AD-766 809

APPROACHES IN SEQUENTIAL DESIGN OF EXPERIMENTS

Herman Chernoff

Stanford University

Prepared for: Office of Naval Research

11 May 1973

DISTRIBUTED BY:

National Technical Information Service

where was

U. S. DEPARTMENT OF COMMERCE 5285 Port Royal Road, Springfield Va. 22151

فالمحافظ فالمحاصر فالمستعد كالمرحلة والمحافظ والمحافين والمحاصر والمحاد والمحاد والمحافظ

44

APPROACHES IN SEQUENTIAL DESIGN OF EXPERIMENTS

SPL O

00

0

₹~s

語を見ていたので、「「

BY HERMAN CHERNOFF

TECHNICAL REPORT NO. 77 MAY 11, 1973



PREPARED UNDER CONTRACT NOO014-67-A-0112-0051 (NR-042-993) FOR THE OFFICE OF NAVAL RESEARCH

DATIONAL PERMITAL



Security CLASSIFICATION		
DOCUM	ENT CONTROL DATA - RAD	an die emerif meest te obsettied.
ORIGINATIN & ACTIVITY (Corporate author)	20. AL	PORT SECURITY & LASSIFICATION
DEPT. OF STATISTICS		
STANFORD UNIVERSITY, CALIF.	21. 00	BUP
REPORT TITLE		
APPROACHES IN SEQUENTI	AL DESIGN OF EXPERIMENT	
DESCRIPTIVE NOTES (Type of report and inclusive of	¢*****)	
AUTHOR(S) (Leet name, Set name, Millet)		
	·	
CHERNOFF, Herman		
REPORT DATE	74. TOTAL NO. OF PASES	76. NO. OF REFS
May 11, 1973	<u>112-44</u>	84
CONTRACT OR GRANT NO. N00074-67-4-0112-0051	SA ONEMATOR'S REPORT N	unser(2)
PROJECT NO.	#77	
NR-042-993		
	PA OTHER REPORT HO(2) (A) Bio mpan)	ny other numbers that may be seeign
	· · · ·	
AVAILABILITY/LIMITATION NOTICES		
Unlimited		
SUPPL ENENTINY MOTES		
	Office of Naval R	esearch .
	Arlington, VA.	
ABSTRACY		
Sequenting design of some		
characterized by the fact that	and deta accumulate the	s of inference
choose whether or not to experi	ment further. If he deci	des to orneriment
further, he can decide which ex	periment to carry out nex	t and if he decides
to stop experimentation, he mus	t decide what terminal de	cision to make.
	non harris diama A	
The literature contains t	wo broad types of general	approach and
The literature contains t several major classes of applic	ations. One general appr	approach and oach is that of
The literature contains t several major classes of applic stochastic approximation. Three	wo broad types of general ations. One general appr e variations are the Robb	approach and oach is that of ins-Monro methods,
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface mether other general approach consists	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down	approach and oach is that of ins-Monro methods, methods. The
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface methother general approach consists other general approach consists	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as	approach and oach is that of ins-Monro methods, methods. The ymptotically
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface metho other general approach consists optimal designs, generally in a Special classes of applic	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as Bayesian decision theore ations include survey som	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context.
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface met other general approach consists optimal designs, generally in a Special classes of applic continuous sampling inspection.	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as Bayesian decision theore ations include survey sam selecting the largest of	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context. pling, multilevel k populations
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface met other general approach consists optimal designs, generally in a Special classes of applica continuous sampling inspection, which includes clinical trials a	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as Bayesian decision theore ations include survey sam selecting the largest of and two-armed bandit-type	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context. pling, multilevel k populations, problems. screening
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface met other general approach consists optimal designs, generally in a Special classes of application continuous sampling inspection, which includes clinical trials a experiments, group testing, and	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as Bayesian decision theore ations include survey sam selecting the largest of and two-armed bandit-type search problems.	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context. pling, multilevel k populations, problems, screening
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface met other general approach consists optimal designs, generally in a Special classes of applic continuous sampling inspection, which includes clinical trials a experiments, group testing, and	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as Bayesian decision theore ations include survey sam selecting the largest of and two-armed bandit-type search problems.	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context. pling, multilevel k populations, problems, screening
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface met other general approach consists optimal designs, generally in a Special classes of application continuous sampling inspection, which includes clinical trials a experiments, group testing, and	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as Bayesian decision theore ations include survey sam selecting the largest of and two-armed bandit-type search problems.	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context. pling, multilevel k populations, problems, screening
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface met other general approach consists optimal designs, generally in a Special classes of applic continuous sampling inspection, which includes clinical trials a experiments, group testing, and	wo broad types of general ations. One general appr e variations are the Robb hods and the up-and-down of finding optimal or as Bayesian decision theore ations include survey sam selecting the largest of and two-armed bandit-type search problems.	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context. pling, multilevel k populations, problems, screening
The literature contains t several major classes of applic stochastic approximation. Thre Box-Wilson response surface metho other general approach consists optimal designs, generally in a Special classes of application continuous sampling inspection, which includes clinical trials a experiments, group testing, and	o o o o o o o o o o o o o o	approach and oach is that of ins-Monro methods, methods. The ymptotically tic context. pling, multilevel k populations, problems, screening

F.

í.

5

North Sector

the second

Ē

ないできたとうないないのとうないできた

道法で

UNCLASSIFIED							
4		LIN	KA	LIN	K D	LIN	KC
KET WORDS		HOLE		ROLE	WT	ROLE	
sequential analysis design of experiments two-armed bandit stochastic approximation survey sampling sampling inspection clinical trials screening experiment group testing search problems							
			Ĺ				
IN I. ORIGINATING ACTIVITY: Fotor the same and address	TRUCTIONS		cleasifi	cation w		مدمده السماد	
 A. REPORT SECURITY CLASSIFICATION: Enter the ovall security classification of the report. Indicate whether "Restricted Date" is included. Marking is to be in accord ance with appropriate ascurity regulations. 24. GROUP: Automatic downgrading is specified in DoD I rective \$200.10 and Armed Forces Industrial Manual. Enter the group number. Also, when applicable, show that option markings have been used for Group 3 and Group 4 as authoused. 3. REPORT TITLE: Enter the complete report title in all capital letters. Titles in all cases should be unclassified. If a meaningful title cannot be selected without classification, show title classification in all capitals in parenthesis immediately following the title. 4. DESCRIPTIVE NOTES. If appropriate, enter the type of report, e.g., interim, progress, summery, annual, or final. Give the inclusive dates when a specific reporting period is covered. 5. AUTHOR(S): Enter the name(s) of suthor(s) as shown of in the report. Enter the date of the report as day, month, year, or month, year. If more than one date appears on the report, use date of publication. 7. TOTAL NUMBER OF PAGES: The total page count should follow normal pagination procedures, i.e., enter the mark of pages containing information. 7. NUMBER OF REFERENCES Enter the total number of references cited in the report. 8. ACT OR GRANT NUMBER: If appropriate, enter the mumber of pages containing information. 7. NUMBER OF REFERENCES Enter the total number of seferences cited in the report. 8. CONTRACT OR GRANT NUMBER: If appropriate, enter the applicable number of the contract or grant under which the report uses written. 8. S. & G. PROJECT NUMBER: Enter the epipopriate entities and project number, etc. 9. ONTRACT OR SREPORT NUMBER: Enter the officient report was written. 9. Chick Add. PROJECT NUMBER: Enter the epipopriate entities appropriate and project number, by a project number, etc. 9. ONT	(1) " (1) " (1) " (2) " (3) " (3) " (4) " (4) " (5) " (5) " (6) " (6) " (7) " (7	 imposed by security classification, using standard stateuch as: "Qualified requesters may obtain copies of the report from DDC." "Foreign aunouncement and dissemination of report by DDC is not authorized." "U. S. Government agencies may obtain copies of this report directly from DDC. Other qualified users shall request through "U. S. military agencies may obtain copies of report directly from DDC. Other qualified user whall request through "U. S. military agencies may obtain copies of report directly from DDC. Other qualified user whall request through "All distribution of this report is controlled ifled DDC users shall request through "All distribution of this report is controlled ifled DDC users shall request through "If the report mas been furnished to the Office of Te Services, Department of Commerce, for sale to the publicate this fact and enter the price, if known. SUPPLEMENTARY NOTES: Use for additional e tory notes. SPONSORING MILITARY ACTIVITY: Enter the mithe departmental project office or laboratory sponsorin ing for) the research and development. Include address it may also repear elsewhere in the body of the technal summery of the document indicative of the report, avera it may also repear elsewhere in the body of the technal point. If additional spece is required, a continuation st be attached. It is highly desirable that the abstract of classifie to the report, formation in the partmaph, represented as (TS), (S), (C). There is no limitation on the length of the abstract of classifiertion of formation in the partmaph, represented as (TS), (S), (C). There is no limitation on the length of the abstract of classifiertion of the military security classification is required. 				les of this stion of this in copies qualified opies of (fied user) troiled. (troiled user) troiled of Tec the public tional exploit the public tional exploit the public ter the me pensoring ter the me pensoring ter the me ter the	emente is ihis s of DDC
 a sny other report numbers (either by the originator or by the sponsor), also enter this number(s). AVAILABILITY/LIMITATION NOTICES: Enter any li itations on further dissemination of the report, other than t 	B-	4.4 a r Rithlig	nt At 1117	. ., 14196 ,	ana Weij	pite 14 Op	uosei.
U 15287. 1473 (BACK)				1h	1 C] 9 885.	fled	

Chilles Willing Harry

APPROACHES IN SEQUENTIAL DESIGN OF EXPERIMENTS

Ъy

Herman Chernoff

1. Introduction

aterationalizationalizationalizationalizationalizationalizationalizationalizationalizationalizationalizationali

Sequential design of experiments refers to problems of inference characterized by the fact that as data accumulate, the experimenter can choose whether or not to experiment further. If he decides to experiment further, he can decide which experiment to carry out next, and if he decides to stop experimentation, he must decide what terminal decision to make.

In principle, ordinary sequential analysis, where there is no choice of experiment but where one must simply decide whether or not to repeat a specified experiment, is a special if slightly degenerate case of sequential design. The same can be said for double sampling, where the experimental choice reduces to selecting the size of the first sample and, given the outcome, the size of the second sample. Indeed double sampling may be regarded as the origin of sequential analysis and hence of sequential design of experiments. With the exception of a few references of special interest, we shall avoid the discussion of these degenerate cases, and we shall concentrate mainly on problem areas and theories where there is a choice of experimentation after each observation. We shall do this in our search for general insights

even though double sampling is probably of more practical interest than the remainder of sequential experimentation.

In recognition of the importance of a theory of sequential design, Robbins [R1] proposed the two-armed bandit problem as a prototype problem of possibly fundamental importance. Two variations of the simplest version are the following. In both there exist probabilities p, and p, corresponding to the probability of success with two arms. Selecting an observation from arm i leads to a success with probability p_1 , 1=1,2. The two alternative hypotheses are H_1 : $(p_1,p_2) =$ (p_{10}, p_{20}) and $H_2((p_1, p_2) = (p_{20}, p_{10})$ where p_{10} and p_{20} are distinct specified probabilities. Thus one knows both probabilities, but one doesn't know which corresponds to which arm. Each hypothesis is assumed to be equally likely. After each observation the experimenter may select the arm to be used nex* until N observations have been taken. In one variation the object is to make the choices so as to maximize the probability of deciding which hypothesis is true after the N^{th} observation. In the other variation the object is to maximize the expected total number of successes in N trials. The second version is the one usually referred to as the two-armed bandit problem and seems to confront the major issue more directly. How does one compromise between the anticipated cost and the value of the information? For in that problem the choice of the arm less likely (according to the posterior probability) to have the larger probability would constitute a sacrifice of immediate gain in the hope of information which could

5

lead to ultimate profit.

A Contract of States and Stat

As a prototype this problem was attacked vigorously, but the results implied that this problem failed as a useful prototype, at least in its immediate interpretation. The main result, which was surprisingly difficult to establish [F5], always calls for the use of the arm most likely to have the higher probability and hence does not yield a useful comparison of cost with information. The variations of this problem where this result does not apply did not seem to have any clearly generalizable interpretation. These variations involve imposing different prior distributions on $(\mathbf{p}_1,\mathbf{p}_2)$. Note that the original problem corresponds to a two-point prior distribution with probability allocated to the two points $(\mathbf{p}_{10},\mathbf{p}_{20})$ and $(\mathbf{p}_{20},\mathbf{p}_{10})$.

A problem which is curtently of considerable interest in pattern recognition problems is fundamentally related to sequential design of experiments, although structly speaking there may be no novel experimentation. Here the question becomes one of which functions of the already collected data should be studied. For example, one may have samples of cardiograms for normal people and for people having had heart attacks. One may wish to develop a method of clausifying a given cardiogram into one of these two categories. What aspect of the cardiogram should one study? One may select first some simple function of the data (called a meature in the pattern recognition literature). To the extent that the use of this feature can only

3

ALTERNAL AND AND CARDON

do part of the job of classifying, one may attempt to look for additional features sequentially. Although the data are completely available, the process of selecting new features is equivalent to the carrying out of additional experiments, as is practiced by the physician who diagnoses an illness by a succession of "tests". Both of these cases have one aspect in common which separates them from the main body of the literature on sequential design of experiments. In both of these the result of the nth "experiment" is statistically dependent on the previous results. However, most of the literature in sequential design of experiments concentrates on problems where once the nth experiment is selected, its outcome is independent of the past. Indeed an experiment can be repeated (independently) several times in such problems, whereas a repetition is useless (except to correct for experimental error) in the cardiogram and diagnosis-type problems. and the second state of the second state of the second second second second second second second second second

۵۰ به شدید ور این و میکرد. وی این این و میکرد و این و میکرد. وی این میکرد و این میکرد و میکرد و میکرد و میکرد و میکرد و میکرد و میکرد. و میکرد و میکرد. و میکرد و میکرد و میکرد و میکرد و میکرد و میکرد و میکرد.

The literature in sequential design contains two broad types of general approach and several major classes of applications. One type of general approach is that of stochastic approximation. Three variations are the Robbins-Monro methods, the Box-Wilson response surface methods, and the up-and-down methods. These variations apply to the estimation of characteristics of a regression function and use the data to determine the next level of the independent variable at which to measure the dependent variable. Typically no attention is paid to a stopping rule. The other general approach consists of finding optimal or asymptotically optimal designs, generally in a Bayesian

decision theoretic context.

Special classes of applications, about some of which little will be said here, are (1) survey sampling, (2) multilevel continuous sampling inspection, (3) selecting the largest of k populations, (4) screening experiments, (5) group testing, and (6) search problems. While one would expect Monte Carlo sampling to be one of these classes, the literature seems to lack interest in the sequential selection of simulation experiments. There are a few miscellaneous categories such as "forcing experiments to be balanced" and some process control problems which also deserve mention.

This paper consists of two major parts. One is devoted to the more general approaches, the other to the classes of applications.

2. Stochastic Approximation

The Robbins-Monro [R4] method applies to the following problem. Corresponding to a choice x of the "independent variable", one observes the dependent variable Y(x) with non-decreasing expectation M(x) = E[Y(x)]. It is desired to estimate θ , that value of x for which $M(x) = \alpha$ for some specified value α . Starting with an initial guess x_1 , successive choices x_2, x_3, \ldots are made according to

 $\mathbf{x}_{n+1} = \mathbf{x}_n - \mathbf{a}_n \{\mathbf{Y}_n(\mathbf{x}_n) - \alpha\}$

for some specified sequence $\{a_n\}$. The sequence $\{x_n\}$ serves both as the successive estimates of θ and as the experimental levels of

x. Since $Y_n(x_n) - \alpha$ tends to reflect how far x_n is from θ , the above iteration represents a correction for overestimates or underestimates. The $\{a_n\}$ sequence represents the extent of the correction. If the a_n were bounded away from zero, the successive terms would tend to fluctuate by an amount determined in part by the variance of $Y(x_n)$. If the $a_n \to 0$ too rapidly, the corrections might not build up fast enough to correct for an initial error. However, if $a_n \to 0$ at a suitable rate, it is possible to show that $x_n \to x$ with probability one under weak assumptions concerning the distribution of Y(x). There is an extensive literature to this effect which indicates that the method requires little but that $M(x) > \alpha$ for $x > \theta$ and $M(x) < \alpha$ for $x < \theta$.

While very little is required of the sequence $\{a_n\}$, what does seem remarkable is that with a proper choice of $\{a_n\}$ this method, which confuses design level with estimate and which ignores the past except for the last estimate and the number of observations, is asymptotically efficient. Hodges and Lehmann [H6] have shown that if Y(x) has mean $M(x) = \beta x + \delta$ and constant variance σ^2 , and $a_n = c/n$, then $\theta = \beta^{-1}(\alpha \cdot \delta)$ and

$$E(x_{n+1}-\theta)^2 \approx \frac{\sigma^2 c^2}{n(2c\beta-1)}$$
 if $c\beta > 1/2$

It follows that if $c = \rho^{-1}$, this method has asymptotic efficiency one for estimating 9 in the normal linear regression problem where the slope β is known but the y-intercept δ is not known.

Some reflection based on the following facts will help explain this result. If the regression is linear and β is known, the efficiency of the conventional estimate $\hat{\delta}_n$ of $\hat{\varepsilon}$ is independent of the design. Indeed $\hat{\delta}_n = \overline{Y}_n - \beta \overline{x}_n$ has variance $n^{-1}\sigma^2$, and the corresponding estimate of θ , $\hat{\theta}_{n+1} = \beta^{-1}[\alpha - (\overline{Y}_n - \beta \overline{x}_n)]$ has variance $n^{-1}\beta^{-2}\sigma^2$. Moreover, if x_n is selected to be $\hat{\theta}_n$, then $\hat{\theta}_{n+1} = \hat{\theta}_n - \frac{1}{n\beta}[Y_n - \alpha]$.

DEPHINE STATES OF STATES

Finally in the case where β is not known, the asymptotic variance of the conventional estimate of θ is $\beta^{-2}\sigma^2 \{1+s_n^{-2}(\bar{x}_n-\theta)^2\}$ where s_n^2 is $n^{-1} \sum_{i=1}^n (x_i - \bar{x}_n)^2$. Thus the results of the known β case can be approximated as long as $\bar{x}_n - \theta$ is small compared to s_n . In the stochastic approximation case using the sequence $a_n = c/n$, there is no prior knowledge of θ to insure that $c = \beta^{-1}$. However, as data accumulate one would hopefully obtain a satisfactory estimate of β providing the successive x_n are not too close to each other. This proviso was achieved by Venter [V1] and Fabian [F2,3] by the expedient of separating the design and estimation functions of x_n . That is, they use z_n as an estimate of θ and select two levels $z_n + c_n$ and $z_n - c_n$ at which to draw successive observations from which an estimate of M'(θ) is derived as well as an estimate of θ .

These revised versions of the Robbins-Monro method have some of the robustness property of the original method. Furthermore, with regularity conditions under which M(x) is locally linear (and smooth) with slope β at x=0, $\sqrt{n}(X_n-\theta)$ is asymptotically normal with mean 0 and variance $\sigma^2/n\beta^2$ where n is the number of observations.

This is the best one could hope for in the case of normal linear regression.

nya analon ing mangana na kata ang mangana na kata ng mangana ng mangana ng mangang na katang ng mangang ng man

The suggestion for separating the design and estimation functions of x_n was implicit in the earlier generalization of the Robbins-Monro method by Kiefer and Wolfowitz [K2] to the problem of locating the value θ of x at which M(x) achieves a maximum. Just as -sgn[$Y(x_n) - \alpha$] estimates $sgn(\theta - x_n)$ and points in the direction of θ from x_n in the R-M problem, so does $sgn M'(x_n)$ point in the direction of θ from x_n in the K-W method. Here $[M'(z_n)]$ is estimated by $[Y(z_n + c_n) - Y(z_n - c_n)]/2c_n$ and the K-W method uses

$$z_{n+1} = z_n + \frac{a_n [Y(z_n+c_n) - Y(z_n-c_n)]}{c_n}$$

where $a_n, c_n \to 0$ so that $\Sigma a_n = \infty$, $\Sigma a_n c_n < \infty$, $\Sigma a_n^2 c_n^{-2} < \infty$ (e.g., $a_n = n^{-1}$, $c_n = n^{-1/3}$).

Venter $\{V_2\}$ and Fabian $\{Fl,3\}$ have also generalized the K-W scheme to obtain procedures which converge in general but which are asymptotically optimal if the local behavior of M(x) at θ is smooth. This work has been extended to several dimensions. Relatively little attention in the literature has been paid to stopping rules.

The price paid for the robustness of these methods is that their behavior depends mainly on the nature of M(x) for x close to θ and do not take advantage of extra knowledge. Thus in problems where Y(x) depends in a known way upon several unknown parameters, it could be possible to develop more efficient if less robust sequential estima-

tion techniques. In particular, for estimation the LD.OL of the Probit model, the efficiency of the "best" Robbins-Monro method relative to a locally optimal design is only 64%.

ne a fa gen ar a strategen den teken an ein fan in trikke riger trikk tij it einer it oan wertertikken.

o hay be de andra andra andra de de antra de la construction de la construction de la construction de la constru

A parallel development to the Robbins-Monrc, Kiefer-Wolfowitz methods was the stochastic approximation methods of Box-Wilson [B11', which gave rise to a literature using the terms "response surface" and "steepest ascent" and "rotatable designs". Principally designed for multivariable applications, one observes Y(x) for a set of points x in k-dimensional space. Approximating EY(x) by a plane surface, one estimates the direction of sieepest ascent (gradient) and moves in that direction. Alternatively, one can appreximate EY(x) by a quadratic surface and estimate the point at which the quadratic is maximized. At each stage the estimated parameters are used not only to estimate the location of the maximum but to suggest another set of values of x at which to take additional observations. Rotatable designs are a special class of designs used around the point of interest [B5,6]. The general approach is rather pragmatic and informal compared with the methods proposed by Robbins and Monro, Kiefer and Wolfowitz, and Fabian and hence are less amenable to systematic analysis and evaluation. On the other hand, as these more formal methods developed they tended to resemble the Box-Walson approach more and more.

A variation of the Box-Wilson approach uses Partan, a method developed by Shah, Buehler and Kempthorne (S2). It replaces the

gradient or steepest ascent approach by a more sophisticated variation which combines two successive gradients in a method which is successful in speeding up convergence for deterministic problems and is apparently effective in the stochastic problems dealt with here. A review of the literature on response surface methodology was given by Hill and Hunter [H5].

A somewhat more specialized method of stochastic approximation applied in quantal response problems is that of the up-and-down method, introduced by Dixon and Mood [D3]. It is desired to estimate the dose x for which the probability of response assumes a certain specified value α . The possible dose levels of the experiment are equally spaced (possibly in a logarithmic scale). If a dose at level x leads to a response, the next dose applied is one step down and if it does not lead to response, the next dose applied is one step up. When the investigator terminates sampling, he estimates the parameters of the model by some method such as maximum-likelihood. A considerable number of variations of the basic approach have developed. See [C7, D1, W2]. For quantal response problems, this approach has a potential advantage over that of the Robbins-Monro method in that the associated estimation procedure makes use of the specific model applied. In doing so it of course loses the all-purpose robustness properties of the Robbins-Monro method.

3. Optimization Approaches

In principle the problems of sequential design of experiments can,

AARAN MANANSINSA KANANSA MANANSA MANANSA MANANSA MANANSA MANANSA MANANSA MANANSA MANANSA. Manansa manansa

by assuming a priori probability distributions and cost functions, be reduced to optimization problems which can be solved by backward induction. This idea has been exploited by Whittle [W2], who used it to set up a functional equation in terms of posterior probability distributions. However, the approach has been effective on very few rather simple problems. The insight provided by this statement ras limited value in most statistical problems.

It is not uncommon for investigators to use a myopic version of backward induction. Here the experimenter asks, after the outcome of each trial, "If I have at most one more experiment to perform, which if any will I perform?" In many cases this method seems to yield satisfactory results. I say "seems" because one seldom compares it with optimal procedures. One case where it has been used is in medical diagnosis [G2]. In principle this idea is also used in stepwise regression techniques for building up a good set of predictor variables.

It will be informative to see how this myophe policy works in a completely different context. To maximize a function f(x), $x \in E_{r^2}$ by the gradient method, one adjusts the n^{th} estimate by

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \mathbf{b} \frac{\partial \mathbf{f}}{\partial \mathbf{x}} (\mathbf{x}_n)$$

where $\partial f/\partial x$ represents the gradient or vector of partial derivatives with respect to the components of x. This method does not specify the value of the scalar of h. A special version called the <u>optimal</u> <u>gradient method</u> selects h to be that value for which $f'_{x_n} + h \frac{\partial f}{\partial x}(x_n)$

assumes its maximum. This can be regarded as a myopic sequential optimization procedure. Applying it to the function $f(x) = -(x_1^2 + cx_2^2)$, the gradient is $(-2x_1, -2cx_2)$ and an initial approximation $x = (x_1, x_2)$ to the point (0,0) which maximizes f is followed by $x^* = x + h \partial f/\partial x$, where $h = \frac{1}{2} (1 + \mu^{-2})/(c + \mu^{-2})$, $x_1^* = x_1(c-1)/(c + \mu^{-2})$, $x_2^* = -x_2(c-1)/(1 + c\mu^2)$, and $\mu = cx_2/x_1$. The value of f is reduced by a factor $f^*/f = (c-1)^2/(c + \mu^{-2})[c + \mu^{-2}]$. Since $\mu^* = cx_2^*/x_1^* = -\mu^{-1}$, this factor does not change in successive iterations even though h alternates between the above value and $h^* = \frac{1}{2}[1 + \mu^2]/[c + \mu^2]$. On the other hand

$$f^{*}(x) = -\{[x_1(1-2h)]^2 + [x_2(1-2ch)]^2\}$$

そうちんのないのないないないで、たちないでいいでいたないでいたいないないというとし、

could be much more rapidly reduced by alternating h between 1/2 and 1/2c .

For this particular function, if we assume no round-off error, alternating h between 1/2 and 1/2c accomplishes the maximization in two steps. In general, when f represents only the main term in the expansion of the function to be maximized, and there are round-off errors, two iterations will not suffice to reach the maximizing point. The above example illustrates that the rate of convergence can be faster than for the myopic policy called the "optimal gradient method" if the values of h are chosen with due attention to the characteristic roots of the quadratic form approximating the function to be maximized.

Two slightly less myopic policies which are probably more effective

and correspondingly more difficult to execute are the following: (1) Look two steps ahead into the future. This involves the mathematics of a two-step backward induction at each stage. (2) At each step ask whether there is an experiment e and a number of repetitions m so that the statistician would prefer m independent repetitions of e to any other (e^{+},m^{+}) and to stopping. If so, select e for the next trial. Apparently until recently this latter approach has been used only to determine reasonable stopping rules in problems with no choice of experimentation [A3, C2]. Recently Gittins and Jones [G1-] have used a variation of this idea effectively to gain new insight in the two-armed bandit problem by evaluating a choice in terms of how good it would be if we had to use that choice thereafter.

4. Asymptotically Optimal Procedures in Testing Hypotheses

Large sample theory provides useful insight in statistical problems for two reasons. First, the derivation and simple expression of appropriate distributions are easiest for sample sizes of 1, 2, and ∞ . Second, as sample size becomes large, many different philosophical approaches lead to results which are similar, and while uniformly best procedures are generally nonexistent for finite sample size, asymptotically optimal procedures do exist. It was hoped that large sample theory would provide insights which might permit one to bypass the need for backward induction. As we shall see latter, thus is relatively trivial in estimation problems where hotally optimal designs yield relatively efficient procedures easily.

In testing problems, the situation seems more difficult. But even here simple asymptotic results yield useful insights. In sequentially testing of a simple hypothesis $H_1: \theta = \theta_1$ versus a simple alternative $H_2: \theta = \theta_2$ where the successive observations $X_1, X_2, \ldots, X_n, \ldots$ are i.i.d. with density $f(x, \theta)$ the admissible procedures are the sequential probability-ratio test with limits A and B, $B \le 1 \le A$, on the likelihood-ratio $\lambda_n = \prod_{i=1}^{N} [f(X_i | \theta_1 / f(X_i | \theta_2)]]$. In a Bayesian framework with initial prior probabilities ξ_1 and $\xi_2 =$ $1 - \xi_1$, a cost of sampling c, and regrets for deciding wrong $r_1 =$ $r(\theta_i) > 0$, i = 1,2 the Bayes procedure is determined by appropriate limit. $A(\xi_1, r_1, r_2, c)$ and $B(\xi_1, r_1, r_2, c)$. As the cost of sampling $c \rightarrow 0$, the appropriate sample size $\rightarrow \infty$ and this is derived from the fact that $\log A \rightarrow \infty$ and $\log B \rightarrow -\infty$. In fact, $\log A \approx \log B \approx -\log c$, the posterior risk upon stopping as well as the posterior probability of being wrong is of the order of magnitude of c and the expected sample size is given by

ntaraalindhisaataa ahkteriida karin hiida sa hiida sa haa sa haraa sa aha waataa ku ku ku ku ku ku ku ku ku ku

$$E_{\theta_1}(n) \approx \frac{-\log c}{I(\theta_1, \theta_2)} = E_{\theta_2}(n) \approx \frac{-\log c}{I(\theta_2, \theta_1)}$$

where $I(\theta,\phi) = E_{\theta}\{\log[f(X,\theta)/f(X,\phi)]\} = \int \log[f(x,\theta)/f(x,\phi)]f(x,\theta)dx$ is the Kullback-Leibler information number. Indeed the main contribution to the risk or expected loss is the cost of sampling, and this is given by

$$R(\theta_1) \approx \frac{-c \log c}{I(\theta_1, \theta_2)}$$
 and $R(\theta_2) \approx \frac{-c \log c}{I(\theta_2, \theta_1)}$

In effect, the importance of the Kullback-Leibler information number derives from the fact that $I(\theta_1, \theta_2)$ measures how fast the posterior probability for θ_2 approaches zero when θ_1 is the true state of nature.

This simple result for sequentially testing simple hypotheses where there is no choice of experimentation suggests that if one had a choice of experiments to perform at each stage, the appropriate choice would depend on $I(\theta_1, \theta_2, e)$ the Kullback-Leibler number corresponding to data from experiment e. Indeed if $I(\theta_1, \theta_2, e_1) > I(\theta_1, \theta_2, e_2)$ and $I(\theta_2, \theta_1, e_1) > I(\theta_2, \theta_1, e_2)$, it seems clear that e_1 is preferable to e_2 . But if the last inequality is reversed, then e_1 is preferable to e_2 only if H_1 is true. The obvious implication is that if the data strongly suggests H_1 is true, one should select the next experiment to maximize $I(\theta_1, \theta_2, e)$ provided the evidence is not so overwhelming that it pays to stop sampling.

Suppose now that we move to the more complex problem which involves composite hypotheses with a fixed experiment. The simplest case is where $H_1: \theta = \theta_1$ and $H_2: \theta = \theta_2$ or θ_3 . Suppose $\theta_1, \theta_2, \theta_3$ start out with initial prior probabilities $\xi_{10}, \xi_{20}, \xi_{30}$. After n observations the posterior probabilities are $\xi_{1n}, \xi_{2n}, \xi_{3n}$, and assuming H_1 is true,

$$-nI(\theta_1, \theta_2) -nI(\theta_1, \theta_3)$$

$$\xi_{2n} \sim e, \xi_{3n} \sim e$$

Thus the rate at which the posterior probability of H, approaches

zero is determined by the minimum of $I(\theta_1, \theta_2)$ and $I(\theta_1, \theta_3)$. This observation leads to the following suggested procedure for the more general problem of testing the composite hypotheses $H_1: \theta \in \omega_1$ vs. $H_2: \theta \in \omega_2$ when there is a choice of experiments. Stop sampling after the nth observation if the posterior probability of one of the hypotheses is of the order of magnitude of c (or if the posterior risk of stopping and making a terminal decision is of this order). Otherwise select the next experiment e to maximize

$$\inf_{\substack{\mathfrak{p} \in \mathbf{a}(\hat{\theta}_n)}} \mathbb{I}(\hat{\theta}_n, \varphi, e)$$

where $\hat{\theta}_{n}$ is the maximum likelihood estimate of θ and $\mathbf{a}(\theta)$, the alternate hypothesis to θ , is defined by

$$\mathbf{a}(\theta) = \begin{cases} \omega_2 & \text{if } \theta \in \frac{\omega_1}{\omega_2} \\ \omega_1 & \omega_2 \end{cases}$$

It should be noted that e is selected from among the class of randomized experiments, and it has been assumed that each of these experiments has the same low cost c. If the cost per experiment varies, then one deals with information per unit cost rather than information.

The method suggested above was shown to be asymptotically optimal under mild conditions $\{C1, C5\}$ as $c \rightarrow 0$ in the sense that for each θ it yields a risk

$$E(\theta) \approx \frac{-c \log c}{I(\theta)}$$

where

and a first way the birth of the data see a rest with the data see and the second second

$$I(\theta) = \sup \quad \inf \quad I(\theta, \varphi, e) \\ e \in \xi^* \quad \varphi \in \mathbf{a}(\theta)$$

and $\boldsymbol{\zeta}^*$ is the class of randomized experiments derived from the class $\boldsymbol{\xi}$ of available or "elementary" experiments. Moreover for any alternative procedure to do better for some value of θ , it must do worse by an order of magnitude for some other value of θ . This result was first proved for the case where ω_1, ω_2 , and $\boldsymbol{\zeta}$ were finite. Bessler [B2] extended the result to the case where $\boldsymbol{\xi}$ is infinite and the problem of choosing between two hypotheses could be replaced by a choice among k actions. Albert (Al) extended this result further to the case where the hypothesis spaces ω_1, ω_2 may be infinite sets.

নতানাক উদ্ধাৰ্থ মান্দ্ৰ প্ৰথম ব্যৱহায় আৰু বিষয়ে হ'ব প্ৰথম কৰা বাবে ব্যৱহায় বিষয়ে কৰে বিষয়ে মান্দ্ৰ কৰা মান

Here a fundamental difficulty appeared. In such a simple problem as testing whether the probability of response to one drug is greater than for another drug, the two hypothesis spaces are adjacent to one another and $I(\theta)$ vanishes on the boundary. Then the asymptotic optimality breaks down. Heuristics indicated that the difficulty arises more from the stopping rule than the experimental design aspect of the problem, and G. Schwarz [S1] attacked that problem by studying optimal sequential procedures for testing that the mean μ of a normal distribution is μ_1 versus the alternative that it is $-\mu_1$ when it is possible that μ could be $\pm \mu_1$ or 0. In the latter case it doesn't matter what terminal decision is made. His results extended to asymptotically optimal and Bayes results for testing that the mean exceeds μ_1 versus the alternative that it is less than μ_2 , $(\mu_2 < \mu_1)$, when it is possible that $\mu_2 \leq \mu \leq \mu_1$, in which case either

decision is equally satisfactory. In other words, this is the case of an <u>indifference</u> zone. Here asymptotically Bayes procedures consist of stopping when the posterior risk of stopping and making a terminal decision is O(c) and yield overall risks of order $O(-c \log c)$.

Finally Kiefer and Sacks [K1] combined these results to obtain an asymptotically optimal procedure. for problems in sequential design where the parameter points for which various actions are preferred are separated by indifference zones. In these results the key information number is expressed by

also de la comparte de la comparte de la comparte de la contraction de la comparte de la comparte de la compart

$$I(\theta) = \sup \sup \inf I(\theta, \varphi, e)$$

$$e \in \mathcal{E}^* \quad i \in G_{\Theta} \quad \varphi \notin \omega_i$$

where ω_i is the set of θ 's on which the ith action is optimal, and G_{θ} is the set of i for which the ith action is optimal when θ is the true state of nature. (In the two action problems, $G_{\theta} = \{1,2\}$ for θ in the indifference zone.) The appropriate experiment is the randomized experiment $e \in \xi^*$ which yields $I(\theta)$ as the supremum in the above expression and

$$R(\theta) \approx \frac{-c \log c}{I(\theta)}$$

Both the proof and the method are simplified considerably in the Kiefer and Sacks paper where a two-stage sampling procedure is used. An initial large sample of size o(-log c) is followed by an estimate of θ and a second sample of appropriate size on an appropriate choice of e.

In principle this approach is extremely successful in bypassing the need for backward induction. Asymptotically optimal results are obtained with recourse only to Kullback-Leibler information numbers and likelihood-ratio statistics. However, there are several shortcomings. First, the role of indifference zones implies that the simple problem of deciding whether the mean μ of a normal distribution is positive or negative with a positive loss such as $|\mu|$ attached to the wrong decision is not covered. Second, the approach is very coarse for moderate sample size problems. Indeed the Kiefer-Sacks two-stage variation sidesteps the issue of how to experiment in the early stages whereas the original Chernol'f approach simply treats the estimate of θ based on a few observations with as much respect as that based on many observations.

the calles in Direction in the of Societ Stream and the Statistic Statistics

On top of these shortcomings the asymptotic analysis distinguishes sharply between terms of order of magnitude of c and of cloge, whereas the difference in most applied examples may be less than overwhelming. (A proper analysis should pay more attention to the fact that log c is dimensionally wrong. The quantity c should be normalized approximately with respect to the costs of making the wrong decision. This normalization occurs naturally if one stops when the posterior risk of stopping is of the order of the cost of stopping. and the second strategies and a second se

In addition to this approach, alternative procedures have been proposed by Lindley [L2], DeGroot [D2], and Box and Hill [B7]. For example Lindley suggested measuring the value of an experiment in

terms of the Shannon Information or Entropy. One may select at each stage the experiment for which the expected reduction in entropy is a maximum. To be more specific, if $\theta_1, \theta_2, \ldots, \theta_r$ are the possible states of nature among which one must decide and ξ_1 is the prior probability of θ_i , the entropy is $-\Sigma \xi_1 \log \xi_1$. After an experiment yielding X_e the prior probabilities ξ_1 are replaced by ξ_1^* proportional to $\xi_1 f(X_a | \theta_i, e)$ and the reduction in entropy is

$$\Sigma[\xi_1 \log \xi_1 - \xi_1^* \log \xi_1^*]$$

whose expectation may be computed to be

والمنافعات المعادية والمعالمة المنافعة المحالية والمحالية والمحالية والمحالية والمحالية والمحالية والمحالية والمحالية

$$\sum_{i} I(\theta^*, \theta_i, e)$$

where θ^* corresponds to an ideal distribution with density $\Sigma \xi_i f(x_i, \theta_i, e)$.

Box and Hill started with the same approach, but to simplify the calculus approximate the expected reduction in entropy by an upper bound

$$\frac{1}{2} \sum_{i \in j} [I(\theta_i, \theta_j, e) + I(\theta_j, \theta_i, e)]$$

which they proceed to use to select the next experiment. Neither of these approaches is asymptotically optimal except in special "symmetric" problems. One may expect the Box-Hill approach to fail to be optimal because it is only an approximation to the method proposed. Apparently the Lindley approach, which seems more reasonable, fails because a myopic one-stage-ahead policy cannot be depended on for optimality as was seen in the illustration of the optimal gradient method.

On the other hand Meeter, Pirie and Blot [M2] carried out some extensive Monte Carlo experimentations on two problems. These were to select the single odd coin in a group of k coins and to identify three normal populations with common variance if the values of the three means are specified but the appropriate order is not known. In both of these the BCx-Hill approach did better than the Chernoff approach for sample sizes that were limited by a stopping rule which led roughly to error probabilities of .05. Apparently the difficulty with the asymptotically efficient approach of Chernoff was that initial experiments which seems to show up in these examples. Blot and Meeter [B4] subsequently attempted to develop an alternative which would be asymptotically optimal and effective in the early stages. Their method seems to be effective an a special class of problems.

ないたいななないのないないないのとうないないです。

At this time the major theoretical problems seem to be the problem of no-indifference zone and finding effective methods of experimentation at the early stages of sampling. For the problem of noindifference zone, the problem of deciding the sign of a normal mean was used as a prototype on the ground that its solution could be extended via logarithm of likelihood-ratio to more general situations. Although this work was done in the context of no experimental choice, one consequence is of sume interest here. Consider the problem where the cost of deciding wrong is $k(\mu)$ and the cost per observation is $c \rightarrow 0$. Then using Bayes procedures the risk for non-sequential pro-

cedures is $R(\mu) = O(c^{1/2})$. The risk for the optimal sequential procedure is $R_{\mu}(\mu) = O(c^{2/3})$. Hald and Keiding [H1,2] have shown that the risk for k-stage sampling procedures satisfy $R_{\mu}(\mu) = \frac{\gamma_{k}}{(\log c)} \frac{2\gamma_{k}-1}{2}$ where $\gamma_{k} = (2^{k}-1)/(3 \cdot 2^{k-1}-1)$.

5. Optimal Design in Estimation

どうぶん あた えい どういをち

As a preliminary to this section we mention results in two types of problems. For sequentially estimating the muan of a normal distribution with known variance, using squared error loss and constant cost per observation, the optimal sample size n_0 is obtained by minimizing on $+ k\sigma^2 n^{-1}$. Thus n_c is $(k\sigma^2/c)^{1/2}$, and the optimal risk is $2(ck\sigma^2)^{1/2}$. If the variance σ^2 is not known, an approach suggested by Robbins [R5] consists of sampling until the sample size n exceeds 2 and the current estimate of $(k\sigma^2/c)^{1/2}$. Thus we stop when $n \ge 3$ and

$$\sum_{i=1}^{n} (X_i - \overline{X}) \leq ck^{-1}n^2(n-1)$$

Results of Starr and Wordroofe [S7] indicate that the difference between the optimal risk and that for this procedure is O(c), i.e., the cost of not knowing the nuisance parameter σ is equivalent to that if a finite number of observations. (This cost is shout the cost of one observation unless σ is extremely small, in which case these observations are excessive). Alvo [A2] has attained precise bounds in a Bayesian context. The point of this discussion is that in estimating, one can expect to do very well using rather simple ideas. That is,

it is easy to find procedures ...ich achieve risks which are $2(ck\sigma^2)^{1/2} + O(c)$ where the first term would be optimal when σ^2 is known. A nontrivial notion of asymptotic optimality must attempt to minimize the O(c) term. On the other hand, the practical use for such a nontrivial optimality may not be great.

and a set have had been a set of the set of

the tertestic the second of the tert of the second state and the second states and the

A second result concerns one-armed bandit problems. This may be stated as follows. Let X_1, X_2, \ldots be independent observations on a random variable X. A player who plays $n \leq N$ times collects $X_1 + X_2 + \cdots + X_n$ whose expectation is nE(X). Determine n sequentially to maximize the expected payoff which is E(n)E(X). If E(X) > 0, it pays to play N times, and the expected payoff is NE(X). If E(X) < 0, it pays not to play. Chernoff and Ray [C6] have given a characterization for the solution of the normal version where the X_i are normal with unknown mean μ and known avariance σ^2 , and μ has a specified normal prior distribution, and N is large. Here it is shown that the expected loss due to ignorance of the sign of μ is of the order of magnitude of $(\log N)^2$. One may conjecture that the two-armed bandit problem would share this property.

A number of papers in optimal design approach the sequential estimation problem from a myopic iterative point of view without much attention to stopping rules [B8, F6, P2, S5, S6]. For example, consider the normal, linear regression problem with

$$y = \theta' x + u$$

where u is normal with mean 0 and constant variance 1, and where x may be selected from some compact set S. The covariance matrix of the estimates of θ based on n observations corresponding to x_1, x_2, \dots, x_n is

$$\Sigma_{n} = \begin{bmatrix} n \\ \Sigma \\ i=1 \end{bmatrix} \cdot \mathbf{x}_{i} \mathbf{x}_{i}^{\prime} \end{bmatrix}^{-1}$$

One approach is to select the $(n+1)^{st}$ experiment, i.e., x_{n+1} , to minimize the generalized variance, $|\Sigma_{n+1}|$. Since

$$\Sigma_{n+1}^{-1} = \Sigma_n^{-1} + J_{n+1}$$

where $J_{n+1} = x_{n+1}x_{n+1}'$ is the Fisher information contributed by the $n+1^{st}$ observation and is of rank one, the matrix identity

$$(A+xx')^{-1} = \left[I - \frac{A^{-1}xx'}{1+x'A^{-1}x} \right] A^{-1}$$

facilitates the minimizing calculation. The iteration involved is independent of the actual data observed and is also used to calculate fixed sample size designs which minimize the generalized variance. See also [N3]. Minor variations of this basic idea apply Bayesian notions and can be used in nonlinear problems.

This approach has two shortcomings. First, the emphasis on the criterion of generalized variance is deplorable. While the criterion of minimizing the generalized variance has the aesthetic property of leading to invariance of optimality under linear transformations of the parameter space, this elegant mathematical property simply disguises the underlying fact that the criterion has no basic statistical justification and simply delegates the scientist's responsibility of selecting the criterion to the vagaries of the mathematical soructure of the problem. Thus in a probit model where one is primarily interested in the LD5 and only slightly interested in the LD50, the use of the generalized variance criterion leads to an efficiency of as little as .56.

なないのでないまでいますとないます。そうできないないとうと

ration was shaked as

It is true that in the linear regression problems where one is concerned with all the unknown parameters, the design which minimizes the generalized variance also minimizes the maximum variance of the estimated regression for all $x \in S$ [K3]. However, this min-max optimality interpretation for interpolation disappears when one is concerned with a subset consisting of several but not all of the unknown parameters.

This criticism of the use of generalized variance (i.e., D-optimality) does not invalidate the general idea of the myopic iteration, which can also be applied to other criteria. However, the second shortcoming is that any asymptotic optimality obtained is basically the cheap one which any locally optimal design attains. What would be more interesting is a demonstration of a more sensitive optimal'ty of the sort suggested in our discussion of the Robbins, Starr, Woodroofe, Alvo results. But once again it is far from clear that a myopic policy will be successful in this more delicate task. On the other hand, one may argue that this task is more of academic than practical value. Once again the issue centers about what constitutes effective procedures of cumulating information rapidly in the early stages of sampling and how important are these early stages. I found little discussion in the literature which was relevant to this problem. An exception consists of a paper by myself [C3] and one by Mallik [M1] which combine the ideas of the bandit problems and the Robbins approach to sequential estimation. I believe that these point in the correct direction to assess appropriate orders of magnitude, and a brief discussion follows.

The two-armed bandit was dismissed early in this paper as a failure as a prototype example to clarify the problems of sequential design of experiments. I now propose to disinter it as a problem of theoretical relevance by considering it in a new context. Incidentally some theoretical insights have been contributed by Gittins and Jones [G1], to whom we referred earlier, and to Vogel [V3] and Fabius Von Zwet [F4] who studied minimax solutions.

Suppose that there are two instruments which can be used to measure a parameter μ , but it isn't known which is more accurate. How should one select between the two instruments, and when ought one to stop sampling? More specifically, suppose X is normally distributed with unknown mean μ and variance σ_1^2 and Y is normally distributed with mean μ and variance σ_2^2 . The cost of sampling is c per unit observation where $c \rightarrow 0$. The cost of estimating incorrectly is $k(\hat{\mu}\cdot\mu)^2$, where $\hat{\mu}$ is the estimate of μ . In one version of this

problem σ_1^2 and σ_2^2 are both unknown. In another we know σ_1^2 but σ_2^2 is unknown. Chernoff [C4] approached the first using an approximation to the solution of the two-armed bandit problem. Mallik [M1] attacked the other by using the solution of the one-armed bandit problem. Let us consider this simpler case.

ga ter generaliset in der State in der State in der State State in der State in der State in der State in der St

While σ_2^2 is unknown, it makes sense to take observations on Y, simultaneously obtaining information on $\,\mu\,$ and an estimate of $\,\sigma_{\!\!\!\!\!2}$. One continues until the Robbins-type procedure suggests stopping, or until the evidence indicates that $\sigma_2 > \sigma_1$, in which case one estimates how many additional observations from X are advisable before terminating the sampling process. A careful computation shows that if $\sigma_1 < \sigma_2$, the loss attributed to taking n observations from Y before switching is roughly proportional to $n(\sigma_p - \sigma_1)$. If $\sigma_1 > \sigma_2$, the appropriate number of observations is $n_0 =$ $(co^2/k)^{1/2}$ on Y and a decision to switch to X after n observations leads to a loss of $(n_c-n)(\sigma_1-\sigma_2)$. But in our one-armed bandit problem the expected loss due to taking n observations when $\mu < 0$ was -n μ , whereas the expected loss due to taking n observations when $\mu > 0$ was (N-n) μ . Relating N and μ to n and $\sigma^{}_1 \text{-} \sigma^{}_2$ suggests Mallik's procedure of applying the solution of the one-armed handit to decide when to switch to X .

Monte Carlo simulations suggest that this method yields a highly efficient design for sequential experimentation. Theoretical considerations, supported only partly by the Monte Carlo simulations, indicate

that while losses due to error of estimation are of the order of magnitude of $c^{1/2}$, the loss attributed to lack of knowledge of σ_2 is of the order of magnitude of $c(\log c)^2$. This is slightly larger than the magnitude O(c) achieved in the nondesign problem of Robbins, Starr, Weedroofe, and Alvo.

6. Applications

The ideas of sequential experimentation appear in one form or another in a variety of fields of application. Some of the most important ones have extensive literatures, and we barely mention these. In particular, survey sampling is one field where double sampling and several-stage sampling have an extensive history. Indeed the origin of sequential analysis can be traced back to the double sampling inspection scheme of Dodge and Romig [D5]. In very few of these fields has a serious attempt been made to explore optimality from a fundamental point of view. Typically an ad hoc class of procedures has been proposed, and sometimes the best among these is characterized. Seldem does one attempt to compare these with some more generally optimal procedure. Thus one is often in the dark about the limits of further possible improvements.

7. Multi-Level Continuous Sampling Inspection

An early form of sequential experimentation was in the multilevel inspection schemes of Dodge [D4]. Lieberman and Solomon [L1] rephrased some previous ambitious optimization problems to formulate

a simpler but highly relevant problem. Imagine a continuous production process yielding many items which can be inspected. As the items bass by, they are inspected with one of several available probabilities $p_1 = 1 > p_2 > \cdots > p_k$. If a defect is found, the rate of inspection is increased. If n_i successive non-defects are found while sampling at level λ_i , the rate of inspection is reduced. When the production process turns out items which are defective independently with constant probability, the "state" of the inspection system describes a simple stationary Markov process whose limiting characteristics are easily evaluated. Thus one can compute the costs and gains of this multi-level inspection scheme for each p. One can easily maintain a minimum level of quality of output. When the production process goes out of control, this system seems to respond sensibly. There is one major aspect in which the Lieberman-Sclomon problem differs from the class of problems with which we have previously been concerned in this paper. Those involved termination in a finite time. This process is stationary and should be thought of as going on indefinitely. Indeed this paper initiated a good deal of subsequent research in Markov decision problems and constituted an early form of stochastic control.

8. Largest of k-Means

New-

" South Set, here will be way

As initially formulated [B1] this problem specifies k normal populations Π_i , i=1,2,...,k with means μ_i and common known variance σ^2 . The object is to decide after n observations on each popula-

tion which has the largest mean. The natural procedure is to select the population corresponding to the largest sample mean. The sample size n is selected so as to assure that the probability of correct selection attains at least a given value 1- α if the largest population mean is at least 5 greater than each of the others. Here α and 5 are specified and n is computed as a function of k, σ^2 , α , and 5. This computation is relatively trivial since there is a "least-favorable" configuration of means $\mu_1, \mu_2, \dots, \mu_k$ where $\mu_1 = 5$ and $\mu_2 = \mu_3 = \dots = \mu_k = 0$.

The problem of sequential experimentation appears when one may decide to proceed sequentially. Bessler [B2] applied the theory of Part II to obtain a procedure which is asymptotically optimal if one can assume that the largest mean exceeds all the others by at least a fixed amount.

This result seems to have been ignored by subsequent workers in the field who applied sequential screams where each population is sampled equally often. Subsequently Guttman [G]] and Paulson [P1] developed some alternative multi-stage procedures where the results of each stage were used to discard some populations from further consideration. Alternative methods have been developed by D. Hoel [H7] and J.W.H. Swanepoel and J. Venter [S9]

There has been an extension laterature extending this problem to other distributions and other parameters. The variation of the two-armed bandit problem, where the payoff conturs only after the last

observation and the experimenter decides which is the better arm, is also an example of this type of problem. The largest of k populations problem corresponds to a k-armed bandit problem subject to two variations. The total sample size is not necessarily fixed. Also, in dealing with the k-armed bandit problem one does not typically apply the rather artificial criterion of maintaining a minimum probability of correct selection at configurations of parameters where the largest exceeds the others by at least a specified 5.

لمعلمهما والمعارفة والمعارفة المتعارفة المعادلة والمعارفة والمعارفة والمعارفة والمعارفة المعارفة والمعارفة وال

Several variations of the two-armed bandit problem occur in application contexts. In connection with medical trials where the arms refer to treatments, various investigators [R2, Z1, S4, C9, F7, F6] have investigated Play the Winner Rules, which continue the use of a treatment as long as it is successful and switch when it fails, as well as other "adaptive" methods. These rules can apply in problems with an infinite horizon of patients to be treated. On the other hand, one-armed bandit variations applied to medical trials were discussed by Chernoff [C3], Colton [C8], and Anscombe [A4]. In Colton's version drugs are tried alternately until there is an implicit decision that one is better and the remainder of a horizon of N patients are treated with the drug that is considered better. The one-armed bandit problem comes up naturally in a rectified sampling inspection problem too [].

いちょうちょう ちょうちょう ちょうちょう ちょうちょう ちょうちょう ちょうちょう ちょうちょうちょうちょうちょうちょうちょう

and a set of the set of

Finally, Hellman and Cover [H3] have exploited randomization in a finite memory two-armed bandit problem where the observer is res-

tricted to knowing only the current sample size n and the value of a k-valued function of the past.

9. Screening Experiments

In pharmaceutical research where one seeks drugs which have antidisease activity, one must screen many possible candidate chemical formulations by testing them first on animals. It is important to devise a system where many drugs are tested and quickly discarded (because of the expense of testing) unless they show indications of activity. In that case they are retested more thoroughly. This procedure passes each drug through several screening experiments, each more elaborate than the preceding. If the drug passes all of these, it is regarded as a candidate for further research and testing on humans. (See $[D_1^r, R_2^r]$).

10. Group Testing

During World War II 1* was noted by Dorfman [D6] that the cost of testing blood specimens of individuals for the presence of a moderately rare disease could be reduced considerably by combining the samples of many individuals. If the combined sample showed no sign of disease, the entire group was passed at the cost of one test. If the combined sample shows signs of disease, the individual specimens could be tested separately. With appropriate grouping depending on the overall frequency of disease, this system and improvements produced considerable savings. This subject is elaborated upon by

Sobel and Nebenzahl [S3] who contributed a thorough bibliography.

11. Search Problems

2

they shall be a shift of the shall be a series with the

Search problems have appeared in a variety of contexts and applications. They deal with the problem of locating an item which may be in any one (or sometimes possibly none) of k locations, each of which may be searched and yield the time, if it is there, with a specified probability. Often these problems are treated as combinatorial problems and k is large. No attempt will be made to elaborate on the topic, which has an extensive literature which was surveyed by Enslow [E3], and some further references are given by Sweat [S10]. A different approach is given by Lipster and Shiryaev [L3], who use diffusion approximations for a variation of the search problem where k is not large.

12. Control Theory

Multi-level sampling inspection is one form of control applied to maintain the quality of a continuous production line. Box and Jenkins [B9, B10] have considered the problem of monitoring a complex chemical production process where slow changes in the underlying environment may require adjustment of inputs to maintain optimality. They suggest perturbing the inputs off the position that seems optimal, to detect and estimate possible changes in the response surface by measuring the efficiency of the system. In this way the estimate of the current optimum is continuously updated. The price of this is

the loss of efficiency involved in perturbing the system to measure the response surface. If the perturbation is too small, the surface and possible changes in it are not measured precisely enough. If the perturbations are too large, the experimentation reduces the efficiency of the system. This type of system may be thought of as a stationary control problem.

13. Forcing Experiments to be Balanced

والملا المقاولة المتعالية حمد ولمالك ولمار والمالية والملائل

In clinical trials as well as in many other scientific investigations, the need to avoid bias requires experimentation where the parties involved do not know whether they are receiving a treatment or a control. Thus assignments may be made by using a fair coin, but in small-sized experiments this may result in a severe imbalance. Hlackwell and Hodges [E5] and Efron [E1] have considered alternative schemes to complete randomization to avoid several kinds of bias, e.g., selection bias and experimental bias. One scheme considered is to assign the treatment with probability p if the treatment has been used more often than the control and (1-p) if the control has been used more often. Efrom indicates a preference for $p \approx 2/3$ and compares the balancing properties of this and other schemes as well as the potentialities for selection bias and experimental bias.

14. Miscellane s

Problems of information storage and retrieval and error-correcting codes involve notions of sequential experimentation in a fashion which

does not fit traditional approaches of statistics very well. Nevertheless, these problems have fundamental statistical aspects.

In clinical problems and control problems there are classes of problems where the response to an experiment is not observed immediately and some theory is required to deal with delayed observations [E2, S8].

A useful bibliography on design of experiments is given by Herzberg and Cox [H4].

urdu den in standarden der bieren sich ein sie einigen under bieren werden under anderen werden der sienen der

35

and the second state of th

References

- [A1] Albert, A.E. (1961). The sequential design of experiments for infinitely many states of nature. <u>Ann. Math. Statist. 32</u> 774-799.
- [A2] Alvo, M. (1972). Bayesian sequential estimation. Stanford University Technical Report No. 47.
- [A3] Amster, S.J. (1965). A modified Bayes stopping rule. <u>Ann.</u> <u>Math. Statist. 34</u> 1404-1413.
- [A4] Anscombe, F.J. (1963). Sequential medical trials. J. Amer. Statist. Assn. 60 584-601.
- [B1] Bechhofer, R.E. (1954). A single sample multiple decision procedure for ranking means of normal populations with known variances. Ano. Math. Statist. 25 16-39.
- [H2] Bessler, S. (1960). Theory and application of the sequential design of experiments, k-actions and infinitely many experiments: Part I - theory. Part II - applications. Stanford University Technical Reports 55 and 56.
- [33] Blackwell, D. and Hedges, J.L. (1957). Design for the control of selection bias. Ann. Math. Statist. 28 449-460.
- [B4] Blot, W.J. and Meeter, D. Sequential experimental design procedures. To be published in <u>J. Amer. Statist. Assn</u>.
- [B5] Box, G.E.P. and Draper N.R. (1959). A basis for the selection of a response surface design. J. Amer. Statist. Assn. 54 622-654.
- [B6] Box, G.E.P. and Draper N.R. (1963). The choice of a second-order rotatable design. <u>Biometrika 50</u> 335-352.
- [B7] Box, G.E.P. and Hill, W.J. (1967). Discrimination among mechanistic models <u>Technometrics</u> 9 57-71.
- [B8] Box, G.E.F. and Hunter, W.G. (1965). Sequential design of experiments for nonlinear models. <u>Proc. IEM Sci. Comp. Symp. on</u> <u>Statistics</u>, 123-135.
- [B9] Box, G.E.P. and Jenkins, G.M. (1962). Some statistical aspects of adaptive optimization and control. <u>1. Roy. Statist. Soc. Ser. B</u> 24 297-331.

[B10]	Box, G.E.P. and Jenkins, G.N. (1970). <u>Time Series Analysis:</u> Forecasting and Control. Holden-Day, San Francisco.
[B11]	Box, G.E.P. and Wilson, K.B. (1951). On the experimental attainment of optimum conditions. J. Roy, Statist. Soc. Ser. B 13 1-45.
[C1]	Chernoff, H. (1959). Sequential design of experiments. Ann. Math. Statist. 30 755-770.
[@]	Chernoff, H. (1965). Sequential tests for the mean of a normal distribution III (small b). Ann. Math. Statist: 36 28-54.
[ᠿ]	Chèrnoff, H. (1967). Sequentiàl models for clinical trials. Proc. Fifth Berkeley Symp. Prob. and Math. Statist. 4 805-812.
[C4]	Charnoff, H. (1971). Thé efficient éstimation of a parametér measurable by two instrumènts of Unknown précisions. <u>Optimizing</u> Méthods in Statistics, J.S. Rustagi, ed., Academic Préss, New York.
[C5]	Chernöff, H. (1972), <u>Sequential Analysis and Optimal Design</u> . Soc. for Industrial and Applied Math. (SIAM), Philadelphia.
[06]	Chernoff, H. and Ray, S.N. (1965). A Bayes sequential sampling plan. <u>Ann. Math. Statist. 36</u> 1387-1407.
[7]	Cochran, W.G. and Davis, M. (1965). The Robbins-Monro method for estimating the median lethal dose. <u>J. Roy. Statist. Soc.</u> <u>Ser. B 27</u> 28-44.
[C8]	Colton, T. (1965). A model for selecting one of two medical treatments. J. Amer. Statist. Assn. <u>58</u> 388-400.
[C9]	Cornfield, J., Halperin, M. and Greenhouse, S. (1969). An adap- tive procedure for sequential clinical trials. <u>J. Amer. Statist</u> . <u>Assn. 64</u> 759-770.
[D1]	Davis, M. (1970). Comparison of sequential bioassays in small samples. J. Roy. Statist. Assn. Ser. B 33 78-87.
[12]	DeGroot, M.H. (1962). Uncertainty, information and sequential experiments. Ann. Matt. Statist. 33 404-419.
[D3]	Dixon, W.J. and Mood, A.M. (1948). A method for obtaining and analyzing sensitivity data. J. Amer. Statist. Assn. 43 109-126.
[D4]	Dodge, H.F. (1943). A sampling inspection plan for continuous production. <u>Ann. Math. Statist. 14</u> 264-279.

¥

à

•

â

her supported the second sec

• •,

and the second second

·······

Mar and the state part of the set of the set

2 1 114

N. L WY dimes th

.

19

́-,

[D5]	Dodge, H.F. and Romig, H.G. (1929); A method of sampling inspection. The Bell System Tech. Journal 8 613-631.
[D6]	Dorfman, R. (1943). The detection of defective members of large populations. Ann. Math. Statist. 14 436-440.
[17]	Dunnett, C.W. (1960). Statistical theory of drug screening. Quantitative Méthods in Pharmacology, H. de Jonge, ed., North Holland, Amsterdam 212-231.
[E1]	Efron, B. (1971). Forcing a sequential experiment to be balanced. <u>Biometrika 58</u> 403-417.
[12]	Ehrenfeld, S. (1970). On a scheduling problem in sequential analysis. Ann. Math. Statist. 41 1206-1216.
[E3]	Enslow, P. Jr. (1966). A bibliography of search theory and reconnaissance theory. <u>Naval Rés. Logist: Quart. 13</u> 117-202.
(ř1]	Fabian, V. (1967). Stochastic approximation of minima with improved asymptotic speed. <u>Ann. Math. Statist. 38</u> 191-200.
[12]	Fabian, V. (1968). On the choice of design in stochastic approximation methods. <u>Ann. Math. Statist. 39</u> , 457-465.
[17]	Fabian, V. (1969). Stochastic approximation for smooth functions. Ann. Math. Statist. 40 299-302.
[F4]	Fabius, J. and Von Zwet, W.R. (1970). Some remarks on the two-armed bandit. <u>Ann. Math. Statist. 41</u> 1906-1916.
[F 5]	Feldman, D. (1962). Contributions to the 'two-armed bandit' problem. <u>Ann. Math. Statist. 33</u> 847-856.
[F 6]	Federov, V.V. (1972). The Theory of Optimal Experiments, Academic Press, New York.
[F7]	Flehinger, B. and Louis, T. (1971). Sequential treatment allo- cation in clinical trials. <u>Biometrika 58</u> 419-426.
[F8]	Fushimi, M. (1973). An improved version of a Sobel-Weiss play- the-winner procedure for selecting the better of two binomial populations. Cornell University Technical Report.
[G1]	Gittins, J.C. and somes, D.M. (1972). A dynamic allocation index for the sequential design of experiments. Cambridge University Technical Reports

[02]	Gorry, G.A. and Barnetts G.O. (1968). Sequential diagnosis by computer. J. Amer. Med. Assn. 205 849-854.
[03]	Guttman, I. (1963). A sequential procedure for the best population. Sankhys Ser, A 25 25-28.
[HL]	Hald, A. and Keiding, N. (1969). Asymptotic properties of Bayesian decision rules for two terminal decisions and multiple sampling I. J. Roy. Statist. Soc. Ser. B 31 455-471.
[H2]	Hald, A. and Keiding, N. (1971). Asymptotic properties of Bayesian decision rules for two terminal decisions and multiple sampling II. J. Roy. Statist. Soc. Ser. B 34 55-74.
[H3]	Hellman, M.E. and Cover, T.M. (1970). Learning with finite memory. <u>Ann. Math. Statist. 41</u> 765-782.
[H4]	Herzberg, A.M. and Cox, D.R. (1969). Recent work on the design of experiments: a bibliography and a review. <u>J. Roy</u> . <u>Statist. Soc. Ser. A 132</u> 29-57.
[H5]	Hill, W.J. and Hunter, W.G. (1966). A review of response surface methodology: a literature survey. <u>Technometrics</u> 8 571-590.
(H6)	Hodges, J.L. Jr. and Lehmann, E.L. (1956). Two approximations to the Robbins-Monro process. <u>Proc. Third Berkeley Symp. on</u> <u>Prob. and Math. Stat. 1</u> 95-104.
[H7]	Hoel, D. (1971). A method for the construction of sequential

A STREET OF A ST

NAME AND ADDRESS OF A DOCUMENT OF A

- selection procedures. Ann. Math. Statist. 42 630-642.
 [K1] Kiefer, J. and Sacks, J. (1963). Asymptotically optimum sequen-
- tial inference and design. Ann. Math. Statist. 34 705-750.
- [K2] Kiefer, J. and Wolfowitz (19%). Stochastic estimation of the maximum of a regression function. <u>Ann. Math. Statist. 23</u> 462-466."
- [K3] Kiefer, J. and Wolfowitz, J. (1959). Optimum designs in regression problems. <u>Ann. Math. Statist.</u> <u>30</u> 271-294.
- [11] Lieberman, G.J. and Solomon, H. (1955). Multi-level continuous sampling plans. <u>Ann. Math. Statist</u>, 26 686-704.
- [12] Lindley, D.V. (1956). On the measure of the information provided by an experiment. <u>Ann. Math. Statist. 27</u> 986-1005.
- [L3] Lipster, R. and Shiraev, A.N. (1965). A Bayesian problem of sequential search in diffusion approximation. <u>Theor. Probability</u> <u>Appl. 10</u> 178-186.

こんかくしょうかんか しかしまた たいびしょう かんだかかい ひかかく 気気 かだかか かんかん ユーステム ひんかん しゅうかん しかん かんしょうしょう

- Mallik, A.K. (1971). Sequential estimation of the common mean [M] of two normal populations. Stanford University Tech: Report No. 42.
- Meeter, D., Pirie, W. and Blot, W. (1970). A comparison of [M2] two model-discrimination criteria. Technometrics 12 457-470.

o by de se fan de brand de fan Arman de Brandel de fan de brander fan de brander fan de brander fan de brander

- Mitchell, T.J. (). Applications of an algorithm for the [NG] construction of 'D-optimal' experimental designs in r runs. Proc. Int. Symp. on Statist. Design and Linear Models.
- [P1] Paulson, E. (1967). Sequential procedures for selecting the best of several binomial populations. Ann. of Math. Statist. 38 117-123.
- Pazman, A. (1968). The sequential design of experiments per-[P2] formed at several measurement points. Theor. Probability Appl. 13 457-467 .
- [R1] Robbins, H.E. (1952). Some aspects of the sequential design of experiments. Bull. Amer. Math. Soc. 55 527-535.
- [R2]Robbins, H.E. (1956). A sequential decision problem with finite memory. Proc. Nat. Acad. Sci. 42 920-923.
- [R3]Robbins, H.E. (1959). Sequential estimation of the mean of a normal population. Probability and Statistics, U. Grenander, ed., Wiley, New York, 235-245.
- [R4] Robbins, H. and Monro, S. (1951). A stochastic approximation method. Ann. Math. Statist. 22 400-407.
- [R5] Roseberry, T.C. and Gehan, E.A. (1964). Operating characteristic curves and accept-reject rules for two and three stage screening procedures. Biometrics 20 73-84.

- [S1] Schwarz, G. (1962). Asymptotic shapes of Bayes sequential testing regions. Ann. Math, Statist. 33 224-236.
- [S2] Shah, B.V., Buehler, R.J. and Kempthorne, O. (1964). Some algorathms for minimizing a function of several variables. J. Soc. Indust. and Appl. Math. 12 74-92.
- [\$3] Sobel, M. and Netenzahl (). Infinite models in group testing. Proc. of Int. Symp. on Stat. Design and Linear Models.
- [S4] Sobel, M. and Weiss, G. (1972). Play-the-winner rule and inverse sampling for selecting the best of $k \geq 3$ binomial populations. Ann. Math. Statist. 43 1808-1826.

[85]	Sokelov, S	.N. (1963).	Continuous	planning c	f regression	exper-
	iments, I.	Theor. Pr	obability Apr	$\underline{21}$. 8 89-96	,),	

- [S6] Sokolov, S.N. (1963). Continuous planning of regression experiments II. Theor. Probability Appl. 8 298-504.
- [S7] Starr, N. and Woodroofe, M.B. (1969). Remarks on sequential point estimation. Proc. Nat. Acad. Sci. 63 285-288.
- [S8] Suzuki, Y. (1966). On sequential decision problems with delayed observations. Ann. of Statist. Math. 18 229-268.
- [S9] Swanepoel, J.W.H. and Venter, J.H. On the construction of sequential selection procedures. Unpublished paper.
- [S10] Sweat, C.W. (1970). Sequential search with discounted income, the discount function of the cell searched. <u>Ann. Math. Statist</u>. <u>41</u> 1446-1455.
- [V1] Venter, J.H. (1967). An extension of the Robbins-Monro procedure. Ann. Math. Statist. 30 181-190.
- [V2] Venter, J.H. (1967). On convergence of the Kiefer-Wolfowitz approximation procedure. Ann. Math. Statist. 38 1031-1036.
- [V3] Vogel, W. (1960). An asymptotic minimax theorem for the twoarmed bandit problem. <u>Ann. Math. Statist. 31</u> 444-451.
- [W1] Wetherill, G.B., Chen, H. and Vasudeva, R.B. (1966). Sequential estimation of quantal response curves: a new method of estimation. Biometrika 53 439-454.

- [W2] Whittle, P. (1965). Some general results in sequential design (with discussion). J. Roy. Statist. Soc. Ser. B. 27 371-394.
- [Z1] Zelen, M. (1969). Play-the-winner rule and controlled clinical trials. J. Amer. Statist. Assn. 64 131-140.