The Annals of Statistics
1982, Vol. 10, No. 2, 634–636

OPTIMAL STOPPING REGIONS WITH ISLANDS AND PENINSULAS

BY DONALD A. BERRY[1] AND PECHENG WANG

University of Minnesota

An urn contains a known number of balls, an unknown number R of which are red. Sequential sampling with replacement is possible and cost is proportional to sample size. The objective is to estimate R with 0–1 loss, given that a priori R has a discrete uniform distribution. It is shown that optimal stopping regions may be disconnected and composed of islands and peninsulas.

1. Introduction. In most sequential sampling problems that can be stated simply, the stopping and continuation regions are connected sets. We present a problem which can have a disconnected stopping region. An unusual example of the problem is described in Section 2.

An urn contains a known finite number N of balls; an unknown number R of the balls are red. The prior distribution for R is uniform on $\{0, 1, \cdots, N\}$. We are allowed to sample sequentially from the urn with replacement; Berry (1974) considers a similar problem for sampling without replacement. The present objective is to guess R with loss 0 if correct, 1 if not. Each observation costs c, and loss and sampling cost are assumed to be additive. An equivalent formulation is in terms of coin-tossing: the prior for the probability of heads is uniform on $\{0, 1/N, 2/N, \cdots, 1\}$, but the loss structure is much less natural in this formulation.

A terminal Bayes rule is any mode of the current posterior distribution of R which, given r red balls in a sample of n, is

$$P(R = j \mid r, n) = j^r(N-j)^{n-r}/\textstyle\sum_{i=0}^{N} i^r(N-i)^{n-r}, \quad j = 0, 1, \cdots, n. \tag{1.1}$$

The mode of this posterior is an integer adjacent to the maximum likelihood estimate Nr/n of R. What is of interest here is the nature of optimal sampling rules.

This is a typical problem in sequential decision theory and can be solved using dynamic programming, or backward induction, provided optimal rules are bounded. However, the usual techniques for finding a bound (Ray, 1965) do not apply; in fact, in general an optimal rule is not bounded. To see this, suppose c is very small and N is moderate in size; take N to be odd for convenience. Suppose the sample gives red and non-red in alternation and indefinitely. Of course, as $n \to \infty$, $P(|r/n - 1/2| < \varepsilon) \to 0$ since N is odd, but obtaining red on every other draw has positive probability for n finite and so must be reckoned with. Along this alternating sequence, for sufficiently large n, the probabilities of $R = (N-1)/2$ and $(N+1)/2$ both become greater than $\frac{1}{2} - \varepsilon$ for any positive ε when n is even and the larger probability is between $\frac{1}{2}(1 + N^{-1}) - \varepsilon$ and $\frac{1}{2}(1 + N^{-1})$ when n is odd. In both cases it is optimal to continue sampling for sufficiently small c, so that sampling can continue indefinitely. (Of course, an intelligent sampler would soon come to question the original assumptions.)

The above discussion applies whenever the sample proportion of reds hovers midway between j/N and $(j+1)/N$. When the sample proportion gets sufficently close to j/N for n large, the probability (1.1) that $R = j$ becomes greater than $1 - c$ and so sampling should not continue—perhaps it should have stopped previously. Therefore, for sufficiently small

Received July 1981; revised November 1981.

[1] This author's research supported by NSF Grant No. MCS81-02477.

AMS 1980 *subject classification.* Primary 60L15; secondary 62C10.

Key words and phrases. Dichotomous populations, sequential sampling, sampling with replacement, stopping regions, stopping island, stopping peninsulas, discrete uniform prior, posterior probabilities, maximum likelihood estimation.

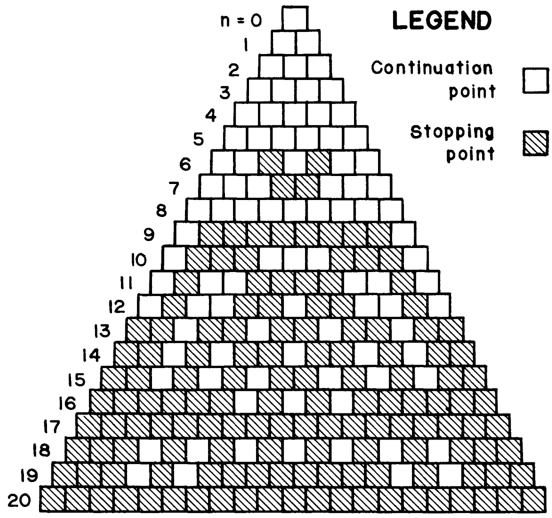

FIG 1. *The optimal stopping rule when* $N = 9$, $c = 0.02$, *and sampling is truncated at* 20 *observations. Each box corresponds to a possible state of information and the boxes in a row correspond to* $r = 0$ *to* $r = n$.

c, the optimal continuation region has N "fingers" which extend indefinitely in the direction of n, and complementary stopping fingers or peninsulas.

Though sampling may continue indefinitely, the probability of being in one of these continuation fingers tends to 0 as n increases. Therefore, there is always an ε-optimal rule that is truncated for any positive ε.

The example presented in the next section does not show these continuation fingers since c is moderately large. It is presented because its stopping region is so unusual.

2. An example. Let $S(r, n)$ denote the stopping risk and $W(r, n)$ the optimal risk after r reds and $n - r$ nonreds. Then

$$S(r, n) = 1 - \max_j P(R = j \mid r, n) + nc,$$

where the posterior probabilities are given by (1.1). When the optimal rule (selection of n) is bounded by M^* then it can be found as follows. For $M \geq M^*$ and $r = 0, 1, \cdots, M$, set

$$W(r, M) = S(r, M). \tag{2.1}$$

Then, for $n = M - 1, M - 2, \cdots, 0$ and $r = 0, 1, \cdots, n$, calculate

$$W(r, n) = \min\{S(r, n), C(r, n)\} \tag{2.2}$$

where C is the continuation risk

$$C(r, n) = \frac{r+1}{n+2} W(r+1, n+1) + \left(1 - \frac{r+1}{n+2}\right) W(r, n+1).$$

Keeping track of whether $W = S$ or $W = C$ (or both) gives all optimal rules. When $M^* = \infty$ then, as indicated in the previous section, solving this system using a particular M gives at best an approximation to an optimal rule and to $W(0, 0)$, the minimum Bayes risk.

By way of example, suppose $N = 9$ and $c = 0.02$. Figure 1 gives the optimal stopping rule when sampling is truncated at $M = 20$. The top-most box corresponds to the starting point. With each observation, the process moves to the next lower row of boxes, the box on the right if the ball is red and the box on the left otherwise. Boxes in which stopping is optimal are shaded. We find $W(0, 0) = 0.728$.

We have checked that the rule given in Figure 1 is actually optimal among all rules, truncated or not, using the following device. Solving (2.1) and (2.2) for any $M \geq 0$ gives lower bounds on the various $W(r, n)$. By replacing (2.1) with

$$W(r, M) = Mc$$

(which would be obtained if an outside source reveals the value of R after M observations) and then solving (2.2) gives upper bounds on the various $W(r, n)$. (This implies that if stopping is optimal in this modified problem it is optimal in the original problem as well.) If the upper and lower bounds and the implicit rules agree (for those (r, n) which can be reached) then the optimal rule is bounded and it is the rule common to both variations. We find this to be the case in our example for $M = 50$ and therefore the optimal rule is the one given in Figure 1.

There are many continuation points in Figure 1 that cannot be reached following the indicated optimal procedure. But the most interesting feature of the procedure is the stopping island between $n = 6$ and $n = 7$. Presumably, one could find values of N and c for which there are many such stopping islands. Stopping islands can also be present when c is sufficiently small for there to be infinite continuation fingers and stopping peninsulas.

3. Conclusion. When sampling exchangeable Bernoulli variables sequentially in an estimation problem, the optimal stopping region can be disconnected and can appear quite strange if the prior distribution is discrete—population size known and finite. The problem considered here assumes 0–1 loss but it is the discreteness of the prior distribution and not the form of the loss function which causes these remarkable stopping regions.

REFERENCES

BERRY, D. A. (1974). A class of optimal stopping problems for sampling without replacement. *Biometrika* **61** 361–368.

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

SCHOOL OF STATISTICS
UNIVERSITY OF MINNESOTA
206 CHURCH ST., S.E.
MINNEAPOLIS, MN 55455