

Two-Sided Search With Experts

Yinon Nahum

Dept. of Computer Science,
Weizmann Institute of Science
Rehovot 76100, Israel

David Sarne

Computer Science Department,
Bar-Ilan University
Ramat-Gan 52900 Israel

Sanmay Das

Dept. of Computer Science,
Virginia Polytechnic Institute
Blacksburg, VA 24060

Onn Shehory

IBM Haifa Research Lab
Mount Carmel
Haifa 31905 Israel

March 24, 2014

Abstract

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 either a fee or a commission, 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. Interestingly, we find that better information can hurt: the presence of the expert, even if the use of her 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 compensating the expert to change the price at which she offers her services.

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 [19, 25], marriage and mating [5], and partnership

formation among artificial agents [38]. 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 which informs her of the value *to her* of that match [6, 7]. 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 others' quality, whether through a one-shot interview process [24] or through repeated interactions like dating [10].

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 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). We start with a simple, but perhaps counter-intuitive, observation about the value of information. In one-sided search, better information can never hurt: a searcher receiving more accurate signals of the values of opportunities is always better off. However, even in a simple model of two-sided search [6, 7], we show that it is possible for all participants in the market to be worse off (in expectation) when signals about the values of possible matches are more precise. This typically happens when search costs are high, in which case more precise information can lead to socially inefficient levels of search which are nevertheless, in equilibrium, individually rational for each participant.

We go beyond this observation to consider more completely the impact of self-interested information brokers who can disambiguate noisy signals upon payment of a fee. We extend the model of Burdett and Wright [6] to allow for noisy signals and an expert (*a la* Chhabra et al.[9]) who will, upon payment of either a fee or a commission, provide the "true" value of the potential match to a searcher. We assume that the expert is honest for reputational or regulational reasons (cf. realtors or mortgage brokers). One-sided search with experts can be modeled as a Stackelberg game, where the expert moves first by setting a 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. We solve for equilibrium for two different scenarios: (1) when experts charge a fixed *a la carte* fee every time they perform the service of disambiguating a signal; (2) when experts operate on a commission basis, charging a percentage of the value received if and only if a successful match is made.

We characterize equilibrium strategies for searchers and experts in these models, and also investigate some phenomenological properties of equilibria for some instantiated distributions of true values and signals which satisfy the properties of our model. Some of these properties, mostly related to the (reservation value) form of the equilibrium strategy, carry over from the model of one-sided search with experts of Chhabra et al. [9]. Other properties are starkly different.

The basis for this difference can be intuitively understood from the example mentioned above, of

how more precise signals can decrease expected utility. The need for agents to reason about the optimal behavior of others changes some of the major systemic properties of two-sided search markets compared to one-sided search markets. These effects can be even more prominent in the presence of a self-interested expert: 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 if the expert 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 to the agents who are searching. The expert thus benefits at the expense of the searching agents. Therefore, 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 subsidizing the expert to *increase* the fees she charges searchers, so that they consult her less frequently. We also investigate the differences between the *a la carte* pricing model and the commission-based model from the perspective of the expert, and find that, in the same setting as above, commission-based pricing may be preferable for experts when search costs are low, while fee-based pricing yields higher profits when search costs are high. The intuition for this extends that above: when search costs are lower, searchers typically search longer and get matches of higher value. In the commission based model, the expert can take advantage of this by charging a percentage of this value instead of a fixed price.

1.1 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 [20, 21, 38, 39], 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 [4, 20].

The underlying foundation for costly search analysis is the theory of sequential search [27, 11, 30]. Utility-maximizing individuals are considered to be sequentially reviewing different opportunities where search incurs a cost [28, 17, 35, and references therein]. Two-sided search models were developed in an effort to understand the effect of simultaneous search on both sides of the market in costly environments [7, 38, 42]. Unlike stable matching scenarios [15, 2] 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 extends the work of Burdett and Wright [6] and of Chade and Ventura [7], 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 [10, 37, 24]). Mediators in two-sided search models have typically taken the form of matchmakers rather than information brokers [5]. Instead we follow the work of Chhabra et al. [9] in focusing on how the presence of self-interested knowledge brokers affects the market.

It has been observed often in the literature that a parameter shift or policy intervention may have negative payoff consequences in a search equilibrium even if it benefits individual searchers when holding others' behavior fixed. For example, Duffie et al. [12] find that providing additional (public) information to individual searchers may lower equilibrium payoffs. Their model differs significantly from ours, as they do not consider two-sided search for the purpose of matching, but rather for information sharing in a large population; agents can control the intensity of meeting other agents, and hence the amount of information shared, though the process of meeting others is costly. Albrecht et al. [1] consider the problem of sequential search when the decision to stop is made by a committee, showing that committee members are less picky and more conservative than the single agent. Here again, there are no two-sided search considerations. Instead, committee members that see a different value in each opportunity need to vote on whether to stop or resume search. Shimer and Smith [41] make use of Pigouvian taxes in a two-sided model where agents are characterized by types that determine the utilities of others from matching with them. They use a tax on search activity for the least productive agents, and a subsidy for the most productive ones. This has two effects: Agents who receive a subsidy increase their search intensity; and they are more reluctant to accept matches, since search is less costly. A search tax creates the opposite incentive, possibly leading some low-productivity agents to forgo search altogether. In addition to the specific differences in models, we note that none of these papers consider the effects introduced by the existence of an external, self-interested information provider.

Some work has considered the role of the platform in two-sided markets. Most of this work focuses on the payments made by either of the sides, justifying different fee-charging schemes (e.g., charging only one side in two-sided markets while the other group is allowed onto the platform for free, as in the case of yellow pages directories [3]), or paying one of the sides to use the platform (e.g., in the case of credit cards, where rewards programs such as contributions to frequent flyer plans are offered [36]). While these works seem to resemble ours in the context of mediated two-sided search, the expert in our model is very different from the platform analyzed in these other works. The expert, in our model, does not have to be used and its sole purpose is to offer a more accurate signal (rather than supplying the matching technology itself).

Partnership formation has also been studied in the context of MAS coalition formation (e.g., by Shehory and Kraus [40]). The work in this paper differs from classical coalition formation in several ways. First, in this work partnership or team formation is gradual and iterative: at each stage only a single partnership is formed (resembling some ideas in early work by Ketchpel [22]); second, values of teams are not known in advance and their revelation is costly (somewhat like elements from the work of Kraus et al. [23]); 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.

Recent literature presents many advances in applying search theory in order to investigate search dynamics in markets, particularly in cases where comparison-shopping principles are applied [43, 18]. Most of these works assume that the interests of the information provider (e.g., comparison shopping agent) and user are identical (the provider's sole purpose is to serve the user's needs) [32]. Naturally, in such cases, the existence of an information provider favors the user (e.g., buyers, in comparison shopping applications) [16, 31]. Those few works that do assume that the information providers are self-interested autonomous entities [20, 21] focus on the use of the information provider for obtaining the signal itself in settings where signals are noiseless (e.g., price quotes) rather than for supplying complementary in-

formation [44]. Another significant difference between that stream of literature and our work is that we are examining two-sided search with nontransferable utility, whereas the papers mentioned above focus on two-sided markets where buyers and sellers interact through a pricing mechanism.

In spirit, our work follows previous results in other settings, in which it has been shown that so-called “market inefficiencies” can increase market performance, under certain circumstances. For example, Masters [26] shows that an increase in the minimum wage, which is often considered by economics as inefficiency, can have positive employment effects. In transportation economics (e.g. congestion games), equilibrium is frequently not the social optimum. In such cases, it has been shown that taxation can change the equilibrium to a more desirable one [34, 33, 14]. Similarly, taxes can facilitate more desirable equilibria in Boolean games [13]. Here we show that a similar phenomenon also occurs in the context of search (and market inefficiency is in the sense of noisy observations), though the model and analysis are, of course, methodologically completely different from the above.

1.2 Roadmap

In the next section, we present the basic model of two-sided search within which we work. In Section 3, without going into solution details, we show an example of a realization of the search model where the entire population is better off in expectation with less precise signals, and discuss the strategies and tradeoffs in this motivating example. After that, in Section 4 we provide details on how to solve for the equilibrium of the search in the presence of an expert who operates on a fee-for-service basis, and show how to extend our analysis to heterogeneous populations (where the two sides of the market are different) as well as to multi-sided search (for example, for team formation rather than pair formation). Section 5 illustrates some of the major phenomenological implications of the model for a particular instantiation of the distribution of signals and the distribution of true values conditional on signals, and also considers a different, second distribution of these as a robustness check. Section 6 analyzes the variant of the model where the expert operates on a commission basis, and Section 7 concludes.

2 Model

The model is based on a standard two-sided distributed search model [6, 7, 38], augmented to include uncertain signals.¹ The model assumes fully rational self-interested agents, searching for appropriate partners to form mutually acceptable pair-wise partnerships.²

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 easier to please than others. The search proceeds in discrete time periods. At period t , the matching technology arranges a meeting between two agents, say i and j , each of whom pays a search cost c_s and receives a different, i.i.d. noisy signal, denoted s_{ijt} , which indicates the estimated value of the match to it. s_{ijt} is a noisy signal of the true value of the match v_{ijt} to i (also i.i.d.), and, in general, $s_{ijt} \neq s_{jit}$. We assume that agents are acquainted with the distribution of signals $f_s(s)$ and the conditional probability density of values given signals, $f_v(v|s)$

¹See Appendix B for a summary of all the notations used in this paper.

²For 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.

(and their corresponding cumulative distribution function (cdf) $F_s(x)$ and $F_v(y|s)$), which are identical for all agents.³ Notationally, throughout this paper we suppress dependence on i, j, t where it is clear from context.

These modeling assumptions are equivalent to each agent receiving 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, it becomes increasingly unlikely that any pair will be evaluated a second time, since potential partnerships are drawn at random from the population; however, even with a relatively small number of agents, we still get a valid model for cases where the utility of a partnership is dependent on the circumstances in which it is formed.

Upon receiving signal s_{ijt} , agent i can either accept the partnership (receiving utility v_{ijt} , decline it (continuing search), or pay a cost c_e to consult an expert who then reveals the (noiseless) true value of the partnership v_{ijt} to agent i . The expert herself incurs a cost of d_e in order to report v_{ijt} to the agent. If the agent does consult the expert, it must decide whether to accept or decline the partnership once it receives the true value (which the expert reports honestly, as described below). If both agents decide to accept the partnership, a match takes place and the agents leave the market and are replaced with identical agents (thus the market is time invariant). 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 with a truthfulness constraint. For the purposes of this paper, we assume that she has the ability to determine the true value and must report it truthfully. Especially given that the expert is a monopolist, this constraint is reasonable – the expert may be regulated or have to consider her reputation, and would not want to lose out on future revenue associated with her privileged position by lying. Keeping in mind this constraint, the expert’s goal is to maximize her own expected utility; if she is queried n times, this is $n(c_e - d_e)$.

3 A Motivating Example

Within the model described above, consider a specific distribution of signals and true values (we return to these distributions, and show how we solve for the example below in detail, in Section 5). First consider the case of noisy signals. Let signals s be distributed uniformly on $[0, 1]$, and let the conditional density of true values be a particular monotonically increasing function in the interval $[0, s]$: $f_v(y|s) = \frac{3\sqrt{y}}{2\sqrt{s^3}}$ (see Figure 1 for an example when $s = 0.8$). Thus, the signal is optimistic, serving as an upper bound on the true value (e.g., people tend to get a good first impression of others).

In the case of perfect signals, instead of seeing the signal s defined above, the searcher gets to actually observe v , but the underlying distribution generating v is exactly the same: it first draws s uniformly from

³Alternatively, 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’ rule.

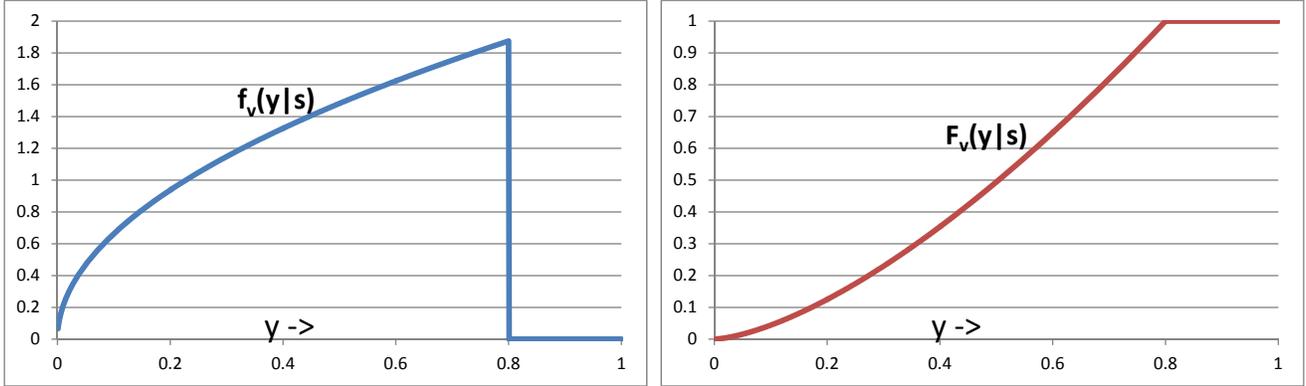


Figure 1: The function $f_v(y|s)$ and its cumulative distribution function, $F_v(y|s)$, when $s = 0.8$ (for $f_v(y|s) = \frac{3\sqrt{y}}{2\sqrt{s^3}}$).

$[0, 1]$ and then v from the density $f_v(y|s)$ defined above (which can also be seen as the case where the true value can be obtained from the expert for free). It seems intuitive that everyone should do better with no noise, when the values of opportunities are perfectly revealed. However, this is not the case. For the moment, consider two different search costs:

| c_s | With Noise | | | | | Without Noise | | | |
|-------|------------|--------|------------|------|----------|---------------|------------|------|----------|
| | Util | Thresh | # Searches | Cost | MatchVal | Util | # Searches | Cost | MatchVal |
| 0.01 | 0.41 | 0.68 | 9.65 | 0.10 | 0.50 | 0.46 | 17.33 | 0.17 | 0.64 |
| 0.1 | 0.18 | 0.31 | 2.08 | 0.21 | 0.39 | 0.17 | 2.59 | 0.26 | 0.43 |

In the table above, we show the agents' search characteristics in equilibrium and the resulting individual expected utility. The equilibrium search pattern is, as proved in the next section, reservation-value based — the agents set a threshold above which they accept other agents and otherwise resume their search. The table specifies the expected utilities (“Util”), reservation values (“Thresh”), the expected number of searches (i.e., the number of opportunities evaluated before an agent accepts a match), the expected cost of these searches, and the expected value obtained from the partnership the agent eventually ends up with (in the case without noise, the expected utility is the same as the threshold, since an agent will only accept an offer if it offers more utility than resuming search). c_s is the cost paid by the searcher to sample another opportunity. If we focus on the expected utility, we see that, while searchers do better with perfect information when $c_s = 0.01$, they are actually better off with noisy signals when $c_s = 0.1$. The reason for this is that they tend to search more when they have more information. This typically leads to them finding a better match in the end, but also to paying higher search costs along the way. If it were just a single agent searching more due to the better information, then the overall utility would have to increase. However, since all agents become more “picky,” the search becomes even longer, and the additional cost is no longer outweighed by the benefit of getting a better match in the end.

The importance of this observation is that it depends on the behavior of others. Without noise, everyone becomes too picky. Given that others are being more picky, it does not benefit you to be less picky. Further, although it is true that if everyone were less picky then everyone would be better off,

any individual would still be better off being more picky, so it is not an equilibrium for everyone to be less picky. To illustrate further, consider what happens when the search cost is 0.1 and there is no noise. Suppose everyone were to agree to use threshold 0. Then no one would ever decline an offer, and the expected utility for everyone would just be the expectation of the value distribution (0.3), minus the cost of seeing a single opportunity, which is 0.1, yielding 0.2, which is higher than the equilibrium outcome values both with and without noise (see Section 5 for details on the distribution and the expected value). However, for any individual, if everyone else were using a threshold of 0, it would make sense to have a higher threshold (and, in fact, the threshold would be higher than 0.2, and would be exactly equal to that agent’s expected utility).

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 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 as well as for cases where agents discount future utility and where agents come from different populations, in a straightforward way, as shown towards the end of this section. For exposition purposes, this section includes only sketch of proofs. The detailed proofs are given in Appendix A.

4.1 Preliminaries

As described in Section 2, 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 [45] and perfect information models have been considered in the two-sided search literature [7, 38]. 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. [9], and for two-sided search with perfect signals we follow Burdett and Wright [6].

One-sided search The optimal strategy in many models of search is a *reservation value* (i.e., threshold-based) strategy, where the searcher accepts any opportunity that is higher than a particular reservation value.⁴ 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 [46, 29], 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:

⁴Meaning that the searcher keeps on reviewing new opportunities as long as the highest-valued opportunity encountered so far is below the reservation value.

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

Corollary 1 [9] 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 [38, 6]. 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 $f_v(v)$ (and corresponding cumulative distribution function (cdf) $F_v(v)$), then they all use the same reservation value in equilibrium. The equilibrium reservation value is the value which 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 which maximizes utility for any individual agent, x^* , satisfies [6]:

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

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 [28]. Consequently, a searcher who 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 $V(x^*)^5$, satisfies $V(x^*) = x^*$,

⁵ In general, in two-sided search, the value function V will also be dependent on the strategies being used by others. Where it is clear from context we suppress such dependence.

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 proven 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:

$$c_s = (1 - F_v(x^*)) \int_{y=x^*}^{\infty} (y - x^*) f_v(y) dy \quad (2)$$

The expected number of search iterations is simply the inverse of the success probability, $1 / (1 - F_v(x^*))^2$, 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):

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

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 ($-c_s$). The second element, $(1 - \Pr(s \in S) \Pr(s^* \in S^*)) V(S, S^*)$, 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 optimality principle and the stationarity of the problem, $V(S, S^*)$. The third element, $\Pr(s^* \in S^*) \Pr(s \in S) E[v|s \in S]$, 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 1 *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:

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

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 when 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. \square

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 $1 / (1 - F_v(t^*))^2$, 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 computed by the function $V()$, and resume otherwise. At optimality, following the Bellman optimality principle, the function $V()$ gives the expected utility from resuming the search and is given by the following modification of Equation 3, given that the other agents use strategy $(S'_{\text{others}}, S''_{\text{others}}, V_{\text{others}}())$:

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

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:

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

The value of $V(S', S'')$ in Equation 5 is 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 $V(S', S'')$. 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 2 For $f_v(y|s)$ satisfying the HSGN assumption (Definition 1), given the search cost c_s and the fee c_e set by the expert, the Bayes Nash Equilibrium for all 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:

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

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

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

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

$$B = A \quad (10)$$

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

Sketch of 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 $s'' > s$. Equations 6, 9, and 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 t_l and t_u which maximize the expected benefit. \square

This results in a set of six equations (6 - 11) with six variables. We can solve these simultaneously to calculate V, t_l , and 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 the searcher queries the expert it receives either the value obtained from the expert if it is above V (with probability B , since this is the probability that the other side accepts the match), or otherwise V :

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

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:

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

which transforms into Equation 8.

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.

Based on the proof for Theorem 2, it is notable that in order to guarantee the double-threshold strategy structure it is enough to require that if $s_2 > s_1$ then the following three conditions hold (of which the

HSGN requirement is a specific case):

$$\begin{aligned}
E[v|s_2] &> E[v|s_1] \\
\int_{y=V}^{\infty} [F_v(y|s_1) - F_v(y|s_2)] dy &> 0 \\
\int_{y=-\infty}^V [F_v(y|s_1) - F_v(y|s_2)] dy &> 0
\end{aligned} \tag{12}$$

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; and (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 $B = A^{k-1}$.⁶

Further, even if the size of the coalition encountered at each stage of the search varies (e.g., entrepreneurs 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 [9], 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 future costs paid). In keeping with the sequential search literature [28], 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:

$$V = \frac{-c_s - c_e(F_s(t_u) - F_s(t_l)) + B \cdot C}{1 - \delta(1 - A \cdot B)} \tag{13}$$

$$C = \delta \left(\int_{s=t_u}^{\infty} f_s(s) E[v|s] ds + \int_{s=t_l}^{t_u} \left(f_s(s) \int_{y=V}^{\infty} y f_v(y|s) dy \right) ds \right) \tag{14}$$

⁶We note that there are situations where some subset of the k might form a team, but we study the case where exactly k agents are needed (e.g., in the case where each agent has a certain capability necessary for a task).

Equating the first derivative of (13), according to t_l and t_u (separately) to zero, obtains the following modifications of Equations 7-8

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

$$c_e = \delta B \int_{y=-\infty}^V (V - y) f_v(y|t_u) dy \quad (16)$$

The intuitive interpretation of Equations 15-16 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) [6], differing in the distribution of values and the costs of search and costs of querying the expert. 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:

$$V^i = \frac{-c_s^i - c_e^i (F_s^i(t_u^i) - F_s^i(t_l^i)) + B^i \cdot C^i}{A^i \cdot B^i} \quad (17)$$

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

$$c_e^i = B^i \int_{y=-\infty}^{V^i} (V^i - y) f_v^i(y|t_u^i) dy \quad (19)$$

$$A^i = 1 - F_s^i(t_l^i) - \int_{s=t_l^i}^{t_u^i} f_s^i(s) F_v^i(V^i|s) ds \quad (20)$$

$$B^i = A^{3-i} \quad \{\text{note that } 3 - i \text{ here is an index rather than power}\} \quad (21)$$

$$C^i = \int_{s=t_u^i}^{\infty} f_s^i(s) E^i(v|s) ds + \int_{s=t_l^i}^{t_u^i} \left(f_s^i(s) \int_{y=V^i}^{\infty} y f_v^i(y|s) dy \right) ds \quad (22)$$

for $i = 1, 2$ i.e., a total of 12 equations.

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

$$\eta_{c_e} = \frac{F_s(t_u) - F_s(t_l)}{A \cdot B} \quad (23)$$

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 using numeric methods in order 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 which allows searchers and experts to interact – for example, an online dating website, or one which 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. [9] 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 (the possibility for which is shown in Section 3, and discussed further in Section 5). In such cases the market designer may want to subsidize the monopolist expert to *increase* her price.

We will analyze such subsidies in the following framework. 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 , resulting in expected profit π'_e to the expert. Presumably the market designer is doing this for the overall benefit of searchers, but she may have to compensate the expert. In an extreme case where the designer has no explicit power over or relationship with the expert, she can still offer a payment β to the expert, which fully compensates the expert for the decreased revenue $(\pi_e - \pi'_e)$, leaving her total profit unchanged. 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):

$$W = V_{c_e^*} + \pi_e \quad (24)$$

When the market designer subsidizes expert queries, the social welfare must also take into account the subsidy. If 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 $W' = V_{c_e'} + \pi_e - \beta = V_{c_e'} + \pi_e'$ (note that β cancels because it is included as a positive term in π_e).

We note that an alternative market design “lever” in this case would be in a situation where the designer has regulatory authority and can impose a tax. In this case the market designer would want to impose a per-unit tax t , levied on the expert, so that

$$\operatorname{argmax}_{c_e^*} (c_e^* - d_e - t)\eta_{c_e^*} = c_e'$$

where c_e' is the price that maximizes social welfare.

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 pairwise partnerships (or k -wise teams, in some settings described in detail where they occur). For our main experimental analysis, we work with a distribution where the signal is an upper bound on the true value (e.g., people tend to get a good first impression of others).⁷ 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 $[0, s]$: $f_v(y|s) = \frac{3\sqrt{y}}{2\sqrt{s^3}}$. The marginal density $f_v(y)$ is given by:

$$f_v(y) = \frac{dF_v(y)}{dy} = \frac{d\left(\int_{s=-\infty}^{\infty} f_s(s)F_v(y|s)ds\right)}{dy} \quad (25)$$

where:

$$F_v(y|s) = \int_{z=-\infty}^y f_v(z|s)dz = \begin{cases} \frac{\sqrt{y^3}}{\sqrt{s^3}} & 0 < y \leq s \\ 0 & y \leq 0 \\ 1 & y > s \end{cases} \quad (26)$$

Substituting from Equation 26 into Equation 25 yields: $f_v(y) = 3 - 3\sqrt{y}$.

⁷It can be assumed that results in this section use this distribution unless otherwise noted.

As part of a robustness analysis, demonstrating that our results generalize, we also consider a second distribution for parts of our analysis, where the generative model is very different: true values (y) come from a Gaussian distribution with mean μ and standard deviation σ_v , and signals for an opportunity are drawn from a Gaussian distribution with mean equal to the true value v and standard deviation σ_s . This reflects another very common noise model [8], and the true value is no longer constrained to be less than the signal. In this case, we can derive the marginal distribution of signals $f_s(s)$ as Gaussian with mean μ and standard deviation $\sqrt{\sigma_v^2 + \sigma_s^2}$ while the distribution of true values conditional on signals $f_v(y|s)$ is Gaussian with mean $(\sigma_s^2\mu + \sigma_v^2s)/(\sigma_v^2 + \sigma_s^2)$ and standard deviation $(\sigma_v\sigma_s)/\sqrt{\sigma_v^2 + \sigma_s^2}$. In numerical experiments, we use $\mu = 50, \sigma_v = 1, \sigma_s = 1$. We only show a few, selected results for this distribution, to demonstrate that our results are robust with respect to the underlying distributions of values and noise; qualitatively, we also find similar results in the cases that we do not show results for both. For the remainder of this section, we assume the expert’s expense d_e to be zero.

5.1 Noisy Signals vs. True Values

First, we demonstrate that having more information is not necessarily beneficial in two-sided markets. Figure 2 depicts the equilibrium expected utility of the agents with perfect signals and with noisy signals as a function of the search cost c_s . The graph on the left shows the searcher’s expected utility in the two settings, i.e., the value of $V(x^*)$ (marked “perfect signals”) and $V(t^*) = E[v|t^*]$ (marked “noisy signals”) for different c_s values (the values of x^* and t^* were calculated using Equations 2 and 4, respectively). The graph on the right shows the difference $V(x^*) - V(t^*)$. Figure 2 demonstrates that agents are better off with perfect signals for some c_s values (for $c_s \leq 0.06$) and with imperfect signals for others (for $c_s > 0.06$).

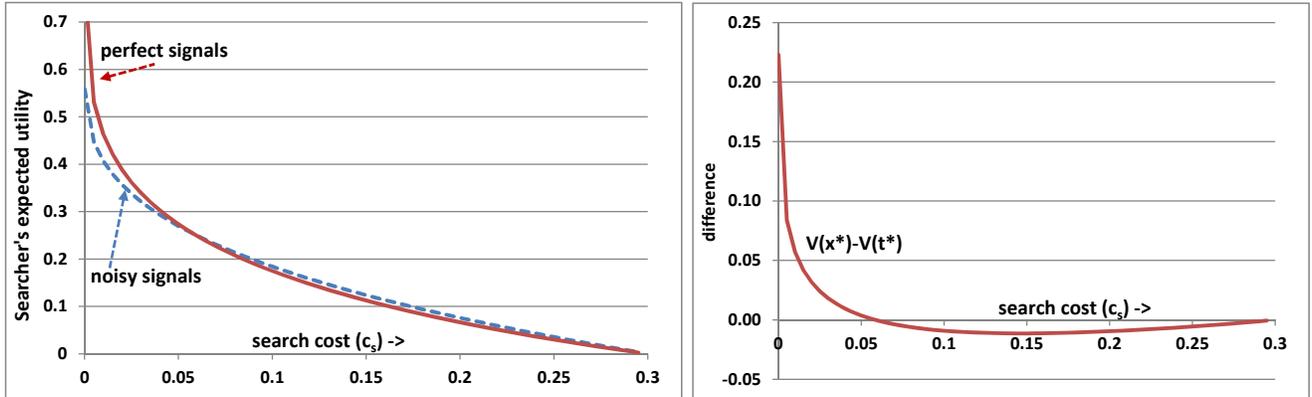


Figure 2: Searcher’s utility with perfect and noisy signals. As search cost increases, expected utility decreases in both model variants. For $c_s \leq 0.06$, agents are better off with perfect signals whereas for $c_s > 0.06$ agents actually benefit (in equilibrium) when signals are noisy.

5.2 Expert Costs and Social Welfare

We first examine the utility of searchers as a function of the search cost c_s . Figure 3 (left) shows the 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 right-hand graph of the figure depicts the difference in the searcher’s utility between any two settings. For any pair of settings, whenever the difference is above zero, having an expert of the first type is better than having an expert of the second type (and vice-versa). Unsurprisingly, the searcher’s utility decreases as the cost of search increases (this effect is obviously true for all expert costs we examined). More surprisingly, once $c_s > 0.06$, searchers are better off in a market with no expert than even a market with an expert who provides services for free!

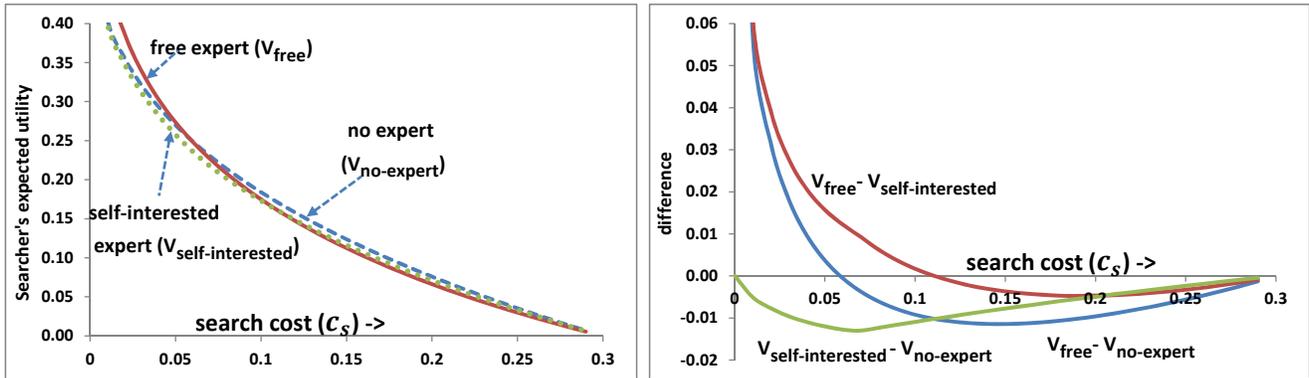


Figure 3: (Left): Expected utility for a searcher decreases as search cost increases. (Right): The differences in expected utilities for a searcher between several different scenarios: (a) where the expert is free ($c_e = 0$); (b) with no expert present; and (c) where the expert is a utility maximizing monopolist.

For these settings, our results imply that (1) for $c_s > 0.06$, the presence of an expert in the market decreases the searcher’s utility, and, (2) for $c_s > 0.11$ it is better for searchers to have a self-interested expert who charges her utility-maximizing fee than an expert that offers her services for free. Qualitatively, these results are robust, and appear even when agents use discounting, and for the case of normally distributed values and signals (see Figure 4). The results seem counter-intuitive: 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. This effect initially becomes more pronounced as the search cost increases, since the expected future search cost factors into the decision, with the free expert hurting all the agents beyond some point. The difference in expected utility between the “no expert” case and the “free expert” case peaks and then declines as search cost increases further, because at some point the agents begin to avoid searching more and more. The intuition is the same as that developed in Section 3. In this case the presence of a self-interested expert can make searchers even worse off, because they may be willing to pay to receive the extra information given that others are also doing so, whereas in Section 3 the extra information is “free” and still hurts the searchers.

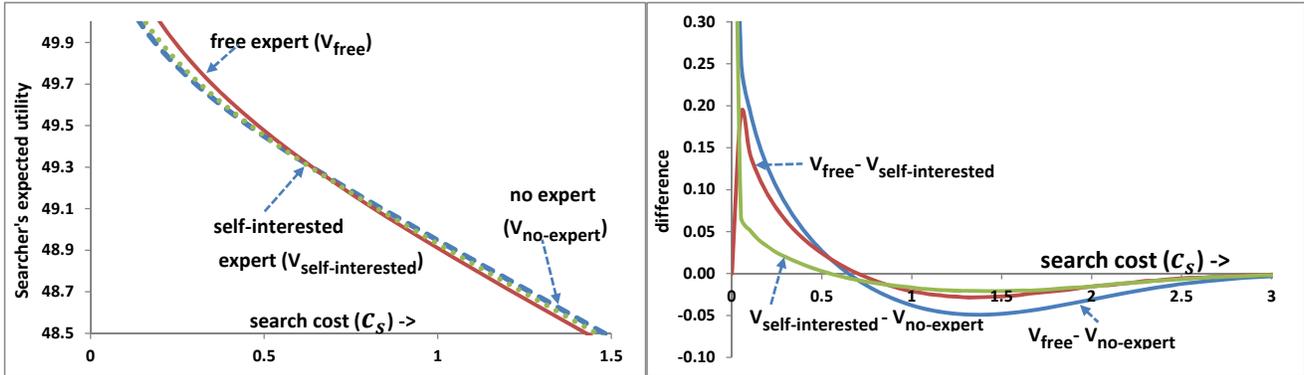


Figure 4: Expected utility for a searcher as a function of search cost (left) and differences in expected utility for different scenarios (right) when values and signals are normally distributed. The results are qualitatively similar to those in Figure 3.

Interestingly, the negative effects of the presence of the expert get worse as the sizes of the teams being formed increase. Figure 5 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. The search cost is taken to be $c_s = 0.02$. 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, but for $k \geq 2$ the presence of the expert becomes harmful because the equilibrium outcome is that agents search longer, paying more in search costs than the extra benefit they gain from receiving a higher value at the end of the search.

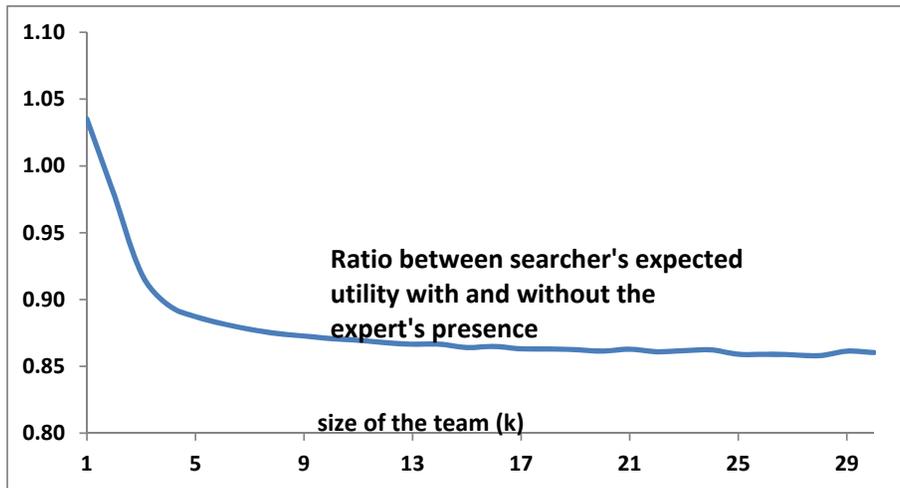


Figure 5: Ratio of the expected utility of a searcher with a self-interested expert in the market versus with no expert in the market as a function of k , the number of agents that have to all agree to form a team. The presence of the expert becomes more damaging as the team size increases.

A market designer 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. This could enhance social welfare; the market designer may have regulatory power over the market, but even if she does not, there are situations

where she may be able to completely compensate the expert for the loss she suffers from the change in query price, while still improving the social welfare of searchers enough to make it practical. For the rest of this section, we assume that the market designer makes side payments to the expert, through the subsidization framework introduced in the previous section.

The process is illustrated in Figure 6, which depicts the searcher’s and expert’s expected utility (and their sum, which is the social welfare) as a function of the fee charged by the expert c_e . The search cost is taken to be $c_s = 0.1$. 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 . In this example, the optimal expert query cost (the value where the lower curve, which demonstrates the expert’s profit as a function of query cost, peaks) is $c_e = 0.0052$ (with an expected utility for the expert of 0.00632).

In this case, social welfare is maximized when the effective query price paid is 0.019 (seen at the upper curve in the figure). The expert’s utility with $c_e = 0.019$, however, is 0.00236, thus she needs to be compensated in the amount of the difference, $0.00632 - 0.00236 = 0.00396$. This compensation is still less than the gain in the searcher’s expected utility (an increase from 0.173 to 0.182).

In the example depicted in Figure 6, the expert is compensated for charging *more* for her services. This result is surprising, since intuitively, one would expect that a reduction rather than an increase in expert query price should increase searcher’s utility and social welfare. This is demonstrated in Figure 7, where the social welfare is maximized when the effective query price paid is zero, i.e., when the expert offers her services for free and thus is used by any searcher throughout her entire search (the search cost in this example is taken to be $c_s = 0$). Yet, in many sensible market settings the contrary holds, such as in Figure 6. Figure 8 (in which $c_s = 0.15$) illustrates a more extreme setting, where the expert is in fact being compensated for increasing her fee to a level for which none of the searchers uses her services (i.e., the expert is actually pushed out of the market). Figure 9 illustrates a setting in which social welfare is maximized (approximately) when the expert charges her expected-utility maximizing fee, and hence there is no need for external interference in the form of subsidy or taxation ($c_s = 0.04$).

The above results hold also for the case where agents’ expected utility is discounted. Figure 10 illustrates such a scenario where searchers have a discount factor of 0.9. The figure depicts the searcher’s expected utility with a self-interested expert who charges a profit-maximizing amount, with an expert who provides her services for free, and with a self-interested expert who is compensated for charging the social-welfare-maximizing amount. The figure shows once again, that market design, in the form of changing 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.

5.3 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 11 (left) shows results for $c_e = 0.01$ in the case of the uniform distribution of signals, and Figure 12 (left) shows the thresholds when values and signals are normally distributed, with $c_e = 2.0$ (variations in the results are minor across several expert query costs). For the uniform distribution, 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. This is likely because of the finite support of signals – as t_l and t_u are both decreasing

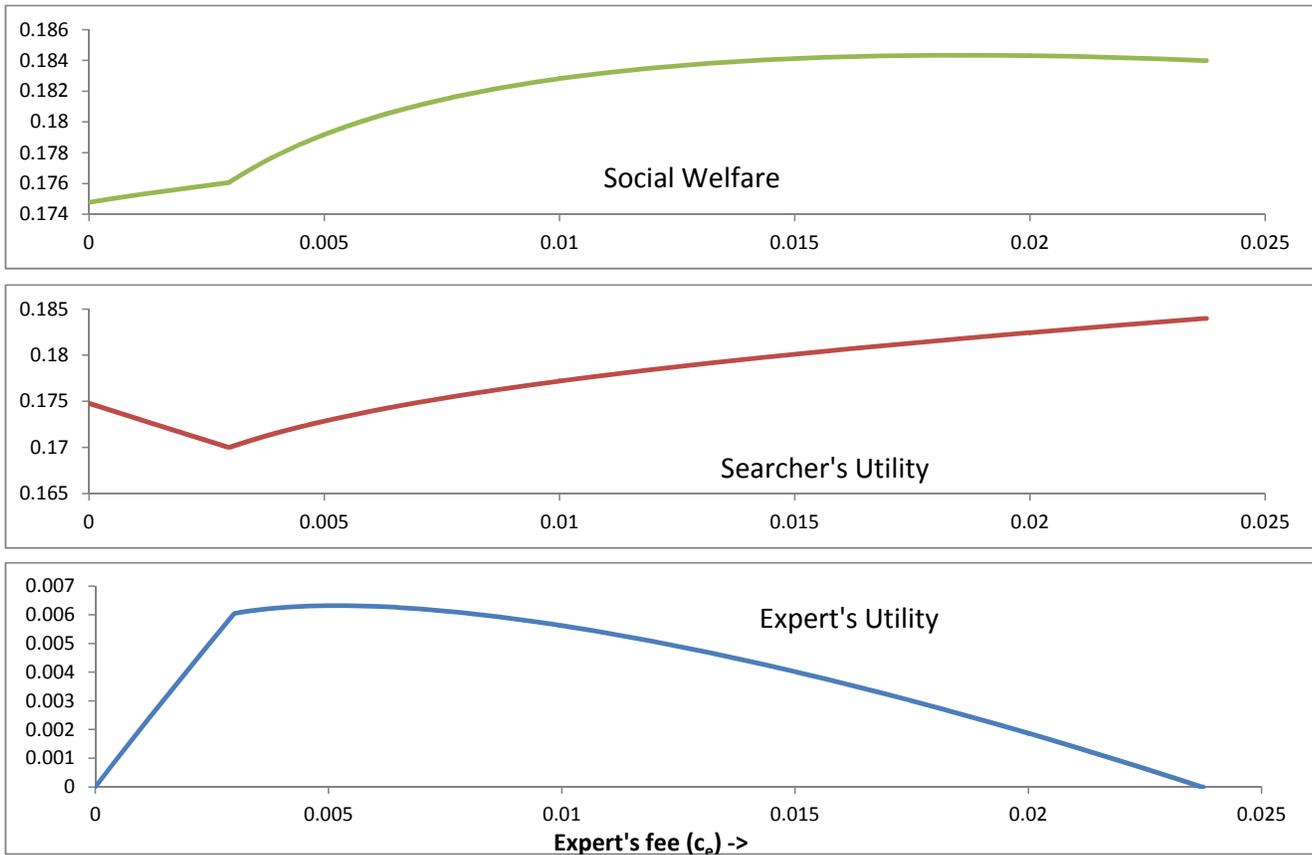


Figure 6: Using subsidy: by increasing the effective query price from 0.0052 to 0.019, a market designer can maximize social welfare. In this example $c_s = 0.1$.

as a function of c_s , they eventually merge. In the case of normally distributed values and noise, the thresholds can both decrease without being constrained by the signal support, and the expert still gets queried in a narrow band (Figure 12 (left)).

We can also examine the effect of expert query cost on the thresholds. Figures 11 and 12 both show similar behavior for the thresholds as a function of expert query costs, with the thresholds merging as query cost increases. In these cases the expert is not queried because it is simply too expensive.

6 A Commission-based Service

As an alternative to imposing a per-query cost, the expert may be able to operate on the basis of charging a commission when the use of her services leads to a successful match being made. This is similar to schemes used by realtors, who typically only get paid when a house is actually bought or sold. The most significant difference when the expert operates on the basis of a commission, that is only charged when a match is successful, is that there is no cost to a searcher to employ the expert's services when the revealed true value ends up being too low; therefore, it is natural that the optimal strategy for a searcher is to always use the expert when the signal is below a certain level. We prove the optimality of this single

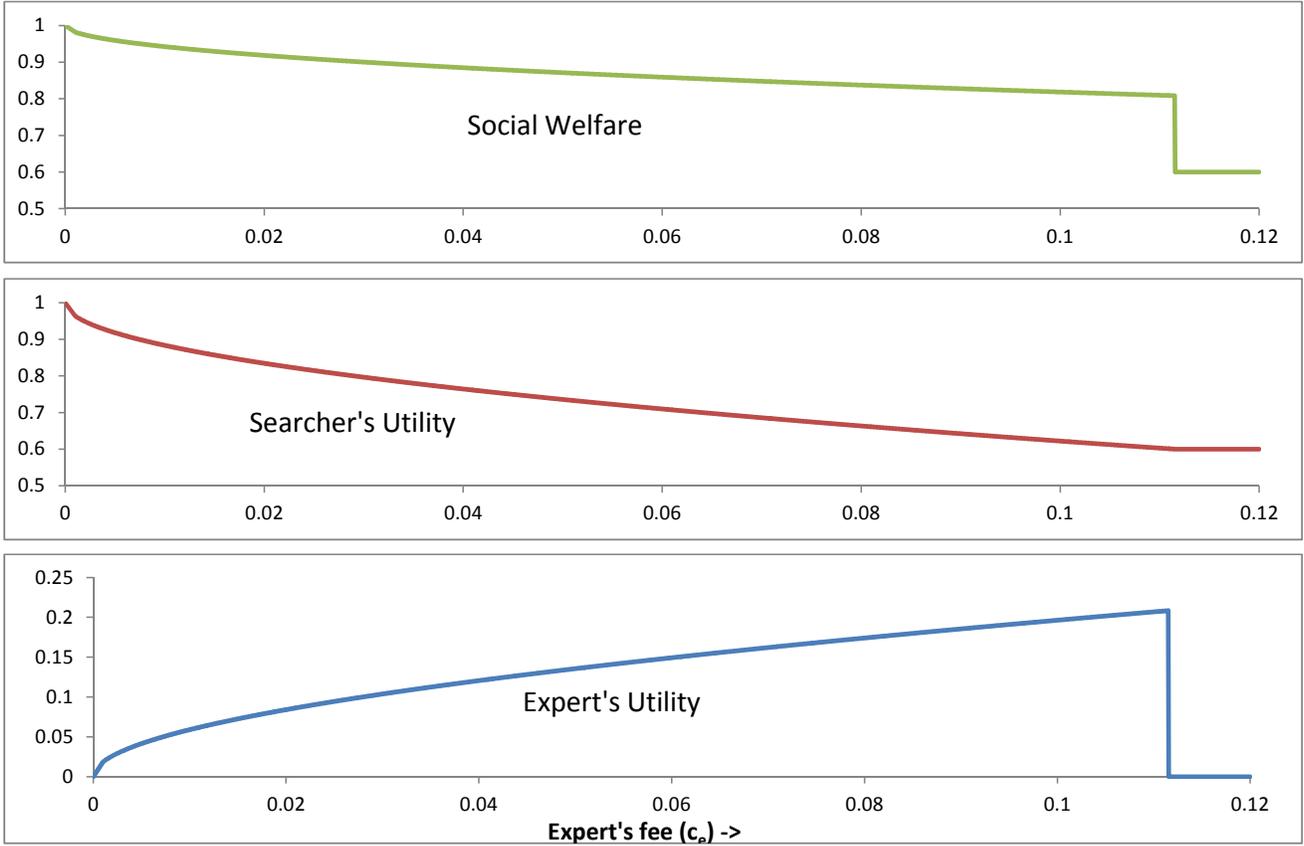


Figure 7: By decreasing the effective query price to 0, a market designer can maximize social welfare. In this example $c_s = 0$.

reservation value strategy below.

It seems that this scheme would not be beneficial to experts, because searchers can choose not to use the expert when the signal is high enough, but always use the expert when the signal is low. However, we show that there are situations where the expert can nevertheless benefit from this kind of pricing.

We now turn to specifying the model for commission-based pricing. Upon receiving a signal, an agent can either accept the partnership, decline it, or consult an expert who then reveals 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. In the latter case, if both agents accept the partnership then the agent will need to pay the expert a commission $\gamma_e v$, where v is the realized value from the evaluated partnership.⁸ We assume that agents truthfully report when a match is made.

We first prove the optimality of a strategy of the form (S, V) where: (a) S is a set of signal intervals for which the searcher should accept the partnership without querying the expert; and (b) for any signal that is not in S the searcher should query the expert, and accept the partnership if the value obtained is above a threshold $V/(1 - \gamma_e)$, and resume otherwise. Note that, in contrast to the fee-per-query model, there is no longer a set of signals where the agent should resume search without querying the expert.

⁸ γ_e denotes the commission, as a fraction of the value of the transaction to the agent, that the agent needs to pay the expert.

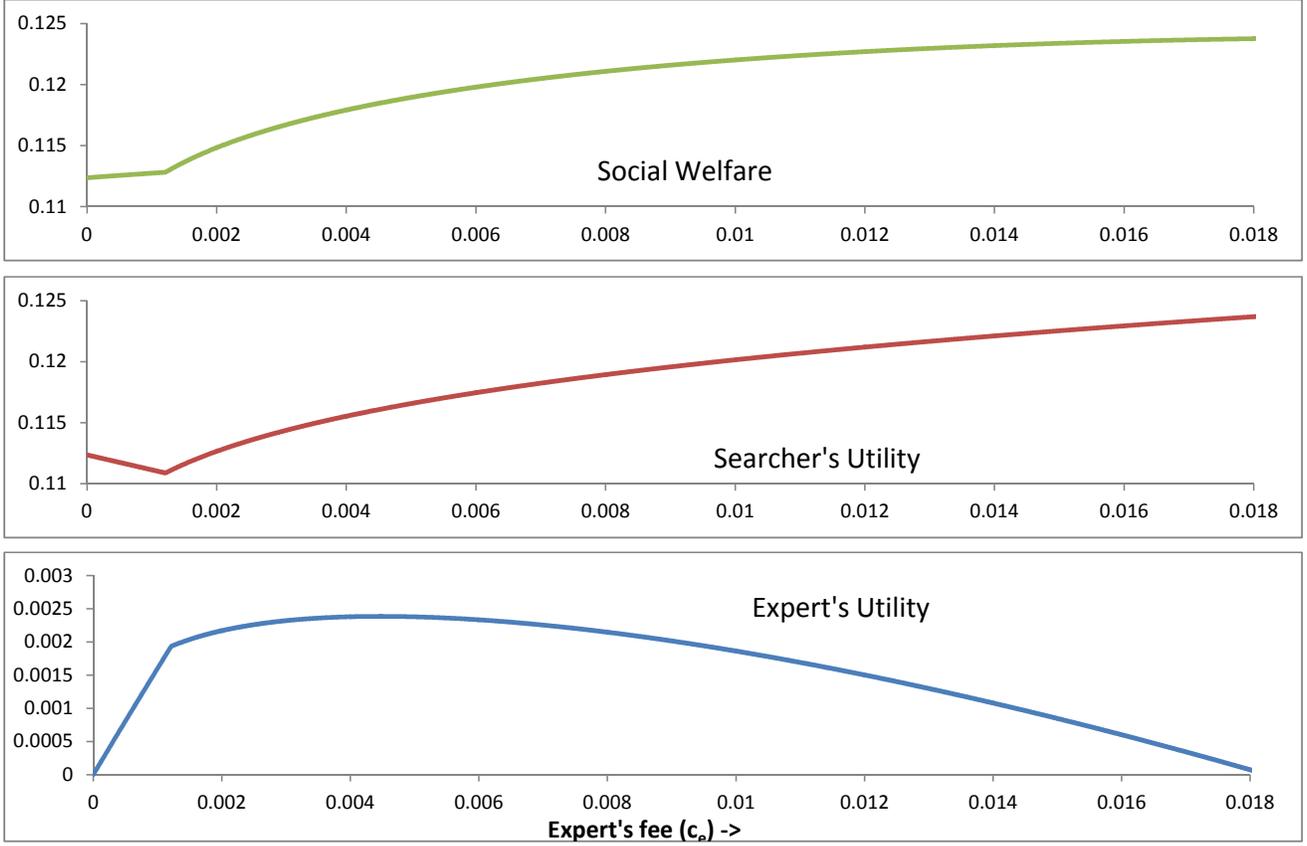


Figure 8: An example where the expert is in fact being compensated for setting a fee for which none of the searchers uses her services (i.e., actually pushed from the market). In this example $c_s = 0.15$.

V , the expected value of resuming search, is given by the following modification of Equation 5, given that the other agents use strategy $(S_{\text{others}}, V_{\text{others}})$:

$$V(S, V) = -c_s + (1 - A \cdot B) \cdot V(S, V) + B \cdot C \quad (27)$$

where A , B , and C are given by:

$$A = \int_{s \in S} f_s(s) ds + \int_{s \notin S} f_s(s) \left(1 - F_v(V/(1 - \gamma_e)|s)\right) ds$$

$$B = \int_{s \in S_{\text{others}}} f_s(s) ds + \int_{s \notin S_{\text{others}}} f_s(s) \left(1 - F_v(V_{\text{others}}/(1 - \gamma_e)|s)\right) ds$$

$$C = \int_{s \in S} f_s(s) E[v|s] ds + \int_{s \notin S} \left(f_s(s) \int_{y=V/(1-\gamma_e)}^{\infty} (1 - \gamma_e) y f_v(y|s) dy \right) ds$$

The value of $V(S, V)$ in Equation 27 is once again derived recursively, considering the next search iteration. The searcher pays c_s for receiving the noisy signal. Unlike in Equation 5, here there is no need to consider γ_e , as the commission is already subtracted from the value obtained in C . The other two

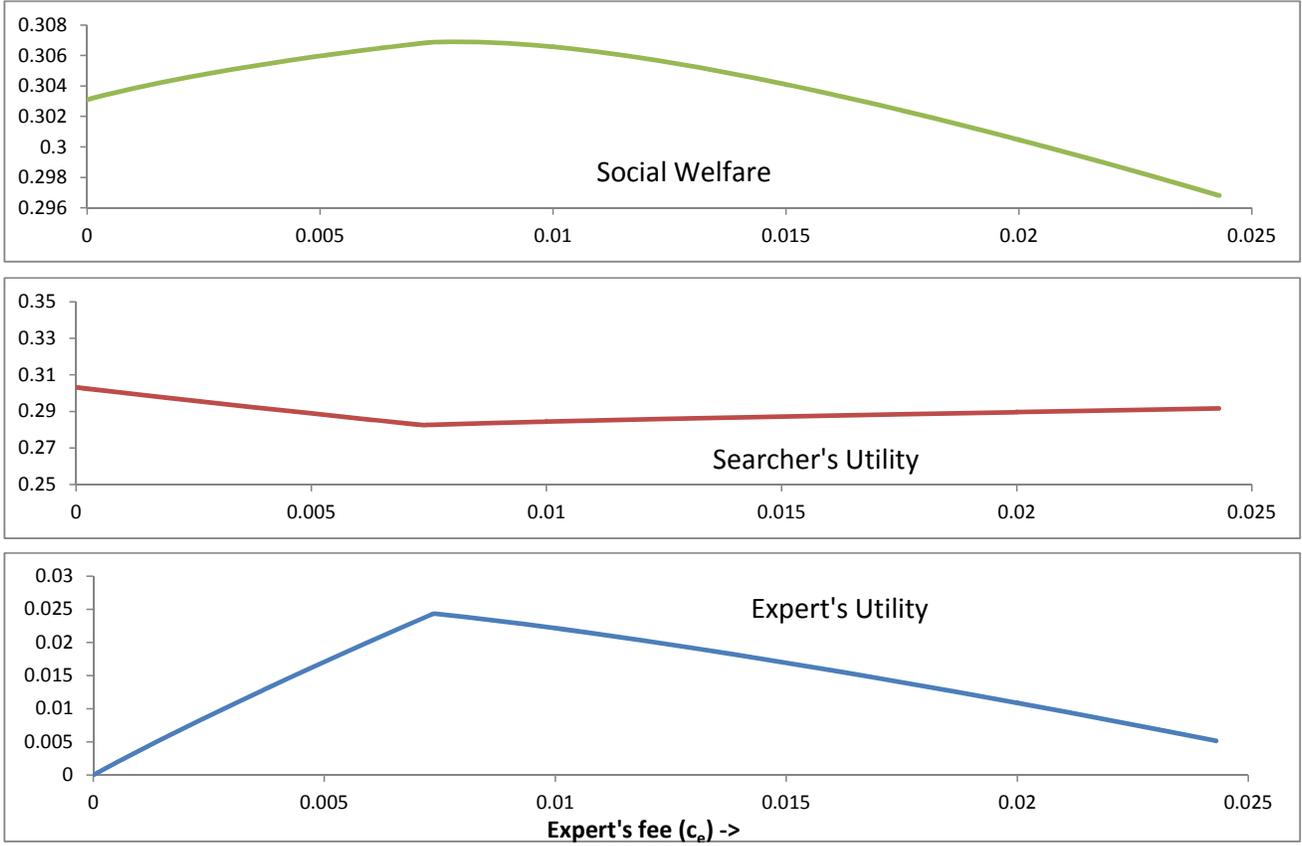


Figure 9: A setting in which the social welfare is maximized when the expert charges her expected-utility-maximizing fee, hence no need for external interference in the form of subsidy or taxation. In this example $c_s = 0.04$.

elements are similar to those used in Equation 5. Similarly, the calculations of A , B and C are similar to those used in Equation 5, except that the searcher accepts the match after querying the expert only if the value obtained is greater than $V/(1 - \gamma_e)$.

Based on the above, we can prove that under the HSGN assumption, the set S actually contains a single interval of signals.

Theorem 3 For $f_v(y|s)$ satisfying the HSGN assumption (Definition 1), given the search cost c_s and the commission γ_e set by the expert, the Bayes Nash Equilibrium for all agents can be described by the tuple (t, V) , where: (a) t is a signal threshold above which the partnership should be accepted; and (b) the expert should be queried given any signal $s < t$ and the partnership should be accepted if the value obtained from the expert is above $V/(1 - \gamma_e)$, where V is the expected value of resuming the search; otherwise search should resume. The equilibrium values t and V can be calculated from solving the set

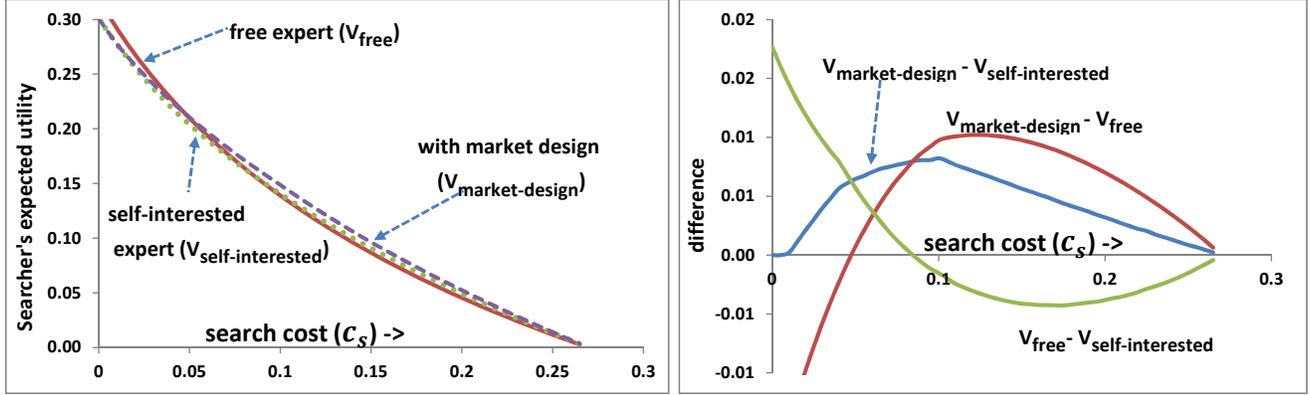


Figure 10: Searcher's utility as a function of search cost, when searchers use a discount factor of 0.9 (left), for the case: (a) where the expert is free to use ($c_e = 0$); (b) when the expert is self-interested and charges its utility-maximizing fee; and (c) when the expert is self-interested however market design is applied (using subsidy). The right-hand graph depicts the differences between any two pairs. Similar to the case without discounting, with the use of subsidy social welfare can be improved, even to a point where it is better than the social welfare when the expert does not charge for her services at all.

of Equations 28-32:

$$V = \frac{-c_s + BC}{AB} \quad (28)$$

$$\int_{y=-\infty}^{V/(1-\gamma_e)} (V - y) f_v(y|t) dy = \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} y f_v(y|t) dy \quad (29)$$

$$A = 1 - \int_{s=-\infty}^t f_s(s) F_v(V/(1-\gamma_e)|s) ds \quad (30)$$

$$B = A \quad (31)$$

$$C = \int_{s=t}^{\infty} f_s(s) E[v|s] ds + \int_{s=-\infty}^t f_s(s) \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e) y f_v(y|s) dy ds \quad (32)$$

Sketch of Proof: The proof resembles that of Theorem 2. We first 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 28, 30, and 32 are obtained after replacing the set of intervals S with the threshold t . Equation 31 represents the fact that the system is symmetric in the sense that ultimately all agents choose the same tuple (t, V) , and so the probability of being accepted is equal to the probability of accepting the match. Finally, we prove the correctness of Equation 29 by taking the derivative of Equation 28 w.r.t. t , equating to zero, and showing that there is a t which maximizes the expected benefit. \square

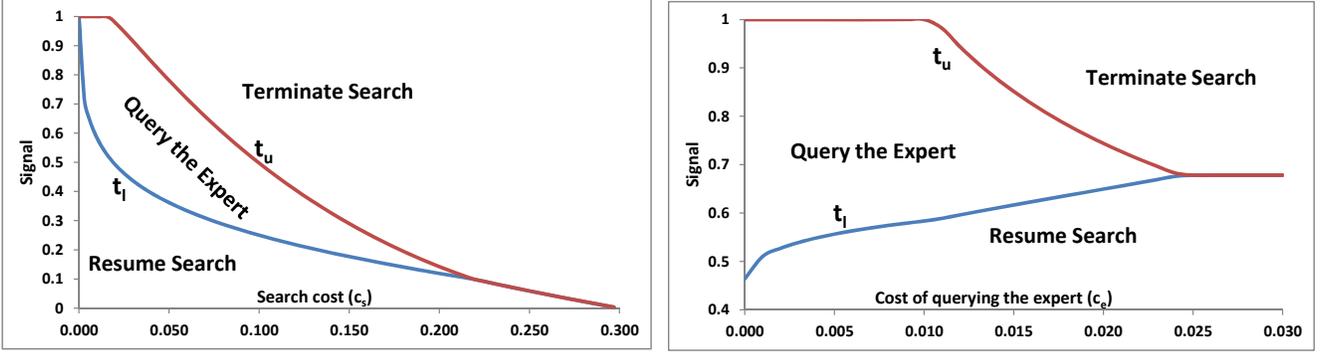


Figure 11: Left: Thresholds decrease and get closer to one another as the cost of search increases with uniformly distributed signals. The thresholds eventually merge, indicating that searchers do not query the expert for $c_s > 0.22$. In this example $c_e = 0.01$. Right: Thresholds get closer to one another and eventually merge as expert cost increases. In this example $c_s = 0.01$.

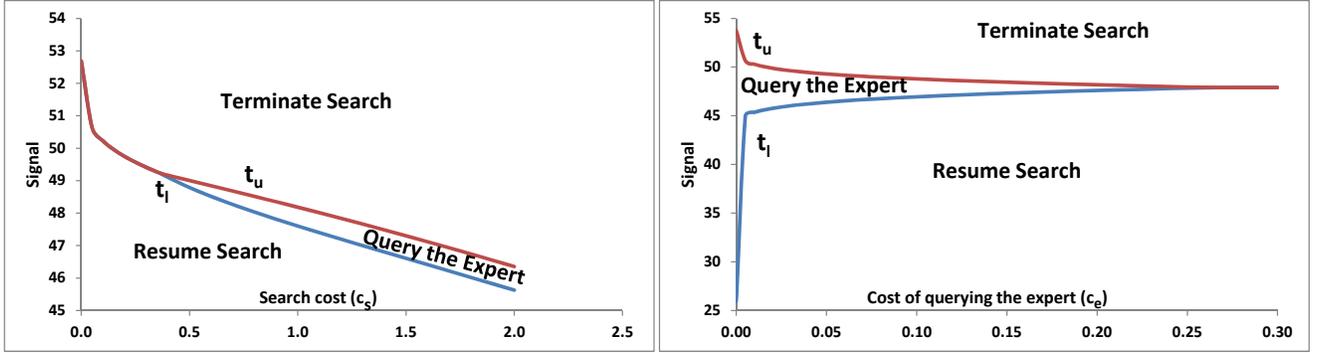


Figure 12: Thresholds for normally distributed values and noise as a function of search cost (left) and expert query cost (right).

We now have 5 equations (28 - 32) in 5 variables. We can solve these simultaneously to calculate the values of V and t .

Equation 29 captures the searcher's indifference, given a signal t , between querying the expert and accepting the partnership without querying the expert. Here, if accepting the partnership without querying the expert the searcher obtains $E[v|t]$ with probability B and V otherwise; however, if the searcher queries the expert it receives either the value obtained from the expert if it is above $V/(1 - \gamma_e)$ (with probability B , since this is the probability that the other side accepts the match), or otherwise V :

$$B \cdot E[v|t] + (1 - B)V = V \left(1 - B(1 - F_v(V/(1 - \gamma_e)|t)) \right) + B \int_{y=V/(1-\gamma_e)}^{\infty} y(1 - \gamma_e)f_v(y|t) dy$$

which transforms into Equation 29.

It is worth noting that the term under the integral in the left hand side of Equation 29 is negative for any value in the interval $V \leq y \leq V/(1 - \gamma_e)$. In these cases, it is optimal for the searcher to resume search even though she observes a true value greater than the expected benefit of resuming the search,

because the commission the searcher will need to pay if she accepts the partnership makes the transaction non-beneficial.

Expected number of queries Once again, the search strategy (t, V) defines the number of times the expert's services are used. The probability of having both sides accept the match is $A \cdot B$, and 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 $F_s(t)$, and so the expected number of expert queries a searcher performs is

$$\eta_{\gamma_e} = \frac{F_s(t)}{A \cdot B} \quad (33)$$

Expected profit of the expert The expected profit is the difference between the expected payment received and the expected expense. The expected expense is given by $d_e \eta_{\gamma_e}$. The expected payment the expert receives is

$$\begin{aligned} V_{expert} &= \sum_{j=0}^{\infty} \int_{y=V/(1-\gamma_e)}^{\infty} \int_{s=-\infty}^t f_s(s) f_v(y|s) \gamma_e y B (1-AB)^j ds dy \\ &= \frac{\int_{y=V/(1-\gamma_e)}^{\infty} \int_{s=-\infty}^t f_s(s) f_v(y|s) \gamma_e y ds dy}{A} \end{aligned} \quad (34)$$

Therefore: $\pi_e = \mathbb{E}(\text{Profit}) = \frac{\int_{y=V/(1-\gamma_e)}^{\infty} \int_{s=-\infty}^t f_s(s) f_v(y|s) \gamma_e y ds dy}{A} - d_e \eta_{\gamma_e}$. The expert can maximize the above expression with respect to γ_e to find the profit maximizing price to charge searchers using the same principles detailed in 4.5.

6.1 Illustrative Results

We again consider some actual numerical results for a particular distribution of values, focusing this time solely on the uniform distribution of signals, where $f_v(y|s) = \frac{3\sqrt{y}}{2\sqrt{s^3}}$ and $d_e = 0$. Many of the results are qualitatively quite similar to those above: for example, agents may be better off in markets where the expert is taxed or subsidized to increase her fees, and the market designer can solve for maximum social welfare in manners analogous to those for the per-query fee model. Therefore, we do not present these results in detail.

Instead, we demonstrate in Figure 13 that there are situations where the expert can make higher profits using a commission-based strategy than she can using the per-query pricing model presented in Sections 4 and 5. The figure shows the difference between the expert's utility when using her utility maximizing commission and when using her utility-maximizing fixed per-query fee, as a function of the search cost c_s , for the particular signal and value distributions we have used throughout. We see that the expert benefits from using commission-based pricing for any search cost $c_s < 0.002$, and from using per-query pricing for $c_s > 0.002$.

Although not definitive, these results are suggestive: perhaps situations with lower search costs offer more upside to experts using commission-based pricing because searchers are less prone to “settle” and end up with higher value items in the end. A percentage of that value is what goes to the expert (search costs are borne entirely by the searcher), so the expert is likely to do better in such situations.

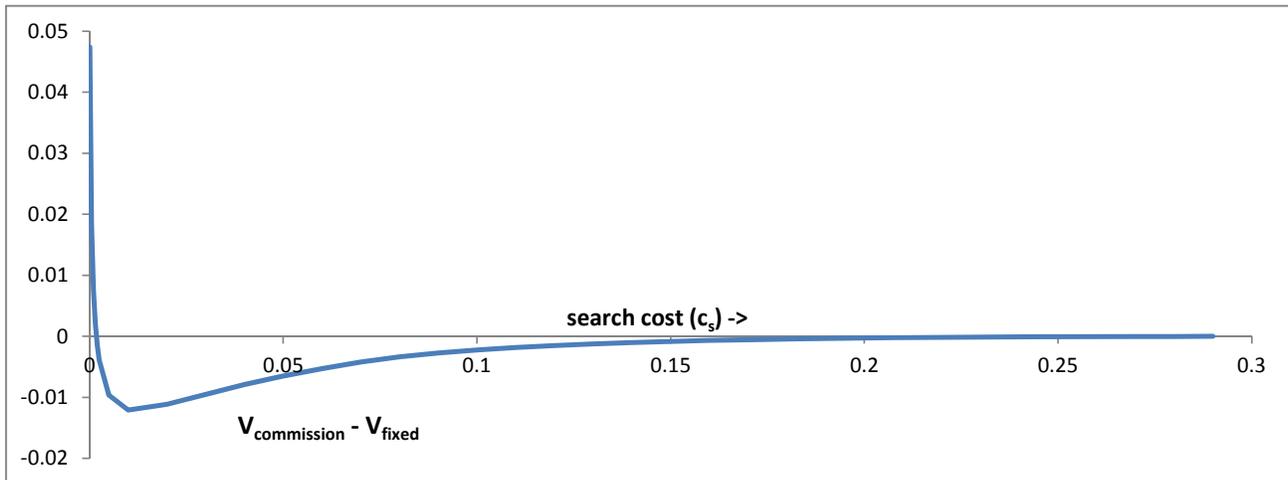


Figure 13: The difference between the expert’s utility when using her utility maximizing commission and when using her utility-maximizing fixed fee, as a function of the search cost c_s . A positive value suggests an advantage for the commission-based scheme and vice-versa. In this case, for a large search cost, the expert is better off using the commission-based scheme.

7 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 we show in Section 4, the analysis can be extended in a straightforward manner to various cases (e.g., the formation of k member teams, discounting of gains, and different population types).

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. First, we show a more general result: there are situations in two- (or many-) sided search (typically environments where search costs are high) where more information can be *bad*. That is, agents would all be better off having noisier signals than having perfect signals of the values of matchings. Generalizing this result to markets with experts, we show that the presence of experts can lead 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 becomes 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 who brings searchers and experts together and seeks to mitigate this effect. We propose that the market designer could incentivize monopolistic experts to *increase* their fees in order to make searchers query the expert less often.

Finally, we investigate, from the perspective of the expert, the differences between commission-based pricing and a per-query pricing model. We show that the commission-based scheme, even though it does not require searchers to pay experts unless a transaction (match or team-formation) is completed, can lead to greater profits for the expert in some situations. We present an example that suggests the commission-based scheme may be better for experts when searchers have to pay low search costs.

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 receives 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, and parallels many of the models used in two-sided matching literature, it would be illuminating to understand what happens when the quality of matches are quite differently determined. Some of the analysis can be extended in some quite straightforward manner, e.g., for the case where two agents get the same idiosyncratic utility for a match (a function of their compatibility, but the same for both) and agents are unaware of others’ use of the expert whenever having to decide on acceptance/rejection. In other cases, e.g., where what one agent gets from a match is only a function of the other agents’ “quality”, this is more complex, though the analysis presented in this paper should provide a good foundation for these extensions. Another interesting direction for future work is the analysis of a model where the expert merely supplies another noisy (though more accurate) signal. This model, however, imposes a modelling challenge, which is the modelling of the noise given the signal obtained and also the distribution of the true value for any pair of signals received by the searcher and the expert.

Acknowledgments

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 1083/13. Das acknowledges additional support from an NSF CAREER Award (IIS-1414452). A preliminary version of this work appears in the *Proceedings of the ACM Conference on Electronic Commerce (EC 2012)*.

References

- [1] J. Albrecht, A. Anderson, and S. Vroman. Search by committee. *Journal of Economic Theory*, 145(4):1386–1407, 2010.
- [2] E. Anshelevich, S. Das, and Y. Naamad. Anarchy, stability, and utopia: creating better matchings. *Autonomous Agents and Multi-Agent Systems*, 26(1):120–140, 2013.
- [3] M. Armstrong. Competition in two-sided markets. *RAND Journal of Economics*, 37(3):668–691, 2006.

- [4] Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, 43(12):1676–1692, 1997.
- [5] F. Bloch and H. Ryder. Two-sided search, marriages, and matchmakers. *International Economic Review*, 41(1):93–115, 2000.
- [6] K. Burdett and R. Wright. Two-sided search with nontransferable utility. *Review of Economic Dynamics*, 1:220–245, 1998.
- [7] H. Chade and G. Ventura. On income taxation and marital decisions. *Review of Economic Dynamics*, 8(3):565–599, July 2005.
- [8] M. Chhabra, S. Das, and D. Sarne. Competitive information provision in sequential search markets. In *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 2014. To appear.
- [9] M. Chhabra, S. Das, and D. Sarne. Expert-mediated sequential search. *European Journal of Operational Research*, 234(3):861–873, 2014.
- [10] S. Das and E. Kamenica. Two-sided bandits and the dating market. In *Proceedings of IJCAI*, pages 947–952, 2005.
- [11] P. Diamond. Aggregate demand management in search equilibrium. *The Journal of Political Economy*, 90(5):881–894, 1982.
- [12] D. Duffie, S. Malamud, and G. Manso. Information percolation with equilibrium search dynamics. *Econometrica*, 77(5):1513–1574, 09 2009.
- [13] U. Endriss, S. Kraus, J. Lang, and M. Wooldridge. Designing incentives for boolean games. In *Proceedings of AAMAS*, pages 79–86, 2011.
- [14] D. Fotakis, G. Karakostas, and S. Kolliopoulos. On the existence of optimal taxes for network congestion games with heterogeneous users. In *Proceedings of SAGT'10*, pages 162–173, 2010.
- [15] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Math. Monthly*, 69:9–15, 1962.
- [16] M. Gorman, D. Salisbury, and I. Brannon. Who wins when price information is more ubiquitous? an experiment to assess how infomediaries influence price. *Electronic Markets*, 19(2-3):151–162, 2009.
- [17] A. Grosfeld-Nir, D. Sarne, and I. Spieglier. Modeling the search for the least costly opportunity. *European Journal of Operational Research*, 197(2):667–674, 2009.
- [18] M. Janssen and J. Moraga-Gonzalez. Strategic pricing, consumer search and the number of firms. *Review of Economic Studies*, 71(4):1089–1118, October 2004.

- [19] B. Jovanovic. Job Matching and the Theory of Turnover. *The Journal of Political Economy*, 87(5):972, 1979.
- [20] J. Kephart and A. Greenwald. Shopbot economics. *Autonomous Agents and Multi-Agent Systems*, 5(3):255–287, 2002.
- [21] J. Kephart, J. Hanson, and A. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32:731–752, 2000.
- [22] S. Ketchpel. Forming coalitions in the face of uncertain rewards. In *In Proceedings of AAAI*, pages 414–419, 1994.
- [23] S. Kraus, O. Shehory, and G. Taase. Coalition formation with uncertain heterogeneous information. In *Proceedings of AAMAS*, pages 1–8, 2003.
- [24] R. Lee and M. Schwarz. Interviewing in two-sided matching markets. NBER Working Paper, 2009.
- [25] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14(2):155–189, 1976.
- [26] A. Masters. Wage posting in two-sided search and the minimum wage. *International Economic Review*, 40(4):809–826, 1999.
- [27] J. McCall. Economics of information and job search. *The Quarterly Journal of Economics*, 84(1):113–126, 1970.
- [28] J. McMillan and M. Rothschild. Search. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. North-Holland, 1994.
- [29] P. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12:380–391, 1981.
- [30] D. Mortensen and C. Pissarides. New developments in models of search in the labor market. *Handbook of labor economics*, 3:2567–2627, 1999.
- [31] M. Nermuth, G. Pasini, P. Pin, and S. Weidenholzer. Price dispersion, search externalities, and the digital divide. Vienna Economics Papers 0916, University of Vienna, Department of Economics, 2009.
- [32] M. Panos and L. Ungar. Shopbots and pricebots in electronic service markets. In S. Parsons, P. Gmytrasiewicz, M. Wooldridge, and G. Weiss, editors, *Game Theory and Decision Theory in Agent-Based Systems*, volume 5 of *Multiagent Systems, Artificial Societies, and Simulated Organizations*, pages 177–195. Springer US, 2002.
- [33] M. Penn, M. Polukarov, and M. Tennenholtz. Random order congestion games. *Mathematics of Operations Research*, 34(3):706–725, 2009.

- [34] M. Penn, M. Polukarov, and M. Tennenholtz. Taxed congestion games with failures. *Annals of Mathematics and Artificial Intelligence*, 56(2):133–151, 2009.
- [35] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4):689–711, 1974.
- [36] M. Rysman. The economics of two-sided markets. *Journal of Economic Perspectives*, 23(3):125–43, 2009.
- [37] P. Sahib and X. Gu. To Tie the Knot or Not: Cohabitation, Marriage and Individuals’ Attitudes to Risk. *Demographic Research*, 6(13):355–382, 2002.
- [38] D. Sarne and S. Kraus. Managing parallel inquiries in agents’ two-sided search. *Artificial Intelligence*, 172(4-5):541–569, 2008.
- [39] D. Sarne, E. Manisterski, and S. Kraus. Multi-goal economic search using dynamic search structures. *Autonomous Agents and Multi-Agent Systems*, 21:204–236, 2010.
- [40] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1):165–200, May 1998.
- [41] R. Shimer and L. Smith. Matching, search, and heterogeneity. *The B.E. Journal of Macroeconomics*, 1(1):1–18, 2001.
- [42] L. Smith. Frictional matching models. *Annual Reviews in Economics*, 3:319–338, 2011.
- [43] R. Waldeck. Search and price competition. *Journal of Economic Behavior & Organization*, 66(2):347–357, May 2008.
- [44] Y. Wan and G. Peng. What’s next for shopbots? *IEEE Computer*, 43:20–26, 2010.
- [45] D. Wiegmann, K. Weinersmith, and S. Seubert. Multi-attribute mate choice decisions and uncertainty in the decision process: A generalized sequential search strategy. *Journal of Mathematical Biology*, 60(4):543–572, 2010.
- [46] R. Wright. Job search and cyclical unemployment. *The Journal of Political Economy*, 94(1):38–55, 1986.

A Proofs

Proof of Theorem 1

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 is given by:

$$\begin{aligned} V(t) &= -c_s + V(t) \left(1 - B \int_{s=t}^{\infty} f_s(s) ds\right) \\ &+ B \int_{s=t}^{\infty} E[v|s] f_s(s) ds = \frac{-c_s + B \int_{s=t}^{\infty} E[v|s] f_s(s) ds}{B(1 - F_s(t))} \end{aligned} \quad (35)$$

where B is the probability of being accepted by the other searcher, and $F_s(s)$ is the cumulative distribution function of the signal s . Setting the first derivative according to t of Equation 35 to zero we obtain: $V(t^*) = E[v|t^*]$ as follows:

$$\begin{aligned} \frac{dV(t)}{dt} &= \frac{dV(t)}{dt} \left(1 - B \int_{s=t}^{\infty} f_s(s) ds\right) \\ &+ V(t) B f_s(t) - B \cdot E[v|t] f_s(t) \end{aligned}$$

$$V(t^*) = E[v|t^*] \quad f_s(t^*) \neq 0 \quad B \neq 0$$

To verify that t^* is global maximum, we calculate the second derivative.

$$\begin{aligned} \frac{d^2V}{dt^2} &= \left(1 - B \int_{s=t}^{\infty} f_s(s) ds\right) \frac{d^2V}{dt^2} + B \cdot f_s(t) \frac{dV}{dt} \\ &+ B \frac{d(V(t) f_s(t))}{dt} - B \frac{d(E[v|t] f_s(t))}{dt} \end{aligned}$$

$$\begin{aligned} \frac{d^2V}{dt^2} |_{t^*} &= \frac{1}{B(1 - F_s(t^*))} B \left(\frac{d((V(t) - E[v|t]) f_s(t))}{dt} |_{t^*} \right) \\ &= - \frac{f_s(t^*)}{1 - F_s(t^*)} \left(\frac{dE[v|t]}{dt} |_{t^*} \right) \end{aligned}$$

Given the HSGN assumption, $E[v|t]$ is an increasing function in t . Therefore $\frac{dE[v|t]}{dt} > 0$, this implies that the value of second derivative evaluated at t^* is less than zero, which confirms that t^* is indeed a global maximum.

Finally substituting $V(t^*) = E[V|t^*]$ in Equation 35 we obtain Equation 4 as follows:

$$\begin{aligned} V(t^*) &= -c_s + V(t^*) \left(1 - B \int_{s=t^*}^{\infty} f_s(s) ds\right) + B \int_{s=t^*}^{\infty} E[v|s] f_s(s) ds \\ E[v|t^*] &= -c_s + E[v|t^*] \left(1 - B \int_{s=t^*}^{\infty} f_s(s) ds\right) + B \int_{s=t^*}^{\infty} E[v|s] f_s(s) ds \\ c_s &= B \int_{s=t^*}^{\infty} (E[v|s] - E[v|t^*]) f_s(s) ds \end{aligned}$$

Since all the searchers end up using the same threshold, we get that:

$$B = \int_{s=t^*}^{\infty} f_s(s) ds$$

Therefore, the value t^* can be calculated using the above equation. \square

Proof of Theorem 2

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 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 36) and querying the expert (Equation 37):

$$V > E[v|s] \tag{36}$$

$$V > VF_v(V|s) + \int_{y=V}^{\infty} yf_v(y|s) dy - c_e \tag{37}$$

Given the HSGN assumption and since $s' < s$, Equation 36 holds also for s' . Similarly, notice that:

$$\begin{aligned} V &> VF_v(V|s) + \int_{y=V}^{\infty} yf_v(y|s) dy - c_e \\ &= V + \int_{y=V}^{\infty} (1 - F_v(y|s)) dy - c_e \\ &> V + \int_{y=V}^{\infty} (1 - F_v(y|s')) dy - c_e \\ &= VF_v(V|s') + \int_{y=V}^{\infty} yf_v(y|s') dy - c_e \end{aligned}$$

and therefore Equation 37 also holds for $s' < s$.

The proof for $s'' > s$ is exactly equivalent: the expected cost of accepting the current opportunity can be shown to dominate both resuming the search and querying the expert. The optimal strategy can thus

be described by the tuple (t_l, t_u, V) as stated above. Equations 6, 9, 11 are obtained after replacing the intervals S, S' with the thresholds t_l, t_u :

$$\begin{aligned}
V &= -c_s - c_e \int_{s=t_l}^{s=t_u} f_s(s) ds + (1 - AB)V + BC \\
V &= \frac{-c_s - c_e (F_s(t_u) - F_s(t_l)) + BC}{AB} \\
A &= \int_{s=t_u}^{\infty} f_s(s) ds + \int_{s=t_l}^{t_u} f_s(s) (1 - F_v(V|s)) ds \\
A &= 1 - F_s(t_l) - \int_{s=t_l}^{t_u} f_s(s) F_v(V|s) ds \\
C &= \int_{s=t_u}^{\infty} f_s(s) E[v|s] ds + \int_{s=t_l}^{t_u} \left(f_s(s) \int_{y=V}^{\infty} y f_v(y|s) dy \right) ds
\end{aligned}$$

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. Now we will show the correctness of Equations 7 and 8:

Taking the derivative of Equation 6 w.r.t. t_l and equating to zero, we obtain a t_l which maximizes the expected benefit.

$$\begin{aligned}
\frac{\partial V}{\partial t_l} &= \frac{(c_e f_s(t_l) + B \frac{\partial C}{\partial t_l}) AB + (-c_s - c_e (F_s(t_u) - F_s(t_l)) + BC) \frac{\partial A}{\partial t_l} B}{A^2 B^2} \\
&= \frac{(c_e f_s(t_l) + B \frac{\partial C}{\partial t_l}) AB - V AB \frac{\partial A}{\partial t_l} B}{A^2 B^2} \\
&= \frac{c_e f_s(t_l) + B \frac{\partial C}{\partial t_l} - V \frac{\partial A}{\partial t_l} B}{AB} \\
&= \frac{c_e f_s(t_l) + B (\frac{\partial C}{\partial t_l} - V \frac{\partial A}{\partial t_l})}{AB}
\end{aligned}$$

Calculating $\frac{\partial C}{\partial t_l} - V \frac{\partial A}{\partial t_l}$ separately, we get:

$$\begin{aligned}
\frac{\partial C}{\partial t_l} &= -f_s(t_l) \int_{y=V}^{\infty} y f_v(y|t_l) dy + \int_{s=t_l}^{t_u} f_s(s) (-V f_v(V|s)) \frac{\partial V}{\partial t_l} ds \\
\frac{\partial A}{\partial t_l} &= -f_s(t_l) + f_s(t_l) F_v(V|t_l) - \int_{s=t_l}^{t_u} f_s(s) f_v(V|s) \frac{\partial V}{\partial t_l} ds \\
&= -f_s(t_l) \int_{y=V}^{\infty} f_v(y|t_l) dy - \int_{s=t_l}^{t_u} f_s(s) f_v(V|s) \frac{\partial V}{\partial t_l} ds
\end{aligned}$$

and so

$$\frac{\partial C}{\partial t_l} - V \frac{\partial A}{\partial t_l} = -f_s(t_l) \int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy$$

and so

$$\begin{aligned} \frac{\partial V}{\partial t_l} &= \frac{c_e f_s(t_l) + B \left(-f_s(t_l) \int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy \right)}{AB} \\ &= \frac{f_s(t_l) (c_e - B \int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy)}{AB} \end{aligned}$$

Substituting $\frac{\partial V}{\partial t_l} = 0$ we get Equation 7.

To confirm that this is a maximum, we compute the second derivative:

$$\begin{aligned} AB \frac{\partial^2 V}{\partial t_l^2} + \frac{\partial A}{\partial t_l} B \frac{\partial V}{\partial t_l} &= \frac{\partial (f_s(t_l))}{\partial t_l} (c_e - B \int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy) \\ &\quad - f_s(t_l) B \frac{\partial \left(\int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy \right)}{\partial t_l} \end{aligned}$$

By substituting $\frac{\partial V}{\partial t_l} = 0$ and $c_e = B \int_V^{\infty} (y - V) f_v(y|t_l) dy$ we get:

$$AB \frac{\partial^2 V}{\partial t_l^2} = -f_s(t_l) B \frac{\partial \left(\int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy \right)}{\partial t_l}$$

Or alternatively:

$$\frac{\partial^2 V}{\partial t_l^2} = \frac{-f_s(t_l) \frac{\partial \left(\int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy \right)}{\partial t_l}}{A}$$

All that is left to show is that $\int_{y=V}^{\infty} (y-V)f_v(y|t_l) dy$ increases with t_l to complete the proof of maximum.

$$\begin{aligned}
\int_{y=V}^{\infty} (y-V)f_v(y|t_l) dy &= \lim_{b \rightarrow \infty} \int_{y=V}^b (y-V)f_v(y|t_l) dy \\
&= \lim_{b \rightarrow \infty} [(y-V)F_v(y|t_l)]_{y=V}^b + \lim_{b \rightarrow \infty} \int_{y=V}^b F_v(y|t_l) dy \\
&= \lim_{b \rightarrow \infty} (b-V)F_v(b|t_l) + \lim_{b \rightarrow \infty} \int_{y=V}^b F_v(y|t_l) dy \\
&= \lim_{b \rightarrow \infty} (b-V) + \lim_{b \rightarrow \infty} \int_{y=V}^b F_v(y|t_l) dy \\
&= \lim_{b \rightarrow \infty} \int_{y=V}^b (1-F_v(y|t_l)) dy = \int_{y=V}^{\infty} (1-F_v(y|t_l)) dy
\end{aligned}$$

and so

$$\begin{aligned}
\frac{\partial \left(\int_{y=V}^{\infty} (y-V)f_v(y|t_l) dy \right)}{\partial t_l} &= \frac{\partial \left(\int_{y=V}^{\infty} (1-F_v(y|t_l)) dy \right)}{\partial t_l} \\
&= - \left((1-F_v(V|t_l)) \right) \frac{\partial V}{\partial t_l} + \int_{y=V}^{\infty} \frac{\partial (-F_v(y|t_l))}{\partial t_l} dy > 0
\end{aligned}$$

The inequality is given by substituting $\frac{\partial V}{\partial t_l} = 0$ and due to HSGN assumption, which determines that $\forall y : \frac{\partial (F_v(y|t_l))}{\partial t_l} < 0$. The proof for Equation 8 is exactly similar to the proof for Equation 7. We have 6 equations (6 - 11) in 6 variables. We can solve these simultaneously to calculate the value of V , t_l and t_u . \square

Proof of Theorem 3

The proof steps resemble those used for Theorem 2. We first show that if, according to the optimal search strategy the searcher should terminate her search (and accept the partnership) given a signal s , then she must also necessarily do so given any other signal $s' > s$. The partnership is accepted based on signal s iff the following two conditions hold:

1. $E[v|s] > V$
2. $E[v|s] > VF_v(V/(1-\gamma_e)|s) + \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e)yf_v(y|s)dy$

Or equivalently:

1. $E[v|s] > V$
2. $\int_{y=-\infty}^{V/(1-\gamma_e)} (y-V)f_v(y|s)dy + \int_{y=V/(1-\gamma_e)}^{\infty} \gamma_e y f_v(y|s)dy > 0$

For every $s' > s$, we get that $E[v|s'] \geq E[v|s] > V$.

We show that for any $s' > s$:

$$\int_{y=-\infty}^{V/(1-\gamma_e)} (y-V)f_v(y|s')dy + \int_{y=V/(1-\gamma_e)}^{\infty} \gamma_e y f_v(y|s')dy \geq \int_{y=-\infty}^{V/(1-\gamma_e)} (y-V)f_v(y|s)dy + \int_{y=V/(1-\gamma_e)}^{\infty} \gamma_e y f_v(y|s)dy > 0$$

Rearranging the above equation we obtain:

$$\begin{aligned} & \int_{y=-\infty}^{V/(1-\gamma_e)} (y-V)[f_v(y|s') - f_v(y|s)]dy + \int_{y=V/(1-\gamma_e)}^{\infty} \gamma_e y [f_v(y|s') - f_v(y|s)]dy \\ &= [(y-V)[F_v(y|s') - F_v(y|s)]]_{y=-\infty}^{V/(1-\gamma_e)} - \int_{y=-\infty}^{V/(1-\gamma_e)} [F_v(y|s') - F_v(y|s)]dy \\ &+ [\gamma_e y [F_v(y|s') - F_v(y|s)]]_{y=V/(1-\gamma_e)}^{\infty} - \int_{y=V/(1-\gamma_e)}^{\infty} \gamma_e [F_v(y|s') - F_v(y|s)]dy \\ &= [(V/(1-\gamma_e) - V)(F_v(V/(1-\gamma_e)|s') - F_v(V/(1-\gamma_e)|s)) - 0] - \int_{y=-\infty}^{V/(1-\gamma_e)} [F_v(y|s') - F_v(y|s)]dy \\ &+ [0 - \gamma_e V/(1-\gamma_e)[F_v(V/(1-\gamma_e)|s') - F_v(V/(1-\gamma_e)|s)]] - \int_{y=V/(1-\gamma_e)}^{\infty} \gamma_e [F_v(y|s') - F_v(y|s)]dy \\ &= - \int_{y=-\infty}^{V/(1-\gamma_e)} [F_v(y|s') - F_v(y|s)]dy - \int_{y=V/(1-\gamma_e)}^{\infty} \gamma_e [F_v(y|s') - F_v(y|s)]dy \\ &\geq 0 + 0 = 0 \end{aligned}$$

where the last inequality is obtained since $F_v(y|s') < F_v(y|s)$ according to the HSGN assumption. Therefore, there is a single threshold for acceptance without querying the expert. Once this is established, Equations 28, 30, 32 are obtained by replacing the set of intervals S with the threshold t . Equation 31 represents the fact that the system is symmetric in the way that ultimately all agents choose the same tuple (t, V) , and so the probability of being accepted is equal to the probability of accepting the match. Finally, the correctness of Equation 29 is proved by taking the derivative of Equation 28 w.r.t. t , equating

to zero, obtaining t which maximize the expected benefit:

$$\begin{aligned}\frac{dV}{dt} &= \frac{B \frac{dC}{dt} AB - (-c_s + BC) \frac{dA}{dt} B}{(AB)^2} \\ &= \frac{B \frac{dC}{dt} AB - (VAB) \frac{dA}{dt} B}{(AB)^2} \\ &= \frac{\frac{dC}{dt} - V \frac{dA}{dt}}{A}\end{aligned}$$

however,

$$\frac{dA}{dt} = -f_s(t)F_v(V/(1-\gamma_e)|t) - \int_{s=t}^{\infty} f_s(s)f_v(V/(1-\gamma_e)|s)ds \frac{dV}{dt}$$

$$\begin{aligned}\frac{dC}{dt} &= -f_s(t)E(v|t) + f_s(t) \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e)yf_v(y|t)dy \\ &\quad + \int_{s=t}^{\infty} f_s(s)(-1)(1-\gamma_e)\frac{V}{(1-\gamma_e)}f_v(V/(1-\gamma_e)|s)ds \frac{dV}{dt} \\ &= -f_s(t)E(v|t) + f_s(t) \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e)yf_v(y|t)dy - \int_{s=t}^{\infty} f_s(s)Vf_v(V/(1-\gamma_e)|s)ds \frac{dV}{dt}\end{aligned}$$

and so:

$$\begin{aligned}
\frac{dV}{dt} &= \frac{\frac{dC}{dt} - V \frac{dA}{dt}}{A} \\
&= \frac{-f_s(t)E(v|t) + f_s(t) \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e)yf_v(y|t)dy + f_s(t)VF_v(V/(1-\gamma_e)|t)}{A} \\
&= \frac{f_s(t) \left(-E(v|t) + \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e)yf_v(y|t)dy + VF_v(V/(1-\gamma_e)|t) \right)}{A} \\
&= \frac{f_s(t) \left(- \int_{y=-\infty}^{\infty} yf_v(y|t)dy + \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e)yf_v(y|t)dy + V \int_{y=-\infty}^{V/(1-\gamma_e)} f_v(y|t)dy \right)}{A} \\
&= \frac{f_s(t) \left(- \int_{y=-\infty}^{V/(1-\gamma_e)} yf_v(y|t)dy - \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy + \int_{y=V/(1-\gamma_e)}^{\infty} (1-\gamma_e)yf_v(y|t)dy + V \int_{y=-\infty}^{V/(1-\gamma_e)} f_v(y|t)dy \right)}{A} \\
&= \frac{f_s(t) \left(\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy - \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy \right)}{A}
\end{aligned}$$

Substituting $\frac{dV}{dt} = 0$ and assuming that $f_s(t) \neq 0$, we obtain Equation 29

To verify maximum we compute the second derivative

$$\begin{aligned}
A \frac{dV}{dt} &= f_s(t) \left(\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy - \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy \right) \\
A \frac{d^2V}{dt^2} + \frac{dA}{dt} \frac{dV}{dt} &= \frac{df_s(t)}{dt} \left(\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy - \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy \right) \\
&\quad + f_s(t) \frac{d \left(\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy - \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy \right)}{dt}
\end{aligned}$$

Substituting $\frac{dV}{dt} = 0$, hence $\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy = \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy$, we obtain

$$\frac{d^2V}{dt^2} = \frac{f_s(t) \frac{d \left(\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy - \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy \right)}{dt}}{A}$$

But since $A, f_s(t) > 0$, it is enough to show that $\left(\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy - \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy \right)$ decreases as t increases. As part of the proof that the optimal strategy is based on reservation value, it was shown that:

$$\int_{y=-\infty}^{V/(1-\gamma_e)} (y-V)f_v(y|t)dy + \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy$$

increases as t increases.

Hence $\int_{y=-\infty}^{V/(1-\gamma_e)} (V-y)f_v(y|t)dy - \gamma_e \int_{y=V/(1-\gamma_e)}^{\infty} yf_v(y|t)dy$ decreases as t increases. \square

B Nomenclature

| Notation | Meaning |
|----------------------|---|
| s | The drawn signal. |
| v | The real value of the match. |
| $f_v(y), F_v(y)$ | The probability density function and cumulative distribution function from which the real values of the match are drawn. |
| $f_s(s), F_s(s)$ | The probability density function and cumulative distribution function from which the signals are drawn. |
| $f_s(s v), F_s(s v)$ | The probability density function and cumulative distribution function from which the signals are drawn, given that the true value of the match is v . |
| $f_v(y s), F_v(y s)$ | The probability density function and cumulative distribution function from which the real values of the matches are drawn, given that the drawn signal is s . |
| $E[v s]$ | The expected value of the match, given that the drawn signal is s . |
| c_s | The search cost. i.e., the cost of drawing another signal. |
| c_e | The cost of querying the expert. |
| γ_e | The commission (in percentages) paid to the expert. |
| V | The expected utility of an agent, using the optimal strategy. By Bellman optimality principle, this value is also the threshold for the real value revealed by the expert. |
| t_l | The lower threshold. Drawing a signal which is less than t_l , the agent rejects the match without querying the expert. |
| t_u | The upper threshold. Drawing a signal which is greater than t_u , the agent accepts the match without querying the expert. |
| $V(S, S^*)$ | The expected utility of an agent in a noisy two-sided search with no expert present, where the agent accepts the match if and only if the signal she sees is in S , and the other agents accept the match if and only if the signal they see is in S^* . |
| $V(S', S'')$ | The expected utility of an agent in a noisy two sided search with an expert present, where the agent accepts the match if the signal she sees, denoted s , is in S' , rejects it if $s \in S''$, queries the expert if $s \notin S' \cup S''$, and whenever the expert is queried, the agent accepts if and only if the revealed value is greater or equal to $V(S, S'')$. This is defined for $S' \cap S'' = \emptyset$. |
| d_e | The cost to the expert of producing the exact value of a match. |
| η_{c_e} | The expected number of expert queries a searcher performs, whenever the expert charges a fixed price c_e . |
| η_{γ_e} | The expected number of expert queries a searcher performs, whenever the expert charges a fixed commission percentage γ_e . |
| π_e | The expected utility of the expert. |
| W | The expected social welfare, defined as $\pi_e + V$. |
| k | The size of the match formed ($k = 1$ implies a one-sided search, $k = 2$ implies a two-sided-search. etc.). |