text{Ave}(|e|)$ or $\text{Max}(|e|)$, by a factor of between 3 and 4 (note that if the factor truly is 4, this implies the size of the errors is proportional to the grid spacing in s). Since refining the grid spacing in s by a factor of 4 involves

8 times as much computational work, this leads to the rough estimate of 2.8 to 3.7 times as much work required to reduce the size of the errors by a factor of 2 when method EX is employed. Although it is clear that the size of the errors made by the other methods will also decrease as the spacing is refined, the actual behavior is unpredictable since no such empirical relationship is obvious for these other methods. The table clearly indicates that while methods CA and EX should not be used with coarse grids, these become the preferred methods with the (more reasonable) refined grids. It should be noted that all four methods provide excellent approximations to the continuous time boundary $\bar{y}(s)$ when reasonable grid spacings are employed.

The optimal risk for the discrete time simple random walk problem was also examined as an approximation to the optimal risk for the continuous time Wiener process problem. A crude summary of the errors in this approximation is presented in Table 2. This summary indicates that refining the grid spacing in s by a factor of 4 leads to a reduction in the size of the errors by a factor of between 3 and 4 also. Further, the table clearly indicates that the risk in the discrete time simple random walk problem provides an excellent approximation to the optimal risk for the continuous time problem, even for quite coarse grids.

We remark that in contrast to the continuous time problem, the random walk problem under consideration here has the property that the continuation region is prematurely truncated; that is, there exists an interval on the s -axis, $(1, s_f(\Delta))$, on which none of the grid points will be classified as continuation points. An easy calculation indicates that, for small values of Δ , the grid point $y = 0$, $s = 1 + k\Delta$ will first (as successive stages of the backward induction are carried out) belong to the continuation region for the random

walk problem when $k = (2\pi\Delta)^{-1/2} + 1 + o(1)$. While this represents a substantial number of successive stages only for very small values of Δ , this feature could also be incorporated to make the computation more efficient for such small values of Δ .

Example 2.6. Van Moerbeke's gambling problem. This gambling problem has a one-sided continuation region $C = \{(x, t): x \geq \bar{x}(t), t \leq 1\}$ with monotonically increasing boundary $\bar{x}(t) = -\alpha(1-t)^{1/2}$, where $\alpha = 0.5061$ is the solution of $\alpha\phi(\alpha) = \phi(\alpha)$. Although this problem is formulated in terms of the function $g(x, t)$ which specifies the reward received by the gambler upon stopping at (x, t) and is given in (2.10), the problem can be equivalently formulated in terms of the stopping cost function

$$d(x, t) = -g(x, t).$$

The appropriate modification of the algorithm (3.3) is then given by

$$(4.1) \quad \begin{aligned} \hat{d}(x, t) &= d(x, t) && \text{for } t = 1, \\ &= \min[d(x, t), (\hat{d}(x+\Delta^{1/2}, t+\Delta) + \hat{d}(x-\Delta^{1/2}, t+\Delta))/2] && \text{for } t < 1. \end{aligned}$$

While the first few stages of this algorithm can be carried out analytically and lead to break-even points $\bar{x}_\Delta(1-\Delta) = 0$, $\bar{x}_\Delta(1-2\Delta) = (-2+3^{1/2})\Delta^{1/2} = -0.268\Delta^{1/2}$, $\bar{x}_\Delta(1-3\Delta) = -(4-10^{1/2})\Delta^{1/2}/2 = -0.419\Delta^{1/2}$, and so on (note that applying the $\frac{1}{2}\Delta^{1/2}$ correction to these exact break-even points would lead to estimates of the continuous time boundary of $\hat{\bar{x}}(1-\Delta) = -0.500\Delta^{1/2}$, $\hat{\bar{x}}(1-2\Delta) = -0.543(2\Delta)^{1/2}$, $\hat{\bar{x}}(1-3\Delta) = -0.531(3\Delta)^{1/2}$, and so on), these exact calculations become unmanageable after a few stages.

Carrying out the algorithm (4.1) proceeds similarly as in the case of Example 2.5 and any of the methods described in the previous section can be employed to obtain an approximation to the continuous time boundary $\bar{x}(t)$. The present problem, however, has its own special features. For the continuous time problem, $\{(x, t): x > 0, t \leq 1\} \subset C$; this fact would be known even if the exact solution were unknown since the "rate of winning", $H(x, t) = \frac{1}{2}g_{xx}(x, t) + g_t(x, t) > 0$ for $x > 0, t \leq 1$. Since it can easily be shown that the sequence of break-even points for the random walk problem inherits the monotonicity property of the continuous time boundary $\bar{x}(t)$, and since $\bar{x}_\Delta(1-\Delta) = 0$, the above result is also true for the discrete time random walk problem. Again, it can be shown that the minimization indicated by the algorithm (4.1) need only be carried out at a single grid level at each stage of the computation.

The fact that all grid points above the x axis are known to be continuation points can be incorporated to reduce the amount of computation required in carrying out the algorithm (4.1). Consider a particular path of the random walk process originating at the point $(x, t) = (c + \Delta^{1/2}, 1 - n\Delta)$. The path of the process could hit the grid level $x = c$ for the first time at $t = 1 - (n-1)\Delta, 1 - (n-3)\Delta, \dots$; alternately, the path could remain above the line $x = c$ all the way to $t = 1$. Since all the grid points $(c, 1 - \ell\Delta)$ for $\ell = 1, 2, \dots$ are continuation points, we have the relation

$$(4.2) \quad \hat{d}(c+\Delta^{1/2}, 1-n\Delta) = \sum_{m=1}^n p_m \hat{d}(c, 1-(n-m)\Delta) + \sum_{k=1}^{\infty} q_{n,k} \hat{d}(c+k\Delta^{1/2}, 1)$$

where p_m is the probability that a simple random walk starting at the origin first passes through the level -1 at the m^{th} step, and $q_{n,k}$ is the probability

that a simple random walk starting at the origin stays above the level -1 for the first n steps and achieves level $k - 1$ at the n th step. Feller (1968, p.89, Theorem 2) provides

$$p_m = \frac{1}{m} \binom{m}{(m+1)/2} 2^{-m} \quad \text{for } m \text{ odd,}$$

$$= 0 \quad \text{for } m \text{ even;}$$

for m positive, $p_{m+2} = mp_m/(m+3)$ with $p_1 = 1/2$ and $p_0 = 0$. Feller (1968, p.73, Ballot Theorem) also provides

$$q_{n,k} = 0 \quad \text{for } n+k \text{ even,}$$

$$= \frac{k!}{n+1} \binom{n+1}{(n+k+1)/2} 2^{-n} \quad \text{otherwise.}$$

The relation (4.2) provides a slight modification for carrying out the backward induction which we will call the truncation modification. At the initial (zero-th) stage, that is at $t = 1$, the risks are specified by $d(x, t)$. At any subsequent stage, corresponding to $t = 1 - n\Delta$ say, compute the risk at the grid level $x = c + \Delta^{1/2}$ by means of (4.2). The risks at the grid levels $x = c + k\Delta^{1/2}$ for $k = 0, -1, -2, \dots$ can be computed using the algorithm (4.1) as described above.

Returning for a moment to the continuous time problem, we have already pointed out that changing the stopping reward function by adding to it any solution of the backward heat equation leaves the optimal continuation region unchanged. For present purposes, it is convenient to consider the new stopping

reward function $g'(x, t)$ defined by

$$g'(x, t) = g(x, t) - 2(1 - t + x^2),$$

or the new stopping cost function $d'(x, t) = -g'(x, t)$. Note that $d'(x, 1) \equiv 0$ for $x \geq 0$. The algorithm (4.1) can be employed to obtain the optimal risk $\hat{d}'(x, t)$ for the discrete time random walk problem corresponding to this version of the continuous time problem; in this case the relation (4.2) simplifies to

$$(4.3) \quad \hat{d}'(c + \Delta, 1 - n\Delta) = \sum_{m=1}^n p_m \hat{d}'(c, 1 - (n - m)\Delta)$$

which results in a reduction in the computation involved in carrying out the algorithm. Limited empirical evidence suggests that the truncation modification reduces the computation time required by a factor of approximately two in those cases where the simplification (4.3) obtains.

In the general case, the transformation

$$g'(x, t) = g(x, t) - \int_{x_1}^{\infty} (t_1 - t)^{-1/2} \phi((x' - x)/(t_1 - t)^{1/2}) g(x', t_1) dx'$$

produces a new stopping reward function with the same optimal continuation region and satisfying $g'(x, t_1) = 0$ for $x \geq x_1$. Unless this integral can be explicitly evaluated, however, no real simplification obtains. For our special function g in (2.10), this integral (with $x_1 = 0$, $t_1 = 1$) does not coincide with $2(1 - t + x^2)$, but the difference is simply a solution of the backward heat equation.

The computations have been carried out for a sequence of grids specified by decreasing values of the grid spacing; in all cases, grids of the form (3.6) with $c = 0$ were employed. Since there was no apparent pattern in the

size of the errors $e = \hat{x} - \bar{x}$ of estimation of the continuous time boundary $\bar{x}(t)$, an overall summary of these errors should be adequate. Such an overall summary for each of the methods CA, LA, QA and EX appears in Table 3.

Table 3 reveals that while methods CA, LA and QA always overestimate the correction required to approximate the continuous time boundary (except at the coarsest grid spacing in the case of CA), such a severe bias is not apparent with EX although the method does tend to underestimate the correction required. The relationship between the size of the errors made and the grid spacing is quite clear for methods CA, LA and QA: refining the grid spacing in s by a factor of 4 appears to reduce the size of the errors by a factor of 2; for method EX the reduction factor appears to be about 3. While all methods provide excellent approximations to the continuous time boundary $\bar{x}(t)$ when employed with reasonable grid spacings, the preferred method would appear to be EX.

A crude summary of the errors in the approximation of the continuous time risk by the optimal risk in the discrete time random walk problem is presented in Table 4; it is apparent that this approximation is excellent even for relatively coarse grids. Further, it is clear that refining the grid spacing in t by a factor of 4 leads to a reduction in the errors by a factor of 4. It is interesting to note that in this example it appears the various discrete time random walk problems are uniformly less favourable than the continuous time problem. Examination of isolated grid points indicates that the risk in the discrete time problem converges monotonically to the continuous time risk. These observations are in contrast to the situation in Example 2.5.

Recall that the methods CA, LA, and QA proceed in two stages: first the break-even points for the discrete time random walk problem are approximated; these are then adjusted by $0.5\Delta^{1/2}$ as suggested by the asymptotic relationship

(3.4). At an early stage of these investigations, the performance of this adjustment of $0.5\Delta^{1/2}$ was investigated in the context of Example 2.6. For a fixed grid spacing Δ , the results of carrying out the backward induction with grids of the form (3.6) with $c = 0(0.0001)\Delta$ were combined to locate the break-even points to within an error of 0.0001 at the expense of a very substantial amount of computing. The errors in the approximation of the continuous time boundary by both the "raw" break-even points and the "adjusted" break-even points (adjusted = raw - $0.5\Delta^{1/2}$) were then evaluated. The results for a few grid spacings are summarized in Table 5. Note that $\text{Ave}(|e|)$ and $\text{Max}(|e|)$ are similar throughout the table; this indicates that the errors are roughly constant at different values of t . As expected on the basis of (3.4), the errors with the raw break-even points are very close to $0.5\Delta^{1/2}$. While the adjustment of $0.5\Delta^{1/2}$ is slightly too large for each grid spacing, this error seems to decrease faster than $0.5\Delta^{1/2}$ as the grid spacing decreases ($393/5000 = 0.079$, $156/2500 = 0.062$, $58/1250 = 0.046$). Comparing these results to those in Table 3, it becomes clear that, for this problem, while method LA does not estimate the break-even points very accurately, QA does reasonably well, particularly for the coarser grid spacings. Method CA always underestimates the break-even points (for this problem and generally) and this compensates for the fact that $0.5\Delta^{1/2}$ is an over-adjustment here. It is important to note that the errors incurred with method EX are very similar to the errors reported in Table 5 (compare especially $\text{Max}(|e|)$); for this problem method EX does as well as any possible method based upon adjusting estimated break-even points.

These methods can be adapted for all of the examples we have discussed. The methods employed in Example 2.5 apply without modification to Example 2.3.

A slight modification was required for Example 2.4; attention could not be restricted to the positive y half-plane since the problem was not symmetric in y . Detailed results for these examples have already appeared in Chernoff & Petkau (1981) and Petkau (1978) respectively. In the remainder of this section we examine the behavior of these methods in Examples 2.1, 2.2, and 2.7.

Example 2.1. Sequential analysis problem. This problem is

symmetric in y with optimal continuation region

$C = \{(y, s): |y| \leq \bar{y}(s), s \geq 0\}$. Asymptotic expansions for the

monotonically increasing boundary demonstrate that $\bar{y}(s) \sim \frac{1}{4}s^2$ as $s \rightarrow 0$ and

$\bar{y}(s) \sim (3s \log s)^{1/2}$ as $s \rightarrow \infty$. The methods employed in Example 2.5 apply

without modification, and the random walk version of this problem is also

naturally truncated; an easy calculation indicates that, for small values of

Δ , the grid point $y = 0, s = k\Delta$ will first belong to the continuation region when $k \approx 2^{1/2}\Delta^{-3/4} + 2^{-5/4}\pi^{-1/2}\Delta^{-3/8} + \dots$.

Although the desired computations can be carried out in a straightforward manner, it is more difficult to examine the performance of the methods since the exact solution to the continuous time problem is unknown. To illustrate behavior as the grid spacing decreased, the approximation to the continuous time solution provided by a given method with the most refined grid spacing was taken as a baseline for that method. The deviation of the approximation obtained with a less refined grid spacing from this baseline is summarized in Table 6. The disparity among the approximations obtained by the different methods with each spacing employed is summarized in Table 7. Table 7 indicates clearly that, in this example, the approximations to the continuation regions for the continuous time problem produced by methods LA and QA are strictly larger than that produced by EX; the same tendency can be noted for CA.

Relative to the size of the grid spacing, the methods CA, QA and EX agree quite well for the smaller grid spacings. Table 6 indicates that while methods CA and EX improve dramatically as the spacing is refined, the improvement is less dramatic for LA and QA. Overall, the patterns here appear to be very similar to those observed in Example 2.5.

The convergence of the optimal risk in the random walk problem as the grid spacing decreased was also examined. The optimal risk with the most refined grid spacing was taken as the baseline. The results as summarized in Table 8 and are not unlike the results obtained in Example 2.5.

Example 2.2. One-armed bandit problem. This problem has a one-sided continuation region $C = \{(y, s): y \geq \bar{y}(s), s \geq 1\}$ with a monotonically decreasing

boundary $\bar{y}(s)$. Asymptotic expansions demonstrate that $\bar{y}(s) \sim -\alpha(s-1)^{1/2}$ as

$s \rightarrow 1$, where $\alpha = 0.63884$ is the solution of $(\alpha^2 - 1)\phi(\alpha) + \alpha^3\Phi(\alpha) = 0$, and

$\bar{y}(s) \sim -(2s \log s)^{1/2}$ as $s \rightarrow \infty$. The first few stages of the backward induction

algorithm lead to break-even points $\bar{y}_\Delta(1 + \Delta) = 0$, $\bar{y}_\Delta(1 + 2\Delta) = -\Delta^{1/2}(1 + 2\Delta)/(3 + 2\Delta)$, $\bar{y}_\Delta(1 + 3\Delta) = -4\Delta^{1/2}(1 + \Delta)(1 + 3\Delta)/(7 + 15\Delta + 6\Delta^2)$, and so on. Addition of any solution

of the forward heat equation to the stopping cost $d(y, s)$ given in (2.4)

leaves the optimal continuation region of the continuous time problem unchanged.

Upon converting to the new stopping cost function $d'(y, s) = d(y, s) + y$, for

which $d'(y, 1) \equiv 0$, the methods employed in Example 2.6 apply to this example

without modification. The results for this example are summarized in Tables

9, 10 and 11. Overall, the results are quite similar to those for Example 2.1.

Example 2.7. The S_n/n problem with finite horizon. Since the "rate of

winning" for this gambling problem is positive for negative x , this region

is contained within the optimal continuation region. As would be anticipated,

this problem has a one-sided continuation region $C = \{(x, t): x \leq \bar{x}(t), 0 < t \leq 1\}$; asymptotic expansions demonstrate that $\bar{x}(t) \sim \alpha_0 t^{1/2}$ as $t \rightarrow 0$, where $\alpha_0 = 0.83991$ is the solution of $\alpha\phi(\alpha) + (\alpha^2 - 1)\phi(\alpha) = 0$, and $\bar{x}(t) \sim \alpha_1(1-t)^{1/2}$ as $t \rightarrow 1$, where $\alpha_1 = 0.63884$ is the same constant which appears in the asymptotic expansion of the boundary for the one-armed bandit problem. The first few break-even points for the random walk problem are given by $\bar{x}_\Delta(1-\Delta) = 0$, $\bar{x}_\Delta(1-2\Delta) = \Delta^{1/2}(1-2\Delta)/(3-2\Delta)$, $\bar{x}_\Delta(1-3\Delta) = 4\Delta^{1/2}(1-\Delta)(1-3\Delta)/(7-15\Delta+6\Delta^2)$, and so on. Since the sequence of break-even points will not be monotone, slightly more detailed calculations are necessary when carrying out the backward induction than was the case in Example 2.6. However, converting to the new stopping reward function $g'(x, t) = g(x, t) - x$, for which $g'(x, 1) = 0$, allows the general technique employed in Example 2.6 to be used here also. The results for this example are summarized in Tables 12, 13 and 14. While the results are qualitatively similar to those in the previous examples, a few features should be noted. Since the optimal boundary is dome-shaped, it is clear that method CA, which approximates this curved surface by a flat surface in the region of the maximum, must do poorly for coarse grid spacings. While both LA and QA produce smooth approximations, method EX produces approximations which occasionally exhibit a lack of smoothness in the neighbourhood of values of t at which the highest continuation level changes; this tendency is most pronounced with coarse grid spacings but persists even with refined grid spacings. Further, since the optimal risk approaches infinity as $t \rightarrow 0$, the deviations summarized in Table 14 become large at the smaller values of t ; indeed, the deviation which is largest in magnitude in each case occurs at $t = 0.04$, $x = 0.1$. In spite of these limitations, the results presented again indicate that the methods perform

quite well. In Table 15, we present an abbreviated table of the approximation to the boundary of the optimal continuation region for the continuous time problem obtained from the computation with the most refined grid spacing. Note the accuracy of the 1-term asymptotic expansions given above as $t \rightarrow 0$ and $t \rightarrow 1$.

5. PRACTICAL IMPLEMENTATION IN STATISTICAL PROBLEMS

The continuous time problems described in Examples 2.1-2.4 arise from statistical problems and share the property that the range of possible s values is infinite. In this section we indicate how the numerical methods which have been described can be employed to obtain estimates of the stopping boundary and the Bayes risk for these problems in the region of large values of s . The properties of the proposed technique will be examined in the context of Example 2.5, the modified Anscombe problem, and summaries of the estimates obtained for both the sequential analysis problem and the one-armed bandit problem will be presented.

While the results of the previous section establish that estimates obtained with the numerical methods are accurate provided that small grid spacings Δ are employed, the use of a small grid spacing in a backward induction designed to obtain estimates for large values of s , say out as far as $s = 10^6$, would require an exorbitant amount of computer time. On the other hand, while the use of a large grid spacing will allow the determination of reasonably good estimates at large values of s , the estimates obtained at smaller values of s would typically be poor. A hybrid technique which uses a small grid spacing at the initial stages of the backward induction and larger grid spacings at larger values of s is required.

A naive technique of this sort would consist of carrying out a number of separate backward inductions, the first with a very small value of Δ and successive ones with successively larger values of Δ . Each of these backward inductions would begin at s_1 , the initial value of s , and if each was carried out to the same number of stages, estimates would be obtained in successively larger overlapping intervals of s . The results of the separate backward

inductions could then be combined; at any fixed value of s , the estimates would be obtained from the backward induction involving the smallest value of Δ to reach this value of s . Thus, in different intervals of s , the estimates of the Bayes risk and the stopping boundary for the continuous time problem are the estimates obtained in different approximating discrete time simple random walk problems. While this simple technique seemed to lead to adequate estimates in Petkau (1978), estimates at large values of s might be unnecessarily crude since these are obtained by backward inductions which use fairly large values of Δ even at the initial stages.

A simple way of avoiding this difficulty is to carry out a single backward induction that incorporates a changing step size as it proceeds. The first phase of this backward induction might execute M_1 stages corresponding to a very small grid spacing Δ_1 , from the initial value s_1 to $s_1 + M_1 \cdot \Delta_1 = s_1^*$ say, and the second phase might execute M_2 stages corresponding to a larger grid spacing Δ_2 , from the initial value for this phase of $s_2 = s_1^*$ to $s_2 + M_2 \cdot \Delta_2 = s_2^*$ say. At the first stage of the second phase, estimates of the risk at all the new grid levels at s_2 could be interpolated from the estimates of the risk at the old grid levels at s_1 . The backward induction could be continued for as many phases as desired; an interpolation of the estimates of the risk would be required at the first stage of each successive phase. Of course, the estimates of the Bayes risk and the stopping boundary for the continuous time problem which are obtained in this way do not correspond, except in first phase, to the estimates which would be obtained from any particular approximating discrete time simple random walk. On the other hand, this technique should lead to more accurate estimates at large values of s than the naive

technique described above; since very small values of Δ could be employed in the initial phases, this would insure that the computations at later phases of the backward induction would be based on excellent approximations to the Bayes risk for the continuous time problem at the earlier phases.

Implementation immediately revealed that this technique led to slight discontinuities in the estimates of both the Bayes risk and the stopping boundary for the continuous time problem at the values of s which marked the transition from one phase to the next. To overcome this difficulty the technique was modified to have successive phases carried out on overlapping intervals of s . Specifically, the first phase is carried out as described above, but the initial value s_2 for the second phase would no longer be s_1^* but rather some value of s intermediate between s_1 and s_1^* . The estimates of the risk obtained at this intermediate value of s would be stored during the course of the computations in the first phase, enabling the interpolation necessary at the first stage of the second phase to be carried out. The estimates of the Bayes risk and the stopping boundary for the continuous time problem at values of s in the overlapping region would be those obtained with the finer grid spacing; that is, those obtained in the earlier phase. This modification would be implemented at the transition from each phase to the next, and the backward induction could be continued for as many phases as desired. Empirical evidence indicated that for all practical purposes this modification removes the observed discontinuities provided that the overlapping interval corresponds to a sufficient number of stages of the next phase. Although what constitutes a sufficient number of stages depends upon the particular problem, our experience suggests that an interval corresponding to a hundred stages of the next phase would certainly be adequate.

This technique was employed in Chernoff and Petkau (1981), and will also be employed here. The overall mechanics of the proposed technique are specified by the grid spacing to be used and the number of stages to be executed in each phase of the backward induction and the extent of overlapping to be employed from each phase to the next. We have not systematically explored the possible versions of the technique, but rather have used the simple version in which the number of stages to be executed is the same in all phases, the grid spacing is increased by a constant multiple from one phase to the next, and the extent of overlapping is a fixed fraction of the interval of s values over which the stages of the previous phase were executed.

The results presented in the following were obtained using the technique with 2080 stages in each phase, the grid spacing Δ increased by a factor of 4 from each phase to the next, and the extent of overlapping corresponding to one-half the interval of s values covered by the previous phase; this extent of overlapping corresponds to 1040 stages of the previous phase or, since Δ is increased by a factor of 4 from each phase to the next, 260 stages of the current phase. Since only grids centered on the y -axis were employed (use of $c=0$ in (3.6)), use of the factor 4 for increasing Δ from phase to phase implies that the grid at the previous phase is a refinement of the grid at the current phase. Consequently, the estimates of the risk at the new grid levels at the value of s corresponding to the first stage of any phase are provided by the estimates of the risk at those same grid levels at that value of s in the previous phase; no interpolation is necessary.

For each example, the grid spacing for the first phase was taken to be $\Delta = 25 \times 10^{-6}$ and estimates were obtained out to $s = 10^6$. The estimates of the risk were obtained as described above. Estimates of the boundary were printed out whenever the grid level corresponding to the last continuation

level changed; the estimates at these values of s were obtained by method EX. Subsequent to the completion of the backward induction, an estimate of the stopping boundary at any fixed value of s can be obtained by interpolation from this listing. Where a tabulation of the stopping boundary is provided in the following, linear interpolation has been employed.

Since the solution of Example 2.5, the modified Anscombe problem, is available in closed form (see (2.8) and (2.9)), the behavior of the above technique can be examined in detail. A crude summary of the accuracy of the estimates of the risk within the continuation region is provided in Table 16. This summary suggests that the relative errors tend to be largest close to $y = 0$, where they are of the order of 10^{-4} ;

detailed examination of the errors on a much finer grid of (s, z) values, where $z = y/s^{1/2}$, suggests the empirical upper bound of 3×10^{-4} on the relative errors in this problem when the proposed technique is employed in the manner described above. A summary of the errors in the estimate of the stopping boundary is provided in Table 17. The largest relative errors (which are still relatively small) occur in the region of s values close to 1 where the stopping boundary $\tilde{y}(s)$ is close to 0; in this region, asymptotic expansions would be available for $\tilde{y}(s)$ and could complement the numerical results. For this problem, the relative errors decrease slightly across phases until the grid spacing exceeds 1 when they begin to increase again. The errors themselves increase across phases roughly in proportion to Δy , the size of grid spacing on the y -axis (at least as long as the grid spacing is less than 1; after this the rate of increase appears to be a bit faster). It is interesting to note that the estimates of the boundary are always overestimates (errors > 0) for phases 1-6 and underestimates (errors < 0) for phases 9-13.

Since this performance of the proposed technique was judged to be adequate for our purposes, the technique was implemented in exactly the same fashion for Examples 2.1 and 2.2. Detailed estimates of the Bayes risk and the stopping boundary for these problems have not been presented in the literature; we summarize the results here.

Example 2.1. Estimates of the stopping boundary for the sequential analysis problem are tabulated in various scales of interest in Table 18. The asymptotic expansions

$$\tilde{x}(s) = \tilde{y}(s)/s - \frac{1}{4} s [1 - \frac{1}{12} s^3 + \frac{7}{240} s^6 - \dots] \quad \text{as } s \rightarrow 0,$$

$$\tilde{z}(s) = \tilde{y}(s)/s^{1/2} - \frac{1}{4} s^{3/2} [1 - \frac{1}{12} s^3 + \frac{7}{240} s^6 - \dots] \quad \text{as } s \rightarrow 0,$$

$$\tilde{\beta}(s) = 1 - \phi(\tilde{z}(s)) - \frac{1}{2} - \frac{1}{\sqrt{2\pi}} \frac{1}{4} s^{3/2} [1 - \frac{7}{96} s^3 + \dots] \quad \text{as } s \rightarrow 0,$$

can be used to extend the table to even smaller values of s . On the other hand, the asymptotic expansions

$$\tilde{x}(s) = \tilde{y}(s)/s - s^{-1/2} \{ \log s^3 - \log(8\pi) - 6(\log s^3)^{-1} + \dots \}^{1/2} \quad \text{as } s \rightarrow \infty,$$

$$\tilde{z}(s) = \tilde{y}(s)/s^{1/2} - \{ \log s^3 - \log(8\pi) - 6(\log s^3)^{-1} + \dots \}^{1/2} \quad \text{as } s \rightarrow \infty,$$

$$\tilde{\beta}(s) = 1 - \phi(\tilde{z}(s)) - 2(s^3 \log s^3)^{-1/2} \{ 1 + [2 + \frac{1}{2} \log(8\pi)] (\log s^3)^{-1} + \dots \} \quad \text{as } s \rightarrow \infty,$$

perform only moderately well at $s = 10^6$.

The Bayes risk in the continuous time version of the sequential analysis problem corresponding to starting at the point (y_0, s_0) in the normalized form of the continuous time version (the coordinates y_0 and s_0 are determined by the parameters μ_0 and σ_0 of the prior distribution for μ together with the parameters k , c and σ ; see the discussion of Example 2.1 in Section 2) is given by

$$k^{2/3} c^{1/3} \sigma^{2/3} [E(d(Y(S), S)) - s_0^{-1}],$$

and the contribution to the Bayes risk of the cost of sampling is given by

$$k^{2/3} c^{1/3} \sigma^{2/3} [E(S^{-1}) - s_0^{-1}],$$

where S is an optimal stopping rule and $d(y, s)$ is given in (2.2). The Bayes expected sample size is obtained by dividing the Bayes expected cost of sampling by c , the cost of sampling associated with a single observation. $E(d(Y(S), S))$ can be approximated by the techniques described above, and the expectation $E(S^{-1})$ can be approximated in a straightforward manner during the execution of the backward induction. For simplicity in tabular presentation we may use the normalization

$$\begin{aligned} \text{BR} &= \text{Bayes risk} / k^{2/3} c^{1/3} \sigma^{2/3} \\ &= E(d(Y(S), S)) - s_0^{-1}, \end{aligned}$$

and

$$\begin{aligned} \text{ECS} &= \text{Bayes expected cost of sampling} / k^{2/3} c^{1/3} \sigma^{2/3} \\ &= E(S^{-1}) - s_0^{-1}, \end{aligned}$$

where both BR and ECS depend only upon the initial values s_0 and y_0 or $t_0 = 1/s_0$ and $z_0 = y_0/s_0^{1/2} = \mu_0/\sigma_0$. Small representative subsets of the normalized risks BR and expected costs of sampling ECS which have been evaluated are presented in Tables 19 and 20 respectively. The behaviour of these properties of the optimal procedure is also illustrated in Figures 1-4. The Bayes risk and

expected cost of sampling are plotted against $\log(t_0)$ for a few values of z_0 in Figures 1 and 2 respectively. In Figures 3 and 4 these quantities are plotted against z_0 for a few values of t_0 . The asymptotic behavior is more clearly illustrated in Figures 5-8 where $\log \text{BR}$ and $\log \text{ECS}$ are plotted against $\log(t_0)$ and z_0 .

The tables and figures reflect the form of the asymptotic expansions

$$\begin{aligned} \text{BR} &\sim K s_0^{-1/2} \phi(z_0) & \text{as } s_0 \rightarrow \infty, \\ \text{ECS} &\sim K' s_0^{-1/2} \phi(z_0) & \text{as } s_0 \rightarrow \infty, \end{aligned}$$

provided by Chernoff (1965a). The values of BR and ECS for $s_0 = 10^6$ and $z_0 = 0$ suggest $K \approx 5.89$, $K' \approx 3.91$; regressing the values of $\text{BR}/s_0^{-1/2} \phi(z_0)$ for $s_0 = (1, 2, \dots, 10) \times 10^5$ and $z_0 = 0$ against the next term of the asymptotic expansion leads to the estimate $K \approx 5.98$.

Example 2.2. Estimates of the stopping boundary for the one-armed bandit problem are provided in Table 21. The asymptotic expansions

$$\bar{x}(s) = \bar{y}(s)/s - (s-1)^{1/2} (c_0 + (c_1 - c_0)(s-1) + \dots) \quad \text{as } s \rightarrow 1,$$

$$\bar{z}(s) = \bar{y}(s)/s^{1/2} - (s-1)^{1/2} (c_0 + (c_1 - \frac{1}{2}c_0)(s-1) + \dots) \quad \text{as } s \rightarrow 1,$$

$$\begin{aligned} \bar{\beta}(s) &= \phi(\bar{z}(s)) - \frac{1}{2} - \frac{1}{\sqrt{2\pi}} (s-1)^{1/2} (c_0 + [c_1 \\ &\quad - c_0(3 + c_0^2)/6](s-1) + \dots) \end{aligned} \quad \text{as } s \rightarrow 1,$$

where $c_0 \approx 0.63883$ and $c_1 \approx 0.23625$ are defined by

$$c_0 \phi(c_0) + \phi(c_0) = 0, \quad c_1 = 2c_0/(5 + c_0^2),$$

fit very well for values of s close to 1. Here, as in the sequential analysis

problem, the asymptotic expansions for large values of s

$$\begin{aligned}\bar{x}(s) &= \bar{y}(s)/s \sim -s^{-1/2} \{ \log s^2 - 2 \log(\log s^2) - \log(8\pi) + \dots \}^{1/2} \quad \text{as } s \rightarrow \infty, \\ \bar{z}(s) &= \bar{y}(s)/s^{1/2} \sim -(\log s^2 - 2 \log(\log s^2) - \log(8\pi) + \dots)^{1/2} \quad \text{as } s \rightarrow \infty, \\ \bar{\beta}(s) &= \phi(\bar{z}(s)) \sim 2s^{-1} \{ 1 + 2/\log s^2 - \frac{1}{8} [\log(\log s^2)/\log s^2]^2 + \dots \} \\ &\quad \text{as } s \rightarrow \infty,\end{aligned}$$

are only moderately accurate at $s = 10^6$.

The Bayes expected payoff in the continuous time version of the one-armed bandit problem corresponding to starting at the point (y_0, s_0) in the normalized form of the continuous time version (the coordinates y_0 and s_0 are determined by the parameters μ_0 and σ_0 of the prior distribution for μ together with the parameters N and σ ; see the discussion of Example 2.2 in Section 2) is given by

$$-\sigma^2 \sigma_0^{-1} s_0^{1/2} [E(d(Y(S), S)) + y_0/s_0],$$

and the Bayes expected sample size is given by

$$\sigma^2 \sigma_0^{-2} s_0 [E(S^{-1}) - s_0^{-1}],$$

where S is an optimal stopping rule and $d(y, s) = -y/s$ for $s \geq 1$ is as given in (2.4). Since the use of $d'(y, s) = d(y, s) + y = y(1 - s^{-1})$ in place of $d(y, s)$ simplified implementation of the truncation modification (see the discussion of Example 2.2 in Section 4), the computations for the Bayes risk employed the identity

$$\text{Bayes expected payoff} = -\sigma^2 \sigma_0^{-1} s_0^{1/2} [E(d'(Y(S), S)) - d'(y_0, s_0)].$$

For simplicity in tabular presentation we may use the normalization

$$\text{BEP} = \text{Bayes expected payoff} / \sigma^2 \sigma_0^{-1},$$

and

$$\text{EN} = \text{Bayes expected sample size} / \sigma^2 \sigma_0^{-2}.$$

where BEP and EN depend only upon the initial values y_0 and s_0 or $z_0 = y_0/s_0^{1/2} = \mu_0/\sigma_0$ and $t_0 = 1/s_0 = \sigma_0^{-2}/(\sigma_0^{-2} + N\sigma^{-2})$, the fraction of the total potential information which is in the prior. Small representative subsets of the normalized quantities BEP and EN are presented in Tables 22 and 23 respectively. The behavior of these properties of the optimal procedure is also illustrated in Figures 9 - 12.

The tables and figures reflect the form of the asymptotic expansions

$$\text{BEP} \sim s_0 \{ \phi(z_0) + z_0 \phi'(z_0) \} \quad \text{as } s_0 \rightarrow \infty,$$

$$\text{EN} \sim s_0 \phi'(z_0) \quad \text{as } s_0 \rightarrow \infty,$$

provided by Chernoff and Ray (1965). This behavior is more clearly illustrated in Figures 13 - 16 where $\log(\text{BEP} + 1)$ and $\log(\text{EN} + 1)$ are plotted against $\log(t_0)$ and z_0 .

6. THE ANSCOMBE PROBLEM WITH ETHICAL COST

Armitage (1963) has argued that the model of Example 2.3, the Anscombe problem, fails to deal adequately with the physician's ethical requirement that he provide his current patient with the treatment he believes to be best. This requirement often frustrates attempts to gain knowledge to benefit future patients. One way to compromise is to modify the model so that an additional ethical cost is attached to each application of a treatment which the physician believes to be inferior. In this section we present results for the special case where this ethical cost is taken to be proportional to the estimate, based on the current posterior distribution, of the inferiority $|\hat{\mu}|$. This consideration of ethical costs introduces a fundamental change in the nature of our optimal stopping problem. Fortunately it can be handled with a minor modification of our methods. The results are compared to those of Chernoff and Petkau (1981) where this ethical cost is not included in the model. We begin with a more detailed discussion of the Anscombe problem without the ethical cost.

The discrete time formulation of the model for the Anscombe problem is described briefly in Section 2. The expected loss or posterior risk associated with stopping after treating n pairs of patients has two components. The first is $E(n|\mu|)$ which represents the expected cost in patient benefit incurred during the experimental phase where n of the $2n$ patients treated were assigned to the inferior treatment, and the second is the expected cost due to the possibility of selecting the inferior treatment for the final stage and thus losing $(N - 2n)|\mu|$.

If μ is assigned an $N(\mu_0, \sigma_0^2)$ prior distribution, upon observing the differences X_1, X_2, \dots, X_n in response for the first n pairs of patients the posterior distribution of μ becomes $N(Y_n^*, s_n^*)$ where

$$(6.1) \quad Y_n^* = (\sigma_0^{-2}\mu_0 + \sigma^{-2} \sum_{i=1}^n X_i) / (\sigma_0^{-2} + n\sigma^{-2}), \quad s_n^* = (\sigma_0^{-2} + n\sigma^{-2})^{-1}.$$

For $n > m$, the marginal distribution of $Y_n^* - Y_m^*$, if we treat μ as random, is $N(0, s_m^* - s_n^*)$ and $Y_n^* - Y_m^*$ is independent of Y_m^* . Thus as sampling continues, Y_n^* behaves like a Gaussian process of independent increments starting from $Y_0^* = \mu_0$. Since the preferred choice of treatment for the remaining $N - 2n$ patients is indicated by the sign of Y_n^* , the expected loss or posterior risk associated with stopping after treating n pairs of patients is $nE(|\mu|) + (N - 2n)E[\max(0, -\text{sgn}(Y_n^*)\mu)]$ where E represents expectation with respect to the posterior distribution of μ . Straightforward calculations then lead to the expression

$$Ns_n^{1/2} \tilde{\psi}(Y_n^* s_n^{*-1/2}) - \frac{1}{2}(N - 2n)|Y_n^*|$$

for the posterior risk, where $\tilde{\psi}(u) = \phi(u) + u(\phi(u) - \frac{1}{2}) = \psi(u) + \frac{1}{2}|u|$ and ψ is defined in (2.3). Using (6.1) the posterior risk can be written as $d_1(Y_n^*, s_n^*)$, where

$$(6.2) \quad d_1(y^*, s^*) = Ns^{1/2} \tilde{\psi}(y^* s^{*-1/2}) - \sigma^2(s_*^{-1} - s^{*-1})|y^*|.$$

Here

$$(6.3) \quad s_*^{-1} = \sigma_0^{-2} + \frac{1}{2}N\sigma^{-2}$$

may be regarded as the total potential information for estimating μ . The problem of selecting the best sequential procedure for terminating the experimental phase is equivalent to the optimal stopping problem where the Gaussian process Y_n^* is observed and one selects the stopping time n to minimize the expected risk $E(d_1(Y_n^*, s_n^*))$.

A natural approximation to this discrete time problem results if the discrete sequence of partial sums ΣX_i is replaced by the continuous time Wiener process $X(t^*)$, with drift μ and variance σ^2 per unit in the t^* scale ($0 \leq t^* \leq \frac{1}{2}N$). The posterior distribution of μ , given $X(t')$ for $0 \leq t' \leq t^*$, is $N(Y^*, s^*)$, where

$$Y^* = Y^*(s^*) = (\sigma_0^{-2} \mu_0 + \sigma^{-2} X(t^*)) / (\sigma_0^{-2} + t^* \sigma^{-2}), \quad s^* = (\sigma_0^{-2} + t^* \sigma^{-2})^{-1}.$$

In parallel with the above, $Y^*(s^*)$ is a Wiener process with drift 0 and variance 1 per unit in the $-s^*$ scale, and originates at the initial point (Y_0^*, s_0^*) , where $s_0^* = \sigma_0^2$, $Y_0^* = Y^*(s_0^*) = \mu_0$. As t^* increases from 0 to $\frac{1}{2}N$, s^* decreases from s_0^* to s_* as defined in (6.3).

The posterior risk associated with stopping at (Y^*, s^*) is $d_1(Y^*, s^*)$, and this continuous time problem is equivalent to an optimal stopping problem for the continuous time process Y^* . Since for $a > 0$ the transformation $Y = aY^*$, $s = a^2 s^*$ replaces $Y^*(s^*)$ by a Wiener process $Y(s)$ in the $-s$ scale, we may select a so that $a^2 s_* = 1$, that is, $a = s_*^{-\frac{1}{2}} = (\sigma_0^{-2} + \frac{1}{2} N \sigma^{-2})^{\frac{1}{2}}$. Then the initial point $(Y_0^*, s_0^*) = (\mu_0, \sigma_0^2)$ is transformed to (y_0, s_0) , where

$$y_0 = \mu_0(\sigma_0^{-2} + \frac{1}{2} N \sigma^{-2})^{\frac{1}{2}}, \quad s_0 = \sigma_0^2(\sigma_0^{-2} + \frac{1}{2} N \sigma^{-2})$$

Setting $y = ay^*$, $s = a^2 s^*$, from (6.2) we have, for $s_0 \geq s \geq 1$,

$$\begin{aligned} d_1(y^*, s^*) &= Na^{-1} s^{\frac{1}{2}} \tilde{\psi}(ys^{-\frac{1}{2}}) - a\sigma^2(1 - s^{-1})|y| \\ (6.4) \quad &= \sigma^2 \sigma_0^{-1} s_0^{\frac{1}{2}} \{2(1 - s_0^{-1}) s^{\frac{1}{2}} \tilde{\psi}(ys^{-\frac{1}{2}}) - (1 - s^{-1})|y|\} \\ &\equiv d_2(y, s). \end{aligned}$$

Since $s^{\frac{1}{2}} \tilde{\psi}(Y(s)s^{-\frac{1}{2}})$ is a martingale, the term involving $s^{\frac{1}{2}} \tilde{\psi}$ satisfies the heat equation and does not affect the solution. Hence the continuous time version of the Anscombe problem is equivalent to the parameter-free problem where the stopping cost is

$$d_3(y, s) = -(1 - s^{-1})|y|.$$

The parameters enter only in the determination of the starting point (y_0, s_0) and the transformation back to the original (X, t^*) scale.

From (6.4) the Bayes risk corresponding to starting at (y_0, s_0) is given by

$$(6.5) \quad E(d_2(Y(S), S)) = \sigma^2 \sigma_0^{-1} s_0^{\frac{1}{2}} [E(d_3(Y(S), S)) + 2(1 - s_0^{-1}) s_0^{\frac{1}{2}} \tilde{\psi}(y_0 s_0^{-\frac{1}{2}})]$$

where S is an optimal stopping rule; the quantity $E(d_3(Y(S), S))$ can be approximated by the techniques described earlier.

After observing the differences in response for the first n pairs of patients, the current estimate of μ is Y_n^* . To incorporate the ethical cost,

the additional cost of $\gamma|Y_n^*|$, where γ is a constant of proportionality, is incurred if the decision is to observe another pair of patients; this additional cost corresponds to the application of an apparently inferior treatment to one of the patients in the pair. During the experimental phase, this ethical cost is accumulated at the rate

$$\gamma|Y_n^*|dn = -\sigma^2\gamma|Y_n^*|s_n^{*-2}ds_n^*,$$

which corresponds in the continuous time problem to

$$\begin{aligned} -\sigma^2\gamma|Y^*|s^{*-2}ds^* &= -\sigma^2\gamma|Y|s^{-2}ds \\ &= -\sigma_0^{-2}s_0^{\frac{1}{2}}\gamma|Y|s^{-2}ds \end{aligned}$$

in the transformed scale.

The introduction of the ethical cost has changed the nature of our problem. Basically we must consider not only the cost of stopping but also the charge for continuing each short period of time. There are two equivalent ways of regarding this problem. One is in terms of the optimizing backward induction and the other is in terms of the diffusion or modified heat equation satisfied by the risk function. The former is, for $s > 1$ and Z a standard normal deviate,

$$\hat{d}_3(y, s) = \min[d_3(y, s), \gamma|y|s^{-2}ds + E\{\hat{d}_3(y+Z(ds)^{\frac{1}{2}}, s - ds)\}].$$

with natural discrete time normal and Bernoulli analogues (compare with (3.1)).

The second term on the right incorporates the novel cost term. It leads to the following free boundary problem in terms of a nonhomogeneous diffusion equation (compare with (2.1)):

$$\begin{aligned} \hat{d}_{3s}(y, s) &= \frac{1}{2}\hat{d}_{3yy}(y, s) + \gamma|y|s^{-2} & \text{for } (y, s) \in C, \\ \hat{d}_3(y, s) &= d_3(y, s) & \text{for } (y, s) \in S, \\ \hat{d}_{3y}(y, s) &= d_{3y}(y, s) & \text{for } (y, s) \in \partial C. \end{aligned}$$

For the discrete time Bernoulli analogue of the backward induction we have

$$\begin{aligned} (6.6) \quad \hat{d}_3(y, s) &= d_3(y, s) & \text{for } s = 1, \\ &= \min[d_3(y, s), \gamma|y|s^{-2}\Delta + \{\hat{d}_3(y+\Delta^{\frac{1}{2}}, s-\Delta) + \hat{d}_3(y-\Delta^{\frac{1}{2}}, s-\Delta)\}/2] & \text{for } s > 1, \end{aligned}$$

and the CA, LA, QA and EX adjustments can be calculated in the same way as before

Using the discrete time versions the ethical cost accumulated over the interval s to $s-\Delta$ is zero when $Y(s) = 0$. But in the continuous time version the expectation of the ethical cost accumulated over the same interval would be (ignoring the constant multiplier $\sigma_0^{-2}s_0^{\frac{1}{2}}$) the positive quantity

$$-\gamma E\left(\int_s^{s-\Delta} |Y(u)|u^{-2}du \mid Y(s)=0\right).$$

In general the difference between $|y|s^{-2}\Delta$ and

$$\begin{aligned} -E\left(\int_s^{s-\Delta} |Y(u)|u^{-2}du \mid Y(s)=y\right) &= -|y|s^{-1} + 2(s-\Delta)^{-1}\Delta^{\frac{1}{2}}\tilde{\psi}(y\Delta^{-\frac{1}{2}}) \\ &\quad + \int_s^{s-\Delta} u^{-1}(s-u)^{-\frac{1}{2}}\phi(y(s-u)^{-\frac{1}{2}})du \end{aligned}$$

represents one source of error in our approximations. This difference can be estimated. Ignoring the constant multiplier $\sigma_0^{-2}s_0^{\frac{1}{2}}$, an asymptotic expansion shows that with $\epsilon = \sqrt{\Delta/s}$, C , the expectation of the ethical cost accumulated over this interval in the continuous time version, is approximated by

$$\begin{aligned} C &\approx 4\gamma\phi(0)s^{-\frac{1}{2}}\left[\frac{1}{3}\epsilon^3 + \frac{2}{5}\epsilon^5 + \frac{3}{7}\epsilon^7 + \dots\right] & \text{for } y = 0, \\ &\approx \gamma|y|s^{-1}[\epsilon^2 + \epsilon^4 + \dots] & \text{for } y \neq 0, \end{aligned}$$

whereas the analogue in the discrete time version is $\gamma|y|s^{-2}\Delta = \gamma|y|s^{-1}\epsilon^2$;

thus our approximation introduces errors of order $O(\epsilon^3) = O(\Delta^{3/2})$.

An alternative approach to the ethical cost problem consists of transforming it to an equivalent stopping problem without the nonhomogeneous cost term. In this particular application that approach is not practical because the transformed problem involves a stopping cost containing an integral, the evaluation of which throughout the course of the backward induction is too expensive to be worthwhile. The general principle may be of some interest and is presented here.

Given $Y(s_0) = y_0$, $s_0 \geq s \geq s_1$, the stopping cost corresponding to stopping at $Y(s) = y$ is

$$d(y, s) + I(s_0, s)$$

where

$$I(s_0, s) = - \int_{s_0}^s c(Y(u), u) du$$

and $c(y, s)$ is the rate of accumulation of the ethical cost when $Y(s) = y$.

This stopping cost depends not only on $Y(s)$ and s but also on the path

$Y(s')$, $s_0 \geq s' \geq s$.

Let F_s denote the sigma algebra containing the history of the process from s_0 to $s \geq s_1$. Then

$$M(s) = E\{I(s_0, s_1) | F_s\} = E\{I(s_0, s) | F_s\} + E\{I(s, s_1) | F_s\}$$

is a martingale. Moreover

$$E\{I(s_0, s) | F_s\} = I(s_0, s)$$

and

$$E\{I(s, s_1) | F_s\} = - \int_s^{s_1} E\{c(Y(u), u) | F_s\} du = h(Y(s), s) \text{ say.}$$

Thus, the stopping cost

$$d(Y(s), s) + I(s_0, s) = d(Y(s), s) - h(Y(s), s) + M(s)$$

may be expressed as a function of $Y(s)$ and s plus a martingale. But the expectation of the martingale is independent of the stopping rule and the optimal stopping rule is the same as for the problem with stopping cost

$$d'(y, s) = d(y, s) - h(y, s),$$

for which our methods apply.

In our special problem, the function h involves an integral of the form

$$\int_s^{s_1} u^{-1} (s-u)^{-\frac{1}{2}} \phi(y(s-u)^{-\frac{1}{2}}) du$$

which would have to be evaluated numerically throughout the course of the backward induction. Since this was judged to be impractical, the first approach was employed here; from (6.5) the approximation to the Bayes risk in the continuous time problem corresponding to the starting point (y_0, s_0) is given by

$$\sigma_{\sigma_0}^2 s_0^{-1} s_0^{\frac{1}{2}} \{ \hat{d}_3(y_0, s_0) + 2(1-s_0^{-1}) s_0^{\frac{1}{2}} \tilde{v}(y_0 s_0^{-\frac{1}{2}}) \}$$

where \hat{d}_3 is evaluated by the backward induction algorithm (6.6).

Properties of the optimal stopping rule in addition to the Bayes risk can also be approximated. For the continuous time problem, direct calculation shows that the contribution to the Bayes risk of the post-experimental phase (where all the remaining patients are assigned to the treatment which is inferred to be superior) is given by

$$(6.7) \quad \sigma_{\sigma_0}^2 s_0^{-1} s_0^{\frac{1}{2}} [2E\{(1-s^{-1}) s^{\frac{1}{2}} \psi(Y(s) s^{-\frac{1}{2}})\}]$$

while the Bayes expected sample size (number of pairs of patients treated during the experimental phase) is given by

$$(6.8) \quad \sigma_{\sigma_0}^2 s_0^{-2} s_0 \{E(s^{-1}) - s_0^{-1}\}.$$

The two expectations appearing in these expressions can be approximated in a straightforward manner during the execution of the backward induction which leads

to the approximation of the Bayes risk.

Some of the results of such computations for the Anscombe problem without ethical cost were reported in Chernoff and Petkau (1981); there

$$\begin{aligned} d_4(y, s) &= d_3(y, s) + 2s^{1/2} \tilde{\psi}(ys^{-1/2}), \\ &= -(1-s^{-1})|y| + 2s^{1/2} \tilde{\psi}(ys^{-1/2}), \\ &= |y|s^{-1} + 2s^{1/2} \tilde{\psi}(ys^{-1/2}) \end{aligned}$$

was employed. The term added to d_3 is a solution of the heat equation and therefore does not affect the optimal policy. It was expected to contribute to numerical stability since for large $|y|s^{-1/2}$ it is approximately $|y|$ and cancels the major part of d_3 and is important when s is large. In this case, d_3 is replaced by d_4 in the algorithm (6.6) and the approximation to the Bayes risk in the continuous time problem corresponding to the starting point (y_0, s_0) is given by

$$(6.9) \quad \sigma_0^2 s_0^{-1} s_0^{1/2} \{d_4(y_0, s_0) - 2s_0^{-1/2} \tilde{\psi}(y_0 s_0^{-1/2})\}.$$

These computations have been carried out for the cases $\gamma = 0$ (the Anscombe problem without ethical cost), 0.1, 1.0 and 10.0. The computations were implemented in exactly the fashion described for Examples 2.5, 2.1 and 2.2 in Section 5; 2080 stages were carried out in each phase, the grid spacing Δ was increased by a factor of 4 from each phase to the next, and the extent of overlapping corresponded to one-half of the interval of s values covered by the previous phase. For each case the grid spacing for the initial phase was taken to be $\Delta = 25 \times 10^{-6}$ and estimates were obtained out to $s = 10^{-6}$. The entire computation for the individual cases, which included evaluation of the Bayes expected sample size and the proportion of the Bayes risk due to the experimental phase as well as the Bayes risk, required between 28 and 35 seconds of CPU time at a cost of between \$2.00 and \$2.50 on the 12-megabyte Amdahl 470 V/8 at the University of British Columbia.

The optimal procedure for the continuous time problem may be described by the stopping boundary $\tilde{\beta}_Y(t) = 1 - \Phi(\tilde{z}_Y(t))$, presented for the cases under consideration in Table 24 and interpreted as follows. Define

$$z = Y(s)/s^{1/2} = Y^*(s^*)/s^{*1/2},$$

the number of standard deviations that the current Bayes estimate of μ is away from zero, and

$$t = 1/s = s^{*-1}/s_*^{-1} = (\sigma_0^{-2} + t^* \sigma^{*-2})/(\sigma_0^{-2} + \frac{1}{2} N \sigma^{*-2}),$$

the currently available proportion of the total potential information. If at any time $\beta = 1 - \Phi(|z|) < \tilde{\beta}_Y(t)$, stop taking observations and for the remaining $N - 2t^*$ units of time use the treatment in accord with the sign of Y^* . Note that β is the observed P value for a one-sided test of $\mu = 0$ based on the data and the prior. At time t , the curve $\tilde{z}_Y(t)$ specifies the number of standard deviations required for stopping and $\tilde{\beta}_Y(t)$ is the corresponding nominal significance level. Thus the optimal procedure may be described as a sequence of repeated significance tests with the nominal significance level varying with the amount of information available; as the proportion of information available increases from 0 to 1, the nominal significance level becomes less stringent, increasing from 0 to 1/2. The optimal boundaries are plotted in the (β, t) scale in Figure 17. Note that for a given value of t , Bayes estimates of μ further from zero are required for stopping for larger values of γ , the ethical cost parameter; that is, larger values of γ imply earlier stopping.

Although Figure 17 provides a clear overall comparison, the exact form of the stopping boundaries near the distinguished points $t = 0$, where few patients have been treated, and $t = 1$, where nearly all the patients

have been treated, is of particular interest. An asymptotic expansion for values of t close to 1 yields

$$\tilde{z}_Y \sim c_Y(1-t)^{1/2}, \quad \tilde{\beta}_Y \sim 1/2 - c_Y(1-t)^{1/2}/\sqrt{2\pi},$$

where c_Y is the unique positive solution of

$$\phi(c) = (1+\gamma)c^2\tilde{\phi}(c),$$

for the values $\gamma = 0.0, 0.1, 1.0$ and 10.0 , $c_Y \approx 0.7642, 0.7401, 0.5972$ and 0.2893 respectively. An asymptotic expansion for small values of t yields

$$-2 \log t \sim \tilde{z}_Y^2 + \log \tilde{z}_Y^2 + \log[2\pi(1+\gamma)^2] + 2\tilde{z}_Y^{-2} + \tilde{z}_Y^{-4},$$

$$\tilde{\beta}_Y \sim (1+\gamma)t(1 + 3(\log t)^{-2}/4).$$

Since small values of t are particularly relevant for problems involving large values of the horizon size N , it is important to note the accuracy of the approximation $\tilde{\beta}_Y(t) \sim (1+\gamma)t$ for small values of t in Table 24.

While comparison of the stopping boundaries indicates the effect of the ethical cost on the optimal stopping rules, of possibly greater interest are the risks incurred when these optimal procedures are employed. These risks depend upon the five parameters $\mu_0, \sigma_0, \sigma, N$ and γ . For simplicity in tabular presentation we may use the normalization

$$BR = \text{Bayes risk}/\sigma^2\sigma_0^{-1},$$

where, as is clear from (6.9), BR depends only upon γ in addition to the

initial values of t and z , namely

$$t_0 = \sigma_0^{-2}/(\sigma_0^{-2} + \frac{1}{2}N\sigma^{-2}), \quad z_0 = \mu_0/\sigma_0.$$

A small representative subset of the normalized risks BR which have been evaluated are presented in Table 25. In each case the normalized Bayes expected sample size

$$EN = \text{Bayes expected sample size}/\sigma^2\sigma_0^{-2}$$

and the proportion, PR , of the Bayes risk resulting from the experimental phase where one-half of the patients are assigned to the inferior treatment are also tabulated; these quantities are computed according to (6.8) and (6.7) respectively.

For fixed values of t_0 and z_0 , the Bayes risk increases monotonically with γ ; the tabulated values provide an indication of the magnitude of the effect of the ethical cost. The tabulated values of EN reflect the differences in the stopping rules which are evident in Table 24 as well as Figure 17, and translate these differences into more meaningful quantities. Note that for small values of t_0 , the ethical cost has little effect on PR , the proportion of the Bayes risk due to the experimental phase. The leading term of an asymptotic expansion for t_0 small and z_0 not large indicates that

$$\text{Bayes risk} \sim \sigma^2\sigma_0^{-1}(1+\gamma)\phi(z_0)(\log t_0)^2.$$

This result explicitly indicates the effect of the ethical cost, and means that the order of magnitude of the optimal Bayes risk is $(\log N)^2$ which may seem surprisingly small. These asymptotic expansions for the Bayes risk and the optimal stopping boundaries can be obtained by the techniques described in Chernoff and Petkau (1981).

The behavior of these Bayes properties of the optimal procedure is illustrated in Figures 18-23. The Bayes risk, Bayes expected sample size, and proportion of the Bayes risk due to the experimental phase at $z_0 = 0$ are plotted against $\log(t_0) = -\log(s_0)$ in Figures 18, 19 and 20 respectively. While the quadratic nature of the dependence of the Bayes risk on $\log s_0$ for large values of s_0 is clearly indicated in Figure 18, Figure 19 indicates that the Bayes expected sample size grows at a considerably faster rate. These trends are even more apparent when the same quantities are plotted against $\{\log(t_0)\}^2$, although such plots are not included here. These same plots for other values of z_0 yielded similar patterns. In Figures 21, 22 and 23 these same quantities at $t_0 = 10^{-6}$ ($s_0 = 10^6$) are plotted against z_0 ; the same plots for other values of t_0 yielded similar patterns.

The results presented were all obtained using the backward induction (6.6) with d_3 replaced by d_4 . This algorithm approximates the expectation of the ethical cost accumulated over the interval s to $s-\Delta$ in the continuous time version by the ethical cost accumulated over the same interval in the discrete time version, thereby introducing errors of order $O(\epsilon^3)$, where $\epsilon^2 = \Delta/s$. Since $\epsilon^2 < 0.003$ in our implementation of this algorithm, these errors should have negligible effect.

The investigation of the convergence properties of two different versions of the backward induction algorithm provides detailed information on the magnitude of this effect. Version 1 is that described above while version 2 is the modification obtained by replacing the term $\gamma|y|s^{-2}\Delta = \gamma|y|s^{-1}\epsilon^2$ in (6.6) with

$$c = \gamma\phi(0)s^{-1/2} [2\epsilon/(1-\epsilon^2) + \log\{(1-\epsilon)/(1+\epsilon)\}] \quad \text{for } y = 0,$$

$$= \gamma|y|s^{-1} [1/(1-\epsilon^2) - 1] \quad \text{for } y \neq 0,$$

except for terms of order $O\{\epsilon^7\phi(y s^{-1/2}\epsilon^{-1})\}$ in the case of $y \neq 0$, c is equal to C , the expectation of the ethical cost accumulated over the interval s to $s-\Delta$ in the continuous time version.

The computations carried out are similar to those described for Examples 2.1 and 2.2 in Section 4; the algorithm is executed over the interval $1 < s < 100$ for each grid in the sequence specified by $\Delta = 4^{-k}$ for $k = 0, 1, 2, 3, 4$. For each version, the approximation to the continuous time solution provided by the results for the most refined grid spacing is taken as baseline and the deviation from this baseline of the approximation obtained with a less refined grid spacing is examined. The results for both versions in the case $\gamma = 1$ are summarized in Table 26 and are qualitatively similar to the results obtained in Examples 2.1 and 2.2. Note that the correction required to approximate the continuous time boundary is underestimated in both versions. On the other hand, while version 1 results in underestimates of the continuous time risk, version 2 results in overestimates.

Of greater interest in the present case is the examination of the behaviour of the differences between the results produced by the two versions as the grid spacing Δ decreases. The differences in the estimates of both the boundary and the risk for the computations described above are summarized in Table 27. The results clearly indicate that the differences in the estimates of the Bayes risk produced by the two versions, as measured by either the maximum or average difference, are directly proportional to Δ , the grid spacing in s . Either version will produce excellent approximations to both the boundary and the risk of the continuous time problem when reasonable grid spacings are employed.

7. SUMMARY AND COMMENTS

We have presented a method for obtaining numerical solutions for optimal stopping problems involving a stopping cost $d(y,s)$ when the Wiener process $Y(s)$ in the $-s$ ($s \geq s_1$) scale stops at $(Y(s), s) = (y, s)$. The main idea of this method is to approximate the Wiener process by a discrete time process with independent Bernoulli increments $Z_i \Delta^{1/2}$, i.e. $Z_i = \pm 1$ with probability $1/2$ and

$$(7.1) \quad \begin{aligned} Y_n &= y_0 + \sum_{i=1}^n Z_i \Delta^{1/2} \\ s_n &= s_0 - n\Delta \end{aligned}$$

The optimal stopping procedure for a stopping cost $d(y,s)$ associated with the above discrete time process may be derived by the backward induction scheme with the following simple recursion equation for the optimal risk $\hat{d}(y,s)$

$$(7.2) \quad \hat{d}(y,s) = \min(d(y,s), [\hat{d}(y+\Delta^{1/2}, s-\Delta) + \hat{d}(y-\Delta^{1/2}, s-\Delta)]/2);$$

the optimal stopping procedure calls for continuation when $\hat{d}(y,s) < d(y,s)$ and stopping otherwise.

If the boundaries for the optimal stopping problems for continuous and discrete time are denoted by \bar{y} and \bar{y}_Δ , then the approximation

$$(7.3) \quad \bar{y} = \bar{y}_\Delta \pm 0.5\Delta^{1/2}$$

furnishes the basis for a considerable improvement in accuracy. Unfortunately a single backward induction calculation provides $\hat{d}(y,s)$ only on a rectangular grid of points and \bar{y}_Δ is not calculated directly and may be in error by as much as $\Delta^{1/2}$. Several alternate continuity correction methods were described to compensate for this difficulty. Three simple methods of estimating \bar{y}_Δ are the crude adjusted (CA), linear adjusted (LA), and quadratic adjusted (QA). A fourth, more refined, method called the extrapolation method (EX) is based on the solution to a simple discrete time stopping problem which also provides the theoretical basis for (7.3). It involves the calculation of $D(y,s) = \hat{d}(y,s) - d(y,s)$ at the two continuation points closest to $\bar{y}_\Delta(s)$ for each s on the grid.

This approach is fundamentally unsound as a numerical method to derive refined approximations with accuracy to many significant digits. To increase accuracy by cutting $\Delta^{1/2}$ in half involves increasing the numerical work by a factor of 8. Without continuity corrections this would increase the accuracy by a factor of 2. As determined by numerous calculations on several examples, one obtains surprisingly good results for crude intervals Δ . Moreover as $\Delta \rightarrow 0$, the use of EX seems to divide the error by 3 to 4 when $\Delta^{1/2}$ is cut in half indicating that doubling the accuracy requires only about three times as much numerical computation.

Several variations of the basic approach are occasionally useful in reducing the computing effort. (1) If results are desired over a very large range of s values, then it was suggested that a small value of Δ be used for a range of s values, followed by a larger value of Δ over an overlapping range of s values, etc. (2) When the optimal continuation region is unbounded in y , a truncation procedure was described where $\hat{d}(y,s)$ need not be calculated for $y > c$ if $y \geq c$ is in the continuation region for all $s > s_1$. This method

depends on the probability that the (Y_n, s_n) process originating at $(c + \Delta^k, s)$ will reach (c, s_n) for some $s_n < s$. It is particularly useful when $d(y, s) = 0$ for $y > c$ and $s = s_1$. Moreover, a transformation of d (to be discussed shortly) which reduces the computational effort and does not affect the optimal stopping boundary can be applied to make $d = 0$ for $y > c$ and $s = s_1$. (3) When $d(y, s)$ is symmetric in y , it is possible to restrict calculations to values of $y \geq 0$ thereby reducing the numerical work by half.

The original continuous time stopping problem has a solution which can be described in terms of a free boundary problem (FBP) involving the heat equation. Related to that is the fact that if one adds a solution of the heat equation to the stopping cost $d(y, s)$, the optimal stopping region is not affected and the risk is increased by this solution of the heat equation. This fact is a special case of the more general fact that if $d(Y(s), s)$ is increased by a martingale $M(s)$, the optimal stopping procedure is not affected. These properties were used in the truncation variation of the proceeding paragraph. They were used in one of the examples where d and \hat{d} became large to reduce d and thereby attain numerical stability. Finally, they were used in the Anscombe problem with ethical cost to transform that problem to a stopping problem with stopping cost $d(Y(s), s)$.

Eight applications were considered. Several consisted of problems with known solutions so that the numerical accuracy of the methods could be evaluated. Several consisted of old problems of importance in the statistical literature so that refined calculations of the solutions could be presented. These include the sequential analysis and one-armed bandit problems. Finally the Anscombe problem with ethical cost represents a new problem whose solution may be regarded as having potential value in applications.

One method of describing the optimal stopping procedure for some of these problems, which derive from observations on a Wiener process with unknown mean with a normal prior distribution, is in terms of a nominal significance level $\tilde{\beta} = 1 - \Phi(|\tilde{y}|s^{-1/2})$. This description can be used to interpret the optimal procedure as a sequence of repeated significance tests where the significance level is not held constant, but depends on the amount of information collected to date.

The general approach is easily adaptable to decomposing the optimal risk into parts representing terms such as the cost of sampling, the cost of error, etc. It may also be applied to evaluate alternative, non-optimal procedures although that was not done in this paper and the refinement due to the correction (7.3) and to the use of EX is not meaningful then.

Many problems originate as discrete time or discrete time and discrete process problems. For example, the rectified sampling inspection problem is such a problem where the fraction defective in a lot is compared to a fixed number p_0 . The continuous time version of that problem is the one-armed bandit problem which is approximated by our approach. But the solution of the latter problem is only an approximation to the solution of the sampling inspection problem which involves a Bernoulli process with increments which have probability p_0 and $1 - p_0$ respectively. The theorem which provides the approximation (7.3) also provides a similar approximation relating y and the optimal boundary for the original discrete time sampling inspection problem. This approximation is discussed in Chernoff and Petkau (1976). The idea of using a discrete approximation to a Wiener process problem which itself approximates a discrete time problem is not as circular as it seems. Our numerical calculation is particularly simple partly because we can choose the intervals in s to suit our taste. Moreover the Wiener process version

of the problem often allows normalizations which permit us to solve many problems at once.

A general purpose program designed to handle a large variety of these stopping problems should be capable of taking advantage of the special features of particular problems which might allow the necessary computational effort to be substantially reduced. Although it is possible to write such a general purpose program, one should anticipate that special versions may occasionally require intelligent intervention to avoid numerical difficulties such as underflows, overflows and round off errors. For example, using these techniques rather careful programming was required to obtain a good approximation to the optimal stopping boundary in the problem with stopping cost

$$d(y,s) = \min(y,0) \exp(-1/s) \quad \text{for } s \geq 0,$$

discussed by Bather (1983).

REFERENCES

- Anscombe, F.J. (1963). Sequential medical trials. J. Am. Statist. Assoc. 58, 365-83.
- Armitage, P. (1963). Sequential medical trials. Some comments on F.J. Anscombe's paper. J. Am. Statist. Assoc. 58, 384-7.
- Bather, J. (1962). Bayes procedures for deciding the sign of a normal mean. Proc. Cambridge Philos. Soc. 58, 599-620.
- Bather, J. (1983). Optimal stopping of Brownian motion: a comparison technique. In Recent Advances in Statistics. Papers in Honor of Herman Chernoff on his Sixtieth Birthday. M.H. Rizvi, J. Rustagi, D. Siegmund, eds., Academic Press, New York, 19-50.
- Bather, J. & Chernoff, H. (1967a). Sequential decisions in the control of a spaceship. Proc. 5th Berkeley Symp. 3, 181-207.
- Bather, J. & Chernoff, H. (1967b). Sequential decisions in the control of a spaceship (finite fuel). J. Appl. Prob. 4, 584-604.
- Begg, C.B. & Mehta, C.R. (1979). Sequential analysis of comparative clinical trials. Biometrika 66, 97-105.
- Breakwell, J. & Chernoff, H. (1964). Sequential tests for the mean of a normal distribution II. Ann. Math. Statist. 35, 162-73.
- Chernoff, H. (1961). Sequential tests for the mean of a normal distribution. Proc. 4th Berkeley Symp. 1, 79-91.
- Chernoff, H. (1965a). Sequential tests for the mean of a normal distribution III. Ann. Math. Statist. 36, 29-54.
- Chernoff, H. (1965b). Sequential tests for the mean of a normal distribution IV. Ann. Math. Statist. 36, 55-68.
- Chernoff, H. (1967). Sequential models for clinical trials. Proc. 5th Berkeley Symp. 4, 805-12.
- Chernoff, H. (1968). Optimal stochastic control. Sankhya A 30, 221-52.
- Chernoff, H. (1972). Sequential Analysis and Optimal Design. SIAM. J.W. Arrowsmith, Bristol.
- Chernoff, H. & Petkau, A.J. (1976). An optimal stopping problem for sums of dichotomous random variables. Ann. Prob. 4, 875-89.
- Chernoff, H. & Petkau, A.J. (1981). Sequential medical trials involving paired data. Biometrika 68, 119-32.
- Chernoff, H. & Ray, S.N. (1965). A Bayes sequential sampling inspection plan. Ann. Math. Statist. 36, 1387-407.

- Chow, Y.S. & Robbins, H. (1965). On optimal stopping rules for S_n/n . Illinois J. Math. 9, 444-54.
- Day, N.E. (1969). A comparison of some sequential designs. Biometrika 56, 301-11.
- Dvoretzky, A. (1965). Existence and properties of certain optimal stopping rules. Proc. 5th Berkeley Symp. 1, 441-52.
- Feder, P.I. & Stroud, T. (1971). Sequential decisions in the control of a spaceship (terminal cost proportional to magnitude of miss distance). J. Appl. Prob. 8, 285-97.
- Feller, W. (1968). An Introduction to Probability Theory and Its Applications. Vol. 1, 3rd edition, John Wiley and Sons, Inc., New York.
- Lai, T.L., Levin, B., Robbins, H. & Siegmund, D. (1980). Sequential medical trials. Proc. Natl. Acad. Sci. U.S.A. 77, 3135-8.
- Lai, T.L., Robbins, H. & Siegmund, D. (1983). Sequential design of comparative clinical trials. In Recent Advances in Statistics, Papers in Honor of Herman Chernoff on his Sixtieth Birthday, M.H. Rizvi, J. Rustagi, D. Siegmund, eds., Academic Press, New York, 51-68.
- Lindley, D.V. (1961). Dynamic programming and decision theory. Applied Statistics 10, 39-51.
- Lindley, D.V. & Barnett, B.N. (1965). Sequential sampling: two decision problems with linear losses for binomial and normal random variables. Biometrika 52, 507-32.
- Moriguti, S. & Robbins, H. (1962). A Bayes test of $p \leq \frac{1}{2}$ versus $p > \frac{1}{2}$. Rep. Statist. Appl. Res., Un. Japan Sci. Engrs. 9, 39-60.
- Petkau, A.J. (1978). Sequential medical trials for comparing an experimental with a standard treatment. J. Am. Statist. Assoc. 73, 328-38.
- Petkau, A.J. (1980). Frequentist properties of three stopping rules for comparative clinical trials. Biometrika 67, 690-2.
- Shepp, L.A. (1969). Explicit solutions to some problems of optimal stopping. Ann. Math. Statist. 40, 993-1010.
- Shiryaev, A.N. (1978). Optimal Stopping Rules. Springer-Verlag, New York.
- Snell, J.L. & Tisdale, H. (1978). private communication.
- Taylor, H.M. (1968). Optimal stopping in a Markov process. Ann. Math. Statist. 39, 1333-44.
- Teicher, H. & Wolfowitz, J. (1966). Existence of optimal stopping rules for linear and quadratic rewards. Z. Wahrscheinlichkeitstheorie und Verw. Gebiete 5, 361-8.
- Van Moerbeke, P.L.J. (1974a). Optimal stopping and free boundary problems. Rocky Mountain J. Math. 4, 539-78.
- Van Moerbeke, P.L.J. (1974b). An optimal stopping problem with linear reward. Acta Mathematica 132, 111-51.
- Van Moerbeke, P.L.J. (1975). On optimal stopping and free boundary problems. Archive for Rat. Mech. Math. 60, 101-48.
- Walker, L.H. (1969). Regarding stopping rules for Brownian motion and random walks. Bull. Amer. Math. Soc. 75, 46-50.

TABLE 1. ERRORS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.5.*

Grid Spacing			Method			
s	y		CA	LA	QA	EX
1	1	Ave(e)	-1690	8	-818	-1676
		Ave(e)	1690	209	818	1676
		Max(e)	3358	813	1276	2857
4 ⁻¹	2 ⁻¹	Ave(e)	-266	283	-91	-419
		Ave(e)	266	283	93	419
		Max(e)	649	466	197	666
4 ⁻²	2 ⁻²	Ave(e)	-14	209	29	-101
		Ave(e)	29	209	29	104
		Max(e)	113	315	55	175
4 ⁻³	2 ⁻³	Ave(e)	19	116	34	-22
		Ave(e)	19	116	34	26
		Max(e)	36	180	48	52
4 ⁻⁴	2 ⁻⁴	Ave(e)	11	70	23	7
		Ave(e)	11	70	23	9
		Max(e)	16	96	30	17

*In each case, the errors summarized are $e = (\hat{y} - \bar{y}) \times 10^4$ at $s = 25(1)100$; for this range, \bar{y} varies between 10 and 26.

TABLE 2. ERRORS IN ESTIMATION OF RISK FOR EXAMPLE 2.5.*

Grid Spacing			Ave(e)	Ave(e)	Max(e)
s	y				
1	1		-587	2391	9162
4 ⁻¹	2 ⁻¹		-243	576	1515
4 ⁻²	2 ⁻²		-118	178	531
4 ⁻³	2 ⁻³		-34	48	148
4 ⁻⁴	2 ⁻⁴		-9	13	39

*In each case, the errors summarized are $10^{-6}xe = \text{optimal risk in discrete time problem} - \text{optimal risk in continuous time problem}$ at all grid points on the intersections of the lines $s = 25(1)100$, $y = 0(1)\infty$ and within the continuation region for the discrete time problem. Note that $\text{Ave}(e) \neq \text{Ave}(|e|)$ demonstrates that it is not the case that the various discrete time random walk problems are uniformly less favourable than the continuous time problem in this example (in fact, these discrete versions are on the average more favourable here); nor is it the case that the convergence is monotone at fixed (y,s) grid points.

Note that the optimal risk for this problem is symmetric in y , always negative, and, for fixed s , becomes increasingly negative as $|y|$ increases; inside the continuation region the risk decreases from -4.0 to -10.3 at $s = 25$ and from -8.0 to -25.8 at $s = 100$.

TABLE 3. ERRORS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.6.*

Grid Spacing**			Method			
t	x		CA	LA	QA	EX
1(-2)	1(-1)	Ave(e)	-65	-1121	-357	194
		Ave(e)	106	1121	357	237
		Max(e)	307	1529	460	436
4 ⁻¹ (-2)	2 ⁻¹ (-1)	Ave(e)	-51	-560	-197	51
		Ave(e)	51	560	197	71
		Max(e)	129	788	257	140
4 ⁻² (-2)	2 ⁻² (-1)	Ave(e)	-37	-281	-102	13
		Ave(e)	37	281	102	24
		Max(e)	64	404	138	62
4 ⁻³ (-2)	2 ⁻³ (-1)	Ave(e)	-17	-140	-52	3
		Ave(e)	17	140	52	8
		Max(e)	23	204	73	20
4 ⁻⁴ (-2)	2 ⁻⁴ (-1)	Ave(e)	-7	-71	-26	1
		Ave(e)	7	71	26	3
		Max(e)	10	104	38	10

*In each case, the errors summarized are $e = (\hat{x} - \bar{x}) \times 10^5$ at $t = 0(0.01)0.75$; for this range, \bar{x} varies between -0.25 and -0.51.

** In this and all following tables a(-n) represents $a \times 10^{-n}$.

TABLE 4. ERRORS IN ESTIMATION OF RISK FOR EXAMPLE 2.6.*

Grid Spacing		Ave(e)	Ave(e)	Max(e)
t	x			
1(-2)	1(-1)	658	658	2126
4 ⁻¹ (-2)	2 ⁻¹ (-1)	148	148	537
4 ⁻² (-2)	2 ⁻² (-1)	36	36	140
4 ⁻³ (-2)	2 ⁻³ (-1)	9	9	37
4 ⁻⁴ (-2)	2 ⁻⁴ (-1)	2	2	7

*In each case, the errors summarized are $10^{-6}xe$ = optimal risk in discrete time problem - optimal risk in continuous time problem at all grid points on the intersections of the lines $t = 0(0.01)0.75$, $x = 0(-0.1)-\infty$ and within the continuation region for the discrete time problem. Ave(e) = Ave(|e|) indicates the various discrete time random walk problems are uniformly less favourable than the continuous time problem.

Note that for the version of the problem being considered the optimal risk is always negative in this portion of the continuation region and, for fixed t , becomes increasingly negative as x decreases from 0; in this portion of the continuation region the risk decreases from -0.18 to -0.83 at $t = 0.75$ and from -0.72 to -3.32 at $t = 0$.

TABLE 5. ACCURACY OF ADJUSTMENT (3.4) FOR EXAMPLE 2.6.*

Grid Spacing			Ave(e)	Ave(e)	Max(e)
t	x				
1(-2)	1(-1)	Raw	4607	4607	4637
		Adjusted	-393	393	436
4 ⁻¹ (-2)	2 ⁻¹ (-1)	Raw	2344	2344	2362
		Adjusted	-156	156	178
4 ⁻² (-2)	2 ⁻² (-1)	Raw	1192	1192	1203
		Adjusted	-58	58	69

*In each case, the errors summarized are $e = (\hat{x} - \bar{x}) \times 10^5$ at $t = 0(0.01)0.75$; for this range, \bar{x} varies between 0.25 and 0.51.

TABLE 6. DEVIATIONS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.1.*

Grid Spacing			Method			
s	y		CA	LA	QA	EX
1	1	Ave(d)	-1454	239	-687	-1869
		Ave(d)	1454	381	687	1859
		Max(d)	3299	880	1474	3592
4 ⁻¹	2 ⁻¹	Ave(d)	-259	298	-76	-473
		Ave(d)	266	300	86	478
		Max(d)	677	540	246	919
4 ⁻²	2 ⁻²	Ave(d)	-14	149	23	-111
		Ave(d)	35	152	28	127
		Max(d)	140	302	58	228
4 ⁻³	2 ⁻³	Ave(d)	9	67	14	-26
		Ave(d)	12	81	15	34
		Max(d)	27	128	34	55

*The computation with grid spacing in $s = 4^{-4}$, in $y = 2^{-4}$ provides the baseline for each method. In each case, the deviations summarized are $d = (\hat{y}_1 - \hat{y}_2) \times 10^4$ at $s = 25(1)100$, where \hat{y}_2 is the approximation to the continuous time boundary at baseline and \hat{y}_1 is the approximation at the grid spacing listed (both computed by the same method). For this range of values of s , \bar{y} varies between 11 and 31.

TABLE 7. DEVIATIONS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.1.*

Grid Spacing			Method		
s	y		CA	LA	QA
1	1	Ave(d)	435	2182	1212
		Ave(d)	794	2182	1212
		Max(d)	2222	3571	2578
4 ⁻¹	2 ⁻¹	Ave(d)	234	845	428
		Ave(d)	309	845	428
		Max(d)	702	125	809
4 ⁻²	2 ⁻²	Ave(d)	117	334	164
		Ave(d)	124	334	164
		Max(d)	238	527	248
4 ⁻³	2 ⁻³	Ave(d)	55	166	70
		Ave(d)	55	166	70
		Max(d)	91	236	100
4 ⁻⁴	2 ⁻⁴	Ave(d)	20	74	30
		Ave(d)	20	74	30
		Max(d)	32	113	45

*In each case, the deviations summarized are $d = (\hat{y} - \hat{y}_{EX}) \times 10^4$ at $s = 25(1)100$, where \hat{y}_{EX} is the approximation obtained using method EX at the grid spacing under consideration. For this range of s values, \hat{y} varies between 11 and 31.

TABLE 8. DEVIATIONS IN ESTIMATION OF RISK FOR EXAMPLE 2.1.*

Grid Spacing		Ave(d)	Ave(d)	Max(d)
s	y			
1	1	210	288	1139
4 ⁻¹	2 ⁻¹	-5	73	457
4 ⁻²	2 ⁻²	-9	20	137
4 ⁻³	2 ⁻³	-3	5	33

*In each case, the deviations summarized are 10^6 times the differences between the optimal risk in the random walk problem with the grid spacing listed and that with grid spacing in $s = 4^{-4}$, in $y = 2^{-4}$ at all grid points on the intersections of the lines $s = 25(1)100$, $y = 0(1)\infty$ and within the continuation region for the discrete time problem with the less refined grid spacing.

Note that the optimal risk for this problem is symmetric in y, always positive, and, for fixed s, decreases as |y| increases; inside the continuation region the risk decreases from 0.30 to 0.06 at $s = 25$ and from 0.17 to 0.01 at $s = 100$.

TABLE 9. DEVIATIONS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.2.*

Grid Spacing			Method			
s	y		CA	LA	QA	EX
1	1	Ave(d)	1107	-308	481	1227
		Ave(d)	1107	330	481	1227
		Max(d)	2417	627	870	2187
4 ⁻¹	2 ⁻¹	Ave(d)	164	-310	19	314
		Ave(d)	168	310	43	317
		Max(d)	469	510	104	475
4 ⁻²	2 ⁻²	Ave(d)	-5	-167	-30	72
		Ave(d)	23	167	30	78
		Max(d)	52	309	64	135
4 ⁻³	2 ⁻³	Ave(d)	-11	-57	-16	13
		Ave(d)	13	74	19	23
		Max(d)	24	134	37	47

*The computation with grid spacing in $s = 4^{-4}$, in $y = 2^{-4}$ provides the baseline for each method. In each case, the deviations summarized are $d = (\hat{y}_1 - \hat{y}_2) \times 10^4$ at $s = 25(1)100$, where \hat{y}_2 is the approximation at baseline and \hat{y}_1 is the approximation at the grid spacing listed (both computed by the same method). For this range of s values, \bar{y} varies between -8 and -21.

TABLE 10. DEVIATIONS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.2.*

Grid Spacing			Method		
y	s		CA	LA	QA
1	1	Ave(d)	-196	-1594	-775
		Ave(d)	494	1594	775
		Max(d)	1441	2414	1716
4 ⁻¹	2 ⁻¹	Ave(d)	-167	-698	-323
		Ave(d)	189	698	323
		Max(d)	420	987	437
4 ⁻²	2 ⁻²	Ave(d)	-94	-313	-130
		Ave(d)	99	313	130
		Max(d)	183	458	186
4 ⁻³	2 ⁻³	Ave(d)	-41	-144	-57
		Ave(d)	41	144	57
		Max(d)	69	221	87
4 ⁻⁴	2 ⁻⁴	Ave(d)	-17	-75	-29
		Ave(d)	17	75	29
		Max(d)	30	109	42

*In each case, the deviations summarized are $d = (\hat{y} - \hat{y}_{EX}) \times 10^4$ at $s = 25(1)100$, where \hat{y}_{EX} is the approximation obtained using method EX at the grid spacing under consideration. For this range of s values, \bar{y} varies between -8 and -21.

TABLE 11. DEVIATIONS IN ESTIMATION OF RISK FOR EXAMPLE 2.2.*

Grid Spacing		Ave(d)	Ave(d)	Max(d)
s	y			
1	1	1205	1474	8487
4 ⁻¹	2 ⁻¹	-22	155	420
4 ⁻²	2 ⁻²	-10	39	131
4 ⁻³	2 ⁻³	-3	8	20

*In each case, the deviations summarized are 10^6 times the differences between the optimal risk in the random walk problem with the grid spacing listed and that with grid spacing in $s = 4^{-4}$, in $y = 2^{-4}$ at all grid points on the intersections of the lines $s = 25(1)100$, $y = 0(-1)-\infty$ and within the continuation region for the discrete time problem with the less refined grid spacing.

Note that for the version of the problem being considered the optimal risk is always negative in this portion of the continuation region and, for fixed s , becomes increasingly negative as y decreases from 0; in this portion of the continuation region the risk decreases from -1.7 to -6.7 at $s = 25$ and from -3.7 to -20.8 at $s = 100$.

TABLE 12. DEVIATIONS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.7.*

Grid Spacing			Method			
t	x		CA	LA	QA	EX
1(-2)	1(-1)	Ave(d)	-6501	523	188	1984
		Ave(d)	6501	523	188	1984
		Max(d)	9688	985	579	5910
4 ⁻¹ (-2)	2 ⁻¹ (-1)	Ave(d)	-830	444	160	216
		Ave(d)	843	444	160	222
		Max(d)	2188	685	245	1050
4 ⁻² (-2)	2 ⁻² (-1)	Ave(d)	-248	204	75	56
		Ave(d)	256	209	76	61
		Max(d)	938	375	138	297
4 ⁻³ (-2)	2 ⁻³ (-1)	Ave(d)	-55	74	28	11
		Ave(d)	64	91	33	15
		Max(d)	313	168	62	64

*The computation with grid spacing in $t = 4^{-4} \times 10^{-2}$, in $x = 2^{-4} \times 10^{-1}$ provides the baseline for each method. In each case, the deviations summarized are $d = (\hat{x}_1 - \hat{x}_2) \times 10^5$ at $t = 0.04(0.01)0.75$, where \hat{x}_2 is the approximation at baseline and \hat{x}_1 is the approximation at the grid spacing listed (both computed by the same method). For this region of t values, \bar{x} varies between .16 and .35.

TABLE 13. DEVIATIONS IN ESTIMATION OF BOUNDARY FOR EXAMPLE 2.7.*

Grid Spacing		Method			
t	x		CA	LA	QA
1(-2)	1(-1)	Ave(d)	-8504	-1400	-1777
		Ave(d)	8504	1575	1783
		Max(d)	13006	5410	5504
4 ⁻¹ (-2)	2 ⁻¹ (-1)	Ave(d)	-1064	289	-36
		Ave(d)	1064	426	196
		Max(d)	2402	866	946
4 ⁻² (-2)	2 ⁻² (-1)	Ave(d)	-322	209	38
		Ave(d)	323	229	87
		Max(d)	1042	359	254
4 ⁻³ (-2)	2 ⁻³ (-1)	Ave(d)	-84	124	36
		Ave(d)	85	128	44
		Max(d)	395	197	69
4 ⁻⁴ (-2)	2 ⁻⁴ (-1)	Ave(d)	-18	61	20
		Ave(d)	20	61	21
		Max(d)	80	103	38

*In each case, the deviations summarized are $d = (\hat{x} - \hat{x}_{EX}) \times 10^5$ at $t = 0.04(0.01)0.75$, where \hat{x}_{EX} is the approximation obtained using method EX at the grid spacing under consideration. For this range of t values, \hat{x} varies between .16 and .35.

TABLE 14. DEVIATIONS IN ESTIMATION OF RISK FOR EXAMPLE 2.7.*

Grid Spacing				
t	x	Ave(d)	Ave(d)	Max(d)
1(-2)	1(-1)	6133	6133	111700
4 ⁻¹ (-2)	2 ⁻¹ (-1)	1378	1378	22643
4 ⁻² (-2)	2 ⁻² (-1)	338	338	7560
4 ⁻³ (-2)	2 ⁻³ (-1)	65	65	1006

*In each case, the deviations summarized are 10^6 times the differences between the optimal risk in the random walk problem with the grid spacing listed and that with grid spacing in $t = 4^{-4} \times 10^{-2}$, in $x = 2^{-4} \times 10^{-1}$ at all grid points on the intersections of the lines $t = 0.04(0.01)0.75$, $x = 0.0(0.1)\infty$ and within the continuation region for the discrete time problem with the less refined grid spacing.

Note that for the version of the problem being considered, the optimal risk is always negative in this portion of the continuation region and, for fixed t , becomes increasingly negative as x increases from 0; in this portion of the continuation region the risk decreases from -.03 to -.07 at $t = 0.75$, from -.15 to -.37 at $t = 0.45$, and from -1.57 to -2.67 at $t = 0.04$.

TABLE 15. APPROXIMATION TO BOUNDARY FOR EXAMPLE 2.7.

t	$\bar{x}(t)$	t	$\bar{x}(t)$
0.001	0.027	0.55	0.341
0.002	0.038	0.60	0.333
0.005	0.059	0.65	0.321
0.01	0.082	0.70	0.306
0.02	0.114	0.75	0.287
0.03	0.138	0.80	0.263
0.04	0.157	0.82	0.252
0.05	0.174	0.84	0.240
0.06	0.188	0.86	0.226
0.07	0.201	0.88	0.211
0.08	0.213	0.90	0.194
0.09	0.223	0.91	0.185
0.10	0.233	0.92	0.175
0.12	0.250	0.93	0.165
0.14	0.265	0.94	0.153
0.16	0.277	0.95	0.140
0.18	0.289	0.96	0.126
0.20	0.298	0.97	0.109
0.25	0.318	0.98	0.090
0.30	0.332	0.99	0.064
0.35	0.341	0.995	0.045
0.40	0.346	0.998	0.028
0.45	0.348	0.999	0.020
0.50	0.346		

TABLE 16. ERRORS IN ESTIMATION OF RISK FOR MODIFIED ANSCOMBE PROBLEM

		$z = y/s^{1/2}$				
s		0	1	2	3	4
10	ER*	-.6(-4)	-.1(-3)			
	RER**	.2(-4)	.3(-4)			
10 ²	ER	-.5(-3)	-.5(-3)	-.2(-4)		
	RER	.6(-4)	.4(-4)	.1(-5)		
10 ³	ER	-.2(-2)	-.3(-2)	.4(-4)	.7(-4)	
	RER	.8(-4)	.7(-4)	-.6(-6)	-.7(-6)	
10 ⁴	ER	-.4(-2)	-.1(-2)	.1(-2)	.2(-3)	
	RER	.5(-4)	.8(-5)	-.5(-5)	-.7(-6)	
10 ⁵	ER	-.2(-1)	-.8(-2)	.3(-2)	.9(-3)	.2(-4)
	RER	.6(-4)	.2(-4)	-.5(-5)	-.1(-5)	-.2(-7)
10 ⁶	ER	-.2(-1)	-.1(-1)	.5(-2)	.1(-2)	-.5(-5)
	RER	.3(-4)	.1(-4)	-.3(-5)	-.4(-6)	.1(-8)

*ER = Error = Estimate of risk - optimal risk

**RER = Relative error = Error/optimal risk

TABLE 17. ERRORS IN ESTIMATES OF BOUNDARY FOR MODIFIED ANSCOMBE PROBLEM

Phase	Last s value	Δy	Maximum AER*	Maximum ARER**
1	1.052	.005	.00010	.004
2	1.234	.010	.00025	.004
3	1.962	.020	.00057	.002
4	4.874	.040	.00124	.001
5	16.522	.080	.00222	.0008
6	63.114	.160	.00443	.0004
7	249.482	.320	.00564	.0003
8	994.954	.640	.01246	.0003
9	3,976.842	1.280	.05500	.0005
10	15,904.394	2.560	.16643	.0007
11	63,614.602	5.120	.50157	.001
12	254,455.434	10.240	1.34060	.001
13	1,017,818.762	20.480	3.29124	.001

*AER = Absolute error = absolute value of $\hat{y} - \bar{y}$

**ARER = Absolute relative error = absolute value of $(\hat{y} - \bar{y})/\bar{y}$

TABLE 18. ESTIMATES OF STOPPING BOUNDARY FOR SEQUENTIAL ANALYSIS PROBLEM

$t = 1/s$	$\bar{x}(s) = \bar{y}(s)/s$	$\bar{z}(s) = \bar{y}(s)/s^2$	$\bar{\beta}(s) = 1 - \Phi(\bar{z}(s))$
10.00	.02499	.0079	.4968
9.50	.02630	.0085	.4966
9.00	.02776	.0093	.4963
8.50	.02939	.0101	.4960
8.00	.03123	.0110	.4956
7.50	.03331	.0122	.4951
7.00	.03569	.0135	.4946
6.50	.03843	.0151	.4940
6.00	.04163	.0170	.4932
5.50	.04541	.0194	.4923
5.00	.04995	.0223	.4911
4.50	.05550	.0262	.4896
4.00	.06243	.0312	.4875
3.50	.07130	.0381	.4848
3.00	.08315	.0480	.4809
2.50	.09961	.0630	.4749
2.00	.1240	.0877	.4651
1.50	.1636	.1336	.4469
1.40	.1744	.1474	.4414
1.30	.1865	.1636	.4350
1.20	.2004	.1830	.4274
1.15	.2080	.1940	.4231
1.10	.2162	.2061	.4184
1.05	.2250	.2196	.4131
1.00	.2344	.2344	.4073
.95	.2451	.2515	.4007
.90	.2563	.2702	.3935
.85	.2679	.2906	.3857
.80	.2805	.3136	.3769
.75	.2940	.3395	.3671
.70	.3085	.3688	.3561
.65	.3240	.4019	.3439
.60	.3409	.4401	.3299
.55	.3590	.4840	.3142
.50	.3781	.5348	.2964
.48	.3862	.5574	.2886
.46	.3943	.5814	.2805
.44	.4026	.6069	.2719
.42	.4111	.6343	.2629
.40	.4198	.6637	.2534
.38	.4285	.6951	.2435
.36	.4373	.7288	.2331
.34	.4461	.7651	.2221
.32	.4550	.8042	.2106
.30	.4637	.8466	.1986

TABLE 18. (continued)

$t = 1/s$	$\bar{x}(s)=\bar{y}(s)/s$	$\bar{z}(s)=\bar{y}(s)/s^{1/2}$	$\bar{\beta}(s)=1-\phi(\bar{z}(s))$
.28	.4724	.8928	.1860
.26	.4809	.9432	.1728
.24	.4892	.9986	.1590
.22	.4969	1.0595	.1447
.20	.5038	1.1264	.1300
.19	.5069	1.1629	.1224
.18	.5097	1.2013	.1148
.17	.5121	1.2421	.1071
.16	.5144	1.2859	.09924
.15	.5159	1.3319	.09144
.14	.5170	1.3818	.08351
.13	.5175	1.4352	.07562
.12	.5170	1.4926	.06778
.11	.5158	1.5552	.05995
.10	.5134	1.6234	.05225
.09	.5096	1.6985	.04471
.08	.5039	1.7816	.03741
.07	.4961	1.8751	.03039
.06	.4855	1.9819	.02375
.05	.4707	2.1052	.01764
.04	.4507	2.2534	.01212
.03	.4222	2.4376	.007392
.02	.3797	2.6848	.003629
.01	.3074	3.0738	.001057
9(-3)	.2969	3.1294	.8759(-3)
8(-3)	.2853	3.1903	.7107(-3)
7(-3)	.2726	3.2583	.5605(-3)
6(-3)	.2583	3.3345	.4273(-3)
5(-3)	.2420	3.4231	.3096(-3)
4(-3)	.2231	3.5276	.2097(-3)
3(-3)	.2004	3.6581	.1271(-3)
2(-3)	.1714	3.8329	.6334(-4)
1(-3)	.1300	4.1118	.1964(-4)
9(-4)	.1246	4.1518	.1650(-4)
8(-4)	.1187	4.1964	.1357(-4)
7(-4)	.1123	4.2461	.1088(-4)
6(-4)	.1054	4.3028	.8440(-5)
5(-4)	.09767	4.3681	.6271(-5)
4(-4)	.08895	4.4473	.4352(-5)
3(-4)	.07874	4.5462	.2733(-5)
2(-4)	.06620	4.6811	.1428(-5)
1(-4)	.04902	4.9020	.4749(-6)

TABLE 18. (continued)

$t = 1/s$	$\bar{x}(s)=\bar{y}(s)/s$	$\bar{z}(s)=\bar{y}(s)/s^{1/2}$	$\bar{\beta}(s)=1-\phi(\bar{z}(s))$
9(-5)	.04681	4.9346	.4021(-6)
8(-5)	.04446	4.9709	.3337(-6)
7(-5)	.04193	5.0113	.2707(-6)
6(-5)	.03918	5.0578	.2124(-6)
5(-5)	.03614	5.1113	.1602(-6)
4(-5)	.03274	5.1773	.1128(-6)
3(-5)	.02881	5.2603	.7204(-7)
2(-5)	.02403	5.3742	.3854(-7)
1(-5)	.01759	5.5638	.1323(-7)
9(-6)	.01678	5.5917	.1127(-7)
8(-6)	.01590	5.6227	.9428(-8)
7(-6)	.01497	5.6581	.7675(-8)
6(-6)	.01396	5.6986	.6056(-8)
5(-6)	.01285	5.7457	.4590(-8)
4(-6)	.01161	5.8031	.3266(-8)
3(-6)	.01018	5.8759	.2110(-8)
2(-6)	.008453	5.9770	.1140(-8)
1(-6)	.006146	6.1456	.3999(-9)

TABLE 19. ESTIMATES OF BAYES RISK FOR SEQUENTIAL ANALYSIS PROBLEM*

$\tau_0 = 1/s_0$	$z_0 = \mu_0/\sigma_0$							
	0	0.5	1.0	1.5	2.0	3.0	4.0	5.0
5.00	.1759							
2.00	.2667							
1.00	.3414							
.50	.3876							
.20	.3765	.3183	.1828					
.10	.3322	.2910	.1934	.9099(-1)				
.05	.2775	.2468	.1734	.9498(-1)	.3737(-1)			
.02	.2072	.1856	.1338	.7816(-1)	.3705(-1)			
.01	.1614	.1447	.1048	.6201(-1)	.3052(-1)	.3763(-2)		
5(-3)	.1234	.1106	.8001(-1)	.4730(-1)	.2343(-1)	.3697(-2)		
2(-3)	.8468(-1)	.7576(-1)	.5447(-1)	.3187(-1)	.1559(-1)	.2642(-2)		
1(-3)	.6284(-1)	.5611(-1)	.4011(-1)	.2321(-1)	.1117(-1)	.1858(-2)	.2094(-3)	
5(-4)	.4619(-1)	.4118(-1)	.2926(-1)	.1674(-1)	.7902(-2)	.1253(-2)	.1833(-3)	
2(-4)	.3040(-1)	.2704(-1)	.1907(-1)	.1076(-1)	.4954(-2)	.7210(-3)	.1167(-3)	
1(-4)	.2199(-1)	.1953(-1)	.1371(-1)	.7660(-2)	.3466(-2)	.4698(-3)	.7563(-4)	
5(-5)	.1583(-1)	.1405(-1)	.9819(-2)	.5441(-2)	.2425(-2)	.3056(-3)	.4697(-4)	.6465(-5)
2(-5)	.1020(-1)	.9032(-2)	.6286(-2)	.3451(-2)	.1512(-2)	.1743(-3)	.2409(-4)	.4874(-5)
1(-5)	.7283(-2)	.6446(-2)	.4475(-2)	.2444(-2)	.1060(-2)	.1149(-3)	.1429(-4)	.3309(-5)
5(-6)	.5191(-2)	.4591(-2)	.3180(-2)	.1729(-2)	.7432(-3)	.7639(-4)	.8430(-5)	.2092(-5)
2(-6)	.3308(-2)	.2924(-2)	.2021(-2)	.1094(-2)	.4660(-3)	.4517(-4)	.4185(-5)	.1071(-5)
1(-6)	.2349(-2)	.2075(-2)	.1432(-2)	.7733(-3)	.3278(-3)	.3066(-4)	.2480(-5)	.6247(-6)

*The quantity tabulated is $BR = \text{Bayes risk}/k^{2/3} c^{1/3} \sigma^{2/3} = E\{d(Y(S), S)\} - s_0^{-1}$.

06

TABLE 20. ESTIMATES OF BAYES EXPECTED COST OF SAMPLING FOR SEQUENTIAL ANALYSIS PROBLEM*

$\tau_0 = 1/s_0$	$z_0 = \mu_0/\sigma_0$							
	0	0.5	1.0	1.5	2.0	3.0	4.0	5.0
5.00	.2501(-2)							
2.00	.1457(-1)							
1.00	.4931(-1)							
.50	.1079							
.20	.1575	.1166	.2389(-1)					
.10	.1606	.1345	.7380(-1)	.1243(-1)				
.05	.1459	.1275	.8386(-1)	.3804(-1)	.4990(-2)			
.02	.1166	.1040	.7382(-1)	.4141(-1)	.1738(-1)			
.01	.9410(-1)	.8437(-1)	.6101(-1)	.3583(-1)	.1710(-1)	.5592(-3)		
5(-3)	.7391(-1)	.6636(-1)	.4824(-1)	.2873(-1)	.1429(-1)	.1805(-2)		
2(-3)	.5209(-1)	.4673(-1)	.3388(-1)	.2013(-1)	.1007(-1)	.1698(-2)		
1(-3)	.3927(-1)	.3517(-1)	.2537(-1)	.1494(-1)	.7384(-2)	.1293(-2)	.5924(-4)	
5(-4)	.2925(-1)	.2614(-1)	.1874(-1)	.1090(-1)	.5293(-2)	.9086(-3)	.1063(-3)	
2(-4)	.1951(-1)	.1739(-1)	.1236(-1)	.7081(-2)	.3347(-2)	.5358(-3)	.8427(-4)	
1(-4)	.1423(-1)	.1266(-1)	.8949(-2)	.5066(-2)	.2347(-2)	.3509(-3)	.5855(-4)	
5(-5)	.1031(-1)	.9163(-2)	.6442(-2)	.3610(-2)	.1642(-2)	.2278(-3)	.3787(-4)	.2344(-5)
2(-5)	.6685(-2)	.5929(-2)	.4144(-2)	.2295(-2)	.1022(-2)	.1286(-3)	.2000(-4)	.3204(-5)
1(-5)	.4793(-2)	.4246(-2)	.2958(-2)	.1627(-2)	.7151(-3)	.8390(-4)	.1200(-4)	.2469(-5)
5(-6)	.3426(-2)	.3033(-2)	.2106(-2)	.1152(-2)	.5006(-3)	.5512(-4)	.7106(-5)	.1669(-5)
2(-6)	.2190(-2)	.1937(-2)	.1341(-2)	.7291(-3)	.3132(-3)	.3208(-4)	.3518(-5)	.8997(-6)
1(-6)	.1558(-2)	.1377(-2)	.9518(-3)	.5155(-3)	.2200(-3)	.2153(-4)	.2068(-5)	.5382(-6)

*The quantity tabulated is $ECS = \text{Bayes expected cost of sampling}/k^{2/3} c^{1/3} \sigma^{2/3} = E(S^{-1}) - s_0^{-1}$.

16

TABLE 21. ESTIMATES OF STOPPING BOUNDARY FOR ONE-ARMED BANDIT PROBLEM

$t = 1/s$	$\bar{x}(s)=\bar{y}(s)/s$	$\bar{z}(s)=\bar{y}(s)/s^{1/2}$	$\bar{B}(s)=\phi(\bar{z}(s))$
.9995	-.01430	-.0143	.4943
.999	-.02026	-.0203	.4919
.995	-.04524	-.0454	.4819
.99	-.06391	-.0642	.4744
.98	-.09022	-.0911	.4637
.97	-.1103	-.1120	.4554
.96	-.1272	-.1298	.4484
.95	-.1420	-.1456	.4421
.94	-.1554	-.1603	.4363
.93	-.1675	-.1737	.4310
.92	-.1789	-.1865	.4260
.91	-.1894	-.1986	.4213
.90	-.1993	-.2101	.4168
.88	-.2177	-.2321	.4082
.86	-.2344	-.2528	.4002
.84	-.2498	-.2725	.3926
.82	-.2640	-.2915	.3853
.80	-.2775	-.3103	.3782
.78	-.2900	-.3284	.3713
.76	-.3017	-.3461	.3646
.74	-.3129	-.3637	.3580
.72	-.3234	-.3811	.3516
.70	-.3333	-.3984	.3452
.68	-.3428	-.4157	.3388
.66	-.3517	-.4329	.3325
.64	-.3602	-.4503	.3262
.62	-.3683	-.4678	.3200
.60	-.3760	-.4854	.3137
.58	-.3832	-.5032	.3074
.56	-.3901	-.5212	.3011
.54	-.3965	-.5396	.2947
.52	-.4026	-.5583	.2883
.50	-.4086	-.5778	.2817
.48	-.4138	-.5973	.2752
.46	-.4187	-.6173	.2685
.44	-.4231	-.6379	.2618
.42	-.4272	-.6592	.2549
.40	-.4309	-.6813	.2479
.38	-.4340	-.7041	.2407
.36	-.4367	-.7278	.2334
.34	-.4388	-.7525	.2259
.32	-.4405	-.7787	.2181
.30	-.4415	-.8061	.2101
.28	-.4419	-.8351	.2018
.26	-.4416	-.8660	.1932
.24	-.4404	-.8990	.1843
.22	-.4383	-.9345	.1750
.20	-.4354	-.9736	.1651

TABLE 21. (continued)

$t = 1/s$	$\bar{x}(s)=\bar{y}(s)/s$	$\bar{z}(s)=\bar{y}(s)/s^{1/2}$	$\bar{B}(s)=\phi(\bar{z}(s))$
.18	-.4311	-1.0160	.1548
.16	-.4252	-1.0630	.1439
.14	-.4176	-1.1161	.1322
.12	-.4076	-1.1767	.1197
.10	-.3946	-1.2477	.1061
.09	-.3866	-1.2886	.09877
.08	-.3773	-1.3340	.09110
.07	-.3665	-1.3854	.08296
.06	-.3540	-1.4450	.07423
.05	-.3387	-1.5146	.06494
.04	-.3198	-1.5988	.05493
.03	-.2956	-1.7069	.04392
.02	-.2626	-1.8569	.03166
.01	-.2108	-2.1081	.01751
9(-3)	-.2035	-2.1454	.01596
8(-3)	-.1956	-2.1868	.01438
7(-3)	-.1869	-2.2336	.01275
6(-3)	-.1772	-2.2871	.01109
5(-3)	-.1662	-2.3500	.009388
4(-3)	-.1534	-2.4258	.007638
3(-3)	-.1381	-2.5222	.005832
2(-3)	-.1188	-2.6554	.003961
1(-3)	-.09090	-2.8744	.002024
9(-4)	-.08720	-2.9067	.001826
8(-4)	-.08321	-2.9420	.001630
7(-4)	-.07892	-2.9829	.001428
6(-4)	-.07420	-3.0293	.001226
5(-4)	-.06895	-3.0836	.001023
4(-4)	-.06297	-3.1487	.8201(-3)
3(-4)	-.05597	-3.2316	.6155(-3)
2(-4)	-.04730	-3.3445	.4122(-3)
1(-4)	-.03533	-3.5327	.2057(-3)
9(-5)	-.03378	-3.5605	.1851(-3)
8(-5)	-.03213	-3.5917	.1643(-3)
7(-5)	-.03034	-3.6263	.1438(-3)
6(-5)	-.02838	-3.6640	.1242(-3)
5(-5)	-.02625	-3.7116	.1030(-3)
4(-5)	-.02383	-3.7680	.8230(-4)
3(-5)	-.02103	-3.8397	.6162(-4)
2(-5)	-.01762	-3.9388	.4095(-4)
1(-5)	-.01297	-4.1013	.2055(-4)
9(-6)	-.01238	-4.1257	.1849(-4)
8(-6)	-.01175	-4.1532	.1640(-4)
7(-6)	-.01107	-4.1839	.1434(-4)
6(-6)	-.01033	-4.2190	.1228(-4)
5(-6)	-.009526	-4.2602	.1022(-4)
4(-6)	-.008620	-4.3101	.8163(-5)
3(-6)	-.007570	-4.3706	.6200(-5)
2(-6)	-.006307	-4.4597	.4107(-5)
1(-6)	-.004608	-4.6077	.2038(-5)

TABLE 22. ESTIMATES OF BAYES EXPECTED PAYOFF FOR ONE-ARMED BANDIT PROBLEM *

$\tau_0 = 1/s_0$	$z_0 = \mu_0/\sigma_0$						
	0	-0.5	-1.0	-1.5	-2.0	-3.0	-4.0
.50	.1774	.0035					
.20	.1059(1)	.2390					
.10	.2769(1)	.8971	.0839				
.05	.6420(1)	.2479(1)	.5175				
.02	.1784(2)	.7773(1)	.2375(1)	.3003			
.01	.3728(2)	.1708(2)	.5961(1)	.1249	.0268		
5(-3)	.7658(2)	.3618(2)	.1362(2)	.3592(1)	.3915		
2(-3)	.1953(3)	.9447(2)	.3758(2)	.1150(2)	.2219(1)		
1(-3)	.3940(3)	.1925(3)	.7838(2)	.2538(2)	.5867(1)		
5(-4)	.7920(3)	.3892(3)	.1607(3)	.5387(2)	.1370(2)	.0233	
2(-4)	.1987(4)	.9812(3)	.4093(3)	.1406(3)	.3824(2)	.6137	
1(-4)	.3981(4)	.1969(4)	.8245(3)	.2861(3)	.7992(2)	.2089(1)	
5(-5)	.7969(4)	.3946(4)	.1657(4)	.5783(3)	.1640(3)	.5440(1)	
2(-5)	.1994(5)	.9878(4)	.4154(4)	.1456(4)	.4176(3)	.1628(2)	
1(-5)	.3988(5)	.1977(5)	.8319(4)	.2920(4)	.8409(3)	.3482(2)	.0493
5(-6)	.7977(5)	.3955(5)	.1665(5)	.5850(4)	.1689(4)	.7247(2)	.4370
2(-6)	.1995(6)	.9890(5)	.4165(5)	.1464(5)	.4232(4)	.1860(3)	.2117(1)
1(-6)	.3989(6)	.1978(6)	.8330(5)	.2929(5)	.8477(4)	.3765(3)	.5354(1)

*The quantity tabulated is $BEP = \text{Bayes expected payoff}/\sigma^2\sigma_0^{-1} = -s_0^{1/2}[E\{d'(Y(S), S)\} - d'(y_0, s_0)]$.

94

TABLE 23. ESTIMATES OF BAYES EXPECTED SAMPLE SIZE FOR ONE-ARMED BANDIT PROBLEM *

$\tau_0 = 1/s_0$	$z_0 = \mu_0/\sigma_0$						
	0	-0.5	-1.0	-1.5	-2.0	-3.0	-4.0
.50	.5786	.0835					
.20	.2217(1)	.1047(1)					
.10	.4851(1)	.2637(1)	.7097				
.05	.1001(2)	.5782(1)	.2256(1)				
.02	.2528(2)	.1513(2)	.6942(1)	.1884(1)			
.01	.5052(2)	.3064(2)	.1480(2)	.5065(1)	.5210		
5(-3)	.1008(3)	.6157(2)	.3055(2)	.1154(2)	.2588(1)		
2(-3)	.2512(3)	.1542(3)	.7795(2)	.3125(2)	.9084(1)		
1(-3)	.5016(3)	.3086(3)	.1571(3)	.6434(2)	.2018(2)		
5(-4)	.1002(4)	.6172(3)	.3155(3)	.1308(3)	.4257(2)	.5929	
2(-4)	.2502(4)	.1543(4)	.7913(3)	.3308(3)	.1103(3)	.4387(1)	
1(-4)	.5003(4)	.3086(4)	.1584(4)	.6643(3)	.2236(3)	.1094(2)	
5(-5)	.1000(5)	.6172(4)	.3171(4)	.1332(4)	.4507(3)	.2417(2)	
2(-5)	.2501(5)	.1543(5)	.7931(4)	.3336(4)	.1132(4)	.6433(2)	
1(-5)	.5001(5)	.3086(5)	.1586(5)	.6677(4)	.2270(4)	.1314(3)	.1032(1)
5(-6)	.1000(6)	.6171(5)	.3173(5)	.1336(5)	.4544(4)	.2659(3)	.4101(1)
2(-6)	.2500(6)	.1543(6)	.7935(5)	.3340(5)	.1137(5)	.6696(3)	.1334(2)
1(-6)	.5000(6)	.3086(6)	.1587(6)	.6681(5)	.2274(5)	.1344(4)	.2907(2)

*The quantity tabulated is $EN = \text{Bayes expected sample size}/\sigma^2\sigma_0^{-2} = s_0[E(S^{-1}) - s_0^{-1}]$.

95

TABLE 24. STOPPING BOUNDARIES $\beta_Y(t)$ FOR ANSCOMBE'S PROBLEM WITH ETHICAL COST*

t	$\gamma=0.0$	$\gamma=0.1$	$\gamma=1.0$	$\gamma=10.0$
1(-6)	1.03(-6)	1.13(-6)	2.05(-6)	1.12(-5)
2(-6)	2.05(-6)	2.25(-6)	4.10(-6)	2.25(-5)
3(-6)	3.07(-6)	3.38(-6)	6.13(-6)	3.37(-5)
4(-6)	4.09(-6)	4.50(-6)	8.17(-6)	4.50(-5)
5(-6)	5.11(-6)	5.62(-6)	1.02(-5)	5.63(-5)
6(-6)	6.14(-6)	6.75(-6)	1.23(-5)	6.76(-5)
7(-6)	7.15(-6)	7.87(-6)	1.43(-5)	7.88(-5)
8(-6)	8.17(-6)	8.98(-6)	1.64(-5)	9.02(-5)
9(-6)	9.20(-6)	1.01(-5)	1.84(-5)	1.01(-4)
1(-5)	1.02(-5)	1.12(-5)	2.05(-5)	1.13(-4)
2(-5)	2.04(-5)	2.25(-5)	4.09(-5)	2.26(-4)
3(-5)	3.07(-5)	3.38(-5)	6.15(-5)	3.39(-4)
4(-5)	4.09(-5)	4.50(-5)	8.20(-5)	4.53(-4)
5(-5)	5.12(-5)	5.65(-5)	1.03(-4)	5.65(-4)
6(-5)	6.16(-5)	6.78(-5)	1.23(-4)	6.79(-4)
7(-5)	7.20(-5)	7.91(-5)	1.44(-4)	7.92(-4)
8(-5)	8.22(-5)	9.04(-5)	1.65(-4)	9.05(-4)
9(-5)	9.25(-5)	1.02(-4)	1.85(-4)	0.00102
1(-4)	1.03(-4)	1.13(-4)	2.06(-4)	0.00113
2(-4)	2.07(-4)	2.27(-4)	4.13(-4)	0.00223
3(-4)	3.11(-4)	3.42(-4)	6.21(-4)	0.00332
4(-4)	4.16(-4)	4.56(-4)	8.28(-4)	0.00437
5(-4)	5.20(-4)	5.71(-4)	0.00103	0.00540
6(-4)	6.25(-4)	6.86(-4)	0.00124	0.00642
7(-4)	7.28(-4)	8.01(-4)	0.00144	0.00741
8(-4)	8.34(-4)	9.15(-4)	0.00165	0.00839
9(-4)	9.39(-4)	0.00103	0.00185	0.00935
0.001	0.00104	0.00114	0.00206	0.0103
0.002	0.00208	0.00228	0.00405	0.0191
0.003	0.00311	0.00340	0.00599	0.0269
0.004	0.00412	0.00451	0.00786	0.0340
0.005	0.00513	0.00561	0.00969	0.0405
0.006	0.00612	0.00668	0.0115	0.0466
0.007	0.00710	0.00775	0.0132	0.0523
0.008	0.00807	0.00880	0.0149	0.0576
0.009	0.00903	0.00984	0.0166	0.0626
0.01	0.00998	0.0109	0.0182	0.0674
0.02	0.0190	0.0206	0.0332	0.1060
0.03	0.0274	0.0295	0.0462	0.1339
0.04	0.0353	0.0378	0.0578	0.1560
0.05	0.0427	0.0457	0.0684	0.1741
0.06	0.0498	0.0531	0.0783	0.1896
0.07	0.0566	0.0602	0.0876	0.2034
0.08	0.0631	0.0670	0.0962	0.2153
0.09	0.0694	0.0736	0.1043	0.2260
0.10	0.0754	0.0799	0.1120	0.2357
0.11	0.0813	0.0860	0.1193	0.2445
0.12	0.0870	0.0918	0.1263	0.2525
0.13	0.0926	0.0976	0.1330	0.2600
0.14	0.0980	0.1032	0.1394	0.2669
0.15	0.1033	0.1085	0.1456	0.2735

TABLE 24 (continued)

t	$\gamma=0.0$	$\gamma=0.1$	$\gamma=1.0$	$\gamma=10.0$
0.16	0.1085	0.1138	0.1515	0.2795
0.17	0.1135	0.1190	0.1573	0.2852
0.18	0.1184	0.1240	0.1629	0.2906
0.19	0.1233	0.1290	0.1683	0.2958
0.20	0.1280	0.1339	0.1736	0.3006
0.22	0.1374	0.1434	0.1839	0.3100
0.24	0.1463	0.1525	0.1935	0.3183
0.26	0.1550	0.1612	0.2026	0.3259
0.28	0.1635	0.1698	0.2114	0.3330
0.30	0.1718	0.1781	0.2198	0.3396
0.32	0.1798	0.1862	0.2279	0.3458
0.34	0.1877	0.1941	0.2358	0.3516
0.36	0.1955	0.2019	0.2434	0.3571
0.38	0.2031	0.2095	0.2508	0.3623
0.40	0.2106	0.2170	0.2579	0.3672
0.42	0.2180	0.2243	0.2650	0.3720
0.44	0.2253	0.2316	0.2718	0.3765
0.46	0.2325	0.2388	0.2786	0.3809
0.48	0.2397	0.2459	0.2852	0.3852
0.50	0.2468	0.2529	0.2917	0.3893
0.52	0.2540	0.2601	0.2982	0.3934
0.54	0.2610	0.2670	0.3045	0.3974
0.56	0.2680	0.2740	0.3108	0.4012
0.58	0.2751	0.2809	0.3170	0.4049
0.60	0.2821	0.2878	0.3231	0.4085
0.62	0.2891	0.2947	0.3292	0.4120
0.64	0.2961	0.3016	0.3353	0.4155
0.66	0.3032	0.3086	0.3414	0.4190
0.68	0.3104	0.3156	0.3475	0.4224
0.70	0.3176	0.3227	0.3536	0.4258
0.72	0.3249	0.3298	0.3598	0.4291
0.74	0.3324	0.3371	0.3660	0.4325
0.76	0.3399	0.3445	0.3723	0.4359
0.78	0.3477	0.3521	0.3787	0.4393
0.80	0.3556	0.3599	0.3853	0.4428
0.82	0.3639	0.3680	0.3921	0.4463
0.84	0.3725	0.3763	0.3990	0.4500
0.86	0.3814	0.3849	0.4063	0.4537
0.88	0.3908	0.3941	0.4139	0.4576
0.90	0.4009	0.4039	0.4219	0.4617
0.92	0.4118	0.4145	0.4307	0.4660
0.94	0.4241	0.4264	0.4403	0.4708
0.95	0.4309	0.4330	0.4458	0.4735
0.96	0.4383	0.4403	0.4517	0.4765
0.97	0.4467	0.4484	0.4583	0.4797
0.98	0.4566	0.4580	0.4660	0.4835
0.99	0.4694	0.4704	0.4761	0.4884
0.995	0.4784	0.4791	0.4831	0.4918
0.999	0.4904	0.4907	0.4924	0.4963
0.9995	0.4932	0.4935	0.4947	0.4975
1.0000	0.5000	0.5000	0.5000	0.5000

*t = currently available proportion of total potential information

 β_Y = nominal significance level

TABLE 25. BAYES PROPERTIES OF OPTIMAL PROCEDURES IN ANSCOMBE'S PROBLEM WITH ETHICAL COST

t_0	γ	z_0											
		0.0		0.5		1.0		1.5		2.0		3.0	
		BR (PR)	EN	BR (PR)	EN	BR (PR)	EN	BR (PR)	EN	BR (PR)	EN	BR (PR)	EN
1(-1)	0.0	1.78(.61)	1.76	1.62(.57)	1.43								
	0.1	1.83(.60)	1.69	1.66(.56)	1.37								
	1.0	2.16(.53)	1.29	1.99(.49)	.99								
	10.0	3.54(.34)	.47	3.18(.24)	.23								
5(-2)	0.0	2.55(.63)	2.91	2.34(.61)	2.47								
	0.1	2.63(.63)	2.80	2.42(.60)	2.36								
	1.0	3.20(.56)	2.15	2.99(.54)	1.76								
	10.0	5.81(.40)	.82	5.45(.35)	.54								
2(-2)	0.0	3.80(.66)	5.31	3.52(.65)	4.61	2.78(.62)	2.96						
	0.1	3.94(.66)	5.11	3.65(.65)	4.42	2.89(.61)	2.82						
	1.0	4.99(.60)	3.93	4.69(.60)	3.35	3.76(.55)	2.01						
	10.0	10.20(.46)	1.57	9.77(.45)	1.21	7.25(.29)	.41						
1(-2)	0.0	4.95(.68)	8.11	4.60(.68)	7.11	3.68(.66)	4.75						
	0.1	5.15(.68)	7.80	4.80(.67)	6.83	3.85(.65)	4.54						
	1.0	6.70(.63)	6.01	6.31(.63)	5.21	5.16(.60)	3.34						
	10.0	14.82(.51)	2.46	14.27(.50)	2.01	11.42(.42)	.98						
5(-3)	0.0	6.31(.70)	12.19	5.87(.70)	10.75	4.73(.69)	7.34	3.30(.65)	3.80				
	0.1	6.59(.70)	11.73	6.14(.69)	10.34	4.97(.68)	7.04	3.48(.64)	3.62				
	1.0	8.77(.65)	9.05	8.27(.65)	7.92	6.82(.64)	5.28	4.80(.59)	2.56				
	10.0	20.82(.54)	3.76	20.06(.54)	3.18	16.65(.50)	1.82	10.33(.31)	.49				
2(-3)	0.0	8.45(.72)	20.53	7.86(.72)	18.17	6.35(.72)	12.59	4.51(.70)	6.79	2.81(.65)	2.77		
	0.1	8.86(.72)	19.76	8.25(.72)	17.48	6.70(.71)	12.09	4.78(.70)	6.50	2.98(.64)	2.63		
	1.0	12.12(.68)	15.27	11.41(.68)	13.47	9.46(.68)	9.21	6.84(.66)	4.81	4.19(.58)	1.79		
	10.0	31.21(.59)	6.45	30.00(.59)	5.58	25.45(.58)	3.55	17.69(.50)	1.49	8.35(.13)	.13		

$$z_0 = \mu_0/\sigma_0$$

t_0 = proportion of total information in prior

$$BR = \text{Bayes risk}/\sigma_0^2\sigma_0^{-1}$$

PR = proportion of Bayes risk due to the experimental phase

$$EN = \text{Bayes expected sample size}/\sigma_0^2\sigma_0^{-2}$$

86

TABLE 25 (continued)

t_0	γ	z_0											
		0.0		0.5		1.0		1.5		2.0		3.0	
		BR (PR)	EN	BR (PR)	EN	BR (PR)	EN	BR (PR)	EN	BR (PR)	EN	BR (PR)	EN
1(-3)	0.0	10.34(.74)	30.15	9.61(.74)	26.72	7.77(.74)	18.61	5.55(.73)	10.18	3.54(.70)	4.34		
	0.1	10.88(.73)	29.03	10.13(.73)	25.72	8.22(.73)	17.90	5.89(.73)	9.77	3.76(.69)	4.14		
	1.0	15.17(.70)	22.47	14.26(.70)	19.87	11.80(.70)	13.72	8.60(.69)	7.37	5.49(.65)	2.98		
	10.0	41.25(.62)	9.59	39.54(.62)	8.38	33.69(.62)	5.54	24.39(.58)	2.64	13.66(.41)	.69		
5(-4)	0.0	12.50(.75)	44.00	11.60(.75)	39.01	9.36(.76)	27.21	6.68(.75)	14.98	4.32(.74)	6.51		
	0.1	13.19(.75)	42.37	12.16(.75)	37.56	9.92(.75)	26.19	7.12(.75)	14.40	4.61(.73)	6.24		
	1.0	18.71(.72)	32.85	17.55(.72)	29.08	14.47(.73)	20.20	10.58(.72)	11.00	6.90(.70)	4.65		
	10.0	53.42(.64)	14.15	51.03(.65)	12.44	43.46(.65)	8.41	32.12(.63)	4.29	19.63(.54)	1.48		
2(-4)	0.0	15.77(.77)	71.90	14.61(.77)	63.74	11.73(.78)	44.47	8.35(.78)	24.51	5.44(.77)	10.74		
	0.1	16.69(.77)	69.26	15.49(.77)	61.40	12.47(.77)	42.83	8.92(.77)	23.59	5.83(.77)	10.33		
	1.0	24.17(.74)	53.81	22.60(.74)	47.68	18.53(.75)	33.21	13.51(.75)	18.22	8.93(.74)	7.90		
	10.0	73.13(.68)	23.43	69.51(.68)	20.69	58.89(.69)	14.23	44.00(.68)	7.58	28.52(.64)	3.02		
1(-4)	0.0	18.57(.78)	103.73	17.19(.78)	91.92	13.74(.79)	64.05	9.74(.79)	35.24	6.35(.79)	15.43		
	0.1	19.71(.78)	99.95	18.26(.78)	88.57	14.64(.79)	61.72	10.42(.79)	33.95	6.81(.79)	14.86		
	1.0	28.94(.76)	77.78	27.00(.76)	68.91	22.01(.77)	47.99	15.98(.77)	26.36	10.60(.77)	11.49		
	10.0	91.05(.70)	34.11	86.22(.70)	30.17	72.62(.71)	20.89	54.30(.71)	11.30	36.00(.69)	4.72		
5(-5)	0.0	21.67(.79)	149.08	20.02(.80)	132.05	15.93(.80)	91.88	11.24(.81)	50.41	7.31(.81)	21.99	2.67(.75)	2.25
	0.1	23.05(.79)	143.68	21.31(.79)	127.27	17.01(.80)	88.56	12.04(.80)	48.59	7.86(.81)	21.19	2.86(.74)	2.16
	1.0	34.28(.77)	111.98	31.90(.77)	99.19	25.86(.78)	69.03	18.66(.79)	37.88	12.38(.79)	16.51	4.38(.70)	1.61
	10.0	111.73(.72)	49.44	105.42(.72)	43.77	88.18(.73)	30.40	65.74(.74)	16.58	44.09(.73)	7.10	12.11(.45)	.42

66

TABLE 25 (continued)

t_0	γ	z_0															
		0.0		0.5		1.0		1.5		2.0		3.0					
		BR	(PR)	EN	BR	(PR)	EN	BR	(PR)	EN	BR	(PR)	EN	BR	(PR)	EN	BR
2(-5)	0.0	26.24	(.81)	239.73	24.19	(.81)	212.20	19.13	(.81)	147.31	13.38	(.82)	80.45	8.66	(.83)	34.84	3.30
	0.1	27.97	(.80)	231.10	25.81	(.81)	204.57	20.46	(.81)	142.02	14.36	(.82)	77.58	9.32	(.83)	33.61	3.55
	1.0	42.25	(.79)	180.44	39.20	(.79)	159.75	31.53	(.80)	110.97	22.54	(.81)	60.68	14.89	(.81)	26.33	5.61
	10.0	143.62	(.74)	80.30	134.90	(.74)	71.11	111.75	(.75)	49.44	82.67	(.76)	27.07	55.75	(.76)	11.72	18.40
1(-5)	0.0	30.06	(.82)	342.35	27.66	(.82)	302.87	21.78	(.82)	209.94	15.13	(.83)	114.31	9.74	(.84)	49.22	3.78
	0.1	32.09	(.81)	330.09	29.56	(.82)	292.03	23.33	(.82)	202.44	16.26	(.83)	110.25	10.50	(.84)	47.49	4.08
	1.0	48.98	(.80)	258.05	45.34	(.80)	228.34	36.25	(.81)	158.40	25.72	(.82)	86.38	16.90	(.83)	37.30	6.56
	10.0	171.31	(.75)	115.45	160.38	(.76)	99.65	131.87	(.77)	71.05	96.83	(.78)	38.89	65.25	(.78)	16.87	23.32
5(-6)	0.0	34.19	(.83)	487.99	31.41	(.83)	431.54	24.62	(.83)	298.64	16.98	(.84)	162.10	10.86	(.85)	69.41	4.27
	0.1	36.56	(.82)	470.57	33.62	(.82)	416.15	26.40	(.83)	288.02	18.27	(.84)	156.36	11.73	(.85)	66.98	4.62
	1.0	56.33	(.81)	368.27	52.04	(.81)	325.75	41.37	(.82)	225.61	29.13	(.83)	122.66	19.02	(.84)	52.68	7.53
	10.0	202.24	(.77)	165.53	188.77	(.77)	146.52	154.06	(.78)	101.73	112.19	(.80)	55.57	75.34	(.80)	24.06	28.37
2(-6)	0.0	40.15	(.84)	777.63	36.81	(.84)	687.32	28.68	(.84)	474.90	19.61	(.85)	256.91	12.43	(.86)	109.29	4.94
	0.1	43.01	(.83)	749.99	39.47	(.84)	662.91	30.81	(.84)	458.08	21.13	(.85)	247.86	13.44	(.86)	105.47	5.35
	1.0	67.02	(.82)	587.63	61.75	(.82)	519.50	48.74	(.83)	359.21	33.96	(.84)	194.62	21.97	(.85)	83.03	8.83
	10.0	248.20	(.79)	265.49	230.80	(.79)	234.89	186.59	(.80)	162.83	134.32	(.81)	88.68	89.52	(.82)	38.19	35.20
1(-6)	0.0	45.03	(.84)	1104.72	41.23	(.84)	976.11	31.98	(.85)	673.67	21.72	(.86)	363.63	13.67	(.87)	154.07	5.44
	0.1	48.31	(.84)	1065.57	44.26	(.84)	941.53	34.39	(.85)	649.86	23.42	(.86)	350.83	14.79	(.87)	148.70	5.91
	1.0	75.86	(.83)	835.48	69.77	(.83)	738.36	54.78	(.84)	509.92	37.87	(.85)	275.62	24.31	(.86)	117.09	9.83
	10.0	286.93	(.80)	378.69	266.13	(.80)	334.91	213.69	(.81)	231.84	152.46	(.82)	125.92	100.92	(.83)	53.99	40.49

100

TABLE 26. DEVIATIONS IN ESTIMATION OF BOUNDARY AND RISK FOR ANSCOMBE'S PROBLEM WITH ETHICAL COST: CASE $\gamma = 1$.

Grid Spacing		Version 1				Version 2			
s	y	Ave(d)	Ave(d)	Max(d)	Ave(d)	Ave(d)	Max(d)	Boundary ^a	
1	1	-674	720	1233	-1151	1151	1863	Risk ^b	
	2 ⁻²	-158	197	312	-281	291	489		
	4 ⁻²	-40	54	100	-68	76	135		
	4 ⁻³	-9	19	40	-15	22	48		
2	1	-326	333	709	397	397	883		
	2 ⁻¹	-105	105	201	80	80	161		
	4 ⁻²	-27	27	52	16	16	33		
	4 ⁻³	-6	6	11	3	3	6		

The computation with grid spacing in $s = 4^{-4}$, in $y = 2^{-4}$ provides the baseline, one for each version of the backward induction algorithm.

^aThe deviations summarized are $d = (\hat{d}_1 - \hat{d}_2) \times 10^4$ at $s = 25(1)100$, where \hat{y}_2 is the approximation to the continuous boundary at baseline and \hat{y}_1 is the approximation at the grid spacing listed (both computed by the method EX). For this range of s values, \hat{y} varies between 8 and 21.

^bThe deviations summarized are $d = (\hat{d}_1 - \hat{d}_2) \times 10^4$ at all the grid points on the intersections of the lines $s = 25(1)100$, $y = 0(1)1$ and within the continuation region for the discrete time problem with the less refined grid spacing; here \hat{d}_2 is the result of the backward induction at baseline and \hat{d}_1 is the result at the grid spacing listed. The function \hat{d}_4 is symmetric in y , always positive, and, for fixed s , decreases as $|y|$ increases; inside the continuation region \hat{d}_4 decreases from 3.6 to 1.7 at $s = 25$ and from 6.7 to 1.6 at $s = 100$.

101

TABLE 27. DIFFERENCES IN TWO VERSIONS FOR ANSCOMBE'S PROBLEM WITH ETHICAL COST: CASE $\gamma = 1$.

Grid Spacing		Ave(d)	Ave(d)	Max(d)
s	y			
Boundary ^a				
1	1	479	479	1797
4 ⁻¹	2 ⁻¹	125	125	442
4 ⁻²	2 ⁻²	30	30	51
4 ⁻³	2 ⁻⁴	8	8	12
4 ⁻⁴	2 ⁻⁴	2	2	3
Risk ^b				
1	1	-727	727	1556
4 ⁻¹	2 ⁻¹	-188	188	359
4 ⁻²	2 ⁻²	-46	46	89
4 ⁻³	2 ⁻³	-12	12	22
4 ⁻⁴	2 ⁻⁴	-3	3	6

^aThe differences summarized are $d = (\hat{y}_1 - \hat{y}_2) \times 10^4$ at $s = 25(1)100$, where \hat{y}_1 and \hat{y}_2 are the approximations to the continuous time boundary produced by versions 1 and 2 respectively (both computed by the method EX).

^bThe differences summarized are $d = (\hat{d}_{41} - \hat{d}_{42}) \times 10^4$ at all the gridpoints on the intersections of the lines $s = 25(1)100$ and $y = 0(1)\infty$ and within the continuation region for both versions; here \hat{d}_{41} and \hat{d}_{42} are the results of the backward inductions for versions 1 and 2 respectively.

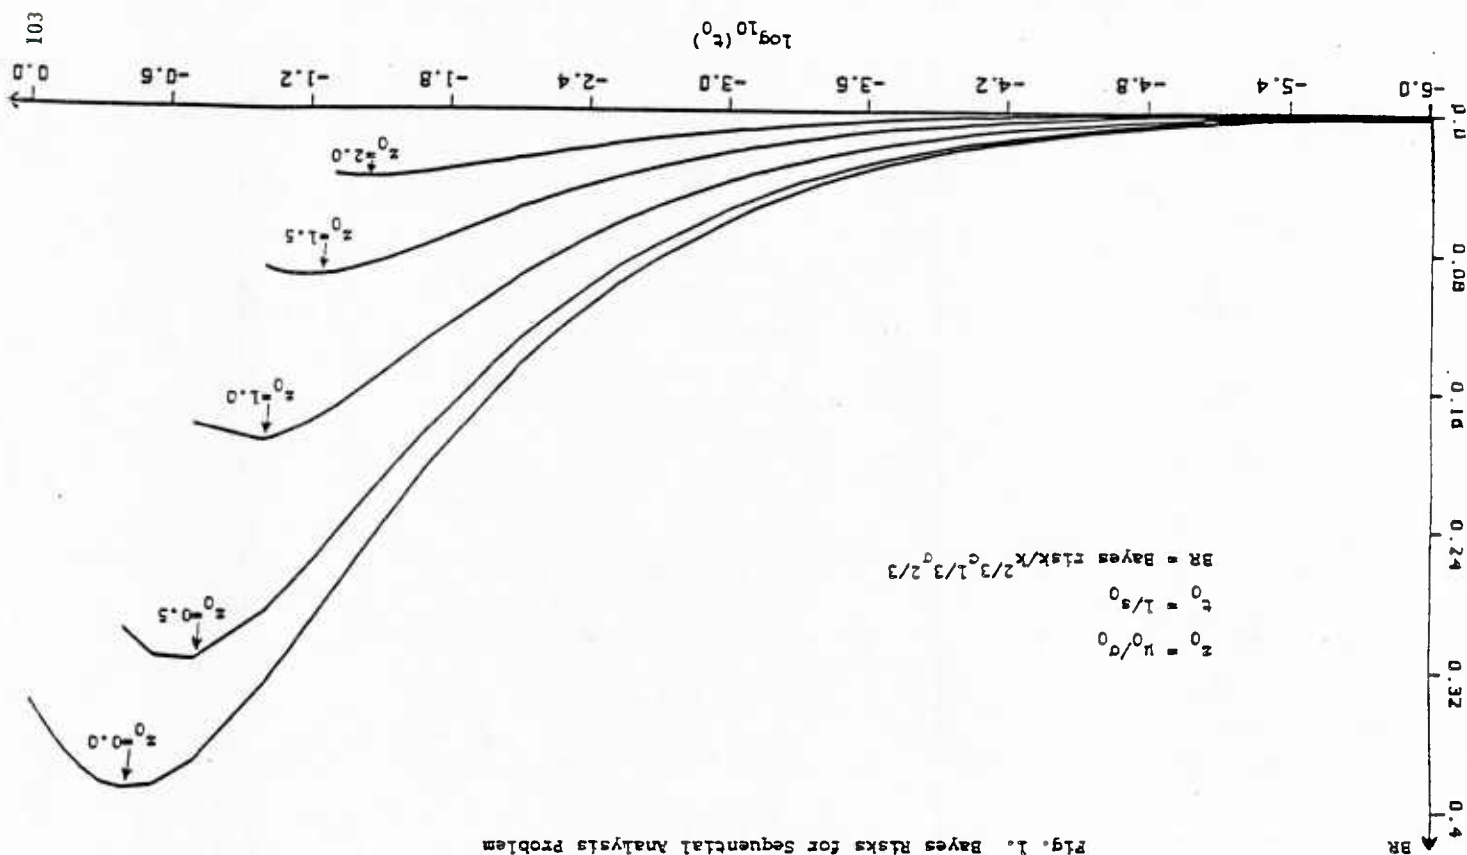
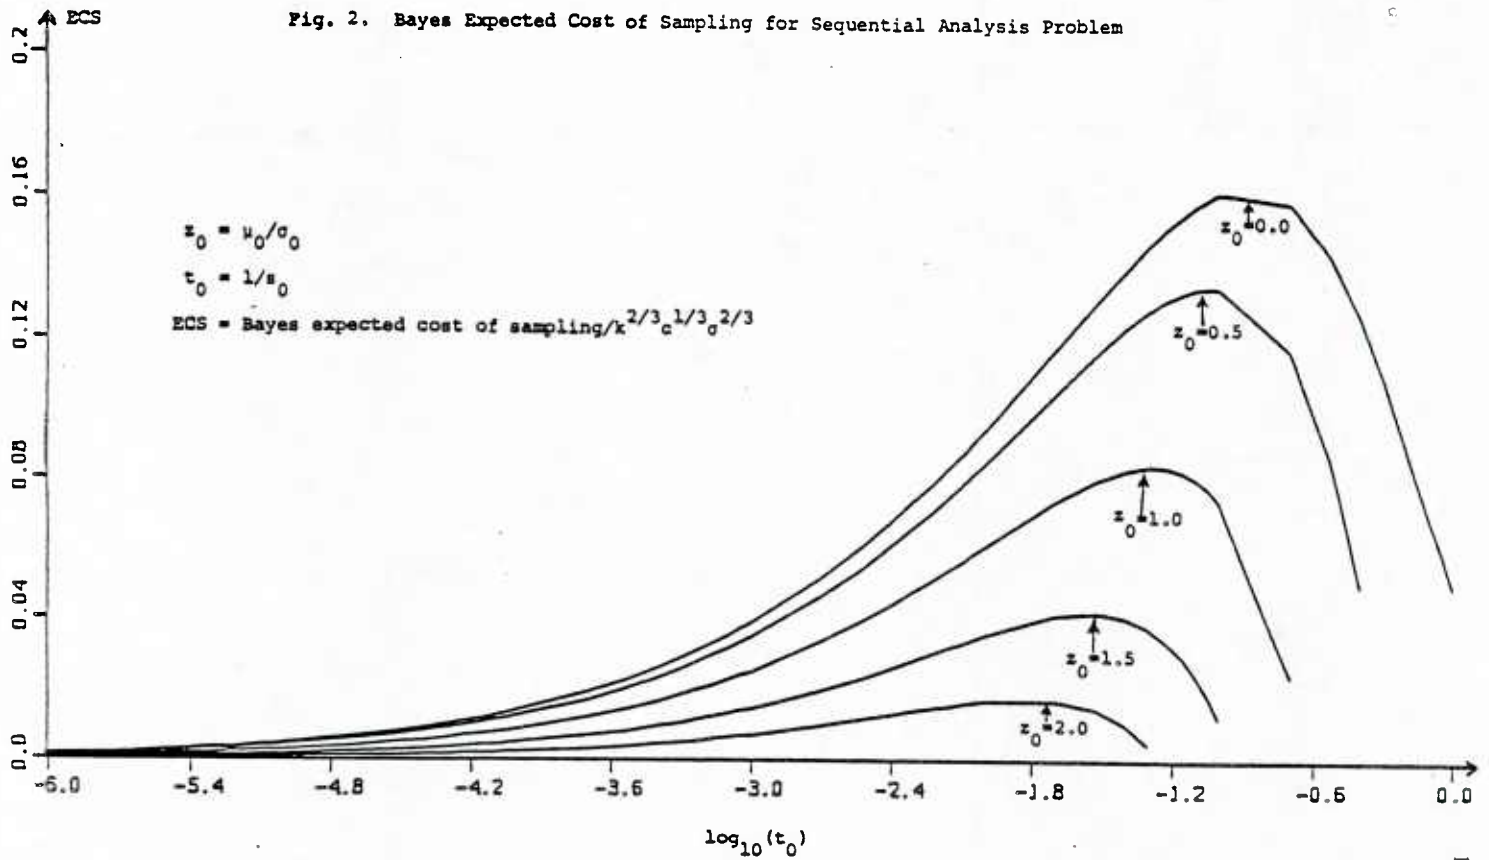


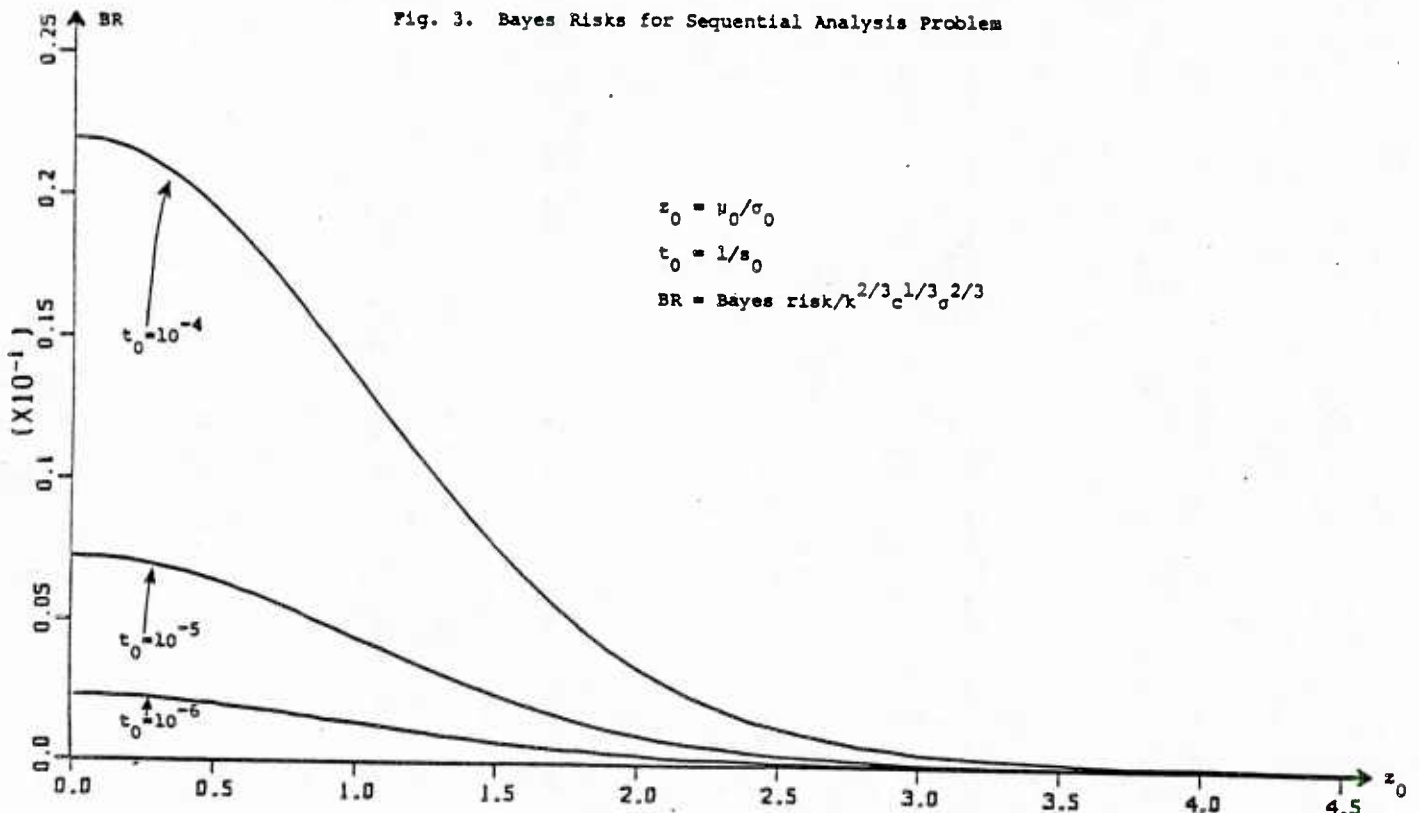
Fig. 1. Bayes Risks for Sequential Analysis Problem

Fig. 2. Bayes Expected Cost of Sampling for Sequential Analysis Problem

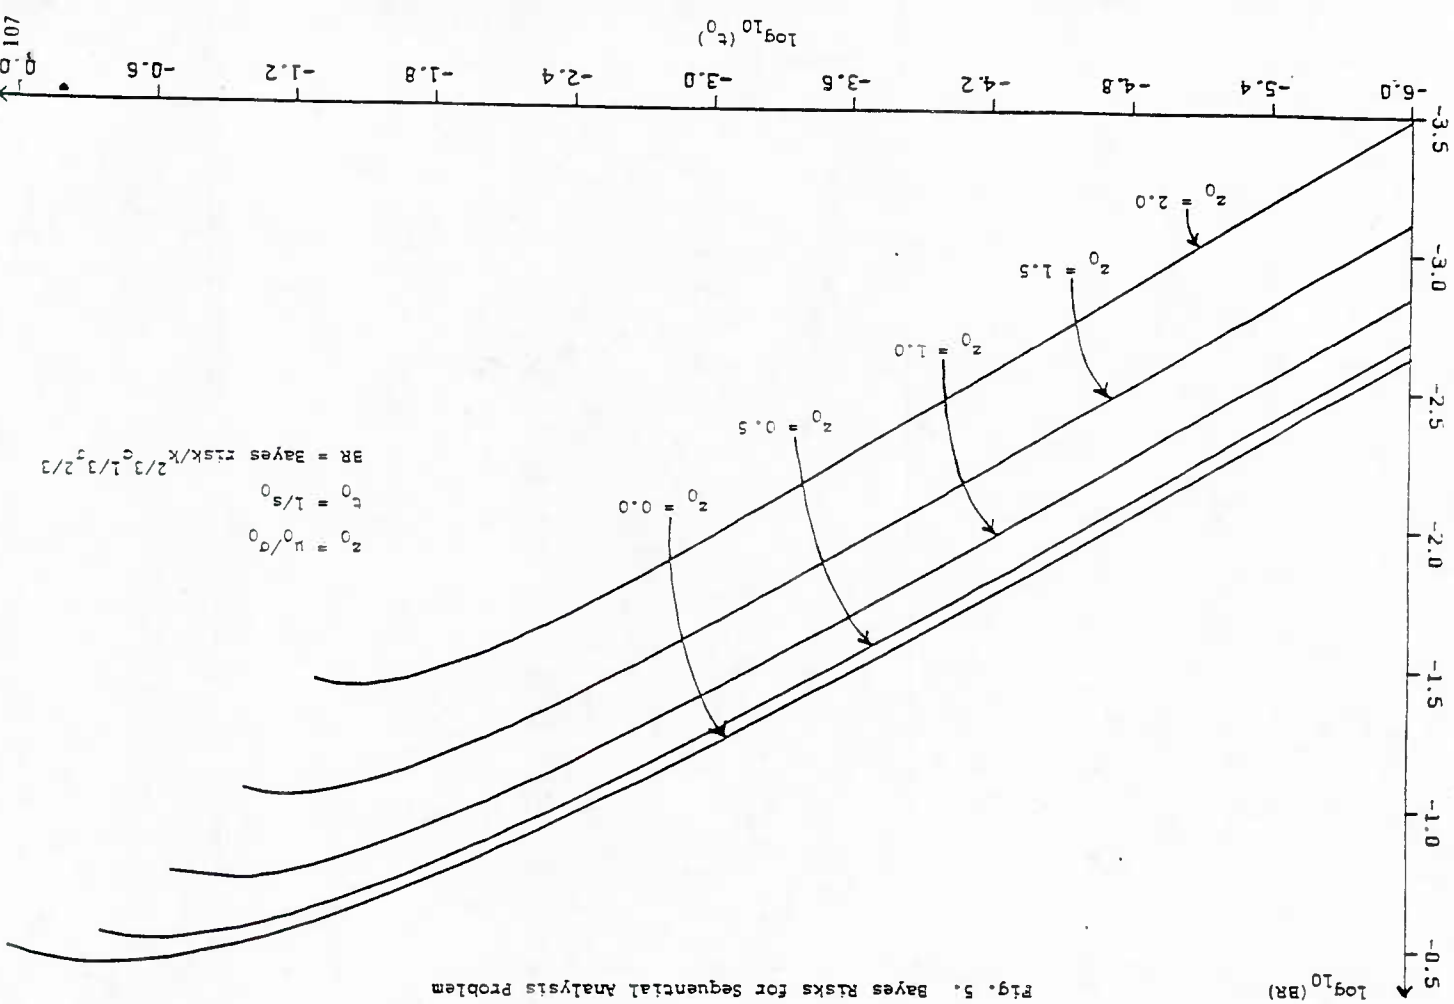
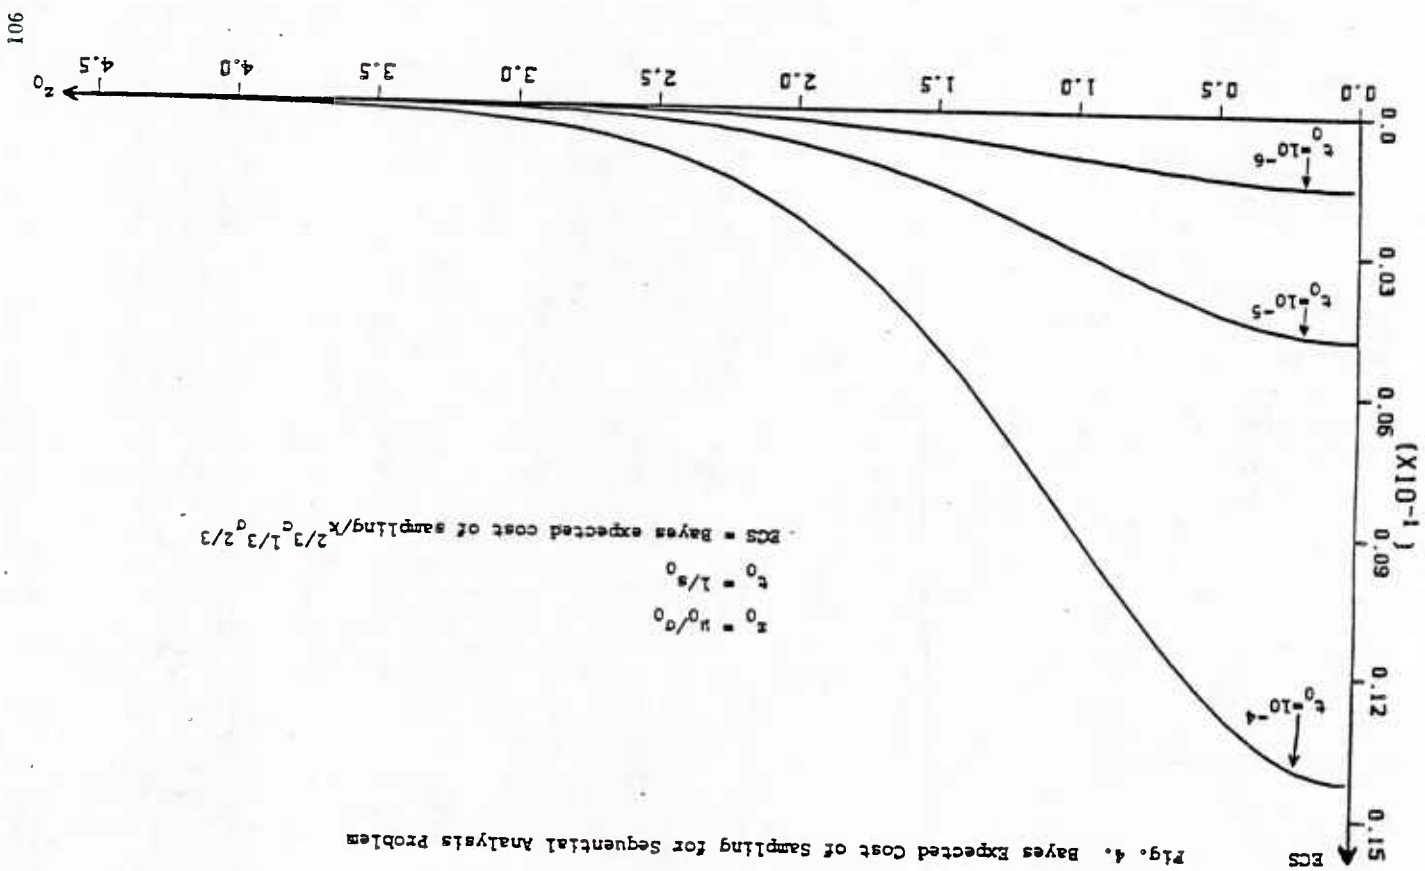


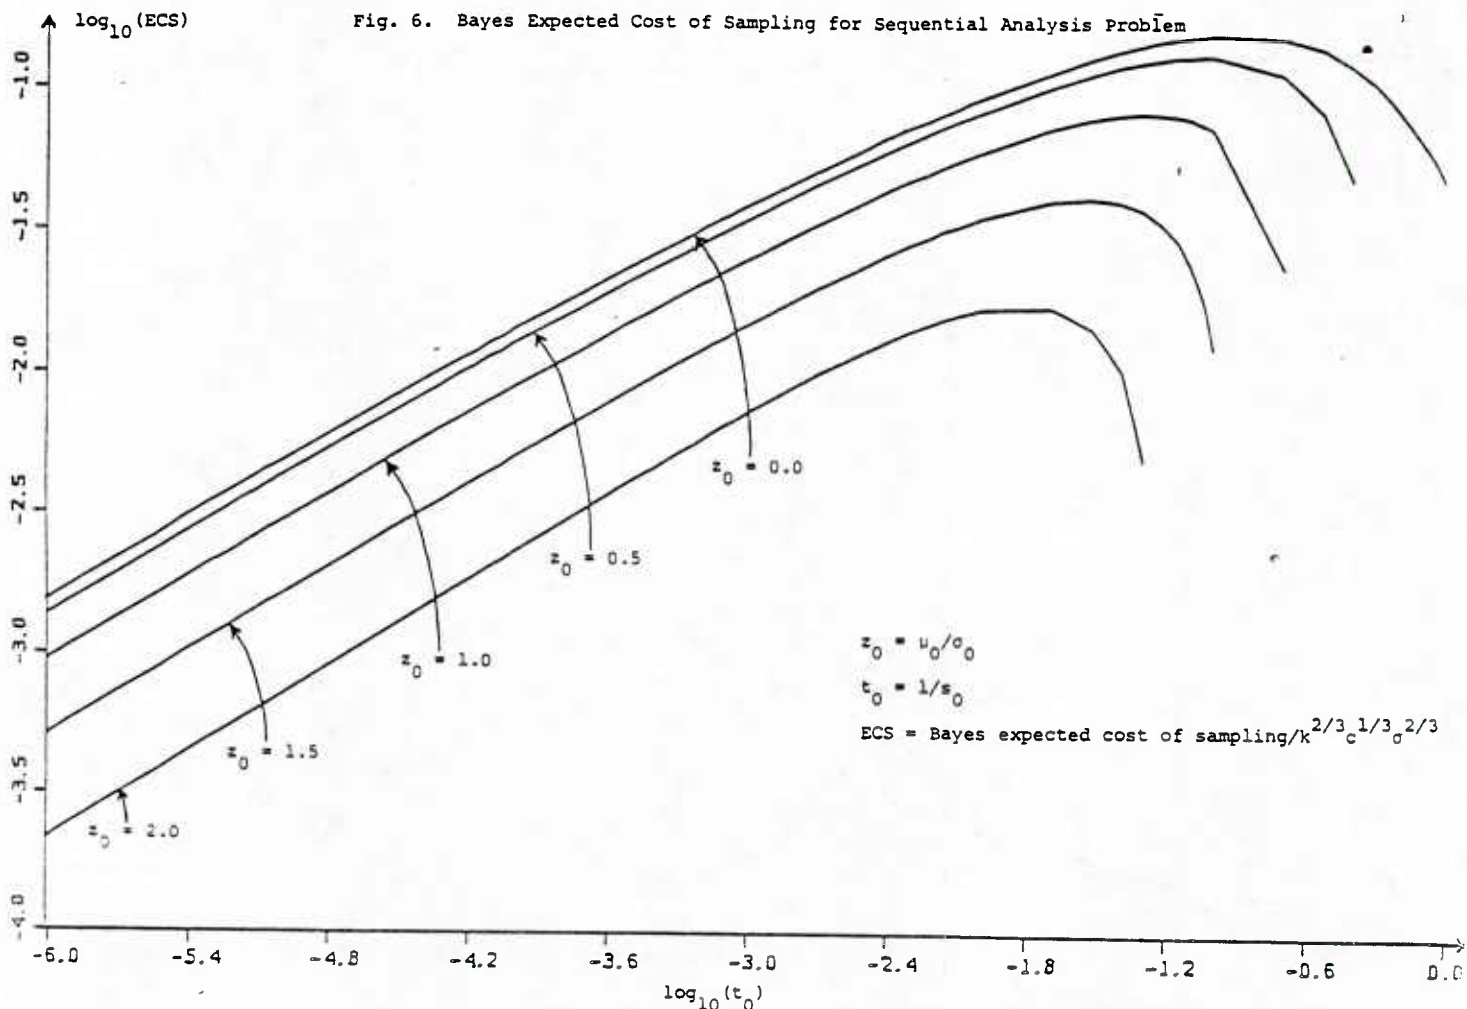
104

Fig. 3. Bayes Risks for Sequential Analysis Problem

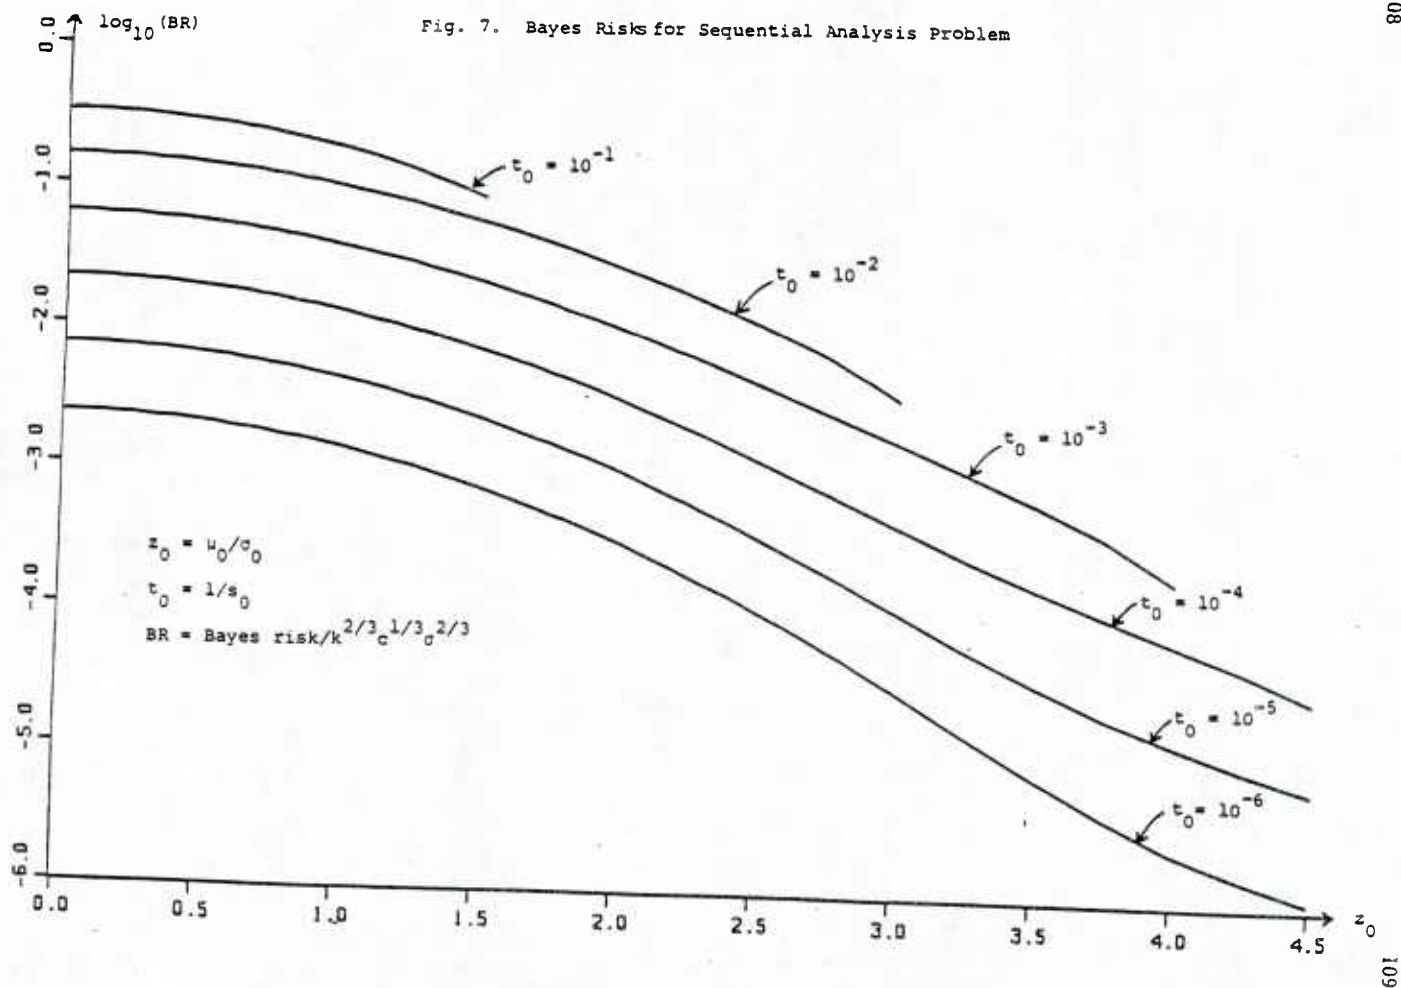


105





108



109

Fig. 8. Bayes Expected Cost of Sampling for Sequential Analysis Problem

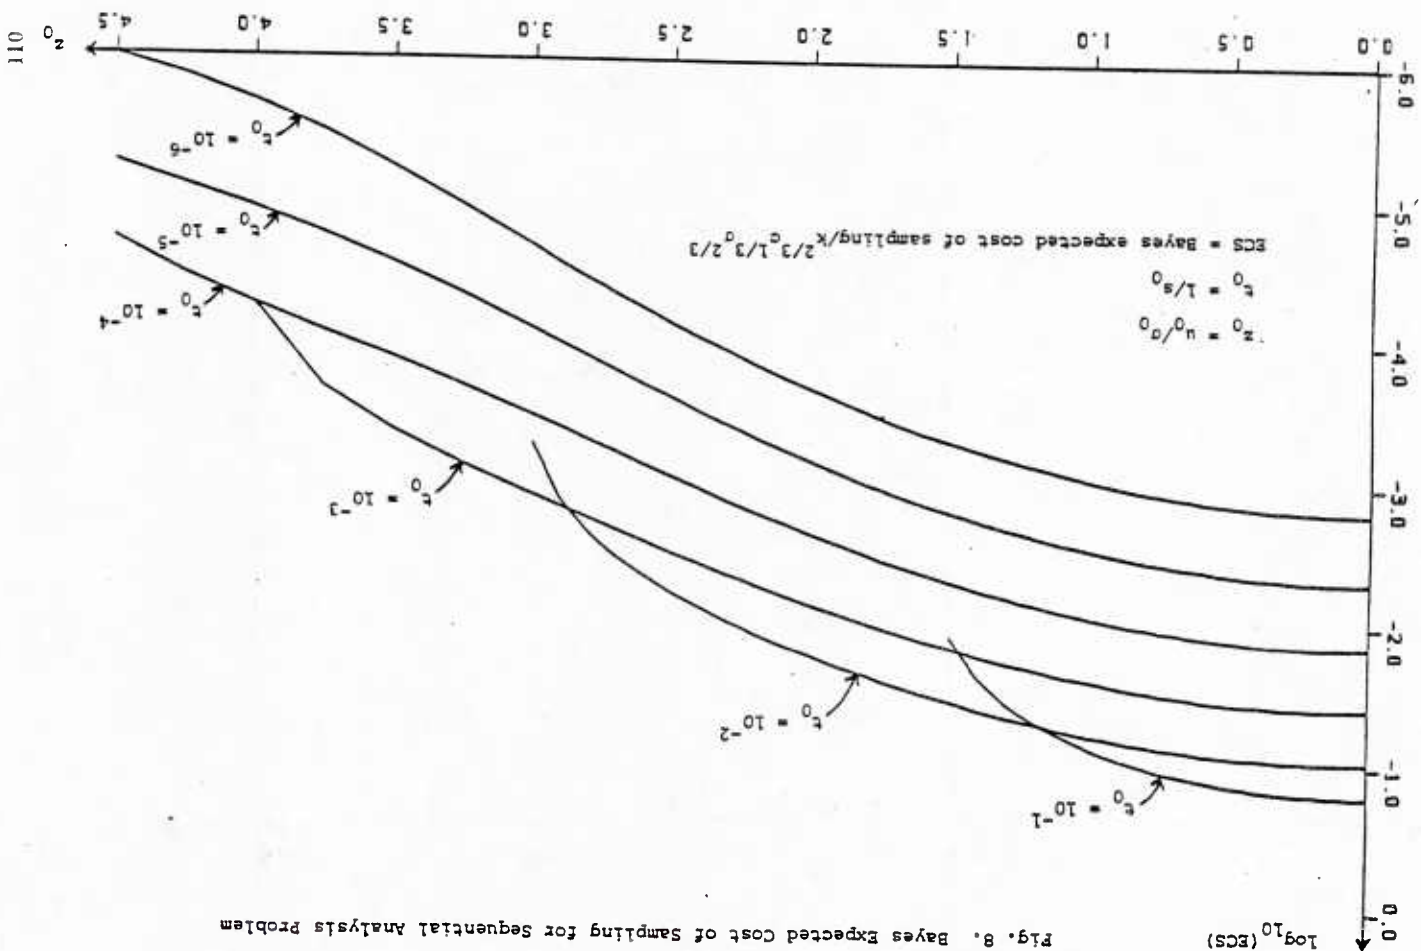


Fig. 9. Bayes Expected Payoffs for One-Armed Bandit Problem

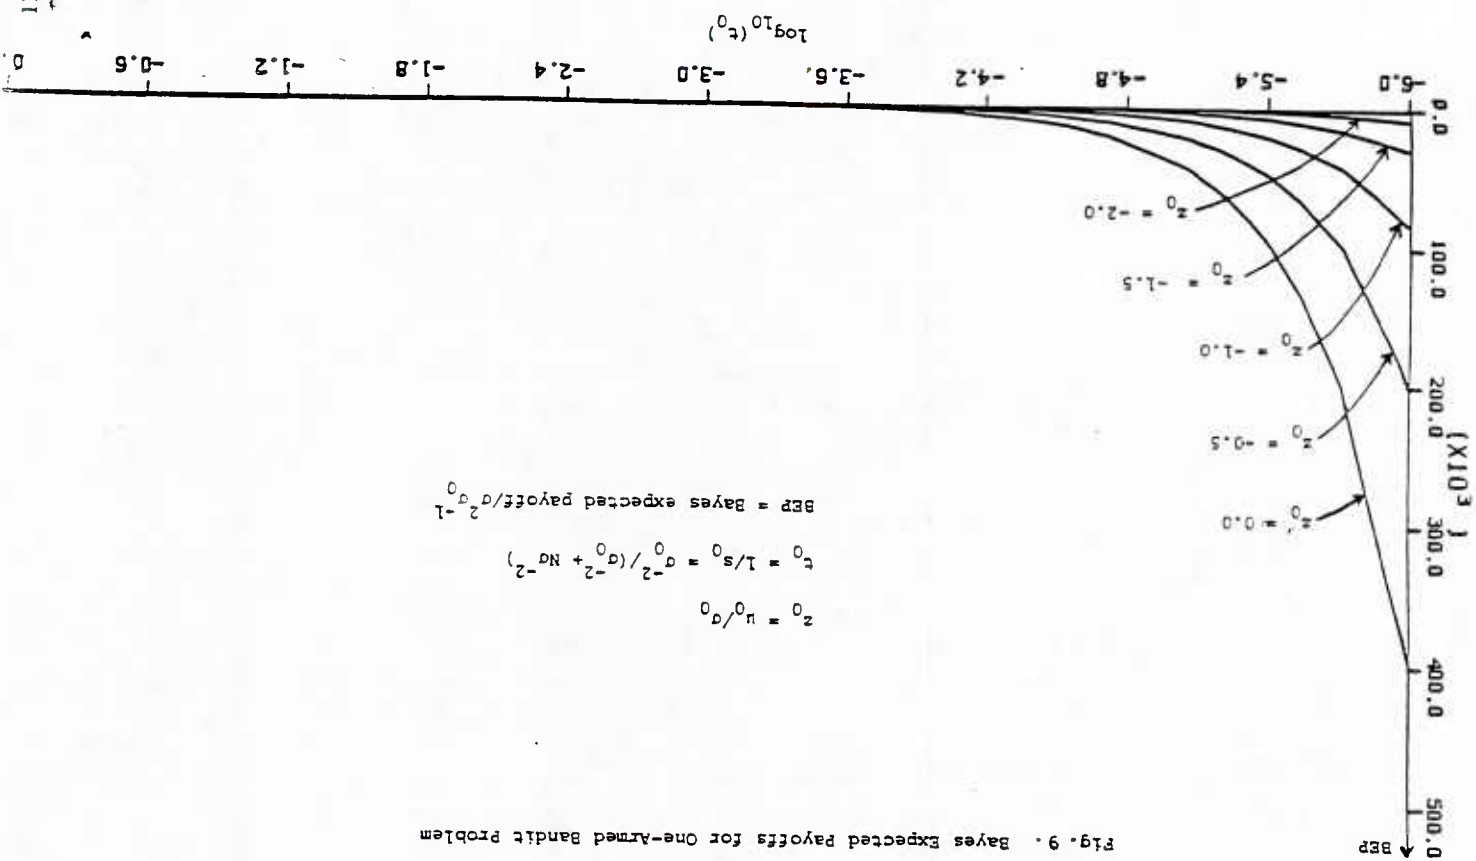
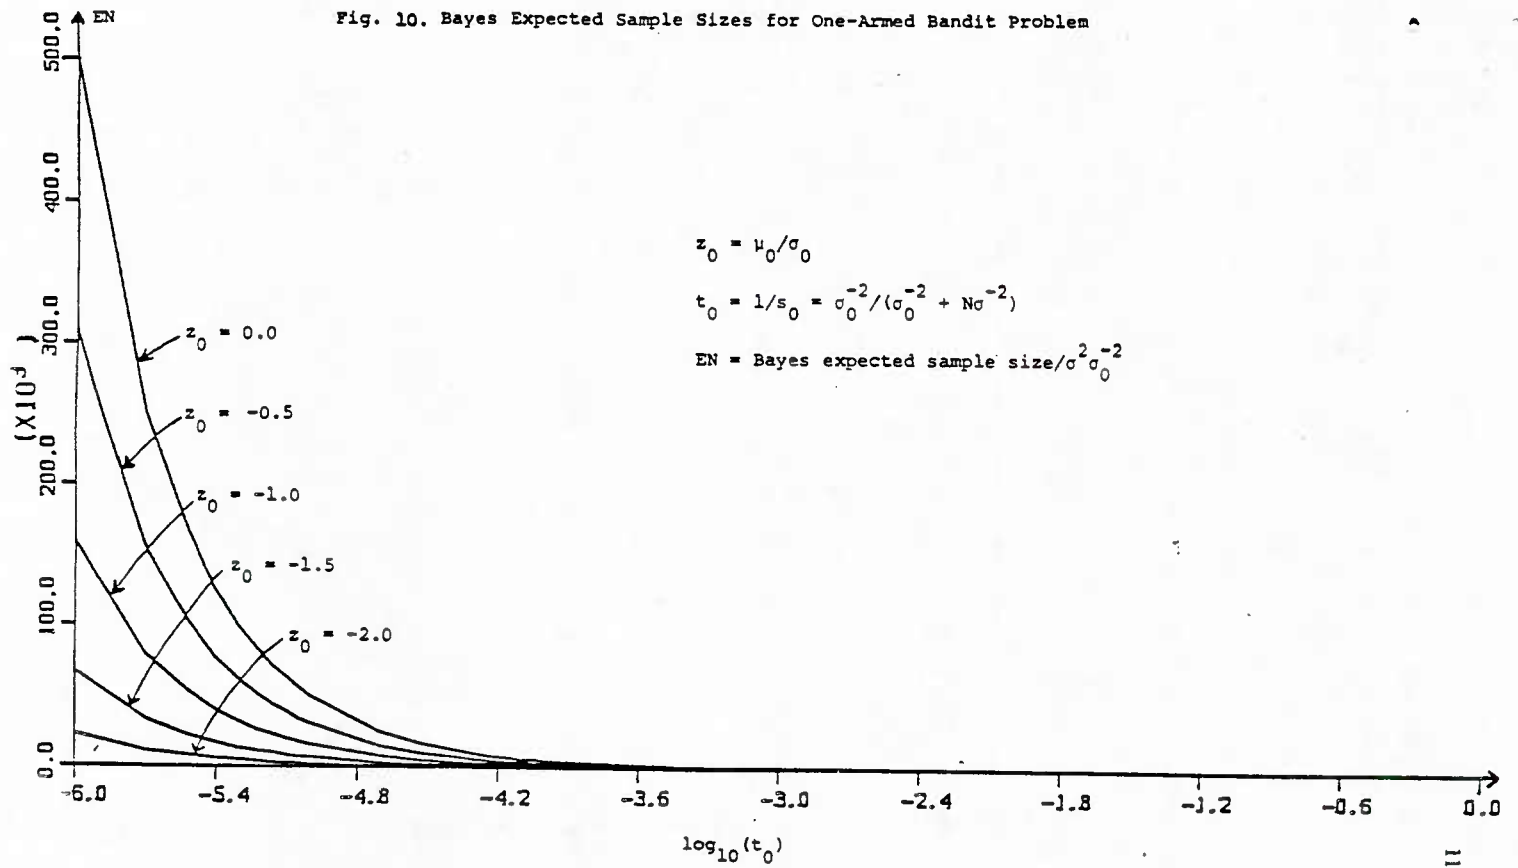
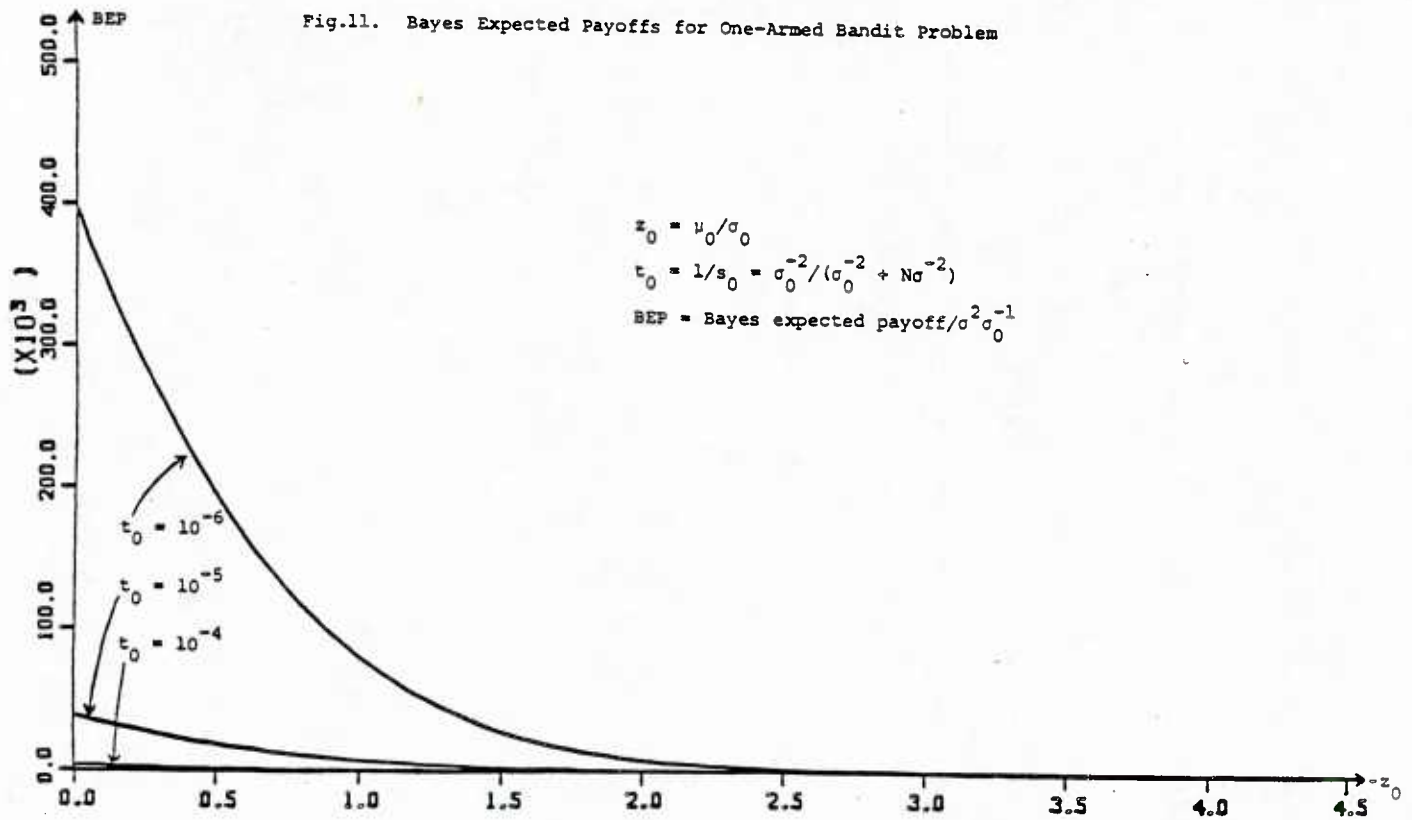


Fig. 10. Bayes Expected Sample Sizes for One-Armed Bandit Problem



112

Fig. 11. Bayes Expected Payoffs for One-Armed Bandit Problem



113

Fig.12. Bayes Expected Sample Sizes for One-Armed Bandit Problem

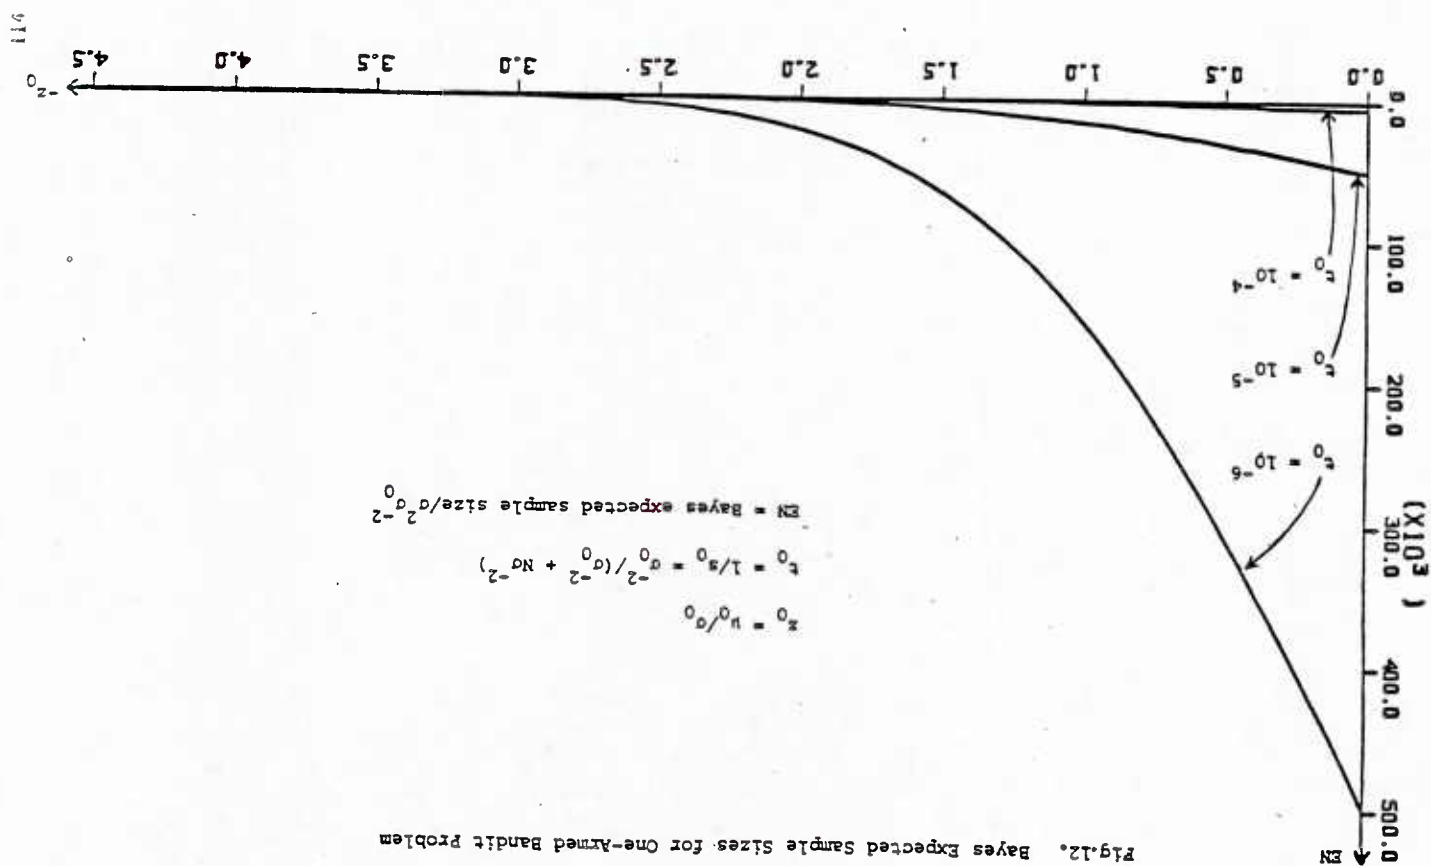


Fig.13. Bayes Expected Payoffs for One-Armed Bandit Problem

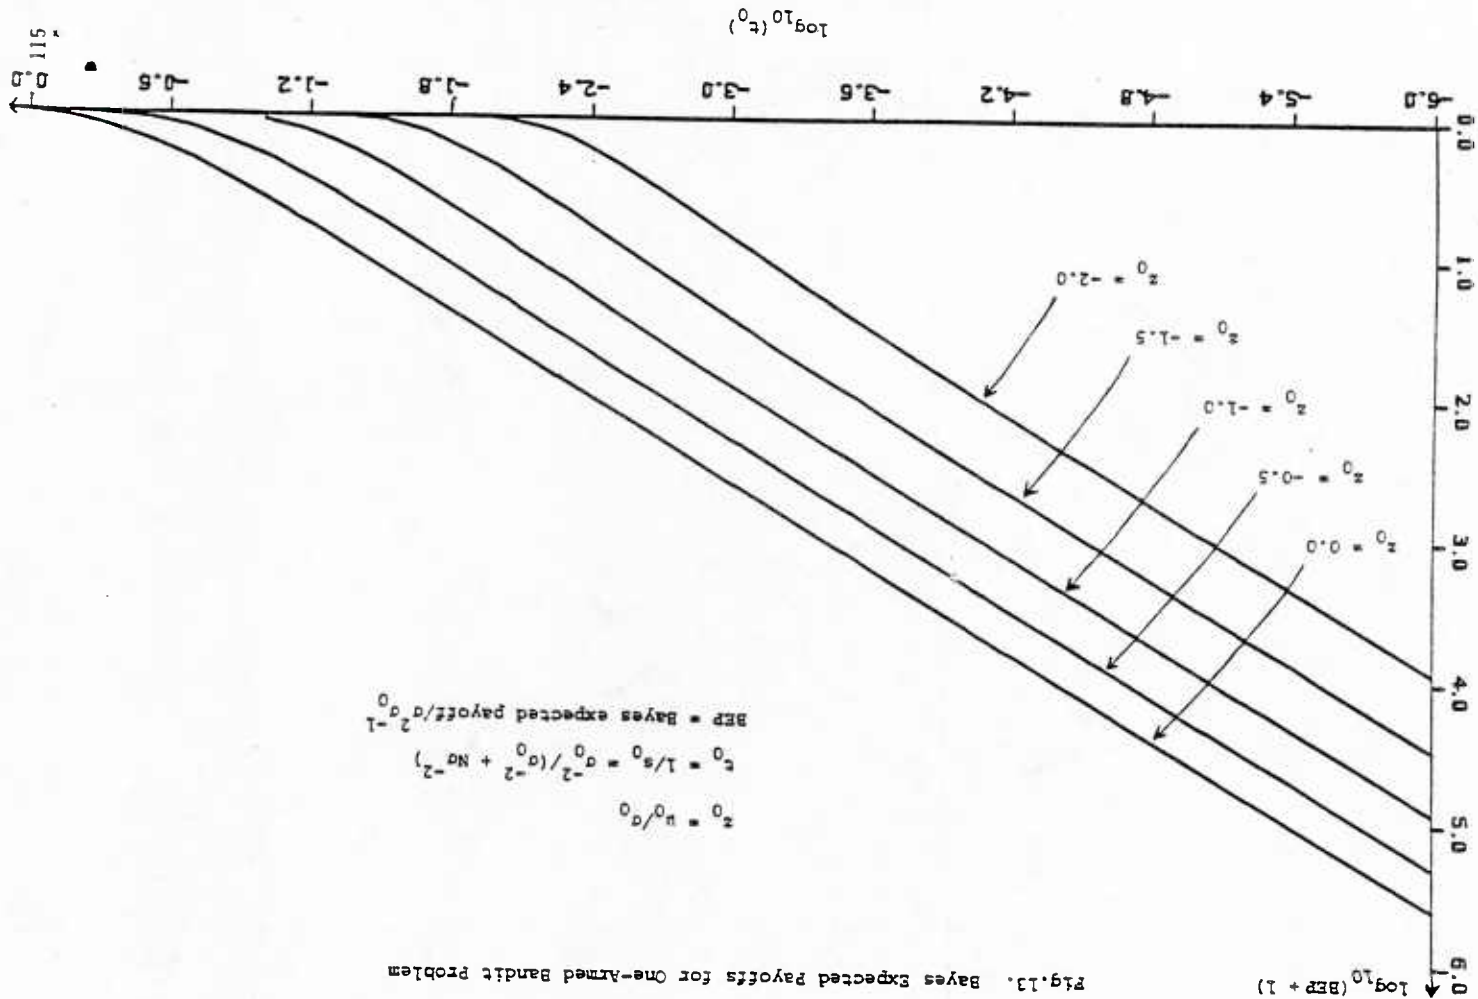


Fig. 14. Bayes Expected Sample Sizes for One-Armed Bandit Problem

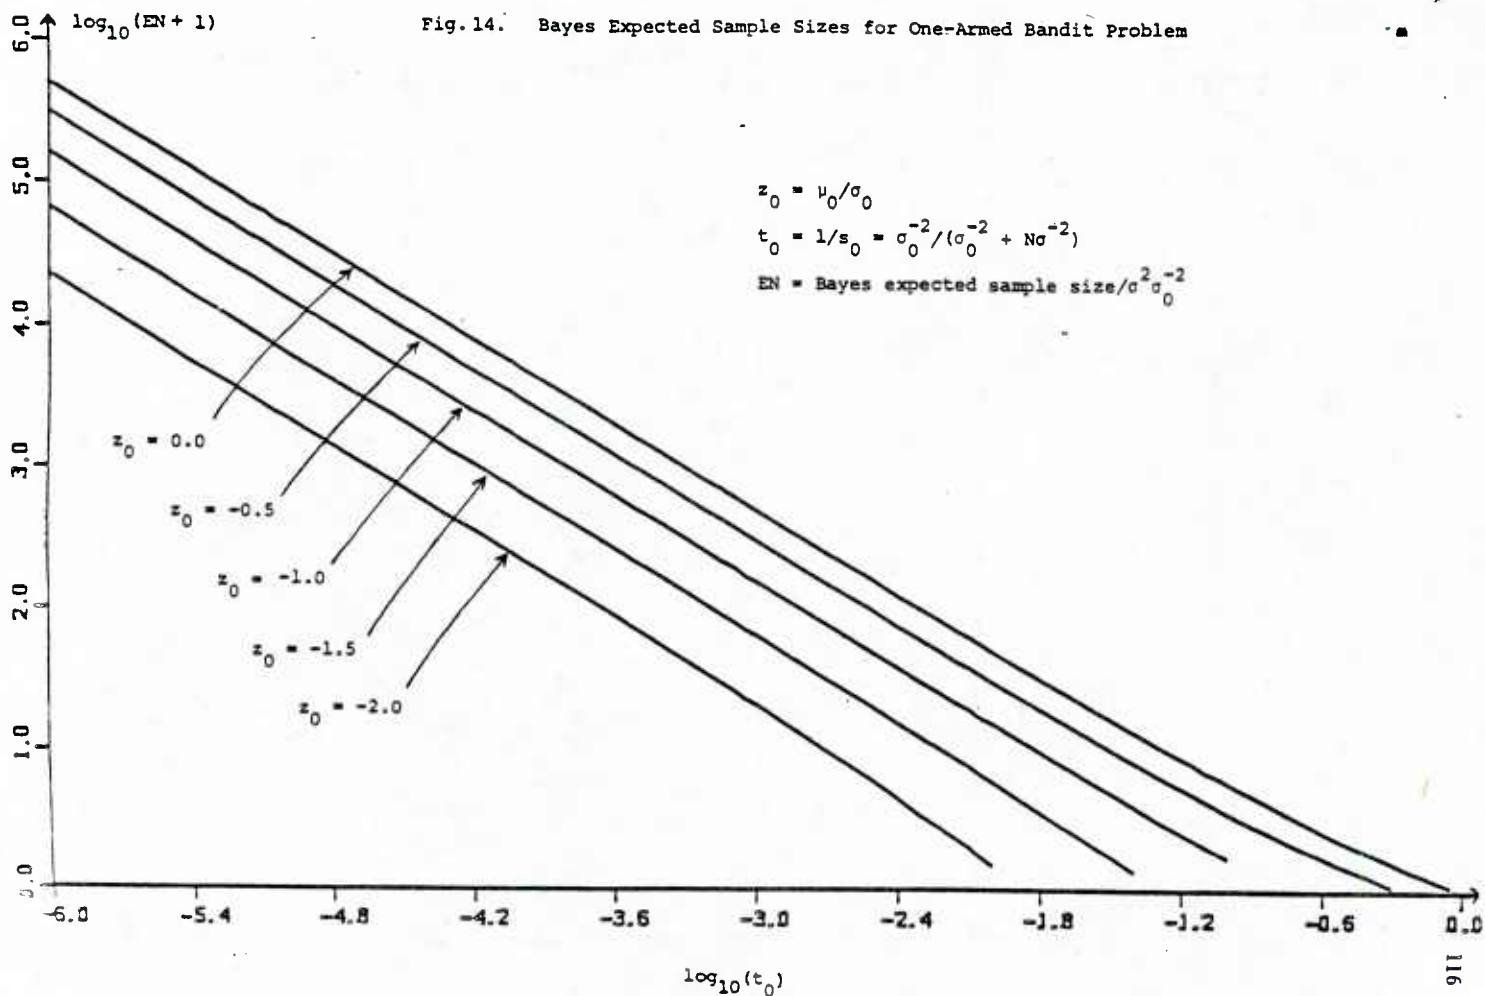


Fig. 15. Bayes Expected Payoffs for One-Armed Bandit Problem

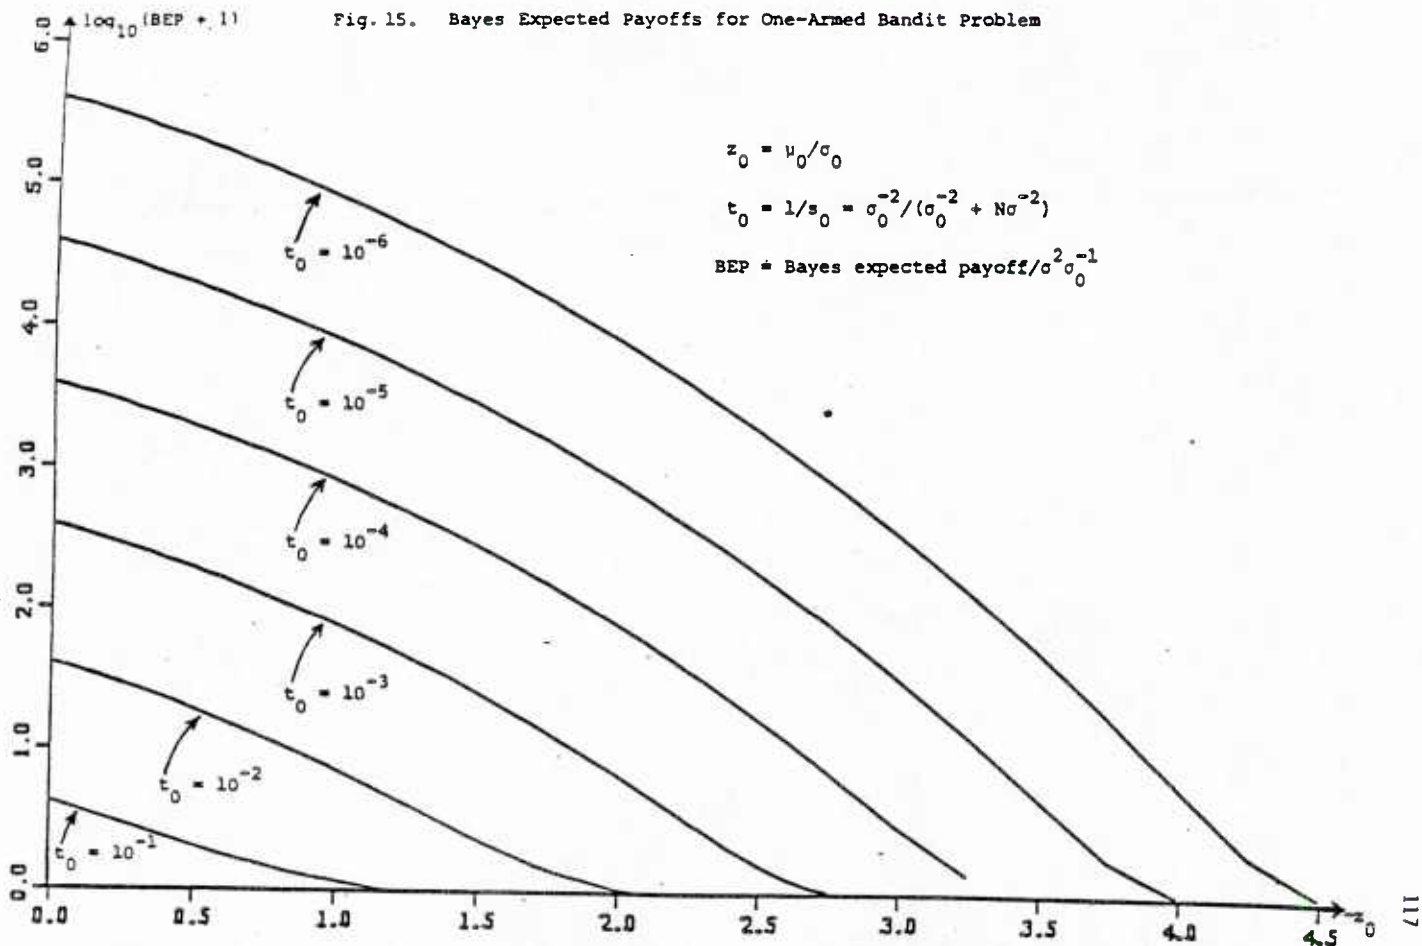


Fig. 16. Bayes Expected Sample Sizes for One-Armed Bandit Problem

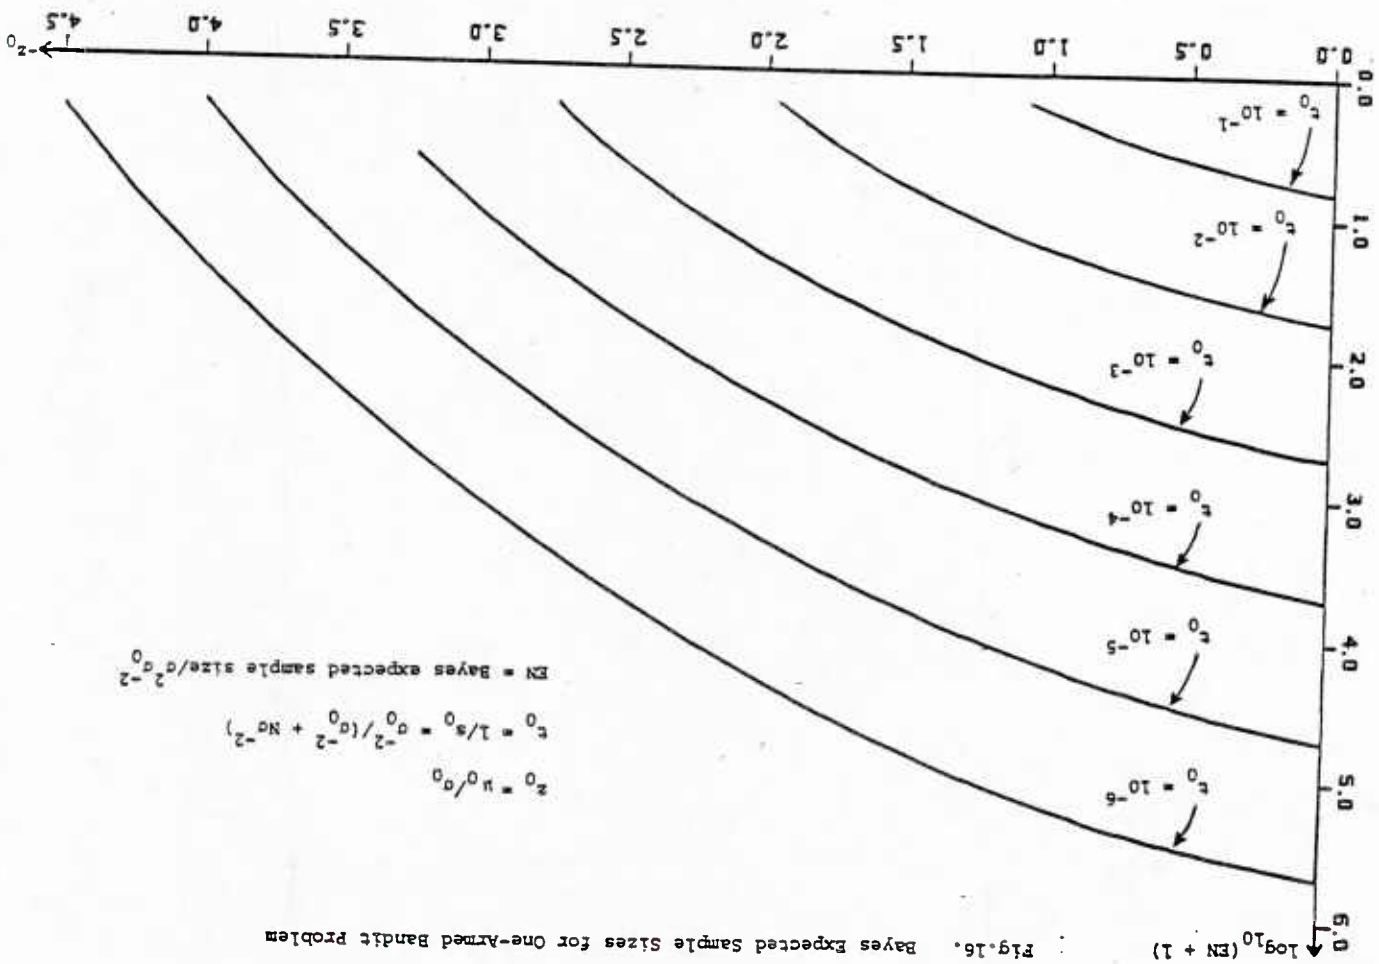
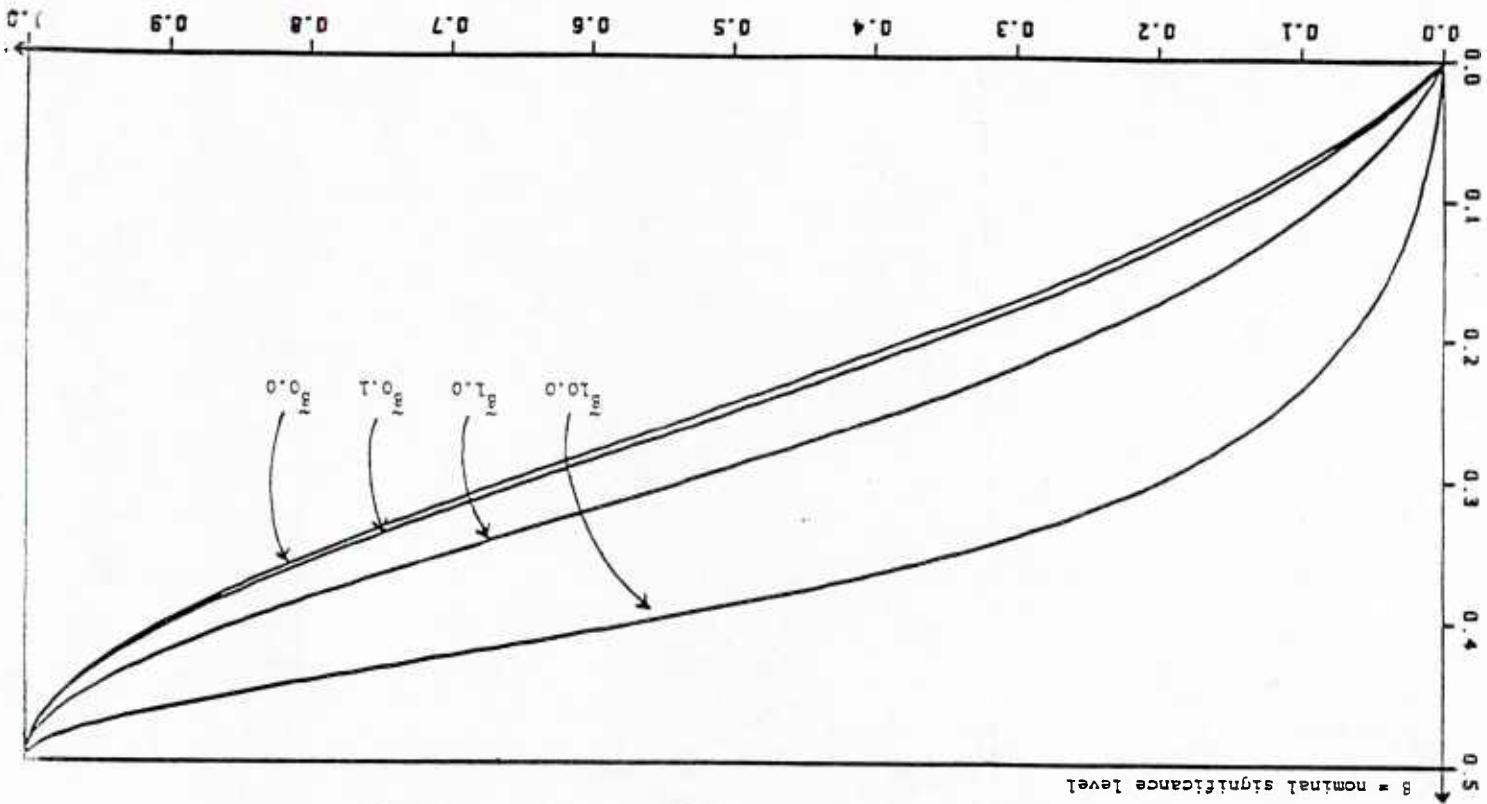


Fig. 17. Optimal Stopping Rules for Anscombe's Problem with Ethical Cost



$c = \text{currently available proportion of total potential information}$

Fig.18. Bayes Risks at $z_0 = 0$ for Anscombe's Problem with Ethical Cost

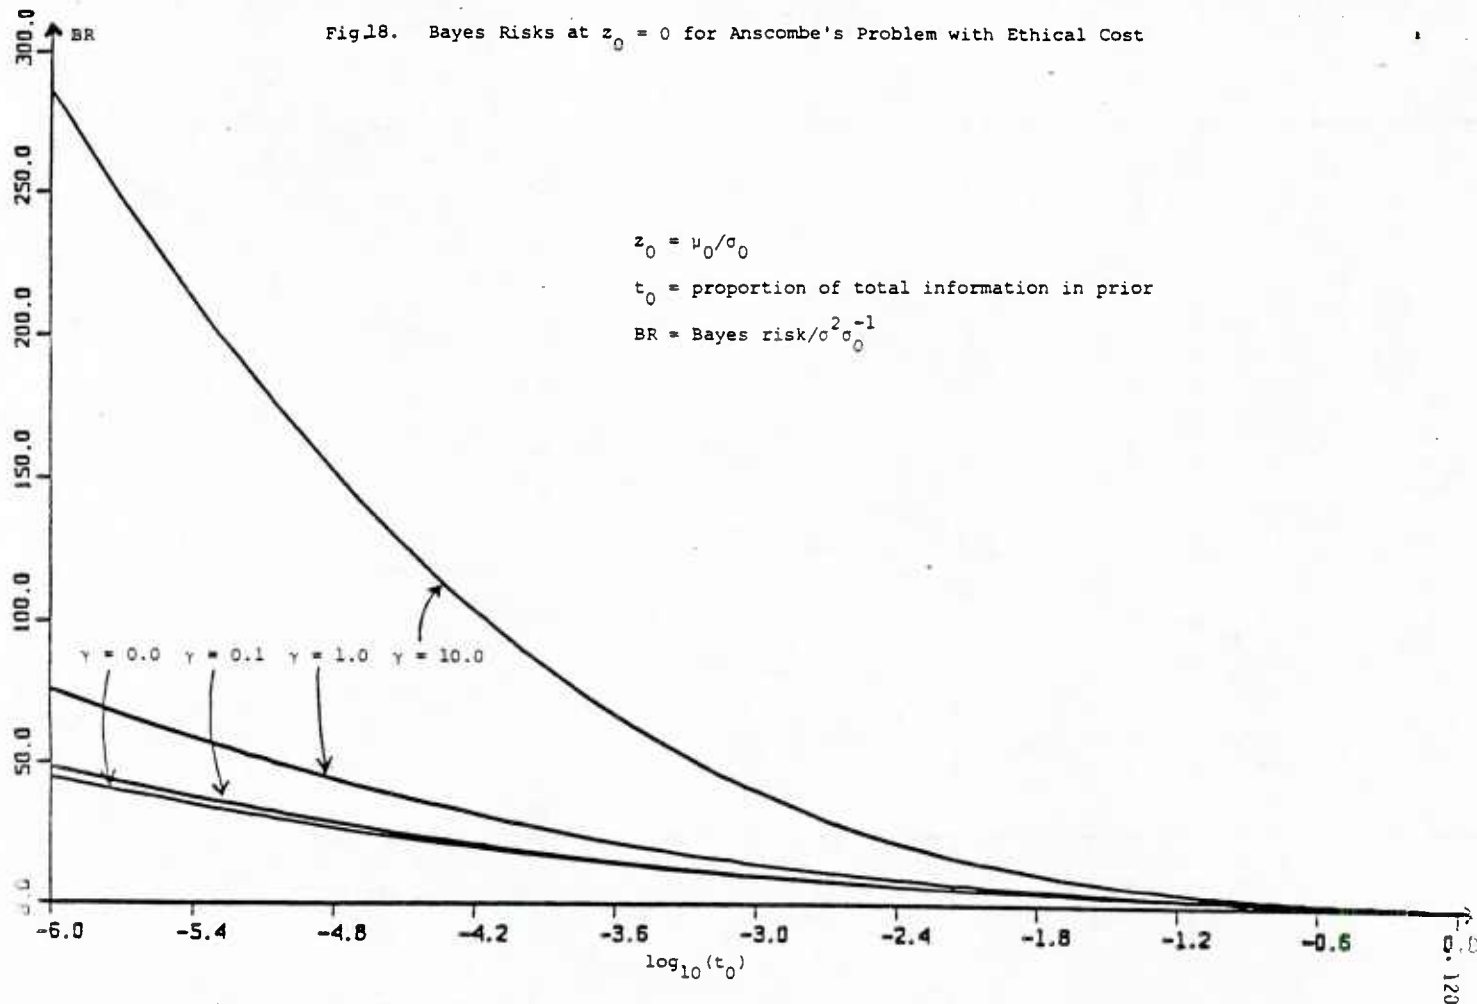
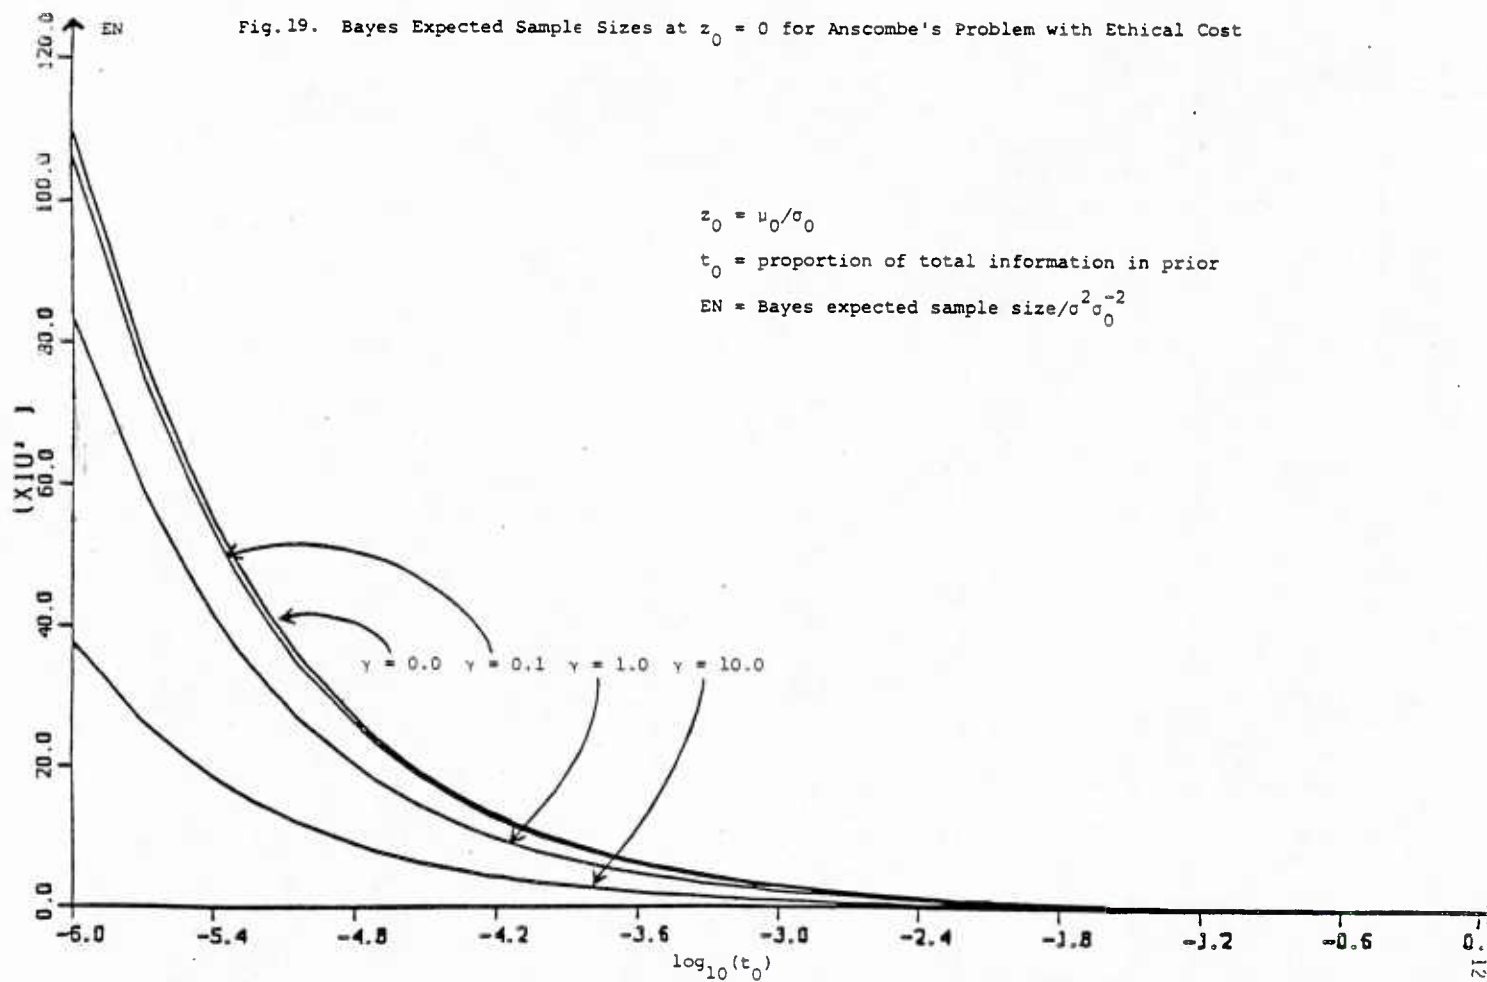
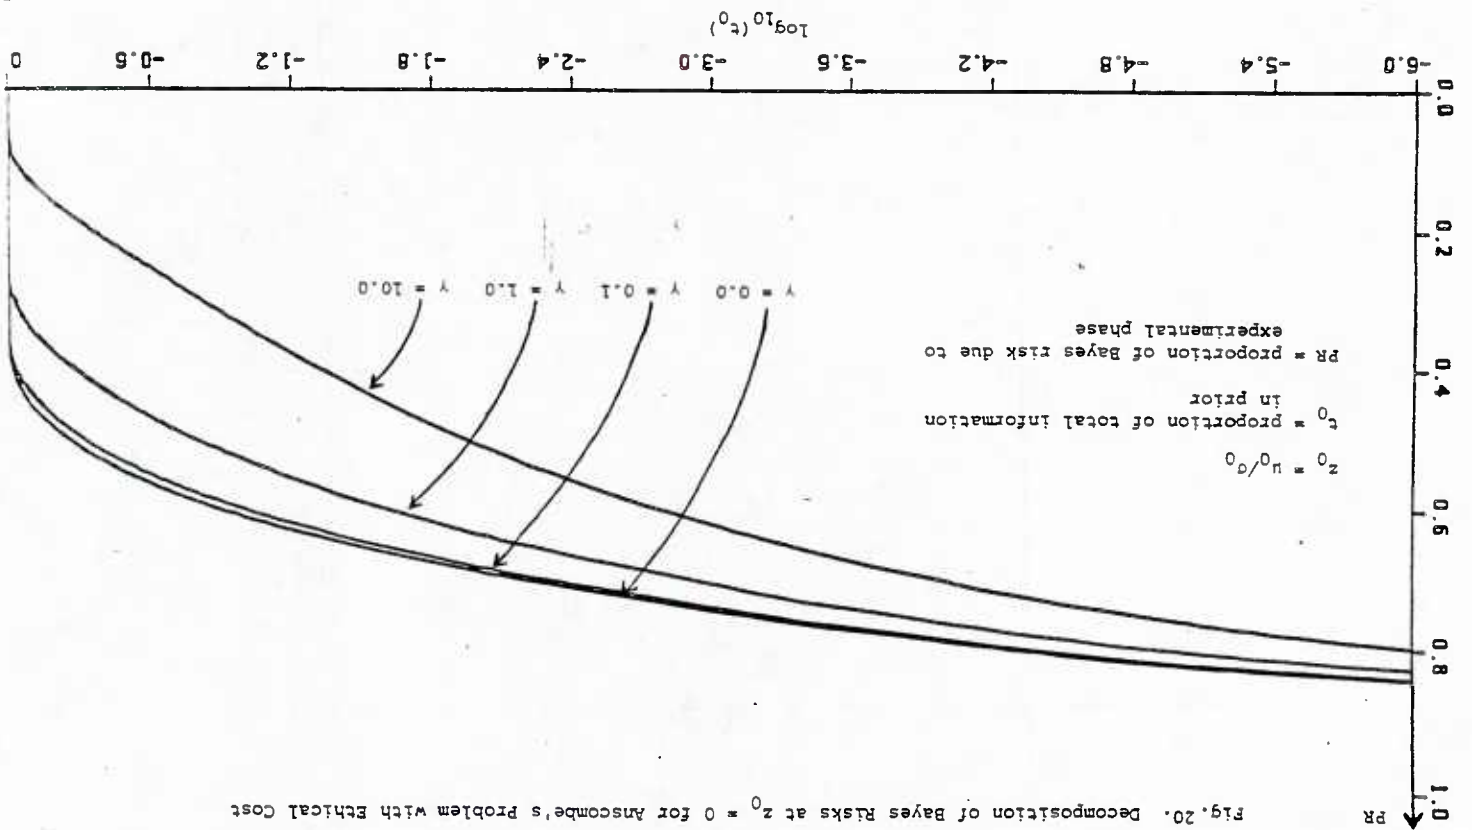
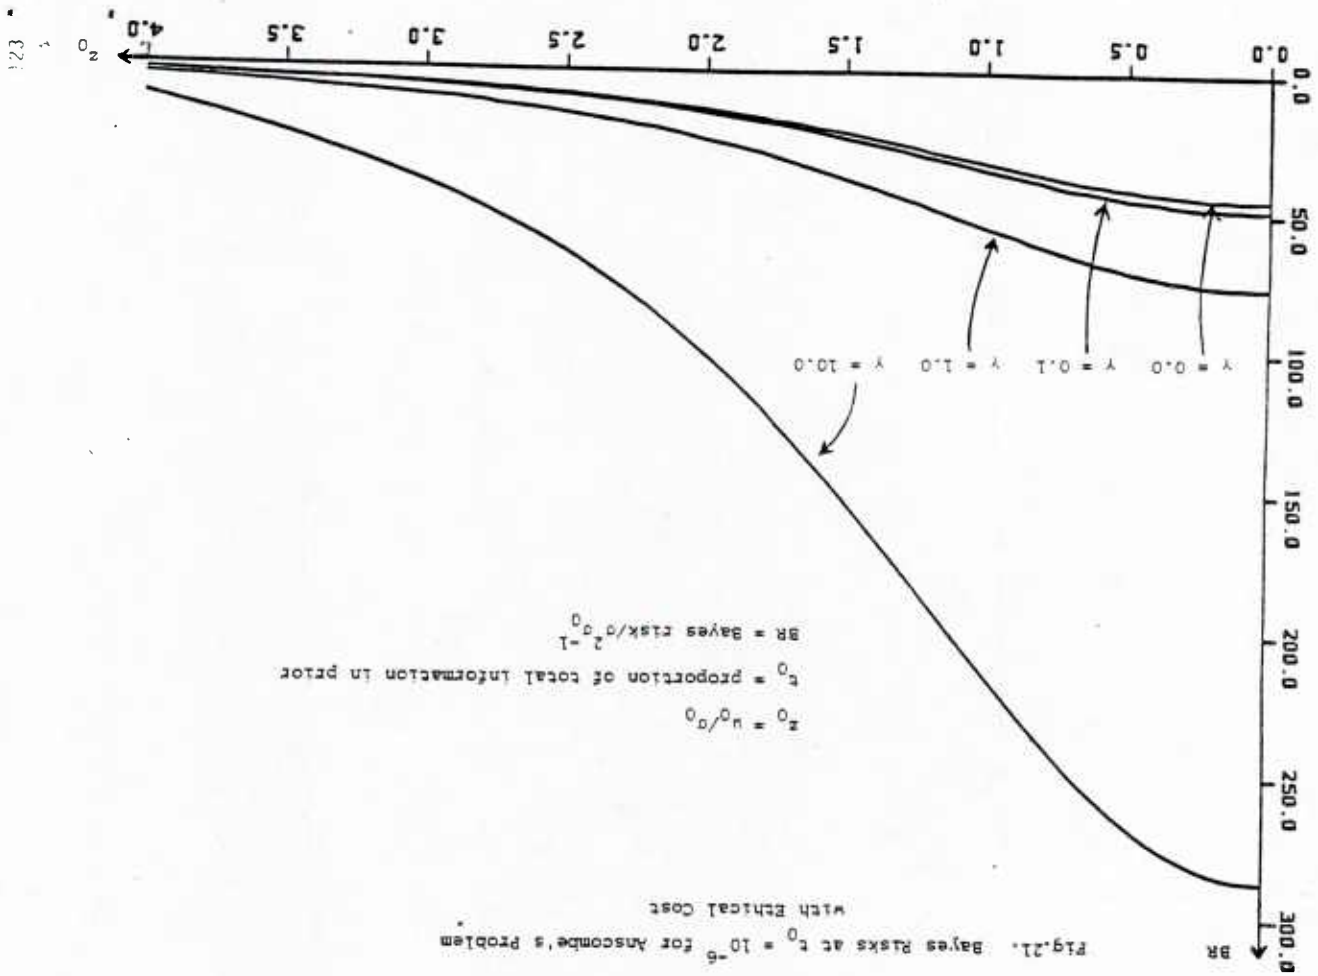
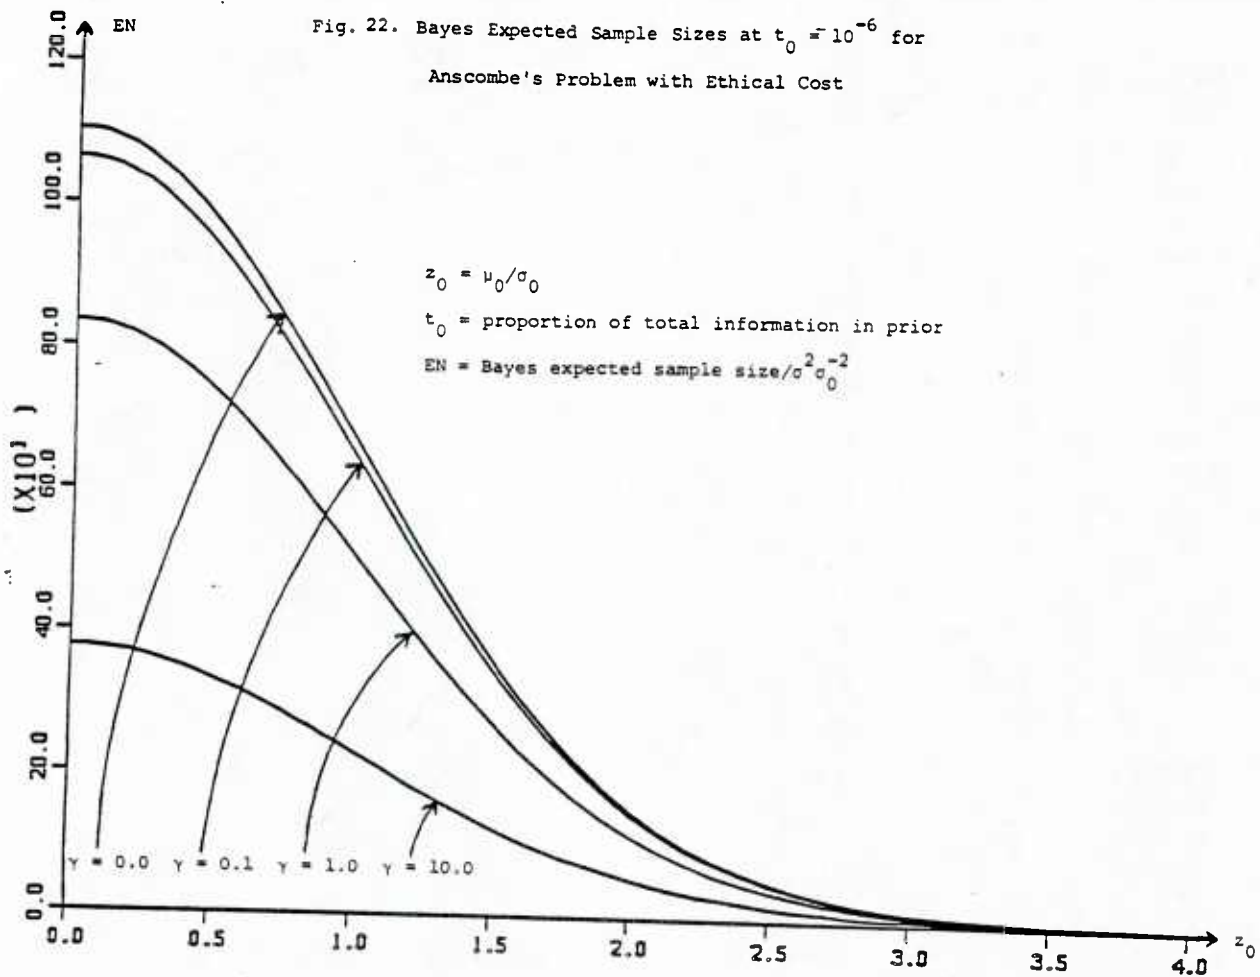


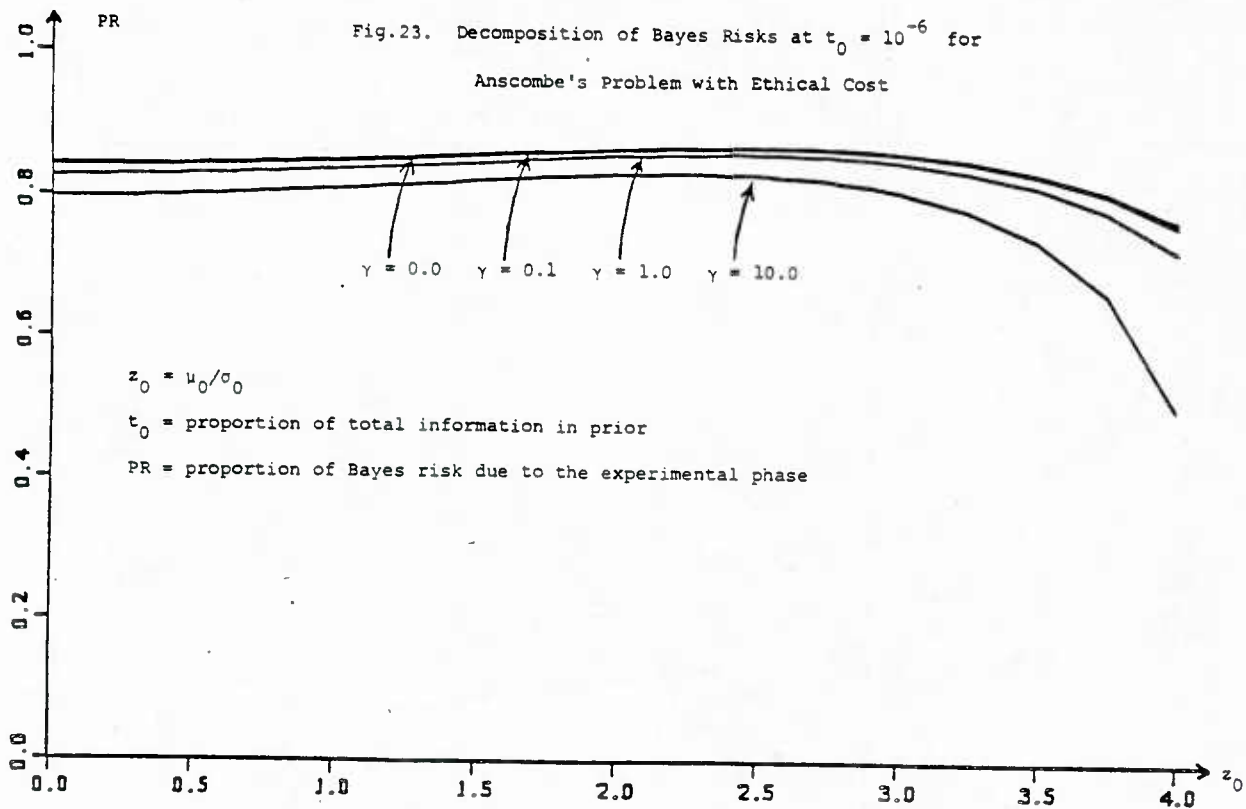
Fig.19. Bayes Expected Sample Sizes at $z_0 = 0$ for Anscombe's Problem with Ethical Cost







124



125

U211616

