TESTING HYPOTHESES FOR MARKOV CHAINS WHEN THE PARAMETER SPACE IS FINITE

By Ken-ichi Yoshihara

1. Summary.

In Chernoff [2], a procedure was presented for the sequential design of experiments where the problem was one of testing a hypothesis. When there were only a finite number of states of nature and a finite number of available experiments, the procedure was shown to be "asymptotically optimal" as the cost of sampling approached zero. An analogous procedure can be applied to the problem of testing a hypothesis with respect to a Markov process, and this procedure will also be shown to be "asymptotically optimal".

2. Assumptions.

Let Θ be a parameter space which consists of finite elements. We shall test the hypothesis $\theta \in H_1$ against the alternative $\theta \in H_2$ where H_1 and H_2 are two non-null disjoint subsets of Θ. In what follows, we make the following assumptions.

A 1. Θ is a finite set. H_1 and H_2 are two non-null subsets of Θ and $H_1 \cup H_2 = \Theta$ and $H_1 \cap H_2 = \phi$.

A 2. $\{X_k: k=0, 1, 2, \cdots\}$ is a Markov chain with state space $(\mathfrak{X}, \mathfrak{A})$ and with stationary transition measures $p_\theta(\xi, \cdot)$ $(\theta \in \Theta)$ which satisfy the following conditions:

(a) For each $\theta \in \Theta$, the transition measures $p_\theta(\xi, \cdot)$ satisfy Doeblin's condition (D) in Doob [3], and there exists only one ergodic set and the transient set is empty;

(b) For each $\theta \in \Theta$, the transition measures $p_\theta(\xi, \cdot)$ admit of a unique stationary probability measure $p_\theta(\cdot)$.

In what follows $E_\theta(\cdot)$ will denote an expected value computed under the assumption that θ is true and that $p_\theta(\cdot)$ is the initial distribution.

A 3. There is a measure λ on $\mathfrak{A}$, not necessarily finite, with respect to which all the transition measures $p_\theta(\xi, \cdot)$ have densities $f(\xi, \eta: \theta)$ and the initial distribution $p_\theta(\cdot)$ has a density $f(\xi: \theta)$ for each θ. These densities satisfy the following conditions:

(a) For each $\theta, f(\xi, \eta: \theta)$ is measurable in ξ and η, and $f(\xi: \theta)$ is measurable in ξ;

(b) If θ, φ are in Θ and $\theta \neq \varphi$, then

$$P_\theta\{f(X_0, X_1: \theta) \neq f(X_0, X_1: \varphi) \mid X_0 = \xi\} > 0 \qquad \text{for almost all } \xi.$$

Received March 15, 1963.