

Steve Alpern, [Thomas Lidbetter](#)

Optimal trade-off between speed and acuity when searching for a small object

Article (Accepted version)
(Refereed)

Original citation:

Alpern, Steven and Lidbetter, Thomas (2015) *Optimal trade-off between speed and acuity when searching for a small object*. [Operations Research](#), 63 (1). pp. 122-133. ISSN 0030-364X

DOI: [10.1287/opre.2014.1331](https://doi.org/10.1287/opre.2014.1331)

© 2015 [Informs](#)

This version available at: <http://eprints.lse.ac.uk/61504/>

Available in LSE Research Online: November 2015

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Abstract

A Searcher seeks to find a stationary Hider located at some point H (not necessarily a node) on a given network Q . The Searcher can move along the network from a given starting point at unit speed, but to actually *find* the Hider she must pass it *while moving at a fixed slower speed* (which may depend on the arc). In this ‘bimodal search game’, the payoff is the first time the Searcher passes the Hider while moving at her slow speed. This game models the search for a small or well hidden object (e.g., a contact lens, improvised explosive device, predator search for camouflaged prey). We define a Bimodal Chinese Postman tour as a tour of minimum time δ which traverses every point of every arc at least once in the slow mode. For trees and *weakly Eulerian* networks (networks containing a number of disjoint Eulerian cycles connected in a tree-like fashion) the value of the bimodal search game is $\delta/2$. For trees, the optimal Hider strategy has full support on the network. This differs from traditional search games, where it is optimal for him to hide only at leaf nodes. We then consider the notion of a *lucky Searcher* who can also detect the Hider with a positive probability q even when passing him at her fast speed. This paper has particular importance for demining problems.

Subject classifications: teams: games/group decisions; search/surveillance; tree algorithms: networks/graphs.

Area of review: Military and Homeland Security

Optimal trade-off between speed and acuity when searching for a small object

Steve Alpern* and Thomas Lidbetter**

Warwick Business School
University of Warwick
Coventry CV4 7AL
United Kingdom
Steve.Alpern@wbs.ac.uk

Department of Management
London School of Economics
London WC2A 2AE
United Kingdom
T.R.Lidbetter@lse.ac.uk

1 Introduction

In traditional models of search games, a Searcher moves around a search space with the aim of *reaching* (and hence finding) a stationary Hider in minimal time. In this paper we make the observation that a Searcher's ability to detect a Hider or hidden object may depend on the speed at which she is traveling, so if she is moving too fast she may pass the Hider without detecting him. For example, when searching for a contact lens it is necessary to move slowly on one's hands and knees, but there is also the option of standing up and walking briskly to a different location. To take a more serious example, there has been a lot of recent interest (Johnson and Ali, 2012) in the detection of landmines and improvised explosive devices (IEDs). An explosives expert searching for an IED may be able to move from place to place quickly in a vehicle, but in order to detect the IEDs she must get out of her vehicle and move at a slower pace. Similarly, specially trained mine detection dogs (MDDs) are used extensively in the detection of mines, and it has been observed that dogs perform this role more effectively when moving at a slower speed, as explained in the following extract from a report for the Geneva International Centre for Humanitarian Demining (McLean, 2001):

For many medium to large sized dogs, the natural walking pace of a human (5-6 km/h) is an uncomfortable speed, leading to pulling or dragging on a lead because the dog wants to walk (3-4 km/h) or lope (8-10 km/h). The loping run can be maintained by dogs for hours, and it is often the gait used by MDDs while searching for mines. Unfortunately, while the loping gait allows the dog to cover a great deal of ground and work for long periods, it is probably moving too fast to conduct an effective search over the entire surface of the ground. To search effectively, the dog must move slowly by reducing its gait to a walk, or even a slow walk. Dogs that do so naturally are likely to be better MDDs.

Another example, taken from biological science is that of a bird of prey searching for prey: in this case the bird can fly quickly to an area of high grass, and then land on the ground and walk at a slower speed when it wants to see the hidden prey. Indeed, experimental studies have shown that in nature high speed can have a detrimental effect on visual acuity (for example, see Gendron 1986).

In this paper, we model the search space Q as a rooted network consisting of a finite set of arcs of given lengths, with a distinguished node O called the *root*. We denote the length (Lebesgue measure) of an arc a by $\lambda(a)$ and the total length of the network by λ . In the traditional model of search, introduced by Isaacs (1969) and studied extensively by Gal (1979, 2000) and others, the Searcher moves around the network at unit speed starting from O until she meets a stationary Hider located somewhere on the network. Here we suppose that on each arc the Searcher can travel at either a slow speed, which may depend on which arc she is on, or a fast speed (which we normalize to 1). It will be convenient to specify the Searcher's slow and fast speeds on an arc a in terms of the *travel times* $\alpha = \alpha_a$ and $\beta = \beta_a$ that it takes her to travel the length of a at her slow and fast speeds, respectively. The fast time β is thus equal to the length $\lambda(a)$ of a .

In the *bimodal search game*, which we denote by $B(Q, O)$ for a network Q with root O , a Hider strategy is simply a point on the network. A Searcher strategy is a path on Q starting from O on which the Searcher can move at either her fast (unit) or slow speed, and which covers every point of Q at least once in the slow mode. In general, we will allow the Searcher to change her speed while in the interior of an arc, though only a finite number of times. We will also assume that the Searcher searches any given arc all at once, so that she never travels part of the way along an arc and then turns back. We follow the convention originally introduced by Gal (1979) of denoting Searcher pure strategies by upper case letters (S) and mixed strategies by lower case (s). Hider pure strategies will usually be given by a point x on the network, and mixed strategies by some probability density $h(x)$. For given Searcher and Hider pure strategies S and x , we define the *search time*, denoted by $T(S, x)$ as the first time that the Searcher passes the Hider while moving at her slow speed. The Searcher wishes to minimize, and the Hider to maximize, the search time. For mixed strategies s and h we write $T(s, h)$ for the expected search time.

Bimodal search games provide a nice continuous link between the classical continuous path search games studied by Gal (1979) and discontinuous path games, called *generalized search* in Alpern and Gal (2003). Suppose the slow speed and fast speed are constant on the network, and $\sigma \leq 1$ denotes the ratio of the slow speed to the fast speed. If σ is very close to 1, that is $\beta \simeq \alpha$ on all arcs, then there is no need for the Searcher to move at her fast speed, so in this case bimodal search approaches Gal's original game. On the other hand, if we normalize the slow speed to 1 and let the fast speed $1/\sigma$

go to infinity (or equivalently let σ go to 0), this is as if the Searcher can instantaneously (or almost instantly) jump to any point of the network at any time, and resume slow search at that point. The latter case approximates *generalized search*, in which the value of the game is half the time required to tour the network at the slow speed. Thus bimodal search, with a speed ratio σ going between 0 and 1, links the known cases of generalized search and classical search by a continuum of search problems.

The paper is organized as follows. Section 2 analyzes some strategies for the players that are available on any network. In Section 3 of the paper we restrict our attention to the bimodal search game played on a single arc I with the root at one end. In Gal's original search game (what we call the classical game), the solution for this case is trivial: the Hider hides at the opposite end of the arc from the root and the Searcher just traverses the arc. However in the bimodal search game the solution is far from trivial, and in particular the optimal strategy for the Hider is a continuous probability distribution (related to the beta distribution) whose support is the whole arc. (If the Hider always hid at the far end, the Searcher would just go there at the fast speed, and return at the slow speed). An optimal strategy of this complexity for the Hider hasn't been seen in the literature before for such a simple network as a single arc.

In Section 4 we give the solution of the bimodal search game on a tree, and explain how this solution can be extended to *weakly Eulerian* networks, networks that consist of a number of disjoint Eulerian cycles which, when contracted to a single point, leave a tree. For trees, the classical games of Gal (1979) have finitely many strategies for each player, so Von Neumann's minimax result suffices in that context. However, for bimodal games the strategy spaces (even for trees) are no longer finite, so a new minimax result is needed, which we present in Section 5. In Section 6 we consider how the game is related to a search game proposed by Kikuta and Ruckle (1994) in which the Searcher moves in a network at constant speed but must pay a *search cost* at each node to check whether or not the Hider is there. Finally, in Section 7 we study a variant on the game in which the Searcher has a positive probability of detecting the Hider even when traveling at his fast speed.

2 Simple Strategies for General Networks

In this section we consider strategies for the bimodal search game that can be used on any rooted network and evaluate the expected search times that they guarantee. These generalize similar strategies introduced by Gal (1979) for traditional (single-speed) search games, and the proofs are very similar. The material that really distinguishes bimodal search will not be seen until the next section.

The Searcher can always adopt the following strategy.

Definition 1 *A Bimodal Chinese Postman (BCP) tour $P(t)$ is a closed path $P(t)$ on Q which traverses every point at least once in the slow mode and takes minimum time $\delta = \delta(Q)$ (called the **bimodal tour time** of Q). A random mixture of an BCP tour and its time-reversed tour $P(\delta - t)$ is called a **Random Bimodal Chinese Postman (Random BCP) tour**.*

As an example, consider a network with two nodes, O (start) and A , that are connected by three arcs which are identical in that they have the same slow and fast traversal times α and β . A BCP tour must travel slowly for time at least 3α and must also traverse one of the arcs fast in time β to construct a tour. As such a tour is clearly possible, we have $\delta = 3\alpha + \beta$. If we write $\Sigma\alpha = \sum_{\text{arcs } a} \alpha_a$ for the sum of the slow travel times of the arcs of a network Q and $\Sigma\beta = \sum_{\text{arcs } a} \beta_a$ for the corresponding sum at the fast times, then a tree clearly has $\delta = \Sigma\alpha + \Sigma\beta$ as each arc must be traversed once fast and once slowly. An Eulerian network clearly has $\delta = \Sigma\alpha$ since a BCP tour consists of an Eulerian tour taken at the slow speed.

By using a Random BCP tour in the bimodal search game $B(Q, O)$ the Searcher can find every point of the network Q in expected time at most $\delta/2$. To see this note that if the BCP tour finds a point x in time t , the reversed tour finds it no later than $\delta - t$, so on average in time not exceeding $(t + (\delta - t)) / 2 = \delta/2$. In summary:

Lemma 2 *A Random Bimodal Chinese Postman tour finds any point in a rooted network Q, O in expected time not exceeding $\delta/2$.*

If the value of the game is $\delta/2$, we say that Q is *simply searchable*.

We now turn to Hider strategies. A strategy always available to the Hider on any network is the *slow-uniform strategy*, which hides in any interval of any arc with a probability proportional to its slow traversal time. By using

this strategy, the same argument as for traditional (single speed) games, Lemma 1 of (Gal 1980), shows that she can ensure the expected search time is at least $\Sigma\alpha/2$. (This is because Hider can ensure that the probability he is found by time t does not exceed $t/\Sigma\alpha$, which implies the expected capture time is at least $\Sigma\alpha/2$.)

Lemma 3 *If the Hider adopts the slow-uniform distribution, then for any bimodal Searcher pure strategy, the expected capture time is at least $\Sigma\alpha/2$.*

Clearly if $\Sigma\alpha$ is equal to δ then we can immediately deduce from Lemmas 2 and 3 that the value of the game is $\delta/2$. If the network Q is Eulerian, so that there is a closed path that visits each arc exactly once, then clearly $\delta = \Sigma\alpha$. In this case a BCP tour is simply a Eulerian network traversed at the slow speed. Hence we have:

Theorem 4 *If Q is an Eulerian network then the value of the game $B(Q, O)$ for any starting point O is $V = \Sigma\alpha/2 = \delta/2$. A Random BCP tour is optimal for the Searcher and the slow-uniform strategy is optimal for the Hider.*

While the results shown in this section are obvious extensions of the original work of Gal (1979), we shall see in the next section that even for a network consisting of a single arc, the bimodal search game can be very different from the traditional single-speed search game.

3 Searching a single arc

We begin by defining and analyzing the bimodal search game played on a single arc network $Q = I = [0, 1]$, with length (and fast travel time) normalized to $\beta = 1$, root at $O = 0$ and the other end point 1 labeled as P . The slow traversal time is given by $\alpha > 1$, with corresponding slow search speed $\sigma = 1/\alpha$. At the end, we undo the normalization and rewrite our results for an interval $[0, \beta]$ of arbitrary fast traversal time (length) β and slow traversal time α .

A single arc network is a very special case of a tree, and for tree networks Q we will assume that the Searcher follows a *depth-first path*. This means that whenever she is at any point x of Q , she searches every point of the subtree starting at x (at least once in slow mode), before searching anywhere else. In the classical model (Gal, 1979) and in more recent search games

models (Alpern, 2010 and Alpern and Lidbetter, 2013), it is always optimal for the Searcher to use depth-first searches. In the present case of a single arc, this means that the Searcher simply goes monotonically from 0 to 1 and then monotonically back again. So the only flexibility in the searcher pure strategy is contained in the subset C of Q on which she searches (travels at slow speed) on her first arrival at a point. That is, points x in C are found by the Searcher on her first arrival at x , points in the complement C' of C are found at the second arrival of the Searcher. We will require throughout that the Searcher can alternate between slow and fast speeds only a finite number of times, that is, C is a finite union of closed intervals. (To obtain the existence of a value in Section 5, we will also put an arbitrarily small but positive lower bound on the length of the intervals comprising C , but this is a purely technical assumption we will ignore for the present.)

More generally, for a measurable set C , define the path S_C which *first* reaches the point $x \in I$ at the *first arrival time* $f(x)$ defined by

$$f(x) = f(x, C) = \alpha \cdot \lambda(C \cap [0, x]) + 1 \cdot (x - \lambda(C \cap [0, x])), \quad (1)$$

where $f'(x)$ equals α on C and equals 1 on C' .

The path $S_C(t)$ reaches x for the second time (on the return journey) after traversing the interval $[x, 1]$ once fast and once slowly, so in time $f(x) + (1 + \alpha)(1 - x)$. So

$$T(S_C, x) = \begin{cases} f(x), & \text{if } x \in C \\ f(x) + (1 + \alpha)(1 - x), & \text{if } x \in C'. \end{cases} \quad (2)$$

The interpretation is that C is searched slowly on the way to $P = 1$ and its complement C' is searched slowly on the way back.

In the traditional search game on an interval with endpoint start, all Hider locations are dominated by hiding at the opposite end to the start, so the Hider in this case has a pure strategy optimum. In the bimodal search game, even on the simple interval network, it turns out that *the Hider must use a full support distribution on the interval!*

Theorem 5 *A rooted tree Q, O consisting of a single interval (arc) $[0, 1]$ with root $O = 0$ is simply searchable. That is, if the slow speed is σ , so slow traversal time is $\alpha = 1/\sigma$, the value of the bimodal search game on it is given by $V = \delta/2 = (1 + \alpha)/2 = (1 + 1/\sigma)/2$. The unique optimal Hider distribution*

has density h^* and cumulative distribution H^* given by

$$H^*(x) = 1 - (1-x)^\gamma \quad \text{and} \quad h^*(x) = \gamma(1-x)^{\gamma-1}, \quad (3)$$

where $\gamma = \frac{\alpha-1}{\alpha+1}$.

More generally, if the slow and fast traversal times are α and β , $V = \delta/2 = (\alpha + \beta)/2$ and

$$H^*(x) = 1 - \left(1 - \frac{x}{\beta}\right)^\gamma, \quad h^*(x) = \gamma \left(1 - \frac{x}{\beta}\right)^{\gamma-1}, \quad \gamma = \frac{\alpha - \beta}{\alpha + \beta} = \frac{1 - \sigma}{1 + \sigma}. \quad (4)$$

Proof. We consider the normalized version ($\beta = 1$) and show that the unique Hider strategy that ensures the same expected search time against any Searcher strategy S_C is h^* . We then show that this expected search time is $\delta/2$, so by Lemma 2, that is the value.

It follows from equation (2) that for any Hider density $h(x)$ with cumulative distribution $H(x)$ and any Searcher strategy S_C , we can write the expected search time $T(h, S_C)$ as

$$\begin{aligned} T(h, S_C) &= \int_C f(x) h(x) dx + \int_{C'} (f(x) + (1 + \alpha)(1 - x)) h(x) dx \\ &= \int_0^1 f(x) h(x) dx + \int_{C'} (1 + \alpha)(1 - x) h(x) dx. \end{aligned} \quad (5)$$

We evaluate the left integral of (5) using integration by parts, giving

$$\begin{aligned} \int_0^1 f(x) h(x) dx &= \int_0^1 -f(x) \frac{d}{dx} (1 - H(x)) dx \\ &= [-f(x)(1 - H(x))]_0^1 + \int_0^1 f'(x)(1 - H(x)) dx \\ &= 0 + \alpha \int_C 1 - H(x) dy + \int_{C'} 1 - H(x) dx. \end{aligned} \quad (6)$$

We evaluate the expected time $T(h, S_C)$ by substituting (6) in (5) to obtain

$$T(h, S_C) = \alpha \int_C 1 - H(x) dx + \int_{C'} (1 - H(x) + (1 + \alpha)(1 - x) h(x)) dx. \quad (7)$$

We seek an $h = h^*$ such that (7) is independent of C . So for all $x \in [0, 1]$, we must have

$$\alpha(1 - H^*(x)) = (1 - H^*(x) + (1 + \alpha)(1 - x)h^*(x)) \quad \text{or} \quad (8)$$

$$H^*(x) = \gamma \frac{1 - H^*(x)}{1 - x}, \quad H^*(0) = 0, \quad H^*(1) = 1, \quad \gamma = \frac{\alpha - 1}{\alpha + 1}.$$

This ODE has the unique solution given by (3). For this $H(x)$, we have

$$\begin{aligned} \alpha(1 - H^*(x)) &= (1 - H^*(x) + (1 + \alpha)(1 - x)h^*(x)) \\ &= \alpha(1 - x)^{\frac{\alpha-1}{\alpha+1}}, \quad \text{so for any } C, \\ T(h^*, S_C) &= \int_0^1 \alpha(1 - x)^{\frac{\alpha-1}{\alpha+1}} dx = \frac{\alpha + 1}{2}, \quad \text{as claimed.} \end{aligned}$$

For an interval $[0, \beta]$ with traversal times α and β , γ is given by (4). Since (3) is the unique solution of the ODE (8), the optimal Hider strategy is unique.

■

The distribution h^* can also be expressed as the beta distribution with parameters 1 and γ . Figure 1 depicts a graph of the density function h^* and the corresponding cumulative distribution H^* , for the traversal time parameters $\alpha = 2$ and $\beta = 1$ (giving $\gamma = 1/3$). Note that the density approaches infinity as x approaches 1.




Figure 1. Graphs of $h^*(x)$ (top) and $H^*(x)$ (bottom), $\gamma = 2/3$.

Note that as $\alpha \rightarrow \infty$ ($\sigma \rightarrow 0$) for fixed $\beta = 1$, the h^* distribution converges to the uniform distribution; as $\alpha \rightarrow \beta = 1$ ($\sigma \rightarrow 1$) it converges to the atomic distribution at $x = 1$. The former case corresponds to the so-called arbitrary start single-speed game, in which the Searcher can start anywhere, and the uniform distribution was shown to be optimal by Dagan and Gal (2008). The latter case corresponds to traditional single-speed game, in which the only undominated Hider strategy is to hide at the opposite end ($x = 1$) to the Searcher starting point. So we have a nice continuous bridge between two well known variations of the single-speed game.

We can vary the game another way by specifying a starting point in the interior of an interval, and even give the two resulting arcs different slow speeds. It turns out that the value of the game is independent of the starting point. This is stated in the following result.

Corollary 6 *Consider the bimodal search game on an interval $Q = [0, \beta_1 + \beta_2]$ with an interior starting point $O = \beta_1$, and slow traversal times α_1 on the arc $a_1 = [0, \beta_1]$ and α_2 on the the arc $a_2 = [\beta_1, \beta_1 + \beta_2]$. This network is simply searchable, with $V = \delta/2 = (\alpha + \beta)/2$, where α and β are the total slow and fast traversal times $\alpha = \alpha_1 + \alpha_2$ and $\beta = \beta_1 + \beta_2$.*

Proof. Suppose the Hider hides along each arc a_i with distribution $p_i h_i^*$, where $p_i = \delta_i/\delta$ and h_i^* is the optimal hiding distribution on a_i, O given in Theorem 5. Fix any bimodal pure Searcher strategy. Since it must be depth-first it either searches a_1 before b_1 or the other way around. Suppose the former. Then if the Hider is in a_1 the expected search time is at least $\delta_1/2$, by Theorem 5 for a single arc; similarly if he is in a_2 , then the expected search time is at least δ_1 (the time to tour a_1) plus $\delta_2/2$ (the expected additional time to find him in a_2). Consequently the expected capture time is at least

$$p_1 \frac{\delta_1}{2} + p_2 \left(\delta_1 + \frac{\delta_2}{2} \right) = \frac{1}{2} \left(\frac{\delta_1}{\delta} (\delta_1) + \frac{\delta_2}{\delta} \left(\delta_1 + \frac{\delta_2}{2} \right) \right) = \frac{(\delta_1 + \delta_2)^2}{2\delta} = \frac{\delta}{2}.$$

The same estimate holds if the search takes place in the reverse order. ■

The construction of the optimal Hider distribution for an interior start game on an interval is illustrated below.

Example 7 *Consider the interval network $Q = [-1/3, 2/3]$ with interior start at $O = 0$. The arc $a_1 = [-1/3, 0]$ has $\alpha_1 = 2/3$ and $\beta_1 = 1/3$ ($\sigma_1 = 1/2$) and the arc $a_2 = [0, 2/3]$ has $\alpha_2 = 2$ and $\beta_2 = 2/3$ ($\sigma_2 = 1/3$).*

So $\delta_1 = 1$, $\delta_2 = 8/3$, and the total weights on a_1 and a_2 are $p_1 = 3/11$ and $p_2 = 4/11$. We have $\gamma_1 = (2/3 - 1/3) / (2/3 + 1/3) = 1/3$ and $\gamma_2 = (2 - 2/3) / (2 + 2/3) = 1/2$. So the weight on the left arc between $-y$ and 0 is $(3/11) \left(1 - \left(1 - \frac{-y}{1/3} \right)^{1/3} \right)$ and on the right arc between 0 and y is $(8/11) \left(1 - \left(1 - \left(\frac{y}{2/3} \right) \right)^{1/2} \right)$. The graph in Figure 2 below shows the optimal probability for the Hider to be between 0 and y on the two arcs a_1 and a_2 .




Figure 2. Optimal probability Hider is between 0 and y , $-1/3 \leq y \leq 2/3$ as described in Example 7.

4 Search on Rooted Trees

Under the continuing assumption that the Searcher uses depth-first paths, the analysis of bimodal games on an interval can be extended to rooted trees. In this section we show that rooted trees are simply searchable – their search values are half their respective bimodal tour times. The following ‘replacement lemma’ shows that the search value of a tree is not changed if we replace a subtree by a leaf arc with the same bimodal tour time.

Lemma 8 *Suppose Q, O is a rooted tree with $\delta(Q) = \delta$ and with a proper subtree Q_1 starting at some node P . Let Q^*, O be the tree which is the same*

as Q, O except that the subtree Q_1 at P is replaced by a subtree Q_2 consisting of a single arc a with $\delta(a) = \delta_1 \equiv \delta(Q_1)$. If Q^*, O and Q_1, P are simply searchable, then so is Q, O .

Proof. Fix an optimal Hider distribution h^* on Q^*, O . Let h be a distribution on Q which is equal to h^* on $Q - Q_1$ and is proportional to an optimal distribution on Q_1 . Let S be any pure bimodal search strategy on Q, O . We will show that $T(S, h) \geq \delta/2$. Let $J = [t_0, t_1]$ denote the closed time interval when S is searching Q_1 - it is an interval by the depth-first assumption. Let S^* be the pure bimodal search strategy on Q^*, O which has $S^*(t) = S(t)$ for $t \notin J$ and which is any bimodal search of Q_1, P during J . We will show that

$$T(S, h) \geq T(S^*, h^*) \geq \delta/2, \text{ by assumption on } h^*.$$

Let F and F^* denote the cumulative distribution for the capture times corresponding to S, h and S^*, h^* , respectively. Note that $T(S, h)$ and $T(S^*, h^*)$ are the means of the respective distributions F and F^* , so it is enough to show that the mean of F is at least the mean of F^* . Clearly $F(t) = F^*(t)$ for $t \notin J$. Furthermore, the conditional mean of F on J is at least the conditional mean of F^* on J , because $E(F(t) - F(t_0)) \geq \delta_1/2$ (by assumption) and $E(F^*(t) - F^*(t_0)) = \delta_1/2$ (by Theorem 5). It follows that the mean of F is at least the mean of F^* , so $T(S, h) \geq T(S^*, h^*) \geq \delta/2$, as claimed. It follows from Theorem 2 that $V(Q, O) = \delta/2$, and hence Q, O is simply searchable. ■

Note that in choosing the traversal times $\alpha(a)$ and $\beta(a)$ of the new arc a in the above construction, we have to satisfy the equation $\alpha(a) + \beta(a) = \delta_1$, where δ_1 is given. This gives us one degree of freedom, which we can exploit.

We will show soon that we can use the above lemma to establish that $V = \delta/2$ for all trees, but to illustrate the idea we first give a special case of an interval with two distinct slow search times. Recall that we equate fast time β with length.

Example 9 Take $Q = [0, 1]$, $O = 0$, as in Section 3 and with an artificial node $P = 1/2$ separating subintervals with distinct slow speeds. Consider a decomposition $Q_1 = [1/2, 1]$, and $Q_0 = Q - Q_1 = [0, 1/2]$. Suppose that α_0 (slow traversal time of Q_0) is 1 and α_1 (slow traversal time for Q_1) is 2; so we have distinct slow speeds on the full interval, $\sigma_0 = 1/2$ (on the left) and $\sigma_1 = 1/4$ (on the right). Note that $\delta(Q_1) = \alpha_1 + \beta_1 = 2 + 1/2 = 5/2$. We use Lemma 8 to replace Q_1 with another arc $Q_2 = a$ (starting at $P = 1/2$)

such that $\delta(a) = 5/2$. If we can also achieve $\sigma_a = 1/4$, the same as σ_0 , then we will have a single arc network Q^* with constant slow speed, so we know from Theorem 5 that $V(Q^*) = \delta(Q)/2 = 2$. Thus our algebraic problem is to have slow and fast traversal times for the new arc α which satisfy

$$\alpha + \beta = 5/2 \text{ and } \beta/\alpha = 1/2, \text{ or } \alpha = 10/6 \text{ and } \beta = 5/6 .$$

The new network Q^* constructed by Lemma 8 has slow and fast traversal times $\alpha^* = 1 + 10/6 = 16/6$ and $\beta^* = 1/2 + 5/6 = 8/6$, with $\gamma^* = (16/6 - 8/6) / (16/6 + 8/6) = 1/3$ by (4) and constant slow speed $\sigma^* = 1/2$. Hence by Theorem 5 the optimal cumulative Hider distribution is given by $H^*(y) = 1 - (1 - y / (8/6))^{1/3}$ on $[0, 8/6]$. The arc $Q_0 = [0, 1/2]$ is mapped by the Lemma isometrically onto the same interval $[0, 1/2]$, which is hidden in with total probability $1 - (1 - (1/2) / (8/6))^{1/3} = 1 - \sqrt[3]{5}/2 \simeq 0.145$. The second arc $Q_1 = [1/2, 1]$ is mapped onto the interval $[1/2, 10/6]$ linearly by $y \rightarrow g(y) = 1/2 + (5/3)(y - 1/2)$. So the optimal Hider distribution on the original two speed interval $Q_0 = [0, 1]$ is given by $H^*(y)$ for $y \leq 1/2$ and by $H^*(g(y))$ for $1/2 \leq y \leq 1$, as illustrated in Figure 3.




Figure 3. Optimal Hider distribution for Example 9.

Clearly this example can be extended more generally, so any tree consisting of two consecutive arcs above the root, with distinct slow speeds, is simply searchable, as stated formally below.

Lemma 10 *Let Q, O be the rooted tree based consisting of an interval $[0, \beta_1 + \beta_2]$ with root $O = 0$, with slow speed α_1/β_1 on arc $[0, \beta_1]$ and distinct slow speed α_2/β_2 on arc $[\beta_1, \beta_1 + \beta_2]$. Then Q, O is simply searchable, i.e. has value $V(Q, O) = \delta/2 = (\alpha_1 + \alpha_2 + \beta_1 + \beta_2) / 2$.*

The two procedures illustrated in Examples 7 and 9 and established in Corollary 6 and Lemma 10 can be used to recursively convert any rooted tree into a single-speed interval with root at the end. We take any penultimate node (leading to two leaf nodes) and convert the two leaf arcs to a single one by Corollary 6. Then we convert the two arcs (going in and out of the former penultimate node) to a single arc with a constant speed by Lemma 10. A slightly more general procedure is outlined in the proof of the following.

Theorem 11 *Every rooted tree is simply searchable. That is, the search value $V(Q, O)$ of a rooted tree Q, O exists and is equal to half its bimodal tour time δ .*

Proof. We prove this result by induction on the number n of arcs of Q , noting that we must put (artificial) nodes of degree 2 at points of an arc where the slow search speed changes, as in Example 9 above. If $n = 1$ or 2 then Q is topologically an interval. If the root is in the interior, it has two arcs and is simply searchable by Corollary 6. If the root is at an end, it has one or two arcs and is simply searchable by Theorem 5 or Lemma 10. Suppose the result is true for trees with up to $n - 1$ arcs, and let Q, O be a rooted tree with $n > 2$ arcs. Let Q_1 be subtree of Q rooted at some node P with at least two and fewer than n arcs. Since Q_1 has fewer than n arcs it is simply searchable by the induction hypothesis. Now apply Lemma 8 to obtain a new rooted tree Q^*, O by replacing Q_1 at P with a single arc a with $\delta(a) = \alpha(a) + \beta(a) = \delta(Q_1)$, with $V(Q, O) = V(Q^*, O)$ and $\delta(Q^*) = \delta(Q) = \delta$. Since Q^* has at most $n - 1$ arcs, it is simply searchable, and hence so is Q, O . ■

We now show how the ideas of this section can be used to construct an optimal Hiding distribution on a rooted tree.

Example 12 *Consider the rooted tree drawn on the left in Figure 4. The lengths (and fast travel times β) of the arcs are indicated to the left of the arcs in the figure and the slow travel times α are indicated on the right. Note that the subtree at P is the same as the interior start interval of Example 7,*

which has a bimodal tour time δ of $11/3$.




Figure 4. Illustration of networks in Example 12 with fast and slow travel times indicated.

According to Lemma 8, we can replace it (as we do on the right) with a single arc of $\delta = 11/3$. We also want this new arc to have the same speed as the root arc, namely $\sigma = 1/4$. So we solve $\alpha + \beta = 11/3$ and $\beta/\alpha = 1/4$ to get times $\alpha = \frac{44}{15}, \beta = \frac{11}{15}$ for the new arc above P . The full tree on the right has $\beta = 1/2 + 11/15 = 37/30, \alpha = 4\beta$, and hence by (4) it has $\gamma = (4\beta - \beta) / (4\beta + \beta) = 3 / (5\beta) = 18/37$. So the cumulative distribution on OP (on the left or right) is given by $1 - (1 - y / (37/30))^{18/37}$, which is the blue curve on the right in Figure 5, drawn up to $y = 1/2$. The total probability of hiding on OP is $1 - (1 - (1/2) / (37/30))^{18/37} \simeq 0.22$. The distribution of the subtree above P is the same as in Example 7, except that the probabilities are reduced by a factor of $.78 = 1 - .22$. The optimal Hider distribution on the tree is shown below in Figure 5, with the lower arc distribution drawn to the

right.




Figure 5. Optimal Hider distribution on tree of Figure 4.

To summarize some of our results, we have shown that both Eulerian networks (Theorem 4) and trees (Theorem 11) are simply searchable. As in Gal (2000), these results can be combined to show that weakly Eulerian networks (roughly speaking, networks which consist of disjoint Eulerian networks arranged in a tree-like fashion) are also simply searchable. The proof is essentially the same as that of Gal, or in the more recent paper of Alpern (2011). In the latter, given any weakly Eulerian network Q_0 , we first identify all the nodes in a given Eulerian component, to produce another network Q_1 . We then replace every resulting loop at such an identified node by an arc of half its length, to arrive at a final network Q_2 , which is a tree. These two

transformations are shown in Figure 6.




Figure 6. A network Q_0 and its transformed networks, Q_1 and Q_2 , illustrating how Q_0 is simply searchable.

Both transitions result in networks that are no worse for the searcher (lower or equal value) and preserve the total length of the network. For bimodal games, we simply note that the transitions are still favorable to the Searcher, and the bimodal tour times δ are preserved. So the same arguments show that $V(Q_2) = \delta/2$ and hence $V(Q_0) \geq V(Q_1) \geq V(Q_2) = \delta/2$. Since we always have $V(Q_0) \leq \delta/2$ by Lemma 2, $V(Q_0) = \delta/2$, as claimed. We do not know if the converse result proved by (Gal, 2000), that simply searchable networks are weakly Eulerian, also holds for bimodal search.

5 Existence of the Value

In other sections of this paper, we always establish the existence of the value directly by exhibiting explicit optimal strategies for the players. However it is nice to know that standard minimax theorems can guarantee the existence of the value of the game played on an arbitrary network Q , in advance of knowing optimal strategies. We outline existence arguments here, beginning with the easier case in which Q has a single arc, and then proceeding to

general networks. The argument is more complicated than for the traditional (single speed) theory.

We first show the value exists when $Q = [0, 1]$ is the single arc network of Section 3, with constant slow speed σ . Let \mathcal{C} denote the family of all ε -thick (as defined in Parks (1979)) closed subsets of $Q = [0, 1]$, that is, sets which are finite unions of closed intervals of length at least 2ε , or contain an endpoint. The set \mathcal{C} is compact with respect to the Hausdorff distance and Lebesgue measure $\lambda(C)$ is continuous for $C \in \mathcal{C}$ (Krantz and Parks, 1999). (Without the ε -thick assumption, the sets $C_n = \{i/n : i = 0, 1, \dots, n\}$ which have Lebesgue measure 0, converge to $[0, 1]$.)

For any closed set $C \in \mathcal{C}$, define S_C to be the path in Q that moves at speed σ when $t \in C$ and at speed 1 when $t \notin C$, while going from 0 to 1. When going back to 0 it does the opposite.

We wish to show that the capture time $T(x, S_C)$ is lower semicontinuous in x and C , so that the existence of a value follows from the Glicksberg (1952) or Alpern-Gal (1988) theorems. We first observe that the first arrival time function $f(x, C)$ of (1) can be written as

$$f(x, C) = x - (\alpha - 1) \lambda(C \cap [0, x]),$$

which as observed above is continuous in x and C (this requires our assumption of ε -thickness). So if $x \in C$, then $T(x, S_C) = f(x, C)$ and a small change to \hat{x} and \hat{C} either (i) keeps $\hat{x} \in \hat{C}$, in which case $T(\hat{x}, S_{\hat{C}}) = f(\hat{x}, \hat{C})$ is close to $T(x, S_C)$ by continuity of f , or (ii) has $\hat{x} \notin \hat{C}$, in which case $T(\hat{x}, S_{\hat{C}}) = f(\hat{x}, \hat{C}) + (1 + \alpha)(1 - \hat{x}) > T(x, S_C)$. If $x \notin C$, then $T(x, S_C) = f(x, C) + (1 + \alpha)(1 - x)$ and since C is compact it follows that for close \hat{x}, \hat{C} we also have $\hat{x} \notin \hat{C}$ and $T(\hat{x}, S_{\hat{C}}) = f(\hat{x}, \hat{C}) + (1 + \alpha)(1 - \hat{x})$, which is close to $T(x, S_C)$ by continuity of f . So T is lower semicontinuous, as claimed.

Since the pure strategy set Q of the Hider is compact and the pure strategy set \mathcal{C} of the Searcher is compact with respect to Hausdorff distance, the lower semicontinuity of the payoff function insures the existence of the value and an optimal mixed Searcher (minimizer) strategy, by the Alpern-Gal Minimax Theorem (Alpern and Gal, 1988).

If Q, O is a general rooted network, a pure strategy for the Searcher is determined by a pair (P, K) , where P is a Lipschitz function from R^+ into

Q with Lipschitz constant 1, K belongs to \mathcal{C} , now taken as subsets of the time space R^+ , and $P(K) = Q$. A Searcher pure strategy $S_{P,K}$ is one that follows the path P , going at the slow speed σ (which is constant on arcs) at times $t \in K$ and at speed 1 otherwise. The condition $P(K) = Q$ now says that every point of Q is slow searched. We endow the set \mathcal{P} of all unit speed paths P with the topology of uniform convergence, under which it is compact. Consequently the Searcher pure strategy set $\mathcal{S} = \mathcal{P} \times \mathcal{C}$ is also compact. The lower semicontinuity of the payoff function follows as before. The only additional argument is that if a point x is not slow searched on the first traversal, that is if $x \notin P(K)$, then for close $\hat{x}, \hat{P}, \hat{C}$ we also have $\hat{x} \notin \hat{P}(\hat{C})$. To see this, note that as $P(K)$ is a compact subset of Q , we have $d(x, P(K)) = \varepsilon > 0$, where d is the metric on Q . Taking each "hatted" variable within $\varepsilon/3$ of its "unhatted" version, with respect to d , uniform metric on \mathcal{P} , and Hausdorff metric on \mathcal{C} , we obtain $\hat{x} \notin \hat{P}(\hat{C})$, as required.

6 Relation to the Kikuta-Ruckle Search Game

We now discuss how the bimodal search game is related to another search game proposed by Kikuta and Ruckle (1994), studied further by Kikuta (1995) and recently extended by Baston and Kikuta (2013), in which the Hider must hide at a node and the Searcher must pay a *search cost* to search any given node. The Searcher's choice of skipping or searching a node corresponds to the idea of moving fast or slowly respectively. This idea of the Searcher choosing whether to "skip or search" can be found in Gluss (1961), though not in a game theoretic context.

Kikuta's game $K = K(Q, O)$ is played on a network Q with root O , like our bimodal search game, but it follows the traditional search game formulation in which the Searcher can move only at unit speed (so $\alpha = \beta$ in our model). However, the Hider can only hide at nodes (including specified nodes of degree 2) and the Searcher must pay a fixed search cost $c(i)$ for searching a node i . When reaching a node the Searcher can either pay the search cost to search it or pass it without searching it. The total payoff is the sum of the time taken for the Searcher to find the Hider and the total cost she has paid to search nodes prior to finding the Hider.

Consider the game $K(Q_n, O)$ played on the network Q_n in Figure 7, in which a single arc of unit length is split into n smaller arcs by inserting n

equally spaced nodes of search cost c/n . In this example $n = 6$.




Figure 7. The network Q_n .

This can be viewed as an approximation of the bimodal search game on the unit arc. The interpretation is that paying the search cost of c/n to search a node in Q_n is equivalent to passing that node at the slow speed (and hence searching it) in the bimodal search game. Traversing the whole of Q_n without paying the search cost at any of the nodes is equivalent to traversing the unit arc at the fast speed of 1 in the bimodal search game; traversing the whole of Q_n and paying the search cost of c/n at each of the n nodes is equivalent to traversing the unit arc at the slow speed of $1/(1+c)$ in the bimodal search game.

The authors showed in Alpern (2010) and Alpern and Lidbetter (2013) that the game $K(Q_n, O)$ is a special case of a search game on a *variable speed network*. In the variable speed model the Hider can hide anywhere but the Searcher's speed on a given arc depends on her direction of travel. For example in the variable speed network \tilde{Q}_n in Figure 8, the Searcher traverses the horizontal arcs in the same time $1/n$ in either direction, and traverses the diagonal arcs in time c/n in the upward direction and time 0 in the downward direction. We use the convention of writing the traversal time away from the root on the left of an arc.




Figure 8. The network \tilde{Q}_n .

It is optimal for the Hider to choose one of the leaf nodes in \tilde{Q}_n (that is either the end of a diagonal arc or the end of the horizontal arc on the far

right hand side), as other hiding points are dominated. It is thus easy to see that the variable speed game played on \tilde{Q}_n is equivalent to the Kikuta game, $K(Q_n, O)$, so we can compare the variable speed game on \tilde{Q}_n to a bimodal game on the unit arc I with slow travel time $\alpha = 1 + c$ and fast travel time $\beta = 1$.

In Alpern and Lidbetter (2014), the authors present a solution of the variable speed game played on tree networks, showing that it is optimal for the Hider to use the *Equal Branch Density* (EBD) distribution, first defined by Gal (1979) for the classical game, and for the Searcher to use a strategy which mixes between depth-first searches. The EBD is the unique distribution on the leaf nodes such that at any branch node, the probability that the Hider is in a branch is proportional to its tour time. For fixed n and $k = 0, 1, \dots, n - 1$, let $q_n(k)$ be the probability under the EBD distribution that the Hider chooses a leaf arc in \tilde{Q}_n whose base is at distance greater than k/n from the root. So $q_n(0)$ is the probability the Hider does not choose the first leaf arc of \tilde{Q}_n (that is the leaf arc on the far left in Figure 8).

The probability the Hider *is* on the first leaf arc of \tilde{Q}_n under the EBD distribution is proportional to the tour time, c/n of that arc. Since the total tour time of \tilde{Q}_n is $2 + c$, this probability is $(c/n) / (2 + c)$. Then $q_n(0)$ is the complementary probability given by

$$q_n(0) = 1 - \frac{c/n}{2 + c} = \frac{n(2 + c) - c}{n(2 + c)} = 1 - \frac{\gamma}{n},$$

where $\gamma = (\alpha - \beta) / (\alpha + \beta) = c / (2 + c)$, as in (4). Recursively calculating the probabilities $q_n(1), q_n(2), \dots$, we obtain the formula

$$q_n(k) = \prod_{j=0}^k \left(1 - \frac{\gamma}{n - j} \right), \text{ for } k = 0, 1, \dots, n - 1.$$

The probabilities $q_n(k)$ also give the optimal strategy for the Kikuta game $K(Q_n, O)$. For the bimodal search game played on the unit arc I , we showed in Section 3 that the optimal probability for the Hider to choose a point after $x \in [0, 1]$ is $1 - H^*(x) = (1 - x)^\gamma$. In Q_n the furthest node from the root before the point x is the $\lfloor nx \rfloor$ th node, and the distance of this node from O converges to x as $n \rightarrow \infty$. In the game $K(Q_n, O)$ it is optimal for the Hider to hide after this point with probability $q_n(\lfloor nx \rfloor)$. We will show that as $n \rightarrow \infty$ this probability $q_n(\lfloor nx \rfloor)$ converges to $(1 - x)^\gamma$ so that the

optimal Hider strategy in Kikuta's game converges to the optimal strategy in the bimodal game.

Theorem 13 *As $n \rightarrow \infty$, the optimal Hider strategy q_n for the game $K(Q_n, O)$ on the unit interval converges to the optimal Hider strategy in the bimodal search game $B([0, 1], O)$. That is*

$$q_n(\lfloor nx \rfloor) \equiv \prod_{j=0}^{\lfloor nx \rfloor} \left(1 - \frac{\gamma}{n-j} \right) \rightarrow (1-x)^\gamma \text{ as } n \rightarrow \infty.$$

We feel sure that Theorem 13 is a known result but are unable to find it in the literature, so we include a sketch of a proof in the online appendix for completeness. We thank an anonymous referee to whom the final simplified proof we present is due.

7 Lucky Fast Searcher

Up until now, we have been assuming that when passing the Hider at the fast speed (taken to be 1), the Searcher has no chance at all of detecting him. However we note that one often finds missing objects by chance (a glint of light, a lucky glance) even at a speed that is not guaranteed to find them. In this section, we attempt to model this possibility by ascribing a probability q , known to both players, to the probability of detecting the Hider when passing him at fast speed. Since this new game $B_q(Q, O)$ is quite complicated, we will only analyze the case where Q is the unit interval $I = [0, 1]$, which was solved (for the original case $q = 0$) in Section 3. Thus the fast traversal time is given by $\beta = 1$ and the slow traversal time is $\alpha = 1/\sigma > 1$, where $\sigma < 1$ is the slow (search) speed.

Observe that when $q = 1$ there is no point in the Searcher adopting the slow search speed, as going at fast speed 1 also captures the Hider upon finding him. Thus the game $B_1(Q, O)$ is equivalent to the usual network search game defined by Isaacs, and the various results of Gal apply. For the simple case $Q = I$, the solution is a trivial one where the Hider always hides at the far end 1, the Searcher goes there at maximum speed, and the capture time (and value) is 1. When $q = 0$, as observed in the previous paragraph, this is just our original bimodal search game $B(Q, O)$. Thus the continuum of games $B_q(Q, O)$, $0 \leq q \leq 1$, forms a nice bridge between bimodal search

games as discussed in this paper and familiar search games ($q = 1$) as studied by Gal.

In our original bimodal games $B(Q, O)$, the Searcher could often (for example, in all weakly Eulerian networks) restrict to *exhaustive* searches, which we define as Bimodal Chinese Postman Tours. For the interval case $Q = I$ these go from one end of the interval to the other and back again. However, with lucky Searcher games, the Searcher may want to exploit the possibility of going back and forth several times at the fast speed. We call this larger family of search paths *non-exhaustive*. It unsurprisingly turns out that exhaustive search is simpler to analyze, so we do this first in Subsection 7.1; then in Subsection 7.2 we analyze the more complicated non-exhaustive form of the game. Whilst we will consider more general search strategies for exhaustive search in 7.1, it turns out that the Searcher constructs his optimal mixed strategy from only the two pure strategies of going fast to the end of I and slowly back again or going slowly to the end of I and fast back again, as in Section 3. In Subsection 7.2, we restrict attention to pure Searcher strategies S^j ($j = 0, 1, \dots$) defined by going from one end of the interval to the other fast the first j times and then back to the other end slowly on the $(j + 1)$ th transit. (So of these strategies, only S^0 and S^1 are permitted in Subsection 7.1).

7.1 Exhaustive lucky Searcher games

In this subsection we discuss the solution to the lucky fast Searcher game on the interval with fast traversal speed 1 and slow traversal speed σ , when the Searcher is forced to use pure strategies which are Bimodal Chinese Postman Tours.

For exhaustive search, it turns out that the Searcher need not use general strategies S_C , which alternate speeds while going to the far end. As in the original bimodal game, he can restrict to S^0 (going out slowly and finding the Hider before reaching the end) and S^1 (going out fast, finding the Hider on the way out with some probability and if not, finding him on the return trip). For q sufficiently close to 1 (how close depends on σ) the pure strategy S^1 is optimal. The optimal probability p_0 of choosing S^0 (with probability

$1 - p_0$ of S^1) turns out to be given by the formula

$$p_0 = p_0(q, \sigma) = \left(\frac{1 - q(1 + \sigma)}{2 - q(1 + \sigma)} \right)^+, \text{ so}$$

$$p_0 = 0 \text{ and } S^1 \text{ is optimal for } q(1 + \sigma) \geq 1, \text{ or } q \geq 1/(1 + \sigma).$$

This extends the case of $p_0 = 1/2$ (what we called *random* bimodal Chinese postman path) for the non-lucky case $q = 0$. As q increases from 0, the slow probability p_0 decreases from $1/2$ until it eventually reaches $p_0 = 0$ for $q = 1/(1 + \sigma)$. This is intuitively plausible, as going fast is more attractive when it includes an increased possibility of capturing the Hider when passing him. For $q \geq 1/(1 + \sigma)$, the pure strategy S^1 ($p_0 = 0$) is optimal.

Next we consider the optimal Hider strategies. In the region $q \geq 1/(1 + \sigma)$ where the Searcher uses S^1 , the Hider adopts the pure strategy $H = 1$ (hide at the far end), as in traditional search games on trees, where the Hider hides at leaf nodes. In the complementary region, the Hider uses a variation of the beta distribution defined earlier in (3), with the probability density given by

$$h_{q,\sigma}^*(x) = \gamma_{q,\sigma} (1 - x)^{\gamma_{q,\sigma} - 1}, \text{ where} \quad (9)$$

$$\gamma_{q,\sigma} = \gamma / (1 - q) = \frac{1 - \sigma}{(1 + \sigma)(1 - q)}. \quad (10)$$

To see how the optimal Hider distribution $\gamma_{q,\sigma}$ adapts to decreasing q , in Figure 9 we plot a series of distributions for fixed $\sigma = 1/2$ and incremental values of q . The black, green, red and blue lines (bottom to top at $x = 0$) correspond to $q = 0, 1/2, 3/5$ and $2/3$, respectively. When $q = 2/3 = 1/(1 + \sigma)$, then the optimal Hider strategy is the uniform distribution (though this is not true for general σ). It is also optimal for the Hider to use the pure strategy $x = 1$.




Figure 9. The optimal Hider distribution for $q = 0, 1/2, 3/5, 2/3$.

We summarize our results in the following theorem, the proof of which is given in the appendix.

Theorem 14 *The solution of the exhaustive form of the lucky Searcher game on the interval $I = [0, 1]$, starting at 0, is as follows:*

- For $q \geq 1/(1 + \sigma)$, the optimal strategies are S^1 (fast to 1) for the Searcher and $H = 1$ for the Hider. The value is simply $V = 1$.
- For $q \leq 1/(1 + \sigma)$, the unique optimal Hider strategy is $h_{q,\sigma}^*$ (as given in (9)) and the optimal Searcher strategy is to play S^0 (slow to 1) with probability $p_0 = p_0(q, \sigma)$ and S^1 (fast to 1, slow back) with probability $1 - p_0$. The value V^e for the exhaustive problem is given by

$$V^e = \frac{(1 - q)(1 + 1/\sigma)}{2 - q(1 + \sigma)}. \quad (11)$$

We note that if the Searcher is restricted to the ‘combinatorial strategies’ S^0 and S^1 , the Hider can adopt any distribution with the same mean $m^e = 1/(1 + \gamma_{q,\sigma}) = \frac{(1-q)(1+\sigma)}{2-q(1+\sigma)}$ as $h_{\sigma,q}^*$.

In Figure 10 we plot q against σ to show how the region is split up according to the optimal strategies of the players. The region labeled (i) corresponds to $q \geq 1/(1 + \sigma)$, and regions (ii) and (iii) correspond to $q \leq 1/(1 + \sigma)$. The distinction between regions (ii) and (iii), while irrelevant in this subsection, will be made clear in the next Subsection 7.2.




Figure 10. Decomposition of the σ - q plane by optimal strategy type.

7.2 Non-exhaustive lucky Searcher games

We now consider the non-exhaustive version of lucky search, where the Searcher may use any ‘combinatorial strategy’ S^0, S^1, S^2, \dots . So the question to address is whether these new strategies help the Searcher and lower the value, or whether the Hider mixed strategies adopted in Subsection 7.1 keep the same minimum search time against these new strategies.

We will consider three regions corresponding to the relationship between q and σ . Region (i), depicted in Figure 10, is defined by the inequality

$q \geq 1/(1 + \sigma)$, and we note that the pure strategy $H = 1$ certainly ensures that $T = 1$ against any strategy of the Searcher, as it takes at least time 1 to reach it. So there is nothing new to say about region (i). The next question is whether the mixed strategy $h_{q,\sigma}^*$ which was formerly optimal for $q \leq 1/(1 + \sigma)$ still guarantees the exhaustive search value V^e of (11) against the new strategies S^2, S^3, \dots . The answer to this question depends on the relative magnitude of q and $2\sigma/(1 + \sigma)$, the region (ii), where q is smaller; and the region (iii), where q is larger. In region (ii), we have $T(h_{q,\sigma}^*, S^j) \geq V^e$ for all j , so there is no change in the analysis of this region. It turns out that the only region where adding non-exhaustive search strategies actually helps the Searcher is the triangular region (iii) where $2\sigma/(1 + \sigma) \leq q \leq 1/(1 + \sigma)$.

In region (iii) the optimal searcher strategy can be described as a stationary behavioral strategy or as a mixture of the S^j where j is odd. We give the former description here and the latter in the statement of Theorem 15. The optimal behavioral search strategy can be described by specifying the probability of going slow or fast on each transit: When at 0 going to 1, always go fast. When at 1 going to 0, go slow with probability p' , independent of previous choices. The formula for the optimal slow probability p' is given by

$$p' = p'(q, \sigma) = \frac{q^2 \sigma}{(1 - q)(1 - q\sigma)}. \quad (12)$$

The optimal strategies for the non-exhaustive form of the lucky Searcher game are stated formally in the theorem below, whose proof can be found in the online appendix.

Theorem 15 *Consider the non-exhaustive form of the lucky Searcher game on the interval $I = [0, 1]$. In the regions (i) given by $q \geq 1/(1 + \sigma)$ and (ii) given by $q \leq 1/(1 + \sigma)$ and $q \leq 2\sigma/(1 + \sigma)$, the solution is the same as for the exhaustive form of the game described in Theorem 14. In the remaining region (iii) given by $2\sigma/(1 + \sigma) \leq q \leq 1/(1 + \sigma)$, the solution is as follows:*

- *The optimal Searcher strategy is to choose each odd strategy S^{2j+1} with probability $p'(1 - p')^j$, where p' is given by (12).*
- *The optimal Hider strategies are those where the mean distance to 0 is given by*

$$m^{ne} = \frac{1 - \sigma - (1 - q)^2(1 + \sigma)}{q(2 - q(1 + \sigma))}.$$

- The value V^{ne} of the non-exhaustive game is given by

$$V^{ne} = \frac{(1 + \sigma)(1 - q)^2 + 1 - \sigma}{q(2 - q(1 + \sigma))}.$$

It is useful to observe that in the region (iii), where the Searcher benefits from non-exhaustive search, the optimal Hider strategy is further from 0, that is $m^{ne} \geq m^e$. The intuitive explanation is that in region (iii), the Searcher's slow speed is quite slow, so he is more inclined to persevere at his fast speed, and hence he may traverse the arc several times at the fast speed before finally resorting to using his slow speed.

We note that since we are restricting the Searcher to pure strategies S^0, S^1, \dots , this means that the Hider strategy of hiding certainly at distance m^{ne} from O is optimal. If we had made the same restriction in Subsection 7.1, the strategy of choosing to hide at m^e would have also been optimal for the Hider.

8 Conclusion

The simple search paradigm introduced by Isaacs (1965), in which target capture occurs when the Searcher merely passes the Hider's location at her usual speed, has proved very useful and generated many important results. In addition to those of Gal already mentioned, we should add those of Reijnen and Potters (1993) and Pavlovic (1995). However this paper observes that this paradigm does overlook some important aspect of search, particularly when the sought hidden object is small or well hidden, and consequently requires a slower speed for actual detection. This paper shows that, when such a slow detection speed is specified, some of the existing theory is maintained or generalized, while other aspects change dramatically. In particular optimal hiding requires use of all of the network, even in some very simple cases (all tree networks, for example). We hope that our work will encourage dialogue between demining practitioners and those working in theoretical search theory. We are optimistic that the theory of bimodal search could be developed to be of some practical use in constructing demining algorithms, so that mines left buried from earlier conflicts can be safely removed.

A useful extension of our model (suggested by an anonymous referee) would be to allow intermediate search speeds, where detection probability is

a decreasing concave function of speed. Such a stochastic model has been studied as a stochastic game (Alpern *et al*, 2014) in the the context of search for a mobile hider.

We also note that we have assumed in our model that the object in question is hidden adversarially, but we did not consider the analogous *Bayesian problem*. That is, if the object is located on the network according to some given probability distribution. In some circumstances this problem is easier to solve. For example, if the Hider is located uniformly on a tree we can show that the Searcher’s optimal strategy is to perform a depth-first search of the tree, traversing an arc slowly the first time it is encountered and quickly the second time.

Another aspect of bimodal search that has not been explored here is the context in which two (or more) agents all want to minimize the meeting time, that is, in the area of rendezvous search, e.g. Alpern (1995) and Weber (2012). One might require that one or both of the searchers are in their slow mode to enable rendezvous. In particular the case of rendezvous on the line (as in Gal (1999), Howard (1999); Chester and Tutuncu (2004) and Han et al (2008)) should be re-evaluated in the context in which the agents must slow down to enable mutual detection. We also believe that the problem of search in a maze (unknown network), as studied in Anderson and Gal (1990) could usefully be attacked from this direction.

Acknowledgements

Steve Alpern wishes to acknowledge support from the Air Force Office of Scientific Research under grant FA9550-14-1-0049.

References

- Alpern S (1995). The rendezvous search problem. *SIAM J. Control Optim.* 33(3):673-683.
- Alpern S (2010) Search games on trees with asymmetric travel times. *SIAM J. Control Optim.* 48(8):5547-5563.
- Alpern S (2011) A new approach to Gal’s theory of search games on weakly Eulerian networks. *Dynam. Games Appl.* 1(2):209-219.

- Alpern S, Fokkink R, Gal S, Timmer M (2014) On search games that include ambush. *SIAM J. Control Optim.* 51(6):4544-4556.
- Alpern S, Gal S (1988) A mixed strategy minimax theorem without compactness. *SIAM J. Control Optim.* 26(6):1357-1361.
- Alpern S, Gal S (2003) *The Theory of Search Games and Rendezvous*. Kluwer International Series in Operations Research and Management Sciences (Kluwer, Boston), 319.
- Alpern S, Lidbetter T (2013) Mining coal or finding terrorists: the expanding search paradigm. *Oper. Res.* 61(2):265-279.
- Alpern S, Lidbetter T (2013) Searching a Variable Speed Network. *Math. Oper. Res.* Published online on 23 December 2013, ISSN 1526-5471.
- Anderson EJ, Gal S (1990) Search in a maze. *Prob. Eng. Inf. Sci.* 4:311-318.
- Baston V, Kikuta K (2013) Search games on networks with travelling and search costs and with arbitrary searcher starting point. *Networks* 62(1):72-79.
- Chester C, Tütüncü RH (2004) Rendezvous search on the labeled line. *Oper. Res.* 52(2):330-334.
- Dagan A, Gal S (2008) Network search games with arbitrary searcher starting point. *Networks* 52(3):156-161.
- Gal S (1979) Search games with mobile and immobile hider. *SIAM J. Control Optim.* 17(1):99-122.
- Gal S (1980) *Search Games* (Academic Press, New York).
- Gal S (1999) Rendezvous search on the line. *Oper. Res.* 47(6):974-976.
- Gal S (2000) On the optimality of a simple strategy for searching graphs. *Internat. J. Game Theory* 29(4):533-542.
- Gendron RP (1986) Searching for cryptic prey: evidence for optimal search rates and the formation of search images in quail. *Anim. Behav.* 34(3):898-912.
- Glicksberg IL (1952) A further generalization of the Kakutani Fixed Point Theorem, with application to Nash Equilibrium points. *Proc. Amer. Math. Soc.* 3(1):170-174.
- Gluss B (1961) Approximately optimal one-dimensional search policies in which search costs vary through time. *Naval Res. Logist. Quart.* 8(3):277-283.
- Han Q, Du D, Vera J, Zuluaga LF (2008) Improved bounds for the symmetric rendezvous value on the line. *Oper. Res.* 56(3):772-782.

- Howard JV (1999) Rendezvous search on the interval and circle. *Oper. Res.* 47(4):550-558.
- Isaacs R (1965) *Differential Games* (Wiley, New York).
- Johnson D, Ali A (2012) Modeling and simulation of landmine and improvised explosive device detection with multiple loops. *J. Def. Model. Simul.* Published online before print September 11, 2012, doi: 10.1177/1548512912457886
- Kikuta K (1995) A search game with travelling cost on a tree. *J. Oper. Res. Soc. Japan*, 38(1):70-88.
- Kikuta K, Ruckle W (1994) Initial point search on weighted trees. *Naval Res. Logist.* 41(6):821-831.
- Krantz S, Parks H (1999) *The geometry of domains in space* (Birkhauser, Boston).
- McLean I (2001) *Designer dogs: improving the quality of mine detection dogs*, Geneva, Switzerland, Geneva International Centre for Humanitarian Demining.
- Parks H (1979) Some new constructions and estimates in the problem of least area. *Trans. A.M.S.* 248(2):311-346.
- Reijnierse JJ and Potters JAM (1993) Search games with immobile hider. *Internat. J. Game Theory* 21:385-394.
- Pavlovic L (1995) A search game on the union of graphs with immobile hider. *Naval Res. Logist.* 42(8):1177-1189.
- Weber R (2012) Optimal symmetric rendezvous search on three locations. *Math. Oper. Res.* 37(1):111-122.

Biographies

Steve Alpern is Professor of Operational Research at the University of Warwick. His area of pure mathematics is ergodic theory and dynamical systems, but his relevant expertise for this article is the area of search games. He has often collaborated with Shmuel Gal of the University of Haifa, and it is that work which has led to this offshoot of Gal's original theory of search games with immobile hidere. He also has an extensive collaboration with Thomas Lidbetter of the London School of Economics.

Thomas Lidbetter recently completed a PhD in mathematics at the London School of Economics (LSE), and went on to secure a Senior Fellowship in the Department of Management at the LSE. His research interests lie in search and game theory, and he has worked as an operational research analyst

in the UK Home Office. He is also interested in the application of search games in areas of mathematical biology, criminology and demining.