

AD-A170 981


COMPUTING OPTIMAL SEQUENTIAL ALLOCATION RULES IN  
CLINICAL TRIALS. (U) STATE UNIV OF NEW YORK AT STONY  
BROOK DEPT OF APPLIED MATHEMA. M N KATEHAKIS ET AL.  
23 SEP 85 AHS-85-39 AFOSR-TR-86-0521 F/G 6/5

1/1

UNCLASSIFIED

NL





MICROCOPY RESOLUTION TEST CHART  
NATIONAL BUREAU OF STANDARDS-1963-A

State University  
of New York at  
**Stony Brook**



AD-A170 801

Department of  
Applied  
Mathematics  
and  
Statistics

DTIC  
ELECTE  
AUG 13 1986

S

D

B

Approved for public release,  
distribution unlimited

Computing Optimal Sequential Allocation  
Rules In Clinical Trials

by

Michael N. Katchakis and Cyrus Derman  
SUNY at Stony Brook Columbia Univ.

Report # AMS-85-59

DTIC FILE COPY

DTIC  
ELECTE  
S AUG 13 1986 D  
B

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (AFOSR)  
SOURCE IS UNCLASSIFIED TO DTIC  
This technical report has been reviewed and is  
approved for public release IAW AFR 190-12.  
Distribution is unlimited.  
MATTHEW J. KENNER  
Chief, Technical Information Division

**Computing Optimal Sequential Allocation  
Rules In Clinical Trials**

by

Michael N. Katehakis and Cyrus Derman  
SUNY at Stony Brook Columbia Univ.

Report # AMS-85-59

**DISTRIBUTION STATEMENT A**

Approved for public release  
Distribution Unlimited

## REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION <b>UNCLASSIFIED</b>		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S) <b>AMS-85-59</b>		5. MONITORING ORGANIZATION REPORT NUMBER(S) <b>AFOSR-TR- 86-0521</b>	
6a. NAME OF PERFORMING ORGANIZATION <b>S.U.N.Y. at Stony Brook</b>	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION <b>Air Force Office of Scientific Research</b>	
6c. ADDRESS (City, State and ZIP Code) <b>Dept. of Applied Mathematics and Statistics S.U.N.Y., Stony Brook, N.Y.11794</b>		7b. ADDRESS (City, State and ZIP Code) <b>Directorate of Mathematical &amp; Information Sciences, Bolling AFB DC 20332-6448</b>	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION <b>AFOSR</b>	8b. OFFICE SYMBOL (If applicable) <b>NM</b>	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER <b>AFOSR-84-0136</b>	
8c. ADDRESS (City, State and ZIP Code) <b>Bolling AFB DC 20332-6448</b>		10. SOURCE OF FUNDING NOS	
		PROGRAM ELEMENT NO <b>61102F</b>	PROJECT NO <b>2304</b>
		TASK NO. <b>H5</b>	WORK UNIT NO.
11. TITLE (Include Security Classification) <b>COMPUTING OPTIMAL SEQUENTIAL ALLOCATION RULES IN CLINICAL TRIALS.</b>			
12. PERSONAL AUTHOR(S) <b>M. N. Katehakis and C. Derman</b>			
13a. TYPE OF REPORT <b>Research Report</b>	13b. TIME COVERED FROM <b>7/84</b> TO <b>9/85</b>	14. DATE OF REPORT (Yr. Mo., Day) <b>1985, September, 23</b>	15. PAGE COUNT <b>10</b>
16. SUPPLEMENTARY NOTATION <b>Also supported by N.S.F. Grants DMS-84-05413 , ECS-85-07671 and ONR contract N00014-84-K-0244.</b>			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB GR	
		<b>Sequential Sampling, Bayes Rules, Clinical Trials.</b>	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)			
<p>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.</p>			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS <input type="checkbox"/>		21. ABSTRACT SECURITY CLASSIFICATION <b>UNCLASSIFIED</b>	
22a. NAME OF RESPONSIBLE INDIVIDUAL <b>Brian Woodruff</b>	22b. TELEPHONE NUMBER (Include Area Code) <b>(202) 767- 5027</b>	22c. OFFICE SYMBOL <b>NM</b>	

# Computing Optimal Sequential Allocation Rules In Clinical Trials\*

by

Michael N. Katehakis and Cyrus Derman  
SUNY at Stony Brook      Columbia University

## Abstract

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  $\theta_n$  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 bandit problem treated successfully by Gittins and Jones (1974). The original formulation of the multi armed bandit problem and the sequential clinical trials problem is due to Robbins (1952). Gittins and Jones showed that there is an index associated with each state of each bandit

---

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

such that an optimal procedure always uses the bandit with the largest current index value. Recently, Katchakis and Veinott (1985) have obtained a new characterization of the index which allows the index to be more easily calculated. The purpose of this paper is to illustrate the calculation of the index in the context of the clinical trials problem using this new characterization.

2. Computing Dynamic Allocation Indices: Suppose  $N$  treatments are available for treating patients with a certain disease. Let  $Y_n(k) = 1$  ( $Y_n(k) = 0$ ) denote the outcome that the  $n^{\text{th}}$  treatment has been successful (unsuccessful) the  $k^{\text{th}}$  time it is used. At times  $t = 1, 2, \dots$ , based on past observations, one has to decide which treatment to allocate to the next patient. At the start of the experiment we assume that  $\theta_n$  is a random variable with Beta prior density with parameter vector  $(a_n, b_n)$ ; i.e.,  $\theta_n$  has the prior density

$$(1) \quad g_n(\theta) = \Gamma(a_n)\Gamma(b_n)\{\Gamma(a_n + b_n)\}^{-1}\theta^{a_n-1}(1-\theta)^{b_n-1}, \quad \forall \theta \geq 0,$$

where in (1)  $a_n, b_n$  are strictly positive constants. Furthermore, we assume that  $\theta_1, \dots, \theta_n$  are independent. If after  $k$  trials using treatment  $n$  we let  $x_n(k) = (s_n(k), f_n(k))$ , where  $s_n(k)$  ( $f_n(k)$ ) denotes the number of successes (the number of failures) then, the posterior density of  $\theta_n$  given  $x_n(k)$  is also Beta with parameter vector  $(a_n + s_n(k), b_n + f_n(k))$ . Thus, the information obtained during the first  $k$  trials from treatment  $n$  is summarized by  $x_n(k)$ . Furthermore,  $\{x_n(k), k \geq 1\}$  is a Markov chain on  $S = \{(s, f), s, f = 0, 1, 2, \dots\}$  with transition probabilities given by

$$(2) \quad \begin{aligned} P(x_n(k+1) = (s+1, f) \mid x_n(k) = (s, f)) \\ &= 1 - P(x_n(k+1) = (s, f+1) \mid x_n(k) = (s, f)) \\ &= P(Y_n(k+1) = 1 \mid x_n(k) = (s, f)) \\ &= \frac{a_n + s}{a_n + b_n + s + f}. \end{aligned}$$



SEARCHED  
SERIALIZED  
INDEXED  
FILED  
MAY 1985  
FBI - MEMPHIS  
A-1

The problem is to determine a policy  $\pi$  which maximizes the expected discounted number of successes; i.e., to maximize  $w(\pi, \alpha)$

$$(3) \quad w(\pi, \alpha) = \int \dots \int E \left( \sum_{t=1}^{\infty} \alpha^{t-1} Y_{\pi(t)} \right) g_1(d\theta_1) \dots g_N(d\theta_N),$$

where  $Y_{\pi(t)}$  is  $Y_n(k)$  if at time  $t$  treatment  $\pi(t)$  is used for the  $k$ th time and  $\alpha \in (0, 1)$  is a discount factor. An interpretation of the discount factor  $\alpha$  is that  $1 - \alpha$  is the probability that at any given time the entire experiment will be terminated. Stated otherwise, there are  $N$  Markov chains; the problem is to sequentially activate one of them, leaving the others inactive, in order to maximize the expected total discounted reward. In this case the expected reward at any time is the expected posterior probability of success associated with the state of the activated Markov chain; i.e., if the  $n$ th chain is activated for the  $k$ th time when  $x_n(k) = (s, f)$ , then, the corresponding reward is

$$(4) \quad r_n(s, f) = E(Y_n(k+1) \mid x_n(k) = (s, f)) \\ = \frac{a_n + s}{a_n + b_n + s + f}.$$

Within the context of this formulation, Gittins and Jones (1974) showed that this problem can be reduced to  $N$  one dimensional problems. Each of the latter problems involves a single Markov chain and its solution is the calculation of a dynamic allocation index  $m_n(s, f)$  associated with the current state  $(s, f)$  of the Markov chain. Then, at each point of time an optimal policy for the original problem is such that it activates the chain with the largest current index value. Based on an earlier characterization of  $(1-\alpha)^{-1} m_n(s, f)$ , Gittins and Jones (1979) used an algorithm for computing optimal policies. Recently, Katehakis and Veinott (1985) have obtained a different characterization of the index. This characterization casts the calculation of the index into the form of a familiar replacement problem, e.g., see Derman



(1970, pp. 121). Namely, if  $C$  is the class of policies  $R$  for controlling  $\{x_n(k), k \geq 1\}$  by either allowing it to continue or to instantaneously restart it at its initial state  $x_n(1) = (s, f)$ , then,

$$(5) \quad m_n(s, f) = \sup_R \{E_R \left( \sum_{k=1}^{\infty} \alpha^{k-1} r_n(x_n(k)) \mid x_n(1) = (s, f) \right)\} .$$

We next show that (5) can be used to evaluate  $m_n(s, f)$  with sufficient accuracy. In the sequel we will be concerned with a single treatment; for notational simplicity we will drop the subscript  $n$ . Since computing  $m(s, f)$  is essentially the same as computing  $m(0, 0)$  - it only involves changing the prior vector from  $(a, b)$  to  $(a + s, b + f)$  - it suffices, without loss of generality, to discuss only the computation of  $m(0, 0)$ . It is well known that solving (5) for a fixed initial state  $(0, 0)$  involves solving the dynamic programming equations

$$(6) \quad V(s, f) = \max \left\{ \frac{a}{a+b} + \alpha \left[ \frac{a}{a+b} V(1, 0) + \frac{b}{a+b} V(0, 1) \right] , \right. \\ \left. \frac{a+s}{a+b+s+f} + \alpha \left[ \frac{a+s}{a+b+s+f} V(s+1, f) + \frac{b+f}{a+b+s+f} V(s, f+1) \right] \right\} ,$$

$$\forall (s, f) \in S .$$

That the above equation (6) is for computing  $m(0, 0)$  is reflected in the appearance of the terms  $V(1, 0)$  and  $V(0, 1)$  in the right side of (6). Given the solution  $\{V(s, f), \forall (s, f) \in S\}$  of (6) then  $m(0, 0) = V(0, 0)$ .

Equation (6) is of the form  $V(s, f) = T_{s,f}V$  or equivalently

$$(7) \quad V = TV$$

where in (7)  $V$  is the vector of values  $\{V(s, f)\}$  and  $T$  is a contraction operator on a complete metric space. Thus, it has a unique bounded solution.

In computing the solution of (7) we consider the finite subset

$$S_L = \{(s, f) \in S : s + f \leq L\} \text{ and the two systems of equations}$$

$$(8a) \quad u_L(s, f) = T_{sf} u_L, \quad , \text{ if } s + f \neq L,$$

$$(8b) \quad u_L(s, f) = \frac{a + s}{a + b + s + f} \frac{1}{1 - \alpha}, \quad , \text{ if } s + f = L,$$

$$(9a) \quad U_L(s, f) = T_{sf} U_L, \quad , \text{ if } s + f \neq L,$$

$$(9b) \quad U_L(s, f) = \frac{1}{1 - \alpha}, \quad , \text{ if } s + f = L.$$

We will use the following more compact notation for (8) and (9)

$$(8c) \quad u_L = T_1 u_L,$$

$$(9c) \quad U_L = T_2 U_L.$$

The transformations  $T$ ,  $T_1$ ,  $T_2$  are monotone contractions, thus, successive approximations will converge to their unique fixed points for any initial points  $v^{(0)}$ ,  $u_L^{(0)}$ ,  $U_L^{(0)}$ . That is,

$$(10) \quad \lim_{n \rightarrow \infty} v^{(n)} = \lim_{n \rightarrow \infty} T v^{(n-1)} = v,$$

$$(11) \quad \lim_{n \rightarrow \infty} u_L^{(n)} = \lim_{n \rightarrow \infty} T_1 u_L^{(n-1)} = u_L,$$

$$(12) \quad \lim_{n \rightarrow \infty} U_L^{(n)} = \lim_{n \rightarrow \infty} T_2 U_L^{(n-1)} = U_L.$$

Moreover, if the points  $v^{(0)}$ ,  $u_L^{(0)}$ ,  $U_L^{(0)}$  are chosen propitiously, the convergence in (10), is from below or above as desired and from below (above) in (11) ((12)).

An algorithm to compute  $v(0,0)$  based on (10) involves an infinite number of variables; however, propositions 1 and 2, below, allow us to use (11) and (12) which involve only a finite number of variables. The proof of proposition 1 is easy and it is omitted.

**Proposition 1:** For equations (7), (8) and (9) we have

$$(13) \quad \frac{a + s}{a + b + s + f} (1 - \alpha)^{-1} \leq v(s, f) \leq (1 - \alpha)^{-1} \quad \text{for all } (s, f) \in S,$$

and

$$(14) \quad u_L(s,f) \leq V(s,f) \leq U_L(s,f) , \text{ for all } (s,f) \text{ such that } s + f \leq L .$$

**Proposition 2:** For any  $\epsilon > 0$  there exist an  $L_0 = L(\epsilon)$  such that

$$(15) \quad U_L(0,0) - u_L(0,0) \leq \epsilon , \text{ for all } L \geq L_0 .$$

**Proof:** Because of (14) it suffices to show that for any positive constants  $\epsilon_1$  and  $\epsilon_2$  there exist  $L_1 = L(\epsilon_1)$  and  $L_2 = L(\epsilon_2)$  such that

$$(16) \quad U_L(0,0) - V(0,0) \leq \epsilon_1 , \text{ for all } L \geq L_1 ,$$

and

$$(17) \quad V(0,0) - u_L(0,0) \leq \epsilon_2 , \text{ for all } L \geq L_2 .$$

We only prove (16) since the proof of (17) is analogous.

If we take  $U_L^{(0)} = V^{(0)} = (1 - \alpha)^{-1}$  in (10) and (12) then, for any  $L$  and all  $n \leq L$  we obtain that

$$(18) \quad U_L^{(n)}(0,0) = V^{(n)}(0,0) ,$$

and the convergence in (10), (12) is from above; thus, using (10) and the fact that  $V(s,f) \geq 0$  we have

$$(19) \quad V^{(n)}(0,0) - V(0,0) \leq \alpha^n \sup_{(s,f)} \{ V^{(0)}(s,f) - V(s,f) \} \leq \alpha^n (1 - \alpha)^{-1} .$$

It follows from (18), (19) that for any  $L$  and for all  $n \leq L$

$$(20) \quad U_L^{(n)}(0,0) - V(0,0) \leq \alpha^n (1 - \alpha)^{-1} .$$

Similar arguments using (12) imply that for all  $n \geq 1$

$$(21) \quad U_L^{(n)}(0,0) - U_L(0,0) \leq \alpha^n (1 - \alpha)^{-1} .$$

Thus, using (20) and (21) it is now easy to complete the proof of (16).

**Remark:** It was assumed that each clinical trial resulted either in a success or in a failure. The methodology described here extends straightforwardly to the case where the outcome of a trial can be classified into  $c$ ,  $c \geq 2$ , classifications. Then the parameter  $\theta_n$ , is a vector  $(\theta_n^1, \dots, \theta_n^c)$  where  $\theta_n^i$

is the probability of the trial resulting in the  $i^{\text{th}}$  classification. The Beta prior is replaced by a Dirichlet prior and the state space becomes  $S = \{(s_1, \dots, s_c) , s_i = 0, 1, \dots\}$  , where  $s_i$  denotes the number of trials resulting in classification  $i$  ( $1 \leq i \leq c$ ) . The reward is a given function of the classification; see ,also, Glazebrook (1978).

3. Computations: For a given  $(a,b)$  in order to compute  $m(0,0) = V(0,0)$  we use transformations  $T_1$  and  $T_2$  starting from

$$u_L^{(0)}(s,f) = \frac{a+s}{a+b+s+f} \frac{1}{1-\alpha} ,$$

and

$$U_L^{(0)}(s,f) = \frac{1}{1-\alpha} .$$

We choose  $L$  sufficiently large according to proposition 1 and iterate until the difference:  $U_L^{(n)}(0,0) - u_L^{(n)}(0,0)$  is small . We, then, take as our approximation to  $V(0,0)$  the mid point of the final interval.

Since there is always an error in computing the indices, the possibility of not using an optimal policy always exists. In our context, here, this can be overcome by doing enough computations to guarantee that in computing the indices the bounding intervals do not overlap. However in general, Katchakis and Veinott (1985) have shown that if the computed indices are close to the exact indices then the expected discounted return of the policy based on the computed indices will be close to the optimal expected discounted return.

In the following tables the results of some calculations are tabulated. There is a separate table for each value of  $\alpha = .5, .75, .9$  . An entry in cell  $(a+s,b+f)$  is the index for a treatment having prior  $(a,b)$  and in state  $(s,f)$  .

Note that the numbers in table 2 (for  $a+s, b+f = 1, 2, \dots, 5$ ) are consistent with those published by Gittins and Jones (1979).

**Table 1 ( $\alpha = .5$ )**

b+f a+s	1	2	3	4	5	10	20	30	40	50	100
1	1.118	.751	.560	.444	.367	.194	.099	.066	.049	.039	.019
2	1.411	1.071	.859	.715	.611	.351	.188	.128	.097	.078	.039
3	1.554	1.257	1.051	.902	.789	.482	.269	.186	.142	.115	.058
4	1.639	1.379	1.187	1.040	.925	.592	.342	.240	.185	.150	.077
5	1.697	1.466	1.288	1.147	1.032	.688	.410	.291	.226	.184	.096
10	1.829	1.683	1.558	1.449	1.354	1.017	.677	.507	.405	.337	.183
20	1.908	1.824	1.747	1.675	1.609	1.344	1.008	.807	.672	.575	.335
30	1.937	1.878	1.822	1.769	1.720	1.507	1.207	1.005	.862	.754	.463
40	1.952	1.906	1.863	1.821	1.781	1.605	1.338	1.148	1.004	.892	.573
50	1.961	1.924	1.888	1.854	1.820	1.670	1.433	1.254	1.115	1.003	.668
100	1.980	1.961	1.942	1.923	1.905	1.819	1.668	1.540	1.430	1.335	1.001

**Table 2 ( $\alpha = .75$ )**

b+f a+s	1	2	3	4	5	10	20	30	40	50	100
1	2.484	1.702	1.272	1.007	.829	.428	.212	.139	.104	.083	.040
2	2.986	2.303	1.856	1.548	1.322	.754	.397	.267	.201	.161	.080
3	3.224	2.642	2.221	1.909	1.672	1.018	.563	.386	.293	.236	.119
4	3.367	2.863	2.476	2.174	1.935	1.240	.712	.497	.381	.308	.157
5	3.463	3.019	2.663	2.378	2.143	1.429	.848	.600	.463	.377	.194
10	3.689	3.410	3.164	2.948	2.758	2.076	1.383	1.034	.824	.685	.370
20	3.827	3.666	3.516	3.375	3.245	2.715	2.039	1.631	1.358	1.163	.676
30	3.880	3.766	3.657	3.554	3.456	3.033	2.431	2.026	1.737	1.519	.933
40	3.908	3.819	3.734	3.652	3.574	3.224	2.691	2.308	2.020	1.795	1.153
50	3.925	3.853	3.783	3.715	3.649	3.351	2.877	2.519	2.240	2.016	1.343
100	3.961	3.923	3.886	3.849	3.813	3.643	3.342	3.087	2.867	2.676	2.008

Table 3 ( $\alpha = .9$ )

b+f a+s	1	2	3	4	5	10	20	30	40	50	100
1	7.028	5.001	3.796	3.021	2.488	1.269	.608	.391	.287	.226	.110
2	7.999	6.346	5.163	4.342	3.721	2.117	1.099	.732	.545	.433	.212
3	8.541	7.071	6.001	5.184	4.562	2.785	1.526	1.039	.784	.629	.313
4	8.721	7.538	6.578	5.809	5.179	3.333	1.906	1.322	1.008	.813	.411
5	8.904	7.868	6.996	6.276	5.676	3.800	2.249	1.585	1.219	.989	.506
10	9.341	8.694	8.103	7.572	7.101	5.373	3.582	2.674	2.129	1.767	.951
20	9.620	9.243	8.883	8.543	8.223	6.905	5.197	4.160	3.462	2.964	1.718
30	9.729	9.461	9.201	8.950	8.710	7.664	6.157	5.135	4.403	3.851	2.363
40	9.789	9.580	9.375	9.177	8.985	8.121	6.792	5.830	5.102	4.537	2.912
50	9.827	9.655	9.486	9.322	9.161	8.426	7.246	6.349	5.647	5.082	3.387
100	9.907	9.816	9.726	9.637	9.549	9.128	8.382	7.745	7.196	6.719	5.042

References.

- Derman, C. (1970). *Finite State Markovian Decision Processes*. Academic Press, New York.
- Gittins, J. C. and D. M. Jones (1974). A Dynamic Allocation Index for the Sequential Design of Experiments. In J. Gani, K. Sarkadi and I. Vince (eds.), *Progress in Statistics*, North Holland, 241-266.
- Gittins, J. C. and D. M. Jones (1979). A Dynamic Allocation Index for the discounted multiarmed bandit problem. *Biometrika*, 66, 561-565.
- Gittins, J. C. (1979). Bandit Processes and Dynamic Allocation Indices. *J. Roy. Statist. Soc. Ser. B* 41, 148-164.
- Glazebrook, K. D. (1978). On the optimal allocation of two or more treatments in a controlled clinical trial. *Biometrika*, 65, 335-340.

Katehakis M. N. and A. F. Veinott Jr. (1985). *The Multi-Armed Bandit Problem: Decomposition and Computation*. Department of Oper. Res., Stanford Univ., Technical Report, 13pp.

Robbins H. (1952). *Some Aspects of the Sequential Design of Experiments*. *Bull. Amer. Math. Monthly*. 58, 527-586.

END

DITIC

9 - 86