

Expert-Mediated Search

Meenal Chhabra
Rensselaer Polytechnic Inst.
Dept. of Computer Science
Troy, NY, USA
chhabm@cs.rpi.edu

Sanmay Das
Rensselaer Polytechnic Inst.
Dept. of Computer Science
Troy, NY, USA
sanmay@cs.rpi.edu

David Sarne
Bar-Ilan University
Dept. of Computer Science
Ramat Gan, Israel
sarned@cs.biu.edu

ABSTRACT

Increasingly in both traditional, and especially Internet-based marketplaces, knowledge is becoming a traded commodity. This paper considers the impact of the presence of knowledge-brokers, or experts, on search-based markets with noisy signals. For example, consider a consumer looking for a used car on a large Internet marketplace. She sees noisy signals of the true value of any car she looks at the advertisement for, and can disambiguate this signal by paying for the services of an expert (for example, getting a Carfax report, or taking the car to a mechanic for an inspection). Both the consumer and the expert are rational, self-interested agents. We present a model for such search environments, and analyze several aspects of the model, making three main contributions: (1) We derive the consumer's optimal search strategy in environments with noisy signals, with and without the option of consulting an expert; (2) We find the optimal strategy for maximizing the expert's profit; (3) We study the option of market designers to subsidize search in a way that improves overall social welfare. We illustrate our results in the context of a plausible distribution of signals and values.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Economics

Keywords

Economically-motivated agents, Modeling the dynamics of MAS

1. INTRODUCTION

In many multi-agent system (MAS) settings, agents engage in *one-sided search* [13, 6]. This is a process in which an agent faces a stream of opportunities that arise sequentially, and the process terminates when the agent picks one of those opportunities. A classic example is a consumer looking to buy a used car. She will typically investigate cars one at a time until deciding upon one she wants. Similar settings

can be found in job-search, house search and other applications [12, 15]. In modern electronic marketplaces, search is likely to become increasingly important as, with the proliferation of possible sellers of a good, consumers will turn to artificial agents, whom we term “searchers,” to find the best prices or values for items they are interested in acquiring.

The decision making complexity in one-sided search usually arises from the fact that there is a cost incurred in finding out the true value of any opportunity encountered. For example, there is a cost to arranging a meeting to test drive a car you are considering purchasing. The searcher thus needs to trade off the potential benefit of continuing to search and seeing a more valuable opportunity with the costs incurred in doing so. The optimal stopping rule for such search problems has been widely studied, and is often a *reservation strategy*, where the searcher should terminate search upon encountering an opportunity which has a value above a certain reservation value or threshold [18, 13]. Most models assume that the searcher obtains the exact true value of the opportunities it encounters. However, in many realistic settings, search is inherently noisy and searchers may only obtain a noisy signal of the true value. For example, the drivetrain of a used car may not be in good condition, even if the body of the car looks terrific. The relaxation of the assumption of perfect values not only changes the optimal strategy for a searcher, it also leads to a niche in the marketplace for new knowledge-brokers. The knowledge brokers, or *experts*, are service providers whose main role is to inform consumers or searchers about the values of opportunities.

An expert offers the searcher the option to obtain a more precise estimate of the value of an opportunity in question, in exchange for the payment of a fee (which covers the cost of providing the service as well as the profit of the expert). To continue with the used car example, when the agent is intrigued by a particular car and wants to learn more about it, she could take the car to a mechanic who could investigate the car in more detail to make sure it is not a lemon. The expert need not be a mechanic – it could be, for example, an independent agency, like Carfax, that monitors the recorded history of transactions, repairs, claims, etc. on cars. It can also be a repeated visit of the searcher to see the car, possibly bringing an experienced friend for a more thorough examination. In all these cases, more accurate information is obtained for an additional cost (either monetary or equivalent).

In this paper we investigate optimal search and mechanism design in environments where searchers observe noisy signals and can obtain (i.e., query the expert for) the ac-

Cite as: Expert-Mediated Search, Meenal Chhabra, Sanmay Das, and David Sarne, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 415–422. Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

tual values for a fee. Our main contributions are threefold: First, we introduce a specific model of one-sided search with noisy signals and prove that the optimal search rule, for a large class of real-life settings, is reservation-value based. Second, we formally introduce the option of consulting an expert in such noisy search environments, and derive the optimal strategy for the searcher given a cost of consulting the expert, as well as the profit-maximizing price for the expert to charge for its services. As part of the analysis we prove that under the standard assumption that higher signals are “good news” (i.e., the distribution of the true value conditional on a higher signal stochastically dominates the distribution conditional on a lower signal), the optimal search strategy is characterized by a “double reservation value” strategy, wherein the searcher rejects all signals below a certain threshold, resuming search, and accepts all signals above another threshold, thus terminating search, without querying the expert; the agent queries the expert for all signals that are between the two thresholds. Finally, we study a market design mechanism (introduction of a subsidy) with the potential to improve social welfare in such domains. These mechanisms may be implemented by either an electronic marketplace that employs the expert (for example a website for used cars that has a relationship with a provider of car history reports), or an entity with regulatory power like the government. Along with the general theory, we illustrate our results in a specific, plausible distributional model, in which the true value is always bounded by the signal value, and the probability monotonically decays as the discrepancy between the two increases.

2. THE GENERAL MODEL

The standard one-sided search problem [13] considers an agent or *searcher* facing an infinite stream of opportunities from which she needs to choose one. While the specific value v of each future opportunity is unknown to the searcher, she is acquainted with the (stationary) probability distribution function from which opportunities are drawn, denoted $f_v(x)$. The searcher can learn the value of an opportunity for a cost c_s (either monetary or in terms of resources that need to be consumed for this purpose) and her goal is to maximize the net benefit, defined as the value of the opportunity eventually picked minus the overall cost incurred during the search. Having no *a priori* information about any specific opportunity, the searcher reviews the opportunities she encounters sequentially and sets her optimal stopping rule. The stopping rule specifies when to terminate and when to resume search, based on the opportunities encountered.

Our model relaxes the standard assumption that the searcher receives the exact true value of an opportunity. Instead, we assume that the searcher receives, at cost c_s , a noisy signal s , correlated with the true value according to a known probability density function $f_s(s|v)$. In addition, the searcher may query and obtain from a third party (the *expert*) the true value v of an opportunity for which signal s was received, by paying an additional fee c_e . The goal of the searcher is to maximize the total utility received i.e., the expected value of the opportunity eventually picked minus the expected cost of search and expert fees paid along the way.

The first question that arises is how to characterize the optimal strategy for the searcher. A second question is how the expert sets her service fee c_e . In this paper we consider a monopolist provider of expert services. The searcher’s opti-

mal strategy is directly influenced by c_e , and thus implicitly determines the expected number of times the services of the expert are required, and thus the expert’s revenue. The problem can be thought of as a Stackelberg game [5] where the expert is the first mover, and wants to maximize her profits with respect to the fee c_e she charges searchers.

The new search model raises interesting new questions about market design. Assuming exogeneity of opportunities, social welfare, denoted W , is a function of the expected value to searchers and the expected profit of the expert. (We abstract away from modeling the existence of “sellers” of opportunities, instead viewing them as exogenous, or else as being offered at some “fair price” by the seller.) We assume that the provider of expert services has already performed the “startup work” necessary, and only pays a marginal cost d_e per query, and that social welfare is additive. Since the process scales up linearly in the number of searchers, we can simply consider the interactions involving a single searcher and the expert. The social welfare is then the sum of the expected net benefit to the searcher and the expected profit of the expert. It turns out that social welfare can be significantly affected (and improved) if the market designer (or a regulator like the government) subsidizes queries by compensating the expert in order to reduce query costs to the searcher.

We now turn to developing the mathematical machinery to address these problems.

3. OPTIMAL POLICIES

In this section, we analyze the searcher’s optimal search strategy and her expected use of the expert’s services, given the fee c_e set by the expert. The analysis builds on the trivial non-noisy model and gradually adds the complexities of signals and having the expert option. From the searcher’s optimal search strategy we derive the expert’s expected benefits as a function of the fee she sets, enabling maximization of the expert’s revenue.

One-Sided Search.

The optimal search strategy for the standard model, where the actual value of an opportunity can be obtained at cost c_s , can be found in the extensive literature of search theory [13, 6]. In this case, the searcher follows a reservation-value rule: she reviews opportunities sequentially (in random order) and terminates the search once a value greater than a reservation value x^* is revealed, where the reservation value x^* satisfies:

$$c_s = \int_{y=x^*}^{\infty} (y - x^*) f_v(y) dy \quad (1)$$

Intuitively, x^* is the value where the searcher is precisely indifferent: the expected marginal benefit from continuing search and obtaining the value of the next opportunity exactly equals the cost of obtaining that additional value. The reservation property of the optimal strategy derives from the stationarity of the problem — resuming the search places the searcher at the same position as at the beginning of the search [13]. Consequently, a searcher that follows a reservation value strategy will never decide to accept an opportunity she has once rejected and the optimal search strategy is the same whether or not recall is permitted. The expected number of search iterations is simply the inverse of the success probability, $\frac{1}{1 - F_v(x^*)}$, since this becomes a Bernoulli

sampling process, as opportunities arise independently at each iteration.

One-Sided Search with Noisy Signals.

Before beginning the analysis of search with noisy signals, we emphasize that, given $f_v(x)$ and $f_s(s|v)$, we can also derive the distribution of the signal received from a random opportunity, $f_s(x)$, and the distribution of true values conditional on signals, $f_v(v|s)$ (the conditionals are interchangeable by Bayes' law). In many domains, it may be easier to assess/learn $f_s(s)$ than $f_v(v|s)$ as most past experience involves signals, with the actual value revealed for only a subset of these signals.

When the searcher receives a noisy signal rather than the actual value of an opportunity, there is no guarantee that the optimal strategy is reservation-value based as in the case where values obtained are certain. Indeed, the stationarity of the problem still holds, and an opportunity that has been rejected will never be recalled. Yet, in the absence of any restriction over $f_s(s|v)$, the optimal strategy is based on a set S of signal-value intervals for which the searcher terminates the search. The expected value in this case, denoted $V(S)$, is given by:

$$\begin{aligned} V(S) &= -c_s + \Pr(s \notin S)V(S) + \Pr(s \in S)E[v|s \in S] \\ &= -c_s + V(S) \int_{s \notin S} f_s(s) ds + \int_{s \in S} f_s(s)E[v|s] ds \end{aligned} \quad (2)$$

The fact that the optimal strategy may not be reservation-value based in this case is because there may be no correlation between the signal and the true value of the opportunity. Nevertheless, in most real-life cases, there is a natural correlation between signals and true values. In particular, a fairly weak and commonly used restriction on the conditional distribution of the true value given the signal goes a long way towards allowing us to recapture a simple space of optimal strategies. This is the restriction that higher signal values are "good news" in the sense that when $s_1 > s_2$, the conditional distribution of v given s_1 first-order stochastically dominates that of v given s_2 [19, 14]. The condition requires that given two signals s_1 and s_2 where $s_1 > s_2$, the probability that the actual value is greater than any particular value v is greater for the case where the searcher receives signal s_1 . Formally:

DEFINITION 1. Higher signals are good news (HSGN) assumption: If $s_1 > s_2$, then, $\forall y, F_v(y|s_1) \leq F_v(y|s_2)$.

This enables us to prove the following theorem.

THEOREM 1. For any probability density function $f_v(v|s)$ satisfying the HSGN assumption, the optimal search strategy is a reservation-value rule, where the reservation value, t^* , satisfies:

$$c_s = \int_{s=t^*}^{\infty} (E[v|s] - E[v|t^*]) ds \quad (3)$$

Proof: The proof is based on showing that, if according to the optimal search strategy the searcher should resume her search given a signal s , then she must necessarily also do so given any other signal $s' < s$. Let V denote the expected benefit to the searcher if resuming the search. Since the optimal strategy given signal s is to resume search, we know $V > E[v|s]$. Given the HSGN assumption, $\int_y y f_v(y|s') dy < \int_y y f_v(y|s) dy$ holds for $s' < s$. Therefore, $V > E[v|s']$,

proving that the optimal strategy is reservation-value. Then, the expected value of the searcher when using reservation signal t is given by:

$$\begin{aligned} V(t) &= -c_s + V(t) \int_{s=-\infty}^t f_s(s) ds + \int_{s=t}^{\infty} E[v|s] f_s(s) ds \\ &= \frac{-c_s + \int_{s=t}^{\infty} E[v|s] f_s(s) ds}{1 - F_s(t)} \end{aligned} \quad (4)$$

where $F_s(s)$ is the cumulative distribution function of the signal s . Setting the first derivative according to t of Equation 4 to zero we obtain: $V(t^*) = E[v|t^*]$. The second derivative for t^* that satisfies the latter equality confirms that this is indeed a global maximum. Finally, using integration by parts over the derivative according to t of Equation 4 we obtain Equation 3, and the value t^* can be calculated accordingly. \square

The social welfare W is the expected gain to the searcher from following the optimal strategy, $V(t^*) = E[v|t^*]$. The expected number of search iterations is $\frac{1}{1 - F_s(t^*)}$, since this is a Bernoulli sampling process.

The Expert Option.

The introduction of an expert extends the number of decision alternatives available to the searcher. When receiving a noisy signal of the true value, she can choose to (1) reject the offer without querying the expert, paying search cost c_s to reveal the signal for the next offer; (2) query the expert to obtain the true value, paying a cost c_e , and then make a decision; or (3) accept the offer without querying the expert, receiving the (unknown) true value of the offer. In case (2), there is an additional decision to be made, whether to resume search or not, after the true value v is revealed.

As in the no-expert case, a solution for a general density function $f_v(v|s)$ dictates an optimal strategy of a complex structure. In our case, the optimal strategy will have the form of (S', S'', V) , where: (a) S' is a set of signal intervals for which the searcher should resume her search without querying the expert; (b) S'' is a set of signal intervals for which the searcher should terminate her search without querying the expert (and pick the opportunity associated with this signal); and (c) for any signal that is not in S' or S'' the searcher should query the expert, and terminate the search if the value obtained is above a threshold V , and resume otherwise. The value V is the expected benefit from resuming the search and is given by the following modification of Equation 2:

$$\begin{aligned} V &= -c_s + V \int_{s \in S'} f_s(s) ds - c_e \int_{s \notin \{S', S''\}} f_s(s) ds + \\ &\quad \int_{s \notin \{S', S''\}} f_s(s) \left(V \int_{-\infty}^V f_v(x|s) dx + \int_{x=V}^{\infty} x f_v(x|s) dx \right) ds + \int_{s \in S''} f_s(s) E[v|s] ds \end{aligned} \quad (5)$$

The first element on the right hand side of the equation applies to the case of resuming search, in which case the searcher continues with an expected benefit V . The second element is the expected payment to the expert. The next elements relate to the case where the search is resumed based on the value received from the expert (in which case the expected revenue is once again V) and where the search is

terminated (with the value $E[v|s]$ obtained as the revenue), respectively. Finally, the last element applies to the case where the searcher terminates the search without querying the expert.

We can show that under the HSGN assumption, each of the sets S' and S'' actually contains a single interval of signals, as illustrated in Figure 1.

THEOREM 2. *For $f_v(y|s)$ satisfying the HSGN assumption (Definition 1), the optimal search strategy can be described by the tuple (t_l, t_u, V) , where: (a) t_l is a signal threshold below which the search should be resumed; (b) t_u is a signal threshold above which the search should be terminated and the current opportunity picked; and (c) the expert should be queried given any signal $t_l < s < t_u$ and the opportunity should be accepted (and search terminated) if the value obtained from the expert is above the expected value of resuming the search, V , otherwise search should resume (see Figure 1). The values t_l , t_u and V can be calculated from solving the set of Equations 6-8:*

$$V = -c_s + V \int_{s=-\infty}^{t_l} f_s(s) ds - c_e \int_{s=t_l}^{t_u} f_s(s) ds + \int_{s=t_l}^{t_u} f_s(s) \left(V \int_{x=-\infty}^V f_v(x|s) dx ds + \int_{x=V}^{\infty} x f_v(x|s) dx ds \right) + \int_{s=t_u}^{\infty} f_s(s) E[v|s] ds \quad (6)$$

$$c_e = \int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy \quad (7)$$

$$c_e = \int_{-\infty}^V (V - y) f_v(y|t_u) dy \quad (8)$$

Proof: The proof extends the methodology used for proving Theorem 1. We first show that if, according to the optimal search strategy the searcher should resume her search given a signal s , then she must also do so given any other signal $s' < s$. Then, we show that if, according to the optimal search strategy the searcher should terminate her search given a signal s , then she must also necessarily do so given any other signal $s'' > s$. Again, we use V to denote the expected benefit to the searcher if resuming the search.

If the optimal strategy given signal s is to resume search then the following two inequalities should hold, describing the superiority of resuming search over terminating search (Equation 9) and querying the expert (Equation 10):

$$V > E[v|s] \quad (9)$$

$$V > -c_e + \int_{y=V}^{\infty} y f_v(y|s) dy + V \int_{y=-\infty}^V f_v(y|s) dy \quad (10)$$

Given the HSGN assumption and since $s' < s$, Equation 9 holds also for s' , and so does Equation 10 (which can be formalized after some mathematical manipulation as: $V > -c_e + V + \int_{y=V}^{\infty} (y - V) f_v(y|s) dy$). The proof for $s'' > s$ is similar: the expected cost of accepting the current opportunity can be shown to dominate both resuming the search and querying the expert. We omit the details because of space considerations. The optimal strategy can thus be described by the tuple (t_l, t_u, V) as stated in the Theorem. Therefore, Equation 5 transforms into Equation 6. Taking the derivative of Equation 6 w.r.t. t_l and equating to zero,

we obtain a unique t_l which maximizes the expected benefit (verified by second derivative), and similarly for t_u . Finally, using integration by parts over the derivatives of Equation 6 w.r.t. t_l and t_u we obtain Equations 7-8. \square

Intuitively, t_l is the point at which a searcher is indifferent between either resuming the search or querying the expert and t_u is the point at which a searcher is indifferent between either terminating the search or querying the expert. The cost of purchasing the expert's services must equal two different things: (1) the expected savings from resuming the search when the actual utility from the current opportunity (which is not known) turns out to be greater than what can be gained from resuming the search (once it is revealed) (this is the condition for t_l); (2) the expected savings from terminating the search in those cases where the actual utility from the current opportunity (once revealed) is less than what can be gained from resuming the search (for t_u).




Figure 1: Characterization of the optimal strategy for noisy search with an expert. The searcher queries the expert if $s \in [t_l, t_u]$ and accepts the offer if the worth is greater than the value of resuming the search V . The searcher rejects and resumes search if $s < t_l$ and accepts and terminates search if $s > t_u$, both without querying the expert.

It is notable that there is also a reasonable degenerate case where $t_l = t_u (= t)$. This happens when the cost of querying is so high that it never makes sense to engage the expert's services. In this case, a direct indifference constraint exists at the threshold t , where accepting the offer yields the same expected value as continuing search, so $V = E[v|t]$. This can be solved in combination with Equation 4, since there are now only two relevant variables.

Expected number of queries: The search strategy (t_l, t_u, V) defines how many times the expert's services are consulted. In order to compute the expected number of queries, we consider four different types of transitions in the system. Let A be the probability that the searcher queries the expert and then does not accept, resuming search, B be the probability that the searcher resumes search without querying, C be the probability that the searcher terminates without querying, and D be the probability that the searcher queries the expert and terminates search. Then:

$$A = \Pr(t_l \leq s \leq t_u \text{ and } v < V) \quad (11)$$

$$B = \Pr(s < t_l) \quad (12)$$

$$C = \Pr(s > t_u) \quad (13)$$

$$D = \Pr(t_l \leq s \leq t_u \text{ and } v \geq V) \quad (14)$$

Let P_j denote the probability that the searcher queries the expert exactly j times before terminating. The searcher can terminate search after exactly j queries in one of two ways: either she makes $j - 1$ queries, then queries the expert and chooses to terminate, or she makes j queries and

then chooses to terminate without querying the expert. Factoring in all the possible ways of interleaving j queries with an arbitrary number of times that the searcher chooses to continue search without querying the expert, we get:

$$P_j = \sum_{m=0}^{\infty} \frac{(m+j-1)!}{m!(j-1)!} A^{j-1} B^m D + \sum_{m=0}^{\infty} \frac{(m+j)!}{m!j!} A^j B^m C$$

$$= \frac{A^j}{(1-B)^j} \left(\frac{D}{A} + \frac{C}{1-B} \right)$$

Then the expected number of queries, when charging an expert fee c_e , is given by $\sum_{j=0}^{\infty} jP_j$, yielding

$$\eta_{c_e} = \mathbb{E}(\text{Number of queries}|c_e) = \frac{(1-B)D + CA}{(1-B-A)^2} \quad (15)$$

Expected number of opportunities examined: Using the same notation as above, we see that the probability of terminating the search at any iteration is $C + D$, and these are independent Bernoulli draws at each opportunity. Therefore the expectation of the number of opportunities examined is simply $\eta_s = 1/(C + D)$.

Expected profit of the expert: Let d_e denote the marginal cost of the service the expert is providing. The expected profit of the expert is then simply

$$\pi_e = \mathbb{E}(\text{Profit}) = (c_e - d_e)\eta_{c_e}$$

The expert can maximize the above expression with respect to c_e (η_{c_e} decreases as c_e increases) to find the profit maximizing price to charge searchers.

Social Welfare: The social welfare is given by the sum of all parties involved, thus far just the searcher and the expert. Of course this generalizes to multiple agents as well, since each search process would be independent. We define:

$$W = V_{c_e^*} + \pi_e \quad (16)$$

where c_e^* is the fee that maximizes the expert's profit.

4. MARKET DESIGN

Above, we have described the basics of search in such expert-mediated markets. In this section, we describe possible uses of the theory described above to improve the design of markets in which such search takes place. The prospect of designing or significantly influencing these markets is not remote. Consider the design of a large scale Internet website like `autotrader.com`. The listings for cars that users see are signals, and they may be unsure of a car's true worth. `autotrader` can partner with a provider of reports like Carfax, to make it easy for users to look up a car's worth. In order to be general, let us refer to `autotrader` as "the market" (or in some instances as "the market designer") and Carfax as the expert. The market wants to attract customers to it, rather than to rival markets. The best way of doing this is to provide customers with a high value shopping experience. The expert wishes to maximize its profits. Since the market and the expert both have significant power, it is reasonable to imagine them coming up with different models of the kinds of relationships they may have. It is notable that while we are thinking about private markets here, this entire discussion is equally relevant to a big player like the government as market designer, and independent providers of expert services.

In order to provide customers with the highest value shopping experience, the market may choose to subsidize the cost of expert services. A typical problem with subsidization is that it often decreases social welfare because the true cost of whatever is being subsidized is hidden from the consumer, leading to overconsumption of the resource. In this instance, however, the natural existence of many monopolies in expert services, combined with the existence of search frictions, make it quite possible that subsidies will in fact increase social welfare. We show in Section 5 that this is in fact the case for some natural distributions.

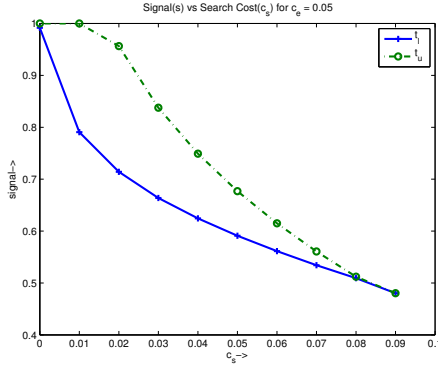
The basic framework of subsidization works as follows. Suppose a monopolist provider of expert services maximizes its profits by setting the querying cost to c_e^* , yielding an expected profit $\pi_e = (c_e^* - d_e)\eta_{c_e^*}$ (this discussion is on a per-consumer basis). The market designer can step in and negotiate a reduction of the fee c_e charged by the expert, for the benefit of the agents. In return for the expert's agreement, the market designer will need to offer a per-consumer payment β to the expert, which fully compensates the expert for the decreased revenue, leaving her total profit the same. Since $c_e' < c_e^*$, $\eta_{c_e'} > \eta_{c_e^*}$ (the consumer queries more often because she has to pay less). The compensation for a requested decrease in the expert's fee from c_e^* to c_e' is thus $\beta = (c_e^* - d_e)\eta_{c_e^*} - (c_e' - d_e)\eta_{c_e'}$. The overall welfare per agent in this case increases by $V_{c_e'} - V_{c_e^*}$, where $V_{c_e'}$ and $V_{c_e^*}$ are the expected value of searchers according to Equation 6-8, when the expert uses a fee c_e' and c_e^* respectively, at a cost β to the market designer. Since the expert is fully compensated for her loss due to the decrease in her fee, the change in the overall social welfare is $V_{c_e'} - V_{c_e^*} - \beta$. Under the new pricing scheme c_e' , and given the subsidy β , the social welfare is given by $W' = V_{c_e'} + \pi_e - \beta$. In the following section we illustrate how such a subsidy β can have a positive change over the social welfare.

An interesting special case to consider is when $d_e = 0$. We can think of this case as "digital services," analogous to digital goods like music MP3s – producing an extra one of these has zero marginal cost. Similarly, producing an extra electronic history of a car, like a Carfax report, can be considered to have zero marginal cost. In this case, there is no societal cost to higher utilization of the expert's services, so subsidy is welfare improving right up to making the service free. These are the cases where it could make sense for the market designer or government to take over offering the service themselves, and making it free, potentially leveraging the increased welfare of consumers by attracting more consumers to their market, or increasing their fees.

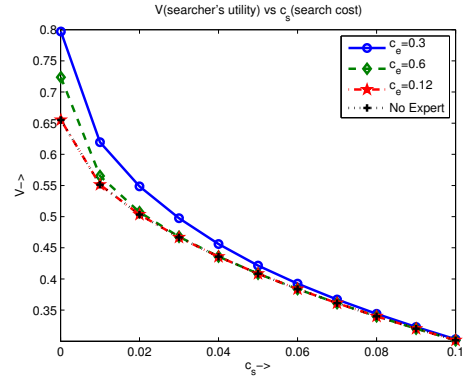
5. A SPECIFIC EXAMPLE

In this section we illustrate the theoretical analysis given in the former section for a particular plausible distribution of signals and values. This case illustrates the general structure of the solutions of the model and demonstrates how interventions by the market designer can increase social welfare.

We consider a case where the signal is an upper bound on the true value. Going back to the used car example, sellers and dealers offering cars for sale usually make cosmetic improvements to the cars in question, and proceed to advertise them in the most appealing manner possible, hiding defects using temporary fixes. Specifically, we assume signals s are



(a) Signal threshold vs c_s



(b) Variation of agent utility w.r.t c_s and c_e

Figure 2: Effect of c_s on the signal thresholds (t_l, t_u) and agent utility V

uniformly distributed on $[0, 1]$, and the conditional density of true values is linear on $[0, s]$. Thus

$$f_s(s) = \begin{cases} 1 & \text{if } 0 < s < 1 \\ 0 & \text{otherwise} \end{cases} \quad f_v(y|s) = \begin{cases} \frac{2y}{s^2} & \text{for } 0 \leq y \leq s \\ 0 & \text{Otherwise} \end{cases}$$

We can substitute in these distributions in Equations 6 through 8 and simplify. From Equation 6:

$$V = -c_s + Vt_l - c_e(t_u - t_l) + \frac{V^3(t_u - t_l)}{3t_u t_l} + \frac{1 - t_l^2}{3}$$

$$V = \frac{2t_l}{3} + \frac{V^3}{3t_l^2} - c_e \quad (\text{from Equation 7})$$

$$c_e = \int_0^V (V - y)f_v(y|t_u) dy = \frac{V^3}{3t_u^2} \quad (\text{from Equation 8})$$

We can find feasible solutions of this system for different parameter values, as long as the condition $t_l < t_u$ holds. Otherwise, when c_e is high enough that querying never makes sense, a single threshold serves as the optimal strategy, as in the case with no expert. In the latter case, we obtain the optimal reservation value to be used by the searcher from Equation 3, yielding $t^* = 1 - \sqrt{3c_s}$.

The other thing to note here is that Equation 8 above is for the case when the support on signal s is unbounded. When there is an upper limit on s i.e $s \leq m$ for some m (as is the case here, where signals are bounded in $[0, 1]$), once t_u reaches m (we never buy without querying), Equation 8 does not hold. Now the system rejects if the signal is below t_l or queries if it is above.

Figure 2(a) illustrates how the reservation values t_l and t_u change as a function of c_s for $c_e = 0.05$. The vertical axis is the interval of signals. As can be seen from the graph, for very small search cost (c_s) values, the searcher never terminates search without querying the expert.¹ Due to the low search cost the searcher is better off only querying the expert when a high signal is received. The expert option is preferred over accepting without querying the expert for those high signals, because if a low value is received from the expert then the cost of finding a new opportunity with a high signal is low. As the search cost c_s increases, there is some

¹When search costs are 0 the problem is ill-defined. The first point on the graph shows an extremely low, but non-zero search cost. In this case $t_u = 1$ and t_l is almost 1, but not exactly, and the expert is again always queried.

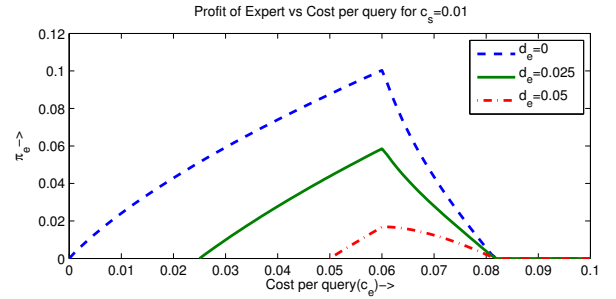


Figure 3: Expert's profit as a function of c_e and d_e for $c_s = 0.01$.

behavior that is not immediately intuitive. The reservation values t_l and t_u become closer to each other until coinciding at $c_s = 0.08$, at which point the expert is never queried anymore. The reason for this is that the overall value of continuing search goes down significantly as c_s increases, therefore the cost of querying the expert becomes a more significant fraction of the total cost, making it comparatively less desirable. This is a good example of the additional complexity of analyzing a system with an expert, because in the static sense the cost of consulting the expert does not change, so the fact that the expert should be consulted less and less frequently is counter-intuitive. Figure 2(b) illustrates the change in the searcher's welfare as a function of the search cost, c_s , for different values of the service fee, c_e , charged by the expert. As expected, the searcher's welfare is better with the expert option than without, and the smaller the fee charged by the expert, the better the searcher's welfare.

Expected number of queries.

We can find the expected number of queries in this case by using our knowledge of the uniform distribution and the noise distribution in Equations 11-15, yielding

$$A = V^2\left(\frac{1}{t_l} - \frac{1}{t_u}\right); \quad B = t_l; \quad C = (1 - t_u);$$

$$D = t_u - t_l - V^2\left(\frac{1}{t_l} - \frac{1}{t_u}\right)$$

which give the final expressions:

$$\eta_{c_e} = \frac{t_l t_u (t_u - t_l)}{t_u t_l^2 - t_l V^2 - t_u t_l + t_u V^2}; \quad \eta_{c_s} = \frac{1}{1 - t_l - V^2\left(\frac{1}{t_l} - \frac{1}{t_u}\right)}$$

The monopolist expert’s optimal strategy.

Using the above derivations, it is now easy to calculate numerically the value of c_e^* that maximizes π_e , trading off decreasing number of queries η and increasing revenue per query c_e . Figure 3 shows examples of the graph of expected profit for the expert as a function of the expert’s fee, c_e , for different values of d_e , the marginal cost to the expert of producing an extra “expert report”, for $c_s = 0.01$.

Subsidizing the expert.

As discussed above, the market designer or the government can guarantee the reduction of the expert’s charge from c_e^* to c_e' , keeping π_e constant, by paying a per-consumer subsidy β to the expert. Figure 4 shows the improvement in social welfare, denoted $\delta(W)$, as a function of the subsidy paid to the expert, β , for various d_e values (where $c_s = 0.01$).

From the graphs, we do indeed find that subsidization can lead to substantial increases in social welfare, even when there is a significant marginal cost of producing an expert report. While this could be from a reduction in search and query costs or an increase in the expected value of the opportunity finally taken, the data in Table 1 indicates that the latter explanation is the dominant factor in this case. It is also worth noting that social welfare is maximized at the point where the searcher pays exactly d_e per query, thus fully internalizing the cost to the expert of producing the extra report. If the searcher had to pay less, it would lead to inefficient overconsumption of expert services, whereas if she had to pay more, the expected decline in the value she receives from participating in the search process would outweigh the savings to the market designer or government from having to pay less subsidy.

6. RELATED WORK

The autonomous agents literature has often considered the problem of search which incurs a cost [2, 9, 10]. The underlying foundation for such analysis is search theory [4, 15], and in particular, its one-sided branch which considers an individual sequentially reviewing different opportunities from which she can choose only one. The search incurs a cost and the individual is interested in minimizing expected cost or maximizing expected utility ([13, 7, 16], and references therein). To the best of our knowledge, none of the one-sided search literature in either search theory or multi-agent systems, has considered the resulting market dynamics when observations are not accurate and more accurate information can be purchased from a self-interested agent.

Relaxation of the perfect signals assumption is typically found in models of two-sided search [1], including marriage or dating markets [3] and markets with interviewing [11]. The literature has not to this point focused on the decision problems faced by self-interested knowledge brokers, or how their presence affects the market.

In terms of market interventions, the two-sided search literature has considered the impact of search frictions on labor markets (the 2010 Nobel Prize in Economics was awarded for this work [15, 4]). One classic regulatory intervention in these models is the introduction of a minimum wage, which can be shown to be welfare increasing in many contexts [8], but we are unaware of any work on subsidizing providers of expert knowledge, as we discuss here.

7. DISCUSSION AND CONCLUSIONS

The power of modeling markets using search theory is well established in the literature on economics and social science [12; 1, *inter alia*]. It has led to breakthroughs in understanding many domains, ranging from basic bilateral trade [17] to labor markets [15]. While knowledge has always been an economically valuable commodity, its role continues to grow in the Internet age. The ubiquity of electronic records and communications means there is an increasing role for knowledge brokers in today’s marketplaces. For example, it is now feasible for agencies like Carfax to collect the available records of every recorded accident, insurance claim, oil change, inspection, and so on for every car. The presence of such knowledge brokers necessitate that we take them into account in modeling the search process of consumers.

This paper takes the first step in this direction. We introduce a search model in which agents receive noisy signals of the true value of an object, and can pay an expert to reveal more information. We show that, for a natural and general class of distributions, the searcher’s optimal strategy is a “double reservation” strategy, where she maintains two thresholds, an upper and a lower one. When she receives a signal below her lower threshold, she rejects it immediately. Similarly, when she receives a signal above her upper threshold, she accepts it immediately. Only when the signal is between the thresholds does she consult the expert, determining whether to continue searching or accept the offer based on the information revealed by the expert.

In such models, there is scope for an authority like a market designer or regulator to improve social welfare by subsidizing the cost of querying the expert. The benefit to the searcher of having less friction in the process could potentially more than offset the cost to the authority. The downside would be that if the authority provided too much subsidy, this could lead to inefficient overconsumption of costly (to produce) expert services. By solving the model for a natural combination of the distribution of signals and the conditional distribution of the true value given the signal, we can analyze such questions more specifically. We show that in our example, subsidies can in fact be welfare-enhancing, and, in fact, social welfare is maximized when the searcher has to pay exactly the marginal production cost of expert services. Both the model and our results are significant for designers of markets in which consumers will search and the need for expert services will arise naturally (like an Internet marketplace for used cars), because by enhancing social welfare, the market designer can take market share away from competitors, or perhaps charge higher commissions, because it is offering a better marketplace for consumers.

There are several directions for future research, from both the expert’s perspective and the market-designer’s. An interesting problem for the monopolist expert is the optimal pricing of bundles of queries, where agents must purchase a bundle, instead of individual reports. More realistic modeling of “startup costs” and hence the average “supply curve” of expert services (instead of the marginal cost considered here) may also explain a richer range of behaviors. The existence of the expert has ramifications beyond one-sided search, our focus in this paper. For example, in two-sided search markets like labor markets, there may be different types of experts: those who conduct background checks, for example, or providers who run independent testing services to vet potential employees. What are their incentives, and

d_e	Without Subsidy						With Subsidy						
	c_e	$\eta_{c_e} c_e$	$c_s \eta_s$	Worth	V	W	subsidy	c_e	$\eta_{c_e} c_e$	$c_s \eta_s$	Worth	V	W
0	0.06	0.10	0.08	0.7472	0.5653	0.6656	0.1003	0	0	0.1036	0.7928	0.6893	0.6893
0.025	0.06	0.10	0.08	0.7472	0.5653	0.6238	0.0585	0.025	0.05	0.09	0.7712	0.6295	0.6295
0.05	0.061	0.09	0.08	0.7365	0.5637	0.5806	0.0169	0.05	0.09	0.08	0.7536	0.5824	0.5824

Table 1: The different components of social welfare with and without subsidy for $c_s = 0.01$. “Worth” is the expected value of the opportunity eventually picked. Initially the decrease in query cost contributes more to the increase in social welfare, but as d_e increases, this contribution becomes less significant. Note that the first two columns in the case without subsidy are similar because the profit-maximizing c_e is the same and the searcher’s cost depends only on value of selected c_e , not d_e .




Figure 4: Increase in social welfare vs subsidy. When there is no marginal cost ($d_e = 0$), it is better for the market designer to make the service available for free but when there is some marginal cost involved, then increase in social welfare is a concave peaked at marginal cost.

how do these affect two-sided search markets? From the market designer’s point of view, new alternatives to subsidization as a means for improving social welfare can be explored, e.g., inducing competition, or provision of expert services by the market designer herself (e.g., a government takes over the role of providing expert services).

8. ACKNOWLEDGMENTS

We are grateful to the US-Israel BSF for supporting this research under Grant 2008404. Das also acknowledges support from an NSF CAREER award (0952918), and Sarne is partially supported by ISF grants 1401/09.

9. REFERENCES

- [1] K. Burdett and R. Wright. Two-sided search with nontransferable utility. *Review of Economic Dynamics*, 1:220–245, 1998.
- [2] S. Choi and J. Liu. Optimal time-constrained trading strategies for autonomous agents. In *Proc. MAMA*, 2000.
- [3] S. Das and E. Kamenica. Two-sided bandits and the dating market. In *Proc. IJCAI*, 2005.
- [4] P. Diamond. Aggregate demand management in search equilibrium. *The Journal of Political Economy*, 90(5):881–894, 1982.
- [5] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [6] J. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73, 1966.
- [7] A. Grosfeld-Nir, D. Sarne, and I. Spiegler. Modeling the search for the least costly opportunity. *European Journal of Operational Research*, 197(2):667–674, 2009.
- [8] L. Ho. Wage subsidies as a labour market policy tool. *Policy Sciences*, 33(1):89–100, 2000.
- [9] J. Kephart and A. Greenwald. Shopbot economics. *JAAMAS*, 5(3):255–287, 2002.
- [10] J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32:731–752, 2000.
- [11] R. Lee and M. Schwarz. Interviewing in two-sided matching markets. *NBER Working Paper*, 2009.
- [12] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189, 1976.
- [13] J. McMillan and M. Rothschild. Search. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. 1994.
- [14] P. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.
- [15] D. Mortensen and C. Pissarides. New developments in models of search in the labor market. *Handbook of labor economics*, 3:2567–2627, 1999.
- [16] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82:689–711, 1961.
- [17] A. Rubinstein and A. Wolinsky. Middlemen. *The Quarterly Journal of Economics*, 102(3):581–593, 1987.
- [18] M. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 47(3):641–654, 1979.
- [19] R. Wright. Job search and cyclical unemployment. *The Journal of Political Economy*, 94(1):38–55, 1986.