

SOME ASPECTS OF THE SEQUENTIAL DESIGN OF EXPERIMENTS

HERBERT ROBBINS

1. Introduction. Until recently, statistical theory has been restricted to the design and analysis of sampling experiments in which the size and composition of the samples are completely determined before the experimentation begins. The reasons for this are partly historical, dating back to the time when the statistician was consulted, if at all, only after the experiment was over, and partly intrinsic in the mathematical difficulty of working with anything but a fixed number of independent random variables. A major advance now appears to be in the making with the creation of a theory of the *sequential design* of experiments, in which the size and composition of the samples are not fixed in advance but are functions of the observations themselves.

The first important departure from fixed sample size came in the field of industrial quality control, with the double sampling inspection method of Dodge and Romig [1]. Here there is only one population to be sampled, and the question at issue is whether the proportion of defectives in a lot exceeds a given level. A preliminary sample of n_1 objects is drawn from the lot and the number x of defectives noted. If x is less than a fixed value a the lot is accepted without further sampling, if x is greater than a fixed value b ($a < b$) the lot is rejected without further sampling, but if $a \leq x \leq b$ then a second sample, of size n_2 , is drawn, and the decision to accept or reject the lot is made on the basis of the number of defectives in the total sample of $n_1 + n_2$ objects. The total sample size n is thus a random variable with two values, n_1 and $n_1 + n_2$, and the value of n is stochastically dependent on the observations. A logical extension of the idea of double sampling came during World War II with the development, chiefly by Wald, of sequential analysis [2], in which the observations are made one by one and the decision to terminate sampling and to accept or reject the lot (or, more generally, to accept or reject whatever statistical "null hypothesis" is being tested) can come at any stage. The total sample size n now becomes a random variable capable in principle of assuming infinitely many values, although in practice a finite upper limit on n is usually set. The advantage of sequential

An address delivered before the Auburn, Alabama, meeting of the Society, November 23, 1951, by invitation of the Committee to Select Hour Speakers for Southeastern Sectional Meetings; received by the editors December 10, 1951.

over fixed-size sampling lies in the fact that in some circumstances the judicious choice of a sequential plan can bring about a considerable reduction in the average sample size necessary to reduce the probability of erroneous decision to a desired low level. The theory of sequential analysis is still very incomplete, and much work remains to be done before optimum sequential methods become available for treating the standard problems of statistics.

The introduction of sequential methods of sampling freed statistics from the restriction to samples of fixed size. However, it is not only the sample *size* that is involved in the efficient design of an experiment. Most statistical problems met with in practice involve more than one population, and in dealing with such problems we must specify *which population* is to be sampled at each stage. An example will serve to clarify this point. Suppose we are dealing with two normally distributed populations with unknown means μ_1, μ_2 and variances σ_1^2, σ_2^2 , and that we wish to estimate the value of the difference $\mu_1 - \mu_2$. In order to concentrate on the point at issue we shall suppose that the total sample size, n , is fixed. There remains the question of how the n observations are to be divided between the two populations. If \bar{x}_1, \bar{x}_2 denote the means of samples of sizes n_1, n_2 from the two populations, then $\bar{x}_1 - \bar{x}_2$ is an unbiased estimator of $\mu_1 - \mu_2$, with variance $\sigma^2 = (\sigma_1^2/n_1) + (\sigma_2^2/n_2)$. For fixed $n = n_1 + n_2$, σ^2 is a minimum when $n_1/n_2 = \sigma_1/\sigma_2$. If the latter ratio is known in advance, all is well. If this ratio is not known, but if the sampling can be done in two stages, then it would be reasonable to draw preliminary samples of some size m from each of the two populations and to use the values so obtained to estimate σ_1/σ_2 ; the remainder of the $n - 2m$ observations could then be allocated to the two populations in accordance with the sample estimate of σ_1/σ_2 . The question then becomes, what is the best choice for m ? If m is small, no accurate estimate of σ_1/σ_2 can be made. If m is large, then the remaining $n - 2m$ observations may be too few to permit full utilization of the approximate knowledge of σ_1/σ_2 . (This kind of dilemma is characteristic of all sequential design problems.) More generally, we could consider schemes in which the observations are made one by one, with the decision as to which population each observation should come from being allowed to depend on all the previous observations; the total sample size n could be fixed or could be a random variable dependent on the observations.

Despite the total absence of theory, a notable pioneering venture in the spirit of sequential design was carried out in 1938 by Mahalanobis [3] to determine the acreage under jute in Bengal. Preliminary sur-

veys were made on a small scale to estimate the values of certain parameters, a knowledge of which was essential to the efficient design of a subsequent large scale census. In a subsequent publication [4] Mahalanobis called attention to the desirability of revising the design of any experiment as data accumulates. The question, of course, is how best to do this.

We are indebted to Wald for the first significant contribution to the theory of sequential design. His book [5] states the problem in full generality and gives the outline of a general inductive method of solution. The probability problems involved are formidable, since dependent probabilities occur in all their complexity, and explicit recipes are not yet available for handling problems of practical interest. Nevertheless, enough is visible to justify a prediction that future results in the theory of sequential design will be of the greatest importance to mathematical statistics and to science as a whole.

In what follows we shall discuss a few simple problems in sequential design which are now under investigation and which are different from those usually met with in statistical literature. Optimum solutions to these problems are not known. Still, it is often better to have reasonably good solutions of the proper problems than optimum solutions of the wrong problems. In the present state of statistical theory this principle applies with particular force to problems in sequential design.

2. A problem of two populations. Let A and B denote two statistical populations (coins, urns, manufacturing processes, varieties of seed, treatments, etc.) specified respectively by univariate cumulative distribution functions $F(x)$ and $G(x)$ which are known only to belong to some class D . We shall suppose that the expectations

$$(1) \quad \alpha = \int_{-\infty}^{\infty} x dF(x), \quad \beta = \int_{-\infty}^{\infty} x dG(x)$$

exist. *How should we draw a sample x_1, x_2, \dots, x_n from the two populations if our object is to achieve the greatest possible expected value of the sum $S_n = x_1 + \dots + x_n$?*

For example, let A and B denote two coins of unknown bias, and suppose that we are allowed to make n tosses, with the promise of getting \$1 for each head but nothing for tails. If $x_i = 1$ or 0 according as heads or tails occurs on the i th toss, then S_n denotes the total sum which we are to receive, and α and β ($0 \leq \alpha, \beta \leq 1$) are the respective probabilities of obtaining heads on a single toss of coins A and B .

As a general intuitive principle, whenever we feel pretty sure from

the results of previous observations that one of the two numbers α , β is the greater, we shall want to devote more of our future observations to that population. Note that there is no terminal decision to make; that is, we are not interested in estimating $\alpha - \beta$ or in testing the hypothesis that $\alpha = \beta$, etc., after the sample is drawn. The whole problem lies in deciding *how to draw the sample*. There certainly exist practical situations in which the present problem represents more nearly what one wants to solve than would any formulation in terms of testing hypotheses, estimating parameters, or making terminal decisions. In fact, the problem represents in a simplified way the general question of how we learn—or should learn—from past experience. A reasonably good solution of the problem must therefore be found if mathematical statistics is to provide a guide to what has been called by Neyman [6] inductive behavior.

To begin with we shall consider the special case already mentioned in which A and B are coins and the unknowns α and β are the respective probabilities of obtaining heads ($x_i = 1$) on a single trial. Let us take as an example of a possible sampling rule the following.

Rule R_1 . For the first toss choose A or B at random. Then, for $i = 1, 2, \dots$, if the i th toss results in heads, stick to the same coin for the $(i+1)$ th toss, while if the i th toss results in tails, switch to the other coin for the $(i+1)$ th toss.

What are the operating characteristics of the rule R_1 ? The successive tosses represent the evolution of a simple Markov chain with four states, (A, H) , (A, T) , (B, H) , (B, T) , and with transition probabilities which are easily written down; for example, the probability of a transition from (A, H) on the i th toss to (A, T) on the $(i+1)$ th toss is $1 - \alpha$. Let p_i denote the probability of obtaining heads on the i th toss. To avoid trivialities we shall suppose that α and β are not both 0 or both 1; then $|\alpha + \beta - 1| < 1$. It is easy to show that

$$(2) \quad p_{i+1} = (\alpha + \beta - 1)p_i + (\alpha + \beta - 2\alpha\beta),$$

from which it follows that

$$(3) \quad p_i = (\alpha + \beta - 1)^{i-1} \left[p_1 - \frac{\alpha + \beta - 2\alpha\beta}{2 - (\alpha + \beta)} \right] + \frac{\alpha + \beta - 2\alpha\beta}{2 - (\alpha + \beta)},$$

and hence that

$$(4) \quad \lim_{i \rightarrow \infty} p_i = \frac{\alpha + \beta - 2\alpha\beta}{2 - (\alpha + \beta)} = \gamma + \frac{\delta^2}{1 - \gamma},$$

where we have set

$$(5) \quad \gamma = \frac{\alpha + \beta}{2}, \quad \delta = \frac{|\alpha - \beta|}{2}.$$

It follows that in using the rule R_1 ,

$$(6) \quad \lim_{n \rightarrow \infty} E\left(\frac{S_n}{n}\right) = \lim_{n \rightarrow \infty} \left(\frac{p_1 + \cdots + p_n}{n}\right) = \gamma + \frac{\delta^2}{1 - \gamma}.$$

Now, if we *knew* which of the two numbers α, β is the greater, then by using the corresponding coin exclusively we could achieve the result

$$(7) \quad E\left(\frac{S_n}{n}\right) = \max(\alpha, \beta) = \gamma + \delta.$$

Hence it is natural to take the difference

$$(8) \quad L(A, B, R_1) = (\gamma + \delta) - \left(\gamma + \frac{\delta^2}{1 - \gamma}\right) = \delta \left[1 - \frac{\delta}{1 - \gamma}\right] \geq 0$$

as a measure of the asymptotic *loss* per toss, by a person who uses R_1 , *due to ignorance of the true state of affairs*. It is easy to show that $L(A, B, R_1)$ has its maximum value, $M_1 = 3 - 2^{3/2} \cong .172$, when $\alpha = 0$ and $\beta = 2 - 2^{1/2} \cong .586$ or vice versa. Thus a person using R_1 will, for large n , lose on the average at most 17.2 cents per toss due to ignorance of which is the better coin. On the other hand, consider the rule R_0 which consists in choosing one of the two coins A, B at random and then sticking to it, come what may (or in tossing the two coins alternately). The corresponding quantity $L(A, B, R_0)$ is easily seen to have the value $(\gamma + \delta) - \gamma = \delta$, which has its maximum, $M_0 = 1/2$, when $\alpha = 0$ and $\beta = 1$ or vice versa. Clearly, R_1 is considerably better than R_0 .

The rule R_0 makes the choice of the coin for the i th toss independent of the results of previous tosses, while R_1 makes this choice depend on the result of the $(i-1)$ th toss only. For the most general rule R the choice of the coin for the i th toss will depend on the results x_1, \cdots, x_{i-1} of all the previous tosses. For any such rule R let

$$(9) \quad L_n(A, B, R) = \max(\alpha, \beta) - E\left(\frac{S_n}{n}\right)$$

where E denotes expectation computed on the basis of α, β , and R , and let

$$(10) \quad M_n(R) = \max_{(\alpha, \beta)} [L_n(A, B, R)],$$

$$(11) \quad \phi(n) = \min_{(R)} [M_n(R)].$$

It would be interesting to know the value of $\phi(n)$ and the explicit description of any "minimax" rule R for which the value $\phi(n)$ is attained.

A much simpler problem is: do there exist rules R such that

$$(12) \quad \lim_{n \rightarrow \infty} L_n(A, B, R) = 0 \quad \text{for every } A, B?$$

We shall see in the next paragraph that the answer is yes, not only in the case of the coins but for any two populations.

Returning to the general case in which A and B are arbitrary statistical populations for which the values (1) exist, consider the sampling rule \bar{R} defined as follows: let

$$(13) \quad \begin{aligned} 1 &= a_1 < a_2 < \cdots < a_n < \cdots, \\ 2 &= b_1 < b_2 < \cdots < b_n < \cdots \end{aligned}$$

be two fixed, disjoint, increasing sequences of positive integers of density 0; that is, such that the proportion of the integers $1, 2, \dots, n$ which are either a 's or b 's tends to 0 as $n \rightarrow \infty$. We define inductively: if the integer i is one of the a 's, take the i th observation, x_i , from population A , if i is one of the b 's, take x_i from B , and if i is neither one of the a 's nor one of the b 's, take x_i from A or B according as the arithmetic mean of all previous observations from A exceeds or does not exceed the arithmetic mean of all previous observations from B . It can be shown to follow from the strong law of large numbers that with probability 1,

$$(14) \quad \lim_{n \rightarrow \infty} \frac{S_n}{n} = \max(\alpha, \beta).$$

This in turn can be shown to imply the relation

$$(15) \quad \lim_{n \rightarrow \infty} E\left(\frac{S_n}{n}\right) = \max(\alpha, \beta),$$

so that

$$(16) \quad \lim_{n \rightarrow \infty} L_n(A, B, \bar{R}) = \max(\alpha, \beta) - \lim_{n \rightarrow \infty} E\left(\frac{S_n}{n}\right) = 0$$

for any A, B such that α, β exist.

3. Some other problems of sequential design. The problem of

§2 can be generalized in various ways. For one thing, we can let the total sample size n be a random variable, either independent of the observations or dependent on them. As an example of the latter case, suppose in the problem of the two coins that we have to pay a fixed amount c for the privilege of making each toss. We may then decide to stop tossing whenever it seems pretty certain that $\max(\alpha, \beta) < c$; this amounts to a special case of a problem of three populations. We can even consider the case of a continuum of populations. Suppose we can apply a certain treatment to some plant or animal at any intensity θ in some interval, and let $F(x, \theta)$ be the cumulative distribution function of the response x to the treatment of intensity θ . The expected value

$$(17) \quad \alpha(\theta) = \int_{-\infty}^{\infty} x dF(x, \theta),$$

the "regression" of x on θ , is assumed to be unknown. Let $\{\theta_i\}$ denote any sequence of θ values, chosen sequentially by the experimenter, and let $\{x_i\}$ denote the corresponding sequence of responses, so that each x_i has the distribution $\Pr [x_i \leq x] = F(x, \theta_i)$. (I) Suppose $\alpha(\theta)$ has a unique maximum at some unknown point θ_0 . How should the experimenter choose the sequence $\{\theta_i\}$ in order to maximize the expected value of the sum $S_n = x_1 + \dots + x_n$ or, alternatively, in order to estimate the value of θ_0 ? (II) Suppose $\alpha(\theta)$ is an increasing function of θ which takes on a given value α_0 at some unknown point θ_0 . How should the experimenter choose the sequence $\{\theta_i\}$ in order to estimate the value of θ_0 ? Problem I is the problem of the experimental determination of the maximum of a function when the observations are subject to a random error; Problem II is fundamental in sensitivity testing and bioassay.

It is clear that in both of these problems the choice of each θ_i should be made to depend on the responses x_1, \dots, x_{i-1} at the previous levels $\theta_1, \dots, \theta_{i-1}$ of the treatment, so that we are dealing with problems of sequential design. The non-sequential study of Problem I was initiated by Hotelling [7] (see also [8]), but no sequential theory has yet been published. Problem II has been considered by Robbins and Monro [9]; their method is as follows. Let $\{a_n\}$ be a sequence of positive constants such that

$$(18) \quad \sum_1^{\infty} a_n^2 < \infty, \quad \sum_1^{\infty} a_n = \infty,$$

let θ_1 be arbitrary, and set

$$(19) \quad \theta_{n+1} = \theta_n + a_n(\alpha_0 - x_n) \quad (n = 1, 2, \dots).$$

Then under certain mild restrictions on $F(x, \theta)$ it can be shown that

$$(20) \quad \lim_{n \rightarrow \infty} \theta_n = \theta_0 \quad \text{in probability.}$$

In this and other problems, any sequential design with reasonably good properties is likely to find an appreciative audience. This will encourage the use of random sampling methods to find empirical approximations to the operating characteristics of sequential designs when a full mathematical solution is difficult. An empirical study of the rapidity of convergence in (20) has been made by Teichrow [10].

4. The problem of optional stopping. To fix the ideas, let x be normally distributed with unknown mean θ and unit variance. Suppose we wish to test the null hypothesis, H_0 , that $\theta = 0$ against the alternative, H_1 , that $\theta > 0$. The standard statistical test based on a sample of fixed size n is the following. Let $S_n = x_1 + \dots + x_n$ and reject H_0 in favor of H_1 if and only if

$$(21) \quad S_n > \alpha n^{1/2},$$

where α is some constant. The probability of rejecting H_0 when it is true will then be

$$(22) \quad \epsilon(\alpha) = 1 - \Phi(\alpha),$$

where we have set

$$(23) \quad \Phi(x) = \frac{1}{(2\pi)^{1/2}} \int_{-\infty}^x e^{-t^2/2} dt,$$

and by choosing α large we can make $\epsilon(\alpha)$ as small as we please. For example, if $\alpha = 3.09$ then $\epsilon(\alpha) \cong .001$.

Suppose now that H_0 is true but that an unscrupulous experimenter wishes to get an unwary statistician to reject it. If the sample size n has not been agreed upon in advance the experimenter could adopt the technique of stopping the sampling as soon as the inequality (21) is verified. The law of the iterated logarithm of probability theory implies that with probability 1 the inequality (21) will hold for infinitely many values of n if the sampling continues indefinitely, no matter how large the value of α . Hence the experimenter is "sure" to come eventually to a value of n for which (21) holds, and by stopping the experiment at this point he will cause the statistician to reject H_0 even though it is true. This fact immediately vitiates the use of (21) as a test of H_0 if there is any possibility that optional stopping may be involved.

The simplest way for the statistician to guard against the effect

of optional stopping is to insist that the size of the sample be fixed in advance of the experimentation. Such a restriction would often be too rigid for practical use. The statistician might therefore content himself with setting limits $n_1 \leq n \leq n_2$ on the sample size which will be flexible enough to meet the contingencies of experimentation but narrow enough to eliminate the worst effects of optional stopping. To this end the statistician would like to know the value of the function

$$(24) \quad g(n_1, n_2, \alpha) = \Pr [S_n > \alpha n^{1/2} \text{ for some } n_1 \leq n \leq n_2],$$

where the x_i are independent and normal $(0, 1)$. It is quite easy to establish the inequality

$$(25) \quad g(n_1, n_2, \alpha) < \frac{1 - \Phi(\alpha)}{1 - \Phi\left(\alpha \cdot \frac{\lambda^{1/2} - 1}{(\lambda - 1)^{1/2}}\right)}, \text{ where } \lambda = \frac{n_2}{n_1},$$

which is useful when λ is not too large, and sharper inequalities can no doubt be devised.

The problem of optional stopping has received little attention in statistical theory. (See, however, [11], especially pp. 286–292.) One need not assume that the experimenter is consciously trying to deceive the statistician—the two are often the same person—to recognize the desirability of devising methods of statistical analysis which would be relatively insensitive to the effects of optional stopping.

BIBLIOGRAPHY

1. H. F. Dodge and H. G. Romig, *A method of sampling inspection*, Bell System Technical Journal vol. 8 (1929) pp. 613–631, and *Single sampling and double sampling inspection tables*, *ibid.* vol. 20 (1941) pp. 1–61.
2. A. Wald, *Sequential analysis*, New York, Wiley, 1947.
3. P. C. Mahalanobis, *A sample survey of the acreage under jute in Bengal with discussion of planning of experiments*, Snakhyā vol. 4 (1940) pp. 511–531.
4. ———, *On large-scale sample surveys*, Philos. Trans. Roy. Soc. London, Ser. B vol. 231 (1944) pp. 329–451.
5. A. Wald, *Statistical decision functions*, New York, Wiley, 1950.
6. J. Neyman, *First course in probability and statistics*, New York, Holt, 1950.
7. H. Hotelling, *Experimental determination of the maximum of a function*, Ann. of Math. Statist. vol. 12 (1941) pp. 20–45.
8. M. Friedman and L. J. Savage, *Planning experiments seeking maxima*, Chap. 3 of *Selected techniques of statistical analysis*, New York, McGraw-Hill, 1947.
9. H. Robbins and S. Monroe, *A stochastic approximation method*, Ann. of Math. Statist. vol. 22 (1951) pp. 400–407.
10. D. Teichrow, unpublished.
11. W. Feller, *Statistical aspects of ESP*, Journal of Parapsychology vol. 4 (1940) pp. 271–298.

UNIVERSITY OF NORTH CAROLINA