SEQUENTIAL DESIGN OF EXPERIMENTS

By Herman Chernoff[1]

Stanford University

1. Introduction. Considerable scientific research is characterized as follows. The scientist is interested in studying a phenomenon. At first he is quite ignorant and his initial experiments are preliminary and tentative. As he gathers relevant data, he becomes more definite in his impression of the underlying theory. This more definite impression is used to construct more informative experiments. Finally after a certain point he is satisfied that his evidence is sufficient to allow him to announce certain conclusions and he does so.

While this sequential searching for relevant and informative experiments is common, very little statistical theory has been directed in this direction. The general problem may reasonably be called that of sequential design of experiments. A truncated variation of this problem called the two-armed bandit problem has attracted some attention (see [1] and [5]). Up to now an optimal solution for the two-armed bandit problem has not been attained. The failure to solve the two-armed bandit problem and certain obvious associated results indicate strongly that while optimal strategies are difficult to characterize, *asymptotically* optimal results should be easily available. Here the term asymptotic refers to large samples. For the sequential design problems, large samples and small cost of experimentation are roughly equivalent.

In this paper we present a procedure for the sequential design of experiments where the problem is one of testing a hypothesis. Formally, we assume that there are two possible actions (terminal decisions) and a class of available experiments. After each observation, the statistician decides on whether to continue experimentation or not. If he decides to continue, he must select one of the available experiments. If he decides to stop he must select one of the two terminal actions.

For the special case where there are only a finite numer of states of nature and a finite number of available experiments this procedure will be shown to be "asymptotically optimal" as the cost of sampling approaches zero. The procedure can be partially described by saying that at each stage the experimenter acts as though he is almost convinced that $\hat{\theta}$, the current maximum likelihood estimate of the state of nature, is actually equal to or very close to the true state of nature.

In problems were the cost of sampling is not small, this procedure may leave something to be desired. More specifically, until enough data are accumulated, the procedure may suggest very poor experiments because it does not sufficiently distinguish between the cases where $\hat{\theta}$ is a poor estimate and where $\hat{\theta}$ is a good

Received July 5, 1958.

[1] This work was sponsored by the Office of Naval Research under Contract N6 onr-25140 (NR-342-022).