A Classical Search Game in Discrete Locations

Jake Clarkson* Kyle Y. Lin† Kevin D. Glazebrook‡

March 5, 2022

Abstract

Consider a two-person zero-sum search game between a hider and a searcher. The hider hides among n discrete locations, and the searcher successively visits individual locations until finding the hider. Known to both players, a search at location i takes t_i time units and detects the hider—if hidden there—independently with probability α_i, for $i = 1, \ldots, n$. The hider aims to maximize the expected time until detection, while the searcher aims to minimize it. We prove the existence of an optimal strategy for each player. In particular, any optimal mixed hiding strategy hides in each location with a nonzero probability, and there exists an optimal mixed search strategy which can be constructed with up to n simple search sequences.

Keywords: Search games, Gittins index, semi-finite games, search and surveillance.

1 Introduction

Consider the following two-person zero-sum game. A hider chooses one of n locations (henceforth boxes for conciseness) to hide in, and a searcher searches these boxes one at a time in order to find the hider. A search in box i takes $t_i > 0$ time units and will find the hider with probability $\alpha_i \in (0, 1)$ if the hider is there, for $i = 1, \ldots, n$. Due to the possibility of overlook, the searcher may need to visit a box many times to find the hider, and the total time until detection can be arbitrarily long. The searcher wants to minimize the expected total time until the hider is found, while the hider wants to maximize it.

*STOR-i Centre for Doctoral Training, Science and Technology Building, Lancaster University, LA1 4YR, j.clarkson@lancaster.ac.uk

†Operations Research Department, Naval Postgraduate School, Monterey, CA 93943, kylin@nps.edu

‡Department of Management Science, Lancaster University, LA1 4YX, k.glazebrook@lancaster.ac.uk

If the hider announces to the searcher the probability with which they will hide in each box at the beginning of the search, then the resulting search model is one that is well studied in the literature. An optimal search strategy, first discovered by Blackwell (reported in Matula (1964)), is to always next search a box with a maximal probability of detection per unit time at that moment. In other words, if presently the hider is believed to be in box i with probability p_i', $i = 1, \ldots, n$—a value updated throughout the search using Bayes' theorem—then it is optimal to next search box i with a maximal $p_i' \alpha_i / t_i$. A comment by Kelly in Gittins (1979) notes that Blackwell's solution is equivalent to a Gittins index policy obtained by modeling the search as a tractable version of the multi-armed bandit problem (Gittins et al., 2011). Other variations of this search model have been studied in Ross (1969); Kadane (1971); Chew (1973); Wegener (1980); Kress et al. (2008); Clarkson et al. (2020).

The search problem becomes a substantially more complicated *search game* if the searcher does not know the hiding strategy. Whilst the hider has n pure strategies to choose from—each corresponding to hiding in a box—a pure search strategy must specify an indefinite, ordered sequence of boxes for the searcher to search, because the search can take arbitrarily long.

The special case of our search game with $t_i = 1$ for $i = 1, \ldots, n$—the case of unit search time—has been studied in the literature with limited results. Bram (1963) proves an optimal search strategy exists, and Ruckle (1991) solves a few special cases and finds the best pure search strategy. Roberts and Gittins (1978) and Gittins (1989) further specialize to $n = 2$ boxes; the former finds an optimal hiding strategy under certain conditions, while the latter shows the existence of an optimal search strategy that randomly chooses between just two simple search sequences. Gittins and Roberts (1979) develops an algorithm to estimate an optimal hiding strategy for $n \geq 3$ boxes. The technical report Lin and Singham (2015) presents an algorithm to estimate each player's optimal strategy by successively bounding the value of the game. However, as discussed in Hellerstein et al. (2019), an algorithm that guarantees convergence in polynomial time remains a challenge, mainly because the searcher's pure strategy set is infinite.

In another special case with $\alpha_i = 1$ for $i = 1, \ldots, n$—the case of perfect detection—each box needs to be searched at most once. The strategy space of the searcher is thus the set of all $n!$ permutations of the n boxes. Lidbetter (2013) derives an analytic solution to the game, and Lidbetter and Lin (2019) further extends the results to the case with two or more hiders.

Subelman (1981) studies a different objective function in which the searcher wants to maximize the probability of finding the hider by an announced deadline, while the hider wants to minimize it. Lin and Singham (2016) extends the results in Subelman (1981) to show that the searcher has a uniformly optimal strategy that maximizes the probability of finding the hider simultaneously for all deadlines.

The main contribution of this paper is to develop a rigorous mathematical framework to extend earlier results to the search game in its full generality. In particular, we allow for an arbitrary number of boxes n, each having its own search time $t_i > 0$, $i = 1, \ldots, n$. We first prove that the value of the game and an optimal hiding strategy both exist by a simple appeal to general game theory. On the other hand, our next proof that an optimal search strategy exists requires considerable technical innovation. We finally develop properties of the searcher's optimal strategies, and show that the searcher can construct an optimal strategy by a careful randomization among n simple search sequences. Based on these properties, we present a practical procedure to test the optimality of a hiding strategy. If the test is positive, it also yields an optimal search strategy. The findings in this paper both strengthen our understanding of the search game of interest and provide insight into effective practice in real-world search for an intruder.

Our work falls in the general area of search theory, where a searcher seeks a hidden target. Besides the aforementioned papers closely related to our work, search theory has a rich literature with a variety of search models. Common choices of search spaces include the real line, a two-dimensional area, or a network of nodes connected by edges. The target may be stationary, or move around the search space via either a known or random path. Some works assume that the searcher detects the target when their paths cross, and some others consider the possibility of overlook. The searcher may aim to find the target as soon as possible, or maximize the probability of detection before a deadline. For a general review of search theory, see Washburn (2002), Stone (2004) and Stone et al. (2016). For a summary of search games in which the target is a hider actively trying to avoid detection in a time-stationary search space, see Book 1 of Alpern and Gal (2003) and Part I of Alpern et al. (2013). See Garrec and Scarsini (2020) for a novel stochastic search game where the search space, a network, changes randomly over time. If the hider is, for example, a survivor of a disaster, then the hider and the searcher both want to meet up, so the rendezvous search in Book 2 of Alpern and Gal (2003) applies.

The rest of the paper proceeds as follows. Section 2 formulates our search game as a semi-finite two-person zero-sum game and establishes some preliminary results, including the existence of both the value of the game and an optimal strategy for the hider. Section 3 proves the existence of an optimal strategy for the searcher. Section 4 discusses several properties of the optimal strategies and shows that there exists an optimal search strategy which involves a careful randomization among n simple search sequences. Section 5 concludes and offers some future research directions.

2 Model and Preliminaries

Consider a two-person zero-sum *search game* G as follows. A hider decides where to hide among n boxes labeled $1, \ldots, n$, and a searcher decides an ordered sequence of boxes to search. A search in box i takes time t_i and will find the hider with detection probability α_i, $i = 1, \ldots, n$, if the hider is indeed hidden there. These quantities are common knowledge to both players. The game proceeds until the searcher finds the hider, with the total time to detection being the payoff of the game. The searcher wishes to minimize the expected payoff—namely the expected time to detection—while the hider wishes to maximize it.

The hider's pure strategy space is $\{1, \ldots, n\}$, where each pure strategy corresponds to a box in which to hide. A mixed hiding strategy is a probability vector $\mathbf{p} \equiv (p_1, \ldots, p_n) \in \Delta^n$, where p_i is the probability that the hider hides in box i and

$$\Delta^n \equiv \left\{ (p_1, \ldots, p_n) : p_i \in [0,1] \text{ for } i = 1, \ldots, n \text{ and } \sum_{i=1}^{n} p_i = 1 \right\}.$$

The searcher's pure strategy space is the infinite Cartesian product $\mathcal{C} \equiv \{1, 2, \ldots, n\}^\infty$. Each pure strategy is a *search sequence*—an infinite, ordered list of boxes to search until the hider is found. A natural way to define a mixed search strategy is by a probability measure over the uncountable set $\mathcal{C}$ of pure search strategies. In this paper, however, we will restrict the searcher to mixed strategies over countable subsets of $\mathcal{C}$. In other words, we define a mixed search strategy as a non-negative function θ with domain $\mathcal{C}$ such that the set $\{\xi \in \mathcal{C} : \theta(\xi) > 0\}$ is countable, and $\sum_{\xi \in \mathcal{C}} \theta(\xi) = 1$. Under strategy θ, the searcher plays search sequence $\xi \in \mathcal{C}$ with probability $\theta(\xi)$, and we say θ is a *mixture* of those ξ with $\theta(\xi) > 0$. As we shall show in the paper, there exists an optimal mixed strategy for the searcher that mixes no more than n search sequences, so this restriction does not hinder the searcher's ability to find an optimal mixed strategy.

For a search sequence $\xi \in \mathcal{C}$, write $u(i, \xi)$ for the expected time to detection if the hider hides in box i, for $i = 1, \ldots, n$. In other words, $u(i, \xi)$ is the expected payoff for the hider-searcher strategy pair (i, ξ). While the hider's pure strategy space is of size n, the searcher's pure strategy space $\mathcal{C}$ is uncountable; therefore, G is a two-person zero-sum semi-finite game.

The hider seeks a mixed strategy to guarantee the highest possible expected time to detection regardless of what the searcher does, so the hider seeks to determine

$$\text{(Hider)} \qquad v_1 \equiv \max_{\mathbf{p}} \inf_{\xi \in \mathcal{C}} \sum_{i=1}^{n} p_i u(i, \xi). \tag{1}$$

Likewise, the searcher seeks to determine

$$\text{(Searcher)} \qquad v_2 \equiv \inf_{\theta} \max_{i \in \{1,\ldots,n\}} \sum_{\xi \in \mathcal{C}} u(i, \xi)\theta(\xi).$$

By definition, it is clear that $v_1 \leq v_2$. Using the standard results for semi-finite games (see, for example, Chapter 13 in Ferguson (2020)), we can establish the following. Because the payoff function—namely the time to detection—is bounded below by 0, it follows that $v_1 = v_2$, which is the value of G, written by v^*. In addition, the hider has an optimal strategy that guarantees an expected time to detection of at least v^*, and the searcher has an ϵ-optimal strategy; that is, for an arbitrarily small $\epsilon > 0$, the searcher can find a strategy to guarantee an expected time to detection of at most $v^* + \epsilon$.

2.1 A Gittins Search Sequence

Recall that a pure search strategy is a search sequence—an infinite, ordered list of boxes. In this section, we define a particular type of search sequence.

Definition 1 A *Gittins search sequence* against a mixed hiding strategy $\mathbf{p} \equiv (p_1, \ldots, p_n)$ is an infinite, ordered list of boxes that meets the following rule. If $m_i \in \mathbb{N}_0 \equiv \{0, 1, 2, \ldots\}$ searches have already been made of box i during the search process, for $i = 1, \ldots, n$, then the next search is some box j satisfying

$$j = \arg\max_{i \in \{1,\ldots,n\}} \frac{p_i(1-\alpha_i)^{m_i}\alpha_i}{t_i}. \tag{2}$$

The terms in (2) are known as *Gittins indices*, and a Gittins search sequence always next searches a box with a maximal Gittins index. Note that there may be multiple Gittins search sequences against the same hiding strategy $\mathbf{p}$ due to ties for the maximum in (2).

If the searcher *knows* that the hider will choose mixed strategy $\mathbf{p}$, several authors (Norris (1962), Bram (1963), Blackwell (reported in Matula (1964)), Black (1965), Ross (1983)) have proved that any Gittins search sequence against $\mathbf{p}$ is the searcher's best response. This result is also recognized by a comment by Kelly on Gittins (1979), which formulates the search game with known $\mathbf{p}$ as a multi-armed bandit problem optimally solved by Gittins indices. Further, the proof of Ross (1983) shows the reverse is also true; in other words, any pure strategy that is a best response to the hiding strategy $\mathbf{p}$ must be a Gittins search sequence against $\mathbf{p}$. Therefore, we write $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}} \subset \mathcal{C}$ for the set of Gittins search sequences against hiding strategy $\mathbf{p}$.

2.2 Preliminary Properties of Optimal Strategies

This section presents a few preliminary results regarding an optimal hiding strategy and, if one exists, an optimal search strategy.

For any $\xi \in \mathcal{C}$, as a function of $\mathbf{p}$, the expected payoff $\sum_{i=1}^{n} p_i u(i, \xi)$ is a hyperplane in $\mathbb{R}^{n+1}$. Combined with (1), it follows that v^* is the maximum of the lower envelope of an uncountable set of hyperplanes, which is a concave function of $\mathbf{p}$. Let Δ_{opt}^n be the set of $\mathbf{p}$ attaining this maximum, so Δ_{opt}^n is the set of optimal hiding strategies. Since Δ_{opt}^n is the set of maxima of a concave function, it is a convex set. If $|\Delta_{\mathrm{opt}}^n| = 1$, then the optimal hiding strategy is unique; an example with $|\Delta_{\mathrm{opt}}^n| > 1$ can be found in Example 13 in Section 4.

For any *finite* two-person zero-sum game, every pure strategy in the support of an optimal mixed strategy of one player must be a best response to every optimal strategy of the other player (see Theorem 9.1 of Ferguson (2020)). This result can be extended to any *semi-finite* two-person zero-sum game with a value, as seen in the following proposition in the context of our search game G.

Proposition 2 Any optimal search strategy θ^*—if one exists—is a mixture of sequences contained in $\cap_{\mathbf{p} \in \Delta_{\mathrm{opt}}^n} \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$.

Proof. Since the search game G has a value v^*, if the searcher chooses θ^* and the hider any $\mathbf{p}^* \equiv (p_1^*, \ldots, p_n^*) \in \Delta_{\mathrm{opt}}^n$, the expected time until detection is v^*; in other words, we have

$$v^* = \sum_{\xi \in \mathcal{C}} \theta^*(\xi) \left(\sum_{i=1}^{n} p_i^* u(i, \xi) \right). \tag{3}$$

Suppose the statement of the proposition is false, so there exist $\bar{\mathbf{p}} \equiv (\bar{p}_1, \ldots, \bar{p}_n) \in \Delta^n_{\text{opt}}$ and $\bar{\xi} \notin \mathcal{C}^{\text{B}}_{\bar{\mathbf{p}}}$ such that $\theta^*(\bar{\xi}) > 0$. Since $\bar{\mathbf{p}}$ is optimal for the hider, $\sum_{i=1}^n \bar{p}_i u(i, \xi) \geq v^*$ for all $\xi \in \mathcal{C}$. Because $\bar{\xi} \notin \mathcal{C}^{\text{B}}_{\bar{\mathbf{p}}}$, however, $\sum_{i=1}^n \bar{p}_i u(i, \bar{\xi}) > v^*$. Consequently, the right-hand side of (3) must be strictly greater than v^* when $\mathbf{p}^* = \bar{\mathbf{p}}$, leading to a contradiction. ■

It is intuitive that adding a new box will increase the value of the game, because the hider has one more place to hide, so the searcher needs to cover more ground, as seen in the next proposition.

Proposition 3 Write $u(i, \theta)$ for the expected time to detection if the hider hides in box i and the searcher uses mixed strategy θ. The following statements are true.

(i) If $\mathbf{p}^* \equiv (p^*_1, \ldots, p^*_n)$ is optimal for the hider, then $p^*_i > 0$ for $i = 1, \ldots, n$.

(ii) If θ^* is optimal for the searcher, then $u(i, \theta^*) = v^*$ for $i = 1, \ldots, n$.

(iii) Adding a new box increases the value of the game.

Proof. We begin by proving (i), which concerns the hider. Write $u(\mathbf{p})$ for the expected time to detection when the hider chooses $\mathbf{p}$ and the searcher chooses any search sequence in $\mathcal{C}^{\text{B}}_{\mathbf{p}}$; therefore, any optimal hiding strategy maximizes $u(\mathbf{p})$. Let $\xi_{\mathbf{p}}$ be the element of $\mathcal{C}^{\text{B}}_{\mathbf{p}}$ which, when multiple boxes satisfy (2), searches the box with the smallest label.

To prove the statement by contradiction, suppose that we have $\mathbf{p}^* \equiv (p^*_1, \ldots, p^*_n) \in \Delta^n_{\text{opt}}$ with $p^*_k = 0$ for some $k \in \{1, \ldots, n\}$. Without loss of generality, relabel the boxes such that we have $p^*_n = 0$, so

$$u(\mathbf{p}^*) = \sum_{i=1}^{n-1} p^*_i u(i, \xi_{\mathbf{p}^*}). \tag{4}$$

Take any $\epsilon \in (0, 1)$ and consider $\bar{\mathbf{p}} \equiv (\bar{p}_1, \ldots, \bar{p}_n)$, where

$$\begin{aligned} \bar{p}_i &= p^*_i(1 - \epsilon), \qquad i = 1, \ldots, n-1; \\ \bar{p}_n &= \epsilon. \end{aligned}$$

Compare $\xi_{\bar{\mathbf{p}}} \in \mathcal{C}^{\text{B}}_{\bar{\mathbf{p}}}$ and $\xi_{\mathbf{p}^*} \in \mathcal{C}^{\text{B}}_{\mathbf{p}^*}$. Both apply the same rule when multiple boxes satisfy (2), and for any $i, j \in \{1, \ldots, n-1\}$, we have

$$\frac{\bar{p}_i}{\bar{p}_j} = \frac{p^*_i(1 - \epsilon)}{p^*_j(1 - \epsilon)} = \frac{p^*_i}{p^*_j}.$$

Therefore, the subsequence of $\xi_{\bar{\mathbf{p}}}$ consisting of searches of boxes $1, 2, \ldots, n-1$ is identical to $\xi_{\mathbf{p}^*}$. In other words, $\xi_{\bar{\mathbf{p}}}$ is just $\xi_{\mathbf{p}^*}$ with searches of box n inserted between some searches of the first $n-1$ boxes. Hence, we must have $u(i, \xi_{\bar{\mathbf{p}}}) > u(i, \xi_{\mathbf{p}^*})$ for any $i \in \{1, \ldots, n-1\}$ with $p_i^* > 0$. Further, we may choose ϵ small enough so that $\xi_{\bar{\mathbf{p}}}$ does not search box n until at least $u(\mathbf{p}^*)$ time units have passed, ensuring $u(n, \xi_{\bar{\mathbf{p}}}) > u(\mathbf{p}^*)$. From these observations and (4), it follows that

$$\begin{aligned} u(\bar{\mathbf{p}}) &= \epsilon u(n, \xi_{\bar{\mathbf{p}}}) + \sum_{i=1}^{n-1} p_i^* (1-\epsilon) u(i, \xi_{\bar{\mathbf{p}}}) \\ &> \epsilon u(\mathbf{p}^*) + (1-\epsilon) \sum_{i=1}^{n-1} p_i^* u(i, \xi_{\mathbf{p}^*}) = u(\mathbf{p}^*), \end{aligned}$$

contradicting the optimality of $\mathbf{p}^*$, and therefore proving (i).

Next, we prove (ii), concerning the searcher. Suppose θ^* is optimal for the searcher and $\mathbf{p}^*$ is optimal for the hider. The expected time to detection under the strategy pair $(\mathbf{p}^*, \theta^*)$ is

$$v^* = \sum_{i=1}^{n} p_i^* u(i, \theta^*). \tag{5}$$

To prove by contradiction, suppose that $u(j, \theta^*) < v^*$ for some $j \in \{1, \ldots, n\}$. By (5), there must either exist $k \in \{1, \ldots, n\}$ such that $u(k, \theta^*) > v^*$, or we must have $p_j^* = 0$. The former cannot happen as θ^* guarantees the searcher an expected time to detection of at most v^*. The latter cannot happen by (i), which leads to a contradiction proving (ii).

Finally, we prove (iii) by showing that $v_{n+1}^* > v_n^*$, where v_n^* is the value of an n-box game, and v_{n+1}^* is the value if a new box is added to the n-box game. In the game with $n+1$ boxes, the hider can guarantee an expected payoff of at least v_n^* by not hiding in the new box, so $v_{n+1}^* \geq v_n^*$. However, any such strategy has $p_{n+1} = 0$ so it is not optimal by (i). Therefore, v_n^* is not the value of the $(n+1)$-box game, so $v_{n+1}^* > v_n^*$, proving (iii). ■

Note that (i) in Proposition 3 is also proved by Bram (1963) for unit-search-time G via a different method to the proof above.

3 Existence of Optimal Search Strategies

The aim of this section is to prove that the searcher has an optimal strategy that guarantees an expected time to detection of at most v^* in the search game G. By using an S-game formulation, we first show that the searcher has an optimal mixed strategy in a modified search game. The results of Section 2.2 then enable us to draw the same conclusion for the search game G.

3.1 S-Game Formulation for Search Game G

We begin by reformulating G as an S-game of Blackwell and Girshick (1954), in which, instead of choosing a pure strategy in $\mathcal{C}$, the searcher chooses a vector in the set

$$\mathcal{S} \equiv \{(u(1,\xi),\ldots,u(n,\xi)) : \xi \in \mathcal{C}\} \subset \mathbb{R}^n.$$

If the hider hides in box $i \in \{1,\ldots,n\}$ and the searcher selects $(u(1,\xi),\ldots,u(n,\xi)) \in \mathcal{S}$, then the payoff is $u(i,\xi)$.

By Theorem 2.4.1 of Blackwell and Girshick (1954), the searcher selecting a mixed strategy is equivalent to choosing a point in $\mathrm{Conv}(\mathcal{S})$, the convex hull of $\mathcal{S}$. By Theorem 2.4.2 of Blackwell and Girshick (1954), if $\mathcal{S}$, or equivalently $\mathrm{Conv}(\mathcal{S})$, is closed, then there exists an optimal search strategy which is a mixture of at most n search sequences.

The intuition behind this result is the following, adapted from Section 13.1 of Ferguson (2020). If $\mathbf{s} \equiv (s_1,\ldots,s_n) \in \mathrm{Conv}(\mathcal{S})$, then there exists a mixed search strategy which, if the hider hides in box i, achieves an expected payoff s_i, $i = 1,\ldots,n$. It follows that the value of the game v^* satisfies

$$v^* = \inf_{\mathbf{s}\in\mathrm{Conv}(\mathcal{S})} \left\{ \max_{i\in\{1,\ldots,n\}} s_i \right\}. \tag{6}$$

If $\mathrm{Conv}(\mathcal{S})$ is closed, then the infimum in (6) is attained, since the payoff function u is bounded below by 0. Consequently, there exists $\mathbf{s}^* \equiv (s_1^*,\ldots,s_n^*) \in \mathrm{Conv}(\mathcal{S})$ with $\max_{i\in\{1,\ldots,n\}} s_i^* = v^*$, so $\mathbf{s}^*$ is an optimal search strategy. See Ruckle (1991) for more on the geometrical interpretation of optimal strategies in the search game, particularly for $n = 2$. The reason why there exists an optimal search strategy that is a mixture of at most n search sequences follows from the fact that the hider has n pure strategies, so $\mathrm{Conv}(\mathcal{S})$ sits in an n-dimensional space. See the proof of Theorem 15 for a full argument.

Bram (1963) concludes that an optimal search strategy exists for the search game with $t_i = 1$ for $i = 1,\ldots,n$ by showing that $\mathrm{Conv}(\mathcal{S})$ is closed. In this paper, we take a different approach, using the modified search game below, to extend the result to arbitrary $t_i > 0$, $i = 1,\ldots,n$.

3.2 A Modified Search Game $G(\epsilon)$

In this section, we introduce a modified search game and use its S-game formulation to show that an optimal search strategy exists in this modified game.

Consider a search game $G(\epsilon)$, parametrized by $\epsilon \in (0, 1/n)$, identical to G in all aspects apart from its set of pure search strategies, which are constructed by the following. For $i = 1, \ldots, n$, write

$$M_i(\epsilon) \equiv \inf\{u(i, \xi) : \xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}} \text{ with } p_i < \epsilon\}. \tag{7}$$

In words, among all Gittins search sequences against hiding strategies with $p_i < \epsilon$, $M_i(\epsilon)$ is the smallest expected time to detection in G if the hider is in box i. Unlike G, in $G(\epsilon)$, a pure strategy $\zeta_i(\epsilon)$ is available to the searcher for $i = 1, \ldots, n$. When selected in $G(\epsilon)$, for $i = 1, \ldots, n$, $\zeta_i(\epsilon)$ results in payoff $M_i(\epsilon)$ if the hider is in box i or payoff 0 otherwise. In addition, available to the searcher in $G(\epsilon)$ are Gittins search sequences against hiding strategies in $\Delta^n(\epsilon)$, where

$$\Delta^n(\epsilon) \equiv \{\mathbf{p} \in \Delta^n : p_i \geq \epsilon,\ i = 1, \ldots, n\}.$$

To summarize, in $G(\epsilon)$, the searcher has the following pure strategy set:

$$\mathcal{C}(\epsilon) \equiv \{\xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}} : \mathbf{p} \in \Delta^n(\epsilon)\} \cup \{\zeta_i(\epsilon) : i = 1, \ldots, n\}. \tag{8}$$

For any $\epsilon \in (0, 1/n)$, since the payoff in $G(\epsilon)$ is bounded below by 0, by the standard results for semi-finite games (see, for example, Section 13 in Ferguson (2020)), we can conclude that $G(\epsilon)$ has a value and optimal hiding strategy. To prove that $G(\epsilon)$ has an optimal search strategy, we consider its S-game formulation, in which the searcher chooses a vector in

$$\mathcal{S}(\epsilon) \equiv \{(u(1, \xi), \ldots, u(n, \xi)) : \xi \in \mathcal{C}(\epsilon)\} \subset \mathbb{R}^n.$$

By Theorem 2.4.2 of Blackwell and Girshick (1954), if $\mathcal{S}(\epsilon)$ is closed, then there exists an optimal search strategy for the game $G(\epsilon)$. Write

$$\bar{\mathcal{S}}(\epsilon) \equiv \mathcal{S}(\epsilon) \setminus \{(u(1, \zeta_i(\epsilon)), \ldots, u(n, \zeta_i(\epsilon))),\ i = 1, \ldots, n\}.$$

Since $\mathcal{S}(\epsilon)$ adds only a finite subset of $\mathbb{R}^n$ to $\bar{\mathcal{S}}(\epsilon)$, if $\bar{\mathcal{S}}(\epsilon)$ is closed, then $\mathcal{S}(\epsilon)$ is closed. The majority of this section is devoted to showing that $\bar{\mathcal{S}}(\epsilon)$ is closed.

Write $U(\xi) \equiv (u(1,\xi),\ldots,u(n,\xi))$; therefore, any element of $\bar{\mathcal{S}}(\epsilon)$ takes the form $U(\xi)$ where ξ is a Gittins search sequence against some $\mathbf{p} \in \Delta^n(\epsilon)$.

By Definition 1, the next box searched by any Gittins search sequence against a hiding strategy $\mathbf{p}$ must satisfy (2). If, at some point whilst following a Gittins search sequence against $\mathbf{p}$, multiple boxes satisfy (2), we say the searcher has encountered a *tie* and $\mathbf{p}$ is a *tie point*. When encountering a tie, any Gittins search sequence must immediately search each of the tied boxes once in some arbitrary order. Thereafter, any two Gittins search sequences will be identical until another tie is encountered.

Note that an equivalent definition of a tie point $\mathbf{p}$ is $|\mathcal{C}^{\mathrm{B}}_{\mathbf{p}}| > 1$. If $|\mathcal{C}^{\mathrm{B}}_{\mathbf{p}}| = 1$, we say $\mathbf{p}$ is a *non-tie point*. If $\mathbf{p}$ is a non-tie point, then there is a unique Gittins search sequence against $\mathbf{p}$, whereas, for a tie point $\mathbf{p}$, a specific Gittins search sequence against $\mathbf{p}$ is entirely determined by how we break ties between boxes.

How the searcher chooses to break ties is important, since knowledge of the searcher's tie-breaking preferences could be used by the hider to their advantage. Write S_n for the set of permutations of $\{1,\ldots,n\}$; a permutation $\sigma \in S_n$ serves as a preference ordering to choose which box to search next if a tie is encountered. For example, a tie between boxes 1, 2 and 4 is broken in the order 2, 1, 4 by the permutation $(2,3,1,4) \in S_4$.

There are a variety of rules for breaking ties. For example, the searcher could roll an $n!$-sided die every time a tie is encountered to determine a preference ordering. For another example, the searcher could use permutation $(1,2,\ldots,n)$ if the tie involves box 1, and $(n,\ldots,2,1)$ otherwise. Still another example, the searcher could use a predetermined sequence of tie-breaking permutations to break ties; in other words, before starting the search, write down an infinite sequence of permutations of $\{1,2,\ldots,n\}$, and use the jth element of the sequence to break the jth tie encountered. Since any Gittins search sequence can be attained by such a predetermined sequence of permutations, we shall write $\mathcal{R} \equiv (S_n)^\infty$, the set of all infinite sequences of elements of S_n, and restrict our attention to tie-breaking rules using a predetermined sequence of permutations.

For example, suppose $n = 5$, and the jth tie encountered following a Gittins search sequence against $\mathbf{p}$ involves boxes 2, 3 and 5. Consider a tie-breaking rule $\mathbf{r} \in \mathcal{R}$ and suppose its jth element is $(5,4,2,3,1)$. Then, under rule $\mathbf{r}$, tie j is split by searching boxes 5, 2 and 3 in that order. Note that changing the jth element of $\mathbf{r}$ to $(4,5,1,2,3)$ does not affect the Gittins search sequence generated, demonstrating that multiple rules in $\mathcal{R}$ can generate the same Gittins search sequence. For convenience, we write $\mathbf{r} \equiv r_1, r_2, \ldots$ such that $r_j \in S_n$ is the jth element in $\mathbf{r}$, for $j \in \mathbb{N} \equiv \{1, 2, \ldots\}$.

For any $\mathbf{r}, \tilde{\mathbf{r}} \in \mathcal{R}$, define their distance by

$$d(\mathbf{r}, \tilde{\mathbf{r}}) \equiv \sup_{j \in \mathbb{N}} \frac{D(r_j, \tilde{r}_j)}{j}, \tag{9}$$

where

$$D(x, y) = \begin{cases} 1, & \text{if } x \neq y, \\ 0, & \text{if } x = y. \end{cases}$$

In other words, for any $k \in \mathbb{N}$, we have $d(\mathbf{r}, \tilde{\mathbf{r}}) = 1/k$ if and only if the first $k-1$ elements of $\mathbf{r}$ and $\tilde{\mathbf{r}}$ are identical and the kth differs. In addition, $d(\mathbf{r}, \tilde{\mathbf{r}}) = 0$ if and only if $\mathbf{r} = \tilde{\mathbf{r}}$. It is straightforward to verify that the distance function in (9) satisfies the three conditions of positivity, symmetry, and the triangle inequality. Hence, $\mathcal{R}$ with the distance function in (9) forms a metric space. Further, because the metric space $\mathcal{R} \equiv (S_n)^\infty$ is a product of compact spaces, by the countable version of Tychonoff's theorem, $\mathcal{R}$ is compact (see, for example, Theorem 7.4.2. in Wilansky (2008)).

Recall that we write

$$\Delta^n \equiv \left\{ (p_1, \ldots, p_n) : p_i \geq 0, \sum_{i=1}^{n} p_i = 1 \right\}$$

for the space of mixed hiding strategies. Write f for the function from $\mathcal{R} \times \Delta^n \to \bar{\mathcal{S}}(\epsilon)$ satisfying $f(\mathbf{r}, \mathbf{p}) = U(\xi(\mathbf{r}, \mathbf{p}))$, where $\xi(\mathbf{r}, \mathbf{p})$ is the Gittins search sequence against $\mathbf{p}$ that breaks ties using rule $\mathbf{r}$. In other words, f maps a tie-breaking rule $\mathbf{r}$ and a hiding strategy $\mathbf{p}$ to the vector of expected times to detection (across all boxes) of the corresponding Gittins search sequence.

The continuity of f in both of its arguments will be fundamental to achieving our aim of showing that $\bar{\mathcal{S}}(\epsilon)$ is closed. We prove this continuity by a series of lemmas whose proofs are deferred to Appendix A. First, we establish continuity in the first argument of f in the following lemma with proof in Appendix A.1.

Lemma 4 For any $\mathbf{p} \in \Delta^n$ with $p_i > 0$ for $i = 1, \ldots, n$, the function $f(\mathbf{r}, \mathbf{p})$ is continuous in $\mathbf{r}$.

Before moving on to continuity in the second argument of f, we first use Lemma 4 to prove the compactness of a subset of $\bar{\mathcal{S}}(\epsilon)$ concerning only Gittins search sequences against a fixed hiding strategy with $p_i > 0$ for $i = 1, \ldots, n$.

Lemma 5 For any $\mathbf{p} \in \Delta^n$ with $p_i > 0$ for $i = 1, \ldots, n$, the set

$$\mathcal{S}_{\mathbf{p}} \equiv \{U(\xi) : \xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}\} \tag{10}$$

is compact.

Proof. To prove this result, we will use the theorem that the image of a compact set under a continuous function is compact (see, for example, Theorem 9.3.4 in Strichartz (1995)).

Recall from its definition that $f(\mathbf{r}, \mathbf{p}) = U(\xi(\mathbf{r}, \mathbf{p}))$, where $\xi(\mathbf{r}, \mathbf{p})$ is the Gittins search sequence against $\mathbf{p}$ that breaks ties using rule $\mathbf{r}$. Therefore, for any fixed $\mathbf{p} \in \Delta^n$, $\mathcal{S}_{\mathbf{p}}$ is the image of $\mathcal{R}$ under f with respect to its first argument. Because $\mathcal{R}$ is compact, and for $\mathbf{p} \in \Delta^n$ with $p_i > 0$ for $i = 1, \ldots, n$ the function $f(\mathbf{r}, \mathbf{p})$ is continuous in $\mathbf{r}$ (Lemma 4), it follows that $\mathcal{S}_{\mathbf{p}}$ is compact. ■

Next, via two lemmas, we investigate the continuity of the function $f(\mathbf{r}, \mathbf{p})$ in its second argument at different $\mathbf{p}$ for any fixed $\mathbf{r} \in \mathcal{R}$. The first lemma, whose proof is deferred to Appendix A.2, deals with the simpler case of continuity at non-tie points in $\Delta^n(\epsilon)$.

Lemma 6 If $\mathbf{p} \in \Delta^n(\epsilon)$ is a non-tie point, then, for any fixed first argument $\mathbf{r} \in \mathcal{R}$, f is continuous in its second argument at $\mathbf{p}$.

Now we consider the continuity of f in its second argument at tie points in $\Delta^n(\epsilon)$, the more challenging case. Informally, if $\mathbf{p} \in \Delta^n(\epsilon)$ is a tie point, then f is only continuous in its second argument at $\mathbf{p}$ for certain fixed first arguments $\mathbf{r} \in \mathcal{R}$, and only approaching $\mathbf{p}$ via certain paths in Δ^n.

To state the continuity conditions precisely, we first need a few definitions. Let $\mathbf{p} \equiv (p_1, \ldots, p_n)$ be a tie point in $\Delta^n(\epsilon)$. Recall S_n as the set of permutations of $\{1, \ldots, n\}$, and let $\Sigma \subseteq S_n$. For $\sigma \in S_n$, write $\sigma(i)$ for the number in the ith position of σ and write

$$\Delta^n(\mathbf{p}, \Sigma) \equiv \left\{\mathbf{x} \equiv (x_1, \ldots, x_n) \in \Delta^n : x_{\sigma(i)}/p_{\sigma(i)} \geq x_{\sigma(i+1)}/p_{\sigma(i+1)},\ i = 1, \ldots, n-1,\ \sigma \in \Sigma\right\}. \tag{11}$$

For $\mathbf{x} \in \Delta^n(\mathbf{p}, \Sigma)$, if there is a tie among Gittins indices calculated using $\mathbf{p}$ and $\sigma \in \Sigma$ is used to break the tie, then replacing $\mathbf{p}$ with $\mathbf{x}$ in the Gittins index calculation—and then applying σ if needed—would produce the same search order for the tied boxes. Clearly $\mathbf{p} \in \Delta^n(\mathbf{p}, \Sigma)$ for any subset Σ of S_n. Further, for each $\sigma \in \Sigma$, the $\mathbf{x}$ in (11) is defined by $n-1$ inequalities, each of which induces a half-space including $\mathbf{p}$. Therefore, $\Delta^n(\mathbf{p}, \Sigma)$ is the intersection of Δ^n and a finite number of half-spaces, so it is a convex set.

Informally, the following lemma says that, if its first argument is fixed to be some $\mathbf{r} \in \mathcal{R}$ containing only elements of $\Sigma \subset S_n$, then f is continuous in its second argument at $\mathbf{p}$ approaching from any path in $\Delta^n(\mathbf{p}, \Sigma)$.

Lemma 7 Suppose $\mathbf{p} \in \Delta^n(\epsilon)$ is a tie point and $\Sigma \subset S_n$. Let $\{\mathbf{x}_a : a \in \mathbb{N}\}$ be a sequence in $\Delta^n(\mathbf{p}, \Sigma)$ with $\lim_{a\to\infty} \mathbf{x}_a = \mathbf{p}$. Then, for any $\mathbf{r} \in \mathcal{R}$ whose elements all belong to Σ, we have $\lim_{a\to\infty} f(\mathbf{r}, \mathbf{x}_a) = f(\mathbf{r}, \mathbf{p})$.

The proof of Lemma 7 is deferred to Appendix A.3. We are now ready to show that $\bar{\mathcal{S}}(\epsilon)$ is closed.

Proposition 8 The set $\bar{\mathcal{S}}(\epsilon)$ is closed.

Proof. Write $\{\mathbf{s}_a : a \in \mathbb{N}\}$ for a convergent sequence in $\bar{\mathcal{S}}(\epsilon)$, and write $\mathbf{s}_0 \equiv \lim_{a\to\infty} \mathbf{s}_a$ for its limit. To show that $\bar{\mathcal{S}}(\epsilon)$ is closed, we need to show that $\mathbf{s}_0 \in \bar{\mathcal{S}}(\epsilon)$.

Since each element of $\bar{\mathcal{S}}(\epsilon)$ corresponds to some mixed hiding strategy $\mathbf{p} \in \Delta^n(\epsilon)$ and tie-breaking rule $\mathbf{r} \in \mathcal{R}$, $\bar{\mathcal{S}}(\epsilon)$ is equal to $f(\mathcal{R} \times \Delta^n(\epsilon))$. Therefore, for all $a \in \mathbb{N}$, we may choose $\mathbf{x}_a \in \Delta^n(\epsilon)$ and $\mathbf{r}_a \in \mathcal{R}$ such that $\mathbf{s}_a = f(\mathbf{r}_a, \mathbf{x}_a)$. Further, since $\Delta^n(\epsilon)$ is bounded, the sequence $\{\mathbf{x}_a : a \in \mathbb{N}\}$ has a convergent subsequence $\{\mathbf{x}_{h(a)} : a \in \mathbb{N}\}$, and, since $\Delta^n(\epsilon)$ is closed, $\mathbf{x}_0 \equiv \lim_{a\to\infty} \mathbf{x}_{h(a)} \in \Delta^n(\epsilon)$. Any infinite subsequence of the convergent sequence $\{\mathbf{s}_a\}$ must converge to the same limit as $\{\mathbf{s}_a\}$; therefore, we have $\mathbf{s}_0 = \lim_{a\to\infty} \mathbf{s}_{h(a)}$.

We consider two cases. First, suppose that $\{\mathbf{x}_{h(a)}\}$ attains its limit $\mathbf{x}_0$. In other words, there exists $A \in \mathbb{N}$ such that $\mathbf{x}_{h(a)} = \mathbf{x}_0$ for all $a \geq A$. In this instance, the sequence $\{\mathbf{s}_{h(a)} : a \geq A\}$, which has limit $\mathbf{s}_0$, is a sequence in the set $\mathcal{S}_{\mathbf{x}_0}$ defined in (10), which was shown to be compact by Lemma 5. Therefore, $\mathbf{s}_0 \in \mathcal{S}_{\mathbf{x}_0} \subset \bar{\mathcal{S}}(\epsilon)$, completing the proof for the first case.

Second, suppose that $\{\mathbf{x}_{h(a)}\}$ does not attain its limit $\mathbf{x}_0$. Since $\mathbf{x}_0 \in \Delta^n(\epsilon)$, we have $f(\mathbf{r}, \mathbf{x}_0) \in \bar{\mathcal{S}}(\epsilon)$ for any $\mathbf{r} \in \mathcal{R}$. To complete the proof for the second case, we show that $\mathbf{s}_0 = f(\mathbf{r}, \mathbf{x}_0)$ for some $\mathbf{r} \in \mathcal{R}$. To do this, we split the second case into two subcases.

First, consider the easier subcase in which $\mathbf{x}_0$ is a non-tie point. Then, for any $\mathbf{r} \in \mathcal{R}$, we have

$$\mathbf{s}_0 = \lim_{a\to\infty} \mathbf{s}_{h(a)} = \lim_{a\to\infty} f(\mathbf{r}_{h(a)}, \mathbf{x}_{h(a)}) = f(\mathbf{r}, \mathbf{x}_0),$$

where the last equality follows by Lemma 6 and the fact that the $f(\mathbf{r}, \mathbf{x}_0)$ are equal for all $\mathbf{r} \in \mathcal{R}$.

The rest of the proof concerns the more challenging subcase, in which $\mathbf{x}_0 \equiv (x_{0,1}, \ldots, x_{0,n})$ is a tie point. Note that, for any $\mathbf{x} \equiv (x_1, \ldots, x_n) \in \Delta^n$, there exists a subset $\Sigma_{\mathbf{x}} \subseteq S_n$ for which $\sigma \in \Sigma_{\mathbf{x}}$ if and only if

$$x_{\sigma(i)}/x_{0,\sigma(i)} \geq x_{\sigma(i+1)}/x_{0,\sigma(i+1)},\ i = 1, 2, \ldots, n-1.$$

Recalling the definition in (11), we have $\mathbf{x} \in \Delta^n(\mathbf{x}_0, \Sigma_{\mathbf{x}})$, and further $\Sigma_{\mathbf{x}}$ is the unique subset of maximal size such that $\mathbf{x} \in \Delta^n(\mathbf{x}_0, \Sigma_{\mathbf{x}})$. Since the elements $\{x_i/x_{0,i},\ i = 1, \ldots, n\}$ must lie in some size order, $\Sigma_{\mathbf{x}}$ is non-empty for any $\mathbf{x} \in \Delta^n(\epsilon)$.

Since there are a finite number of subsets of S_n, there must exist $\Sigma^* \subset S_n$ and a convergent subsequence of $\{\mathbf{x}_{h(a)}\}$, say $\{\mathbf{x}_m : m \in \mathbb{N}\}$, such that $\Sigma_{\mathbf{x}_m} = \Sigma^*$ for all $m \in \mathbb{N}$. In other words, $\{\mathbf{x}_m\}$ is a sequence in $\Delta^n(\mathbf{x}_0, \Sigma^*)$.

Since $\{\mathbf{x}_m\}$ is a convergent subsequence of $\{\mathbf{x}_{h(a)}\}$, we have $\lim_{m\to\infty} \mathbf{x}_m = \mathbf{x}_0$, and

$$\lim_{m\to\infty} f(\mathbf{r}_m, \mathbf{x}_m) = \lim_{m\to\infty} \mathbf{s}_m = \mathbf{s}_0.$$

To complete the proof, it remains to show $\mathbf{s}_0 = f(\mathbf{r}, \mathbf{x}_0)$ for some $\mathbf{r} \in \mathcal{R}$. We split into two further subcases, numbered below.

1. Suppose that there are finitely many tie points in $\{\mathbf{x}_m\}$. In this case, we may choose M such that there are no tie points in the sequence $\{\mathbf{x}_m : m \geq M\}$. Therefore, for all $m \geq M$, we have $f(\mathbf{r}_m, \mathbf{x}_m) = f(\mathbf{r}, \mathbf{x}_m)$ for all $\mathbf{r} \in \mathcal{R}$. Let $\mathbf{r}^* \in \mathcal{R}$ contain only elements in Σ^*. Then we have

$$\mathbf{s}_0 = \lim_{m\to\infty} f(\mathbf{r}_m, \mathbf{x}_m) = \lim_{m\to\infty} f(\mathbf{r}^*, \mathbf{x}_m) = f(\mathbf{r}^*, \mathbf{x}_0),$$

   where the last equality follows by Lemma 7.

2. Now suppose that there are infinitely many tie points in $\{\mathbf{x}_m\}$. We begin with two observations for any $\mathbf{x} \in \Delta^n(\epsilon)$.

   (i) The position of any equalities in the ordering of the terms $\{x_i/x_{0,i},\ i = 1, \ldots, n\}$ completely determines $\Sigma_{\mathbf{x}}$. In particular, for any pair of boxes $i, j \in \{1, \ldots, n\}$, we have $x_i/x_{0,i} = x_j/x_{0,j}$ if and only if there exists $\sigma_1, \sigma_2 \in \Sigma_{\mathbf{x}}$ with $\sigma_1(i) > \sigma_1(j)$ and $\sigma_2(j) > \sigma_2(i)$. Also, we have $x_i/x_{0,i} > x_j/x_{0,j}$ if and only if $\sigma(i) > \sigma(j)$ for all $\sigma \in \Sigma_{\mathbf{x}}$.

(ii) For any pair of boxes $i, j \in \{1, \ldots, n\}$ and any $y, z \in \mathbb{N}_0 \equiv \{0, 1, 2, \ldots\}$ write

$$k_{i,j}(y, z) \equiv \frac{\alpha_j (1 - \alpha_j)^z t_i}{\alpha_i (1 - \alpha_i)^y t_j},$$

recalling that α_i (resp. t_i) is the detection probability (resp. search time) of box i, $i = 1, \ldots, n$. Then, inspection of (2) shows that, following a Gittins search sequence against $\mathbf{x}$, there is a tie between boxes i and j after y (resp. z) searches of box i (resp. j) have been made if and only if $x_i/x_j = k_{i,j}(y, z)$.

Now consider the sequence $\{\mathbf{x}_m\}$, and write $\mathbf{x}_m \equiv (x_{m,1}, \ldots, x_{m,n})$ for the mth term, $m \in \mathbb{N}$. Since $\{\mathbf{x}_m\}$ has limit $\mathbf{x}_0$, for any two boxes $i, j \in \{1, \ldots, n\}$, we have

$$\frac{x_{m,i}}{x_{m,j}} \to \frac{x_{0,i}}{x_{0,j}} \quad \text{as} \quad m \to \infty. \tag{12}$$

Now choose $\mathbf{x} \in \{\mathbf{x}_m\}$ and suppose that $c, d \in \{1, \ldots, n\}$ satisfy $x_c/x_{0,c} \neq x_d/x_{0,d}$. Then, by (i), the same must be true for every element of $\{\mathbf{x}_m\}$ since $\Sigma_{\mathbf{x}_m} = \Sigma^*$ for all $m \in \mathbb{N}$. Hence, the limit in (12) is never attained for $i = c$ and $j = d$. In other words, $x_{m,c}/x_{m,d}$ approaches but never reaches $x_{0,c}/x_{0,d}$ as $m \to \infty$.

Let $y, z \in \mathbb{N}_0$ and consider $k_{c,d}(y, z)$ defined in (ii). There are two scenarios. First, we may have $k_{c,d}(y, z) = x_{0,c}/x_{0,d}$; in this scenario, since the limit in (12) is never attained, in no Gittins search sequence against any element of $\{\mathbf{x}_m\}$ is there a tie between boxes c and d after y (resp. z) searches of box c (resp. d) have been made. Second, if $k_{c,d}(y, z) \neq x_{0,c}/x_{0,d}$, by the limit in (12), there exists a finite smallest element of $\mathbb{N}$, say $M_{c,d}(y, z)$, such that the same statement holds after the $M_{c,d}(y, z)$th term of $\{\mathbf{x}_m\}$; in other words, in no Gittins search sequence against any element of $\{\mathbf{x}_m : M_{c,d}(y, z) \geq m\}$ is there a tie between boxes c and d after y (resp. z) searches of box c (resp. d) have been made.

For any $b \in \mathbb{N}$, write

$$M_{c,d}^b \equiv \max\{M_{c,d}(y, z) : y, z \in \mathbb{N}_0 \text{ with } y + z \leq b\}.$$

It follows that, whilst the total number of searches of boxes c and d is no larger than b, no ties involving *both* boxes c and d are encountered in a Gittins search sequence against any element of $\{\mathbf{x}_m : m \geq M_{c,d}^b\}$. Clearly $\{M_{c,d}^b : b \in \mathbb{N}\}$ forms an increasing sequence; therefore, as $m \to \infty$, any tie involving both boxes c and d occurs increasingly later and later into the search, so the effect on the expected time to detection of how such a tie is broken decreases to 0.

Therefore, as $m \to \infty$, only the manner in which ties involving *only* boxes i and j satisfying $x_i/x_{0,i} = x_j/x_{0,j}$ are broken has any effect on the expected time to detection under a Gittins search sequence against an element of $\{\mathbf{x}_m\}$. Without a loss of generality, suppose such a tie is between boxes $1, \ldots, y$, for some $y \in \{2, \ldots, n\}$. Suppose the tie is broken by $\sigma \in S_n$. By (i), since $x_1/x_{0,1} = \cdots = x_y/x_{0,y}$, there exists $\sigma^* \in \Sigma_{\mathbf{x}}$ which ranks boxes $1, \ldots, y$ in the same order as σ. Therefore, breaking the tie using σ^* leads to boxes $1, \ldots, y$ being searched in the same order as breaking the tie using σ. It follows that there exists $\mathbf{r}^*$ with only elements in Σ^* such that, if, for all $m \in \mathbb{N}$, we replace $\mathbf{r}_m$ with $\mathbf{r}^*$, as $m \to \infty$, the effect on the expected time to detection tends to 0. In other words,

$$\mathbf{s}_0 = \lim_{m\to\infty} f(\mathbf{r}_m, \mathbf{x}_m) = \lim_{m\to\infty} f(\mathbf{r}^*, \mathbf{x}_m) = f(\mathbf{r}^*, \mathbf{x}_0),$$

where the last equality follows by Lemma 7.

The proof is completed. ■

We conclude this section by showing that there exists an optimal search strategy in the game $G(\epsilon)$.

Proposition 9 For any $\epsilon \in (0, 1/n)$, the game $G(\epsilon)$ has an optimal search strategy which is a mixture of at most n search sequences.

Proof. Since $\mathcal{S}(\epsilon)$ adds only a finite subset of $\mathbb{R}^n$ to $\bar{\mathcal{S}}(\epsilon)$ and $\bar{\mathcal{S}}(\epsilon)$ is closed (Proposition 8), it follows that $\mathcal{S}(\epsilon)$ is closed. It then follows from Theorem 2.4.2 of Blackwell and Girshick (1954) that there exists an optimal search strategy for the game $G(\epsilon)$. ■

3.3 The Connection Between G and $G(\epsilon)$

In this section, we show that the existence of an optimal search strategy in $G(\epsilon)$ implies the existence of an optimal search strategy in G. The following result draws upon the properties in Section 2.2 to conclude that, for small enough ϵ, the games G and $G(\epsilon)$ are almost equivalent.

Lemma 10 Consider G and its set of optimal hiding strategies Δ^n_{opt}. For any $\mathbf{p}^* \in \Delta^n_{\text{opt}}$, there exists $\epsilon_{\mathbf{p}^*} \in (0, 1/n]$ such that, for all $\epsilon \in (0, \epsilon_{\mathbf{p}^*})$, the games G and $G(\epsilon)$ share the same value, $\mathbf{p}^*$ is optimal in $G(\epsilon)$ as well as in G, and a search strategy is optimal in G if and only if it is optimal in $G(\epsilon)$.

Proof. Let $\mathbf{p}^* \equiv (p^*_1, \ldots, p^*_n) \in \Delta^n_{\text{opt}}$ and $\epsilon_1 \equiv \min_{i\in\{1,\ldots,n\}} p^*_i$; we have $\epsilon_1 > 0$ by (i) in Proposition 3. Further, under any mixed hiding strategy, some box is chosen with at most probability $1/n$, so $\epsilon_1 \le 1/n$.

The function $M_i(\epsilon)$ in (7) decreases in ϵ, since the set over which the infimum is taken grows with ϵ. Write $\mathbf{p} \equiv (p_1, \ldots, p_n)$. If $p_i = 0$, then any $\xi \in \mathcal{C}^{\text{B}}_{\mathbf{p}}$ never searches box i, so $u(i, \xi)$ is infinite, and hence $M_i(\epsilon) \uparrow \infty$ as $\epsilon \downarrow 0$. On the other hand, if $p_i = 1$, then any $\xi \in \mathcal{C}^{\text{B}}_{\mathbf{p}}$ only searches box i, so $u(i, \xi) = t_i/\alpha_i$, and hence $M_i(\epsilon) \downarrow t_i/\alpha_i \le v^*$ as $\epsilon \uparrow 1$, where v^* is the value of G. Combining the above information, we may conclude that

$$\epsilon_2 \equiv \sup\{\epsilon : M_i(\epsilon) > v^*/p^*_i,\ i = 1, \ldots, n\}$$

exists, and $M_i(\epsilon) > v^*/p^*_i$ for all $\epsilon \in (0, \epsilon_2)$, $i = 1, \ldots, n$.

Let $\epsilon_{\mathbf{p}^*} \equiv \min(\epsilon_1, \epsilon_2)$; we show that $\epsilon_{\mathbf{p}^*}$ satisfies the conditions of the lemma. For any $\bar{\mathcal{C}} \subset \mathcal{C}$, write

$$u(\mathbf{p}, \bar{\mathcal{C}}) \equiv \inf_{\xi\in\bar{\mathcal{C}}} u(\mathbf{p}, \xi),$$

where $u(\mathbf{p}, \xi)$ is the expected time to detection if the hider chooses mixed strategy $\mathbf{p}$ and the searcher chooses pure strategy ξ. Throughout the following, let $\epsilon \in (0, \epsilon_{\mathbf{p}^*})$.

Recall from (8) that $\mathcal{C}(\epsilon)$ is the pure strategy set in $G(\epsilon)$. In G, a best response to a hiding strategy $\mathbf{p}$ is any sequence in $\mathcal{C}^{\text{B}}_{\mathbf{p}}$, which leads to an expected time to detection of $u(\mathbf{p}, \mathcal{C})$, where $\mathcal{C}$ is the pure strategy set in G. Bearing the above in mind, we show that $u(\mathbf{p}, \mathcal{C}) \ge u(\mathbf{p}, \mathcal{C}(\epsilon))$ for any hiding strategy $\mathbf{p}$ by considering two cases.

1. $\mathbf{p} \in \Delta^n(\epsilon)$. In this case, $\mathcal{C}^{\text{B}}_{\mathbf{p}}$ is contained in $\mathcal{C}(\epsilon)$; therefore, if the hider chooses $\mathbf{p}$, the searcher does no worse when $\mathcal{C}$ is replaced with $\mathcal{C}(\epsilon)$, so $u(\mathbf{p}, \mathcal{C}) \ge u(\mathbf{p}, \mathcal{C}(\epsilon))$.

2. $\mathbf{p} \notin \Delta^n(\epsilon)$. In this case, there exists $j \in \{1, \ldots, n\}$ such that $p_j < \epsilon$. By the construction of the sequences $\zeta_j(\epsilon)$ defined at the start of Section 3.2, we have $u(i, \zeta_j(\epsilon)) \le u(i, \xi)$ for any

$\xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$, $i = 1, \ldots, n$. Therefore, $\zeta_j(\epsilon) \in \mathcal{C}(\epsilon) \setminus \mathcal{C}$ dominates any sequence in $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$. It follows that $u(\mathbf{p}, \mathcal{C}) \geq u(\mathbf{p}, \mathcal{C}(\epsilon))$.

Now consider $u(\mathbf{p}^*, \mathcal{C}(\epsilon))$; by the above, $v^* = u(\mathbf{p}^*, \mathcal{C}) \geq u(\mathbf{p}^*, \mathcal{C}(\epsilon))$. Since $\epsilon < \epsilon_1$, we have $\mathbf{p}^* \in \Delta^n(\epsilon)$, and hence $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}} \subset \mathcal{C}(\epsilon)$. The only pure search strategies in $\mathcal{C}(\epsilon)$ that, when the hider chooses $\mathbf{p}^*$, could achieve a lower expected time to detection than a sequence in $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}}$ are those not in $\mathcal{C}$, namely $\{\zeta_i(\epsilon),\ i = 1, \ldots, n\}$. Therefore, we have

$$
\begin{aligned}
u(\mathbf{p}^*, \mathcal{C}(\epsilon)) &= \min \left(v^*, \min_{i \in \{1,\ldots,n\}} \left\{ \sum_{j=1}^{n} p_j^* u(j, \zeta_i(\epsilon)) \right\} \right) \\
&= \min \left(v^*, \min_{i \in \{1,\ldots,n\}} p_i^* M_i(\epsilon) \right) = v^*,
\end{aligned}
$$

where the final equality holds since $\epsilon < \epsilon_2$.

To conclude, we have $u(\mathbf{p}, \mathcal{C}) \geq u(\mathbf{p}, \mathcal{C}(\epsilon))$ for all hiding strategies $\mathbf{p}$, and $v^* = u(\mathbf{p}^*, \mathcal{C}) = u(\mathbf{p}^*, \mathcal{C}(\epsilon))$. It follows that $\mathbf{p}^*$ is optimal in $G(\epsilon)$ and the value of $G(\epsilon)$ is v^*.

As for the searcher, since the set of pure hiding strategies and the value are the same for G and $G(\epsilon)$, any optimal search strategy in G is optimal in $G(\epsilon)$ if it is available to the searcher in $G(\epsilon)$ and vice versa. By Proposition 2, any optimal search strategy in G chooses only search sequences in $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}}$, available in $G(\epsilon)$ since $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}} \subset \mathcal{C}(\epsilon)$. Further, for $i = 1, \ldots, n$, we have $p_i^* M_i(\epsilon) > v^*$. Therefore, it is suboptimal in $G(\epsilon)$ for the searcher to choose any strategy in $\mathcal{C}(\epsilon) \setminus \mathcal{C} = \{\zeta_i(\epsilon),\ i = 1, \ldots, n\}$. It follows that a search strategy is optimal in G if and only if it is optimal in $G(\epsilon)$, completing the proof. ■

We conclude this section with its main result.

Theorem 11 In the search game G, for any optimal hiding strategy $\mathbf{p}^*$, there exists an optimal search strategy which is a mixture of at most n elements of $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}}$.

Proof. By Lemma 9, there exists a search strategy θ^*, optimal in $G(\epsilon)$, which is a mixture of at most n search sequences. By Lemma 10, θ^* is also optimal in the search game G. By Proposition 2, for any optimal hiding strategy $\mathbf{p}^*$, the search sequences mixed by θ^* must all belong to $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}}$, completing the proof. ■

4 Properties of Optimal Strategies

While we have shown that each player has an optimal strategy in the search game G, it turns out that each player's optimal strategy need not be unique. In this section, we demonstrate how to identify an optimal hider-searcher strategy pair, present an example where the hider has multiple optimal strategies, and show that the searcher may always choose a simple optimal strategy among the many available. These findings will underpin the development of a practical optimality test for any hiding strategy.

We begin by combining Propositions 2 and 3 to identify simple conditions on a hider-searcher strategy pair which are both necessary and sufficient for optimality.

Theorem 12 Write $u(\mathbf{p}, \theta)$ for the expected time to detection if the hider and the searcher use mixed strategies $\mathbf{p}$ and θ, respectively. The mixed strategy $\mathbf{p}$ (resp. θ) is optimal for the hider (resp. searcher) if and only if

(i) θ is a mixture of some subset of $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$.

(ii) $u(i, \theta) = u(\mathbf{p}, \theta)$, for $i = 1, \ldots, n$.

Proof. First, we prove the forwards implication. If $\mathbf{p}$ is optimal for the hider and θ is optimal for the searcher, then (i) follows from Proposition 2. Further, if $\mathbf{p}$ and θ are optimal, then $u(\mathbf{p}, \theta)$ is the value of the game, so (ii) follows from Proposition 3 (ii).

Second, we prove the backwards implication. By (i), $\mathbf{p}$ guarantees an expected time to detection of at least $u(\mathbf{p}, \theta)$ regardless of what the searcher does. In addition, by (ii), θ guarantees an expected time to detection equal to $u(\mathbf{p}, \theta)$ regardless of what the hider does. Therefore, neither the hider nor the searcher can obtain a better guarantee than $u(\mathbf{p}, \theta)$; it follows that $\mathbf{p}$ and θ are an optimal strategy pair. ■

Note that the backwards implication of Theorem 12 is Theorem 8.3 of Gittins (1989) applied to the search game.

Theorem 12 shows that any optimal search strategy θ^* is an *equalizing* strategy; in other words, whenever the searcher plays θ^*, the expected time to detection is the same no matter where the hider hides. Theorem 5.2 of Ruckle (1991) shows that, in the search game with unit search times, there exists an equalizing pure search strategy ξ. However, ξ is not necessarily optimal by Theorem 12, because it is not necessarily a Gittins search sequence against any hiding strategy. Further, the proof of Theorem 5.2 of Ruckle (1991) does not extend to the search game with arbitrary search times, as it relies on $u(k,\xi)$ being unaffected when the positions of a search of box i and box j in ξ are switched.

We next use Theorem 12 to demonstrate that it is possible for the hider to have multiple optimal strategies.

Example 13

Consider a two-box search game where box i has search time t_i and detection probability α_i, $i = 1, 2$, with $\alpha_1 < \alpha_2$ and $t_1 > t_2$. Write p for the probability that the hider hides in box 1. Inspection of (2) shows that if

$$p \in \left[\frac{\alpha_2/t_2}{\alpha_2/t_2 + \alpha_1/t_1}, \frac{\alpha_2/t_2}{\alpha_2/t_2 + \alpha_1(1-\alpha_1)/t_1} \right], \tag{13}$$

then there exists a Gittins search sequence against the hiding strategy $(p, 1-p)$ that begins by searching box 1, followed by box 2, and then box 1 again.

Suppose that

$$(1-\alpha_2) = (1-\alpha_1)^2; \tag{14}$$

therefore, for any p, if the searcher makes, in any order, two unsuccessful searches of box 1 and one unsuccessful search of box 2, then the posterior probability that the hider is in box 1 returns to p, and hence the problem has reset itself. It follows that the sequence ξ that repeats the cycle of boxes $1, 2, 1$ indefinitely is a Gittins search sequence against any hiding strategy $(p, 1-p)$ with p satisfying (13).

Calculate

$$u(1,\xi) = \sum_{k=1}^{\infty} (1-\alpha_1)^{2(k-1)} \alpha_1 \left[(k-1)(2t_1 + t_2) + t_1 + (1-\alpha_1)k(2t_1 + t_2) \right],$$

$$u(2,\xi) = \sum_{k=1}^{\infty} (1-\alpha_2)^{k-1} \alpha_2 \left[(k-1)(2t_1 + t_2) + t_1 + t_2 \right].$$

By rewriting

$$u(1,\xi) = \sum_{k=1}^{\infty} (1-\alpha_1)^{2(k-1)} [\alpha_1 t_1 + \alpha_1(1-\alpha_1)(2t_1 + t_2) + (1-(1-\alpha_1)^2)(k-1)(2t_1 + t_2)],$$

and using (14), we have

$$u(1,\xi) - u(2,\xi) = [\alpha_1(t_1 + (1-\alpha_1)(2t_1+t_2)) - \alpha_2(t_1+t_2)] \sum_{k=1}^{\infty}(1-\alpha_2)^{k-1},$$

from which it follows that $u(1,\xi) = u(2,\xi)$ if and only if

$$\alpha_1 = \frac{t_1 - t_2}{t_1}. \tag{15}$$

By choosing α_2 to satisfy (14), it follows from Theorem 12 that ξ is optimal for the searcher and $(p, 1-p)$ for any p satisfying (13) is optimal for the hider.

For a numerical example, if $\alpha_1 = 0.4$, $\alpha_2 = 0.64$, $t_1 = 1$ and $t_2 = 0.6$, then $(p, 1-p)$ for any $p \in [8/11, 40/49]$ is an optimal choice for the hider. □

By Theorem 12, it is sufficient for the searcher to consider Gittins search sequences against any optimal hiding strategy. Therefore, if there exists an optimal $\mathbf{p}$ against which there is a unique Gittins search sequence ξ (so $|\mathcal{C}^{\mathrm{B}}_{\mathbf{p}}| = 1$), then the pure strategy ξ is optimal for the searcher. One example can be found in Example 13; any $(p, 1-p)$ with p satisfying (13) (aside from the two endpoints) is both optimal for the hider and has $\mathcal{C}^{\mathrm{B}}_{(p,1-p)} = \{\xi\}$, so the pure strategy ξ is optimal for the searcher.

Interestingly, the condition $|\mathcal{C}^{\mathrm{B}}_{\mathbf{p}}| = 1$ for some optimal $\mathbf{p}$ is not always necessary for the existence of an optimal pure search strategy. Ruckle (1991) shows that in the search game with $\alpha_i = 0.5$ and $t_i = 1$, $i = 1, \ldots, n$, the unique optimal hiding strategy $\mathbf{p}^*$ selects each box with probability $1/n$, and therefore any search sequence repeatedly passing through the n boxes is in $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$, meaning $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$ is of infinite size. Ruckle (1991) further shows that one such search sequence is optimal, namely

$$1, 2, \ldots, n; n, n-1, \ldots, 1; n, n-1, \ldots, 1; \ldots,$$

which passes through the boxes once in ascending order, then in descending order ad infinitum. Since all boxes are identical, by symmetry, any sequence beginning by permuting $1, 2, \ldots, n$ before applying the reverse permutation ad infinitum is also optimal.

Other than specifically constructed cases, however, it is more common that any optimal search strategy is a mixed strategy when $|\mathcal{C}^{\mathrm{B}}_{\mathbf{p}}| > 1$ for all optimal $\mathbf{p}$, a more challenging case on which we focus henceforth. Recall that, by Definition 1, the next box searched by any Gittins search sequence against a hiding strategy $\mathbf{p}$ must satisfy (2). If $|\mathcal{C}^{\mathrm{B}}_{\mathbf{p}}| > 1$, at some point in the search, the searcher must encounter a *tie* where some $k \in \{2, \ldots, n\}$ boxes satisfy (2). At such a tie, any Gittins search sequence must search the k tied boxes next in some arbitrary order. Thereafter, any two Gittins search sequences will be identical until another tie is encountered. Therefore, elements of $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$ differ from one another only in how they break ties.

The way that the searcher breaks ties is important, however, as information about the searcher's tie breaking rule could be exploited by the hider. A mixed tie-breaking strategy adds some unpredictability to the searcher's behavior. Recall S_n as the set of permutations of $\{1, \ldots, n\}$. As discussed in Section 3.2, whilst there are a variety of rules the searcher can use to choose how to break ties, any Gittins search sequence can be attained by some rule which writes down an infinite sequence of elements in S_n and uses the jth element of the sequence as a preference ordering to choose which box to search next when the jth tie is encountered, for $j = 1, 2, \ldots$. For example, a tie between boxes 1, 2 and 4 is broken in the order 2, 1, 4 by the permutation $(2, 3, 1, 4) \in S_4$.

One particularly simple tie-breaking rule to generate a Gittins search sequence is to break every tie using the same preference ordering. Write $\xi_{\sigma,\mathbf{p}}$ for the Gittins search sequence against $\mathbf{p}$ that breaks every tie encountered using $\sigma \in S_n$. We define the following subset of $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$:

$$\widehat{\mathcal{C}}_{\mathbf{p}}^{\mathrm{B}} \equiv \{\xi_{\sigma,\mathbf{p}} : \sigma \in S_n\}.$$

Whilst $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$ could be an infinite set, $|\widehat{\mathcal{C}}_{\mathbf{p}}^{\mathrm{B}}| \leq n!$ since $|S_n| = n!$. By Theorem 11, there exists an optimal search strategy that is a mixture of at most n elements of $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}}$ for any optimal hiding strategy $\mathbf{p}^*$. The aim of the remainder of this section is to show that the same holds true if we replace $\mathcal{C}_{\mathbf{p}^*}^{\mathrm{B}}$ with $\widehat{\mathcal{C}}_{\mathbf{p}^*}^{\mathrm{B}}$.

For any search strategy θ (pure or mixed), write $U(\theta) \equiv (u(1, \theta), \ldots, u(n, \theta))$. For any hiding strategy $\mathbf{p}$, recall

$$\mathcal{S}_{\mathbf{p}} \equiv \{U(\xi) : \xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}\}$$

from (10), and additionally define

$$\widehat{\mathcal{S}}_{\mathbf{p}} \equiv \{U(\xi) : \xi \in \widehat{\mathcal{C}}_{\mathbf{p}}^{\mathrm{B}}\},$$

noting that $\widehat{\mathcal{S}}_{\mathbf{p}} \subset \mathcal{S}_{\mathbf{p}} \subset \mathbb{R}^n$. Clearly, if some search strategy θ is a mixture of a subset of $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$, then $U(\theta)$ can be written as a convex combination of the elements in $\mathcal{S}_{\mathbf{p}}$. The following lemma shows that the same statement is true replacing $\mathcal{S}_{\mathbf{p}}$ with $\widehat{\mathcal{S}}_{\mathbf{p}}$ if $\mathbf{p}$ hides in each box with a nonzero probability.

Lemma 14 For any hiding strategy $\mathbf{p} \equiv (p_1, \ldots, p_n)$ with $p_i > 0$, $i = 1, \ldots, n$, the convex hull of $\mathcal{S}_{\mathbf{p}}$ is equal to the convex hull of $\widehat{\mathcal{S}}_{\mathbf{p}}$.

Proof. If $|\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}| = 1$ then $\mathcal{S}_{\mathbf{p}} = \widehat{\mathcal{S}}_{\mathbf{p}}$ and the result is trivially true. For the rest of the proof, assume $|\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}| > 1$.

Write $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ for the convex hull of $\mathcal{S}_{\mathbf{p}}$. By definition, $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ is convex, and, since $\mathcal{S}_{\mathbf{p}}$ is compact (Lemma 5) in the finite-dimensional vector space $\mathbb{R}^n$, $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ is also compact. (See Corollary 5.33 of Charalambos and Aliprantis (2013).) Therefore, we may apply the Krein-Milman theorem to deduce that $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ is equal to the convex hull of its extreme points. To prove Lemma 14, we show that $\widehat{\mathcal{S}}_{\mathbf{p}}$ is the set of extreme points of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$. We first note the following useful facts.

- By definition, a point $\mathbf{x} \in \mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ is extreme if and only if, for any $\mathbf{y}, \mathbf{z} \in \mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ and $\lambda \in (0,1)$ satisfying $\mathbf{x} = \lambda \mathbf{y} + (1-\lambda)\mathbf{z}$, we have $\mathbf{x} = \mathbf{y} = \mathbf{z}$. In other words, the only way we can express $\mathbf{x}$ as a convex combination of elements of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ is by $\mathbf{x}$ itself.

- For any $\mathbf{s} \equiv (s_1, \ldots, s_n) \in \mathcal{S}_{\mathbf{p}}$, the weighted average $\sum_{i=1}^n s_i p_i$ is equal to the expected time to detection if the hider chooses $\mathbf{p}$ and the searcher any best response $\xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$. Therefore, all elements of $\mathcal{S}_{\mathbf{p}}$ lie on the same hyperplane, say H, in $\mathbb{R}^n$.

- By the definition of $\mathcal{S}_{\mathbf{p}}$, we have

$$\mathcal{S}_{\mathbf{p}} \subset R \equiv \left\{ (v_1, \ldots, v_n) \in \mathbb{R}^n : \min_{\xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}} u(i,\xi) \leq v_i \leq \max_{\xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}} u(i,\xi),\ i = 1, \ldots, n \right\},$$

  where R is a hyperrectangle in n-dimensional space (also known as an n-orthotope). Further, since $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ is the smallest convex set containing $\mathcal{S}_{\mathbf{p}}$, and R is also a convex set containing $\mathcal{S}_{\mathbf{p}}$, we have $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}}) \subseteq R$.

The proof will be done by double inclusion. In the first half of the double inclusion proof, we show that any point in $\widehat{\mathcal{S}}_{\mathbf{p}}$ is an extreme point of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$. Let $\mathbf{x} \equiv (x_1, \ldots, x_n) \in \widehat{\mathcal{S}}_{\mathbf{p}}$. Then there exists $\xi_\sigma \in \widehat{\mathcal{C}}_{\mathbf{p}}^{\mathrm{B}}$ such that $\mathbf{x} = U(\xi_\sigma)$, where ξ_σ breaks all ties using some $\sigma \in S_n$. Without loss of generality, let $\sigma = (1, 2 \ldots, n)$. Suppose that

$$\mathbf{x} = \lambda \mathbf{y} + (1-\lambda)\mathbf{z} \tag{16}$$

for some $\mathbf{y}, \mathbf{z} \in \mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ and $\lambda \in (0,1)$. To prove that $\mathbf{x}$ is extreme, we show that we must have $\mathbf{x} = \mathbf{y} = \mathbf{z}$.

Since, when breaking any tie, ξ_σ gives preference to box 1 over any other box, no other search sequence in $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$ makes the jth search of box 1 any sooner than ξ_σ, $j = 1, 2, \ldots$; therefore, we have $x_1 = \min_{\xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}} u(1, \xi)$. For any $\mathbf{v} \equiv (v_1, \ldots, v_n) \in \mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$, since $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}}) \subset R$, we must have $v_1 \geq x_1$. It follows that, for (16) to hold, we must have $y_1 = z_1 = x_1$.

We now demonstrate how this argument may be repeated to show that $y_2 = z_2 = x_2$. Let $\mathcal{C}_{1,\mathbf{p}}^{\mathrm{B}}$ be the elements of $\mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$ which, when breaking any tie involving box 1, give preference to box 1; therefore, for any $\xi \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$, we have $u(1, \xi) = x_1$ if and only if $\xi \in \mathcal{C}_{1,\mathbf{p}}^{\mathrm{B}}$. Recall $U(\xi) \equiv (u(1, \xi), \ldots, u(n, \xi))$ and write $\mathcal{S}_{1,\mathbf{p}} \equiv \{U(\xi) : \xi \in \mathcal{C}_{1,\mathbf{p}}^{\mathrm{B}}\}$; therefore, for any $\mathbf{v} = (v_1, \ldots, v_n) \in \mathcal{S}_{\mathbf{p}}$, we have $v_1 = x_1$ if and only if $\mathbf{v} \in \mathcal{S}_{1,\mathbf{p}}$. Since $\mathcal{S}_{1,\mathbf{p}} \subset \mathcal{S}_{\mathbf{p}}$, we have $\mathrm{Conv}(\mathcal{S}_{1,\mathbf{p}}) \subset \mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$. Note that $\mathbf{y}, \mathbf{z} \in \mathrm{Conv}(\mathcal{S}_{1,\mathbf{p}})$ since $y_1 = z_1 = x_1$.

Write

$$R_1 \equiv \left\{ (x_1, v_2, \ldots, v_n) \in \mathbb{R}^n : \min_{\xi \in \mathcal{C}_{1,\mathbf{p}}^{\mathrm{B}}} u(i, \xi) \leq v_i \leq \max_{\xi \in \mathcal{C}_{1,\mathbf{p}}^{\mathrm{B}}} u(i, \xi),\ i = 2, \ldots, n \right\}.$$

Then R_1, a hyperrectangle in $(n-1)$-dimensional space, is a convex set containing $\mathcal{S}_{1,\mathbf{p}}$, so $\mathrm{Conv}(\mathcal{S}_{1,\mathbf{p}}) \subset R_1$. Since, when breaking any tie, ξ_σ gives preference to box 1 over box 2, but then to box 2 over any other box i, $i = 3, \ldots, n$, no other sequence in $\mathcal{C}_{1,\mathbf{p}}^{\mathrm{B}}$ makes the jth search of box 2 sooner than ξ_σ, $j = 1, 2, \ldots$; therefore, we have $x_2 = \min_{\xi \in \mathcal{C}_{1,\mathbf{p}}^{\mathrm{B}}} u(2, \xi)$. Since any $\mathbf{v} \equiv (v_1, \ldots, v_n) \in \mathrm{Conv}(\mathcal{S}_{1,\mathbf{p}})$ also belongs to R_1, we must have $v_2 \geq x_2$. It follows that, for (16) to hold, we must have $y_2 = z_2 = x_2$.

We may repeat the above argument a further $n-3$ times to conclude that $y_i = z_i = x_i$ for $i = 1, \ldots, n-1$. Finally, since $\mathbf{y}, \mathbf{z}$ and $\mathbf{x}$ all lie in the same hyperplane H in $\mathbb{R}^n$, we must have $y_n = z_n = x_n$, so $\mathbf{y} = \mathbf{z} = \mathbf{x}$, and $\mathbf{x}$ must be an extreme point of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$.

In the second half of the double inclusion proof, we show that any extreme point of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ is in $\widehat{\mathcal{S}}_{\mathbf{p}}$. We will prove the contrapositive of this statement; i.e., we show that any $\mathbf{a} \in \mathrm{Conv}(\mathcal{S}_{\mathbf{p}}) \setminus \widehat{\mathcal{S}}_{\mathbf{p}}$ is not an extreme point of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$.

To begin, note that, by definition, any element of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ can be written as a convex combination of some $m \in \mathbb{N}$ elements of $\mathcal{S}_{\mathbf{p}}$. If $\mathbf{b} \in \mathrm{Conv}(\mathcal{S}_{\mathbf{p}}) \setminus \mathcal{S}_{\mathbf{p}}$, then any such convex combination must contain at least $m \geq 2$ elements of $\mathcal{S}_{\mathbf{p}}$, so $\mathbf{b}$ is not an extreme point of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$. Hence, our task is reduced to showing that any $\mathbf{a} \in \mathcal{S}_\mathbf{p} \setminus \widehat{\mathcal{S}}_\mathbf{p}$ is not an extreme point of $\mathrm{Conv}(\mathcal{S}_\mathbf{p})$.

Consider $\mathbf{a} \in \mathcal{S}_\mathbf{p} \setminus \widehat{\mathcal{S}}_\mathbf{p}$, and write $\xi_\mathbf{a}$ for the corresponding search sequence in $\mathcal{C}^\mathrm{B}_\mathbf{p} \setminus \widehat{\mathcal{C}}^\mathrm{B}_\mathbf{p}$ satisfying $\mathbf{a} = U(\xi_\mathbf{a})$. Since $\xi_\mathbf{a} \notin \widehat{\mathcal{C}}^\mathrm{B}_\mathbf{p}$, there exists no permutation in S_n with which $\xi_\mathbf{a}$ breaks every tie it encounters. It follows that there must exist some $k \in \{2, \dots, n\}$ boxes (without loss of generality boxes $1, \dots, k$) and k ties encountered by $\xi_\mathbf{a}$ such that no permutation of $\{1, \dots, k\}$ serves as a preference ordering for how all k ties are broken. Suppose, again with no loss of generality, tie m involves (at least) boxes m and $m+1$, with box m searched before box $m+1$ by $\xi_\mathbf{a}$, $m = 1, \dots, k-1$, and tie k involves (at least) boxes k and 1, with box k searched before box 1 by $\xi_\mathbf{a}$. Note that ties $1, \dots, k$ are not necessarily consecutive nor in chronological order.

We now aim to construct a mixture of elements of $\mathcal{C}^\mathrm{B}_\mathbf{p}$ which mimics the performance of $\xi_\mathbf{a}$; it will follow that $\mathbf{a}$ is not an extreme point of $\mathrm{Conv}(\mathcal{S}_\mathbf{p})$.

Consider tie k, which is broken by $\xi_\mathbf{a}$ using the preference ordering

$$j_1, \dots, j_{\beta-1}, k, j_{\beta+1}, \dots, j_{\gamma-1}, 1, j_{\gamma+1}, \dots, j_n,$$

bearing in mind that box j_i is not necessarily involved in tie k, $i \in \{1, \dots, n\} \setminus \{\beta, \gamma\}$. In other words, box k ranks in the βth position, and box 1 ranks in the γth position, for some $1 \leq \beta < \gamma \leq n$.

Let $\xi_{\mathbf{a},\beta} \in \mathcal{C}^\mathrm{B}_\mathbf{p}$ break tie k using preference ordering

$$j_1, \dots, j_{\beta-1}, 1, k, j_{\beta+1}, \dots, j_{\gamma-1}, j_{\gamma+1}, \dots, j_n,$$

and all other ties in the same order as $\xi_\mathbf{a}$. Similarly, let $\xi_{\mathbf{a},\gamma} \in \mathcal{C}^\mathrm{B}_\mathbf{p}$ break tie k using preference ordering

$$j_1, \dots, j_{\beta-1}, j_{\beta+1}, \dots, j_{\gamma-1}, 1, k, j_{\gamma+1}, \dots, j_n,$$

and all other ties in the same order as $\xi_\mathbf{a}$. Note that if $\gamma = \beta + 1$ (so there are no boxes searched between boxes k and 1 by $\xi_\mathbf{a}$ when breaking tie k), then $\xi_{\mathbf{a},\beta} = \xi_{\mathbf{a},\gamma}$, but the following argument is still valid.

Note that $u(i, \xi_{\mathbf{a},\beta}) = u(i, \xi_{\mathbf{a},\gamma}) = u(i, \xi_\mathbf{a})$ for any box $i \in \{j_1, \dots, j_{\beta-1}, j_{\gamma+1}, \dots, j_n\}$. Recall t_i is the search time of box i, and let

$$\theta_k \equiv \frac{t_k}{t_1 + t_k} \xi_{\mathbf{a},\beta} \oplus \frac{t_1}{t_1 + t_k} \xi_{\mathbf{a},\gamma},$$

where $p\xi_1 \oplus (1-p)\xi_2$ denotes the mixture of ξ_1 and ξ_2, which selects ξ_1 with probability p and ξ_2 with probability $1 - p$. Clearly $u(i, \theta_k) = u(i, \xi_\mathbf{a})$ for $i \in \{j_1, \dots, j_{\beta-1}, j_{\gamma+1}, \dots, j_n\}$. For $i \in$

$\{j_{\beta+1}, \ldots, j_{\gamma-1}\}$, let w_i be the probability that the hider is found on the first search of box i after the kth tie is reached, conditional on the hider being in box i. Then we have $u(i, \xi_{\mathbf{a},\beta}) = u(i, \xi_{\mathbf{a}}) + w_i t_1$ and $u(i, \xi_{\mathbf{a},\gamma}) = u(i, \xi_{\mathbf{a}}) - w_i t_k$ for $i \in \{j_{\beta+1}, \ldots, j_{\gamma-1}\}$. It follows that

$$u(i, \theta_k) = \frac{t_k}{t_1 + t_k} (u(i, \xi_{\mathbf{a}}) + w_i t_1) + \frac{t_1}{t_1 + t_k} (u(i, \xi_{\mathbf{a}}) - w_i t_k) = u(i, \xi_{\mathbf{a}})$$

for $i \in \{j_{\beta+1}, \ldots, j_{\gamma-1}\}$.

Because $u(k, \xi_{\mathbf{a},\gamma}) > u(k, \xi_{\mathbf{a},\beta}) > u(k, \xi_{\mathbf{a}})$, we have $u(k, \theta_k) > u(k, \xi_{\mathbf{a}})$. Since $U(\theta_k)$ and $U(\xi_{\mathbf{a}})$ lie in the same hyperplane, H, in $\mathbb{R}^n$, we must have $u(1, \theta_k) < u(1, \xi_{\mathbf{a}})$. To summarize, we have

$$u(k, \theta_k) > u(k, \xi_{\mathbf{a}}), \quad u(1, \theta_k) < u(1, \xi_{\mathbf{a}}), \quad \text{and} \quad u(i, \theta_k) = u(i, \xi_{\mathbf{a}}) \text{ for } i \neq 1, k. \tag{17}$$

Now, for $m \in \{1, \ldots, k-1\}$, we repeat the same procedure with tie m for boxes m and $m+1$ to create θ_m which satisfies:

$$u(m, \theta_m) > u(m, \xi_{\mathbf{a}}), \quad u(m+1, \theta_m) < u(m+1, \xi_{\mathbf{a}}), \quad \text{and} \quad u(i, \theta_m) = u(i, \xi_{\mathbf{a}}) \text{ for } i \neq m, m+1. \tag{18}$$

To complete the proof, we develop a mixture of $\{\theta_1, \ldots, \theta_k\}$ which mimics the performance of $\xi_{\mathbf{a}}$. To begin, by (18) with $m = 1, 2$, for any $\lambda \in (0, 1)$, the mixture

$$\theta_{1,2}(\lambda) \equiv \lambda\theta_1 \oplus (1 - \lambda)\theta_2 \tag{19}$$

satisfies

$$u(1, \theta_{1,2}(\lambda)) > u(1, \xi_{\mathbf{a}}), \quad u(3, \theta_{1,2}(\lambda)) < u(3, \xi_{\mathbf{a}}), \quad \text{and} \quad u(i, \theta_{1,2}(\lambda)) = u(i, \xi_{\mathbf{a}}) \text{ for } i = 4, \ldots, n.$$

Also by (18), there exists $\lambda^* \in (0, 1)$ such that $\theta_{1,2} \equiv \theta_{1,2}(\lambda^*)$ satisfies $u(2, \theta_{1,2}) = u(2, \xi_{\mathbf{a}})$. Therefore, we have

$$u(1, \theta_{1,2}) > u(1, \xi_{\mathbf{a}}), \quad u(3, \theta_{1,2}) < u(3, \xi_{\mathbf{a}}), \quad \text{and} \quad u(i, \theta_{1,2}) = u(i, \xi_{\mathbf{a}}) \text{ for } i \neq 1, 3. \tag{20}$$

By (18) with $m = 3$ and (20), there exists a mixture, $\theta_{1,2,3}$, of θ_3 and $\theta_{1,2}$ satisfying

$$u(1, \theta_{1,2,3}) > u(1, \xi_{\mathbf{a}}), \quad u(4, \theta_{1,2,3}) < u(4, \xi_{\mathbf{a}}), \quad \text{and} \quad u(i, \theta_{1,2,3}) = u(i, \xi_{\mathbf{a}}) \text{ for } i \neq 1, 4.$$

We may repeat this process of mixing θ_m and $\theta_{1,\ldots,m-1}$ to create $\theta_{1,\ldots,m}$ for $m = 4, \ldots, k-1$, with the resulting $\theta_{1,\ldots,k-1}$ satisfying

$$u(1, \theta_{1,\ldots,k-1}) > u(1, \xi_{\mathbf{a}}), \quad u(k, \theta_{1,\ldots,k-1}) < u(k, \xi_{\mathbf{a}}), \quad \text{and} \quad u(i, \theta_{1,\ldots,k-1}) = u(i, \xi_{\mathbf{a}}) \text{ for } i \neq 1, k. \tag{21}$$

Finally, by (17) and (21), we may mix θ_k and $\theta_{1,\ldots,k-1}$ to create $\theta_{1,\ldots,k}$ satisfying $u(i,\theta_{1,\ldots,k}) = u(i,\xi_{\mathbf{a}})$ for $i = 2,\ldots,n$. Yet, since $U(\theta_{1,\ldots,k})$ and $U(\xi_{\mathbf{a}})$ both lie in H, we must also have $u(1,\theta_{1,\ldots,k}) = u(1,\xi_{\mathbf{a}})$. It follows that, for some $\bar{\lambda} \in (0,1)$, we have

$$\mathbf{a} = U(\xi_{\mathbf{a}}) = U(\theta_{1,\ldots,k}) = \bar{\lambda} U(\theta_{1,\ldots,k-1}) + (1-\bar{\lambda}) U(\theta_k).$$

By the construction of the θ_m for $m = 1,\ldots,k$ as mixtures of elements of $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}}$, $U(\theta_{1,\ldots,k-1})$ and $U(\theta_k)$ are two distinct elements in $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ different from $\mathbf{a}$, showing that $\mathbf{a}$ is not an extreme point of $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}})$ and completing the proof. ■

Lemma 14 allows us to present the main theorem of this section, which strengthens Theorem 11 by replacing $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$ with $\widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}^*}$.

Theorem 15 In search game G, for any optimal hiding strategy $\mathbf{p}^*$, there exists an optimal search strategy which is a mixture of at most n elements of $\widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}^*}$.

Proof. Let $\mathbf{p}^* \equiv (p_1^*,\ldots,p_n^*)$ be an optimal hiding strategy. By Theorem 11, there exists an optimal search strategy θ^* which is a mixture of elements of $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$. Therefore, $U(\theta^*)$ can be written as a convex combination of elements of $\mathcal{S}_{\mathbf{p}^*}$ and hence belongs to $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}^*})$, the convex hull of $\mathcal{S}_{\mathbf{p}^*}$.

By (i) in Proposition 3, $p_i^* > 0$ for $i = 1,\ldots,n$. Therefore, by Lemma 14, $\mathrm{Conv}(\mathcal{S}_{\mathbf{p}^*}) = \mathrm{Conv}(\widehat{\mathcal{S}}_{\mathbf{p}^*})$. It follows that $U(\theta^*) \in \mathrm{Conv}(\widehat{\mathcal{S}}_{\mathbf{p}^*})$ and hence can be written as a convex combination of elements of $\widehat{\mathcal{S}}_{\mathbf{p}^*}$.

By Carathéodory's theorem, $U(\theta^*)$ can be written as a convex combination of at most $n+1$ elements in $\widehat{\mathcal{S}}_{\mathbf{p}^*}$. Yet, for any hiding strategy $\mathbf{p} \equiv (p_1,\ldots,p_n)$ and $\mathbf{s} \equiv (s_1,\ldots,s_n) \in \mathcal{S}_{\mathbf{p}}$, the weighted average $\sum_{i=1}^n s_i p_i$ is equal to the expected time to detection if the hider chooses $\mathbf{p}$ and the searcher any best response $\xi \in \mathcal{C}^{\mathrm{B}}_{\mathbf{p}}$. Therefore, all elements of $\mathcal{S}_{\mathbf{p}}$ lie on the same hyperplane in $\mathbb{R}^n$, and hence so do all elements of $\widehat{\mathcal{S}}_{\mathbf{p}} \subset \mathcal{S}_{\mathbf{p}}$. It follows that the number of elements in the convex combination of $U(\theta^*)$ can be reduced to at most n, so θ^* is a mixture of at most n strategies in $\widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}^*}$. ■

Proposition 8.5 of Gittins (1989) proves a special case of Theorem 15 with $n = 2$ and $t_1 = t_2 = 1$, but that proof does not extend directly to $n \geq 3$; see Appendix B for some discussion. Our result applies to an arbitrary number of boxes and to general search times.

The significance of Theorem 15 is twofold. First, it shows that, in order to construct an optimal search strategy, it is sufficient to consider only the simple type of Gittins search sequence which uses the same preference ordering to break every tie. Second, while Theorem 11 shows that an optimal search strategy exists and can be formed among some collection of (at most) n search strategies, it gives no indication of how to find a suitable collection among $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$, the possibly uncountable set of best responses to an optimal hiding strategy $\mathbf{p}^*$. Theorem 15 shows that, once $\mathbf{p}^*$ is found, the searcher only needs to consider the search sequences in $\widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}^*}$—of which there are at most $n!$—reducing the search game G to a finite matrix game. Then, by computing $u(i, \xi)$ to required accuracy for $i = 1, \ldots, n$ and $\xi \in \widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}^*}$, one can use linear programming methods (see Section 10.4 in Ferguson (2020)) to compute, to an arbitrary degree of accuracy, both the value of the game and an optimal search strategy which is a mixture of at most n search sequences.

It also follows from Theorem 15 that if $\mathbf{p}^*$ is optimal when the set of pure search strategies is $\mathcal{C}$, then $\mathbf{p}^*$ is optimal if the set of pure search strategies is $\widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}^*}$. Below, we show that the reverse implication is also true, creating an optimality test for any hiding strategy requiring only the solution to a finite matrix game.

Proposition 16 Consider $\mathbf{p} \equiv (p_1, \ldots, p_n)$ with $p_i > 0$ for $i = 1, \ldots, n$. Write $G_{\mathcal{D}}$ for the finite game where the searcher's pure strategies are $\mathcal{D} \equiv \widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}} \subset \mathcal{C}$. Let θ be an optimal search strategy in $G_{\mathcal{D}}$. The following three statements are equivalent.

(i) $\mathbf{p}$ is optimal in $G_{\mathcal{D}}$;

(ii) θ is optimal in G;

(iii) $\mathbf{p}$ is optimal in G.

Proof. First, we prove that (i) implies (ii), so suppose $\mathbf{p}$ is optimal in $G_{\mathcal{D}}$. Since θ is optimal in $G_{\mathcal{D}}$, then θ guarantees the searcher an expected time to detection of at most $v^*_{\mathcal{D}}$, the value of $G_{\mathcal{D}}$, no matter which box the hider hides in. Therefore, $u(i, \theta) \leq v^*_{\mathcal{D}}$ for $i = 1, \ldots, n$. By the minimax theorem for finite games, when the hider plays $\mathbf{p}$ and the searcher plays θ, the expected time to detection is $v^*_{\mathcal{D}}$. In other words, we have

$$\sum_{i=1}^{n} p_i u(i, \theta) = v^*_{\mathcal{D}}.$$

Since $p_i > 0$ for $i = 1, \ldots, n$, we must have $u(i, \theta) = v^*_{\mathcal{D}}$ for $i = 1, \ldots, n$. In addition, θ is a mixture of strategies in $\mathcal{D} \subseteq \mathcal{C}^{\mathrm{B}}_{\mathbf{p}}$. By Theorem 12, θ is optimal in G.

Second, we prove that (ii) implies (iii), so suppose θ is optimal in G. By (ii) in Proposition 3, $u(i, \theta) = v^*$ for $i = 1, \ldots, n$. Further, because θ is available to the searcher in $G_{\mathcal{D}}$, it is a mixture of strategies in $\mathcal{D} \subseteq \mathcal{C}^{\mathrm{B}}_{\mathbf{p}}$. By Theorem 12, $\mathbf{p}$ is optimal in G.

Finally, we prove that (iii) implies (i), so suppose $\mathbf{p}$ is optimal in G. By Theorem 15, there exists a search strategy θ^* which is both (a) optimal in G, so guarantees the searcher an expected time to detection of at most v^*, and (b) a mixture of strategies in $\widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}} = \mathcal{D}$, so is available to the searcher in the game $G_{\mathcal{D}}$. Since $\mathbf{p}$ is optimal in G, $\mathbf{p}$ guarantees the hider at least v^* in G; since $\mathcal{D} \subset \mathcal{C}$, $\mathbf{p}$ has the same guarantee for the hider in $G_{\mathcal{D}}$. By (a) and (b) above, the searcher can guarantee at most v^* with θ^* in $G_{\mathcal{D}}$. It follows that $\mathbf{p}$ is optimal in $G_{\mathcal{D}}$, completing the proof. ■

The hiding strategy $\mathbf{p}_0 \equiv (p_{0,1}, \ldots, p_{0,n})$ with

$$p_{0,i} \equiv \frac{t_i/\alpha_i}{\sum_{j=1}^n t_j/\alpha_j}, \quad i = 1, \ldots, n, \tag{22}$$

is of particular interest, since it creates a tie between the Gittins indices of all n boxes in (2) at the start of the search, giving the searcher no preference over which box to search first. Roberts and Gittins (1978) and Gittins and Roberts (1979) both numerically find that $\mathbf{p}_0$ is optimal for the hider in many (but not all) unit-search-time problems. Further, the former proves $\mathbf{p}_0$ is optimal in a two-box problem with $(1 - \alpha_1)^m = (1 - \alpha_2)^{m+1}$ if and only if $m \leq 12$. Because of these earlier findings, the hiding strategy $\mathbf{p}_0$ is a prime candidate to test for optimality via Proposition 16.

5 Conclusion

This paper develops very significantly the existing literature on a search game in discrete boxes where the searcher may overlook a well-concealed hider. There is a theoretical connection between our search game and a related search problem where the hider is replaced by an inanimate object hidden randomly according to a known probability distribution. In the search problem with an inanimate object, the searcher can use a predetermined search sequence to exploit boxes attractive at the beginning of the search. In our search game, however, an intelligent hider will try to make each box equally attractive before the search starts, and the searcher will need to randomize their strategy to avoid behavior predictable by the hider.

Since a pure strategy for the searcher is an indefinite list of boxes to search until the hider is found, the search game is semi-finite and hence difficult to analyze. As a result, most work in the current literature is limited to two boxes or boxes searched in unit time. Using novel proof techniques, we develop a comprehensive theory for the fully-general search game by extending much of the existing work and uncovering new properties along the way.

By making an adjustment to the set of search strategies, we provide a rigorous proof that an optimal search strategy exists, extending a result of Bram (1963). We next develop properties of an optimal search strategy, and, extending a two-box result of Gittins (1989), we show that the searcher can construct an optimal strategy by randomly choosing between some n of $n!$ known, simple search strategies. Based on these properties, we present a novel optimality test for any hiding strategy.

Further work on this search game may involve the combination of the theoretical results in this paper with the algorithm in Lin and Singham (2015) to produce a practical procedure for estimating optimal strategies for each player. When the number of boxes is large, design of heuristic polices may be relevant. Extensions to the search game may involve a network structure rather than discrete boxes. Such an extension is relevant if the geography of the search space prevents the searcher from moving quickly between any pair of hiding locations, for example, a structure of roads. Search games on networks are well studied in the literature, but less so with a chance of overlook.

Acknowledgments

The authors are grateful for the support of the EPSRC funded EP/L015692/1 STOR-i Centre for Doctoral Training, and for many helpful suggestions from Steve Alpern. They also thank an associate editor and two anonymous reviewers for their comments, which have led to a significantly stronger paper.

Appendices

A Proofs of the Continuity of the Function f

A.1 Proof of Lemma 4

Since $\mathbf{p}$ is fixed, to ease notation we drop the second argument from f and ξ and write $f(\mathbf{r})$ and $\xi(\mathbf{r})$ instead. We need to show that for every $\epsilon > 0$ and every $\mathbf{r}_0 \in \mathcal{R}$, there exists $\delta > 0$ such that $d(\mathbf{r}, \mathbf{r}_0) < \delta$ implies that $|f(\mathbf{r}) - f(\mathbf{r}_0)| < \epsilon$.

First, consider the case in which a Gittins search sequence against $\mathbf{p}$ encounters a finite number of ties, and denote this number by k. By the definition of d in (9), any $\mathbf{r}$ with $d(\mathbf{r}, \mathbf{r}_0) < 1/k$ will break the first k ties exactly the same way as $\mathbf{r}_0$, so choosing $\delta = 1/k$ leads to $|f(\mathbf{r}) - f(\mathbf{r}_0)| = 0 < \epsilon$, showing that $f(\mathbf{r})$ is continuous.

Second, consider the case in which a Gittins search sequence against $\mathbf{p}$ encounters infinitely many ties. Because $p_1 > 0$, the Gittins search sequence $\xi(\mathbf{r}_0)$ will visit box 1 infinitely many times. Write a_j for the time point when $\xi(\mathbf{r}_0)$ completes its jth search in box 1. Letting $a_0 \equiv 0$, we can compute the expected time to detection of the strategy pair $(1, \xi(\mathbf{r}_0))$ by

$$u(1, \xi(\mathbf{r}_0)) = \sum_{j=0}^{\infty} (a_{j+1} - a_j)(1 - \alpha_1)^j,$$

where we have used the formula $E[X] = \int_0^\infty P\{X > x\}dx$ for any nonnegative-valued random variable X.

Now consider some Gittins search sequence $\eta \in \mathcal{C}_{\mathbf{p}}^{\mathrm{B}}$ and write b_j for the time point when η completes its jth search in box 1. While b_j need not be the same as a_j, careful examination reveals that their difference is capped by $\sum_{i=1}^{n} t_i$, because both η and $\xi(\mathbf{r}_0)$ are best responses against $\mathbf{p}$, and when encountering a tie in Gittins indices—regardless of the tie-breaking rule adopted—the searcher will visit all boxes involved in the tie in some order before moving on.

Suppose that $b_j = a_j$ for $j = 1, 2, \ldots, k$, but $b_{k+1} \neq a_{k+1}$ for some $k \in \mathbb{N}$. We can bound

$|u(1,\eta)-u(1,\xi(\mathbf{r}_0))|$ as follows:

$$\begin{aligned}|u(1,\eta)-u(1,\xi(\mathbf{r}_0))| &\le \sum_{j=k}^{\infty}|(b_{j+1}-b_j)-(a_{j+1}-a_j)|(1-\alpha_1)^j\\ &\le \sum_{j=k}^{\infty}(|b_{j+1}-a_{j+1}|+|b_j-a_j|)(1-\alpha_1)^j\\ &\le \sum_{j=k}^{\infty}\left(2\sum_{i=1}^{n}t_i\right)(1-\alpha_1)^j\\ &= 2\left(\sum_{i=1}^{n}t_i\right)\frac{(1-\alpha_1)^k}{\alpha_1}.\end{aligned}$$

Define k_1 to be the smallest integer such that the preceding is less than $\epsilon/\sqrt{n}$. Therefore, if $b_j = a_j$ for $j = 1, 2, \dots, k_1$, then $|u(1,\eta)-u(1,\xi(\mathbf{r}_0))| < \epsilon/\sqrt{n}$.

Let $h(m)$ denote the total number of visits to box 1 by $\xi(\mathbf{r}_0)$ from the beginning until all searches involved in the mth index tie are completed, for $m \in \mathbb{N}$. It is clear that $h(m)$ increases weakly in m, and $\lim_{m\to\infty} h(m) = \infty$. Therefore, we can choose some m large enough so that $h(m) \ge k_1$. Write $m_1 \equiv \min\{m : h(m) \ge k_1\}$. In other words, if the first m_1 elements in $\mathbf{r}$ match those in $\mathbf{r}_0$, then

$$|u(1,\xi(\mathbf{r}))-u(1,\xi(\mathbf{r}_0))| < \frac{\epsilon}{\sqrt{n}}.$$

In a similar fashion, for $i = 2, \dots, n$, we can define m_i such that if the first m_i elements in $\mathbf{r}$ match those in $\mathbf{r}_0$, then

$$|u(i,\xi(\mathbf{r}))-u(i,\xi(\mathbf{r}_0))| < \frac{\epsilon}{\sqrt{n}}.$$

Let $m_0 \equiv \max\{m_1, m_2, \dots, m_n\}$. We can conclude that for any $\mathbf{r} \in \mathcal{R}$ whose first m_0 elements match those in $\mathbf{r}_0$, the Euclidean distance

$$\begin{aligned}|f(\mathbf{r})-f(\mathbf{r}_0)| &= |U(\xi(\mathbf{r}))-U(\xi(\mathbf{r}_0))|\\ &= \left(\sum_{i=1}^{n}[u(i,\xi(\mathbf{r}))-u(i,\xi(\mathbf{r}_0))]^2\right)^{1/2}\\ &< \left(n\left(\frac{\epsilon}{\sqrt{n}}\right)^2\right)^{1/2} = \epsilon.\end{aligned}$$

Finally, because any $\mathbf{r} \in \mathcal{R}$ whose first m_0 elements match those in $\mathbf{r}_0$ has $d(\mathbf{r},\mathbf{r}_0) < 1/m_0$, the proof is completed by taking $\delta = 1/m_0$.

A.2 Proof of Lemma 6

Write ξ for the unique Gittins search sequence against $\mathbf{p}$. For $k \in \mathbb{N}$, write $\Delta^n_k \subset \Delta^n$ for the set of mixed hiding strategies $\mathbf{x}$ for which every Gittins search sequence against $\mathbf{x}$ is identical to ξ for the first k searches. Clearly, for any $k \in \mathbb{N}$, we have $\Delta^n_{k+1} \subseteq \Delta^n_k$ and $\mathbf{p} \in \Delta^n_k$.

For any $\delta > 0$, write $B(\mathbf{p}, \delta)$ for the open ball with radius δ centered at $\mathbf{p}$. Since $\mathbf{p}$ is a non-tie point in $\Delta^n(\epsilon)$, for any $k \in \mathbb{N}$, it is possible, in any direction, to move a small-enough (Euclidean) distance in Δ^n away from $\mathbf{p}$ and not disrupt the order of the Gittins indices that generate the first k searches of ξ. Therefore,

$$\delta_k \equiv 0.5 \times \sup\{\delta : B(\mathbf{p}, \delta) \subseteq \Delta^n_k\} > 0, \tag{23}$$

with $\delta_k \geq \delta_{k+1}$, and 0.5 chosen arbitrarily in $(0, 1)$ to ensure that $B(\mathbf{p}, \delta_k) \subseteq \Delta^n_k$ for all $k \in \mathbb{N}$.

Write $\delta^* \equiv \lim_{k\to\infty} \delta_k$. There are two cases. First suppose that $\delta^* > 0$. If $\mathbf{x} \in B(\mathbf{p}, \delta^*)$, then $\mathbf{x} \in \Delta^n_k$ for all $k \in \mathbb{N}$, so any Gittins search sequence against $\mathbf{x}$ is identical to ξ. It follows that ξ, the unique Gittins search sequence against $\mathbf{p}$, is also the unique Gittins search sequence against $\mathbf{x}$. Hence, f is constant on $\mathcal{R} \times B(\mathbf{p}, \delta^*)$, so f is continuous in its second argument at $\mathbf{p}$.

Second, suppose that $\delta^* = 0$. Consider a sequence $\{\mathbf{x}_a\}$ in Δ^n with $\lim_{a\to\infty} \mathbf{x}_a = \mathbf{p}$. To show f is continuous in its second argument at $\mathbf{p}$, we show that

$$\lim_{a\to\infty} f(\mathbf{r}, \mathbf{x}_a) = f(\mathbf{r}, \mathbf{p}) \tag{24}$$

for any fixed $\mathbf{r} \in \mathcal{R}$.

For any $k \in \mathbb{N}$, since $\delta_k > 0$, there must exist a smallest number $g(k) \in \mathbb{N}$ such that every term in the sequence $\{\mathbf{x}_a\}$ after $\mathbf{x}_{g(k)}$ belongs to the ball $B(\mathbf{p}, \delta_k)$. Formally, for any $k \in \mathbb{N}$, write

$$g(k) \equiv \min\{A : \mathbf{x}_a \in B(\mathbf{p}, \delta_k),\ a \geq A\}. \tag{25}$$

Since $\delta_k \geq \delta_{k+1}$, we have $B(\mathbf{p}, \delta_{k+1}) \subseteq B(\mathbf{p}, \delta_k)$ and hence $g(k) \leq g(k+1)$, so the sequence $\{g(k) : k \in \mathbb{N}\}$ increases weakly.

Consider the sequence $\{\mathbf{x}_{g(k)} : k \in \mathbb{N}\}$. We have $\lim_{a\to\infty} \mathbf{x}_a = \mathbf{p}$ by assumption; our next aim is to show that $\lim_{k\to\infty} \mathbf{x}_{g(k)} = \mathbf{p}$ also. To do this, we show that, for any $\epsilon > 0$, we can choose K such that $\mathbf{x}_{g(k)} \in B(\mathbf{p}, \epsilon)$ for all $k \geq K$. Choose $\epsilon > 0$. Since $\lim_{k\to\infty} \delta_k = 0$, there exists K such that $\delta_K < \epsilon$. By the definition of g in (25), we have $\mathbf{x}_a \in B(\mathbf{p}, \delta_K)$ for all $a \geq g(K)$. Since g is increasing, we have $\mathbf{x}_{g(k)} \in B(\mathbf{p}, \delta_K) \subset B(\mathbf{p}, \epsilon)$ for all $k \geq K$, showing that $\{\mathbf{x}_{g(k)}\}$ has limit $\mathbf{p}$.

By the definitions in (23) and (25), we have $\mathbf{x}_{g(k)} \in B(\mathbf{p}, \delta_k) \subseteq \Delta_k^n$. Recall ξ as the unique Gittins search sequence against $\mathbf{p}$, and, for $b \in \mathbb{N}$, write ξ_b for an arbitrary Gittins search sequence against $\mathbf{x}_b$. Since $\mathbf{x}_{g(k)} \in \Delta_k^n$, as k increases, the first time when $\xi_{g(k)}$ and ξ may differ becomes increasingly later and later into the search. Hence, no matter where the hider is hidden, the effect on the expected time to detection of this difference decreases to 0; in other words, for $i = 1, \ldots, n$, $u(i, \xi_{g(k)}) \to u(i, \xi)$ as $k \to \infty$, $i = 1, \ldots, n$, so

$$\lim_{k\to\infty} f(\mathbf{r}, \mathbf{x}_{g(k)}) = f(\mathbf{r}, \mathbf{p})$$

for any $\mathbf{r} \in \mathcal{R}$. Since $\lim_{k\to\infty} \mathbf{x}_{g(k)} = \lim_{a\to\infty} \mathbf{x}_a = \mathbf{p}$, then (24) follows, completing the proof.

A.3 Proof of Lemma 7

First, note that if $\Delta^n(\mathbf{p}, \Sigma) = \{\mathbf{p}\}$, then any sequence in $\Delta^n(\mathbf{p}, \Sigma)$ is constant, and the result is trivially true. The rest of the argument, which is similar to the proof of Lemma 6, deals with the case where $\Delta^n(\mathbf{p}, \Sigma)$ contains elements in addition to $\mathbf{p}$. Let $\mathbf{r} \in \mathcal{R}$ contain only elements from $\Sigma \subset S_n$. For $k \in \mathbb{N}$, let $\Delta_{k,\mathbf{r}}^n \subset \Delta^n$ contain precisely those mixed hiding strategies $\mathbf{x}$ for which the Gittins search sequence against $\mathbf{x}$ under rule $\mathbf{r}$ is identical to $\xi(\mathbf{r}, \mathbf{p})$ (the Gittins search sequence against $\mathbf{p}$ under rule $\mathbf{r}$) for the first k searches. Clearly, for any $k \in \mathbb{N}$, we have $\Delta_{k+1,\mathbf{r}}^n \subseteq \Delta_{k,\mathbf{r}}^n$ and $\mathbf{p} \in \Delta_{k,\mathbf{r}}^n$.

For any $\delta > 0$, write $B(\mathbf{p}, \delta)$ for the open ball with radius δ centered at $\mathbf{p}$. Note that any two points in Δ^n must be within Euclidean distance $\sqrt{2}$ of eachother. Therefore, for any $\mathbf{p} \in \Delta^n$, we must have $B(\mathbf{p}, \sqrt{2}) = \Delta^n$. For $\delta \in [0, \sqrt{2}]$, write $\Delta^n(\mathbf{p}, \delta, \Sigma) \equiv B(\mathbf{p}, \delta) \cap \Delta^n(\mathbf{p}, \Sigma)$. In other words, $\Delta^n(\mathbf{p}, \delta, \Sigma)$ is the subset of mixed hiding strategies in $\Delta^n(\mathbf{p}, \Sigma)$ strictly less than (Euclidean) distance δ from $\mathbf{p}$.

Write

$$\delta_{k,\mathbf{r}} \equiv 0.5 \times \sup\left\{\delta : \Delta^n(\mathbf{p}, \delta, \Sigma) \subseteq \Delta_{k,\mathbf{r}}^n\right\}, \tag{26}$$

with 0.5 arbitrarily chosen in $(0, 1)$ to ensure that $\Delta^n(\mathbf{p}, \delta_{k,\mathbf{r}}, \Sigma) \subseteq \Delta_{k,\mathbf{r}}^n$ for all $k \in \mathbb{N}$. Note that $\delta_{k,\mathbf{r}} \geq 0$ for all $k \in \mathbb{N}$ since $\Delta^n(\mathbf{p}, 0, \Sigma) = \{\mathbf{p}\} \subset \Delta_{k,\mathbf{r}}^n$. The aim of the following is to show that $\delta_{k,\mathbf{r}} > 0$ for all $k \in \mathbb{N}$.

Let $k \in \mathbb{N}$. We examine two cases. First, suppose that, in the first k searches of $\xi(\mathbf{r}, \mathbf{p})$, no ties are encountered. Then, it is possible, in any direction, to move a small enough (Euclidean) distance in Δ^n away from $\mathbf{p}$ and not disrupt the order of the Gittins indices in (2) that generate the first k searches of $\xi(\mathbf{r},\mathbf{p})$. Therefore, we may choose $\delta > 0$ such that $B(\mathbf{p},\delta) \subset \Delta^n_{k,\mathbf{r}}$. It follows that $\Delta^n(\mathbf{p},\delta,\Sigma) \subset \Delta^n_{k,\mathbf{r}}$, and hence that $\delta_{k,\mathbf{r}} \geq \delta/2 > 0$.

Second, suppose that, in the first k searches of $\xi(\mathbf{r},\mathbf{p})$, we do encounter ties between boxes. Suppose such a tie involves b boxes. By (2), whilst the order of the next b boxes searched may depend on the tie-breaking rule, the set of b boxes searched will not. Therefore, after the tie has been broken, the Gittins indices in (2) will be the same no matter how the tie was broken. Hence, it is possible, in any direction, to move a small enough (Euclidean) distance away in Δ^n from $\mathbf{p}$ and not disrupt the order of the Gittins indices that generate the first k searches of $\xi(\mathbf{r},\mathbf{p})$ at any point where there is not a tie between boxes. It follows that we may choose $\delta > 0$ such that, for any $\mathbf{x} \in B(\mathbf{p},\delta)$, any Gittins search sequence against $\mathbf{x}$ differs only in the first k searches to $\xi(\mathbf{r},\mathbf{p})$ for those searches where $\xi(\mathbf{r},\mathbf{p})$ is in the process of breaking a tie. Now suppose additionally that $\mathbf{x} \in \Delta^n(\mathbf{p},\Sigma)$, so $\mathbf{x} \in \Delta^n(\mathbf{p},\delta,\Sigma)$. Since $\mathbf{x} \in \Delta^n(\mathbf{p},\Sigma)$, when a tie is reached by $\xi(\mathbf{r},\mathbf{p})$, if we instead were following a Gittins search sequence against $\mathbf{x}$, the Gittins indices of any boxes involved in the tie will either still be tied, or lie in the ordering determined by σ for all $\sigma \in \Sigma$. Therefore, since $\mathbf{r}$ contains only elements of Σ, the Gittins search sequence against $\mathbf{x}$ that breaks ties using $\mathbf{r}$ will break the tie using the same preference ordering as $\xi(\mathbf{r},\mathbf{p})$, so will be identical to $\xi(\mathbf{r},\mathbf{p})$ for the first k searches. In other words, $\Delta^n(\mathbf{p},\delta,\Sigma) \subset \Delta^n_{k,\mathbf{r}}$, and hence $\delta_{k,\mathbf{r}} \geq \delta/2 > 0$.

Now we have shown $\delta_{k,\mathbf{r}} > 0$ for all $k \in \mathbb{N}$, we are in a position to finish the proof in a similar style to Lemma 6. Write $\delta^*_{\mathbf{r}} \equiv \lim_{k\to\infty} \delta_{k,\mathbf{r}}$, and let $\{\mathbf{x}_a : a \in \mathbb{N}\}$ be a sequence in $\Delta^n(\mathbf{p},\Sigma)$ with $\lim_{a\to\infty} \mathbf{x}_a = \mathbf{p}$. There are two cases.

First suppose that $\delta^*_{\mathbf{r}} > 0$. Let $\mathbf{x} \in \Delta^n(\mathbf{p},\delta^*_{\mathbf{r}},\Sigma)$; then $\mathbf{x} \in \Delta^n_{k,\mathbf{r}}$ for all $k \in \mathbb{N}$, so the Gittins search sequence against $\mathbf{x}$ which breaks ties using rule $\mathbf{r}$ is identical to $\xi(\mathbf{r},\mathbf{p})$. It follows that f is constant on $\Delta^n(\mathbf{p},\delta^*_{\mathbf{r}},\Sigma)$ when its first argument is fixed at $\mathbf{r}$. Furthermore, since $\lim_{a\to\infty} \mathbf{x}_a = \mathbf{p}$, there must exist A such that $\mathbf{x}_a \in B(\mathbf{p},\delta^*_{\mathbf{r}})$ for all $a \geq A$. Yet, since $\{\mathbf{x}_a\}$ is a sequence in $\Delta^n(\mathbf{p},\Sigma)$, we also, for all $a \geq A$, have $\mathbf{x}_a \in \Delta^n(\mathbf{p},\delta^*_{\mathbf{r}},\Sigma)$ and hence $f(\mathbf{r},\mathbf{x}_a) = f(\mathbf{r},\mathbf{p})$ for any $\mathbf{r}$ with all elements in Σ; proving the result for the first case.

Second, suppose that $\delta^*_{\mathbf{r}} = 0$. As in the proof of Lemma 6, since $\{\mathbf{x}_a\}$ has limit $\mathbf{p}$ and, for any $k \in \mathbb{N}$, $\delta_{k,\mathbf{r}} > 0$, there must exist a smallest number $g_{\mathbf{r}}(k) \in \mathbb{N}$ such that every term in $\{\mathbf{x}_a\}$ after $\mathbf{x}_{g(k)}$ belongs to the ball $B(\mathbf{p},\delta_{k,\mathbf{r}})$. Further, since $\{\mathbf{x}_a\}$ is a sequence in $\Delta^n(\mathbf{p},\Sigma)$, every term in

$\{\mathbf{x}_a\}$ after $\mathbf{x}_{g(k)}$ also belongs to $\Delta^n(\mathbf{p}, \delta_{k,\mathbf{r}}, \Sigma)$. Formally, for any $k \in \mathbb{N}$, we write

$$g_{\mathbf{r}}(k) \equiv \min\{A : \mathbf{x}_a \in \Delta^n(\mathbf{p}, \delta_{k,\mathbf{r}}, \Sigma),\ a \geq A\}. \tag{27}$$

Note from (26) that since $\Delta^n_{k+1,\mathbf{r}} \subseteq \Delta^n_{k,\mathbf{r}}$, we have $\delta_{k,\mathbf{r}} \geq \delta_{k+1,\mathbf{r}}$. It follows that $B(\mathbf{p}, \delta_{k+1,\mathbf{r}}) \subseteq B(\mathbf{p}, \delta_{k,\mathbf{r}})$, and hence $g_{\mathbf{r}}(k) \leq g_{\mathbf{r}}(k+1)$ for all $k \in \mathbb{N}$, so the sequence $\{g_{\mathbf{r}}(k) : k \in \mathbb{N}\}$ increases weakly. An identical argument to that in the proof of Lemma 6 for $\{\mathbf{x}_{g(k)}\}$ can be applied to $\{\mathbf{x}_{g_{\mathbf{r}}(k)}\}$ to show that $\lim_{k\to\infty} \mathbf{x}_{g_{\mathbf{r}}(k)} = \mathbf{p}$.

By the definitions in (26) and (27), we have $\mathbf{x}_{g_{\mathbf{r}}(k)} \in \Delta^n(\mathbf{p}, \delta_{k,\mathbf{r}}, \Sigma) \subseteq \Delta^n_{k,\mathbf{r}}$. Recall $\xi(\mathbf{r}, \mathbf{x})$ is the Gittins search sequence against $\mathbf{x}$ which breaks ties using rule $\mathbf{r}$. Since $g_{\mathbf{r}}$ is increasing and $\mathbf{x}_{g_{\mathbf{r}}(k)} \in \Delta^n_{k,\mathbf{r}}$, as k increases, the first time when $\xi(\mathbf{r}, \mathbf{x}_{g_{\mathbf{r}}(k)})$ and $\xi(\mathbf{r}, \mathbf{p})$ differ becomes increasingly later and later into the search. Hence, no matter where the hider is hidden, the effect on the expected time to detection of this difference decreases to 0. Therefore, $u(i, \xi(\mathbf{r}, \mathbf{x}_{g_{\mathbf{r}}(k)})) \to u(i, \xi(\mathbf{r}, \mathbf{p}))$ as $k \to \infty$, $i = 1, \ldots, n$. Combined with $\lim_{k\to\infty} \mathbf{x}_{g_{\mathbf{r}}(k)} = \lim_{a\to\infty} \mathbf{x}_a = \mathbf{p}$, we have

$$\lim_{k\to\infty} f(\mathbf{r}, \mathbf{x}_{g_{\mathbf{r}}(k)}) = \lim_{a\to\infty} f(\mathbf{r}, \mathbf{x}_a) = f(\mathbf{r}, \mathbf{p}),$$

for any $\mathbf{r}$ with all elements in Σ, proving the result for the second case.

B Gittins' Proposition 8.5 Extension

For any (pure or mixed) search strategy θ and $i, j \in \{1, \ldots, n\}$, write $D_{i,j}(\theta) \equiv u(i, \theta) - u(j, \theta)$. If $D_{i,j}(\theta) = 0$ for some $i, j \in \{1, \ldots, n\}$, we say θ *equalizes* boxes i and j. Let θ_1 and θ_2 be search strategies, and choose $i, j \in \{1, \ldots, n\}$ such that $D_{i,j}(\theta_1) \leq 0$. Then, by direct computation, it is easy to show that there exists a mixture of θ_1 and θ_2 that equalizes boxes i and j if any only if $D_{i,j}(\theta_2) \geq 0$.

Let $\mathbf{p}^*$ be an arbitrary optimal hiding strategy. Recall from the notation of Section 4 that $\xi_{\sigma,\mathbf{p}^*}$ is the Gittins search sequence against $\mathbf{p}^*$ that breaks every tie encountered using σ; write $\xi_\sigma \equiv \xi_{\sigma,\mathbf{p}^*}$. When $n = 2$, we have $\widehat{\mathcal{C}}^{\mathrm{B}}_{\mathbf{p}^*} = \{\xi_{12}, \xi_{21}\}$, where ξ_{12} (resp. ξ_{21}) breaks any tie in favor of box 1 (resp. box 2). The proof of Gittins' Lemma 8.4 shows that $D_{1,2}(\xi_{12}) \leq 0$ and $D_{1,2}(\xi_{21}) \geq 0$. In other words, there exists a mixture θ^* of ξ_{12} and ξ_{21} which equalizes boxes 1 and 2; by Theorem 12, θ^* is optimal for the searcher.

By extending Gittins' method to an n-box problem, for any $i, j \in \{1, \ldots, n\}$, we can show that $D_{i,j}(\xi_{i\cdots j}) \leq 0$ and $D_{i,j}(\xi_{j\cdots i}) \geq 0$, where $x \cdots y$ is any permutation of $\{1, \ldots, n\}$ with first element x and last element y. If either of these two $D_{i,j}$ terms is equal to 0, then the corresponding search sequence equalizes i and j. Otherwise, since, for any $\xi \in \mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$, $D_{i,j}(\xi)$ must lie one side of 0, there exists a mixture of ξ and either $\xi_{i\cdots j}$ or $\xi_{j\cdots i}$ that equalizes boxes i and j. Therefore, many mixtures of pairs of sequences in $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$ can be constructed that equalize boxes i and j.

To obtain an optimal search strategy using Theorem 12, we need a mixture of elements of $\mathcal{C}^{\mathrm{B}}_{\mathbf{p}^*}$ that equalizes all n boxes. Yet, problems occur when a third box, say k, is introduced. Suppose θ_1 and θ_2 both equalize boxes i and j; then, any mixture of θ_1 and θ_2 that equalizes boxes i and k (or j and k) will equalize boxes i, j and k. However, there is no guarantee that $D_{i,k}(\theta_1)$ and $D_{i,k}(\theta_2)$ will have opposing signs, so there is no guarantee that such a mixture exists. Whilst we managed to prove that such θ_1 and θ_2 with $D_{i,k}(\theta_1) \leq 0 \leq D_{i,k}(\theta_2)$ exist when $n = 3$ (thus finding an optimal search strategy for the three-box case), the proof cannot be generalized to $n \geq 4$.

References

Alpern, S., Fokkink, R., Gasieniec, L., Lindelauf, R., and Subrahmanian, V. (2013). *Search Theory: A Game Theoretic Perspective*. Springer Publishing Company, Incorporated.

Alpern, S. and Gal, S. (2003). *The Theory of Search Games and Rendezvous*. International Series in Operations Research and Managment Science. Kluwer Academic Publishers, Boston, Dordrecht, London.

Black, W. L. (1965). Discrete sequential search. *Information and Control*, 8(2):159–162.

Blackwell, D. and Girshick, M. A. (1954). *Theory of Games and Statistical Decisions*. John Wiley & Sons, New York.

Bram, J. (1963). A 2-player n-region search game. IRM-31, Operations Evaluation Group, Center for Naval Analysis.

Charalambos, D. and Aliprantis, B. (2013). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag Berlin and Heidelberg GmbH & Company KG.

Chew, Jr, M. C. (1973). Optimal stopping in a discrete search problem. *Operations Research*, 21(3):741–747.

Clarkson, J., Glazebrook, K. D., and Lin, K. Y. (2020). Fast or slow: Search in discrete locations with two search modes. *Operations Research*, 68(2):552–571.

Ferguson, T. S. (2020). *A Course in Game Theory*. World Scientific.

Garrec, T. and Scarsini, M. (2020). Search for an immobile hider on a stochastic network. *European Journal of Operational Research*, 283(2):783–794.

Gittins, J. and Roberts, D. (1979). The search for an intelligent evader concealed in one of an arbitrary number of regions. *Naval Research Logistics Quarterly*, 26(4):651–666.

Gittins, J., Weber, R., and Glazebrook, K. D. (2011). *Multi-Armed Bandit Allocation Indices*. John Wiley and Sons, Ltd.

Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 41(2):148–177.

Gittins, J. C. (1989). *Multi-armed bandit allocation indices*. Wiley.

Hellerstein, L., Lidbetter, T., and Pirutinsky, D. (2019). Solving zero-sum games using best-response oracles with applications to search games. *Operations Research*, 67(3):731–743.

Kadane, J. B. (1971). Optimal whereabouts search. *Operations Research*, 19(4):894–904.

Kress, M., Lin, K. Y., and Szechtman, R. (2008). Optimal discrete search with imperfect specificity. *Mathematical Methods Of Operations Research*, 68(3):539–549.

Lidbetter, T. (2013). Search games with multiple hidden objects. *SIAM Journal on Control and Optimization*, 51(4):3056–3074.

Lidbetter, T. and Lin, K. Y. (2019). Searching for multiple objects in multiple locations. *European Journal of Operational Research*, 278(2):709–720.

Lin, K. Y. and Singham, D. I. (2015). Robust search policies against an intelligent evader. Technical Report NPS-OR-15-009, Monterey, California. Naval Postgraduate School.

Lin, K. Y. and Singham, D. I. (2016). Finding a hider by an unknown deadline. *Oper. Res. Lett.*, 44(1):25–32.

Matula, D. (1964). A periodic optimal search. *The American Mathematical Monthly*, 71(1):15–21.

Norris, R. C. (1962). Studies in search for a conscious evader. *MIT Lincoln Laboratory Technical Report Number 279*.

Roberts, D. M. and Gittins, J. C. (1978). The search for an intelligent evader: Strategies for searcher and evader in the two-region problem. *Naval Research Logistics Quarterly*, 25(1):95–106.

Ross, S. M. (1969). A Problem in Optimal Search and Stop. *Operations Research*, 17(6):984–992.

Ross, S. M. (1983). *Introduction to Stochastic Dynamic Programming: Probability and Mathematical*. Academic Press, Inc., USA.

Ruckle, W. H. (1991). A discrete search game. In *Stochastic Games And Related Topics*, pages 29–43. Springer.

Stone, L. D. (2004). *Theory of Optimal Search*. INFORMS, 2nd edition.

Stone, L. D., Royset, J. O., and Washburn, A. R. (2016). *Optimal Search for Moving Targets*. International Series in Operations Research and Management Science. Springer, New York, United States.

Strichartz, R. S. (1995). *The Way of Analysis*. Jones and Bartlett.

Subelman, E. J. (1981). A hide-search game. *Journal of Applied Probability*, 18:628–640.

Washburn, A. R. (2002). *Search and Detection*. INFORMS, 4th edition.

Wegener, I. (1980). The Discrete Sequential Search Problem with Nonrandom Cost and Overlook Probabilities. *Mathematics Of Operations Research*, 5(3):373–380.

Wilansky, A. (2008). *Topology for Analysis*. Dover Books on Mathematics. Dover Publications.