

OPTIMAL SEQUENTIAL SEARCH: A BAYESIAN APPROACH

BY DAVID ASSAF AND SHMUEL ZAMIR

Hebrew University of Jerusalem

To the classical model of searching for one object out of n , we add uncertainty about the parameters π of the distribution of the n objects among the m boxes. Adopting a Bayesian approach, we study the optimal sequential search strategy. For the case $n = 1$, we obtain a generalization of the fundamental result of Blackwell: the strategy which searches at each stage in the "most inviting" box is optimal. This strategy is also optimal for $m = 2$ and arbitrary n . However, for $n > 1$ the optimal strategy may be very different from that of the classical model, even when the uncertainty about π is very small.

1. Introduction. Suppose n objects are hidden in m boxes. Both n and m are assumed to be fixed and known. Denote by $M = \{1, \dots, m\}$ the set of boxes. Each of the objects is hidden in box i with probability π_i , ($\pi_i \geq 0$, $1 \leq i \leq m$, $\sum_{i=1}^m \pi_i = 1$) and independent of the other objects. The probability vector $\pi = (\pi_1, \dots, \pi_m)$ is unknown but has a known prior distribution. Let X_i denote the number of objects in box i . Then, given the value of π , the vector $\mathbf{X} = (X_1, \dots, X_m)$ has a multinomial distribution with parameters n and π .

Associated with each box i are two quantities, both assumed fixed and known: c_i the cost of searching for objects in box i and α_i the (conditional) probability of finding an object which is in box i when box i is searched. It is assumed that searches for different objects hidden in the same box are independent in the sense that if there are k objects in box i then the probability of finding at least one of them when searching that box equals $1 - (1 - \alpha_i)^k$. To avoid trivialities assume $0 < \alpha_i < 1$, $P(\pi_i = 0) < 1$ and $c_i > 0$ for all $1 \leq i \leq m$. The goal is to search the boxes in succession until at least one object is found, and to do so with a minimal expected total cost.

A search strategy is a sequence $S = (s_1, s_2, \dots)$ where $1 \leq s_d \leq m$ for all $d \geq 1$. Applying the strategy S , we search in box s_d at stage d if no object was found in the first $d - 1$ searches. At each stage of the search denote by h_i the "current" probability of finding an object in box i if we look there, i.e., the conditional probability of finding at least one object there, given the history of unsuccessful searches up to that stage. A box i is said to be *most inviting* (at a certain stage of the search) if $h_i/c_i \geq h_j/c_j$ for all $j \in M$. A strategy which searches at each stage in a most-inviting box will be called a most-inviting strategy. Due to possibilities of equalities $h_i/c_i = h_j/c_j$ there may be many most-inviting strategies; but for the sake of simplicity we shall refer to each of them as *the* most-inviting strategy.

Using general results of negative dynamic programming with finite action

Received September 1983; revised November 1984.

AMS 1980 subject classification. Primary 90B40.

Key words and phrases. Optimal sequential search, most inviting strategy, Bayesian approach.

space, it may be shown that there exists an optimal (stationary nonrandomized) strategy. For a formal derivation see, Smith and Kimeldorf (S-K hereforth) [8] and the results of Strauch [11].

The problem presented above, in different variations, has been studied in several papers. A basic assumption in all previous papers has been that the prior distribution of π is degenerate at some known value. Blackwell [2], Chew [3], Kadane [5], Kelly [6], [7], Sweat [12] and many others have studied the problem for $n = 1$ and showed the optimality of the most-inviting strategy for that case. Smith and Kimeldorf [8] derived results when n is an integer-valued random variable (and in this respect their model is more general than ours, but they assumed that π is known). For other recent results for $n > 1$, see [1], [4], [9] and [13].

The Bayesian (random π) model studied in this paper is appropriate for what seems to be a most common situation; namely, some information regarding π is known, but its exact value is not. This may be the case, for example, if π is basically known but some disturbances of random nature are present or when π is estimated from some previous observations. It turns out that random fluctuations, or statistical inaccuracies, even small ones, may give rise to results totally different from the results obtained when π is perfectly known. These differences become extreme when n is large and the prior distribution of π has supports in the vicinity of the sets $\pi_i = 0$. In Section 3 we provide an example (Example 2) with $m = 2$ in which the optimal strategy for degenerate π is to search box 1 in succession 12 times, while the optimal strategy when a small uncertainty about π is assumed is to search box 1 first and then if no object was found, to search box 2 next.

In Section 2 we show that the most-inviting strategy is optimal in two cases: the case $n = 1$ (Theorem 1) and the case $m = 2$ and any fixed n (Theorem 2). Examples are discussed in Section 3. These are mainly designed to illustrate the differences between the Bayesian and the non-Bayesian approaches. An example considered in Section 3 shows that, even when all costs are equal, the Bayesian model differs from the non-Bayesian one.

2. Results for the Bayesian model. We begin the section with several useful formulas, all of which may be verified using elementary methods. Expectation is denoted as usual by E and, unless specified otherwise, is the expectation with respect to the original prior distribution of π .

The probability of finding at least one object when searching box i is given by

$$(1) \quad h_i = 1 - E(1 - \alpha_i \pi_i)^n.$$

The conditional distribution of X_i given π , following an unsuccessful search of box k , is binomial with parameters n and $\pi_i^{(k)}$, given by

$$(2) \quad \pi_i^{(k)} = \pi_i / (1 - \alpha_k \pi_k), \quad i \neq k$$

$$(3) \quad \pi_k^{(k)} = (1 - \alpha_k) \pi_k / (1 - \alpha_k \pi_k).$$

The probability of finding at least one object in box i , after an unsuccessful

search of box k , is given by

$$(4) \quad h_i^{(k)} = \frac{E(1 - \alpha_k \pi_k)^n - E(1 - \alpha_i \pi_i - \alpha_k \pi_k)^n}{E(1 - \alpha_k \pi_k)^n}, \quad i \neq k$$

$$(5) \quad h_k^{(k)} = \frac{E(1 - \alpha_k \pi_k)^n - E(1 - 2\alpha_k \pi_k + \alpha_k^2 \pi_k)^n}{E(1 - \alpha_k \pi_k)^n}.$$

REMARK 1. An alternative representation may be obtained by considering the joint posterior distribution of the vector $\boldsymbol{\pi}^{(k)} = (\pi_1^{(k)}, \dots, \pi_n^{(k)})$. With this representation it is easy to verify that, following an unsuccessful search of box k , we have a new problem which differs from the original one only in the distribution of $\boldsymbol{\pi}$. Additional searches now give rise to other distributions of $\boldsymbol{\pi}$ and one may consider the family of all distributions which may be obtained in this manner. We will usually prefer the representations in terms of the original prior of $\boldsymbol{\pi}$ but it is useful to keep this remark in mind for some applications (see Lemmas 1 and 2 and Example 1). This also clarifies the relation of our problem to dynamic programming.

REMARK 2. It is worth noting that all formulas remain valid for the more general case in which n is also random. In that case E is to be interpreted as expectation with respect to the prior joint distribution of $(\boldsymbol{\pi}, n)$.

For any search strategy $S = (s_1, s_2, \dots)$, denote by $C(S)$ the expected total cost until an object is found when using the strategy S . The following lemma indicates one reason for the most-inviting strategy to be a natural candidate for optimality.

LEMMA 1 (Compare to Lemma 1 in [8]). *Box i is more inviting than box j if and only if $C(i, j, s_3, s_4, \dots) \leq C(j, i, s_3, s_4, \dots)$ for any (s_3, s_4, \dots) .*

PROOF.

$$C(i, j, s_3, s_4, \dots) = c_i + c_j(1 - h_i) + D(i, j)q(i, j)$$

$$C(j, i, s_3, s_4, \dots) = c_j + c_i(1 - h_j) + D(j, i)q(j, i)$$

where $D(i, j)$ is the total expected cost of the search sequence (s_3, s_4, \dots) when the distribution of $\boldsymbol{\pi}$ is the posterior one after unsuccessful searches in box i and then box j (see Remark 1), and $q(i, j)$ is the probability of unsuccessful searches in boxes i and j . Clearly $D(i, j) = D(j, i)$ and $q(i, j) = q(j, i)$ so that $C(i, j, s_3, s_4, \dots) \leq C(j, i, s_3, s_4, \dots)$ if and only if $c_i + c_j(1 - h_i) \leq c_j + c_i(1 - h_j)$ which is equivalent to box i being more inviting than box j (i.e., $h_i/c_i \geq h_j/c_j$). \square

Given a distribution F of $\boldsymbol{\pi}$, denote by $\Phi(F)$ the family of all distributions of $\boldsymbol{\pi}$ which are obtainable from F through sequences of unsuccessful searches (see Remark 1).

The following lemma provides a sufficient condition for the optimality of the most-inviting strategy.

LEMMA 2. *If for each $G \in \Phi(F)$*

$$(6) \quad h_i/c_i \geq h_j/c_j \quad \text{for all } j \in M,$$

implies

$$(7) \quad h_i^{(k)}/c_i \geq h_j^{(k)}/c_j \quad \text{for all } j \in M, \quad k \neq i,$$

then the most-inviting strategy is optimal (for that F). Here h_i and $h_i^{(k)}$ are given by (1), (4) and (5) with all expectations taken with respect to G .

REMARK 3. This condition is actually the key element in Blackwell's original proof which was based on the fact that this condition is satisfied (trivially) for $n = 1$ and degenerate F (in which case $\Phi(F) = \{F\}$). Later this condition was used again by S-K who verified it for degenerate F when n is distributed $pp(\lambda)$ (see [8], Lemma 2).

PROOF. It is easily seen that an optimal strategy S^* satisfies $C(S^*) < \infty$ and hence must search infinitely often in each box. Let $S^* = (s_1^*, s_2^*, \dots, s_{t-1}^*, s_t^*, \dots)$, ($t \geq 1$) be an optimal strategy in which s_1^*, \dots, s_{t-1}^* are most inviting in their respective stages. We will show that there exists an optimal $S' = (s_1^*, s_2^*, \dots, s_{t-1}^*, s_t', \dots)$, in which s_t' is the most inviting at stage t , which clearly proves the lemma.

To show this, let $\Sigma = \{S = (s_1, s_2, \dots, s_t, \dots) \mid S \text{ is optimal and } s_\ell = s_\ell^* \text{ for } 1 \leq \ell \leq t-1\}$. Let i be a most-inviting box following unsuccessful searches in $s_1^*, s_2^*, \dots, s_{t-1}^*$. For each $S \in \Sigma$, let $m = m(S)$ be the smallest integer such that $s_{t+m} = i$ (such an m exists since an optimal strategy searches infinitely often in box i). Let $M = \min\{m(S) \mid S \in \Sigma\}$. We need to show that $M = 0$.

Suppose $M \geq 1$ and let $S = (s_1^*, s_2^*, \dots, s_{t-1}^*, s_t, \dots, s_{t+M-1}, i, \dots)$ be optimal. Since $s_\ell \neq i$ for $\ell = t, \dots, t+M-1$, it follows from the conjecture of the lemma (applied $M-1$ times) that i remains most inviting at stage $t+M-1$ (since it is most inviting at stage t). Denote $s_{t+M-1} = j \neq i$. From Lemma 1 it follows that the strategy $\hat{S} = (s_1^*, s_2^*, \dots, s_{t-1}^*, s_t, \dots, i, j, \dots)$ which is obtained from S by interchanging i and j has $C(\hat{S}) \leq C(S)$, and hence $\hat{S} \in \Sigma$. But $m(\hat{S}) = M-1$, a contradiction. Thus $M = 0$ and the proof is complete. \square

The following theorems prove the optimality of the most-inviting strategy in two cases. The first theorem extends well-known results for the case $n = 1$ to the Bayesian model.

THEOREM 1. *For $n = 1$ and any prior distribution F of π , the most-inviting strategy is optimal.*

PROOF. Applying (1), (4) and (5) for the case $n = 1$, we have, for $k \neq i, j$

$$\frac{h_i^{(k)}}{h_j^{(k)}} = \left(\frac{\alpha_i E\pi_i}{1 - \alpha_k E\pi_k} \right) / \left(\frac{\alpha_j E\pi_j}{1 - \alpha_k E\pi_k} \right) = \frac{\alpha_i E\pi_i}{\alpha_j E\pi_j} = \frac{h_i}{h_j}$$

in which case (6) and (7) are equivalent. For $k = j$ (and $i \neq j$), we have

$$\frac{h_i^{(k)}}{h_j^{(k)}} = \frac{\alpha_i E\pi_i}{\alpha_j(1 - \alpha_j)E\pi_j} \geq \frac{\alpha_i E\pi_i}{\alpha_j E\pi_j} = \frac{h_i}{h_j}$$

so that (6) implies (7). In any case, the conditions of Lemma 2 are satisfied and the most-inviting strategy is optimal. \square

THEOREM 2. *For $m = 2$ and for any prior distribution F , the most-inviting strategy is optimal.*

PROOF. From Lemma 2 and obvious symmetry considerations, it is sufficient to show that

$$h_2^{(1)}/h_1^{(1)} \geq h_2/h_1,$$

that is, using (1), (4) and (5)

$$(8) \quad \frac{E(1 - \alpha_1\pi_1)^n - E(1 - \alpha_1\pi_1 - \alpha_2\pi_2)^n}{E(1 - \alpha_1\pi_1)^n - E(1 - 2\alpha_1\pi_1 + \alpha_1^2\pi_1)^n} \geq \frac{1 - E(1 - \alpha_2\pi_2)^n}{1 - E(1 - \alpha_1\pi_1)^n}.$$

Now

$$(9) \quad \begin{aligned} E(1 - \alpha_1\pi_1 - \alpha_2\pi_2)^n &\leq E(1 - \alpha_1\pi_1 - \alpha_2\pi_2 + \alpha_1\alpha_2\pi_1\pi_2)^n \\ &= E[(1 - \alpha_1\pi_1)(1 - \alpha_2\pi_2)]^n \\ &\leq E(1 - \alpha_1\pi_1)^n E(1 - \alpha_2\pi_2)^n. \end{aligned}$$

To see the second inequality, let $U = 1 - \alpha_1\pi_1$ and $V = 1 - \alpha_2\pi_2 = 1 - \alpha_2 + \alpha_2\pi_1$. ($\pi_2 = 1 - \pi_1$ when $m = 2$.) U and V are negatively associated (recall that every random variable is associated with itself) so that $\text{cov}(U^n, V^n)$ is nonpositive. From (9) we now have

$$(10) \quad E(1 - \alpha_1\pi_1)^n - E(1 - \alpha_1\pi_1 - \alpha_2\pi_2)^n \geq E(1 - \alpha_1\pi_1)^n [1 - E(1 - \alpha_2\pi_2)^n]$$

On the other hand, since $\pi_1^2 \leq \pi_1$ and $EZ^2 \geq (EZ)^2$ it follows that

$$(11) \quad \begin{aligned} E(1 - 2\alpha_1\pi_1 + \alpha_1^2\pi_1)^n &\geq E(1 - 2\alpha_1\pi_1 + \alpha_1^2\pi_1^2)^n \\ &= E[(1 - \alpha_1\pi_1)^2]^n \geq [E(1 - \alpha_1\pi_1)^n]^2, \end{aligned}$$

so that

$$(12) \quad \begin{aligned} E(1 - \alpha_1\pi_1)^n - E(1 - 2\alpha_1\pi_1 + \alpha_1^2\pi_1)^n \\ \leq E(1 - \alpha_1\pi_1)^n [1 - E(1 - \alpha_1\pi_1)^n]. \end{aligned}$$

Inequality (8) now follows directly from (10) and (12). \square

REMARK 4. Reflection upon the proof of Theorem 2 shows that the proof remains valid for any m , provided $1 - \alpha_i\pi_i$ and $1 - \alpha_j\pi_j$ ($i \neq j$) are negatively associated. What is really needed is a condition which ensures the inequality

$$E[(1 - \alpha_i\pi_i)(1 - \alpha_j\pi_j)]^n \leq E(1 - \alpha_i\pi_i)^n E(1 - \alpha_j\pi_j)^n.$$

This does not seem to be the case even in some special convenient examples, as will be demonstrated in the next section.

REMARK 5. Theorem 2 above holds for any prior F , in particular for degenerate ones. Thus, the most-inviting strategy is optimal in the S-K model when $m = 2$ and n is fixed. In [1] we show that this is not true for $m = 2$ when n is random (S-K leave this as an open problem in [8]).

3. Examples. The first two examples compare the Bayesian and the deterministic model under similar basic problem structure. Both examples have $m = 2$ and fixed n so that by Theorem 2 the most-inviting strategy is optimal. The first is a simple discrete case.

EXAMPLE 1. Let $m = 2$, $\alpha_1 = 0.9$ (α_2 , c_1 and c_2 unspecified). Suppose $\pi = \pi_1$ takes the values 0.5 with probability 0.98, and the values 0 and 1 with probability 0.01 each. Naturally $\pi_2 = 1 - \pi$. After an unsuccessful search of box 1, the number of objects in box 1 has a new binomial distribution with a probability of “success” $\pi^{(1)}$ given by (2) and (3). The probability function of $\pi^{(1)}$ is tabulated for some values of n as follows (the support of $\pi^{(1)}$ is $\{0, \frac{1}{11}, 1\}$ independently of n).

$\pi^{(1)}$	n		
	5	10	20
0	0.17	0.80	0.9994
$\frac{1}{11}$	0.83	0.20	0.0006
1	17×10^{-6}	10^{-10}	10^{-21}

This should be compared with a value of $(1 - 0.9) \times 0.5 / (1 - 0.9 \times 0.5) = \frac{1}{11}$ (with probability 1) if π is assumed to equal 0.5 in the deterministic model. We see that the Bayesian approach in fact does what is was expected to do: the beliefs on the value of π are revised after the unsuccessful search of box 1, and when n is large this revision becomes extreme.

EXAMPLE 2. Let $m = 2$ and let $\pi = \pi_1$ have a uniform distribution on $(0, 1)$. Let $X = X_1$ denote the number of objects in box 1. Initially $EX = n/2$. Following an unsuccessful search of box 1, the expected number of objects in box 1, $E^{(1)}X$, is given by

$$E^{(1)}X = nE^{(1)}\left(\frac{(1 - \alpha_1)\pi}{1 - \alpha_1\pi}\right) = n(1 - \alpha_1) \frac{E\pi(1 - \alpha_1\pi)^{n-1}}{E(1 - \alpha_1\pi)^n}.$$

For π uniform on $(0, 1)$ this quantity is easily computed in closed form as

$$E^{(1)}X = \frac{1 - \alpha_1}{\alpha_1} - \frac{(n + 1)(1 - \alpha_1)^{n+1}}{1 - (1 - \alpha_1)^{n+1}}.$$

Thus for large values of n , $E^{(1)}X \simeq (1 - \alpha_1)/\alpha_1$, i.e., a finite constant ($E^{(1)}X \simeq 1$ for $\alpha_1 = 0.5$). This quantity may be compared with the value of $E^{(1)}X$ when π is

degenerate at 0.5. In this latter case we have

$$E^{(1)}X = n(1 - \alpha_1)/(2 - \alpha_1)$$

which tends to infinity linearly with n (for $n = 300$ and $\alpha_1 = 0.5$, $E^{(1)}X \simeq 1$ in the Bayesian model while $E^{(1)}X = 100$ in the deterministic model).

A second comparison between the Bayesian and the deterministic models is the following. Applying (5) for π uniform on $(0, 1)$ we have

$$h_1^{(1)} = (1 - \alpha_1 - (1 - \alpha_1)^{n+1})/(2 - \alpha_1).$$

As n approaches infinity, this quantity is strictly less than 1 ($h_1^{(1)} \rightarrow 1/3$ as $n \rightarrow \infty$ for $\alpha_1 = 0.5$).

Based on this last observation it is not difficult to construct striking examples, which are very similar in their basic formulation but differ entirely in their optimal search strategies. As a specific example, consider the case where π is a mixture of a uniform distribution on $(0, 1)$ with probability ε and a degenerate distribution at 0.5 with probability $1 - \varepsilon$. For small ε this may be interpreted as basically degenerate at 0.5, but allowing a small chance for unknown disturbances.

Let $\alpha_1 = \alpha_2 = 0.5$, $c_1 = 1$, $c_2 = 2$, $\varepsilon = 0.1$ and $n = 10,000$. Since h_1 and h_2 are both almost equal to 1, it follows that box 1 is more inviting than box 2. After an unsuccessful search of box 1, however, it is readily verified that $h_1^{(1)} \simeq 1/3$ while $h_2^{(1)} \simeq 1$. Box 2 is now more inviting so that by Theorem 2 it follows that the optimal strategy is to search box 1 first and then box 2.

Consider now the case $\varepsilon = 0$ (i.e., π is degenerate at 0.5). After k unsuccessful searches of box 1, the probability of finding an object in it is given by

$$h_1^{(1, \dots, 1)} = 1 - ((1 - (1 - (1 - \alpha_1)^{k+1})\pi)/(1 - (1 - (1 - \alpha_1)^k)\pi))^n$$

which tends (very) slowly to 0. For $\alpha_1 = \alpha_2 = 0.5$, $\pi = 0.5$ and $n = 10,000$, it equals about 0.99 when $k = 10$ and about 0.5 when $k = 13$ (the corresponding probability for box 2 is approximately 1 for all k). Since the most-inviting strategy is optimal (by Theorem 2) it follows that box 1 should be searched about 12 times in succession before searching box 2. It seems disturbing that one would keep searching in a box after not finding a single object in it (and which should contain 5,000 of them on the average) without suspecting that something is wrong.

EXAMPLE 3. Assume $c_i = c$ for all i (without loss of generality take $c = 1$). For any search sequence $S = (s_1, s_2, \dots)$, let $T = T(S)$ be the first stage in which an object was found. The total expected cost associated with the sequence S in this case equals

$$ET = \sum_{d=1}^{\infty} P(T \geq d).$$

One way of minimizing ET is by minimizing $P(T \geq d)$ for every $d = 1, 2, \dots$. A sequence for which this is accomplished may be referred to as a *uniformly optimal strategy* (see [3]). It is easy to verify that a necessary condition for a strategy to be uniformly optimal is for it to be most inviting. The most-inviting strategy on the other hand may or may not be uniformly optimal.

The case of equal c_i is very simple in most models studied thus far, and it is

relatively easy to show that the most-inviting strategy is uniformly optimal for degenerate π as well as for many other models ([2], [3], [8] and [10]).

Surprisingly, as the following counterexample shows, it does not hold for the Bayesian case, even for the most appealing prior for π —the Dirichlet distribution.

Consider the case with $m = 3$, $n = 2$. The detection parameters for the three boxes are: $\alpha_1 = 0.9958$, $\alpha_2 = 0.4877$, $\alpha_3 = 0.2959$. The prior on π is the three-dimensional Dirichlet distribution (i.e., density proportional to $\pi_1^{\nu_1-1} \pi_2^{\nu_2-1} \pi_3^{\nu_3-1}$ over the region $\pi_i \geq 0$; $\pi_1 + \pi_2 + \pi_3 = 1$), with parameters $\nu_1 = 2.273$, $\nu_2 = 4.545$, $\nu_3 = 7.410$. Straightforward (though somewhat tedious) calculations reveal the following:

- At the first stage box 1 is most inviting.
- After an unsuccessful search in box 1, box 2 becomes most inviting.
- The sequence 12 is not optimal for a two-stage search since the probability for finding at least one object in two searches is higher for the sequence 23.

(The figures were found by checking the consistency of several nonlinear inequalities using a Monte Carlo method and approximately a million randomized trials were needed.)

It should be noted that the example merely shows that the most-inviting strategy is not necessarily uniformly optimal. The question regarding optimality of the most-inviting strategy remains open.

4. Concluding remarks. One possible extension of our results may be obtained by taking n to be an integer-valued random variable. This was done in [8] but only for degenerate π . A basic result in [8] is that if n has a $pp(\lambda)$ distribution (a Poisson distribution conditioned on $n \geq 1$), then the most-inviting strategy is optimal and conversely, if $m \geq 3$ and n is not $pp(\lambda)$, then there exist values of α , c and π for which the most-inviting strategy is not optimal.

Taking both n and π as random, a generalized question arises as follows: What are the joint distributions of (n, π) for which the most-inviting strategy is optimal? S-K roughly state that if π is degenerate then n must be $pp(\lambda)$. Theorems 1 and 2 of this paper may be viewed as partial answers to the more general problem.

If the family of distributions of π considered includes the degenerate ones, then the conditional distribution of n given π must be $pp(\lambda)$ for an affirmative answer ($m \geq 3$). Fortunately in this case the distribution of n given π remains $pp(\lambda)$ following each unsuccessful search (parameter depends both on π and on the box searched). If the family of distributions of π does not include the degenerate ones, then n (or n given π) need not be $pp(\lambda)$ (at least not by any known results). Care should be taken, however, since many families of distributions contain the degenerate ones in their closure. This is the case, for example, for the family of Dirichlet distributions.

In spite of this last fact and the counterexample for equal c_i in the previous section, we tend to believe that the Dirichlet distributions are “natural” candidates for this problem. We were not able, however, neither to prove nor to find

a counterexample to the conjecture that if π is Dirichlet and n given π is $pp(\lambda)$, then the most-inviting strategy is optimal.

Acknowledgements. The authors would like to thank Israel Einot of the Hebrew University for writing the computer program for Example 3. Thanks are also due to an Associate Editor and two referees for their many helpful comments and criticisms of a previous version of this paper.

REFERENCES

- [1] ASSAF, D. and ZAMIR, S. (1985). Optimal sequential search for one of many objects. Technical Report. Submitted for publication.
- [2] BLACKWELL, D. (1962). Notes on dynamic programming. Unpublished notes, University of California, Berkeley.
- [3] CHEW, M. C. (1967). A sequential search procedure. *Ann. Math. Statist.* **38** 492–502.
- [4] HASSIN, R. and HENIG, M. (1984). Dichotomous search for random objects on an interval. *Math. Oper. Res.* **9** 301–308.
- [4] KADANE, J. B. (1968). Discrete search and the Neyman-Pearson lemma. *J. Math. Anal. Appl.* **22** 156–171.
- [6] KELLY, F. P. (1982). An optimal search with unknown detection probabilities. *Journal Math. Anal.* **88** 422–432.
- [7] KELLY, F. P. (1982). A remark on search and sequencing problems. *Math. Oper. Res.* **7** 154–157.
- [8] SMITH, F. H. and KIMELDORF, G. (1975). Discrete sequential search for one of many objects. *Ann. Statist.* **3** 906–915.
- [9] SMITH, F. H. and KIMELDORF, G. (1979). Binomial searching for a random number of multinomially hidden objects. *Management Sci.* **25** 1115–1126.
- [10] STONE, L. (1975). *Theory of Optimal Search*. Academic, New York.
- [11] STRAUCH, R. E. (1966). Negative dynamic programming. *Ann. Math. Statist.* **37** 871–890.
- [12] SWEAT, C. W. (1970). Sequential search with discounted income, the discount function of the cell searched. *Ann. Math. Statist.* **41** 1446–1455.
- [13] WEGENER, I. (1982). The discrete search problem and the construction of optimal allocations. *Naval Res. Logist. Quart.* **29** 202–212.
- [14] WILKS, S. S. (1962). *Mathematical Statistics*. Wiley, New York.

DEPARTMENT OF STATISTICS
HEBREW UNIVERSITY
91905 JERUSALEM
ISRAEL