2016

Finding a hider by an unknown deadline

Lin, Kyle Y.; Singham, Dashi I.

Elsevier

Lin, Kyle Y., and Dashi I. Singham. "Finding a hider by an unknown deadline."
Operations Research Letters 44.1 (2016): pp. 25-32.
http://hdl.handle.net/10945/52672

Operations Research Letters 44 (2016) 25–32

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

Finding a hider by an unknown deadline

Kyle Y. Lin *, Dashi I. Singham

Operations Research Department, Naval Postgraduate School, Monterey, CA 93943, United States

ARTICLE INFO

Article history:
Received 15 May 2015
Received in revised form
5 November 2015
Accepted 6 November 2015
Available online 14 November 2015

Keywords:
Search theory
Two-person zero-sum game
Robust strategy

ABSTRACT

An object is hidden among several locations. Each search at the object's location independently finds the object with some location-dependent probability. The goal is to find the object by a deadline, but the deadline is unknown. Assuming the worst-case scenario, where Nature knows the deadline and uses this knowledge to hide the object to hinder the search, this paper shows that there is a randomized search strategy that simultaneously maximizes the probability of finding the object by any deadline.

Published by Elsevier B.V.

1. Introduction

In many search and rescue missions, it is of prime importance to find an object before a crucial deadline. For instance, when a ship is lost at sea, the search team needs to find overboard sailors before they die of hypothermia. After an earthquake, the rescue team may need to find survivors underneath collapsed buildings before they die. A bomb squad may need to find a time bomb before the bomb explodes. In many cases, the crucial deadline is not known to the search team. To study this problem, we consider a search model, where the object is hidden in one of several heterogeneous locations. For each location, there is a known probability that a search at the location will find the object, if the object is hidden there. Assuming the worst-case scenario, where Nature knows the deadline and uses this knowledge to choose where to hide the object to hinder the search, we show that there is a randomized search strategy that simultaneously maximizes the probability of finding the object by an arbitrary deadline. In other words, the searcher can still carry out the optimal search without ever knowing the deadline.

To obtain this result, we treat the object (henceforth the hider) as someone who actively chooses which location (henceforth cell) to hide in order to evade the search. Consider a game-theoretic model as follows. A hider chooses where to hide among n cells. Independent of previous search results, each search in cell i will find the hider with probability α_i, if the hider is there, for $i = 1, \ldots, n$. For a positive integer t, denote the game by $G(t)$, where the searcher gets t searches and wins if he finds the hider, or loses otherwise. The payoff is the probability of finding the hider in t searches, which the searcher wishes to maximize and the hider wishes to minimize. The game $G(t)$ is a two-person zero-sum game. The hider's pure strategy space is $\{1, 2, \ldots, n\}$, where each pure strategy corresponds to hiding in one of the n cells. The searcher's pure strategy space is the Cartesian product $\{1, 2, \ldots, n\}^t$, where each pure strategy corresponds to a search sequence of length t. Since each player's pure strategy space is finite, by the duality theorem of linear programming, the game $G(t)$ has a value, which is denoted by $v(t)$.

Write $\mathbb{N} \equiv \{1, 2, \ldots\}$ for the set of positive integers. We call a player's strategy *uniformly optimal* for $T \subseteq \mathbb{N}$, if the strategy is simultaneously optimal for all games $G(t)$, for $t \in T$. The main contribution of this paper is to show that the searcher always has a uniformly optimal strategy for $G(t),\ t \in \mathbb{N}$. By using the uniformly optimal search strategy, the searcher maximizes the probability of finding the hider by the deadline, without having to know the deadline. On the hider's side, it turns out that the hider has a uniformly optimal strategy for $G(t),\ t \in \mathbb{N}$, if and only if α_i are the same for all $i = 1, \ldots, n$.

This search game has been studied in the literature. Subelman [14] showed how to compute each player's optimal strategy for $G(t)$, for any $t \in \mathbb{N}$. Gittins and Roberts [6,11] considered a different objective function, and showed how to compute the optimal hiding probability to maximize the expected time to detection. In the classical search model, the hiding probabilities are revealed to the searcher

* Corresponding author.
E-mail addresses: kylin@nps.edu (K.Y. Lin), dsingham@nps.edu (D.I. Singham).

http://dx.doi.org/10.1016/j.orl.2015.11.003
0167-6377/Published by Elsevier B.V.

at the beginning. To minimize the expected time to detection, the optimal policy is to search, at any time, the cell that has the largest present value of $\alpha_i p_i$, with the hiding probabilities $(p_1, \ldots, p_n)$ being updated in the Bayesian fashion after each unsuccessful search. This result was first attributed to Blackwell in his notes on dynamic programming; see Black [3] and Matula [10]. Chew [4] and Kadane [7] showed that the same policy maximizes the probability of finding the hider within t searches, for every $t \in \mathbb{N}$. For variants of this search problem, please see Chew [5], Kadane [8], Kress et al. [9], Ross [12], and Wegener [16]. For a general survey on search theory, please see Alpern et al. [1], Alpern and Gal [2], Stone [13], and Washburn [15].

The rest of this paper proceeds as follows. Section 2 presents an example with $n = 2$ cells to illustrate the idea. Section 3 proves our main result that the searcher always has a uniformly optimal strategy for all games $G(t),\ t \in \mathbb{N}$, and Section 4 proves that the hider has a uniformly optimal strategy if and only if α_i are the same for $i = 1, \ldots, n$. Section 5 offers two corollaries to conclude the paper.

2. An example with two cells

We illustrate the main idea of the paper with a simple example with $n = 2$ cells, where $\alpha_1 = 0.2$ and $\alpha_2 = 0.3$. For the game $G(t)$, the hider has two pure strategies (hide in cell 1 or cell 2), and the searcher has a pure strategy space of $\{1, 2\}^t$.

For the game $G(1)$, the payoff matrix is given by

	Hide in 1	Hide in 2
Search (1)	0.2	0
Search (2)	0	0.3

It is straightforward to verify that the hider's optimal mixed strategy is (0.6, 0.4), and the searcher's optimal mixed strategy is also (0.6, 0.4).

For the game $G(2)$, the searcher has 4 pure strategies, and the payoff matrix is given by

	Hide in 1	Hide in 2
Search (1, 1)	$1 - 0.8^2 = 0.36$	0
Search (1, 2)	0.2	0.3
Search (2, 1)	0.2	0.3
Search (2, 2)	0	$1 - 0.7^2 = 0.51$

It is straightforward to verify that the hider's optimal mixed strategy is (15/23, 8/23). Write $x_{i,j}$ for the probability of using the search sequence (i, j) in the searcher's optimal mixed strategy in $G(2)$, and one can obtain

$$x_{1,1} = \frac{5}{23}, \qquad x_{1,2} + x_{2,1} = \frac{18}{23}, \qquad x_{2,2} = 0. \tag{1}$$

It is clear that the hider does not have a uniformly optimal strategy for $T = \{1, 2\}$, since $(0.6, 0.4) \neq (15/23, 8/23)$. If the searcher wants to have a uniformly optimal strategy for $T = \{1, 2\}$, then we need

$$x_{1,1} + x_{1,2} = 0.6, \qquad x_{2,1} + x_{2,2} = 0.4. \tag{2}$$

Solving Eqs. (1) and (2) yields

$$x_{1,1} = \frac{5}{23}, \qquad x_{1,2} = \frac{44}{115}, \qquad x_{2,1} = 0.4, \qquad x_{2,2} = 0.$$

It turns out that for any n and $\alpha_1, \ldots, \alpha_n$, the searcher always has a uniformly optimal strategy for $G(t),\ t \in \mathbb{N}$. Section 3 proves this result by a forward induction. On the other hand, the hider has a uniformly optimal strategy for $G(t),\ t \in \mathbb{N}$, if and only if α_i are the same for $i = 1, \ldots, n$, which will be proved in Section 4.

3. The searcher's uniformly optimal strategy

Theorem 1. *For any number of cells n and detection probabilities $\alpha_i,\ i = 1, \ldots, n$, the searcher has a mixed strategy on $\{1, \ldots, n\}^\infty$ that is uniformly optimal for the games $G(t),\ t \in \mathbb{N}$.*

This section proves Theorem 1. Section 3.1 presents the structure of the searcher's optimal strategy for $G(t)$, for given $t \in \mathbb{N}$. Sections 3.2–3.4 prove Theorem 1 for $n = 2$, $n = 3$, and $n \geq 4$, respectively, by showing that we can randomly generate a pure search sequence such that the mixed strategy is uniformly optimal for $G(t),\ t \in \mathbb{N}$.

3.1. Structure of the optimal mixed strategy

For given $t \in \mathbb{N}$, recall that in the game $G(t)$, the searcher wants to maximize the probability of finding the hider within t searches. If a search sequence contains x_i searches in cell i, then it will find the hider with probability

$$1 - (1 - \alpha_i)^{x_i}, \tag{3}$$

if the hider hides in cell i, for $i = 1, \ldots, n$. In particular, two search sequences are equivalent in value, if they have the same number of searches in each cell, even if the orders of those searches are different. The search sequence (searcher's pure strategy) thus produces the payoff value

$$\min_i \; 1 - (1 - \alpha_i)^{x_i}.$$

Since Eq. (3) is an increasing and concave function in x_i, and the searcher has only t searches to allocate among the n cells, it follows that a greedy allocation algorithm is sufficient to find the best pure strategy for the searcher, with which the searcher sequentially allocates one search to the cell that has the minimal value in Eq. (3) with x_i representing the number of searches already allocated. It is then intuitive that the optimal mixed strategy involves a set of pure strategies, whose numbers of searches in each cell differ by at most 1. Subelman [14] formalized this idea and presented an algorithm to compute the optimal mixed strategy for both the searcher and the hider, which is paraphrased below.

Theorem 2 (*Subelman [14]*). *In the game $G(t)$, the searcher's optimal mixed strategy takes the following form. There exist nonnegative integers $\mu_i(t)$, and fractional numbers $\Psi_i(t) \in [0, 1)$, for $i = 1, \ldots, n$, with*

$$\sum_{i=1}^{n} \mu_i(t) + \sum_{i=1}^{n} \Psi_i(t) = t,$$

such that cell i will receive $\mu_i(t)$ searches with probability $1 - \Psi_i(t)$, or $\mu_i(t) + 1$ searches with probability $\Psi_i(t)$, for $i = 1, \ldots, n$. In other words, cell i will receive at least $\mu_i(t)$ searches, and will receive one additional search with probability $\Psi_i(t)$.

Our goal is to show that the searcher has a uniformly optimal mixed strategy for $G(t)$, for all $t \in \mathbb{N}$. To prove this result, we use mathematical induction to show that if we can construct a mixed strategy that is uniformly optimal for $G(1), G(2), \ldots, G(t)$, then we can construct a mixed strategy that is uniformly optimal for $G(t + 1)$.

To carry out this forward induction, we need to describe a policy based on states and actions. At any stage of the search, the *state* of the search can be delineated by $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$ if the searcher has searched x_i times in cell i. While the sum of the components of $\boldsymbol{x}$ gives t, we sometimes use $\boldsymbol{x}_t$ to denote the state after t searches. A *randomized policy* π maps from a state to a probability distribution over n cells, such that in state $\boldsymbol{x}$ the searcher allocates the next search in cell i with probability $\pi(\boldsymbol{x}, i) \in [0, 1]$, with $\sum_{i=1}^{n} \pi(\boldsymbol{x}, i) = 1$.

For a given randomized policy π, we can generate a random search sequence, denoted by $\{X(t), t = 1, 2, \ldots\}$, such that $X(t) = i$ indicates that cell i is searched at time t. Subelman [14] showed that, for any $t \in \mathbb{N}$, a random search sequence is optimal for $G(t)$, if the joint distribution of $X(1), X(2), \ldots, X(t)$ meets the following two conditions:

$$P\left\{\sum_{s=1}^{t} \mathbf{1}(X(s) = i) = \mu_i(t)\right\} = 1 - \Psi_i(t), \qquad \text{for } i = 1, \ldots, n; \tag{4}$$

$$P\left\{\sum_{s=1}^{t} \mathbf{1}(X(s) = i) = \mu_i(t) + 1\right\} = \Psi_i(t), \qquad \text{for } i = 1, \ldots, n; \tag{5}$$

where $\mathbf{1}(A)$ is the indicator function, returning 1 or 0 depending on whether event A is true or false. To prove that a random search sequence is uniformly optimal for all $t \in \mathbb{N}$, we need to show that $\{X(t), \ t \in \mathbb{N}\}$ meets conditions (4) and (5) simultaneously for all $t \in \mathbb{N}$.

For notational convenience, let $\mu_i(0) = 0$ and $\Psi_i(0) = 0$, for $i = 1, \ldots, n$. Define

$$\begin{aligned} \Delta_i(t) &\equiv P\{X(t) = i\} = E\left[\sum_{s=1}^{t} \mathbf{1}(X(s) = i)\right] - E\left[\sum_{s=1}^{t-1} \mathbf{1}(X(s) = i)\right] \\ &= \mu_i(t) + \Psi_i(t) - (\mu_i(t-1) + \Psi_i(t-1)), \end{aligned} \tag{6}$$

so that $\Delta_i(t)$ is the (unconditional) probability that the tth search is allocated to cell i. In other words, based on how the first $t-1$ searches have been allocated, we need to find a way to allocate the tth search, such that the overall probability of allocating it to cell i is $\Delta_i(t)$, and that conditions (4) and (5) are met at time t. Note that given $\Delta_i(t), \ t = 1, 2, \ldots$, it is straightforward to compute

$$\mu_i(t) = \left\lfloor \sum_{s=1}^{t} \Delta_i(s) \right\rfloor, \qquad \Psi_i(t) = \sum_{s=1}^{t} \Delta_i(s) - \mu_i(t),$$

for $i = 1, \ldots, n$.

3.2. The case with two cells

Consider the case with $n = 2$ cells. Given $\Delta_i(t)$, for $i = 1, \ldots, n$ and $t \in \mathbb{N}$, our goal is to determine a state-dependent distribution for which cell to search next, such that conditions (4) and (5) are met for every $t \in \mathbb{N}$.

To begin, consider $t = 1$ (the first search). We can simply set

$$\pi((0, 0), 1) = P\{X(1) = 1\} = \Delta_1(1), \qquad \pi((0, 0), 2) = P\{X(1) = 2\} = \Delta_2(1)$$

to meet conditions (4) and (5). To show that we can generate $X(t)$ to meet conditions (4) and (5) for $t \in \mathbb{N}$, we use mathematical induction on t. Suppose that we have generated $X(s)$ for $s = 1, \ldots, t$, which meet conditions (4) and (5). We next show how to do so for time $t + 1$. For notational convenience, let $\psi_i = \Psi_i(t)$, $\psi_i' = \Psi_i(t+1)$, and $\delta_i = \Delta_i(t+1)$, for $i = 1, 2$. Recall that, by definition, in Eq. (6), for $i = 1, 2$,

$$\psi_i' = \begin{cases} \psi_i + \delta_i, & \text{if } \psi_i + \delta_i < 1, \\ \psi_i + \delta_i - 1, & \text{if } \psi_i + \delta_i \geq 1. \end{cases}$$

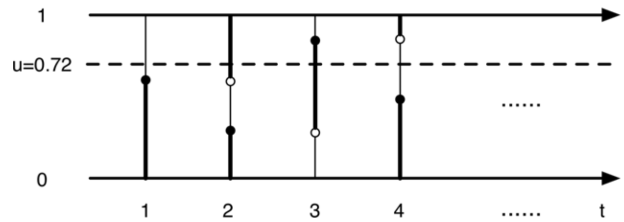

Fig. 1. Generating the search policy with one uniform (0, 1) random variable when $n = 2$. The bold line segments correspond to $I(t)$ defined in Eq. (7). Generate a uniform (0, 1) random variable and draw the corresponding dashed line. At time t, search in cell 1 if the dashed line crosses a bold line segment, or search in cell 2 if it does not. In this example, with $u = 0.72$, the searcher follows the search sequence 2, 1, 1, 2, ...

In addition, $\psi_i, \psi_i' \in [0, 1)$, and $\delta_i \in [0, 1]$, for $i = 1, 2$. Consider three cases:

1. $\psi_1 + \psi_2 = 0$.
   In this case, we have that $\psi_1 = \psi_2 = 0$. At time $t + 1$, search in cell 1 with probability δ_1 and in cell 2 with probability δ_2. The corresponding policy is

$$\pi((\mu_1(t), \mu_2(t)), 1) = \delta_1, \qquad \pi((\mu_1(t), \mu_2(t)), 2) = \delta_2.$$

2. $\psi_1 + \psi_2 = 1$, and $\psi_1' + \psi_2' = 0$.
   Because $\psi_1' = \psi_2' = 0$, we have that $\psi_i + \delta_i = 1$ for $i = 1, 2$. If cell 1 has received $\mu_1(t)$ searches through time t ($\mu_2(t) + 1$ searches in cell 2), then at time $t + 1$ search in cell 1. If cell 1 has received $\mu_1(t) + 1$ searches through time t ($\mu_2(t)$ searches in cell 2), then at time $t + 1$ search in cell 2. The corresponding policy is

$$\pi((\mu_1(t), \mu_2(t) + 1), 1) = 1, \qquad \pi((\mu_1(t), \mu_2(t) + 1), 2) = 0;$$
$$\pi((\mu_1(t) + 1, \mu_2(t)), 1) = 0, \qquad \pi((\mu_1(t) + 1, \mu_2(t)), 2) = 1.$$

3. $\psi_1 + \psi_2 = 1$, and $\psi_1' + \psi_2' = 1$. Through time t, cell i has received at least $\mu_i(t)$ searches, for $i = 1, 2$, with $\mu_1(t) + \mu_2(t) = t - 1$. The additional search is either in cell 1 with probability ψ_1, or in cell 2 with probability ψ_2.
   Since $\psi_1' + \psi_2' = 1$, either $\psi_1 + \delta_1 < 1$ and $\psi_2 + \delta_2 > 1$, or $\psi_1 + \delta_1 > 1$ and $\psi_2 + \delta_2 < 1$; we assume the former case without loss of generality. It then follows that $\mu_1(t + 1) = \mu_1(t)$, $\psi_1' = \psi_1 + \delta_1$, and $\mu_2(t + 1) = \mu_2(t) + 1$, $\psi_2' = \psi_2 + \delta_2 - 1$.
   To allocate the search at time $t + 1$, we need to ensure that through time $t + 1$, cell 1 receives at least $\mu_1(t)$ searches and cell 2 at least $\mu_2(t) + 1$ searches, with the additional search either in cell 1 with probability ψ_1', or in cell 2 with probability ψ_2'. It is then straightforward to verify that the joint probability distribution between $\boldsymbol{x}_t$ and $\boldsymbol{x}_{t+1}$ below meets conditions (4) and (5).

	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$		
$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	(1, 1)	(0, 2)	
(1, 0)	ψ_1	0	ψ_1
(0, 1)	$\psi_2 - \psi_2'$	ψ_2'	ψ_2
	ψ_1'	ψ_2'	

   From the joint probability table, we can determine the corresponding randomized policy:

$$\pi((\mu_1(t) + 1, \mu_2(t)), 1) = 0, \qquad \pi((\mu_1(t) + 1, \mu_2(t)), 2) = 1;$$
$$\pi((\mu_1(t), \mu_2(t) + 1), 1) = 1 - \frac{\psi_2'}{\psi_2}, \qquad \pi((\mu_1(t), \mu_2(t) + 1), 2) = \frac{\psi_2'}{\psi_2}.$$

In summary, if the searcher has a mixed strategy that is uniformly optimal for $\{1, 2, \ldots, t\}$, then we can use the preceding policy $\pi(\boldsymbol{x}_t, \cdot)$ to construct a mixed strategy that is uniformly optimal for $\{1, 2, \ldots, t, t + 1\}$. Hence, Theorem 1 is proved for $n = 2$.

Remark 3. The method just presented can be extended to a search problem with $n \geq 3$ cells, which will be explained in Section 3.3. In the case of $n = 2$ cells, however, there is a rather elegant method that requires only one uniform (0, 1) random variable to produce a search policy that meet conditions (4) and (5) simultaneously for all $t \in \mathbb{N}$. Suppose the generated uniform (0, 1) random variable is u. At $t = 1$, search in cell 1 if $u \leq \Delta_1(1)$; otherwise, search in cell 2. For $t \geq 2$, define

$$I(t) = (\Psi_1(t - 1), \min(1, \Psi_1(t - 1) + \Delta_1(t))] \cup [0, \max(0, \Psi_1(t - 1) + \Delta_1(t) - 1)]. \tag{7}$$

Fig. 1 illustrates this method, where bold line segments indicate $I(t)$, for $t = 1, 2, \ldots$. At $t \geq 2$, let the searcher search in cell 1 if $u \in I(t)$; otherwise, search in cell 2.

As seen in Eq. (7), the probability of searching in cell 1 at time t is $|I(t)| = \Delta_1(t)$. In addition, with the construction as illustrated in Fig. 1, through time t the searcher will search in cell 1 either $\mu_1(t)$ times or $\mu_1(t) + 1$ times. Hence, such a search policy meets conditions (4) and (5) simultaneously for all $t \in \mathbb{N}$.

3.3. The case with three cells

We are given $\Delta_i(t)$, $\mu_i(t)$, and $\Psi_i(t)$, for $i = 1, 2, 3,\ t \in \mathbb{N}$, and need to generate $X(t)$ to meet conditions (4) and (5) for $t \in \mathbb{N}$. Recall that, for $i = 1, 2, 3$,

$$\mu_i(t) = \left\lfloor \sum_{s=1}^{t} \Delta_i(s) \right\rfloor,$$

$$\Psi_i(t) = \sum_{s=1}^{t} \Delta_i(s) - \mu_i(t),$$

where $\mu_i(t)$ is the integral part and $\Psi_i(t)$ is the fractional part of the desired expected number of searches.

To begin, consolidate cells 1 and 2 into a single cell, referred to as cell A, with

$$\Delta_A(t) = \Delta_1(t) + \Delta_2(t),$$
$$\mu_A(t) = \mu_1(t) + \mu_2(t) + \lfloor \Psi_1(t) + \Psi_2(t) \rfloor,$$
$$\Psi_A(t) = \Psi_1(t) + \Psi_2(t) - \lfloor \Psi_1(t) + \Psi_2(t) \rfloor,$$

for $t \in \mathbb{N}$. Between cells A and 3, we can apply the method in Section 3.2 to generate a search sequence that meets conditions (4) and (5) for $t \in \mathbb{N}$. If at time t, the search is allocated to cell 3, then we set $X(t) = 3$. If at time t, the search is allocated to cell A, then we need to decide whether $X(t) = 1$ or $X(t) = 2$.

To show that we can generate $X(t)$ that meet conditions (4) and (5), for $t \in \mathbb{N}$, we use mathematical induction on t. At $t = 1$, if the first search goes to cell A, then assign i to $X(t)$ with probability $\Delta_i(1)/(\Delta_1(1) + \Delta_2(1))$, for $i = 1, 2$. Suppose that we have generated $X(s)$ for $s = 1, \ldots, t$, such that the search policy meets conditions (4) and (5). We next show how to do so at time $t + 1$. For notational convenience, let $\psi_i = \Psi_i(t)$, $\psi_i' = \Psi_i(t + 1)$, and $\delta_i = \Delta_i(t + 1)$, for $i = 1, 2$. Recall that, by definition, for $i = 1, 2$,

$$\psi_i' = \begin{cases} \psi_i + \delta_i, & \text{if } \psi_i + \delta_i < 1, \\ \psi_i + \delta_i - 1, & \text{if } \psi_i + \delta_i \geq 1. \end{cases}$$

In addition, $\psi_i, \psi_i' \in [0, 1)$, and $\delta_i \in [0, 1]$, for $i = 1, 2$. Consider two cases based on the value of $\psi_1 + \psi_2$.

3.3.1. The case $0 \leq \psi_1 + \psi_2 < 1$

In this case, we have $\Psi_A(t) = \psi_1 + \psi_2$ and $\mu_A(t) = \mu_1(t) + \mu_2(t)$. With the induction hypothesis, through time t, cells 1 and 2 have received $(\mu_1(t), \mu_2(t))$ searches with probability $1 - \Psi_A(t) = 1 - \psi_1 - \psi_2$, or $(\mu_1(t) + 1, \mu_2(t))$ searches with probability ψ_1, or $(\mu_1(t), \mu_2(t) + 1)$ searches with probability ψ_2. Denote these three cases at time t by (0, 0), (1, 0), and (0, 1), respectively. Consider three cases based on δ_1 and δ_2.

1. $\psi_1 + \delta_1 < 1$, $\psi_2 + \delta_2 < 1$, and $\psi_1 + \delta_1 + \psi_2 + \delta_2 < 1$.
   In this case, we have $\psi_1' = \psi_1 + \delta_1$, $\psi_2' = \psi_2 + \delta_2$, and $\psi_1' + \psi_2' < 1$. Hence, $\mu_A(t+1) = \mu_A(t) = \mu_1(t) + \mu_2(t)$ and $\Psi_A(t+1) = \psi_1' + \psi_2'$. Through time $t + 1$, cell A will receive either $\mu_A(t)$ searches with probability $1 - \Psi_A(t + 1) = 1 - \psi_1' - \psi_2'$, or $\mu_A(t) + 1$ searches with probability $\Psi_A(t + 1) = \psi_1' + \psi_2'$. At time $t + 1$, the search allocation between cells 1 and 2 needs to be either $(\mu_1(t), \mu_2(t))$, $(\mu_1(t) + 1, \mu_2(t))$, or $(\mu_1(t), \mu_2(t) + 1)$. Denote these three cases at time $t + 1$ by (0, 0), (1, 0), and (0, 1), respectively. It is then straightforward to verify that the joint probability distribution below between $\boldsymbol{x}_t$ and $\boldsymbol{x}_{t+1}$ meets all these constraints. By allocating the search at time $t + 1$ according to this joint probability distribution, the resulting search policy meets conditions (4) and (5) at time $t + 1$.

$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$: (0, 0)	(1, 0)	(0, 1)	
(0, 0)	$1 - \psi_1' - \psi_2'$	$\psi_1' - \psi_1$	$\psi_2' - \psi_2$	$1 - \psi_1 - \psi_2$
(1, 0)	0	ψ_1	0	ψ_1
(0, 1)	0	0	ψ_2	ψ_2
	$1 - \psi_1' - \psi_2'$	ψ_1'	ψ_2'	

2. $\psi_1 + \delta_1 < 1$, $\psi_2 + \delta_2 < 1$, and $\psi_1 + \delta_1 + \psi_2 + \delta_2 \geq 1$.
   In this case, we have $\psi_1' = \psi_1 + \delta_1$, $\psi_2' = \psi_2 + \delta_2$, and $\psi_1' + \psi_2' \geq 1$. Hence, $\mu_A(t + 1) = \mu_A(t) + 1 = \mu_1(t) + \mu_2(t) + 1$ and $\Psi_A(t + 1) = \psi_1' + \psi_2' - 1$. The joint probability distribution below produces a search policy that meets conditions (4) and (5) at time $t + 1$.

$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$: (1, 0)	(0, 1)	(1, 1)	
(0, 0)	$1 - \psi_2' - \psi_1(1 - \frac{\psi_1' + \psi_2' - 1}{\psi_1 + \psi_2})$	$1 - \psi_1' - \psi_2(1 - \frac{\psi_1' + \psi_2' - 1}{\psi_1 + \psi_2})$	0	$1 - \psi_1 - \psi_2$
(1, 0)	$\psi_1(1 - \frac{\psi_1' + \psi_2' - 1}{\psi_1 + \psi_2})$	0	$\psi_1 \frac{\psi_1' + \psi_2' - 1}{\psi_1 + \psi_2}$	ψ_1
(0, 1)	0	$\psi_2(1 - \frac{\psi_1' + \psi_2' - 1}{\psi_1 + \psi_2})$	$\psi_2 \frac{\psi_1' + \psi_2' - 1}{\psi_1 + \psi_2}$	ψ_2
	$1 - \psi_2'$	$1 - \psi_1'$	$\psi_1' + \psi_2' - 1$	

3. $\psi_1 + \delta_1 < 1$, $\psi_2 + \delta_2 \geq 1$ (the case $\psi_1 + \delta_1 \geq 1$, $\psi_2 + \delta_2 < 1$ can be dealt with in a similar way).

In this case, we have $\psi_1' = \psi_1 + \delta_1$, $\psi_2' = \psi_2 + \delta_2 - 1$, and $\psi_1' + \psi_2' < 1$. Hence, $\mu_A(t+1) = \mu_A(t) + 1 = \mu_1(t) + \mu_2(t) + 1$ and $\Psi_A(t+1) = \psi_1' + \psi_2'$. The joint probability distribution below produces a search policy that meets conditions (4) and (5) at time $t+1$.

	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$			
$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	$(0, 1)$	$(1, 1)$	$(0, 2)$	
$(0, 0)$	$1 - \psi_1 - \psi_2$	0	0	$1 - \psi_1 - \psi_2$
$(1, 0)$	0	ψ_1	0	ψ_1
$(0, 1)$	$\psi_1 + \psi_2 - \psi_1' - \psi_2'$	$\psi_1' - \psi_1$	ψ_2'	ψ_2
	$1 - \psi_1' - \psi_2'$	ψ_1'	ψ_2'	

3.3.2. The case $1 \le \psi_1 + \psi_2 < 2$

In this case, $\Psi_A(t) = \psi_1 + \psi_2 - 1$, and $\mu_A(t) = \mu_1(t) + \mu_2(t) + 1$. With the induction hypothesis through time t, cells 1 and 2 have received $(\mu_1(t)+1, \mu_2(t))$ searches with probability $1-\psi_2$, or $(\mu_1(t), \mu_2(t)+1)$ searches with probability $1-\psi_1$, or $(\mu_1(t)+1, \mu_2(t)+1)$ searches with probability $\Psi_A(t) = \psi_1 + \psi_2 - 1$. Denote these three cases at time t by $(1, 0)$, $(0, 1)$, and $(1, 1)$, respectively. Consider four cases based on δ_1 and δ_2.

1. $\psi_1 + \delta_1 < 1$, $\psi_2 + \delta_2 < 1$.
In this case, we have $\psi_1' = \psi_1 + \delta_1$, $\psi_2' = \psi_2 + \delta_2$, and $1 \le \psi_1' + \psi_2' < 2$. Hence, $\mu_A(t+1) = \mu_A(t) = \mu_1(t) + \mu_2(t) + 1$ and $\Psi_A(t+1) = \psi_1' + \psi_2' - 1$. The joint probability distribution below produces a search policy that meets conditions (4) and (5) at time $t+1$.

	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$			
$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	$(1, 0)$	$(0, 1)$	$(1, 1)$	
$(1, 0)$	$1 - \psi_2'$	0	$\psi_2' - \psi_2$	$1 - \psi_2$
$(0, 1)$	0	$1 - \psi_1'$	$\psi_1' - \psi_1$	$1 - \psi_1$
$(1, 1)$	0	0	$\psi_1 + \psi_2 - 1$	$\psi_1 + \psi_2 - 1$
	$1 - \psi_2'$	$1 - \psi_1'$	$\psi_1' + \psi_2' - 1$	

2. $\psi_1 + \delta_1 < 1$, $\psi_2 + \delta_2 \ge 1$, and $\psi_1' + \psi_2' < 1$ (the case $\psi_1 + \delta_1 \ge 1$, $\psi_2 + \delta_2 < 1$ can be dealt with in a similar way).
In this case, we have $\psi_1' = \psi_1 + \delta_1$, $\psi_2' = \psi_2 + \delta_2 - 1$, and $\psi_1' + \psi_2' < 1$. Hence, $\mu_A(t+1) = \mu_A(t) = \mu_1(t) + \mu_2(t) + 1$ and $\Psi_A(t+1) = \psi_1' + \psi_2'$. The joint probability distribution below produces a search policy that meets conditions (4) and (5) at time $t+1$.

	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$			
$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	$(0, 1)$	$(1, 1)$	$(0, 2)$	
$(1, 0)$	0	$1 - \psi_2$	0	$1 - \psi_2$
$(0, 1)$	$1 - \psi_1' - \psi_2'$	$\psi_1' - \psi_1$	ψ_2'	$1 - \psi_1$
$(1, 1)$	0	$\psi_1 + \psi_2 - 1$	0	$\psi_1 + \psi_2 - 1$
	$1 - \psi_1' - \psi_2'$	ψ_1'	ψ_2'	

3. $\psi_1 + \delta_1 < 1$, $\psi_2 + \delta_2 \ge 1$, and $\psi_1' + \psi_2' \ge 1$ (the case $\psi_1 + \delta_1 \ge 1$, $\psi_2 + \delta_2 < 1$ can be dealt with in a similar way).
In this case, we have $\psi_1' = \psi_1 + \delta_1$, $\psi_2' = \psi_2 + \delta_2 - 1$, and $1 \le \psi_1' + \psi_2' < 2$. Hence, $\mu_A(t+1) = \mu_A(t) + 1 = \mu_1(t) + \mu_2(t) + 2$ and $\Psi_A(t+1) = \psi_1' + \psi_2' - 1$. The joint probability distribution below produces a search policy that meets conditions (4) and (5) at time $t+1$.

	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$			
$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	$(1, 1)$	$(0, 2)$	$(1, 2)$	
$(1, 0)$	$1 - \psi_2$	0	0	$1 - \psi_2$
$(0, 1)$	$\psi_1' - \psi_1$	$1 - \psi_1'$	0	$1 - \psi_1$
$(1, 1)$	$\psi_1 + \psi_2 - \psi_1' - \psi_2'$	0	$\psi_1' + \psi_2' - 1$	$\psi_1 + \psi_2 - 1$
	$1 - \psi_2'$	$1 - \psi_1'$	$\psi_1' + \psi_2' - 1$	

4. $\psi_1 + \delta_1 \ge 1$, $\psi_2 + \delta_2 \ge 1$.
In this case, we have $\psi_1' = \psi_1 + \delta_1 - 1$, $\psi_2' = \psi_2 + \delta_2 - 1$, and $\psi_1' + \psi_2' < 1$. Hence, $\mu_A(t+1) = \mu_A(t) + 1 = \mu_1(t) + \mu_2(t) + 2$ and $\Psi_A(t+1) = \psi_1' + \psi_2'$. The joint probability distribution below produces a search policy that meets conditions (4) and (5) at time $t+1$.

	$\boldsymbol{x}_{t+1} - \boldsymbol{\mu}(t)$			
$\boldsymbol{x}_t - \boldsymbol{\mu}(t)$	$(1, 1)$	$(2, 1)$	$(1, 2)$	
$(1, 0)$	$1 - \psi_2$	0	0	$1 - \psi_2$
$(0, 1)$	$1 - \psi_1$	0	0	$1 - \psi_1$
$(1, 1)$	$\psi_1 + \psi_2 - \psi_1' - \psi_2' - 1$	ψ_1'	ψ_2'	$\psi_1 + \psi_2 - 1$
	$1 - \psi_1' - \psi_2'$	ψ_1'	ψ_2'	

3.4. The case with four or more cells

Denote the consolidation of cells 1 through k as A_k, for $k = 1, 2, 3, \ldots, n-1$. First, consider cell A_{n-1} and cell n, and use the method in Section 3.2 to generate a search sequence between these 2 cells. Second, use the method in Section 3.3 to split the searches in cell A_{n-1} between cell A_{n-2} and cell $n-1$. Repeat the procedure to split the searches in cell A_k between cell A_{k-1} and cell k, for $k = n-1, n-2, \ldots, 2$, to have a complete search sequence.

4. The hider's uniformly optimal strategy

Theorem 4. *For any number of cells n, the hider has a uniformly optimal mixed strategy for the games $G(t)$, $t \in \mathbb{N}$, if and only if the detection probabilities α_i are the same for $i = 1, \dots, n$.*

Proof. If α_i are the same for $i = 1, \dots, n$, then the hider's mixed strategy with $p_i = 1/n$, for $i = 1, \dots, n$, is optimal for $G(t)$ for any t, thus uniformly optimal for $G(t)$, $t \in \mathbb{N}$.

Now suppose that $\alpha_i < \alpha_j$ for some i, j, and relabel them as $\alpha_1 < \alpha_2$, without loss of generality. For $G(t)$, Subelman [14] showed that the hider's optimal strategy is to hide in cell i with probability

$$p_i(t) = \frac{(1-\alpha_i)^{-\mu_i(t)}/\alpha_i}{\sum_{k=1}^{n}(1-\alpha_k)^{-\mu_k(t)}/\alpha_k}, \tag{8}$$

where $\mu_i(t)$ is the minimal number of searches in cell i in the searcher's optimal mixed strategy for $G(t)$, as defined in Theorem 2. In particular, for $G(1)$, since the searcher must randomize his only search, it follows that $\mu_i(1) = 0$ for $i = 1, \dots, n$, so Eq. (8) reduces to

$$p_i(1) = \frac{1/\alpha_i}{\sum_{k=1}^{n} 1/\alpha_k},$$

which is also the hider's strategy that ensures the same value, regardless of where the searcher chooses to allocate his only search.

In addition, since $\alpha_1 < \alpha_2$, it is straightforward to verify that $\mu_1(t) \geq \mu_2(t)$ for all $t \in \mathbb{N}$. Pick a time s such that $\mu_1(s) = 1$. Whether $\mu_2(s) = 0$ or 1, using Eq. (8) we can conclude that

$$\frac{p_1(s)}{p_2(s)} \neq \frac{1/\alpha_1}{1/\alpha_2} = \frac{p_1(1)}{p_2(1)},$$

showing that the hider's optimal strategy for $G(s)$ is different from that for $G(1)$. Consequently, the hider does not have a uniformly optimal strategy for $G(t)$, $t \in \mathbb{N}$. □

5. Concluding remarks

Since the searcher always has a uniformly optimal strategy for $G(t)$, $t \in \mathbb{N}$, the searcher can maximize the probability of finding the hider by an arbitrary deadline, even if the deadline is not known. This result has crucial significance in practice. The search and rescue team can maximize the probability of finding the survivors before they die, even if Nature knows the search deadline and hides the survivors to minimize that probability. The following corollary is an immediate result from Theorem 1.

Corollary 5. *Consider the game where Nature chooses the time horizon t with probability b_t and reveals it to the hider alone. The searcher wins if the searcher finds the hider by this time t (payoff 1); otherwise the hider wins (payoff 0). The value of this game is given by*

$$\sum_{t=1}^{\infty} b_t\, v(t),$$

where $v(t)$ is the value of $G(t)$, $t \in \mathbb{N}$. In this game, the hider plays optimally in $G(t)$ for the chosen t, while the searcher plays his uniformly optimal strategy for $G(t)$, $t \in \mathbb{N}$. □

Consider the same search model with a different objective function—the expected time to detection. Write $\Gamma(\infty)$ for the two-person zero-sum game where the searcher wants to minimize the expected time until finding the hider, while the hider wants to maximize it. Gittins and Roberts [6,11] derived the optimal hider's mixed strategy for some special cases, and presented algorithms to compute it for other cases. Generally speaking, the searcher's uniformly optimal strategy for $G(t)$, $t \in \mathbb{N}$ need not be optimal for $\Gamma(\infty)$. The next corollary presents a special case, which, according to Theorem 4, applies only when α_i are the same for $i = 1, \dots, n$.

Corollary 6. *If the hider has a uniformly optimal strategy for games $G(t)$, $t \in \mathbb{N}$, then the uniformly optimal strategy for either player is also optimal for the game $\Gamma(\infty)$.*

Proof. Write $v(t)$ for the value of $G(t)$, which is the probability of finding the hider in t searches. Since the searcher's uniformly optimal strategy attains $v(t)$ in $G(t)$, regardless where the hider hides, we can define X as the corresponding number of searches taken until the searcher finds the hider in $\Gamma(\infty)$. Regardless where the hider hides, the searcher attains an expected time to detection

$$E[X] = \sum_{t=0}^{\infty} P\{X > t\} = 1 + \sum_{t=1}^{\infty}(1 - v(t)).$$

If the hider also has a uniformly optimal strategy for games $G(t)$, $t \in \mathbb{N}$, then the hider can guarantee that the probability of getting found by t is at most $v(t)$. Hence, the hider can guarantee that the expected time to detection is at least $1 + \sum_{t=1}^{\infty}(1 - v(t))$. Consequently, the value of the game $\Gamma(\infty)$ is $1 + \sum_{t=1}^{\infty}(1 - v(t))$, and the result is proved. □

Acknowledgment

The authors are indebted to an anonymous reviewer whose insightful suggestions substantially strengthened this paper.

References

[1] S. Alpern, R. Fokkink, L. Gasieniec, R. Lindelauf, V.S. Subrahmanian (Eds.), Search Theory, Springer, 2013.
[2] S. Alpern, S. Gal, The Theory of Search Games and Rendezvous, Kluwer Academic Publishers, 2003.
[3] W.L. Black, Discrete sequential search, Inf. Control 8 (2) (1965) 159–162.
[4] M.C. Chew Jr., A sequential search procedure, Ann. Math. Statist. 38 (2) (1967) 494–502.
[5] M.C. Chew Jr., Optimal stopping in a discrete search problem, Oper. Res. 21 (3) (1973) 741–747.
[6] J.C. Gittins, D.M. Roberts, The search for an intelligent evader concealed in one of an arbitrary number of regions, Naval Res. Logist. Q. 26 (4) (1979) 651–666.
[7] J.B. Kadane, Discrete search and the Neyman–Pearson Lemma, J. Math. Anal. Appl. 22 (1) (1968) 156–171.
[8] J.B. Kadane, Optimal whereabouts search, Oper. Res. 19 (4) (1971) 894–904.
[9] M. Kress, K.Y. Lin, R. Szechtman, Optimal discrete search with imperfect specificity, Math. Methods Oper. Res. 68 (3) (2008) 539–549.
[10] D. Matula, A periodic optimal search, Amer. Math. Monthly 71 (1) (1964) 15–21.
[11] D. Roberts, J. Gittins, The search for an intelligent evader: Strategies for searcher and evader in the two-region problem, Naval Res. Logist. Q. 25 (1) (1978) 95–106.
[12] S.M. Ross, A problem in optimal search and stop, Oper. Res. 17 (6) (1969) 984–992.
[13] L. Stone, Theory of Optimal Search, INFORMS, 2004.
[14] E. Subelman, A hide-search game, J. Appl. Probab. 18 (3) (1981) 628–640.
[15] A.R. Washburn, Search and Detection, fourth ed., INFORMS, 2002.
[16] I. Wegener, The discrete sequential search problem with nonrandom cost and overlook probabilities, Math. Oper. Res. 5 (3) (1980) 373–380.