Two-Sided Search With Experts

YINON NAHUM, Weizmann Institute of Science
DAVID SARNE, Bar-Ilan University
SANMAY DAS, Rensselaer Polytechnic Institute
ONN SHEHORY, IBM Research

In this paper we study distributed agent matching in environments characterized by uncertain signals, costly exploration, and the presence of an information broker. Each agent receives information about the potential value of matching with others. This information signal may, however be noisy, and the agent incurs some cost in receiving it. If all candidate agents agree to the matching the team is formed and each agent receives the true unknown utility of the matching, and leaves the market. We consider the effect of the presence of information brokers, or experts, on the outcomes of such matching processes. Experts can, upon payment of a fee, perform the service of disambiguating noisy signals and revealing the true value of a match to any agent. We analyze equilibrium behavior given the fee set by a monopolist expert and use this analysis to derive the revenue maximizing strategy for the expert as the first mover in a Stackelberg game. Surprisingly, we find that better information can hurt: the presence of the expert, even if the use of its services is optional, can degrade both individual agents' utilities and overall social welfare. While in one-sided search the presence of the expert can only help, in two-sided (and general k-sided) search the externality imposed by the fact that others are consulting the expert can lead to a situation where the equilibrium outcome is that everyone consults the expert, even though all agents would be better off if the expert were not present. As an antidote, we show how market designers can enhance welfare by taxing use of expert services.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Algorithms, Economics

Additional Key Words and Phrases: Distributed Matching, Information Brokers

1. INTRODUCTION

This paper studies agent partnership or team formation using the framework of search theory. Two-sided search has been used to model labor markets [Jovanovic 1979; Lippman and McCall 1976], marriage and mating [Bloch and Ryder 2000], and partnership formation among artificial agents [Sarne and Kraus 2008]. The typical assumption in models of two-sided search is that potential partners are matched through some mechanism, and then each of the partners receives a signal that informs her of the value *to her* of that match [Burdett and Wright 1998; Chade and Ventura 2005]. For example, in the case of employers and workers, a worker is informed of the wage and the relevant non-wage characteristics of the job, while the employer is informed of the productivity of the worker. However, in many realistic situations, this information is not available when the initial matching occurs. Therefore, many have recently tried to model the explicit process of pairs learning about each other, whether through a one-shot interview

This work is supported in part by a US-Israel BSF Grant (#2008-404) to Das and Sarne. Sarne acknowledges additional support from ISF grant 1401/09. Das acknowledges additional support from an NSF CAREER Award (IIS-0952918).

process [Lee and Schwarz 2009] or through repeated interactions like dating [Das and Kamenica 2005].

Agents may also learn about the quality of a matching by paying an information intermediary, or an *expert* to conduct research on the quality of a potential match and share the information with that agent prior to the agent having to decide on whether to accept the match. Examples of these kinds of experts abound in real life. For example, headhunters for corporations, dating services, private investigators, or contractors that conduct extensive background checks all serve as experts (some of them perform the additional function of matchmaking by being part of the technology for arranging potential pairings).

We explicitly analyze the impact of the presence of such experts on two-sided search markets (and also extend our analysis to general k-sided search for team formation). The central question here is whether the presence of such experts helps or hurts individuals who participate in such markets. We start from a standard model of two-sided search [Burdett and Wright 1998; Chade and Ventura 2005], in which agents are randomly paired and then have to decide whether or not to accept the matching once their potential value from the partnership is revealed to them. We consider the impact on search when agents receive noisy signals of the value of the match which can be disambiguated into perfect signals upon payment of a fee to the expert. Such experts have previously been considered in the context of one-sided search (i.e., when the agent does not depend on the decision of others). There, the presence of experts results in increased social welfare, even if they function as monopolists [Chhabra et al. 2011]. Such one-sided search with experts can be modeled as a relatively simple Stackelberg game, where the expert moves first by setting the price for her services, and searchers respond by following their optimal search strategies. In two-sided search models, however, the outcome of the system is more complex, because one has to consider equilibrium behavior of the searchers under a particular cost structure, rather than just solving a single-agent optimization problem. In this study we characterize the equilibrium for any given cost of consulting the expert when agents share the same exploration costs, and the distribution of valuations is common to the agents; we use our equilibrium characterization to derive the expert's optimal cost structure.

Some of our results are qualitatively similar to those found when considering the impact of experts in one-sided search. For example, the form of the optimal strategy is a similar "double reservation" one: the searcher only queries the expert when signals are between two thresholds, automatically rejecting opportunities with signals below the lower threshold and accepting those with signals above the higher threshold; further, the two reservation values get closer to each other as search costs and query costs increase.

However, some results are starkly different from those found in one-sided search. The need for agents to reason about the optimal behavior of others changes some of the major systemic properties of two-sided search markets. Surprisingly, we show that the additional information provided by the expert can sometimes hurt market participants: it can be an equilibrium for everyone to consult the expert even though everyone would be better off if they agreed that no one should consult the expert, or the option did not exist. The negative effects are not distributional: it is not the case that some agents are benefiting at the expense of others. Instead, the new search frictions introduced by experts outweigh any benefits they provide. Thus it is important to recognize this negative effect, and we suggest a means of correcting it from a market designer's perspective: introducing a (Pigovian) tax on expert services, or equivalently subsidizing the expert to *increase* the fees she charges searchers, so that they consult her less frequently. Using a synthetic environment we demonstrate how the reverse subsidy can actually improve overall social welfare.

2. RELATED WORK

This paper touches on several different literatures, but is primarily grounded in the theory of sequential distributed two-sided matching. The autonomous agents literature has engaged with the problem of costly search [Kephart and Greenwald 2002; Kephart et al. 2000; Sarne and Kraus 2008; Chhabra et al. 2011], in particular in the absence of a central information source which provides instant reliable information on other agents, their availability and states, and the environment. The introduction of search costs into multiagent systems (MAS) models leads to a more realistic description of MAS environments. In particular, search costs are known to be important in electronic commerce environments where agents need to invest/consume some of their resources in order to obtain information concerning the good or the transaction offered by other prospective agents [Bakos 1997; Kephart and Greenwald 2002].

The underlying foundation for costly search analysis is the theory of sequential search [McCall 1970; Diamond 1982; Mortensen and Pissarides 1999]. Individuals are considered to be sequentially reviewing different opportunities where search incurs a cost and the individual is interested in minimizing expected cost or maximizing expected utility ([McMillan and Rothschild 1994; Grosfeld-Nir et al. 2009; Rothschild 1974], and references therein). In an effort to understand the effect of dual search activities in costly environments, "two-sided" search models have been developed [Chade and Ventura 2005; Sarne and Kraus 2008; Smith 2011]. Unlike stable matching scenarios [Gale and Shapley 1962; Anshelevich et al. 2011] which do not involve costly search, equilibrium in two-sided search stems from the existence of search costs because of which searchers are reluctant to resume their search for potentially better outcomes. The assumptions we rely on are standard in the two-sided search literature, and in particular our noisy signal model augments the work of [Burdett and Wright 1998; Chade and Ventura 2005], integrating uncertain observations.

While there has been some work on noisy signals in two-sided search, this has typically involved the assumption that agents can get better signals through repeated interactions (e.g., dating, cohabitation, or interviews [Das and Kamenica 2005; Sahib and Gu 2002; Lee and Schwarz 2009]). Mediators in two-sided search models have typically taken the form of matchmakers rather than information brokers [Bloch and Ryder 2000]. Instead we follow the work of Chhabra et al. [2011] in focusing on self-interested knowledge brokers, and how their presence affects the market.

Partnership formation has also been studied in the context of MAS coalition formation (e.g., in [Shehory and Kraus 1998]). The work in this paper differs from classical coalition formation in several ways. First, here formation is gradual and iterative: at each stage only a single partnership is formed (resembling some ideas in early work by Ketchpel et al [Ketchpel 1994]); second, values of teams are not known in advance and their revelation is costly (somewhat resembling elements of [Kraus et al. 2003]); third, each team is formed in isolation, disregarding externalities and other formations; fourth, the market is large and possibly infinite, unlike typical coalition formation work.

3. MODEL

The model is based on a standard two-sided distributed search model [Burdett and Wright 1998; Chade and Ventura 2005; Sarne and Kraus 2008], augmented to include uncertain signals. The model assumes fully rational self-interested agents, searching for appropriate partners to form mutually acceptable pair-wise partnerships.[1]

[1] For simplicity of presentation, and in line with the literature, we focus on partnerships of size two. The extension to teams of any size k is straightforward, as described in Section 4.3.

The number of agents may be either infinite or finite and all agents are ex ante identical, in that there are no individuals who are "naturally" better than others or more easy to please than others. However, when a potential match is formed, each agent gets some idiosyncratic utility from the particular qualities of that partnership (each agent's utility is drawn independently). This utility is drawn anew each time a partnership with the same agent is evaluated in later stages of the search (as the number of agents in the population grows large, this becomes increasingly unlikely, since potential partnerships are drawn at random from the population; however, even with a relatively small number of agents, it models cases where the utility of a partnership is dependent on the circumstances in which it is formed).

At any period, the matching technology arranges a meeting between two agents, each of whom pays a search cost c_s and receives a different, independent noisy signal, denoted s, that indicates the estimated value of the match to it. We assume that agents are acquainted with the distribution of signals $f_s(s)$ and the conditional probability density of values given signals, $f_v(v|s)$.[2]

Upon receiving a signal, an agent can either accept the partnership, decline it, or pay a cost c_e to consult an expert who then reveals to that agent the (noiseless) true value of the partnership to that agent. If the agent does consult the expert, it must decide whether to accept or decline the partnership once it receives the true value. If both agents decide to accept the partnership, a match takes place and the agents leave the market. If either one of the agents declines the partnership, the agents go back into the searching population and continue their search by sampling another partnering opportunity at search cost c_s, and so on.

Since the agents are fully rational and self-interested, their goal is to maximize their expected utility, defined as the value they receive from the partnership they eventually form minus the accumulated costs of querying the expert and interacting with other agents along the search path. In addition, the expert is a rational, self-interested monopolist; her goal is to maximize her own expected utility, defined as the accumulated payment she receives from the agents minus her expenses, denoted d_e, which are a function of the cost of producing the information required to inform agents of the exact values of matches.

4. ANALYSIS

In this section we derive the agents' and expert's strategies. We first derive the individual agent's utility-maximizing strategy given the strategies used by the other agents and the fee set by the expert. From that we derive the equilibrium strategy of the agents given the expert's fee, and finally the utility-maximizing strategy for a monopolist expert. We start from the framework of Chhabra *et al* [Chhabra et al. 2011], who characterize optimal search strategies in a one-sided search model. We then consider the effect of the two-sided nature of the search on searchers' strategy. The analysis is augmented to the general partnership-size case in a straightforward way, as shown towards the end of this section.

4.1. Preliminaries

As described in Section 3, we consider a specific model of noisy search in which searchers encounter opportunities (partnerships) sequentially, and receive noisy signals of the true value of each opportunity. Similar noisy signal models have been considered in the one-sided search literature [Chhabra et al. 2011; Wiegmann et al. 2010]

[2]Alternatively, one can assume that searchers are acquainted with the distribution of values from which a partnership's values are drawn, $f_v(v)$, and the conditional distribution of signals given the values, $f_s(s|v)$. These are interchangeable by Bayes' Law.

and perfect information models have been considered in the two-sided search literature [Burdett and Wright 1998; Chade and Ventura 2005; Sarne and Kraus 2008]. Our analysis builds upon this line of work. In this subsection we briefly summarize existing definitions and results. For the one-sided search results with experts, we follow Chhabra *et al* [Chhabra et al. 2011], and for two-sided search with perfect signals we follow Burdett and Wright [Burdett and Wright 1998].

One-sided search. The optimal strategy in many models of search is a *reservation value* (i.e., a threshold-based) strategy, where the searcher accepts any opportunity that is higher than a particular reservation value – intuitively, the reservation value is the expected utility of rejecting an opportunity and continuing search. When agents receive noisy signals instead of perfect information about the value of an opportunity, the optimal strategy need not be a reservation value strategy, because the correlation structure between signals and true values may be peculiar. However, assuming a simple stochastic dominance assumption on the signal structure [Wright 1986; Milgrom 1981], it can be shown that the optimal strategy is, in fact, a reservation value strategy in the one-sided case. We restate the assumption and a useful corollary:

Definition 4.1. **Higher signals are good news (HSGN) assumption:** If $s_1 > s_2$, then, $\forall y, F_v(y|s_1) \leq F_v(y|s_2)$.

where $F_v(y|s)$ is the cumulative distribution function (cdf) of values given the signal.

COROLLARY 4.2. *For noisy environments satisfying the HSGN assumption, if* $s_1 > s_2$*, then,* $E[v|s_1] \geq E[v|s_2]$.

The introduction of an expert, who can provide for a fee a perfect signal of the true value of an opportunity, extends the number of decision alternatives available to the agent performing the search based on the noisy signal. This agent can now (1) reject the opportunity without querying the expert, paying search cost c_s to reveal the signal for the next potential opportunity; (2) query the expert to obtain the true value v, paying a cost c_e, and then decide whether to resume search or not; or (3) accept the current opportunity without querying the expert, receiving the (unknown) true value of the opportunity. Under the HSGN assumption, Chhabra *et al* show that the optimal strategy for a searcher is characterized by a tuple (t_l, t_u, V) such that the searcher rejects the opportunity for all signals below t_l, accepts the opportunity without querying the expert for all signals above t_u, and queries the expert for all signals between t_l and t_u, accepting if and only if the revealed value v satisfies $v > V$.

Equilibrium in two-sided search . For the one-sided results described above, it is assumed that opportunities arise exogenously, and the searcher is free to take an opportunity once it arises. In the case of partnership or team formation, however, the process of matching is dependent on both parties (assuming pairwise partnerships) agreeing to the match. Therefore, even if an opportunity is acceptable to an agent, the match may not form, since the agent may not be acceptable to the proposed partner.

Under reasonable assumptions, it can be shown that, in equilibrium, the optimal strategy for agents engaging in distributed two-sided matching with perfect signals is a reservation value strategy [Sarne and Kraus 2008; Burdett and Wright 1998]. Opportunities arise sequentially, in random order; each agent reviews these and terminates the search once a value greater than a reservation value x^* is revealed. If agents are homogeneous in the sense that they all share the same search cost c_s and their values from a partnership derive from the same distribution function $f_v(v)$, then they all use the same reservation value in equilibrium. The equilibrium reservation value is the value that maximizes utility when all other agents are using that value. If all other agents are using a reservation value x^*_{others}, then the reservation value that maximizes utility for any individual agent, x^*, satisfies [Burdett and Wright 1998]:

$$c_s = \left(1 - F(x^*_{\text{others}})\right) \int_{y=x^*}^{\infty} (y - x^*) f(y) dy \tag{1}$$

Equation 1 can be interpreted as comparing the cost of any additional search round with the expected marginal utility from obtaining an additional value. The reservation property of the optimal strategy derives from the stationarity of the problem – since the searcher is not limited by the number of opportunities it can explore, resuming search places her at the same position as at the beginning of the search [McMillan and Rothschild 1994]. 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 utility from the search when using x^*, denoted $V(x^*)$, satisfies $V(x^*) = x^*$, because the reservation value x^* is the value where the searcher is indifferent between accepting the current value x^* and resuming the search process (yielding expected utility $V(x^*)$). This can formally be proved by solving Equation 1 using integration by parts.

Due to symmetry, all the agents use the same reservation value in equilibrium and therefore $x^* = x^*_{\text{others}}$, resulting in:

$$c_s = \left(1 - F(x^*)\right) \int_{y=x^*}^{\infty} (y - x^*) f(y) dy \tag{2}$$

The expected number of search iterations is simply the inverse of the success probability, $1/\left(1 - F_v(x^*)\right)^2$, since this becomes a Bernoulli sampling process, as opportunities arise independently at each iteration.

4.2. Two-Sided Search with Noisy Signals

We can incorporate noisy signals into the two-sided search model above by first characterizing the utility-maximizing strategy of each individual searcher and then finding the resulting equilibrium. As in the one-sided case discussed above, when the searcher receives a noisy signal rather than a perfect one, there is no guarantee that the optimal strategy is reservation-value based. The problem is still stationary though, and an opportunity that has been rejected will never be recalled. In the absence of restrictions 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 utility of search, denoted $V(S, S^*)$, can then be written as (S^* is the signal-value intervals for which the other agents terminate search):

$$\begin{aligned} V(S,S^*) &= -c_s + \left(1 - \Pr(s \in S)\Pr(s^* \in S^*)\right)V(S,S^*) + \Pr(s^* \in S^*)\Pr(s \in S)E[v|s \in S] \\ &= -c_s + V(S,S^*)\left(1 - \left(\int_{s\in S} f_s(s)\,ds\right)\left(\int_{s^*\in S^*} f_s(s^*)\,ds^*\right)\right) \\ &\quad + \left(\int_{s^*\in S^*} f_s(s^*)\,ds^*\right)\int_{s\in S} f_s(s)E[v|s]\,ds \end{aligned} \tag{3}$$

Here, the value of $V(S, S^*)$ is derived recursively from performing one additional search iteration. The first element on the right is the cost of the search iteration. The second element applies to the case where search continues, and is composed of the probability that at least one of the sides rejects the match, multiplied by the expected value of the continued search, which is again, due to the stationarity of the problem, $V(S, S^*)$. The third element applies to the case where search terminates, and is composed of the probability of being accepted by the other side, multiplied by the expected value of accepting the match.

Assuming that higher signals are good news enables us to prove that the equilibrium strategy is a reservation rule.

THEOREM 4.3. *If the conditional distribution of values given signals, $f_v(v|s)$, satisfies the HSGN assumption, then:*
(a) The equilibrium search strategy of any individual agent is a reservation-value rule, where the reservation value, t^*, satisfies:

$$c_s = (1 - F_s(t^*)) \int_{s=t^*}^{\infty} \Big(E[v|s] - E[v|t^*]\Big) f_s(s)\, ds \tag{4}$$

where $F_s(t^*)$ is the cumulative distribution function (cdf) of signals.
(b) The equilibrium expected utility to an agent of using the optimal search strategy satisfies: $V(t^*) = E[v|t^*]$.

Sketch of 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 if signal s is obtained. Since the optimal strategy given signal s is to resume search, we know $V > E[v|s]$. Given the HSGN assumption, $E[v|s] \geq E[v|s']$ 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 can be explicitly stated. Setting the first derivative according to t of the new equation to zero we obtain: $V(t^*) = E[v|t^*]$ (and verifying that t^* is global maximum by calculating the second derivative). Finally substituting $V(t^*) = E[V|t^*]$ in the expected value of the searcher equation obtains Equation 4. □

Note that the condition $V(t^*) = E[v|t^*]$ implies that the reservation value t^* is the signal for which the searcher's utility of resuming search is equal to the expected value of the opportunity associated with that signal. The expected number of search iterations in this case is $1/\left(1 - F_v(t^*)\right)^2$, since this is again a Bernoulli sampling process.

4.3. Two-Sided Search With an Expert

Suppose that any searcher can query an expert at cost c_e to find out the true value (to her) of a potential partner. Now, as in the one-sided case above, the searcher has 3 alternatives. She can (1) reject the current potential partnership without querying the expert, paying search cost c_s to reveal the signal for the next potential partnership; (2) query the expert to obtain the true value v, paying a cost c_e, and then decide whether to accept the partnership with the other searcher; or (3) accept the current partnership without querying the expert. If both potential partners accept then the search terminates. Case (2) termination provides the searcher with the true value v. Case (3) termination provides the searcher with the (unknown) true value of the partnership. With no mutual acceptance, the search resumes.

As in the no-expert case, a solution for a general density function $f_v(v|s)$ dictates an optimal strategy with a complex structure of 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 accept the partnership without querying the expert; and (c) for any signal that is not in S' or S'' the searcher should query the expert, and accept the partnership if the value obtained is above a threshold V, and resume otherwise. The value V is the expected utility from resuming the search and is given by the following modification of Equation 3, given that the other agents use strategy $(S'_{\text{others}}, S''_{\text{others}}, V_{\text{others}})$:

$$V(S', S'', V) = -c_s - c_e \int_{s \notin \{S', S''\}} f_s(s)\, ds + (1 - A \cdot B) \cdot V(S', S'', V) + B \cdot C \tag{5}$$

where A is the probability that the searcher accepts the partnership eventually (either directly or after querying the expert), B is the probability that the potential partner accepts the match, and C is the searcher's expected utility if both sides accept the partnership; these are given by:

$$A = \int_{s \in S''} f_s(s)\, ds + \int_{s \notin \{S', S''\}} f_s(s)\Big(1 - F_v(V|s)\Big)\, ds$$
$$B = \int_{s \in S''_{\text{others}}} f_s(s)\, ds + \int_{s \notin \{S'_{\text{others}}, S''_{\text{others}}\}} f_s(s)\Big(1 - F_v(V_{\text{others}}|s)\Big)\, ds$$
$$C = \int_{s \in S''} f_s(s) E[v|s]\, ds + \int_{s \notin \{S', S''\}} \left(f_s(s) \int_{y=V}^{\infty} y f_v(y|s)\, dy \right) ds$$

The value of $V(S', S'', V)$ in Equation 5 is once again derived recursively, considering the next search iteration. The searcher pays c_s for receiving the noisy signal. The next element is the expected expert query cost, incurred whenever receiving a signal $s \notin \{S', S''\}$. The third element applies to the case of resuming search, when at least one of the sides rejects the partnership, in which case the searcher continues with an expected utility $V(S', S'', V)$. The last element applies to the case where the search is terminated, since both sides accepted the opportunity. Similarly, the first element in A and B applies to a case where the searcher accepted the match without querying the expert and the second applies to a case where the searcher accepted the match after querying the expert. The first element in C applies to a case where the searcher accepted the match without querying the expert, in which case the expected revenue is $E[v|s]$. The second element applies to the case where the searcher accepted the match after querying the expert.

Based on the above, we can prove that, similar to the one-sided case, under the HSGN assumption, each of the sets S' and S'' actually contains a single interval of signals.

THEOREM 4.4. *For $f_v(y|s)$ satisfying the HSGN assumption (Definition 4.1), the utility-maximizing search strategy of an agent, given the search cost c_s, the fee c_e set by the expert and the search strategies used by the other agents, 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 partnership should be accepted; and (c) the expert should be queried given any signal $t_l < s < t_u$ and the partnership should be accepted if the value obtained from the expert is above the expected value of resuming the search, V, otherwise search should resume. The equilibrium values t_l, t_u*

and V can be calculated from solving the set of Equations 6-11:

$$V = \frac{-c_s - c_e\big(F_s(t_u) - F_s(t_l)\big) + B \cdot C}{A \cdot B} \tag{6}$$

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

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

$$A = 1 - F_s(t_l) - \int_{s=t_l}^{t_u} f_s(s) F_v(V|s) ds \tag{9}$$

$$B = A \tag{10}$$

$$C = \int_{s=t_u}^{\infty} f_s(s) E[v|s] ds + \int_{s=t_l}^{t_u} \Big(f_s(s) \int_{y=V}^{\infty} y f_v(y|s) dy \Big) ds \tag{11}$$

Sketch of Proof: The proof extends the methodology used for proving Theorem 4.3. 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 $s'' > s$. Equations 6, 9, 11 are obtained after replacing the intervals S, S' with the thresholds t_l, t_u. Equation 10 represents the fact that the system is symmetric in the way that ultimately all agents choose the same tuple (t_l, t_u, V), and so the probability of being accepted is equal to the probability of accepting the match. Finally, the correctness of Equations 7 and 8 is proved by taking the derivative of Equation 6 w.r.t. t_l and t_u, equating to zero, obtaining unique t_l and t_u which maximize the expected benefit. □

We now have 6 Equations 6 - 11 in 6 variables. We can solve these simultaneously to calculate the value of V, t_l, t_u.

The intuitive interpretation of each of the above equations is as follows. Equation 6 captures the expected utility of searchers if resuming their search with strategy (t_l, t_u, V). Equation 7 captures the indifference of the searcher between querying the expert and resuming the search when receiving a signal t_l — if the searcher resumes search it receives V, however if querying the expert then it receives either $\int_{y=V}^{\infty} y f_v(y|t_l)$ (with probability B) or otherwise V:

$$V = -c_e + V\Big(1 - B\big(1 - F_v(V|t_l)\big)\Big) + B \int_{y=V}^{\infty} y f_v(y|t_l)\, dy$$

which transforms into Equation 7.

Similarly, Equation 8 captures the searcher's indifference, given a signal t_u, between querying the expert and accepting the partnership without querying the expert. Here, if accepting the partnership without querying the expert the searcher obtains $E[v|t_u]$ with probability B and V otherwise:

$$B \cdot E[v|t_u] + (1 - B)V = -c_e + V\Big(1 - B\big(1 - F_v(V|t_u)\big)\Big) + B \int_{y=V}^{\infty} y f_v(y|t_u)\, dy$$

which transforms into Equation 8.

Equations 9 and 11 are the appropriate modifications of A and C (from Equation 5) for the case of using strategy (t_l, t_u, V). Equation 10 derives from the fact that the agents are homogeneous, thus they all use the same set of reservation thresholds in equilibrium.

There is also a degenerate but plausible 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 partnership 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.

It is straightforward to extend the above analysis in order to encompass additional model assumptions and variations. We demonstrate how this is done for the following cases: (a) the teams formed are of a general size; (b) agents discount future utility; (c) agents come from different populations, differing in search costs and/or the distribution of valuations of partnerships.

Extension to k-sided search. Assume that instead of getting acquainted with one other agent at a particular time instant, the agent meets $k-1$ other agents at a time, interested in forming a group of size k (e.g., instead of pairs, students need to divide into groups of four). Assuming the group will be formed only if all agents accept it, then the only required change in the equations is to Equation 10 which turns into $B = A^{k-1}$.

Further, even if the size of the coalition encountered at each stage of the search varies (e.g., entrepeneurs meet each other at each time period to consider a new start-up) and can be captured by the probability function $P_{\text{size}}(i)$, then the equilibrium can be calculated using a simple modification of the above, by changing Equation 10 to $B = \sum_{i=1}^{\infty} P_{\text{size}}(i) A^{i-1}$.

For $i = 1$ we recover the equations of expert-mediated one-sided search [Chhabra et al. 2011], confirming that the latter is a specific case of our model where the searcher is always accepted.

Discounting future utility. Adding time discounting to the model is straightforward and does not qualitatively change the results. Assume that gains from the partnerships formed are discounted according to a discount factor δ (and so are costs paid). In keeping with the sequential search literature [McMillan and Rothschild 1994], we assume that gains are received at the end of a search round whereas search costs are paid at the beginning of a search round. In this case, we can prove that agents will use double-reservation strategies, based on the tuple (t_l, t_u, V) however with different values. While the discounting does not explicitly affect Equations 9 and 10 it does affect Equations 6 and 11 which become:

$$V = \frac{-c_s - c_e\big(F_s(t_u) - F_s(t_l)\big) + B \cdot C}{1 - \delta(1 - A \cdot B)} \tag{12}$$

$$C = \delta\left(\int_{s=t_u}^{\infty} f_s(s) E[v|s] ds + \int_{s=t_l}^{t_u} \Big(f_s(s) \int_{y=V}^{\infty} y f_v(y|s) dy\Big) ds\right) \tag{13}$$

Equating the first derivative of (12), according to t_l and t_u (separately) to zero, obtains the following modifications of Equations (7)-(8)

$$c_e = \delta B \int_{y=V}^{\infty} (y - V) f_v(y|t_l)\, dy \tag{14}$$

$$c_e = \delta B \int_{y=-\infty}^{V} (V - y) f_v(y|t_u)\, dy \tag{15}$$

The intuitive interpretation of Equations (14)-(15) remains similar to the interpretation above for Equations (7)-(8).

Extension to different agent populations. Assume that partnerships are formed between agents of different populations (e.g., men and women, employers and employees) [Burdett and Wright 1998], differing in the distribution of values and the costs of search and costs of querying the expert the agents of each population are associated with. We use c_s^i, c_e^i, $f_s^i(s)$ and $f_v^i(v|s)$ to denote the search cost, the expert querying cost, the signal probability distribution and the value distribution given a signal of searchers of population i (for $i = 1, 2$), respectively. In this case, for the same considerations used above (when all agents are of the same population) we can prove that agents from population i will be using the double-reservation strategy, based on the tuple (t_l^i, t_u^i, V^i). The equilibrium set of strategies $\{(t_l^1, t_u^1, V^1), (t_l^2, t_u^2, V^2)\}$ will be obtained by solving the augmented set of equations:

$$V^i = \frac{-c_s^i - c_e^i\big(F_s^i(t_u^i) - F_s^i(t_l^i)\big) + B^i \cdot C^i}{A^i \cdot B^i} \tag{16}$$

$$c_e^i = B^i \int_{y=V^i}^{\infty} (y - V^i) f_v^i(y|t_l^i)\, dy \tag{17}$$

$$c_e^i = B^i \int_{y=-\infty}^{V^i} (V^i - y) f_v^i(y|t_u^i)\, dy \tag{18}$$

$$A^i = 1 - F_s^i(t_l^i) - \int_{s=t_l^i}^{t_u^i} f_s^i(s) F_v^i(V^i|s) ds \tag{19}$$

$$B^i = A^{3-i} \qquad \{\text{note that } 3 - i \text{ here is an index rather than power}\} \tag{20}$$

$$C^i = \int_{s=t_u^i}^{\infty} f_s^i(s) E^i(v|s) ds + \int_{s=t_l^i}^{t_u^i} \Big(f_s^i(s) \int_{y=V^i}^{\infty} y f_v^i(y|s) dy\Big) ds \tag{21}$$

for $i = 1, 2$

4.4. Expert's Profit Maximization

We can view the search process as a Stackelberg game, where a monopolist expert moves first by setting her query cost c_e. Searchers respond by following their equilibrium strategies described above. Therefore, the expert can solve for searcher behavior, given knowledge of the search cost c_s and the signal and value distributions. The expert should set her fee to maximize profit, defined as the product of the expected number of times her services are used by the searchers, and the profit she makes per query.

Expected number of queries:. The search strategy (t_l, t_u, V) defines the number of times the expert's services are used. For an agent's search to end, both sides need to accept the match. The probability of that happening is $A \cdot B$, using the notation above. Since this is a geometric distribution, the expected number of search iterations an agent performs is $\frac{1}{A \cdot B}$. In each search iteration, the probability of a searcher querying the expert is $F_s(t_u) - F_s(t_l)$, and so the expected number of expert queries a searcher performs, denoted η_{c_e} is

$$\eta_{c_e} = \frac{F_s(t_u) - F_s(t_l)}{A \cdot B} \tag{22}$$

Expected profit of the expert:. The expected profit of the expert is: $\pi_e = \mathbb{E}(\text{Profit}) = (c_e - d_e)\eta_{c_e}$. The expert can maximize the above expression with respect to c_e to find the profit maximizing price to charge searchers (note that the expert only needs to perform the computation for a single searcher).

4.5. Market Design: Subsidization and Taxation

One possible use of the theory described above is to improve the design of markets in which such two-sided search takes place. Consider a platform that allows searchers and experts to interact – for example, an online dating website, or one that brings together developers looking to invest in a housing project. The platform has a privileged position and can either act as the expert itself, or outsource the expertise function, but use its position to negotiate the conditions under which expert services are provided and used.

Chhabra et al. [2011] consider the possible impact of subsidizing the expert in one-sided search markets. In their framework, a platform or market designer can pay the expert to decrease her query cost with the goal of increasing social utility. This makes sense because in one-sided search the presence of an expert is necessarily beneficial, because a searcher can simply ignore the option of consulting the expert if it is not beneficial. However, in two (or more)-sided search, game theoretic considerations come into play, and we must consider the possibility that the presence of an expert may not be helpful (in fact, we will show that this is possible in the next section). In such cases, the platform may choose to tax the interaction between the searcher and the expert in order to increase the cost to the searcher of consulting the expert.

A general framework that encompasses both these market design "levers" is as follows. Suppose a monopolist provider of expert services maximizes her profits by setting the query cost to c_e^*, yielding an expected profit $\pi_e = (c_e^* - d_e)\eta_{c_e^*}$. The market designer can effectively change the fee paid by the searcher from c_e to a value c_e'. If $c_e' < c_e$, this could be accomplished through a subsidy, while if $c_e' > c_e$ this could be accomplished through a transaction tax. Presumably the market designer is doing this for the overall benefit of searchers, but may have to compensate the expert. The market designer can offer a per-query payment β to the expert, which fully compensates the expert for the decreased revenue, leaving her total profit unchanged (in the case of a tax, the entire amount of the tax could also be paid to the expert). The compensation for a requested change in the searcher's payment 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 Equations 6-11, when the expert fees are $c_e^{'}$ and c_e^* respectively, at a cost β to the market designer.

The social welfare is given by the sum of utilities of all parties involved. Thus far, we have just considered two: the searcher and the expert (this generalizes to multiple searchers as well):

$$W = V_{c_e^*} + \pi_e \tag{23}$$

When the market designer subsidizes or taxes expert queries, the social welfare must also take into account the subsidy. Since the expert is fully compensated for her loss due to the decrease or increase 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 $W^{'} = V_{c_e'} + \pi_e - \beta$.

5. ILLUSTRATIVE EVALUATION

In this section we illustrate the properties of our model by numerically examining and depicting its behavior in different settings. Some of the results are quite surprising, and may enable more efficient market designs, beneficial to searchers, experts, and the interaction platform as a whole.

For the numerical study, we use a synthetic environment, where agents form pair-wise partnerships (or k-wise teams, in some settings described in detail where they occur). The signal is an upper bound on the true value (e.g., people tend to get a good first impression of others). Specifically, we assume signals s are uniformly distributed on $[0, 1]$ ($f_s(s) = 1$ if $0 < s < 1$ and zero otherwise) and the conditional density of true values is a monotonic increasing function in the interval $[0, s]$: $f_v(y|s) = \frac{3\sqrt{y}}{2\sqrt{s^3}}$.

5.1. Expert Costs and Social Welfare

We first examine the utility of searchers as a function of the search cost c_s. Figure 1 show's expected searcher utility in a few different cases: with a self-interested expert who charges a profit-maximizing amount, with an expert who provides her services for free, and with no expert present (the interval of c_s values used was divided into two parts, and the vertical interval was broken accordingly, as searchers' utilities substantially change with c_s). Unsurprisingly, searcher utility decreases as the cost of search increases (this effect is obviously true for all expert costs we examined). More surprisingly, once $c_s > 0.07$, searchers are better off in a market with no expert than even a market with an expert who provides services for free!

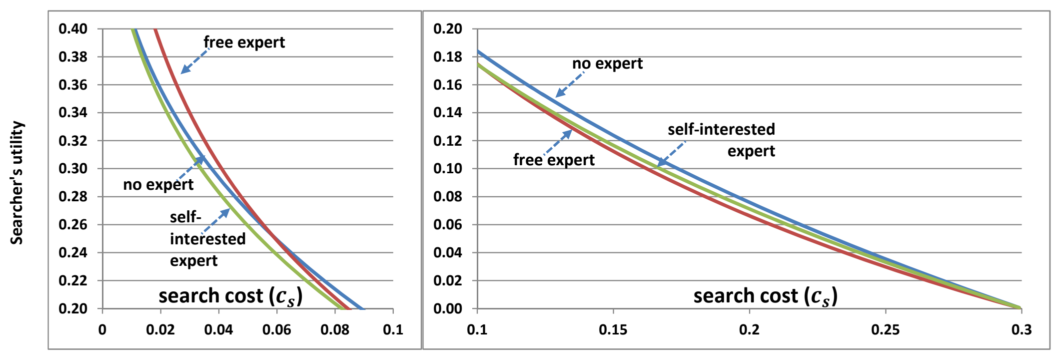

Fig. 1. Searcher's utility decreases as search cost increases. For $c_s > 0.07$, it is better for searchers to have no expert in the market than even an expert providing free services.

This implies that for $c_s > 0.07$, the presence of an expert in the market, regardless of c_e, decreases the searcher's utility. These results are robust, and appear even when agents use discounting. The results seem counterintuitive: why should the presence of an expert, in particular one who provides her services for free, lead to a decrease in utility for searchers? Indeed, such behavior would never occur in one-sided search, but it turns out that equilibrium behavior when many different agents are making decisions complicates the matter significantly. An agent's decision on whether or not to consult an expert could be significantly affected by what it expects others to do: in fact, the very fact that others can consult the expert makes it optimal for an agent to also consult the expert in many cases, *even though everyone would be better off if the expert were not present*. The source of the higher costs is that the presence of the expert induces everyone to stay in the market searching for longer, incurring higher search costs, without achieving a sufficient compensatory benefit in the value of the final match received. It is interesting to note that the effect holds to a greater extent when search costs are higher. This is in part because better information makes agents more picky, and everyone searches longer when the option of consulting the expert is available. Even if the expert is free or low-cost, searchers will end up wasting money due to the higher search costs.

Interestingly, the negative effects of the presence of the expert get worse as the sizes of the teams being formed increase. Figure 2 shows the ratio of the expected utility received by a searcher when there is a profit-maximizing monopolist expert present in the market versus when there is no expert present in the market. We can see that the ratio declines as the number of agents forming a team increases (the k sides in the search). For $k = 1$ the expert is beneficial, and then it becomes harmful for $k \geq 2$ by presenting the opportunity for agents to have to search longer, because it is an equilibrium for the other agents to consult the expert, and given that they are doing so, it makes sense for any individual agent to do so as well. As a result all agents pay more in search costs than they would if there were no expert present.

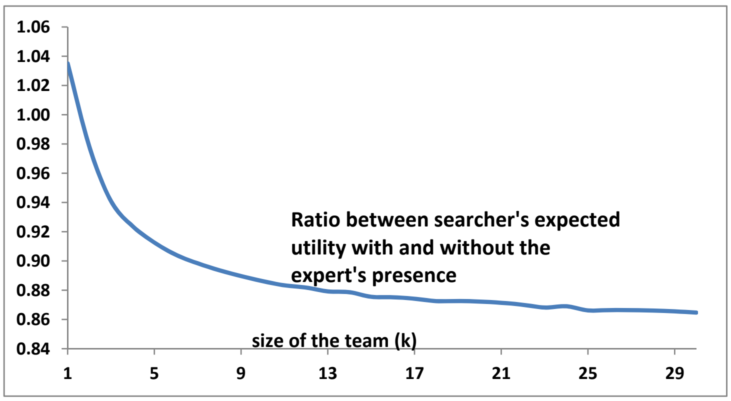

Fig. 2. Ratio of the expected utility of a searcher with a self-interested expert in the market versus with no expert in the market as a function of k, the number of agents that have to all agree to form a team. The presence of the expert becomes more damaging as the team size increases.

A market designer that seeks social welfare maximization can change the effective query cost paid by searchers by either subsidizing the expert to reduce her price, or by instituting a tax on transactions to increase the price paid by searchers. For the purposes of this paper, we assume that the expert must be compensated completely by the market platform for the loss she suffers from the change in query price. This is a weak assumption: in many cases the market designer may be able to impose even stronger change (and higher social utility) by using her special position. In the case of subsidization, we assume that the market designer must make a side payment to the expert. In the case of a tax, we assume that the market designer pays the entire amount of the tax plus some extra amount to the expert (the amount collected by taxation alone is necessarily less than the expert would have made from setting her optimal price).

In order to maximize expected profit, the expert computes her optimal cost c_e given that the individual agents are playing their optimal search strategies subject to c_s and c_e. For instance, for $c_s = 0.1$ in our example, the optimal expert query cost is $c_e = 0.0065$ (see Figure 3, where the lower curve, which demonstrates the expert's profit as a function of query cost, peaks at 0.0065; note, however, that social welfare is not maximized at $c_e = 0.0065$).

Figure 3 presents an example where taxation is helpful in two-sided search. In this case, social welfare (taking into account the tax) is maximized when the effective query price paid is 0.0237 (seen at the upper curve in the figure). Intuitively, one would expect that a reduction in expert query price should increase social welfare. However, in many sensible market settings the contrary holds.

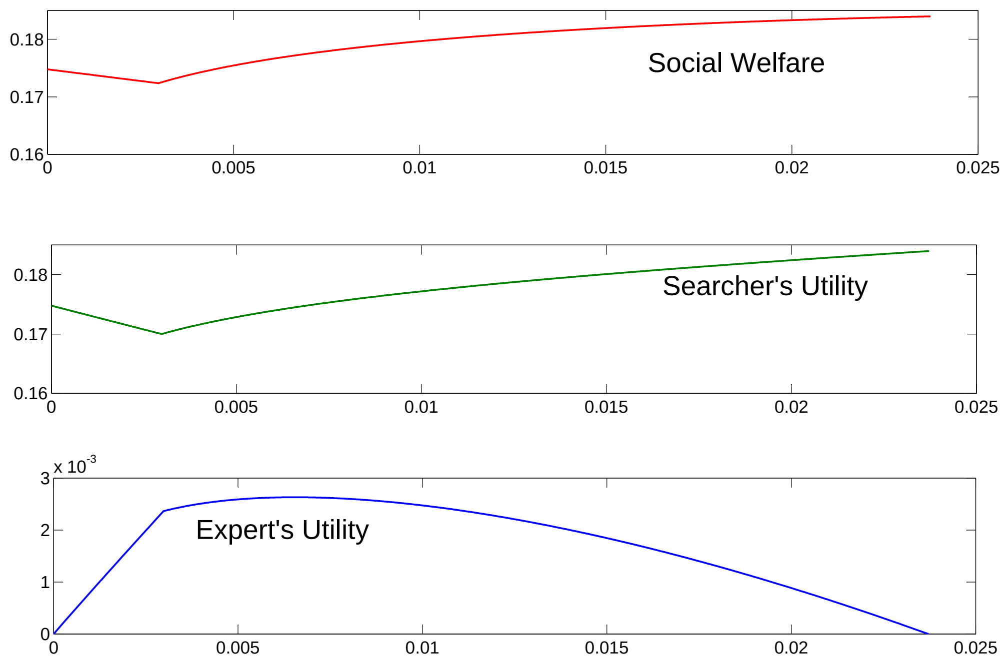

Fig. 3. Taxing transactions between searchers and experts: by increasing the effective query price from 0.0065 to 0.0237, a market designer can maximize social welfare. In this example $c_s = 0.1$.

The above result holds, once again, for the case where agents use discounting. Figure 4 illustrates such a scenario for the same setting, except that searchers use a discount factor of 0.9. As observed from the figure, market design in the form of taxing transactions between searchers and experts (thus increasing the effective query price) and compensating the expert accordingly, improves social welfare, and for a large portion of the c_s interval even improves the social welfare to a value greater than the resulting social welfare when the expert does not charge at all for her services.

5.2. Characteristics of Optimal Searcher Behavior

As discussed in Section 4, the searcher's strategy is characterized by two thresholds, t_l and t_u. We can study the effect of search cost on these thresholds. Figure 5 (left) shows results for $c_e = 0.01$ (variations in the results are minor across several expert query costs). One can observe that, as the cost of search increases, the thresholds get closer to one another, and eventually merge when $c_s > 0.22$. At search costs higher than that the searchers do not query the expert.

Avoiding the expert at high search costs is another seemingly counterintuitive result, as one may expect higher reliance on experts when search costs are high. However, in this case it is in keeping with results from one-sided search. To understand this observation, we revisit the definition of the searcher strategy and the meaning

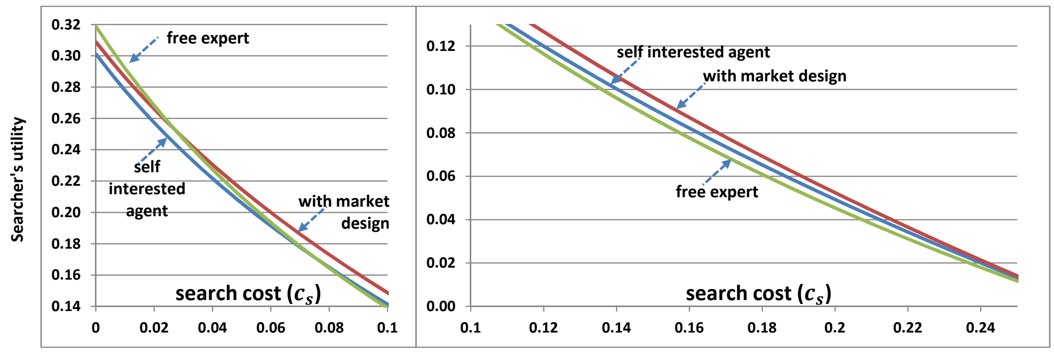

Fig. 4. Searcher's utility (Y axis) as a function of search cost, when searchers use a discount factor of 0.9. Similar to the case without discounting, taxing transactions between searchers and experts improves social welfare, even to a point where it is better than the social welfare when the expert does not charge for her services at all.

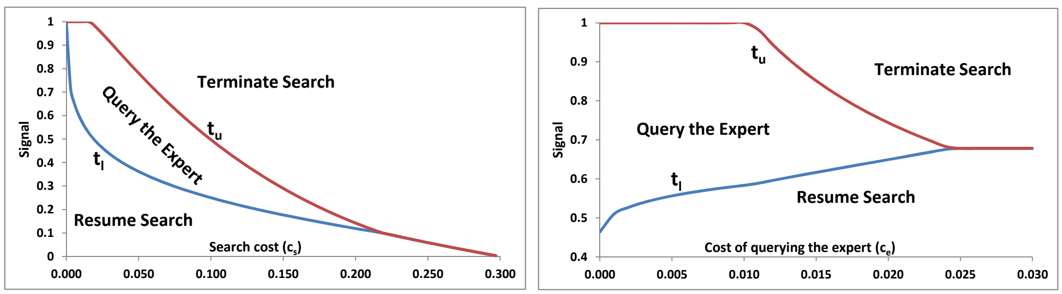

Fig. 5. Left: Thresholds decrease and get closer to one another as the cost of search increases. They eventually merge, indicating that searchers do not query the expert for $c_s > 0.22$. In this example $c_e = 0.01$. Right: Thresholds get closer to one another and eventually merge as expert cost increases. In this example $c_s = 0.01$.

of these thresholds. We specifically recall that above t_u, the searcher accepts the observed (noisy) value without further search. Note that threshold merge occurs only when the thresholds are both small. One can thus conclude that when search cost is high, searchers almost always choose to accept the current opportunity and avoid any additional expenses, be it further search or expert advice.

We can also examine the effect of expert query cost on the thresholds. A typical result at $c_s = 0.01$ is presented in Figure 5 (right). Similar to the threshold dependency on search cost, we observe a merge of the thresholds at some point (specifically, at $c_e = 0.025$). However, an interesting difference is that t_l does not decrease, and at the merge point $t_l = t_u = 0.68$. Here, the expert is not queried because it is simply too expensive.

6. CONCLUSIONS

This paper is the first to look at the impact of experts on two-sided search markets. We generalize some results from the consideration of experts in one-sided search, notably that of the optimal strategy. As demonstrated in the analysis section, the analysis can be extended in a quite straightforward manner to various cases (e.g., the formation of k member teams, discounting of gains and different populations).

Phenomenologically, we find ways in which the behavior of two-sided search markets is drastically different from the more intuitive behavior seen in one-sided search markets with experts. Perhaps most surprisingly, there are perfectly reasonable market settings (typically those with high search costs) where information can be *bad*. The presence of experts actually leads to socially suboptimal outcomes compared to cases in which they are absent. With experts, the equilibrium behavior implies that agents should query the expert because other agents may be doing so, even though they would all be better off if they agreed in advance not to consult the expert (or if the expert was not present). This effect gets worse if we consider many-sided search: in fact, as the number of agents required to form a team increases, the presence of the expert makes individual searchers relatively worse off compared with what they could have expected with no expert present.

We also study the problem faced by the designer of a market platform that brings searchers and experts together and seeks to mitigate this effect. We propose that the market designer could institute a Pigovian tax on transactions between searchers and experts in order to make searchers query the expert less often, and show how to compute the optimal tax.

The model of two-sided search used in this paper assumes agents are ex ante identical, in that there are no individuals who are "naturally" better than others or more easy to please than others. However, when a potential match is formed, each agent gets some idiosyncratic utility from the particular qualities of that match. The two agents will in general have different values for a match. While this is useful, it would be illuminating to understand what happens when the quality of matches are quite differently determined. In future work, we will consider two other common models: one where two agents get the same idiosyncratic utility for a match (a function of their compatibility, but the same for both), and another where what one agent gets from a match is only a function of the other agents' "quality."

REFERENCES

ANSHELEVICH, E., DAS, S., AND NAAMAD, Y. 2011. Anarchy, stability, and utopia: Creating better matchings. *JAAMAS*. Forthcoming.

BAKOS, Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science 43,* 12, 1676–1692.

BLOCH, F. AND RYDER, H. 2000. Two-sided search, marriages, and matchmakers. *International Economic Review 41,* 1, 93–115.

BURDETT, K. AND WRIGHT, R. 1998. Two-sided search with nontransferable utility. *Review of Economic Dynamics 1*, 220–245.

CHADE, H. AND VENTURA, G. 2005. On income taxation and marital decisions. *Review of Economic Dynamics 8,* 3, 565–599.

CHHABRA, M., DAS, S., AND SARNE, D. 2011. Expert mediated search. In *Proceedings of AAMAS*. 415–422.

DAS, S. AND KAMENICA, E. 2005. Two-sided bandits and the dating market. In *Proceedings of IJCAI*. 947–952.

DIAMOND, P. 1982. Aggregate demand management in search equilibrium. *The Journal of Political Economy 90,* 5, 881–894.

GALE, D. AND SHAPLEY, L. 1962. College admissions and the stability of marriage. *American Math. Monthly 69*, 9–15.

GROSFELD-NIR, A., SARNE, D., AND SPIEGLIER, I. 2009. Modeling the search for the least costly opportunity. *European Journal of Operational Research 197,* 2, 667–674.

JOVANOVIC, B. 1979. Job Matching and the Theory of Turnover. *The Journal of Political Economy 87,* 5, 972.

KEPHART, J. AND GREENWALD, A. 2002. Shopbot economics. *JAAMAS 5(3)*, 255–287.

KEPHART, J. O., HANSON, J. E., AND GREENWALD, A. R. 2000. Dynamic pricing by software agents. *Computer Networks 32*, 731–752.

KETCHPEL, S. 1994. Forming coalitions in the face of uncertain rewards. In *In Proc. of AAAI94*. 414–419.

KRAUS, S., SHEHORY, O., AND TAASE, G. 2003. Coalition formation with uncertain heterogeneous information. In *Proceedings of AAMAS '03*. 1–8.

LEE, R. AND SCHWARZ, M. 2009. Interviewing in two-sided matching markets. NBER Working Paper.

LIPPMAN, S. AND MCCALL, J. 1976. The economics of job search: A survey. *Economic Inquiry 14,* 2, 155–189.

MCCALL, J. 1970. Economics of information and job search. *The Quarterly Journal of Economics 84,* 1, 113–126.

MCMILLAN, J. AND ROTHSCHILD, M. 1994. Search. In *Handbook of Game Theory with Economic Applications*, R. Aumann and S. Hart, Eds. North-Holland, 905–927.

MILGROM, P. 1981. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics 12*, 380–391.

MORTENSEN, D. AND PISSARIDES, C. 1999. New developments in models of search in the labor market. *Handbook of labor economics 3*, 2567–2627.

ROTHSCHILD, M. 1974. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy 82,* 4, 689–711.

SAHIB, P. AND GU, X. 2002. To Tie the Knot or Not: Cohabitation, Marriage and Individuals' Attitudes to Risk. *Demographic Research 6,* 13, 355–382.

SARNE, D. AND KRAUS, S. 2008. Managing parallel inquiries in agents' two-sided search. *Artif. Intell. 172,* 4-5, 541–569.

SHEHORY, O. AND KRAUS, S. 1998. Methods for task allocation via agent coalition formation. *Artificial Intelligence 101,* 1, 165–200.

SMITH, L. 2011. Frictional matching models. *Annual Reviews in Economics 3*, 319–338.

WIEGMANN, D., WEINERSMITH, K., AND SEUBERT, S. 2010. Multi-attribute mate choice decisions and uncertainty in the decision process: A generalized sequential search strategy. *Journal of Mathematical Biology 60,* 4, 543–572.

WRIGHT, R. 1986. Job search and cyclical unemployment. *The Journal of Political Economy 94,* 1, 38–55.