

COMPUTING OPTIMAL SEQUENTIAL ALLOCATION RULES IN CLINICAL TRIALS*

Michael N. Katehakis

State University of New York at Stony Brook

and

Cyrus Derman

Columbia University

The problem of assigning one of several treatments in clinical trials is formulated as a discounted bandit problem that was studied by Gittins and Jones. The problem involves comparison of certain state dependent indices. A recent characterization of the index is used to calculate more efficiently the values of these indices.

1. Introduction.

We consider the well known problem of optimal allocation of treatments in clinical trials. A simple version of the problem is as follows. There are several possible treatments for a given disease. When a particular treatment n is used it is either effective with unknown probability θ_n or not effective with probability $1 - \theta_n$. The problem is to find a sequential sampling procedure which maximizes a measure of the expected total number of treatment successes. When the planning horizon is infinite, prior distributions are assigned to the unknown parameters, and one takes the expected total discounted number of successes as the relevant measure of performance of a sequential sampling procedure, the problem can be put into the form of a discounted version of the

* Work supported by USAF Contract AFOSR 840136, NSF Grants DMS-84-05413, ECS-85-07671 and ONR Contract N00014-84-K-0244.

AMS 1980 subject classifications: Primary 62L15; Secondary 60J10.

Key words and phrases: dynamic programming, multi armed bandit problem, sequential allocations, successive approximations.