Increasing Threshold Search for Best-Valued Agents

Simon Shamoun^{a,*}, David Sarne^a

^aDepartment of Computer Science, Bar Ilan University, Ramat Gan, 52900, Israel

Abstract

This paper investigates agent-search for the agent with the best value in a multi-agent system, according to some value assignment. In the type of setting considered, agent values are independent of one another. Under this condition, classic state-space search methods are not very suitable solutions since they must probe the values of all agents in order to determine who the best-valued agent is. The method considered in this paper refines the number of agents that need to be probed by iteratively publishing thresholds on acceptable agent values. This kind of agent-search is applicable to various domains, including auctions, first responders, and sensor networks. In the model considered, there is a fixed cost for publishing the thresholds and a variable cost for obtaining agent values that increases with the number of values obtained. By transforming the threshold-based sequence to a probability-based one, the sequence with minimum expected cost is proven to consist of either a single search round or an infinite sequence of increasing thresholds. This leads to a simplified characterization of the optimal thresholds sequence from which the sequence can be derived. The analysis is extended to the case of search for multiple agents. One important implication of this method is that it improves the performance of legacy economic-search methods that are commonly used in "search theory". Within this context, we show how a threshold based search can be used to augment existing economic search techniques or as an economic search technique itself. The effectiveness of the methods for both best-value search and economic-search is demonstrated numerically using a synthetic environment.

^{*}Preliminary results of this work appeared in Proceedings of AAAI 2010

^{*}Corresponding author

Email addresses: srshamoun@yahoo.com (Simon Shamoun), sarned@cs.biu.ac.il (David Sarne)

1. Introduction

In many multi-agent systems (MAS), it is often necessary to find the agent associated with the *best* value (lowest or highest, depending on the application) amongst a set of agents, each associated with some value to the searcher. If there are no predefined relationships between the agents' values, state-based search methods like BFS, IDDFA, and A* must probe all agents for their values in order to determine which agent is associated with the best value. The number of agents that need to be probed can be reduced only by using some mechanism that distinguishes the set of agents that comply with some condition. One such mechanism is the publication of a request to all agents to communicate their values to the searcher if they comply with some threshold (i.e., whose values are below/above it). An intuitive search method based on this mechanism, denoted *increasing threshold search*, publishes a higher (or lower) threshold each time a request fails to yield any replies. A search strategy in this case defines the sequence of thresholds to use until the best-valued agent is found.

We illustrate the above by considering a government agency searching for the cheapest contractor to provide a service. The agency can issue a call for bids to all potential contractors, who then prepare and submit their bids to the agency, after which it selects the lowest bid amongst all those it received. While this protocol guarantees finding the lowest bidder, it is not necessarily the most cost-efficient one. For example, consider the case in which there is a cost to the contractors for preparing and submitting their bids and to the agency for processing them. Since the agency only needs to learn the value of the lowest bid and the corresponding bidder, the costs of preparing, submitting, and processing all bids are extraneous. However, it is still necessary to include all potential bidders in the search process in order to guarantee finding the lowest bidder.

If the agency is interested in minimizing the expected overall costs to all parties involved, from the social-welfare perspective, or merely in minimizing its own costs, it can apply increasing threshold search instead. The agency publishes a threshold on allowable bids (e.g., a reservation price), requesting that contractors submit their bids only if they are below this threshold. Assuming the contractors can and do estimate what the value of their bids will be once they prepare them, only those whose estimates are below the threshold actually prepare and submit their bids. If the agency does not receive any bids, then it publishes a higher threshold with a similar request. It repeats this process until it obtains a non-empty set of bids, which necessarily includes the lowest bid, and selects the lowest bid from that set. This potentially reduces the number of bids made and the associated costs while ensuring that all bidders are included in the search process. However, if there is a cost to the agency for publishing a call for bids, then additional costs are incurred each time a threshold is published. Despite these additional costs, the right sequence of thresholds can result in lower costs than the original protocol, thereby minimizing search costs to the agency and improving social welfare.

Many multi-agent systems can benefit from increasing threshold search. It is applicable in environments in which: (a) each agent is associated with a single value, which it knows, and which may be an intrinsic value associated with the agent or a mapping from a combination of values associated with the agent; and (b) the searcher can publish a request to all other agents. For example, a volunteer ambulance corps dispatcher needs to find the closest volunteer to an emergency. She must page the volunteers and request that they call back to learn their locations. Instead of requesting that all volunteers call back, she can request that only volunteers within a certain distance of the emergency call back, and repeat the request with greater distances until at least one volunteer calls. Similarly, in sensor networks, only extreme values may be of interest [38, 39]. Requesting that all sensors send their data significantly depletes the sensors' power supply. Rather, the user can request only readings above a certain threshold and iteratively decrease the threshold until at least one reply is received.

This paper thoroughly analyzes the problem of deriving the optimal threshold-based search sequence in settings similar to those above. It is assumed that all agent values are associated with a common distribution, which is known to the searching agent and remains constant over time [1, 2, 7, 29, 30, 31]. Learning the actual value of an agent incurs some cost to the searcher, the searched agent, or both (e.g., consuming some of the agents's resources for communicating with each other, querying for relevant information, and processing the information received). Such costs are commonly incurred in multi-agent systems in the absence of a central source that can supply full, immediate, and reliable information on the environment and the state of the other agents that can be found [2, 12, 23]. Additionally, there are possible publishing costs, which remain constant throughout the search process.

Although the problem of optimizing search sequences has been extensively analyzed in other contexts [3, 7, 8, 10], the results are not applicable to the problem analyzed here, as discussed in Section 4. One important result in this paper is that the thresholds in the optimal search sequence are characterized by a common probabilistic property. This is the because of the way in which the costs of potentially terminating a search are balanced with the expected costs of continuing the search. This result is important because it enables the extraction of a distribution-independent solution. This way, one generic solution can be used to extract the optimal strategy for a large set of problems that only differ in the distribution with which the agent values are associated. In addition, distribution-independent observations of the properties of the optimal sequence can be made, such as expected costs and expected number of rounds. In particular, it is proven that the optimal sequence of thresholds contains a single or an infinite number of thresholds. This is significant, since these types of properties are often analyzed for only a few cases in other problems.

In the next two sections, we formally introduce and analyze the optimal increasing threshold search for the class of environments described above. In Section 4, we illustrate the properties of the optimal strategy with different values of the problem parameters. The benefits of using the optimal strategy are illustrated by comparing it to adaptations of three well studied expanding ring strategies [3, 9, 19] to the problem considered in this paper. In Section 5, we extend the analysis to the case where a group of the best-valued agents needs to be found. This, too, has many real-life applications; for example, extending the applications discussed above, a government agency may need to choose a few of the least costly contractors for a project (if it is too big for any single contractor), or an emergency event may require the arrival of several volunteers to a specific location. We show that a similar method of analysis can be applied to this case.

In Section 6, we show how increasing threshold search is useful for more than just finding the best-valued agent. It is also applicable to *economic search* [2, 27, 30], in which the searcher is not necessarily constrained to finding the best-valued agent, but rather attempts to optimize a function that integrates both search costs and the value of the agent ultimately found. Search theory is an important research domain, flourishing in many disciplines, and best known perhaps for its applications to labor markets, marriage markets, monetary economics, and information theory [26, 33, 44]. By finding the lowest valued agent with a minimal search cost, increasing threshold search potentially achieves this goal, and in many settings can lead to a better overall performance than traditional economic search methods from the economic search point of view. We also show how economic search strategies can be combined with threshold-based searches to further reduce overall costs. We conclude the paper with a review of the relevant literature in Section 7 and a discussion of the results in Section 8. Throughout the paper, we illustrate the properties of the various search techniques using data from a synthetic environment.

2. Model Formulation

This section formally defines the environment in which an increasing threshold search is conducted, the process of conducting an increasing threshold search, and the parameters that define its cost. A summary of the complete set of notations presented in this and all other sections can be found in the appendix.

We consider an agent searching in an environment where N other agents, applicable to its search, can be found. Each of the N agents is characterized by its value to the searcher. Each agent knows its value from the beginning of the search and does not need to consume resources in order to calculate it. As in most search-related models, the values are assumed to be associated with a common continuous distribution described by a PDF f(x) and a CDF F(x), defined over the interval $[x_{min}, x_{max}]$ [1, 2, 7, 11, 29, 30, 31, 40]. The searcher agent is assumed to be ignorant of the value associated with each of the N agents, but acquainted with the overall distribution of values. The searcher is interested in finding the agent associated with the *best* value, which, depending on the application, is either the minimum or the maximum value. Without loss of generality, we assume that the best-valued agent is the one associated with the minimum value.

In order for the searcher to refine the population of agents whose values it obtains, it can publish a maximum threshold r on the agents' values, denoted a *reservation value*, requesting to communicate only with agents that comply with that threshold. If at least one agent complies with r and communicates its value to the searcher, the search process terminates. Otherwise, the searcher sets a new reservation value r' > r and repeats the process. This continues until at least one agent replies, after which the agent associated

Application	Reservation values	Fixed cost (α)	Variable cost $(\beta(j))$
Contract	Bidding scale values	Cost of issuing a	Resources to prepare, place,
bidding		call for bids	and evaluate each bid
Emergency	Distance from event	Disruption to vol-	Time to answer each call
response		unteers and wait	
		time for responses	
Sensor	dB scale values	Energy to broad-	Energy to transmit and pro-
network		cast the request	cess all responses

Table 1: Mapping of sample applications to the increasing threshold search model

with the minimum value is chosen. No constraints are placed on the number of rounds. Accordingly, a strategy S is either a finite sequence $[r_1, \ldots, r_M]$ $(x_{min} < r_i < r_{i+1} < r_M = x_{max}, \forall i < M)$ or an infinite sequence $[r_1, r_2, \ldots]$ $(x_{min} < r_i < r_{i+1} < x_{max}, \forall i)$, where r_i denotes the reservation value to be used in the i^{th} search round. Note that $r_M = r_{max}$ in all finite sequences in order to guarantee search completeness, since it necessarily results in the retrieval of all agent values.

The process of initiating a search round and publishing the next reservation value is associated with a cost α whose value is fixed (e.g., the cost of issuing a new call for bids or the cost of broadcasting a message). Note that this cost may actually be a function of N, but since N remains constant during the search, we can also consider this cost to be constant. Obtaining the actual value of an agent also incurs some cost. In its most general form, the cost of simultaneously obtaining the values of j other agents is $\beta(j)$ (where $\beta(0) = 0$ and $\beta(j)$ is strictly increasing in j) [4, 15, 32]. The overall cost of a search round i is thus $\alpha + \beta(j)$, where j is the number of agents that comply with r_i . The expected accumulated cost of finding the best-valued agent when using strategy S is denoted V(S). The searcher's goal is therefore to derive a strategy S^* that minimizes V(S).

Table 1 maps the problems described in the introduction to the general model introduced in this section.

3. Analysis

In this section, we formulate the expected cost of using a finite search sequence; show how to derive the optimal sequence when the thresholds can only be selected from a finite set of values; prove that the optimal sequence is an infinite sequence of values with a common probabilistic property when the thresholds can be selected from a continuous range of values; show how



Figure 1: Example distribution

to derive such an infinite sequence; and prove some basic properties about the optimal solution.

Consider a searcher agent using a strategy $S = [r_1, \ldots, r_M = x_{max}]$. If the agent has to start the i^{th} search round, then there is necessarily no agent whose value is below r_{i-1} . The *a priori* probability of such a scenario is $(1 - F(r_{i-1}))^N$. Furthermore, upon reaching the i^{th} round, the searcher agent can update its beliefs concerning the PDF and CDF of the values of the Nagents, as it knows that these are necessarily in the interval $(r_{i-1}, x_{max}]$. The PDF of the agents' values after publishing r_{i-1} , denoted $f(x|r_{i-1})$ $(0 < i \leq M)$, can thus be calculated as $(x_{min} \leq x \leq x_{max})$:

$$f(x|r_{i-1}) = \begin{cases} f(x) & i = 1\\ \frac{f(x)}{1 - F(r_{i-1})} & i > 1 \land x > r_{i-1}\\ 0 & i > 1 \land x \le r_{i-1} \end{cases}$$
(1)

Similarly, the CDF of any of the agents' values after publishing r_{i-1} , denoted $F(x|r_{i-1})$ $(0 < i \leq M)$, can be calculated as $(x_{min} \leq x \leq x_{max})$:

$$F(x|r_{i-1}) = \begin{cases} F(x) & i = 1\\ \frac{F(x) - F(r_{i-1})}{1 - F(r_{i-1})} & i > 1 \land x > r_{i-1}\\ 0 & i > 1 \land x \le r_{i-1} \end{cases}$$
(2)

See Figure 1 for an example PDF and CDF. The expected cost of the i^{th}

round is thus:

$$\alpha + \sum_{j=1}^{N} \beta(j) \binom{N}{j} F(r_i | r_{i-1})^j (1 - F(r_i | r_{i-1}))^{N-j}$$
(3)

as it takes into account the cost of initiating the new search round and the expected cost of obtaining any possible number of agent values j ($0 < j \le N$) in round i. This latter cost is the sum, for all j, of the cost $\beta(j)$ of obtaining j agent values multiplied by all $\binom{N}{j}$ combinations of j agents whose values can be obtained, multiplied by the probability $F(r_i|r_{i-1})^j(1-F(r_i|r_{i-1}))^{N-j}$ of obtaining exactly j agent values in round i if none were found in previous rounds.

The probability of starting round i is $(1 - F(r_{i-1}))^N$, which is the probability that no agent values are below the previous reservation value r_{i-1} . The expected cost of using strategy S is thus the sum of the expected cost of each of the M search rounds weighted by the probability of reaching that round:

$$V(S) = \sum_{i=1}^{M} \left(\alpha + \sum_{j=1}^{N} \beta(j) \binom{N}{j} F(r_i | r_{i-1})^j (1 - F(r_i | r_{i-1}))^{N-j} \right) (1 - F(r_{i-1}))^N$$
(4)

The probability of starting the i^{th} search round can alternatively be formulated as the probability that no values were obtained in each of the i-1 previous rounds, expressed as $\prod_{j=1}^{i-1} (1 - F(r_j|r_{j-1}))^N$. Therefore, (4) transforms into:

$$V(S) = \sum_{i=1}^{M} \left(\alpha + \sum_{j=1}^{N} \beta(j) {N \choose j} F(r_i | r_{i-1})^j (1 - F(r_i | r_{i-1}))^{N-j} \right)$$

$$\cdot \prod_{j=1}^{i-1} (1 - F(r_j | r_{j-1}))^N$$
(5)

Note that the probability that the search terminates in round *i* can be calculated as the probability that all agent values are greater than r_{i-1} , $(1 - F(r_{i-1}))^N$, minus the probability that all agent values are greater than r_i ,

 $(1 - F(r_i))^N$. Thus, the expected number of search rounds is

$$\sum_{i=1}^{M} i((1 - F(r_{i-1}))^N - (1 - F(r_i))^N)$$

$$= \sum_{i=1}^{M} (1 - F(r_{i-1}))^N - M(1 - F(r_M))^N = \sum_{i=1}^{M} (1 - F(r_{i-1}))^N$$
(6)

in which $F(r_0) = 0$ and, by definition, $F(r_M) = F(r_{max}) = 1$. Accordingly, (4) can be understood differently by factoring the term $(1 - F(r_{i-1}))^N$ into the remainder of the expression, resulting in the equation

$$V(S) = \alpha \sum_{i=1}^{M} (1 - F(r_{i-1}))^{N} + \sum_{j=1}^{N} \beta(j) {N \choose j} \sum_{i=1}^{M} (F(r_{i}) - F(r_{i-1}))^{j} (1 - F(r_{i}))^{N-j}$$
(7)

Here, the first term is the expected contribution of the fixed cost of publishing the threshold to the total search cost, which is α times the expected number of rounds. The second term is the expected cost of obtaining the values of all agents that comply with the last published reservation value. This is calculated by summing the cost of obtaining j values, over $1 \leq j \leq N$, times the probability that j and only j agents comply with the reservation value published in the last search round.

For the specific case in which the reservation values are chosen from a finite set of L values $\{x_1, \ldots, x_L\}$, $x_{min} < x_i < x_{i+1} < x_L = x_{max}$ for all i, the optimal strategy can be derived with the following dynamic programming formulation in $O(L^2N)$ time:

$$C(L) = 0$$

$$C(l) = \min_{l+1 \le i \le L} \left\{ \alpha + \sum_{j=1}^{N} \beta(j) \binom{N}{j} F(x_i | x_l)^j (1 - F(x_i | x_l))^{N-j} + C(i) (1 - F(x_i | x_l))^N \right\} \quad \forall 0 \le l < L \quad (8)$$

Here, C(l) is the cost of continuing the search if a search using reservation value x_l failed to obtain any applicable agent values. C(l) is solved for

using backwards induction, each time determining what would be the best reservation value larger than x_l to use if x_l is part of the search sequence. The optimal sequence and its expected cost is determined by C(0).

For the general case in which the interval $[x_{min}, x_{max}]$ is continuous, the optimal search strategy must be derived using a different methodology since, as we prove in Theorem 1, the optimal search sequence is either a single search round in which the values of all agents are obtained or an infinite sequence of reservation values.

Theorem 1. The sequence of reservation values that minimizes the expected cost of an iterative threshold search is either $[r_1 = x_{max}]$ or an infinite sequence $[r_1, r_2, \ldots], r_i < r_{i+1} < x_{max}, \forall i \ge 1$.

Proof. Assume that the finite sequence $S_1 = [r_1, \ldots, r_M]$ is the sequence that minimizes the expected cost. We will show that the assumption $V(S_1) < V([x_{max}])$ leads to a contradiction, implying that the sequence that minimizes the expected cost is either $[x_{max}]$ or infinite in length. We use $S_i = [r_i, \ldots, r_M]$ to denote the optimal strategy to be used from round i $(1 < i \leq M)$ onwards if no agent is found in round i - 1 and denote its expected cost from that point on by $V^{(r_{i-1})}(S_i)$. Using S_i , we construct an alternative strategy $S'_i =$ $[r'_i, \ldots, r'_M]$ to be applied from the first round, where $F(r'_j) = F(r_j|r_{i-1}) = P'_j$, $\forall i \leq j \leq M$ and for some P'_j (Figure 2(a)). The expected cost of the new strategy S'_i is denoted $V(S'_i)$. Note that

$$F(r'_{j}|r'_{j-1}) = \frac{F(r_{j}|r_{i-1}) - F(r_{j-1}|r_{i-1})}{1 - F(r_{j-1}|r_{i-1})} = \frac{\frac{F(r_{j}) - F(r_{i-1})}{1 - F(r_{i-1})} - \frac{F(r_{j-1}) - F(r_{i-1})}{1 - F(r_{i-1})}}{\frac{1 - F(r_{j-1})}{1 - F(r_{i-1})}}$$
$$= \frac{F(r_{j}) - F(r_{j-1})}{1 - F(r_{j-1})} = F(r_{j}|r_{j-1})$$
(9)

By substituting $F(r_i|r_{i-1})$ with $F(r'_i|r'_{i-1})$ in (5), we obtain $V(S'_i) = V^{(r_{i-1})}(S_i)$. Since S_1 is the optimal strategy, $V(S_1) \leq V(S'_i)$.

Now consider a new strategy $S''_i = [r''_1, \ldots, r''_M]$ to be applied from round ionwards when r_{i-1} was applied in round i-1, where $F(r''_j|r_{i-1}) = F(r_j) = P''_j$, $\forall 1 \leq j \leq M$ (Figure 2(b)). According to this, r''_j is the reservation value to be used in round i + j - 1 when r_{i-1} was used in round i - 1. We denote the expected cost of S''_i from round i onwards $V^{(r_{i-1})}(S''_i)$. Note that $F(r''_j|r''_{j-1}) = F(r_j|r_{j-1}) = P''_j$, as above. According to (5), we obtain



Figure 2: Alternative strategies when j = 2 in Theorem 1

 $V^{(r_{i-1})}(S''_i) = V(S_1)$. Since S_i is the optimal strategy from round *i* onwards, $V^{(r_{i-1})}(S_i) \leq V^{(r_{i-1})}(S''_i)$, resulting in $V(S_1) \leq V(S'_i) = V^{(r_{i-1})}(S_i) \leq V^{(r_{i-1})}(S''_i) = V(S_1)$, which can hold only if $V^{(r_{i-1})}(S_i) = V(S_1)$.

This logic applies to all search rounds $2 \leq i \leq M$. In particular, this implies that $V(S_1) = V^{(r_{M-1})}(S_M)$. However, we necessarily find all agents in the last (M^{th}) round, since $r_M = x_{max}$; thus, $V(S_M) = \alpha + \beta(N)$. Therefore, we obtain $V(S_1) = \alpha + \beta(N) = V([x_{max}])$, which contradicts the assumption that S_1 minimizes the expected cost, unless $S_1 = [x_{max}]$.

The argument that $V^{(r_{j-1})}(S_j) = V(S_1)$ holds even for infinite sequences, establishing the following corollary.

Corollary 1. When the sequence of reservation values that minimizes the expected cost of an iterative threshold search is an infinite sequence, the cost of continuing the search from any round onwards is a constant value.

By the same argument, it is clear that the reservation values can be set such that $F(r_i|r_{i-1}) = P$, $\forall i > 1$ and some value P. The next theorem establishes this is necessarily the case.

Theorem 2. When the optimal sequence of reservation values is the infinite sequence $[r_1, r_2, \ldots]$, $x_{min} < r_i < x_{max}$, $\forall i > 0$, the values of r_i are necessarily set such that $F(r_{i+1}|r_i) = F(r_1) = P$, for some 0 < P < 1 and $\forall i > 0$.

Proof. First note that we can formulate the expected cost from any round i

onwards as follows:

$$V(P_i) = \alpha + \sum_{j=1}^{N} \beta(j) \binom{N}{j} P_i^j (1 - P_i)^{N-j} + (1 - P_i)^N V^*$$

where V^* is the expected cost of continuing the search, which is constant according to Corollary 1. We prove that only one value of P_i minimizes this formula by showing that the first derivative with respect to P_i is monotonically increasing, and that this value is the same for all *i*.

Taking the first derivative, we have:

$$\begin{split} &\sum_{j=1}^{N} \beta(j) \binom{N}{j} \left[j P_i^{j-1} (1-P_i)^{N-j} - (N-j) P_i^j (1-P_i)^{N-j-1} \right] \\ &- N(1-P_i)^{N-1} V^* = \sum_{j=1}^{N} \beta(j) \binom{N}{j} j P_i^{j-1} (1-P_i)^{N-j} \\ &- \sum_{j=1}^{N} \beta(j) \binom{N}{j} (N-j) P_i^j (1-P_i)^{N-j-1} - N(1-P_i)^{N-1} V^* \\ &= \sum_{j=2}^{N} N\beta(j) \frac{(N-1)!}{(j-1)!(N-j)!} P_i^{j-1} (1-P_i)^{N-j} \\ &- \sum_{j=1}^{N-1} N\beta(j) \frac{(N-1)!}{j!(N-1-j)!} P_i^j (1-P_i)^{N-j-1} - N(1-P_i)^{N-1} (V^* - \beta(1)) \\ &= \sum_{k=1}^{N-1} N\beta(k+1) \frac{(N-1)!}{k!(N-1-k)!} P_i^k (1-P_i)^{N-k-1} \\ &- \sum_{j=1}^{N-1} N\beta(j) \frac{(N-1)!}{j!(N-1-j)!} P_i^j (1-P_i)^{N-j-1} - N(1-P_i)^{N-1} (V^* - \beta(1)) \\ &= \sum_{j=1}^{N-1} N(\beta(j+1) - \beta(j)) \binom{N-1}{j} P_i^j (1-P_i)^{N-j-1} - N(1-P_i)^{N-1} (V^* - \beta(1)) \\ &= N(1-P_i)^{N-1} \left[\sum_{j=1}^{N-1} N(\beta(j+1) - \beta(j)) \binom{N-1}{j} \binom{N-1}{j} \binom{N-1}{j} \binom{N-1}{j} (V^* - \beta(1)) \right] \end{split}$$

The transition from the third line to the fourth line is made by setting k = i - 1. The first term in the big square brackets of the last line is a polynomial in $s = \frac{P_i}{1-P_i}$ of degree N - 1, with vanishing constant term and other coefficients positive. Thus as P_i increases from 0 to 1, causing $s = \frac{P_i}{1-P_i}$ to increase monotonically from 0 to ∞ , this first term increases monotonically from 0 to ∞ . Thus for small P_i , the derivative is less than 0; there is a unique value of P_i , P_i^* , at which the derivative equals 0; and for $P_i > P_i^*$, the derivative is greater than 0. Since all other terms are the same for all values of i, P_i^* is also the same for all values of i.

With these two theorems, it is now possible to show how to derive the optimal strategy. A unique result of the above analysis is that the optimal increasing threshold search strategy can be expressed as a single value $0 < P \leq 1$, denoted the *reservation probability*. Based on this observation, the optimal search strategy can be derived by calculating P and then setting each reservation value r_i such that $F(r_1) = F(r_i|r_{i-1}) = P$, $\forall i > 1$. First we show how to derive the reservation probability P, and then we show how to calculate each r_i . Since the optimal sequence is infinite and the expected cost from each round onwards is stationary, the expected cost of using P can be expressed with the following equation:

$$V(P) = \alpha + \sum_{j=1}^{N} \beta(j) {\binom{N}{j}} P^{j} (1-P)^{N-j} + (1-P)^{N} V(P)$$
(10)

Here, the first term is the fixed cost per round, the second term is the expected cost of obtaining any values, and the last term is the expected cost of continuing the search if necessary. Consequently:

$$V(P) = \frac{\alpha + \sum_{j=1}^{N} \beta(j) {N \choose j} P^{j} (1-P)^{N-j}}{1 - (1-P)^{N}}$$
(11)

The value $P = P^*$ that minimizes V(P) in (11) is the optimal reservation probability. If P^* cannot be solved for directly, then it can be solved for by numerical approximation. Based on (2), each r_i corresponding to P^* can be calculated by solving for r_i in the equation $P^* = \frac{F(r_i) - F(r_{i-1})}{1 - F(r_{i-1})}$, which yields the equation

$$r_i = F^{-1}(P^*(1 - F(r_{i-1})) + F(r_{i-1}))$$
(12)

Just as in (7), (11) can be decomposed into two parts: $\frac{\alpha}{1-(1-P)^N}$ and $\frac{\sum_{j=1}^N \beta(j) {\binom{N}{j}} P^j (1-P)^{N-j}}{1-(1-P)^N}$, respectively representing the expected contribution of the fixed cost to the total search cost and the expected cost of obtaining the agent values. The expected number of search rounds is $\frac{1}{1-(1-P)^N}$. The probability of obtaining j agent values is $\frac{{\binom{N}{j}} P^j (1-P)^{N-j}}{1-(1-P)^N}$, since this becomes a Bernoulli sampling process with a success probability of $1 - (1-P)^N$.

An important result of the above analysis is that P is distribution independent. This means that the searcher only needs to solve for P once for fixed values of N, α , and $\beta(j)$. This also means that certain properties of the optimal strategy, such as the expected cost and the expected number of rounds, do not depend on the probability distribution. P is distribution independent because of the infinite nature of the optimal sequence. P takes the place of the cumulative conditional probability distribution in (5). It is set to a value that balances the cost obtaining several agent values with the expected cost of continuing a search. Since the optimal strategy is an infinite sequence, the searcher is faced with the same problem each round and therefore attempts to make the same tradeoff each round. The infinite nature of the optimal sequence is due to the existence of unlimited opportunities to effect this tradeoff and not incur the full cost of obtaining all agent values.

We now analyze several properties of the optimal solution. Proposition 1 establishes when the optimal strategy is to set $r_1 = x_{max}$; specifically, when $\beta(i) = 0$.

Proposition 1. The optimal strategy is $[x_{max}]$ if and only if $\beta(i) = 0$.

Proof. First we prove that $[x_{max}]$ is the optimal strategy if $\beta(i) = 0$. Substituting 0 for $\beta(j)$ in (11) obtains $V(P) = \frac{\alpha}{1 - (1 - P)^N}$. Setting P = 1 minimizes this expression, which is equivalent to using x_{max} in the first search round.

Next we prove that it is necessarily the case that $\beta(i) = 0$ if the optimal strategy is $[x_{max}]$. We show that there always exists a value P for which the expected cost of using the infinite sequence is less than the expected cost of using a single search when $\beta(j) > 0$. Formally, using (11), we need to show the following inequality holds for some value of 0 < P < 1:

$$\frac{\alpha + \sum_{j=1}^{N} \beta(j) {N \choose j} P^{j} (1-P)^{N-j}}{1 - (1-P)^{N}} < \alpha + \beta(N)$$

We multiply both sides by $1 - (1 - P)^N$, yielding:

$$\alpha(1-P)^{N} < \beta(N)(1-(1-P)^{N}) - \sum_{j=1}^{N} \beta(j) \binom{N}{j} P^{j}(1-P)^{N-j}$$

Notice that $\beta(N)$ can be written as $\beta(N)(1 - P + P)^N = \beta(N)(1 - P)^N + \sum_{j=1}^N \beta(N) {N \choose j} P^j (1 - P)^{N-j}$. Therefore:

$$\alpha (1-P)^{N} < \sum_{j=1}^{N} \beta(N) \binom{N}{j} P^{j} (1-P)^{N-j} - \sum_{j=1}^{N} \beta(j) \binom{N}{j} P^{j} (1-P)^{N-j}$$
$$= N(\beta(N) - \beta(1)) P(1-P)^{N-1} + \sum_{j=2}^{N} (\beta(N) - \beta(i)) \binom{N}{j} P^{j} (1-P)^{N-j}$$

Notice that both terms on the right hand side are positive for any P. Therefore, if we prove that there exists a value P for which $\alpha(1-P)^N = N(\beta(N) - \beta(1))P(1-P)^{N-1}$ then the inequality holds. If $\alpha(1-P)^N = N(\beta(N) - \beta(1))P(1-P)^{N-1}$, we have:

$$\alpha(1-P) = N(\beta(N) - \beta(1))P$$
$$\frac{\alpha}{N(\beta(N) - \beta(1))} = \frac{P}{1-P}$$

The left hand side is positive and finite. $\frac{P}{1-P} = 0$ when P = 0 and approaches infinity when $P \to 1$. Therefore, there must exist such a P.

In contrast to Proposition 1, Proposition 2 states that the optimal strategy when there is no cost for initiating a new search round ($\alpha = 0$) is to increment the reservation value by the smallest amount possible, such that the expected number of values obtained is minimized.

Proposition 2. If $\alpha = 0$, the optimal strategy is to use $P \rightarrow 0$, i.e., to increment the reservation value by $\epsilon \rightarrow 0$ each round.

Proof. Notice that $\lim_{P\to 0} \frac{\beta(j)\binom{N}{j}P^{j}(1-P)^{N-j}}{1-(1-P)^{N}}$ equals $\beta(j)$ for j = 1 and 0 for j > 1, according to L'Hôpital's rule. Substituting $\alpha = 0$ in (11), we obtain $\lim_{P\to 0} V(P) = \beta(1)$. Since $\beta(i)$ is increasing, according to the model assumption, the result obtained is the minimum expected cost possible. \Box

Finally, Proposition 3 highlights how the tradeoff between obtaining a certain number of agents and continuing a search is affected by P in the case where the cost of obtaining the values of j agents is linear in j.

Proposition 3. When $\beta(j)$ is linear in j, the reservation probability that minimizes V(P), $P = P^*$, satisfies $c = (1 - P^*)^{N-1}V(P^*)$.

Proof. First, we simplify (11) by substituting $\beta(j)$ with cj. The expression $\sum_{j=1}^{N} j\binom{N}{j} P^j (1-P)^{N-j}$ in the numerator of (11) is the mean of a binomially distributed random variable, which equals NP. Therefore,

$$V(P) = \frac{\alpha + cNP}{1 - (1 - P)^N}$$
(13)

Differentiating (13) with respect to P and setting it to zero obtains:

$$\frac{cN(1-(1-P)^N) - N(1-P)^{N-1}(\alpha+cNP)}{(1-(1-P)^N)^2} = 0$$
(14)

Notice that $V(P)(1-(1-P)^N) = \alpha + cNP$ according to (13). Substituting the latter expression into (14), we observe that the value P^* which satisfies the equation is given by $cN(1-(1-P)^N) - N(1-P)^{N-1}V(P)(1-(1-P)^N) = 0$. Solving for c gets $c = (1-P)^{N-1}V(P)$.

The explanation of Proposition 3 requires understanding the tradeoff associated with any increase in P. By increasing P, the chance of finding each of the agents increases. Each agent found due to the increased chance will incur a cost c. The benefit is that if the agent found due to the increase is the only agent found in that round, then the increase actually saves the expected cost V(P) associated with continuing the search. The probability that the latter case holds is $(1-P)^{N-1}$ (i.e., when all other agents are characterized with a value above the reservation value set using (12)). Otherwise, the search just ends. Since the incurred cost c is fixed and the expected benefit $(1-P)^{N-1}V(P)$ decreases as P increases, the optimal P value satisfies $(1-P)^{N-1}V(P) = c$, i.e., when the additional benefit due to the potential saving is offset by the cost incurred by finding that agent.

4. Comparative Illustration

In this section, we illustrate the behavior and performance of the optimal increasing threshold search strategy derived in the previous section. We show the effect of N, α , and β on the optimal strategy and its associated cost. Additionally, we demonstrate the improvement achieved by the optimal search strategy over several base strategies. The magnitude of improvement in some instances illustrates the importance of choosing the right strategy, while other instances demonstrate how these base strategies can sometimes be close to optimal. Since this problem has not been well addressed in the literature, we adapt three well-studied expanding ring search strategies to our problem [3, 9, 19]. Expanding ring search is used to find routes in ad hoc networking and to locate files in peer-to-peer networking. In this method, the searcher assigns a query a time-to-live (TTL) value, which determines the number of hops the query is forwarded. If the search target is not met, the searcher repeats the query with a larger TTL value. The cost structure of expanding ring search is different than the cost structure of the problem addressed in this paper, since the cost per round increases with the search extent, and the search extents are typically drawn from a range of discrete values. Still, in the absence of more suitable alternatives, expanding ringbased strategies are a natural basis for comparison to increasing threshold search. In the following paragraphs, we describe the three strategies and then compare their performance in our context.

4.1. Two-step strategy

If an increasing threshold search is to improve search costs, then a twostep strategy $S = [r_1, r_2 = x_{max}]$ alone will provide some improvement [9]. The expected total cost when using this strategy is:

$$V(S) = \alpha + \sum_{j=1}^{N} \beta(j) \binom{N}{j} P^{j} (1-P)^{N-j} + (1-P)^{N} (\alpha + \beta(N))$$
(15)

in which $P = F(r_1)$. The first two terms are the costs associated with the use of the first reservation value r_1 , and the third term is the cost of obtaining all agent values if no values were obtained using r_1 . The optimal two-step strategy is obtained by determining the value P that minimizes (15).

4.2. Fixed increment strategy

A common design of a multi-round expanding ring search strategy is to use a fixed increment between search extents, searching up to some cutoff value before searching the entire search range [19]. The search sequence is thus of the form $\{x_{min} + \delta, x_{min} + 2\delta, \ldots, x_{min} + \mu\delta, x_{max}\}$. The expected cost of using such a sequence can be calculated using (4). The increment and cutoff value pair with the lowest expected cost is selected from all possible pairs.

4.3. California split rule strategy

Another well-studied strategy is the California split rule [3]. According to this strategy, the search extent is doubled each round. We adapt this method to our problem by including a cutoff value above which the reservation value is set to x_{max} , similar to the fixed increment strategy. The search sequence is thus of the form $\{x_{min} + \delta, x_{min} + 2\delta, \ldots, x_{min} + 2^{\mu}\delta, x_{max}\}$. The expected cost of using such a sequence can be calculated using (4). The increment and cutoff value pair with the lowest expected cost is selected from all possible pairs.

4.4. Evaluation

Since the domination of the reservation probability based strategy over the different expanding ring based strategies is unquestionable due to its proven optimality, the goal of this evaluation is merely to demonstrate the effect of different parameters on performance. For this purpose, a synthetic environment is an ideal testbed. In the environment used, the agents values are associated with a truncated normal distribution of values [18], with $\mu =$ 0.5 and $\sigma = 0.125$, over the interval (0,1); that is, $f(x) = \frac{8\phi(\frac{x-0.5}{0.125})}{\phi(\frac{1-0.5}{0.125})-\phi(\frac{-0.5}{0.125})}$ for $0 \le x \le 1$, and 0 otherwise. The cost of simultaneously obtaining the values of j other agents is set to $\beta = 0.01w(j)$, where $w(j) \in \{log(j), j, (j)^2\}$ (denoted "log cost", "linear cost" and "square cost", respectively). The use of the three different cost functions enables testing different rates of increase in the marginal cost of obtaining the value of an additional agent. The cost of initiating a search round and publishing the reservation value is set to $\alpha \in \{0.001, 0.01, 0.1, 1\}$ in order to capture the effects of the magnitude of the ratio $\alpha/\beta(j)$. Table 2 summarizes these parameters, and Figure 3 illustrates the truncated normal distribution function f(x) and the set of cost functions assigned to $\beta(j)$.

Note again that the optimal reservation probability and the expected cost of the optimal strategy are distribution independent; only the actual reservation values used are determined by the probability function. We illustrate how the reservation values are derived using the settings N = 20, $\alpha = 0.01$, $\beta(j) = 0.01i$, and f(x) as defined above. Substituting 0.01



Figure 3: Truncated normal distribution function, f(x), with $\mu = 0.5$ and $\sigma = 0.125$ over the interval (0,1) (left) and the three variants of the cost function $\beta(j)$.

Table 2: Parameter settings in the synthetic environment

parameter	value
f(x)	$\frac{8\phi(\frac{x-0.5}{0.125})}{\phi(\frac{1-0.5}{0.125})-\phi(\frac{-0.5}{0.125})}$
α	0.001, 0.01, 0.1, 1
$\beta(j)$	$0.01w(j), w(j) \in \{\log(j), j, (j)^2\}$

for c and 20 for N in (13) and using numerical approximation, we find that the value of P that minimizes (13) is approximately 0.058. In this case, $r_1 = F^{-1}(0.058) \approx 0.304$. Substituting 0.058 for $F(r_1)$ in (12) yields $r_2 = F^{-1}(0.058(1 - 0.058)) + 0.058)) \approx 0.348$. Then substituting 0.112 for $F(r_2)$ into (12) yields $r_3 = F^{-1}(0.058(1 - 0.112)) + 0.112) \approx 0.377$. This continues indefinitely.

Figure 4 shows the expected cost, the expected number of search rounds, and the 99.9th percentile of the number of search rounds required to find the best-valued agent as a function of the number of agents N, for the different values of α and function assignments to $\beta(j)$. Note that the scale of the yaxis in the first set of charts varies with α in order to improve visualization. The figure also shows the reservation values to be used in each search round when N = 20. As expected, α has a substantial effect on the search strategy and expected cost. The reservation value increases as a function of α , since the tradeoff associated with each reservation value accounts for a greater cost of continuing the search if no agent is found. Although the expected number of search rounds decreases as a consequence, the expected cost still increases, due to the greater cost per round. For large values of α and all $\beta(j)$ cost functions, the reservation probability is set such that the expected number



Figure 4: Increasing threshold search for the settings in Table 2: expected cost (top row); expected number of rounds (second row); 99.9th percentile of the number of rounds (third row); reservation values for N = 20 (bottom row)

of search rounds is close to 1 due to the substantial cost of initiating a new search round.

Amongst the three cost functions, the expected cost is minimal for the log cost function and maximal for the square cost function, in direct correlation with the order of growth of the three functions. The inverse relationship holds for the reservation values because of the higher cost of finding more than one agent associated with the higher order functions. As a result, the expected cost increases due to the increase in the expected number of rounds.

Another important observation from Figure 4 is that the expected cost increases with the number of agents N. Although the expected number of rounds decreases with N for all reservation probabilities, the expected number of agents found also increases. The increase in cost associated with N indicates that, in these settings, the reduction in the number of search

rounds does not offset the cost of finding additional agents.

From the graphs of the expected number of search rounds (second row), we observe that the expected number of search rounds is below five for all combinations of parameters. This supports the applicability of the optimal strategy: Despite the fact that it is an infinite sequence, the search time in practice is comparable to search with competitive finite sequences. The third row complements this observation. Each curve shows the upper bound on the number of search rounds required 99.9% of the time for each of these settings. We see that even when the fixed cost per round is low ($\alpha = 0.001$) and the cost function is high $\beta(j) = 0.001j^2$, this number is less than thirty-one.

Figure 5 shows the expected cost of using the (optimal) reservation probability based strategy and the three expanding ring based strategies as a function of N, for different values of α and function assignments to $\beta(j)$. The graphs illustrate the domination of the optimal strategy and the dominance relationships amongst the remaining strategies. The two step strategy never dominates the fixed increment and California split rule strategies, since it is merely a specific instance of both. Although the fixed increment strategy dominates the California split rule strategy in these settings, it is possible to construct specific settings in which the latter dominates the first. It is interesting that while the expected cost of the optimal strategy always increases with the number of agents, the expected costs of the fixed increment and California split rule strategies decrease in some settings. This is because the increased probability of finding an agent with each reservation value sometimes reduces the inefficiency imposed by the constraints on the patterns of increase in the reservation value. Another interesting observation is that the expected costs of the fixed increment and California split rule strategies converge as α increases.

The upper row of graphs in Figure 5 corresponds to the logarithmic $\beta(j)$ function. All four strategies use higher reservation values as α increases, as the cost of finding more than one agent in a search round becomes negligible in comparison to the cost of requiring an additional search round. Consequently, the three alternative strategies quickly converge to the same expected cost, as their sequences in those cases contain only two reservation values. The optimally infinite strategy, however, produces observably better results, even for $\alpha = 0.1$. This property, that the California split and fixed increment strategies converge more quickly to the performance of the worst strategy (two step) with respect to α than the optimal solution, occurs for the linear cost (middle row) and the square cost (bottom row) functions as



Figure 5: Performance comparison of the optimal strategy and the three expanding ring based strategies

well, for larger values of α . In particular, it is interesting to see that for the square cost function, the fixed increment strategy, which performs quite close to the optimal strategy for small α values, becomes substantially worse as α increases. Overall, in these settings, the improvement of the optimal strategy over the other three methods increases as the order of growth of $\beta(j)$ is increased. Since the optimal search sequence is not constrained to any pattern, it makes a better tradeoff between the cost of finding more than one agent and initiating additional search rounds.

5. Multiple Agent Search

In various settings, search for multiple agents is preferable and even necessary. For example, a government agency seeking bids for a project may prefer to run a second price auction in order to encourage truthful bidding. In such an auction, the agency awards the lowest bidder with the second lowest bid; thus, the agency must obtain the two lowest bids rather than the lowest bid alone. Alternatively, the agency may request a best and final offer from a set of the top bidders after an initial call for bids. In the case of emergency services, several units may be required to attend to an emergency situation. Likewise, a sink in a sensor network may require readings from several sensors to accurately analyze conditions in the field. In all of these examples, the searcher needs to find several best agents. This section extends the analysis of the single agent threshold search to a multiple best-valued agents search.

We apply the analysis methodology described in Section 3 to the model described in Section 2, except that the searcher is interested in finding the K best agents. Without loss of generality, the K best agents are those associated with the K lowest values. The searcher continues its search as long as only k < K agents have been found so far. The searcher's state at the beginning of any round is therefore denoted by the pair (r, k), in which r is the last reservation value the searcher used and k is the number of agents found so far. Accordingly, the initial state is $(x_{min}, 0)$.

A search strategy is defined by a function $S: (r, k) \to r', r' > r$, in which r' is the reservation value to use in a round with initial state (r, k). Since the reservation value r' depends on the number of agents found so far, a search strategy cannot be defined *a priori* by a sequence of reservation values as in the case of a single agent search. Instead, it can be defined by a decision tree, in which each node represents a state (r, k) and has child nodes for all possible $k \leq k' < K$ total number of agents found by the end of that round. Since many states may be repeated throughout the tree, it can be represented more compactly as a directed graph. The searcher terminates its search when at least K - k agents are found in a round with state (r, k)or when $S(r,k) = x_{max}$, in which case all K agents are necessarily found. It is realistically assumed that agents only reply the first time they comply with a reservation value. Therefore, the searcher only searches among the remaining N-k agents that did not comply with any reservation value less than or equal to r. It is assumed that the publishing cost is constant (α) for all k, as opposed to assuming that the publishing cost is a function of k. The ensuing analysis is still valid even if these last two assumptions are not made by substituting N for N-k and $\alpha(k)$ for α . Given these assumptions, the expected cost of continuing the search from state (r, k) when using strategy

S, denoted $V^{(r,k)}(S)$, is defined recursively as follows:

$$V^{(r,k)}(S) = \alpha + \sum_{j=0}^{N-k} \left(\beta(j) + V^{(S(r,k),k+j)}(S) \right) \binom{N-k}{j}$$

$$\cdot F(S(r,k)|r)^{j} (1 - F(S(r,k)|r))^{N-k-j}$$
(16)

in which $V^{(r,k)}(S) = 0$ for all $k \ge K$. The total cost of the search is $V^{(x_{min},0)}(S)$.

We again divide the problem into the discrete and continuous cases. In the discrete case, the reservation values can only be selected from a finite set of L values $\{x_1, \ldots, x_L\}$, $x_{min} < x_i < x_{i+1} < x_L = x_{max}$. The optimal strategy can be derived with the following dynamic programming formulation, in which C(l, k) is the cost of continuing the search after using reservation value x_l and after a total of k agents have been found so far. The optimal cost is determined by C(0, 0). The runtime complexity is $O(L^2KN)$.

$$C(L, k) = 0 \qquad \forall 0 \le k \le K$$

$$C(l, K) = 0 \qquad \forall 0 \le l \le L$$

$$C(l, k) = \min_{l+1 \le i \le L} \left\{ \alpha + \sum_{j=0}^{N-k} (\beta(j) + C(i, k+j)) \binom{N-k}{j} \\ \cdot F(x_i | x_l)^j (1 - F(x_i | x_l))^{N-k-j} \right\} \qquad \forall 0 \le k, l < K, L$$
(17)

As the following analysis shows, the increasing threshold search strategy that minimizes the expected search costs is based on a set of reservation probabilities $S = \{P_0, ..., P_{K-1}\}$, where P_k is the reservation probability to be used in any state (r, k), $x_{min} \leq r < x_{max}$, $0 \leq k < K$. The reservation value r' in state (r, k) can be calculated from the equation $r' = F^{-1}(P_k(1 - F(r)) + F(r))$. As in the single agent search, we prove that this is the optimal strategy by first establishing that the optimal strategy to apply as long as only k < K agents have been found is an infinite sequence of reservation values.

Theorem 3. The increasing threshold search strategy that minimizes the expected search costs is to use the reservation value r' in any state (r, k) such that $F(r'|r) = P_k$, for some $0 < P_k \le 1$.

Proof. We prove that the expected cost of continuing the search from any state (r, k) onwards is some value V_k using backwards induction on k. By assuming that this is true for all k' > k, we can prove that the optimal sequence of reservation values to use as long as no new agent values are obtained is the infinite sequence r_1, r_2, \ldots such that $F(r_1|r) = F(r_{i+1}|r_i) = P_k$, for some $0 < P_k \leq 1$. This result is in turn used to prove the inductive step, since the expected cost is independent of the actual sequence of reservation values.

In the base cases, $V_{k'} = 0 \ \forall k' \geq K$ by definition, since the search is terminated in such cases. Consider the finite sequence $S_1 = [r_1, \ldots, r_M]$, $r_i > r \ \forall 1 \leq i \leq M$. The expected cost of using S_1 from state (r, k) onwards is:

$$V^{(r,k)}(S_1) = \sum_{i=1}^{M} \left(\alpha + \sum_{j=1}^{N-k} (\beta(j) + V_{k+j}) \binom{N-k}{j} \right)$$

$$\cdot F(r_i | r_{i-1})^j (1 - F(r_i | r_{i-1}))^{N-k-j} \prod_{j=1}^{i-1} (1 - F(r_j | r_{j-1}))^{N-k} \quad (18)$$

where $F(r_1|r_0) = F(r_1|r)$.

Using the same reasoning as in Theorem 1, we can show that the assumption that any finite sequence minimizes the expected cost leads to a contradiction, proving that the optimal sequence is the infinite sequence r_1, r_2, \ldots . Likewise, we can formulate the expected costs similar to the way we did in Theorem 2 and prove that only one value of P (that is, P_k) minimizes the expected search costs. Finally, just as we did in Equation (11), we can formulate the expected cost of using this reservation probability as:

$$V(P_k) = \frac{\alpha + \sum_{j=1}^{N-k} (\beta(j) + V_{k+j}) {\binom{N-k}{j}} P_k{}^j (1 - P_k)^{N-k-j}}{1 - (1 - P_k)^{N-k}}$$
(19)

This expression is independent of any reservation values used during the search, thus completing the inductive step that shows that the expected cost of continuing the search from any state (r, k) is some constant V_k .

The optimal values of P_k $(0 \le k < K)$ can be calculated using backward induction: Given the optimal P_j for all j > k, calculate P_k $(0 \le k < K)$ with (19). The actual reservation values are derived in a similar manner as in the single agent search. The first reservation value r_0 is calculated with



Figure 6: The reservation value used by the optimal strategy (P_k) and alternative strategy (P) as a function of k, for different values of K and N = 20

 $r_0 = F^{-1}(P_0)$. Then, the reservation value to be used in any state (r, k) is calculated by substituting P_k for P^* in (12).

We illustrate the properties of the optimal search for multiple agents under the settings in Table 2 and compare it to an alternative strategy based on the single agent search. In the alternative strategy, the searcher searches for one agent at a time, using the optimal P from (11) for N-k agents when k agents have already been found. That is, the searcher begins by searching for one agent using P from (11). Upon finding k > 1 agents, the searcher begins a new search for one agent using P from (11), modified to reflect that only N-k agents are left to reply, and that all remaining agent values are above the last reservation value used. This is continued until the total number of values obtained is greater than or equal to K. Figure 6 shows the optimal values of P_k according to (19) and P according to (11) as functions of the number of agents k already found for N = 20; different numbers of agents K that need to be found; $\alpha = 0.1$; and the different $\beta(j)$ cost functions. As can be observed from the figure, P_k decreases as k increases for the log and linear cost functions and all values of K, at a greater rate for the linear cost function. This can be attributed to two conflicting effects of the increase in k on the search: On one hand, as the number of available agents decreases, the expected number of agents found with any reservation value decreases, possibly supporting an increase in the reservation probability. On the other hand, the probability of finding more than K-kagents with each reservation probability increases, supporting a decrease in the reservation probability. Since the latter effect is more influential than the first, the reservation probability decreases as a function of k.

The behavior of the square cost function differs from the other two. The most obvious difference is that the reservation probability is initially small,

increases as the number of agents found increases, then decreases. In particular, when all agents need to be found (K = 20), the reservation probability is set to a value less than 1, even though no unnecessary agents will be found if it is set to 1. This is because the cost function is concave; thus, there is a high penalty of finding a large number of agents at once. A careful analysis of the data reveals that the reservation probabilities in this case are indeed set such that the expected number of agents found is almost the same each round. However, as the number of remaining agents approaches 0, the reservation probability decreases accordingly to avoid finding unnecessary agents.

As expected, the optimal reservation probability P_k for the multi agent search is substantially greater than the optimal reservation probability for a single agent search. In most cases, P_k increases as the total number of agents K that need to be found increases. The square cost function is an exception. As evident from Figure 6, the reservation probability for the case (K = 18, k = 14) is greater than the one for the case (K = 20, k = 14). This is attributed to the nature of the cost function, as discussed above. In particular, the tradeoff between the expected number of rounds and agents found differs for the case of (K = 20, k = 14), since in this case there is no problem of finding more agents than necessary.

Figure 7 shows the percentage by which the optimal strategy reduces the expected cost of the alternative strategy as a function of K, for different cost functions and values of N. The cost reduction highly depends on the cost function $\beta(j)$ and can be quite substantial, up to 20% for the square cost function and 80% for the log cost function. Furthermore, the improvement increases as the total number of agents K that need to be found increases, as the savings by finding several agents in one round increases. This improvement is mild for the square cost function for the reasons discussed above. It is noteworthy that N has only a minor effect on the cost reduction. An increase in N results in a small decrease in the cost reduction, at a decreasing rate.

6. Application to Economic Search

While the goal of increasing threshold search is to find the best-valued agent, the goal of searchers in many applications may be to find a suitable agent while optimizing the process as a whole [4, 20, 32]. For example, consider a buyer agent that is interested in purchasing a product and that obtaining posted prices from seller agents incurs a cost (e.g., communication



Figure 7: Percentage by which the optimal strategy reduces the expected cost of the alternative strategy as a function of K

costs). The buyer agent can purchase the product from any of the seller agents. As the buyer increases the number of sellers with which it communicates, the price it expects to pay for the product decreases, but its overall communication costs increase. Thus, the optimal search strategy is a tradeoff between the marginal saving of each additional price obtained and the cost of obtaining it.

The research domain in which such problems are studied is called *search* theory ([27, 30], and references therein). Within the framework of search theory, two main clusters of search models can be found: (a) the sequential search model and (b) the fixed sample size model. In the sequential search model [27, 36], the searcher obtains a single agent value at a time, allowing multiple search stages. An example of a sequential search is a buyer who checks prices at different stores until the available options are satisfactory, and then returns to the store with the best option. In the fixed sample size model, the searcher obtains a large set of agent values in a single search round [42], and then chooses the agent associated with the best value. This is most applicable when some time constraint limits the searcher to a single search round. For example, when applying to college, one must typically apply to several institutions at the same time and then choose from the best offer after all the applications have been reviewed.

While economic search strategies are inapplicable to our problem, as they do not guarantee finding the best-valued agent, increasing threshold search alone and the extensions to it analyzed in this section are useful contributions to economic search theory. This is because increasing threshold search, whenever applicable, can result in an overall reduced cost, even in comparison to the optimal economic search strategy. In this section, we show how the searching agent can benefit from threshold based searches. We begin by introducing the optimal sequential search and fixed sample strategies as known from search theory. We then show how increasing threshold search can be used as an alternative to sequential search, both on its own and in conjunction with sequential search. Finally, we show how the fixed sample search can be augmented with a threshold-based search to improve performance.

6.1. Optimal economic search strategies

In the sequential search model [25, 27] a searcher is given N possible opportunities $B = \{B_1, ..., B_N\}$ (e.g., to buy a product) out of which she can choose only one. Each opportunity B_i encapsulates a value to the searcher (e.g., expense, reward, utility). While this value is unknown to the searcher, she is acquainted with the probability distribution function f(x) with which all the values are associated. The true value of any opportunity B_i can be obtained by paying a fee, denoted c.

The searcher can continue to obtain the value of any of the opportunities in B, paying each time the cost c. Once the searcher decides to terminate her search (or once she has obtained the value of all opportunities), she collects the minimum value among those revealed up until that time (assuming that she seeks the minimum value). The goal of the searcher is therefore to find the optimal strategy, i.e., a stopping rule that minimizes the expected expense, defined as the value eventually obtained plus the accumulated costs incurred throughout the process.

The optimal search strategy for this model is reservation-value based [30]. Note that the term "reservation value" has a slightly different meaning in this context than in the context of increasing threshold search. The searcher sets a reservation value r (i.e., a threshold) and sequentially obtains the value of different opportunities, incurring a cost c, until revealing a value less than the reservation value (or until the values of all opportunities are obtained). Since all of the opportunities are associated with the same probability distribution function f(x), the searcher can obtain the values in any order. This strategy is preferred in particular when the cost of obtaining the value of j agents is linear or super-linear in j, since the searcher does not benefit from obtaining the value of more than one agent at a time. The optimal reservation value, r, is based on the distribution of agent values and the cost of obtaining a value, c [27, 30, 43]. It is the value by which the searcher is indifferent between terminating the search and obtaining r, and resuming the search.

The optimal reservation value in this case is derived from [30]:

$$c = \int_{x_{min}}^{x_{max}} (r - \min(y, r)) f(y) \, dy = \int_{x_{min}}^{r} F(y) \, dy \tag{20}$$

Search theory focuses merely on the optimal stopping rule and does not place much importance on the expected cost of using this rule. Therefore, to compare increasing threshold search with sequential search, we ought to explicitly develop the expected overall cost V of using the latter method. When there are N agents, the expected overall cost of this process is:

$$V = E[X|X \le r] \sum_{i=1}^{N} F(r)(1 - F(r))^{i-1} + E_N[X|X > r](1 - F(r))^N$$

+ $c \sum_{i=1}^{N} (1 - F(r))^{i-1} = E[X|X \le r](1 - (1 - F(r))^N)$
+ $E_N[X|X > r](1 - F(r))^N + c \cdot \frac{1 - (1 - F(r))^N}{F(r)}$ (21)

where $E[X|X \leq r]$ is the expected value of an agent whose value is known to be in the range $[x_{min}, r]$, and $E_N[X|X > r]$ is the expected minimum value in a sample of size N when the minimum value is in the range $[r, x_{max}]$. The first two terms in (21) reflect the expected value of the opportunity returned by the search and the third term is the expected cost of obtaining the values. The first term is for when the search terminates with an agent with a value less than r (with probability $\sum_{i=1}^{N} F(r)(1-F(r))^{i-1}$), while the second term is when all of the agent values are above r (with probability $(1-F(r))^N)$, in which case the smallest value amongst all N agents is selected. $E[X|X \leq r]$ can be calculated using $f(x|x \leq r) = \frac{f(x)}{F(r)}$ and $F(x|x \leq r) = \frac{F(x)}{F(r)}$, such that $E[X|X \leq r] = \int_{x_{min}}^{r} y \frac{f(y)}{F(r)} dy = r - \frac{1}{F(r)} \int_{x_{min}}^{r} F(y) dy$. $E_N[X]$ can be calculated using the PDF $f_N(x)$ and CDF $F_N(x)$ of the minimum of a N-size sample. $F_N(x)$ is the probability that at least one agent in a sample of size N has the value x or less, which can be expressed as $F_N(x) = 1 - (1 - F(x))^N$. Thus, $E_N[X]$ can be calculated as follows, using integration by parts in the second step:

$$E_N[X] = \int_{x_{min}}^{x_{max}} y f_N(y) \, dy = x_{max} - \int_{x_{min}}^{x_{max}} F_N(y) \, dy$$

= $x_{min} + \int_{x_{min}}^{x_{max}} (1 - F(y))^N \, dy$ (22)

 $E_N[X|X > r]$ can be calculated by replacing F(x) with F(x|x > r) and x_{min} with r in (22).

As opposed to the sequential search, in the fixed sample size model [42], the searcher is limited to the selection of one sample of agents overall in a single period of time. The searcher then selects the agent with the lowest value from this sample. The expected cost of this strategy as a function of the number of agents simultaneously sampled, $0 < K \leq N$, is given by:

$$V = E_K[X] + \beta(K) \tag{23}$$

The optimal sample size is the value of K that minimizes (23). For general $\beta(i)$, we need to check K = 1, ..., N for the value with the lowest expected cost. If $\beta(i)$ is linear in i (that is, $\beta(i) = ci$) then it is only necessary to solve for K in the equation $E_K[X] - E_{K+1}[X] = c$ and then choose $\lfloor K \rfloor$ [42].

6.2. Increasing threshold-limiting sequential search

Increasing threshold search by itself may be a good alternative to economic search. Since increasing threshold search is only applicable when multiple search rounds are allowed, its most straightforward use is as an alternative to sequential search [27, 36]. The expected cost of the increasing threshold search under the economic search model is composed of the expected cost of search (from (11)) and the expected minimum value obtained (from (22)):

$$V = \frac{\alpha + \sum_{j=1}^{N} \beta(j) {N \choose j} P^{j} (1-P)^{N-j}}{1 - (1-P)^{N}} + E_{N}[X]$$
(24)

The performance of increasing threshold search when used as an alternative to economic search can be improved by processing the agents found using sequential search. According to this improvement, called "increasing threshold-limiting sequential search" from here on, the searcher publishes the thresholds as before until at least one agent responds. Then, instead of processing all of the responses, the searcher processes the responses sequentially as long as the revealed value so far is above some reservation value. Since the responses are all associated with the same probability distribution function, they can be processed in any order, as explained above. In some ways, the searcher is following the original sequential search method; however, its sample space is more targeted, as it is limited by the last threshold to which agents comply. The expected cost of applying a strategy $S = [r_1, \ldots, r_m = x_{max}]$ is thus similar to (4), except that $\beta(j)$ is replaced with the cost $V(j, r_{i-1}, r_i)$ of conducting the optimal sequential economic search on j items whose values are between the last two reservation values used, r_{i-1} and r_i . In this case, Theorem 1 and the resulting solution relying on a fixed reservation probability may no longer be applicable, since the transformation from S_2 to S'_1 and like transformations may not be possible. The complexity of deriving the optimal solution is beyond the scope of this paper. Instead, we show how to solve a discrete version of the problem, in which the reservation values are chosen from a finite set $\{x_1, \ldots, x_L\}$, $x_{min} < x_i < x_{i+1} < x_L = x_{max}$, just as in Section 3. Replacing $\beta(j)$ with $V(j, x_i, x_l)$ in (8) results in the following dynamic programming formulation: C(L) = 0

$$C(l) = \min_{l+1 \le i \le L} \left\{ \alpha + \sum_{j=1}^{N} V(j, x_l, x_i) \binom{N}{j} F(x_i | x_l)^j (1 - F(x_i | x_l))^{N-j} + C(x_i)(1 - F(x_i | x_l))^N \right\}$$
(25)

For any pair of values x_i and x_l , r and $V(j, x_i, x_l)$ can be calculated using (20) and (21), respectively, replacing F(x) with $\frac{F(x)-F(x_i)}{F(x_l)-F(x_i)}$, x_{min} with x_i , and x_{max} with x_l .

In Figure 8, we illustrate the properties and benefits of using increasing threshold search and its improved form as an alternative to sequential search. We use the synthetic environment described in Section 4 and Table 2, except that agents values are associated with the uniform distribution $(f(x) = 1 \text{ for } 0 \le x \le 1 \text{ and } f(x) = 0 \text{ otherwise})$. Figure 8 shows the overall cost of all three strategies (sequential, increasing threshold, and increasing threshold-limiting sequential search) as a function of the number of agents in the environment (horizontal axis), for different values of the ratio $\alpha/\beta(1)$, different cost functions (linear, logarithmic, and square), and different values of the coefficient c in $\beta(i)$. As can be observed from the figure, increasing threshold-limiting sequential search results in a lower overall cost than increasing threshold search, and this improvement increases as a function of α . While this characteristic always holds for linear and concave cost functions, it may not hold when the cost is convex (e.g., logarithmic), since the benefit of evaluating all agents found might outweigh the savings obtained by evaluating only part of the set one by one. Overall, both forms of increasing threshold search perform better than sequential search for low α and c values. In these cases, the improvement in value of the selected agent is greater



Figure 8: Expected overall cost of economic and increasing threshold search

than the additional costs α incurred during the search. Whenever the ratio between the fixed cost of each search iteration and the cost of evaluating an agent $(\alpha/\beta(1))$ is sufficiently low, increasing threshold search results in a lower expected overall cost than sequential search. This is a result of the low search costs due to the low $\alpha/\beta(1)$ ratio. As α increases, the advantages of the threshold-based methods diminish because the sequential search does not incur the cost α .

One interesting observation is that while the expected search cost of increasing threshold search increases as a function of the available agents N(see Figure 4), its overall cost (search cost and value obtained) decreases as N increases. This is because the reduction in the expected minimum value due to the additional agents outweighs the increased search costs. Likewise, we observe that as N increases, the two increasing threshold-based searches dominate sequential search for large $\alpha/\beta(1)$ ratios.

Finally, we note that the difference between the sequential search and the increasing threshold-limiting sequential search does not depend on the cost function $\beta(j)$. This is because both methods process agent values sequentially, incurring a cost $\beta(1)$ for each agent sample. This property is especially advantageous for these search methods if the cost function $\beta(j)$ is convex (e.g., square cost).

6.3. Combined fixed sample size and threshold search

The fixed sample size search [42] can also benefit from threshold based searches. While the restriction to a single search round precludes the use of an increasing threshold search to its full extent, the searcher can still improve expected search costs by publishing a single reservation value r besides sampling a fixed number of agent values, K. By publishing a reservation value, the searcher reduces the expected minimum value found by sampling alone, although with the additional expected cost of a one round threshold search. Sampling is still necessary, since it is possible that no agents will comply with the published reservation value. Because only one search round is allowed, both the threshold-based sampling and fixed size sampling must be conducted simultaneously. Returning to the example of contract bidding, this is like a case in which the agency must make a decision in the length of time it takes for contractors to prepare and submit their bids. The agency can simultaneously post a call for all bids under some threshold while soliciting bids from select contractors, and then choose the lowest of all bids. The expected cost in this case is:

$$V = \beta(K) + \alpha + \sum_{j=1}^{N-K} \beta(j) {N \choose j} F(r)^j (1 - F(r))^{N-K-j} + E_N[X|X \le r] (1 - (1 - F(r))^N) + E_K[X|X > r] (1 - F(r))^N$$
(26)

The first three terms in the above equation reflect the search costs, while the last two terms reflect the expected value of the agent found. The first term is the cost of sampling K elements; the next two terms are for the expected cost of the performing a threshold search on the remaining N - Kelements; the fourth term is the expected value of the agent found if its value is below the search threshold r, which will be found either by the sample or the threshold search; and the fifth term is the expected value of the agent if its value is above r, in which case it is necessarily one of the K agents sampled. $E_N[X|X \leq r]$ can be calculated using $f_N(x|x \leq r) = \frac{f_N(x)}{F_N(r)}$ as follows:

$$E_N[X|X \le r] = \int_{y=x_{min}}^r y \frac{f_N(y)}{F_N(r)} dy = \frac{x_{min} + \int_{y=x_{min}}^r (1 - F(y))^N dy}{1 - (1 - F(r))^N}$$
(27)



Figure 9: Expected overall cost of the fixed sample size and combined fixed sample size and threshold searches as a function of ${\cal N}$

Figure 9 shows the overall cost of the optimal fixed sample size search and combined fixed sample size and threshold search, using the same settings as in the previous section, as a function of the number of agents in the environment (horizontal axis) for different values of the ratio $\alpha/\beta(1)$, different cost functions (linear, logarithmic, and square), and different values of the coefficient c in $\beta(i)$. The combined search substantially improves performance when the ratio $\alpha/\beta(1)$ or the value of c is small. In these cases, the improvement from possibly finding a more targeted set of agents (that comply with the threshold) is greater than the additional cost α incurred. As α increases, the benefit of the combined search diminishes, and for large α values, there is no benefit to the combined search. This is also true for large values of c, since the expected number of samples obtained according to the optimal strategy is reduced, resulting in an increase in the expected value obtained. Similar to increasing threshold-limiting sequential search, the accumulated search cost of the combined search decreases as a function of N. In contrast, the difference in cost between the combined search and the fixed sample size search highly depends on the cost function used. The improvement is greater for concave functions, such as the square cost function, because the combined search uses a small fixed sample and sets the reservation value to minimize the number of agents found.

7. Related Work

While relevant literature is extensively cited throughout the paper whenever applicable, we use this section to introduce some additional work that may seem relevant to increasing threshold search and emphasize the differences between these works and ours.

Search by iterative expanding extents is used in many contexts; the distinguishing factor in determining optimal search sequences is the cost model. Iterative deepening [24, 41] and iterative broadening [17] are used to find goal states in artificial intelligence. Expanding ring search is used for route discovery in mobile networks, data discovery in sensor networks, and file location in peer-to-peer networks [3, 7, 8, 9, 10, 13, 19, 22, 28, 34, 45]. Iterative deepening, iterative broadening, and expanding ring search share the same cost model: The cost of each round includes the cost of revisiting all previously visited nodes as well as the cost of visiting new nodes. The cost model differs in peer-to-peer networking, since queries are processed only once by each node, but may be forwarded multiple times as the TTL value is increased [45]. Blocking expanding ring search [35] has a similar cost model, since nodes only forward messages once, waiting long enough to allow the source to stop the search if the target was found.

The expected search costs were modeled for many of these techniques, but optimal search sequences were primarily analyzed in the context of expanding ring search. Chang and Liu [7] show how to solve the problem optimally using dynamic programming for arbitrary cost and probability distribution functions that are known *a priori*. They also prove that the optimal strategy when the probability distribution of the minimum hop count to the target is not known *a priori* is a randomized strategy, derive such a strategy, and prove that it has a tight worst-case approximation factor of e [8]. Baryshnikov, et al. [3], prove that the California split rule is the optimal *deterministic* strategy when the probability distribution is not known, which has a tight worst-case approximation factor of 4. Cheng and Heinzelmann [10] design a novel heuristic called *ring splitting* to optimize search sequences for multiple targets. The peer-to-peer networking community has employed a technique called *dynamic querying* in multi-target queries [22], in which the searcher uses intermediate results to determine the next TTL value. The results of all these studies are inapplicable to the problem addressed in this paper because, as stated throughout, the cost models differ.

A remarkably long list of articles has been dedicated to variations of the "secretary problem" [14], which is a classical optimal-stopping online problem. Yet, the latter does not involve search costs and the goal is to maximize the probability of finding the best candidate rather than minimizing cost.

As discussed in Section 6 the problem of a searcher operating in a costly environment, seeking to maximize its long-term utility, is widely addressed in classical economic search theory ([27, 30, 36]) and references therein). Many variants of this model were considered, differing in the decision horizon (finite versus infinite) [27], the presence of the recall option [30], and the distribution of values. Some model variants assume findings are valid for a limited time, and with some probability may become obsolete and irrelevant for the search [25]. Over the years, several attempts were made to adopt search theory concepts for agent-based environments associated with search costs [12, 21, 23]. A few studies attempted to extend the search problem beyond a single search goal, e.g., attempting to purchase several commodities while facing imperfect information concerning prices [5, 6, 16], and to physical domains where the search cost derives from the agent's location along the search [20]. Some even considered multi-agent cooperative search for multiple goals [37]. Other work considered the benefit of sequential search as an alternative or a complementary means to mechanism design [1, 29]. For example, McAfee and McMillan [29] showed that a combination of reservation-price search and auction is, with costly communication, the optimal procurement mechanism of a monopolist, when the potential sellers have different production costs. In this case the monopolist sets a reservation value (reservation price) and sequentially offers the firms at this reservation price. If the set of potential bidders is exhausted without anyone accepting, then an auction is held. This idea was extended to the case where there is discounting and several homogeneous units are sought [1]. The main difference between economic search models and the model presented in this paper is that economic-search strategies attempt to find a suitable agent while optimizing the process as a whole, taking into consideration both the value of the agent found and the costs associated along the search, and do not guarantee finding the best-valued agent; thus, they are inapplicable to our problem. Instead, as demonstrated in Section 6, increasing threshold search is useful to economic search theory. To the best of our knowledge, no attempt has been made in search theory literature to use increasing threshold search as an alternative to or as a complementary mechanism to economic search strategies.

8. Discussion and Conclusion

As illustrated throughout the paper, increasing threshold search is highly effective in settings in which the cost of finding the agent associated with either the maximum or minimum value of any specific property needs to be minimized. This problem often arises in MAS, mostly due to their distributed nature. Despite its wide applicability, this method has not received adequate attention in the literature. Most research of techniques utilizing expanding search extents, as discussed in Sections 1, 4, and 7, has focused on models in which the costs associated with the search are correlated with the extent of search (rather than with the number of agents applicable to the extent used, as in our case). Consequently, the strategies most studied in those domains are very different from the optimal strategy to our problem, both in structure and quality of the solution obtained.

The paper introduces a solution to the problem in both the discrete and continuous domains. While the first uses a standard dynamic programming method, the analysis of the latter domain is quite non-standard: By correlating the reservation values to the respective probabilities in the proof of Theorem 1, the essence of the optimal solution in the form of a single reservation probability is revealed, enabling a solution to the problem. Furthermore, both the optimal reservation probability, from which the appropriate thresholds can be calculated, and the corresponding expected cost of search are distribution independent. This means that the solution to one problem instance can be used to derive the optimal search sequence for any other instance that only differs in its distribution of values by merely applying a simple transformation. It also enables making distribution-independent observations given the remaining parameters, such as the expected cost of the optimal strategy and expected number of rounds.

The same analysis methodology is replicated for the important and often necessary case of searching for multiple agents. Here, the single reservation probability property is preserved to the extent of any number of agents already found, and all the remaining advantages of the method (e.g., distribution-independence) hold.

As evident in Section 6, increasing threshold search can also be useful in economic search settings where the searcher is not constrained to finding the best-valued agent. This result is interesting, since increasing threshold search is by definition constrained by the need to find the best-valued agent. In economic search settings, increasing threshold search can either be used as an alternative to traditional economic search methods (e.g., sequential search), possibly resulting in lower expected overall cost, or to directly augment those methods, making use of the threshold publication aspect as part of or alongside the regular search routine. This has many implications in the evolving research area known also as *search theory* [30]. The paper defines and demonstrates the effectiveness of two such extensions to economic search, though many more can be considered. The determination of which amongst these search strategies will result in the lowest cost is setting dependent. Yet, given the probability distribution function, the search cost function, and the number of agents available, a searcher facing an economic search problem can follow the analysis given in this paper to calculate the expected cost of each method and choose accordingly.

Future work can consider more general assumptions regarding costs and cooperation of the different players. This may require modification of the analysis. In some cases, it may also impose external modeling challenges and even require the integration of buyers' strategic behavior. For example, the values of the agents can change in time. This requires external modeling of the way the distributions that characterize the agents values change over time. We note that in such a case, one may find it more beneficial to use a search technique that is based on increasing and decreasing thresholds. Such a strategy may be useful in cases in which agent values may actually decrease with time. Another example is the possibility that agents do not know their true value and will not readily calculate their value. (The current model assumes that agents know their values or are willing to calculate them-taxi drivers will always know their locations, sensors will know their readings). For example, if the threshold on a call for bids is very low, some contractors may decide that it is not worth preparing a bid and hence will not know their actual value. This requires modeling of the a priori value distribution of different agents, based on which their decision is made. More important, the strategic behavior of the agents, in terms of willingness to prepare a bid

based on the threshold they receive, should be modeled and taken into consideration. Another possible extension of this work includes the investigation of competition and cooperation models for two or more searchers when operating in settings where search is costly and one or more of them is capable of using increasing threshold search.

Acknowledgments

This work was supported in part by the Israeli Ministry of Industry and Trade under project RESCUE and ISF/BSF grants 1401/09 and 2008-404. We would like to especially thank Jeremy Schiff of the Bar Ilan University math department for his assistance with the proof of Theorem 2. Finally, we would like to thank the three anonymous reviewers for their valuable comments and suggestions to improve this paper.

Appendix

symbol	explanation	
N	number of agents in the system	
f(x)	PDF (probability density function) of agent values	
F(x)	CDF (cumulative distribution function) of agent values	
$[x_{min}, x_{max}]$	range of agent values	
eta(j)	cost of obtaining j agent values	
r	published reservation (threshold) value	
$S = [r_1, r_2, \ldots]$	search strategy	
α	fixed cost of publishing a reservation value	
V(S)	expected cost of using strategy S	
S^*	optimal strategy	
$f(x r_{i-1})$	PDF of agent values if all values are above r_{i-1}	
$F(x r_{i-1})$	CDF of agent values if all values are above r_{i-1}	
$\{x_1,\ldots,x_L\}$	discretized set of potential reservation values	
C(i)	expected cost of continuing a search after publishing reservation value x_i	
c	cost of obtaining the value of one agent in the linear cost model	
P	reservation probability used in the optimal strategy	
δ	increment of reservation value in the fixed increment and California split rule	
	strategies	
μ	parameter defining the number of search rounds in the fixed increment and	
	California split rule strategies	
K	number of agents that the searcher is interested in finding in a multi agent	
-	search; sample size in the fixed sample size economic search model	
k	number of agents found so far	
(r,k)	state at the beginning of a multi-agent search round	
$V^{(r,k)}(S)$	expected cost of continuing a search from state (r, k) when using strategy S	
P_i	reservation probability to use when i agents have already been found	
C(l,k)	expected cost of continuing a multi-agent search after using reservation value	
	x_l and a total of k agents have been found so far	
$\{B_1,,B_N\}$	possible opportunities in the sequential search model	
$E[X X \le r]$	expected value of an agent whose value is known to be in the range $[x_{min}, r]$	
$E_N[X X > r]$	expected minimum value in a sample of size N when the minimum value is	
	in the range $[r, x_{max}]$	
$f_N(x)$	PDF of the minimum of a N -size sample	
$F_N(x)$	CDF of the minimum of a N -size sample	
$f_N(x x \le r)$	PDF of the minimum of a N -size sample when the minimum is less than r	

References

- [1] M. Arnold and S. Lippman. Selecting a selling institution: Auctions versus sequential search. *Economic Inquiry*, 33(1):1–23, 1995.
- [2] Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. Mgmt. Science, 42:1676–92, 1997.
- [3] Y. Baryshnikov, E. Coffman Jr., P. Jelenkovic, P. Momcilovic, and D. Rubenstein. Flood search under the california split rule. *Oper. Res. Lett.*, 32(3):199–206, 2004.
- [4] J. Benhabib and C. Bull. Job search: The choice of intensity. Journal of Political Economy, 91:747–764, 1983.
- [5] K. Burdett and D. Malueg. The theory of search for several goods. Journal of Economic Theory, 24:362–376, 1981.
- [6] J. Carlson and R. McAfee. Joint search for several goods. Journal of Economic Theory, 32:337–345, 1984.
- [7] N. Chang and M. Liu. Revisiting the TTL-based controlled flooding search: optimality and randomization. In Zygmunt J. Haas, Samir R. Das, and Ravi Jain, editors, *MOBICOM*, pages 85–99. ACM, 2004.
- [8] N. Chang and M. Liu. Controlled flooding search in a large network. IEEE/ACM Trans. Netw., 15(2):436–449, 2007.
- [9] Z. Cheng and W. Heinzelman. Flooding strategy for target discovery in wireless networks. *Wireless Networks*, 11(5):607–618, 2005.
- [10] Z. Cheng and W. Heinzelman. Searching strategies for target discovery in wireless networks. Ad Hoc Networks, 5(4):413–428, 2007.
- [11] M. Chhabra, S. Das, and D. Sarne. Expert-mediated search. In AAMAS, pages 415–422, 2011.
- [12] S. Choi and J. Liu. Markov decision approach for time-constrained trading in electronic marketplace. *International Journal of Information Technology and Decision Making*, 1(3):511–524, 2002.

- [13] J. Deng and S. Zuyev. On search sets of expanding ring search in wireless networks. Ad Hoc Networks, 6(7):1168–1181, 2008.
- [14] T. Ferguson. Who solved the secretary problem? Statistical Science, 4(3):282–289, 1989.
- [15] S. Gal, M. Landsberger, and B. Levykson. A compound strategy for search in the labor market. *International Economic Review*, 22(3):597– 608, 1981.
- [16] J. Gatti. Multi-commodity consumer search. Journal of Economic Theory, 86(2):219–244, 1999.
- [17] M. Ginsberg and W. Harvey. Iterative broadening. Artif. Intell., 55(2):367–383, 1992.
- [18] W. Greene. *Econometric Analysis*. Prentice Hall, 5 edition, 2003.
- [19] J. Hassan and S. Jha. On the optimization trade-offs of expanding ring search. In N. Das, A. Sen, S. Das, and B. Sinha, editors, *IWDC*, volume 3326 of *Lecture Notes in Computer Science*, pages 489–494. Springer, 2004.
- [20] N. Hazon, Y. Aumann, and S. Kraus. Collaborative multi agent physical search with probabilistic knowledge. In Craig Boutilier, editor, *IJCAI*, pages 167–174, 2009.
- [21] J. Huang and A. Kazeykina. Optimal strategies for reviewing search results. In Maria Fox and David Poole, editors, AAAI. AAAI Press, 2010.
- [22] H. Jiang and S. Jin. Exploiting dynamic querying like flooding techniques in unstructured peer-to-peer networks. In *ICNP*, pages 122–131. IEEE Computer Society, 2005.
- [23] J. Kephart and A. Greenwald. Shopbot economics. JAAMAS, 5(3):255– 287, 2002.
- [24] R. Korf. Depth-first iterative-deepening: An optimal admissible tree search. Artif. Intell., 27(1):97–109, 1985.

- [25] M. Landsberger and D. Peled. Duration of offers, price structure, and the gain from search. *Journal of Economic Theory*, 16(1):17–37, 1977.
- [26] R. Lentz. Sorting by search intensity. Journal of Economic Theory, 145(4):1436 – 1452, 2010.
- [27] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189, 1976.
- [28] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and replication in unstructured peer-to-peer networks. In *ICS*, pages 84–95, 2002.
- [29] R. P. McAfee and J. McMillan. Search mechanisms. Journal of Economic Theory, 44(1):99 – 123, 1988.
- [30] J. McMillan and M. Rothschild. Search. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. 1994.
- [31] P. R. Milgrom and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):pp. 1089–1122, 1982.
- [32] P. Morgan and R. Manning. Optimal search. Econometrica, 53(4):923– 944, 1985.
- [33] G. Moscarini and R. Wright. Introduction to search theory and applications. Journal of Economic Theory, 145(4):1319 – 1324, 2010.
- [34] N. Pham, V. Zalyubovskiy, and H. Choo. Threshold analysis of expanding ring search for multi-hop wireless networks. In *The 3rd Asia Pacific International Conference on Information Science and Technol*ogy (APIC-IST), pages 93–102, 2008.
- [35] I. Pu and Y. Shen. Analytical studies of energytime efficiency of blocking expanding ring search. *Mathematics in Computer Science*, 3:443–456, 2010.
- [36] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4):689–711, 1974.

- [37] D. Sarne, E. Manisterski, and S. Kraus. Multi-goal economic search using dynamic search structures. Autonomous Agents and Multi-Agent Systems, 21:204–236, 2010.
- [38] A. Silberstein, R. Braynard, C. Schlatter Ellis, K. Munagala, and J. Yang. A sampling-based approach to optimizing top-k queries in sensor networks. In L. Liu, A. Reuter, K. Whang, and J. Zhang, editors, *ICDE*, page 68. IEEE Computer Society, 2006.
- [39] A. Silberstein, K. Munagala, and J. Yang. Energy-efficient monitoring of extreme values in sensor networks. In S. Chaudhuri, V. Hristidis, and N. Polyzotis, editors, SIGMOD Conference, pages 169–180. ACM, 2006.
- [40] L. Smith. Frictional matching models. Annual Reviews in Economics, 3:319–338, 2011.
- [41] M. Stickel and M. Tyson. An analysis of consecutively bounded depthfirst search with applications in automated deduction. In *IJCAI*, pages 1073–1075, 1985.
- [42] G. Stigler. The economics of information. Journal of Political Economy, 69(3):213–225, 1961.
- [43] M. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–54, May 1979.
- [44] M. Wybo, J. Robert, and P. Léger. Using search theory to determine an applications selection strategy. *Information & Management*, 46(5):285– 293, 2009.
- [45] B. Yang and H. Garcia-Molina. Improving search in peer-to-peer networks. In *ICDCS*, pages 5–14, 2002.