

Optimal Pure Strategies for a Discrete Search Game

Thuy Bui*

Thomas Lidbetter†

Kyle Y. Lin §

Abstract

Consider a two-person zero-sum search game between a Hider and a Searcher. The Hider chooses to hide in one of n discrete locations (or “boxes”) and the Searcher chooses a search sequence specifying which order to look in these boxes until finding the Hider. A search at box i takes t_i time units and finds the Hider—if hidden there—Independently with probability q_i , for $i = 1, \dots, n$. The Searcher wants to minimize the expected total time needed to find the Hider, while the Hider wants to maximize it. It is shown in the literature that the Searcher has an optimal search strategy that mixes up to n distinct search sequences with appropriate probabilities. This paper investigates the existence of optimal pure strategies for the Searcher—a single deterministic search sequence that achieves the optimal expected total search time regardless of where the Hider hides. We identify several cases in which the Searcher has an optimal pure strategy, and several cases in which such optimal pure strategy does not exist. An optimal pure search strategy has significant practical value because the Searcher does not need to randomize their actions and will avoid second guessing themselves if the chosen search sequence from an optimal mixed strategy does not turn out well.

Keywords: search games, zero-sum games, semi-infinite games, Gittins index

*Rutgers Business School, 1 Washington Park, Newark, NJ 07102, USA, tb680@business.rutgers.edu

†Rutgers Business School, 1 Washington Park, Newark, NJ 07102, USA, tlidbetter@business.rutgers.edu

‡Department of Systems and Information Engineering, University of Virginia, VA 22903, USA, tlidbetter@virginia.edu

§Operations Research Department, Naval Postgraduate School, Monterey, CA 93943

1 Introduction

A major criticism of game theory is the notion of randomized (or *mixed*) strategies: see, for example, Binmore (2008) and Gintis (2014). In a Nash equilibrium, players randomize between strategies in a precisely specified way, yet they are indifferent between each of the strategies between which they randomize, given the behavior of the other players. Intuitively, it is difficult to explain such precision in the players’ random actions. Harsanyi (1973) had an elegant solution to this criticism, by defining a Bayesian game where the payoffs of a given game are slightly perturbed by a random amount. Any equilibrium in the original game can “almost always” be obtained as a limit of a pure strategy equilibria in the Bayesian game, as the perturbations go to zero.

This is all very well, but in practice, if a Nash equilibrium requires a player to randomize, there are also issues about the practical implementation of such randomized strategies. This is particularly salient in search games, which can be used to model military operations where lives may be at stake and randomized actions can be hard to justify to decision makers. When using an optimal mixed strategy, there is a chance that the chosen pure strategy turns out to fare poorly against the Hider’s chosen location, which may give the false appearance to a third party that the Searcher made a *wrong* decision. If there is a choice between a mixed strategy solution and a pure strategy (non-randomized) solution, we argue that the latter would be preferable, since it is more easily implementable and explainable. Decision makers do not need to rely on the roll of a die or outcome of a roulette wheel to make their choices, and would save the trouble second guessing the outcome of the randomization device if the chosen pure strategy performs poorly against the other player’s choice.

In this paper we consider the problem of finding optimal pure strategies for a generalization of a classical search game introduced by Bram (1963), where a Hider chooses one of a finite number n of locations (or “boxes”) in which to hide and a Searcher inspects the boxes one-by-one until finding the Hider. When a given box is searched, if the Hider is in that box, there is a specified probability (the “detection probability”) that the Hider is found. The Searcher aims to minimize the expected number of boxes to search before finding the Hider, whereas the Hider aims to maximize it. The game is zero-sum, and Bram (1963) proved that an optimal (min-max) mixed search strategy exists and suggested a numerical way to compute it. There has been sporadic interest in the game, but not much is known in terms of closed form solutions. Indeed, it seems unlikely that such solutions

exist in general. Roberts and Gittins (1978) found an optimal Hider strategy in the case of two boxes under certain, rather specific, conditions and Ruckle (1991) found an optimal search strategy for the case that the detection probabilities are the same for all boxes.

There has been some recent interest in a more general version of the game described above, where each box takes a certain amount of time to search. The payoff becomes the total expected time to find the Hider. Clarkson *et al.* (2023) showed that optimal strategies exist and that there exists an optimal Searcher strategy that mixes between at most n pure strategies. Clarkson and Lin (2023) presented an algorithm to numerically calculate optimal strategies for the players.

The question of whether there exist optimal *pure* strategies in this game has received little attention. Ruckle (1991) pointed out that for the case of equal search times and detection probabilities all equal to $1/2$, there is an optimal pure strategy, but prior to this paper, there were no other known cases for which there is an optimal pure strategy. We rectify this by showing that there are optimal pure strategies for several classes of the game, and we show how to calculate them. We also give conditions under which there is no optimal pure strategy. Roughly speaking, this is the case when the detection probabilities are high. In this case, there is a high probability of discovery after the first box has been searched, which immediately skews the expected search time (in the negative direction) towards this box.

The game can be motivated by problems of national security. For example, suppose the military are trying to use a satellite to locate a “safe house” chosen by a terrorist group. Several possible locations have been identified, and a search at a location corresponds to taking a fixed number of photographs, which will reveal the safe house (if it is there) with some probability. The objective is to minimize the expected time to find the safe house. Another example, which may be considered as a game against Nature, is the search and rescue problem of finding a plane that has lost communication before crashing into the sea or a desert. There are several possible crash sites which take varying times to search. Another game against Nature that may fit into this framework is that of a scatter hoarder, such as a squirrel attempting to find a cache of nuts that he or some other squirrel has previously hidden in one of a number of possible locations.

This paper is laid out as follows. In Section 2, we recall the definition of the game and some previous results. Although this paper is primarily on the topic of optimal pure strategies, we also give a full solution to the game in mixed strategies for the case of equal detection probabilities in Section 3. This case was previously unsolved. In Section 4, we use a series of lemmas to prove our

main result, Theorem 10, which gives conditions under which a pure strategy solution to the game exists. We then apply Theorem 10 to several classes of the game. We end Section 4 by proving conditions under which no pure strategy solution exists, and we explore the consequences of this result. In Section 5, we restrict attention to the case of two boxes, taking advantage of the simpler structure to strengthen our results about the existence of pure strategy solutions. We go on to consider the question of what is a “best” pure strategy in the case that no optimal pure strategy exists. In Section 6 we conclude.

2 Preliminaries

In this section, we recall the definition of the game G and state some previous results that we will use later on in the paper.

2.1 Definition of the Search Game

The game G is played between a Hider (the maximizer) and a Searcher (the minimizer). The Hider’s pure strategies are a set of n boxes, denoted $[n] \equiv \{1, \dots, n\}$, in which he may hide. Hence, the set of mixed strategies for the Hider is

$$\Delta^n \equiv \{(p_1, \dots, p_n) : p_i \geq 0 \text{ for } i = 1, \dots, n \text{ and } \sum_{i=1}^n p_i = 1\}.$$

Here, p_i is the probability the Hider chooses to hide in box i .

A pure strategy for the Searcher is a search sequence—an infinite sequence of the boxes corresponding to the order in which the Searcher searches them. Thus, the Searcher’s pure strategy set is the set $\mathcal{C} \equiv [n]^\infty$, an infinite set. As in Clarkson *et al.* (2023), we define a mixed strategy for the Searcher as a probabilistic choice of a *countable* subset of pure strategies. More precisely, it is a function $\theta : \mathcal{C} \rightarrow [0, 1]$ such that $\{\xi \in \mathcal{C} : \theta(\xi) > 0\}$ is countable and

$$\sum_{\xi \in \mathcal{C}} \theta(\xi) = 1.$$

The condition that $\{\xi \in \mathcal{C} : \theta(\xi) > 0\}$ is countable ensures that this sum is well-defined.

Each box i has a *search time* $t_i > 0$, which is the time required to complete one search of the box. Box i also has a *detection probability* $q_i \in (0, 1]$, which is the probability the Searcher will find the Hider after one search of box i if the Hider is hidden there. Write $r_i = 1 - q_i$ for the overlook

probability of box i . The outcome of each search is independent of all previous search outcomes. For each pure strategy pair (i, ξ) —when the Hider hides in box i and the Searcher uses the search sequence ξ —we write $u(i, \xi)$ for the payoff of the game, which is the expected time needed for the Searcher to find the Hider under the pure strategy pair (i, ξ) . We refer to the payoff as the *expected search time*. When the players use mixed strategies \mathbf{p} and θ , we extend u to denote the expected payoff as follows.

$$u(\mathbf{p}, \theta) \equiv \sum_{i=1}^n \sum_{\xi \in \mathcal{C}} p_i \theta(\xi) u(i, \xi).$$

We similarly extend the definition of u for cases when one player uses a mixed strategy and the other uses a pure strategy.

For a given Hider mixed strategy \mathbf{p} , the problem of finding a best response—that is, a Searcher strategy ξ that minimizes the expected search time $\sum_{i=1}^n p_i u(i, \xi)$ against \mathbf{p} —is well understood. The solution, first discovered by Blackwell (reported in Matula, 1964), can be found recursively. The Searcher should first search any box i that maximizes the probability of detection per unit time, namely $p_i q_i / t_i$. After each subsequent search, the next box to be searched is found by updating the hiding probabilities p_i according to Bayes' law and repeating the calculation. It is straightforward to show that after box i has been searched m_i times for $i = 1, \dots, n$, the next box to be searched should be some i that maximizes the index

$$\psi_i \equiv \frac{p_i q_i r_i^{m_i}}{t_i}. \quad (1)$$

In order to study optimal pure strategies, in this paper, we assume there are some positive coprime integers k_1, \dots, k_n and some $r > 0$ such that $r_i^{k_i} = r$ for all $i \in [n]$. It is worth noting that this assumption is not at all restrictive because if it were not the case, we could find k_1, \dots, k_n and r such that $r_i^{k_i}$ is arbitrarily close to r for each i . With $r_i^{k_i} = r$ for all $i \in [n]$, the index in (1) becomes

$$\psi_i \equiv \frac{p_i q_i r^{m_i/k_i}}{t_i}. \quad (2)$$

Following the terminology of Clarkson *et al.* (2023), we refer to any sequence that can be produced in this way (and is hence a best response to some Hider strategy \mathbf{p}) as a *Gittins search sequence* (against \mathbf{p}). The name comes from a comment by Kelly in Gittins (1979) noting that Blackwell's solution is equivalent to a Gittins index policy obtained by modeling the search as a tractable version of the multiarmed bandit problem (see Gittins *et al.*, 2011).

There may be multiple Gittins search sequences against a given \mathbf{p} in the case that there are ties between the indices during the search process. Any permutation σ of $[n]$ can be used to define a tie-breaking rule, where each tie is broken by choosing the box that appears first in σ . We call a Gittins search sequence *consistent* if all ties are broken using the same tie-breaking rule given by some permutation σ of $[n]$. There are at most $n!$ consistent Gittins search sequences against any given \mathbf{p} , and two different permutations may result in the same consistent Gittins search sequence.

A particular Hider strategy of interest is the *equalizing strategy*, defined as the strategy that makes the Searcher indifferent between all the boxes at the beginning of the search. We define the Hider's equalizing strategy below.

Definition 1 The Hider's equalizing strategy \mathbf{p}^* is given by

$$p_i^* = \frac{t_i/q_i}{\sum_{j=1}^n t_j/q_j}, \quad i \in [n].$$

The Hider's equalizing strategy \mathbf{p}^* appears to be a strong strategy because it makes all boxes equally attractive (or unattractive) to the Searcher at the beginning of the search. In particular, there are $n!$ distinct consistent Gittins search sequences against \mathbf{p}^* . Roberts and Gittins (1978) showed that the Hider's equalizing strategy is indeed optimal in certain cases. We recite the theorem below, which will be used later.

Theorem 2 (Theorem 2 of Roberts and Gittins (1978)) Consider the game G with $n = 2$ and $t_1 = t_2$. If $r_1^k = r_2^{k+1}$ for some positive integer k , where $r_i \equiv 1 - q_i$ for $i = 1, 2$, then the Hider's equalizing strategy \mathbf{p}^* is optimal if and only if $k \leq 12$.

If the Hider uses the equalizing strategy, then the indices $p_i^* q_i / t_i$ are all equal to some λ , and the index ψ_i from (2) reduces to $\psi_i = \lambda r^{m_i/k_i}$. In this case, it is easy to see that any Gittins search sequence against \mathbf{p}^* will be of the form s_1, s_2, \dots , where each $s_j, j = 1, 2, \dots$ is a subsequence containing box i exactly k_i times for each $i \in [n]$. Any *consistent* Gittins search sequence against \mathbf{p}^* will be an infinite repetition of one such subsequence s_j . We denote such a strategy (s_j) . The length of time each subsequence s_j spends searching is equal to $\sum_{i \in [n]} k_i t_i$, and we denote this quantity by T .

2.2 The Value of the Game and Optimal Strategies

As discussed in Clarkson *et al.* (2023), it follows from standard results on semi-infinite games that the game G has a value. That is, there exists some V such that

$$V = \max_{\mathbf{p} \in \Delta^n} \inf_{\xi \in \mathcal{C}} u(\mathbf{p}, \xi) = \inf_{\theta} \max_{i \in [n]} u(i, \theta),$$

where the infimum on the right-hand side is taken over all Searcher mixed strategies θ .

It also follows from standard results that there exists an optimal Hider strategy in G (that is, a mixed strategy for the Hider that guarantee a payoff of at least V) and an ε -optimal Searcher strategies (that is, mixed strategies for the Searcher that guarantee a payoff of at most $V + \varepsilon$, for any $\varepsilon > 0$).

A major result from Clarkson *et al.* (2023), which we will use later in this paper, is a stronger statement on the existence of optimal Searcher strategies.

Theorem 3 (Theorem 3 of Clarkson *et al.* (2023)) There exists an optimal mixed strategy for the Searcher in G that mixes between at most n pure strategies that are all consistent Gittins search sequences against any optimal Hider strategy.

3 The Case of Equal Detection Probabilities

In this section, we consider the game with $q_i = q$ for all $i \in [n]$, for some $q \in (0, 1)$, but arbitrary search times t_1, \dots, t_n . In the special case where $t_1 = \dots = t_n$, Ruckle (1991) found that the Hider's equalizing strategy $\mathbf{p}^* = (1/n, 1/n, \dots, 1/n)$ is optimal. For the Searcher, it is optimal to choose with equal probability each of the search sequences $(i, i+1, \dots, n, 1, 2, \dots, i-1)$, for $i = 1, \dots, n$.

In another special case with $q = 1$, each box need only be searched once to ensure that the Hider is found with probability 1. In this case, the Hider's equalizing strategy \mathbf{p}^* chooses each box i with probability proportional to t_i , and any permutation of $[n]$ is a Gittins search sequence against \mathbf{p}^* . The parameter T , defined at the end of Subsection 2.1 reduces to $T = \sum_{i \in [n]} t_i$. This special case was first considered in Condon *et al.* (2009), and independently in Alpern and Lidbetter (2013) and Lidbetter (2013). The optimal strategies found for the Searcher were different in each of these works, but the optimal Hider strategy is unique (see Alpern and Lidbetter, 2013). We summarize the solution below, giving the optimal Searcher strategy from Condon *et al.* (2009), which is arguably the most elegant of the three.

Theorem 4 In the game G with $q_i = 1$ for all $i \in [n]$, it is optimal for the Hider to hide in box i with probability proportional to t_i . It is optimal for the Searcher to choose the first box i she looks in with probability proportional to t_i , then to search the remaining boxes in the order $i+1, i+2, \dots, n, 1, 2, \dots, i-1$. The value V of the game is given by

$$V = T - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j.$$

Our contribution in this paper is to extend Theorem 4 to the case in which $q < 1$ in the next theorem.

Theorem 5 Denote by $\theta(\sigma)$ the probability of choosing permutation σ of $[n]$ for some arbitrary optimal Searcher strategy for the game G with $q = 1$. For the game with equal detection probabilities $q \in (0, 1)$, the Searcher strategy that chooses each search sequence (σ) —which repeats σ indefinitely—with probability $\theta(\sigma)$ is optimal for the Searcher. The Hider’s equalizing strategy \mathbf{p}^* —choosing box i with probability $p_i^* = t_i/T$ —is optimal. The value of the game is

$$V = \frac{T}{q} - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j.$$

Proof. First suppose the Hider uses the strategy \mathbf{p}^* , and let ξ be an arbitrary Searcher pure strategy which is a best response to \mathbf{p}^* . It must be a Gittins search sequence against \mathbf{p}^* , so it must be of the form $\sigma_1, \sigma_2, \dots$, where each $\sigma_k, k \geq 1$ is a permutation of $[n]$, and $\sigma_k(i)$ denotes the position of box i in the permutation σ_k .

Note that the expected search time of any permutation σ against \mathbf{p}^* in the case that the detection probabilities are all equal to 1 is given by Theorem 4. Also, by definition of this expected search time, we obtain the following identity, for any permutation σ .

$$\sum_{i=1}^n p_i^* \sum_{\sigma(j) \leq \sigma(i)} t_j = T - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j. \quad (3)$$

Now, the expected search time of ξ against \mathbf{p}^* is

$$\begin{aligned} u(\mathbf{p}^*, \xi) &= \sum_{k=1}^{\infty} (1-q)^{k-1} q \left((k-1)T + \sum_{i=1}^n p_i^* \sum_{\sigma_k(j) \leq \sigma_k(i)} t_j \right) \\ &= \sum_{k=1}^{\infty} (1-q)^{k-1} q \left(kT - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j \right) \quad (\text{by (3)}) \\ &= \frac{T}{q} - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j. \end{aligned}$$

Thus, the value satisfies $V \geq T/q - (1/T) \sum_{1 \leq i < j \leq n} t_i t_j$.

We now turn to the Searcher strategy given by θ . Again, Theorem 4 gives an expression for the expected search time of the strategy θ against any fixed Hider strategy i when all the detection probabilities are equal to 1, and using the definition of this expected search time we obtain the following.

$$\sum_{\sigma} \theta(\sigma) \sum_{\sigma(j) \leq \sigma(i)} t_j = T - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j, \quad (4)$$

where the sum on the left is over all permutations σ . But the expected search time of θ against i is

$$\begin{aligned} u(i, \theta) &= \sum_{k=1}^{\infty} (1-q)^{k-1} q \left((k-1)T + \sum_{\sigma} \theta(\sigma) \sum_{\sigma(j) \leq \sigma(i)} t_j \right) \\ &= \sum_{k=1}^{\infty} (1-q)^{k-1} q \left(kT - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j \right) \quad (\text{by (4)}) \\ &= \frac{T}{q} - \frac{1}{T} \sum_{1 \leq i < j \leq n} t_i t_j. \end{aligned}$$

This provides an upper bound on V , and we have equality. \square

Theorem 5 generalizes a result of Ruckle (1991) that solves the special case of equal detection probabilities and unit search times. In this case, as Ruckle (1991) found, it is optimal for the Searcher to choose with equal probability each of the Searcher strategies $(i, i+1, \dots, n, 1, 2, \dots, i-1)$, for $i = 1, \dots, n$. Our next result shows that, for this special case, the Searcher has an optimal strategy that mixes between only two pure strategies.

Proposition 6 In the game G with $q_i = q$ for all $i \in [n]$ and $t_i = t$ for all $j \in [n]$, it is optimal for the Searcher to choose with equal probability each of the search sequences $(1, 2, \dots, n)$ and $(n, n-1, \dots, 1)$.

Proof. By Theorem 5, it is sufficient to show that in the case $q = 1$, it is optimal to choose with equal probability each of the sequences $1, 2, \dots, n$ and $n, n-1, \dots, 1$. Indeed for any box i , this strategy θ guarantees an expected search time of

$$u(i, \theta) = \frac{1}{2}(it) + \frac{1}{2}(n-i+1)t = \left(\frac{n+1}{2} \right) t.$$

But by Theorem 4, the value of the game in this case is

$$V = nt - \frac{1}{nt} \binom{n}{2} t^2 = \left(\frac{n+1}{2} \right) t.$$

Hence, θ is optimal. \square

4 Optimal Pure Strategy Solutions

Because the optimal Hider strategy must be a mixed strategy, the study of optimal pure Searcher strategies—deterministic search sequences that produce the same optimal expected total search time regardless of where the Hider hides—has been largely overlooked in the literature. In the special case with $q_i = 1/2$ and $t_i = 1$ for $i \in [n]$, Ruckle (1991) shows that the search sequence $n, n-1, \dots, 1, (1, 2, \dots, n)$ is optimal. Ruckle (1991) also shows that if $n = 2$ and $q_1 = q_2 \geq 0.8$, then there does not exist an optimal search sequence. In this section, we prove much more general conditions for the existence of optimal pure strategy in Subsection 4.1, and also conditions for the non-existence of optimal pure strategy in Subsection 4.2. These new findings allow us to extend the results in Ruckle (1991) substantially.

To begin, we say that two search strategies θ_1 and θ_2 are *equivalent* if their expected search time against each box $i \in [n]$ is the same. That is,

$$u(i, \theta_1) = u(i, \theta_2)$$

for each $i \in [n]$.

4.1 Existence of Optimal Pure Strategies

For a search strategy ξ , let $T_i^\ell(\xi)$ denote the sum of the search times of the boxes opened under ξ , up to and including the ℓ^{th} search of box i , for $\ell = 1, 2, \dots$. For example, if $\xi = 1, 2, 2, 1, \dots$, then $T_1^1(\xi) = t_1$ and $T_1^2(\xi) = 2t_1 + 2t_2$. Recalling $r_i = 1 - q_i$ is the overlook probability, the payoff function $u(i, \xi)$ can be written as

$$u(i, \xi) = \sum_{\ell=1}^{\infty} r_i^{\ell-1} q_i T_i^\ell(\xi). \quad (5)$$

To establish the existence of optimal pure strategies, we will focus our attention on the cases in which the Hider's equalizing strategy \mathbf{p}^* is optimal. Because we assume there are some positive coprime integers k_1, \dots, k_n and some $r > 0$ such that $r_i^{k_i} = r$ for all $i \in [n]$, a Gittins search sequence against \mathbf{p}^* can be broken up into cycles, where each cycle consists of k_i searches in box i , $i \in [n]$, and has length $T = \sum_{i=1}^n k_i t_i$. Taking advantage of this observation, the next lemma

develops the right-hand side of (5) further if ξ is a Gittins search sequence against the Hider's equalizing strategy \mathbf{p}^* .

Lemma 7 Let $\xi_1 = (s_1), \dots, \xi_M = (s_M)$ denote M distinct consistent Gittins search sequences against \mathbf{p}^* , where each subsequence $s_j, j = 1, \dots, M$, consists of k_i searches in box i , for $i \in [n]$. Let x_1, x_2, \dots be a sequence taking values in $[M]$ and let ξ be the Gittins search sequence s_{x_1}, s_{x_2}, \dots . The following results hold, where $I(\cdot)$ is the indicator function.

(i) The payoff function $u(i, \xi)$ can be written as

$$u(i, \xi) = \frac{Tr}{1-r} + \sum_{j=1}^M w_i(s_j) \sum_{k=1}^{\infty} I(x_k = j) r^{k-1},$$

where $w_i(s_j) = \sum_{\ell=1}^{k_i} r_i^{\ell-1} q_i T_i^\ell(\xi_j)$ and $T = \sum_{i=1}^n k_i t_i$.

(ii) If $\xi = (s)$ for some sequence s then

$$u(i, \xi) = \frac{Tr}{1-r} + \frac{w_i(s)}{1-r}.$$

Proof. For part (i), we observe that if $\ell = (k-1)k_i + \ell'$ for some $k = 1, 2, \dots$ and some $1 \leq \ell' \leq k_i$, then $T_i^\ell(\xi_j) = (k-1)T + T_i^{\ell'}(\xi_j)$. Therefore, (5) reduces to

$$\begin{aligned} u(i, \xi) &= \sum_{k=1}^{\infty} \sum_{\ell=1}^{k_i} \left((k-1)T + \sum_{j=1}^M I(x_k = j) T_i^\ell(\xi_j) \right) r_i^{(k-1)k_i + (\ell-1)} q_i \\ &= T q_i \sum_{\ell=1}^{k_i} r_i^{\ell-1} \sum_{k=1}^{\infty} (k-1) r^{k-1} + \sum_{k=1}^{\infty} \sum_{\ell=1}^{k_i} \sum_{j=1}^M I(x_k = j) r_i^{\ell-1} q_i T_i^\ell(\xi_j) r^{k-1} \quad (\text{since } r = r_i^{k_i}) \\ &= T q_i \frac{1 - r_i^{k_i}}{1 - r_i} \frac{r}{(1-r)^2} + \sum_{k=1}^{\infty} \sum_{j=1}^M I(x_k = j) w_i(s_j) r^{k-1} \\ &= \frac{Tr}{1-r} + \sum_{j=1}^M w_i(s_j) \sum_{k=1}^{\infty} I(x_k = j) r^{k-1} \end{aligned}$$

For part (ii), we simply set $\xi_1 = (s)$ and $x_k = 1$ for all k . Applying part (i), we then have that

$$u(i, \xi) = \frac{Tr}{1-r} + w_i(s) \sum_{k=1}^{\infty} r^{k-1} = \frac{Tr}{1-r} + \frac{w_i(s)}{1-r},$$

which completes the proof. \square

A key step to establish the optimality of a pure search strategy—or a search sequence—is to establish the equivalence between a pure search strategy and a mixed search strategy. Our next

lemma shows that, for any Gittins search sequence against \mathbf{p}^* , it is always possible to construct an equivalent mixed strategy that consists of only consistent search sequences.

Lemma 8 Let $\xi_1 = (s_1), \dots, \xi_M = (s_M)$ denote $M \leq n!$ distinct consistent Gittins search sequences against \mathbf{p}^* and let x_1, x_2, \dots denote a sequence taking values in $[M]$. The search sequence $\xi^* \equiv s_{x_1}, s_{x_2}, \dots$ is equivalent to the mixed Searcher strategy θ that chooses ξ_j with probability

$$\theta(\xi_j) = \sum_{k=1}^{\infty} I(x_k = j)(1-r)r^{k-1},$$

for $j \in [M]$.

Proof. First note that θ is a well-defined Searcher strategy, since

$$\sum_{j=1}^M \theta(\xi_j) = \sum_{k=1}^{\infty} (1-r)r^{k-1} = 1.$$

By Lemma 7, part (ii), the expected payoff of box i against ξ_j is

$$u(i, \xi_j) = \frac{Tr}{1-r} + \frac{w_i(s_j)}{1-r}.$$

Hence, the expected search time of box i under θ is

$$u(i, \theta) = \sum_{j=1}^M \theta(\xi_j) u(i, \xi_j) = \frac{Tr}{1-r} + \sum_{j=1}^M \frac{\theta(\xi_j)}{1-r} w_i(s_j)$$

Applying Lemma 7, part (i), the expected search time of box i against ξ^* is

$$\begin{aligned} u(i, \xi^*) &= \frac{Tr}{1-r} + \sum_{j=1}^M w_i(s_j) \sum_{k=1}^{\infty} I(x_k = j) r^{k-1} \\ &= \frac{Tr}{1-r} + \sum_{j=1}^M w_i(s_j) \frac{\theta(\xi_j)}{1-r}, \end{aligned}$$

by definition of θ . Hence, $u(i, \theta) = u(i, \xi^*)$. \square

Before presenting our main result, we present another lemma that will help us to show it is possible—under certain conditions—to construct a search sequence that is equivalent to a mixed search strategy that mixes between consistent Gittins search sequences against \mathbf{p}^* . Though this lemma is crucial to our main result, it stands alone as a lemma in probability theory. We were not able to find it in the literature, and it may be of independent interest.

Lemma 9 Let M be a positive integer. If $1 - 1/M \leq r < 1$, then for any $\lambda_1, \dots, \lambda_M \geq 0$ with $\lambda_1 + \dots + \lambda_M = 1$, there exists a sequence x_1, x_2, \dots taking values in $[M]$ such that

$$\sum_{k=1}^{\infty} I(\{x_k = j\})(1 - r)r^{k-1} = \lambda_j, \quad (6)$$

for each $j \in [M]$, where I is the indicator function. Moreover, if $r^d \geq 1 - \min_j \lambda_j$, then there are at least M^d sequences that satisfy condition (6).

Proof. We define the sequence x_1, x_2, \dots recursively and will prove by induction on K that for each $K = 0, 1, \dots$, we have

$$\sum_{k=1}^K I(\{x_k = j\})(1 - r)r^{k-1} \leq \lambda_j, \quad (7)$$

for each $j \in [M]$. This is trivially true for $K = 0$. Suppose (7) is true for some $K \geq 0$, and note that

$$\sum_{j=1}^M \left(\lambda_j - \sum_{k=1}^K I(\{x_k = j\})(1 - r)r^{k-1} \right) = 1 - \sum_{k=1}^K (1 - r)r^{k-1} = r^K. \quad (8)$$

By the induction hypothesis, every term in the outer sum on the left-hand side of (8) is non-negative, and it follows that for some $j \in [M]$, we have

$$\lambda_j - \sum_{k=1}^K I(\{x_k = j\})(1 - r)r^{k-1} \geq r^K/M \geq (1 - r)r^K,$$

since $r > 1 - 1/M$. Then we can take $x_{K+1} = j$, and (7) is clearly satisfied for $K + 1$. Equation (6) follows from taking limits in (8) as $K \rightarrow \infty$.

If $1 - r \leq \min_j \lambda_j$, there are M choices for x_1 , and consequently there are at least M distinct sequences. In general, if $1 - r^d = \sum_{k=1}^d (1 - r)r^{k-1} \leq \min_j \lambda_j$ for some $d \geq 1$, there are M choices for each x_i , $i = 1, \dots, d$, so at least M^d distinct sequences meet condition (6). \square

It is interesting to point out that Lemma 9 relies on the axiom of choice in the induction step.

Note that sequence x_1, x_2, \dots can be constructed using a greedy algorithm. Indeed, for any $K = 0, 1, \dots$, let $\delta_j(K) = \sum_{k=1}^K I(\{x_k = j\})(1 - r)r^{k-1}$ for each $j \in [M]$. Notice that $\delta_j(0) = 0$ for all j . Then, x_1, x_2, \dots can be computed recursively by setting $x_{K+1} = i$ for any $K = 0, 1, \dots$, where $i \in [M]$ is chosen such that $\delta_i(K) + (1 - r)r^K \leq \lambda_i$.

Also, note that we cannot strengthen the bound $r \geq 1 - 1/M$ in Lemma 9, since for any $r < 1 - 1/M$, if we take $\lambda_1 = \dots = \lambda_M = 1/M$, then for any choice $x_1 = i$, we have

$$\sum_{k=1}^{\infty} I(\{x_k = j\})(1-r)r^{k-1} \geq (1-r) > 1/M = \lambda_j.$$

We now use Lemmas 8 and 9 to prove our main results in the next theorem, which gives sufficient conditions under which an optimal pure strategy for G can be found.

Theorem 10 Let ξ_1, \dots, ξ_M be M consistent Gittins search sequences against \mathbf{p}^* and suppose θ is some Searcher mixed strategy with support $\{\xi_1, \dots, \xi_M\}$. If $r = r_i^{k_i} \geq 1 - 1/M$ for all $i \in [n]$, then

- (i) there exists some search sequence ξ^* such that

$$u(i, \xi^*) = u(i, \theta) \quad (9)$$

for all $i \in [n]$,

- (ii) moreover if $r^d \geq 1 - \min_j \theta(\xi_j)$ for some $d = 1, \dots$ then there are at least M^d distinct search sequences satisfying (9), and
- (iii) in particular, if \mathbf{p}^* is optimal for the Hider and $r \geq 1 - 1/n$, then the Searcher has an optimal pure strategy.

Proof. To prove part (i) of the theorem, let $\lambda_j = \theta(\xi_j)$ for $j \in [M]$, so that $\sum_{j=1}^M \lambda_j = 1$. By definition, we have that

$$u(i, \theta) = \sum_{j=1}^M \lambda_j u(i, \xi_j).$$

Let x_1, x_2, \dots be a sequence taking values in $[M]$ such that

$$\sum_{k=1}^{\infty} I(x_k = j)(1-r)r^{k-1} = \lambda_j \quad (10)$$

for all $j \in [M]$. By Lemma 9, such a sequence exists because $r \geq 1 - 1/M$.

Since each ξ_j is a consistent Gittins search sequence against \mathbf{p}^* , it can be written as $\xi_j = (s_j)$, for some finite subsequence s_j . Let ξ^* be the pure Searcher strategy s_{x_1}, s_{x_2}, \dots . Then part (i) of the theorem follows from Lemma 8.

For part (ii), if $1 - r^d \leq \min_j \theta(\xi_j)$ for some $d \geq 1$, we know from the second statement of Lemma 9 that at least M^d distinct sequences x_1, x_2, \dots exist. So, by the construction of ξ^* , it is straightforward that there are at least M^d distinct pure strategies ξ^* which satisfy (9).

Part (iii) of the theorem follows from the fact that if \mathbf{p}^* is optimal, then by Theorem 3, there is an optimal Searcher strategy θ that mixes between n consistent Gittins search sequences against \mathbf{p}^* . So by part (i) of this theorem, there must be a Searcher pure strategy whose expected search time against each $i \in [n]$ is the same as that of θ , and which is therefore optimal. \square

We can use Theorem 10 to show that the Searcher has optimal pure strategies in cases of the game G for which we know \mathbf{p}^* is optimal. As indicated in Theorem 10, for there to exist an optimal pure search strategy, r —the probability of not finding the Hider in the first cycle of k_i searches in box i , $i = 1, \dots, n$ —needs to be sufficiently large. Intuitively, with a large r , the subsequence chosen in the first cycle plays a less important role, which makes it possible to achieve (9) by prioritizing the other boxes in the following cycles.

Corollary 11 If the detection probabilities q_1, \dots, q_n are all equal to some $q \leq 1/n$, then the Searcher has an optimal pure strategy. In addition, if $1 - (1 - q)^d \leq \min_j t_j/T$ for some $d \geq 1$, then there are at least n^d distinct optimal pure strategies.

Proof. By Theorem 5, the Hider's equalizing strategy \mathbf{p}^* is optimal and the Searcher strategy θ which chooses each sequence $\xi_i = (i+1, i+2, \dots, n, 1, 2, \dots, i-1)$, $i = 1, \dots, n$, with probability t_i/T is optimal. Since $r = r_1 = \dots = r_n = 1 - q \geq 1 - 1/n$, the existence of an optimal pure search strategy follows Theorem 10, part (iii). Moreover, by Theorem 10 part (ii), if $1 - (1 - q)^d = 1 - r^d \leq \min_i \theta(\xi_i) = \min_i t_i/T$, there exist at least n^d optimal pure strategies. \square

Corollary 12 If the detection probabilities are all equal to some $q \leq 1/2$ and the search times are all equal, then the Searcher has an optimal pure strategy. Furthermore, if $1 - (1 - q)^d \leq 1/2$ for some $d \geq 1$, then there are at least 2^d distinct optimal pure strategies.

Proof. Again, by Theorem 5, the Hider's equalizing strategy \mathbf{p}^* is optimal. By Proposition 6, the Searcher has an optimal strategy that chooses two pure strategies each with probability $1/2$. Since $r = r_1 = \dots = r_n = 1 - q \geq 1/2$, the existence of an optimal pure search strategy follows

Theorem 10, part (iii). Moreover, by Theorem 10 part (ii), if $1 - (1 - q)^d = 1 - r^d \leq \min_i \theta(\xi_i) = 1/2$, there exist at least 2^d optimal pure strategies. \square

Corollary 13 If there are $n = 2$ boxes with $t_1 = t_2$ and $r_1^k = r_2^{k+1} \geq 1/2$ for some $k = 1, \dots, 12$, then the Searcher has an optimal pure strategy.

Proof. By Theorem 2, the Hider's equalizing strategy \mathbf{p}^* is optimal, so the corollary follows immediately from Theorem 10, part (iii). \square

We illustrate Corollary 11 with the following example.

Example 14

Consider $n = 3$ with $q_1 = q_2 = q_3 = q = 0.2$ and $t_1 = 2.00, t_2 = 2.88, t_3 = 5.12$. Let $s_1 = 1, 2, 3, s_2 = 2, 3, 1$ and $s_3 = 3, 1, 2$. We have $r = 1 - q = 0.8$ and $T = t_1 + t_2 + t_3 = 10$. By Corollary 11, the Searcher has an optimal strategy. We can construct the optimal strategy as follows.

First, following Lemma 9, we construct a sequence $\sigma = x_1, x_2, \dots$ taking values in $\{1, 2, 3\}$ such that $\sum_{k=1}^{\infty} I(x_k = i)(1 - r)r^{k-1} = t_i/T$ for $i = 1, 2, 3$. Let $\delta_i(K) = \sum_{k=1}^K I(x_k = i)(1 - r)r^{k-1}$ for $K = 0, 1, \dots$. Note that $\delta_i(0) = 0$ for all i . For any $K = 1, 2, \dots$, we define x_K recursively such that if $\delta_i(K-1) + (1 - r)r^{K-1} \leq t_i/10$, then $x_K = i$.

- For $K = 1$, we have $\delta_i(0) + (1 - r)r^{K-1} = 1 - r = 0.2 \leq t_i/10$ for all i , so x_1 can be 1, 2 or 3.
- For $K = 2$, if we let $x_1 = 1$, it is easy to check that x_2 just can be 2 or 3.
- Continuing this way, we obtain the sequence $\sigma = 1, 2, 2, 3, (3)$.

Second, let $\xi^* = s_1, s_2, s_2, s_3, (s_3) = 1, 2, 3, 2, 3, 1, 2, 3, 1, 3, 1, 2, (3, 1, 2)$. Then, ξ^* is optimal by Theorem 10. Note that if we change the choice of x_1 or x_2 , we will get a different sequence σ which leads to another different optimal pure strategy.

4.2 Non-existence of Optimal Pure Strategies

We now consider what conditions ensure that there will be no optimal pure strategy for the Searcher. As in the previous section we focus on the case that the Hider's equalizing strategy \mathbf{p}^* is optimal.

For a particular instance of the game G , suppose θ is an optimal Searcher strategy that is a mixture of distinct consistent Gittins search sequences ξ, \dots, ξ_M , for some $M \leq n!$. Write $\lambda_{\max}(\theta) = \max_{1 \leq j \leq M} \theta(\xi_j)$, and let λ^* denote the supremum of $\lambda_{\max}(\theta)$ over all such optimal Searcher strategies.

Lemma 15 If \mathbf{p}^* is optimal and $r_1^{k_1} = \dots = r_n^{k_n} = r < 1 - \lambda^*$, then there does not exist an optimal pure strategy for the Searcher.

Proof. Suppose there is an optimal pure Searcher strategy ξ^* . Let $\xi_1 = (s_1), \dots, \xi_{n!} = (s_{n!})$ be the consistent Gittins search sequences against \mathbf{p}^* . Since \mathbf{p}^* is optimal, ξ^* must be a Gittins search sequence against \mathbf{p}^* . Hence, ξ can be written as s_{x_1}, s_{x_2}, \dots , where $x_k \in [n!]$, for $k = 1, 2, \dots$. Let θ be the search strategy that chooses ξ_j with probability

$$\theta(\xi_j) = \lambda_j \equiv \sum_{k=1}^{\infty} I(x_k = j)(1 - r)r^{k-1},$$

for $j \in [n!]$.

By Lemma 8, strategies θ and ξ^* are equivalent, so θ is optimal. But

$$\lambda^* \geq \lambda_{\max}(\theta) = \max_j \lambda_j \geq \lambda_{x_1} = 1 - r + \sum_{k=2}^{\infty} I(x_k = x_1)(1 - r)r^{k-1} \geq 1 - r > \lambda^*,$$

a contradiction. \square

Lemma 15 does not give much clue as to how one might calculate λ^* , and it is not even clear that λ^* is strictly less than 1. We now consider the two classes of solutions for which we know that \mathbf{p}^* is optimal and show that there is a non-empty interval of values of r for which there is no optimal pure strategy.

Theorem 16 For equal detection probabilities $q_1 = \dots = q_n = q$, there is no optimal pure strategy for $q > q^*$, where

$$q^* \equiv 1 - \frac{1}{T(T - t_{\min})} \sum_{\{a,b\} \subset [n]} t_a t_b, \quad (11)$$

and $t_{\min} = \min_i t_i$.

Proof. Since the detection probabilities are equal, according to Theorem 5 the equalizing strategy \mathbf{p}^* is optimal. The consistent Gittins search sequences against \mathbf{p}^* are $\xi_1 = (\sigma_1), \dots, \xi_{n!} = (\sigma_{n!})$, where $\sigma_1, \dots, \sigma_{n!}$ are the permutations of $[n]$. Let θ be an optimal Searcher strategy that mixes between $\xi_1, \dots, \xi_{n!}$ and let $\lambda_j = \theta(\xi_j)$, $j \in [n!]$.

Without loss of generality suppose the last term of σ_1 is i . Using Lemma 7, part (ii) and the fact that $1 - r = q$, the expected search time of box i under θ can be written as

$$u(i, \theta) = \frac{T(1 - q)}{q} + \sum_{j=1}^{n!} \frac{\lambda_j w_i(\sigma_j)}{q}.$$

Since θ is optimal, $u(i, \theta)$ is equal to the value V of the game, which is given in Theorem 5. Hence,

$$\begin{aligned} 0 = V - u(i, \theta) &= \left(\frac{T}{q} - \frac{1}{T} \sum_{\{a,b\} \subset [n]} t_a t_b \right) - \left(\frac{T(1-q)}{q} + \sum_{j=1}^{n!} \frac{\lambda_j w_i(\sigma_j)}{q} \right) \\ &= T - \frac{1}{T} \sum_{\{a,b\} \subset [n]} t_a t_b - \sum_{j=1}^{n!} \frac{\lambda_j w_i(\sigma_j)}{q} \end{aligned} \quad (12)$$

Since the last term of σ_1 is i , we must have $w_i(\sigma_1) = qT$. Also, $w_i(\sigma_j) \geq qt_i$ for all $j \neq 1$, so

$$\sum_{j=1}^{n!} \frac{\lambda_j w_i(\sigma_j)}{q} \geq \lambda_1 T + \sum_{j \neq 1} \lambda_j t_i = (T - t_i) \lambda_1 + t_i.$$

Combining the preceding with (12), we obtain

$$\lambda_1 \leq 1 - \frac{1}{T(T - t_i)} \sum_{\{a,b\} \subset [n]} t_a t_b \leq q^*,$$

by definition of q^* . By a similar analysis, the inequality above holds when λ_1 is replaced by any λ_j , $2 \leq j \leq n!$, so $\lambda_{\max}(\theta) \leq q^*$ and $\lambda^* \leq q^*$, by definition of λ^* .

Hence, if $q > q^*$, then $q > \lambda^*$, so $r = 1 - q < 1 - \lambda^*$, and by Lemma 15, there is no optimal pure strategy. \square

Note that q^* is always at least $1/2$, because

$$\sum_{\{a,b\} \subset [n]} t_a t_b = \frac{1}{2} \left(T^2 - \sum_{i \in [n]} t_i^2 \right) \leq \frac{1}{2} \left(T^2 - \sum_{i \in [n]} t_i t_{\min} \right) = \frac{1}{2} T(T - t_{\min}).$$

It is straightforward to verify that in the case of equal search times as well as equal detection probabilities, q^* is equal to $1/2$, so we obtain the following corollary of Theorem 16.

Corollary 17 For equal detection probabilities q and equal search times, there is no optimal pure strategy for $q > 1/2$.

The bound q^* is tight in this case, since there is an optimal pure strategy for all $q \leq 1/2$, by Corollary 12. Intuitively, if $q > 1/2$ then the first few searches affect the expected total search time so significantly that it is necessary for the Searcher to use a mixed strategy to achieve optimality. If $q \leq 1/2$, then it is possible to achieve optimality with a single search sequence—a pure strategy—by prioritizing the other boxes later on to neutralize the advantage of the first few boxes in the search sequence.

For general search times, the bound of q^* may still be tight. For example, for the three-box game, $q_1 = q_2 = q_3 = q$, $t_{\min} = t_1$, and $t_2^2 = t_1 t_3$, we have $q^* = \frac{t_2^2 + t_2 t_3 + t_3^2}{(t_1 + t_2 + t_3)(t_2 + t_3)}$. It is straightforward to show that the search sequence 3, 2, 1, (1, 2, 3) is optimal when $q = q^*$.

We now turn to the case that $t_1 = t_2 = 1$ and $r = r_1^k = r_2^{k+1}$ for $k = 1, \dots, 12$. The Hider's equalizing strategy \mathbf{p}^* is optimal, by Theorem 2. There are two consistent Gittins search sequences $\xi_1 = (s_1)$, $\xi_2 = (s_2)$ against \mathbf{p}^* such that $s_1 = 1, 2, \dots$, $s_2 = 2, 1, \dots$ and the last $2k - 1$ terms of s_1 and s_2 are the same. For example, $k = 3$, then $s_1 = 1, 2, 2, 1, 2, 1, 2$ and $s_2 = 2, 1, 2, 1, 2, 1, 2$. It is straightforward to verify that $w_1(s_1) = w_1(s_2) - q_1$ and $w_2(s_1) = w_2(s_2) + q_2$.

Let θ be an optimal Searcher strategy that chooses ξ_j with probability λ_j for $j = 1, 2$. By Lemma 7, the expected search time of box 1 under θ can be written as

$$\begin{aligned} u(1, \theta) &= \frac{Tr}{1-r} + \frac{1}{1-r}(\lambda_1 w_1(s_1) + \lambda_2 w_1(s_2)) \\ &= \frac{Tr}{1-r} + \frac{1}{1-r}(\lambda_1(w_1(s_2) - q_1) + \lambda_2 w_1(s_2)) \\ &= \frac{Tr}{1-r} + \frac{1}{1-r}(w_1(s_2) - q_1 \lambda_1) \end{aligned}$$

Similarly, the expected search time of box 2 under θ can be written as

$$\begin{aligned} u(2, \theta) &= \frac{Tr}{1-r} + \frac{1}{1-r}(\lambda_1 w_2(s_1) + \lambda_2 w_2(s_2)) \\ &= \frac{Tr}{1-r} + \frac{1}{1-r}(\lambda_1(w_2(s_2) + q_2) + \lambda_2 w_2(s_2)) \\ &= \frac{Tr}{1-r} + \frac{1}{1-r}(w_2(s_2) + q_2 \lambda_1) \end{aligned}$$

Since θ is optimal, we have $u(1, \theta) = u(2, \theta)$ or $w_1(s_2) - q_1 \lambda_1 = w_2(s_2) + q_2 \lambda_1$. Therefore,

$$\lambda_1 = \frac{w_1(s_2) - w_2(s_2)}{q_1 + q_2} = \frac{w_1(s_2) - w_2(s_2)}{2 - r_1 - r_2}. \quad (13)$$

By Lemma 15, there is no optimal pure strategy if $r < 1 - \max(\lambda_1, \lambda_2)$ or equivalently,

$$r < 1 - \lambda_1 \quad \text{and} \quad r < 1 - \lambda_2 = \lambda_1 \quad (14)$$

Solving (13)-(14) computationally for $k = 1, \dots, 12$, $r = r_1^k = r_2^{k+1}$, we get r^* such that there is no optimal pure strategy for $r < r^*$ as follows.

k	r^*	k	r^*
1	0.216757	7	0.399087
2	0.233196	8	0.426468
3	0.267022	9	0.451679
4	0.302809	10	0.474914
5	0.337219	11	0.496363
6	0.369366	12	0.023657

Note that for any arbitrary $t_1 = t_2 = t$, the bounds r^* do not change.

5 The Case of Two Boxes

In this section we restrict our attention to the case of $n = 2$ boxes. In Subsection 5.1, we examine conditions under which there are (or are not) optimal pure strategies. In Subsection 5.2, we pose the question: if there is no optimal pure strategy, what is the best possible pure strategy we can find?

5.1 Optimal Pure Strategies for Two Boxes

We first investigate the existence of optimal pure Searcher strategies for the two-box game, $n = 2$, making the standing assumption that $t_1 \leq t_2$. We start with a general result that follows from Theorem 10.

Theorem 18 If $n = 2$ and $r_1^{k_1} = r_2^{k_2} \geq 1/2$ then the Searcher has an optimal pure strategy.

Proof. Let \mathbf{p} denote the optimal Hider strategy and let \mathbf{p}^* denote the Hider's equalizing strategy. If $\mathbf{p} = \mathbf{p}^*$, then the theorem follows from Theorem 10, so assume that $\mathbf{p} \neq \mathbf{p}^*$.

By Theorem 3, there is an optimal Searcher strategy θ whose support is two consistent Gittins search sequences ξ_1 and ξ_2 against \mathbf{p} . If $\xi_1 = \xi_2$, then there is nothing to prove, so assume that $\xi_1 \neq \xi_2$. Then ξ_1 and ξ_2 must be identical up until the j th term, for some j , at which point, the Searcher is indifferent between searching the two boxes: that is, the new conditional hiding distribution is \mathbf{p}^* . So each $\xi_i, i = 1, 2$ can be written $\xi_i = s, \xi'_i$ for some finite sequence s and some consistent Gittins search sequences ξ'_1 and ξ'_2 against \mathbf{p}^* . Let θ' be the search strategy that

chooses ξ'_i with probability $\theta(\xi_i)$ for $i = 1, 2$. By Theorem 10, there is a pure strategy ξ^* such that $u(i, \xi^*) = u(i, \theta')$ for $i = 1, 2$. Then, if we set $\xi^{**} = s, \xi^*$, it is easy to see that $u(i, \xi^{**}) = u(i, \theta)$ for $i = 1, 2$, so ξ^{**} is optimal. \square

We now consider separately the special case when $q_1 = q_2 = q$. By Corollary 11 (or Theorem 18), the Searcher has an optimal pure strategy for $q \leq 1/2$ and by Theorem 16, the Searcher has no optimal pure strategy for $q > q^*$, where q^* reduces to $t_2/(t_1 + t_2)$. This bound is tight, since a straightforward calculation shows that the strategy $2, 1, (1, 2)$ is optimal for the Searcher when $q = t_2/(t_1 + t_2)$. Also, note that $2, 1, (1, 2)$ is a unique optimal pure strategy when $q = t_2/(t_1 + t_2)$ and $t_2 > t_1$. Indeed, any optimal pure strategy must be a Gittins sequence $\xi = s_{x_1}, s_{x_2} \dots$ where each $x_k \in \{1, 2\}$, $s_1 = 1, 2$, and $s_2 = 2, 1$. Suppose $s_{x_1} = 1, 2$, then it is obvious $u(2, \xi) \geq u(2, \xi')$ where $\xi' = 1, 2, (2, 1)$. However,

$$u(2, \xi) \geq u(2, \xi') = (t_1 + t_2)/q - t_1^2/(t_1 + t_2) > (t_1 + t_2)/q - (t_1 t_2)/(t_1 + t_2) = V.$$

So, s_{x_1} must be $2, 1$. Then, it is straightforward that $s_{x_k} = 1, 2$ for all $k = 2, \dots$, otherwise $u(1, \xi) > V$. Hence, the strategy $2, 1, (1, 2)$ is unique.

One might conjecture that there is an optimal pure strategy for all the values of q between $1/2$ and $t_2/(t_1 + t_2)$. However, it turns out that this conjecture is not true in general. We show that, at least for some choices of t_1 and t_2 , there is a non-empty interval of values of q between $1/2$ and $t_2/(t_1 + t_2)$ for which there is no optimal pure strategy. For example, if $t_2/(t_1 + t_2) = 0.8$, then there is no optimal pure strategy for $0.553 < q < 0.723$. In general, we have the following result.

Lemma 19 Consider the case with $n = 2$ and equal detection probabilities $q_1 = q_2 = q$ and search times t_1 and t_2 with $t_1 \leq t_2$. If

$$\frac{t_1}{t_1 + t_2} < \frac{1}{4} \quad \text{and} \quad 1 - \sqrt{\frac{t_1}{t_1 + t_2}} < q < \frac{1}{2} \left(\sqrt{1 - \frac{4t_1}{t_1 + t_2}} + 1 \right),$$

then there is no optimal pure strategy.

Proof. We first derive necessary conditions for a pure strategy to be optimal. By Theorem 5, the Hider's equalizing strategy \mathbf{p}^* is optimal and the value of the game is

$$V = \frac{t_1 + t_2}{q} - \frac{t_1 t_2}{t_1 + t_2}. \quad (15)$$

Let ξ be an optimal pure strategy. It must be a Gittins search sequence against \mathbf{p}^* , so ξ must be of the form $\xi = s_{x_1}, s_{x_2}, \dots$, where $s_1 = (1, 2)$ and $s_2 = (2, 1)$ and $x_k \in \{1, 2\}$. Let $\gamma_k = I(x_k = 1)$ (in other words, γ_k is equal to 1 if $s_{x_k} = (1, 2)$, otherwise it is equal to 0). Note that $w_1(s_1) = qt_1$ and $w_1(s_2) = q(t_1 + t_2)$. Similarly for $w_2(s_1)$ and $w_2(s_2)$.

By Lemma 7, part (i),

$$\begin{aligned} u(1, \xi) &= \frac{T(1-q)}{q} + \sum_{k=1}^{\infty} qt_1 \gamma_k r^{k-1} + \sum_{k=1}^{\infty} q(t_1 + t_2)(1 - \gamma_k) r^{k-1} \\ &= \frac{T(1-q)}{q} + q(t_1 + t_2) \sum_{k=1}^{\infty} r^{k-1} - qt_2 \sum_{k=1}^{\infty} \gamma_k r^{k-1} \\ &= \frac{T}{q} - qt_2 \sum_{k=1}^{\infty} \gamma_k r^{k-1}. \end{aligned} \tag{16}$$

Similarly,

$$u(2, \xi) = \frac{T}{q} - qt_1 \sum_{k=1}^{\infty} (1 - \gamma_k) r^{k-1} \tag{17}$$

Since ξ is optimal, we must have $V = u(1, \xi) = u(2, \xi)$, and setting (15), (16) and (17) to be equal gives

$$\begin{cases} q \sum_{k=1}^{\infty} \gamma_k r^{k-1} = t_1/(t_1 + t_2), \\ q \sum_{k=1}^{\infty} (1 - \gamma_k) r^{k-1} = t_2/(t_1 + t_2). \end{cases} \tag{18}$$

Let $y_1 = 1 - \sqrt{\frac{t_1}{t_1+t_2}}$ and $y_2 = \frac{1}{2}(\sqrt{1 - \frac{4t_1}{t_1+t_2}} + 1)$. Observe that $1/2 < y_1 < y_2 < t_2/(t_1 + t_2)$ when $t_1/(t_1 + t_2) < 1/4$. It is easy to see that γ_1 must be 0 because if $\gamma_1 = 1$, then $q \sum_{k=1}^{\infty} \gamma_k r^{k-1} \geq q > 1/2 \geq t_1/(t_1 + t_2)$, contradicting (18).

We consider two cases for γ_2 . First, if $\gamma_2 = 0$ then $q \sum_{k=1}^{\infty} (1 - \gamma_k) r^{k-1} \geq q(1 + r) = q(2 - q)$. It is straightforward to check that $q(2 - q) > t_2/(t_1 + t_2)$ for $1 - \sqrt{\frac{t_1}{t_1+t_2}} < q < \frac{1}{2}(\sqrt{1 - \frac{4t_1}{t_1+t_2}} + 1)$. Second, if $\gamma_2 = 1$, then $q \sum_{k=1}^{\infty} \gamma_k r^{k-1} \geq qr = q(1 - q)$, and it is straightforward to check that $q(1 - q) > t_1/(t_1 + t_2)$ for $1 - \sqrt{\frac{t_1}{t_1+t_2}} < q < \frac{1}{2}(\sqrt{1 - \frac{4t_1}{t_1+t_2}} + 1)$. In either case, we have a contradiction, so we conclude that there does not exist an optimal pure strategy. \square

5.2 Best Pure Strategies

We have seen that for equal search times and equal detection probabilities, there is no optimal pure strategy solution for $q > 1/2$ when $n = 2$. A natural question to ask is what the *best* pure strategy

is for this range of values of q . More precisely, we wish to find a search sequence ξ that minimizes $\max\{u(1, \xi), u(2, \xi)\}$. Ruckle (1991) shows that for equal search times, there always exists ξ such that $u(1, \xi) = u(2, \xi) = \dots = u(n, \xi)$. However, a search sequence that results in the same expected total search time regardless of where the Hider hides is not necessarily the best search sequence. Below is a counterexample.

Example 20

Consider $n = 2$ with $q_1 = q_2 = 0.5$ and $t_1 = t_2 = 1$. According to Corollary 13, the Searcher has an optimal pure strategy. It is straightforward to verify that the value of the game is $7/2$ and the search sequence

$$1, 2, (2, 1)$$

is optimal. However, one can also verify that a different search sequence

$$1, 2, 1, 2, 2, 2, 2, 1, 1, 2, 1, 2, 1, 1, 1, 2, (2, 1)$$

results in an expected search time $923/256$ whether the Hider hides in box 1 or box 2. In other words, it is possible that there exists more than one search sequence that results in the same expected search time regardless of where the Hider hides, so finding one such search sequence does not mean it is the best search sequence that minimizes $\max\{u(1, \xi), u(2, \xi)\}$. \square

While $u(1, \xi) = \dots = u(n, \xi)$ is not a sufficient condition for ξ to be the best pure strategy, it is clearly a necessary condition, for otherwise it is possible to make small tweaks to ξ to reduce $\max_{i=1, \dots, n} u(i, \xi)$. For the case with $n = 2$, $q_1 = q_2 = q = 2/3$ so that $r = 1 - q = 1/3$, a simple calculation shows that the search sequence $\xi = 1, 2, (2, 2, 1, 1)$ guarantees an expected search time of

$$u(1, \xi) = u(2, \xi) = 2 + r + \frac{r^2 + 3r^3}{1 - r^2} = \frac{31}{12}.$$

We conjecture that this is the best pure strategy for $q = 2/3$. Computational results indicate that any pure strategy with expected search time at most $31/12$ must begin with the sequence 1, 2 followed by 15 cycles of 2, 2, 1, 1 (or the same sequence with 1 and 2 swapped). Below we generalize the conjecture to the case with $n = 2$ and $q_1 = q_2 = (m - 1)/m$, for $m = 3, 4, \dots$

Conjecture 21 Suppose $n = 2$ and $q_1 = q_2 = (m - 1)/m$, for some integer $m \geq 3$. Let $r = 1 - q_1 = 1/m$. The best pure strategy is $\xi = 1, 2, (2, 2, \dots, 2, 1, 1, \dots, 1)$, in which the Searcher

begins with 1, 2, and then indefinitely repeats the same pattern that consists of $(m - 1)$ searches in box 2 followed by $(m - 1)$ searches in box 1. The resulting expected search time is

$$u(1, \xi) = u(2, \xi) = 2 + r + \frac{\sum_{j=2}^{m-1} r^j + mr^m}{1 - r^{m-1}} = 2 + \frac{1}{m} + \frac{m^{m-2} + m - 2}{(m - 1)(m^{m-1} - 1)}.$$

We note that for $m = 2$, Conjecture 21 is true because in that case the search sequence 1, 2, (2, 1) is optimal as seen in Example 20.

6 Conclusion

This paper studies a discrete search game and investigates the existence of optimal pure strategies for the Searcher—a single deterministic search sequence that achieves the optimal expected total search time regardless of where the Hider hides. An optimal pure search strategy has significant practical value because it is straightforward to execute without the need of randomization. It would also avoid the potential criticism from lay persons that the chosen course of action does not turn out well even though the chosen action is properly selected from a set of pure strategies that compose the optimal mixed strategy.

For a two-person zero-sum game, typically if one player’s optimal mixed strategy includes several pure strategies, the other player’s optimal mixed strategy also includes the same number of pure strategies. In our search game, the Hider’s optimal mixed strategy always instructs the Hider to hide in each of the n boxes with a strictly positive probability, so it may come as a surprise that the Searcher has an optimal pure strategy in several nontrivial cases. In particular, because the first few boxes in a search sequence have a profound effect on the conditional expected total search time for each box, intuitively it is necessary for the Searcher to randomize the first few boxes in order to achieve optimality. However, if the detection probabilities q_i , $i = 1, \dots, n$, are sufficiently small, then the search tends to take a long time with high probability, so the effect of the first few searches becomes less significant, which makes it possible for the Searcher to prioritize the other boxes later on in a search sequence to still achieve optimality. The several cases in which an optimal pure strategy exists for the Searcher we find in this paper all meet this general observation.

A natural question to ask in general is whether we can determine the best search sequence that minimizes the expected total search time regardless of where the Hider hides. In addition, what is the gap between an optimal mixed search strategy and the best search sequence when the latter

is not optimal? In Section 5.2 we make a conjecture of the best search sequence for a very special case but it appears rather difficult to determine the best search sequence in general.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1935826 and upon work supported by the Air Force Office of Scientific Research under award number FA9550-23-1-0556.

References

Alpern S, Lidbetter T (2013) Mining coal or finding terrorists: The expanding search paradigm. *Oper. Res.* 61(2):265–279.

Binmore K (2008) *Rational Decisions*. Princeton University Press.

Bram J (1963) *A 2-player N-region search game*. Center for Naval Analyses Alexandria Va Operations Evaluation Group.

Clarkson J, Lin KY, Glazebrook KD (2023) A classical search game in discrete locations. *Math. Oper. Res.* 48(2):687–707.

Clarkson J, Lin KY (2023) Computing Optimal Strategies for a Search Game in Discrete Locations. arXiv preprint arXiv:2305.10342.

Condon A, Deshpande A, Hellerstein L, Wu N (2009) Algorithms for distributional and adversarial pipelined filter ordering problems. *ACM Transactions on Algorithms (TALG)* 5(2):1–34.

Gintis H (2014) *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences-Revised Edition*. Princeton University Press.

Gittins JC (1979) Bandit processes and dynamic allocation indices. *J. Roy. Statist. Soc. B* 41(2):148–177.

Gittins J, Glazebrook KD, Weber R (2011) *Multi-Armed Bandit Allocation Indices*. John Wiley and Sons, New York.

Harsanyi JC (1973) Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International journal of game theory* 2(1):1–23.

Lidbetter T (2013) Search games with multiple hidden objects. *SIAM J. Control Optim.* 51(4):3056–3074.

Lin KY, Singham DI (2015) Robust search policies against an intelligent evader. Technical Report NPS-OR-15-009, Naval Postgraduate School, Monterey, CA.

Matula D (1964) A periodic optimal search. *Amer. Math. Monthly* 71(1):15–21.

Roberts DM, Gittins JC (1978) The search for an intelligent evader: Strategies for Searcher and evader in the two-region problem. *Naval Res. Logist. Quart.* 25(1):95–106.

Ruckle WH (1991) A discrete search game. Raghavan TES, Ferguson TS, Parthasarathy T, Vrieze O.J., eds. *Stochastic Games and Related Topics*, Theory and Decision Library, vol. 7 (Springer, Dordrecht, Netherlands), 29–43.