

Strategic Free Information Disclosure for Search-Based Information Platforms

Shani Alkoby
Bar-Ilan University
Ramat Gan, 5290002 Israel
shani.alkoby@gmail.com

David Sarne
Bar-Ilan University
Ramat Gan, 5290002 Israel
sarned@cs.biu.ac.il

Sanmay Das
Washington Univ. in St. Louis
St. Louis MO 63130
sanmay@seas.wustl.edu

ABSTRACT

We investigate information platforms that enable and support user search. Consider users engaged in a sequential search process (e.g. for used cars or consumer goods in e-commerce, or partners on a dating website). Many platforms provide basic information on opportunities of interest for free, while also offering, at a price, premium services that can offer more information to the user on the potential values of different opportunities. Prior research has focused on the question of how to price such services. Here we investigate a novel strategic option: can the platform provide some of the premium services for free, and increase its profit in doing so? By analyzing game theoretic equilibria in such a model, we show that there are cases where the platform can indeed benefit by sometimes providing information for free. The underlying mechanism is that sometimes offering free services leads to more extensive usage of the expert's paid services. A robustness analysis shows that even if the population of users is heterogeneous and a large portion of it a priori does not use the premium services, offering parts of the service for free can still be beneficial for the platform despite the potential misuse.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Economics

Keywords

Economics of information; Sequential search; Information provider; Information broker; Free information disclosure

1. INTRODUCTION

One of the great successes of the Internet has been to reduce the costs inherent in acquiring information of all kinds. Information platforms of different kinds connect users with the types of opportunities that they are potentially interested in. These platforms often, either implicitly or explicitly, guide a process of search carried out by users. For example, e-commerce platforms like eBay make it easy to search for consumer goods; AutoTrader makes it easy to

search for used cars; Match.com makes it easy to search for romantic partners. The ease with which these Internet-based platforms allow users to locate relevant opportunities has led to a resurgence of research in the theory and applications of sequential search, with the understanding that the order-of-magnitude reduction in search costs (particularly the opportunity cost of time) changes the game and necessitates new methodologies for analyzing these markets [2, 29, 13].

Another concomitant development has been the emergence of a new class of information brokers that serve as intermediaries, typically by helping users to evaluate the relative values of different opportunities that may be available to them (for example, Carfax and Autocheck in the used-car space, reputation systems in eBay and other auction sites, electronic and human “compatibility consultants” in dating sites).

In many cases, the platform itself offers these information services as part of a “premium package.”

The typical model of these premium information services is one where users receive noisy signals of the true values of opportunities, and can pay for a premium feature (or external service) that provides more information, helping to disambiguate the uncertainty in the original signal [5, 26].

The study of the strategic behavior of these information intermediaries, whether independent or provided by the platform, has focused primarily on how they should price their services [15, 30, 20, 24, 35, 4, 12, 8, 33]. When intermediaries are paid on a per-use basis (rather than, for example, in commission upon the completion of a transaction), their incentives can become complicated. This is because, for a given user, when the intermediary reveals to the user that an opportunity is a good fit, and the user stops searching and leaves the market, she does not use the intermediary's services any further, cutting off the revenue stream. Therefore, it is typically assumed that the intermediary must be honest for reputation reasons. However, even this, and the literature on this problem thus far, fails to take into account other ways in which the intermediary can remain honest but still increase the probability of extending a user's search process: specifically, it is theoretically possible that the intermediary could sometimes offer to provide extra information for *free* (say for some range of signals received by the user), and, in doing so, actually increase the probability that the user does not terminate her search process and leave the market.

In this paper, we show that this theoretical possibility is realizable. Our contributions are threefold. First, we provide an equilibrium analysis for a model of sequential search where the platform or external information provider, in addition to choosing the single price it usually charges for its services, can also offer its services for free whenever approached by the searcher. We prove the existence of a unique equilibrium structure in this model and provide the set of equations from which it can be extracted for any given settings. Sec-

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May, 4–8, 2015, Istanbul, Turkey.* Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

ond, we provide a proof-by-example that free information disclosure can be beneficial. Third, we provide an important robustness-check of the result that free information disclosure can increase profits. The first-order concern, when providing free services, is that a misspecified model of the population can have disastrous consequences – for example, if there exists a group that is characterized by a very low search cost, and members of this group never use the intermediary’s services (because the intermediary charges a cost which is too high for them), the intermediary may be unaware of their existence. However, by offering some services for free, the intermediary may expose itself to much higher costs from this group it was previously unaware of. We demonstrate that our example is quite robust to this concern, by showing the percentage of this hidden population would have to be very large to make it unprofitable to use the free information revelation strategy. Taken together, our results suggest that information intermediaries in search-based electronic marketplaces may benefit from disclosing some extra information for free, and that this should be part of the strategic arsenal in algorithmic pricing of information services.

2. MODEL

We consider a standard searcher-platform model (e.g., [15]) in which users, denoted *searchers*, login to the information platform in order to gain access to information about opportunities of the type they seek (e.g., cars, mortgages, consumer products, dates). Due to the high rate of new opportunities arriving to the platform, in practice, we can view it as enabling access to an unlimited stream of opportunities. Each searcher is interested in finding the single best opportunity for them (for example, a searcher would be looking to buy just one used car), so, once they decide on one, we model them as leaving the platform. While unaware of the specific value v of each opportunity listed in the platform, the searcher does know the (stationary) probability distribution function from which opportunities values are drawn, denoted $f_v(x)$. For a cost c_s (monetary, opportunity cost, etc.), the searcher can acquire a signal s , which is correlated with the true value v of an opportunity according to a (known) probability density function $f_s(s|v)$. We assume that higher signals are good news (HSGN), i.e., that if $s_1 > s_2$ then $\forall y, F_v(y|s_1) \leq F_v(y|s_2)$ [25].

The searcher may query and obtain the true value v of an opportunity for which signal s was received, by paying an additional fee c_e . This true value could be obtained from either the information platform that lists the different opportunities or from an external expert (e.g., Carfax, or a mechanic). We assume that the platform or expert pays a marginal cost d_e per query (i.e., a “production cost”). For exposition purposes we will use “platform” or “expert” interchangeably to denote this information provider. The goal of the searcher is to maximize the total utility received i.e., the expected value of the opportunity eventually picked minus the expected cost of search and expert fees paid along the way. Thus far, this model is quite standard in prior work [5, 34, 26, 21].

The main departure from previous work in terms of our model is that the expert is allowed to disclose the true value v for free if she determines that this is beneficial. So, for example, if a potential buyer comes to a mechanic with a Carfax report indicating a certain set of flaws, the mechanic may decide to do a free check-up for that car.

We note that the signals received by the searcher are the only form of price discrimination allowed in the model, and thus the only basis on which the free service can be provided in place of the paid service.

Therefore, the model now is as follows. At the very beginning, the expert determines the price she is willing to sell her services for (c_e). Then the search process begins. The searcher receives a sig-

nal s ; he reveals the signal s to the expert, who must then decide whether to offer her services either for free, or at cost c_e . If she offers the information for free, the searcher takes advantage of the offer, finds out the true value v , and then must decide whether to terminate search and take that opportunity, or to continue search, receiving a new signal s and repeating the process. If the expert chooses not to offer the information for free, the searcher must decide whether to purchase the expert’s services at cost c_e . If he does purchase the services, he again finds out the true value v , and then must decide whether to terminate search and take that opportunity, or to continue search. If he does not, then he must decide whether to terminate search and take that opportunity without knowing the true value v , only the signal s , or whether to decline the opportunity and continue search. Figure 1 shows the process in the form of a flowchart.

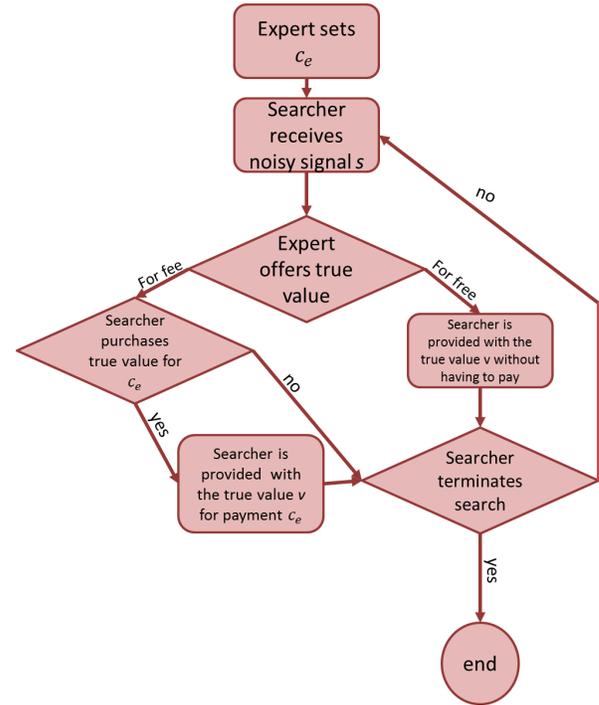


Figure 1: Flowchart of the sequential model where the expert may choose to disclose information for free.

3. EQUILIBRIUM ANALYSIS

No Free Information Disclosure. When the true value is offered by the expert for a fixed fee the game can be solved as a simple Stackelberg game where the expert is the leader, setting the service fee and the searchers are the followers, setting their search strategy accordingly. The searcher in this case, upon evaluating an opportunity and receiving its noisy signal s , can either: (a) reject it and continue search by evaluating a new opportunity; (b) accept it and terminate search; or (c) query the expert to know the true value of the opportunity, incurring a cost c_e , and, based on the value received, either accept it (terminating the search) or reject it and continue search as before. The optimal strategy for a searcher in this case can be found in prior work (e.g., [21, 5]): it is based on a tuple (t_l, t_u, V) (see Figure 2) such that for any signal s : (a) the search should resume if $s \leq t_l$; (b) the opportunity should be accepted if $s \geq t_u$;

and (c) the expert should be queried if $t_l \leq s \leq t_u$ and the opportunity accepted (and search terminated) if the value obtained from the expert is above the expected utility of resuming the search, V , otherwise search should resume. This is where V denotes the expected utility-to-go of following the optimal search strategy. The values of t_l, t_u and V can be extracted by solving a set of equations capturing two key indifference situations. The first is where the searcher is indifferent between resuming search and querying the expert (for t_l) and the second when he is indifferent between terminating search and querying the expert (for t_u) [21, 5].

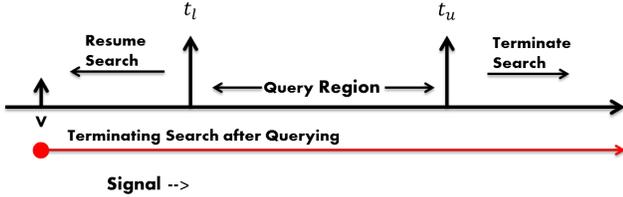


Figure 2: Characterization of the optimal strategy for search with an expert (taken from [5]). The searcher queries the expert if $s \in [t_l, t_u]$ and accepts the offer if its value is greater than the value of resuming the search V . The searcher rejects and resumes search if $s < t_l$ and accepts and terminates search if $s > t_u$, both without querying the expert.

With Free Information. When the expert is allowed to offer the true value for some of the signals for free, the equilibrium dynamics become more complex—when setting its service price c_e the expert needs to consider the equilibrium of the simultaneous game resulting from her decision, in which the searcher decides on its search strategy and the expert on the signals for which she will provide the true value for free. The key for solving the problem is therefore understanding the structure of the equilibrium of the resulting simultaneous game given the price c_e set by the expert. Theorem 1 provides the structure of the equilibrium for the simultaneous game, showing that it can be compactly represented in the form of four thresholds.

Theorem 1. *The equilibrium when the expert is able to disclose information for free, by choice, can be characterized according to the tuple (t_l, t_u, V, t_k) where (see Figure 3): (a) the information is offered for free for any signal $t_u \leq s \leq t_k$; (b) the searcher resumes its search for any signal s such that $s \leq t_l$; (c) the searcher accepts any opportunity associated with a signal $s \geq t_k$ and terminate its search right after; (d) the searcher queries the expert for any signal $t_l \leq s \leq t_k$, either for free (if $s > t_u$) or for a cost c_e (otherwise) and accept the opportunity (and terminate search) if the value obtained from the expert is above the expected utility of resuming the search, V , otherwise search is resumed. The values of t_l, t_u, V, t_k can be extracted by solving the set of equations:*

$$V = \frac{-c_s - c_e(F_s(t_u) - F_s(t_l)) + C}{A} \quad (1)$$

$$c_e = \int_{y=V}^{\infty} (y - V) f_v(y|t_l) dy \quad (2)$$

$$c_e = \int_{y=-\infty}^V (V - y) f_v(y|t_u) dy \quad (3)$$

$$d_e = \pi_e(F_v(V|t_k)) \quad (4)$$

where:

$$A = 1 - F_s(t_l) - \int_{s=t_l}^{t_k} f_s(s) F_v(V|s) ds \quad (5)$$

$$C = \int_{s=t_k}^{\infty} f_s(s) E[v|s] ds + \int_{s=t_l}^{t_k} f_s(s) \int_{y=V}^{\infty} y f_v(y|s) dy ds \quad (6)$$

$$\pi_e = \frac{(c_e - d_e)(F_s(t_u) - F_s(t_l)) - d_e(F_s(t_k) - F_s(t_u))}{A} \quad (7)$$

Proof. We distinguish between three sets of signals. The first, denoted S_{resume} , is the set of signals for which if information is not received for free then the searcher's best response strategy is to resume search without querying the expert. The second, denoted S_{query} , is the set of signals for which even if the information is not free, the searcher's best response strategy is to query the expert, and finally the set $S_{\text{terminate}}$ denoting the set of signals for which if the information is not free, the searcher's best response is not to query the expert but rather to accept the opportunity and terminate the search. We first prove that from the expert's point of view, if the best response to the searcher's strategy is not to offer the information for free for a signal $s \in S_{\text{terminate}}$ then so is the case for any other $s' \in S_{\text{terminate}}$ as long as $s' > s$. By providing the information for free when the signal is s the expert incurs a cost d_e , however gains π_e if instead of terminating her search (as is the searcher's strategy for a signal $s \in S_{\text{terminate}}$) the searcher, based on the true value received, decides to resume the search. The searcher will decide to resume search only if realizing that the true value is less than the expected benefit of further searching, i.e., if the true value is smaller than V . The probability of the latter event is given by $F_v(v|s)$, hence if the expert prefers not to provide the information for free given signal s then the following must hold:

$$d_e \geq \pi_e(F_v(V|s)) \quad (8)$$

Notice that $F_v(v|s) > F_v(v|s')$ for $s' > s$ (due to the HSGN assumption), hence $d_e \geq \pi_e(F_v(v|s)) > \pi_e(F_v(v|s'))$ and therefore the expert necessarily finds it beneficial not to offer the information for free for s' . This is in fact all that needs to be proved for the expert's strategy structure. Obviously there is no benefit from the expert's point of view to offer the information for free for any signal $s' \in S_{\text{query}} \cup S_{\text{resume}}$ as doing so has no immediate benefit and can only potentially eliminate further search rounds (if the reported true value is greater than V) and future profits.

Moving on to the searcher, we prove that given the above strategy structure of the expert, the searcher's best response strategy is of the (t_l, t_u, V, t_k) structure. First, we prove that given a signal $s \in S_{\text{resume}}$, any other signal $s' < s$ also belongs to S_{resume} . The proof is quite straightforward: 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 in this case is resuming the search.

Next, we prove that given a signal $s \in S_{\text{terminate}}$, any other signal $s' > s$ also belongs to $S_{\text{terminate}}$. This proof is also quite straightforward: the searcher decides to terminate the search in case where $E[v|s] > V$. According to the HSGN assumption it is clear that for every $s' > s$ we get that $E[v|s'] \geq E[v|s] > V$.

The structure of the searcher's strategy, for cases where the information is not offered for free, is thus based on three continu-

ous intervals, represented by (t_l, t_u) , where all signals $s < t_l$ belong to S_{resume} , all signals $s > t_u$ belong to $S_{\text{terminate}}$ and all signals $t_l < s < t_u$ belong to S_{query} .

At this point, we have everything we need in order to prove that the information will be provided for free only for signals belonging to the continuous interval (t_u, t_k) . We have already established the fact that the information provider will never offer the information for free for signals belonging to S_{query} and S_{resume} . Now assume there are signals s and s' , such that $s' > s > t_u$ and the expert's best response strategy is to offer the information for a fee for s' and not for free for s . We have already shown that if $s' > s > t_u$ then both signals belong to $S_{\text{terminate}}$. However, if both belong to $S_{\text{terminate}}$ and the best response strategy of the information provider is not to provide the information for free for s then, as shown at the beginning of the proof, so is her strategy for s' , which leads to a contradiction. Therefore, the set of signals for which information is provided for free is necessarily a continuous interval that starts at t_u .

The searcher therefore will receive the information for free for all signals in the interval (t_u, t_k) and will query the expert (for a fee) for all signals in the interval (t_l, t_u) . In both cases, if the information obtained indicates a value greater than her expected benefit from resuming the search the process will be terminated and otherwise resumed.

Once establishing the general (t_l, t_u, V, t_k) structure we can now formally express the expected profit for the searcher, V , and use optimization for deriving his best response set (t_l, t_u) . The searcher's expected profit is given by (1). Here the numerator captures the expected profit within a single search round. This is composed by the cost of receiving the signal, c_s , the expected cost of querying the expert, $c_e(F_s(t_u) - F_s(t_l))$, and the expected benefit of the searcher when stopping the search (without taking into consideration the cost of the search or the cost of using the expert), C , as calculated in (6). The calculation of C in (6) is based on three cases: (i) in case where the value of the signal s is higher than t_k , the searcher's expected profit will be the expectancy of V given the signal ($\int_{s=t_k}^{\infty} f_s(s)E[V|s] ds$) (ii) in the case where the value of the signal s is in the range of $[t_l, t_u]$ the searcher will stop the search only if the true value of the item is greater than V and in those cases will gain this value ($\int_{s=t_l}^{t_u} f_s(s) \int_{y=V}^{\infty} y f_v(y|s) dy ds$) (iii) in the case where the value of the signal s is in the range $[t_u, t_k]$ the searcher again will only stop the search if the item's true value is greater than V and will then gain this true value ($\int_{s=t_u}^{t_k} f_s(s)$

$\int_{y=V}^{\infty} y f_v(y|s) dy ds$). We note that since the choice of t_u does not affect C , cases (ii) and (iii) were merge to one integral in Equation 3, as will be done in the last two cases of Equation 5 to be described. The denominator in (1), A , calculated according to (5), is the probability that the searcher will terminate the search and purchase the offered item. The searcher will terminate search unless: (i) the value of the signal s is smaller than the value t_l (i.e., with probability $F_s(t_l)$); (ii) the value of the signal s is in the range $[t_l, t_u]$ and the true value of the item is smaller than V (i.e., with probability $\int_{s=t_l}^{t_u} f_s(s)F_v(V|s) ds$); (iii) the value of the signal is in the range $[t_u, t_k]$ and the true value of the item is smaller than V (i.e. with probability $\int_{s=t_u}^{t_k} f_s(s)F_v(V|s) ds$).

Setting the first derivative of V according to t_l and t_u to zero obtains Equations 2 and 3. Finally, Equation 4 represents the best response strategy for the auctioneer as explained above.

To conclude the proof we note that there are ultimately 4 strategy parameters: t_l and t_k for the searcher, and t_u and c_e for the expert. Equation 2 gives t_l , Equation 3 gives t_u , Equation 4 gives t_k , and c_e is found by optimizing the expert's profit. \square

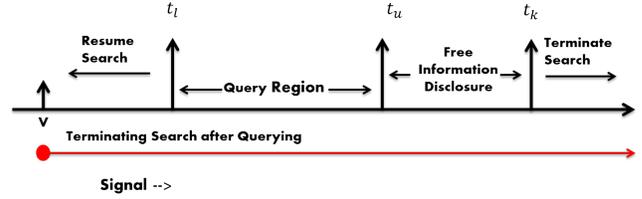


Figure 3: Characterization of the optimal strategy for search with an expert when the expert has the option of disclosing part of the information for free. The searcher queries the expert for a fee if $s \in [t_l, t_u]$ and the expert will disclose the opportunity's true value if $s \in [t_u, t_k]$. In both cases the searcher accepts the offer if its value is greater than the value of resuming the search V , and otherwise resumes search. The searcher rejects and resumes search if $s < t_l$ and accepts and terminates search if $s > t_k$, both without querying the expert.

We note that the above theorem and its proof can be trivially extended for the case where the expert provides a noisy (yet more accurate) signal rather than the true value of the opportunity, using a transformation proposed by MacQueen for the case without the free information disclosure option [21].

Equations 2-4 that characterize the searcher's and the expert's optimal thresholds, can also be derived from their indifference conditions at signals t_l , t_u , and t_k respectively. For example, t_l is the signal at which a searcher is indifferent between either resuming the search or querying the expert, i.e., $V = \int_{y=V}^{\infty} y f_v(y|t_l) dy + V F_v(V|t_l) - c_e$, which transforms into Equation 2; alternatively, t_l can also be interpreted as a point where cost of purchasing the expert's service is equal to the expected increase in utility from consulting the expert when the searcher would otherwise reject and resume search. Similarly, t_u is the signal at which the searcher is indifferent between querying the expert and terminating the search without querying the expert (in case the information is offered for a fee c_e). Finally, t_k is the signal for which the expert is indifferent between providing the information for free and having the searcher terminate its search, i.e., $0 = -d_e + \pi_e(F_v(v|t_k))$, which transforms into Equation 4.

Using the set of Equations 1-7 we can now solve for (t_l, t_u, V, t_k) , and in particular Equation 7 provides us with the resulting expected profit for the platform. Therefore, the expert can solve for the expected-profit-maximizing c_e (e.g., numerically).

3.1 Numerical illustration

We can use the characterization of the equilibrium strategies to solve for the expert's optimal service fee and derive implications for how experts should price their services. Equilibrium in expert-mediated search derives from a complex set of dynamics. Many parameters affect the equilibrium, including the distribution of values, the correlation between signals and values, search frictions and the cost of querying the expert. Uncovering phenomenological properties of the model is therefore difficult and restricted using a static analysis. Instead, we turn to an illustrative model that uses a particular, plausible distribution of signals and values. For this purpose we adopt the setting used in Chhabra et al [5]. The setting uses the signal as an upper bound on the true value. So the signal could be thought of as the searcher's optimistic estimate upon observing the opportunity (e.g., sellers and dealers offering cars for sale usually make cosmetic improvements to the cars in question, and proceed to advertise them in the most appealing manner possible, hiding defects using temporary fixes; mortgage lenders may advertise their

most appealing features, such as a low introductory rate, while keeping troublesome terms and conditions hidden). Specifically, the signals s are uniformly distributed on $[0, 1]$, and the conditional density of true values is linear on $[0, s]$. Thus

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

Figure 4 depicts the expert’s expected profit with and without free information disclosure as a function of the service fee it sets, c_e . The setting used for the graph takes the searcher’s search cost to be $c_s = 0.17$ and the expert’s production cost $d_e = 0.00019$. Obviously, when $c_e = 0$ the expert makes no profit regardless of whether or not she offers some of the information for free. However as c_e increases, and in particular when $c_e > d_e$ the expert makes profit and, as can be observed from the graph, the option to provide information for free results in a greater expected profit. For larger c_e values ($c_e > 0.028$) the expert becomes too costly and is not being used anymore, i.e., the equilibrium is characterized by $t_l = t_u$.

To get a better understanding of the equilibrium dynamics in the resulting simultaneous game once the expert has set its service fee c_e , in particular the effect of free information disclosure on the equilibrium, we present Figure 5. The figure depicts the searcher’s expected benefit V , the thresholds t_l and t_u and the difference between the two, as a function of the percentage of the interval of signals $(t_u, 1)$ for which information is offered for free, denoted w (i.e., $w = (t_k - t_u)/(1 - t_u)$). We show w on the horizontal axis rather than t_k because an increase in t_k per-se has no actual meaning, as it does not say anything about the higher threshold nor the range of signals used by the searcher for using the costly service (t_u). These result from the equilibrium dynamics of the simultaneous game. The use of w as defined above resolves the problem and enforces an equilibrium in which t_k is constrained in terms of a portion of the resulting $(t_u, 1)$ interval. One possible interpretation for w is therefore the extent to which the expert is willing to provide free information in cases where the searcher receives a favorable signal for which the benefit from knowing the true value does not justify paying c_e for it. The setting used for this figure is the same as the one used for Figure 4 ($c_s = 0.17, d_e = 0.00019$), except that here we also fix $c_e = 0.01$, i.e., the expert is not attempting to maximize profits over c_e in that specific market.¹ Figure 5 is complemented by Figure 7, which depicts the expected number of searches (i.e., expected number of opportunities for which a signal was received by the searcher) and the expected number of times the expert was queried by the searcher in the costly-service mode.

As can be seen from Figure 5, the increase in w results in an increase in the searcher’s expected profit. This is expected, as the searcher now receives the true value for free for some of the signals and therefore, since c_e has not changed, it cannot possibly do worse than in the case where the information is always costly. The increase in V results in an increase in t_l and t_u as the searcher will now become indifferent to querying the expert for greater signals. While the increase in t_u is beneficial, from the expert’s point of view, as it increases the interval of signals for which the service is used for a payment, the increase in t_l has the exact opposite effect. Fortunately, since in this example we use a uniform distribution of signals, we can rely on the measure $t_u - t_l$ to determine whether or not the probability the expert will be queried for a fee increased. From the figure we can see that indeed the increase in w results, in this example, in an increase in $t_u - t_l$ and consequently an increase in the chance the expert is used for a payment c_e in

¹This is often the case whenever the expert is operating in parallel markets and needs to set a fixed fee, or cannot distinguish users coming from this market from others.

each search round. Overall, we can see from Figure 7 that the increase in w results both in an increase in the expected number of search rounds and in the expected number of queries made for a fee. The increase in the first measure suggests that the searcher has become more picky. This is interesting especially since with the increase in $t_u - t_l$ and the increase in the portion of $1 - t_u$ for which free information is received the searcher receives/purchases more information overall and seemingly can identify favorable opportunities more easily. Yet, at the same time the improvement in the searcher’s ability to distinguish the favorable opportunities from the non-favorable ones translates to a greater expected benefit from resuming the search process, resulting in a longer search. This also explains the increase in the overall number of paid queries made to the expert. While the increase in this latter measure is beneficial for the expert, it comes with a price—the expert is also experiencing an increase in the overall number of queries she is providing for free. Therefore, supplying information for free for all signals $s > t_u$ is not beneficial and the expert should take into consideration the production cost d_e . Figure 6 shows the expected profit of the expert as a function of w (see Equation 7). Indeed, the expected profit increases as w increases; however using $w = 1$ is not the best response strategy for the expert. The expert should offer the service for free only when the signal is such that the expected benefit from providing it (taking into consideration the chance the true value will indeed turn out to be poor and an additional search round will be initiated and the expected profit from having the searcher resume its search) is greater than the cost of providing the service for free. Formally, this is expressed as $\pi_e F_v(V|s) > d_e$ and depicted in the right graph of Figure 6.

We emphasize that this result (both the expert and the searcher benefitting from the fact that some of the information is offered for free) is limited to the simultaneous game induced after the price is set by the expert. This does not mean that the searcher benefits overall from the expert changing her strategy to provide some information for free—at the end of the day the expert is setting c_e strategically, and it is possible that the searcher does worse overall in the world where the expert has the added flexibility to offer its services for free sometimes. For example, in the setting analyzed above (where $c_s = 0.17$ and $d_e = 0.00019$, with the expert’s expected-profit-maximizing $c_e = 0.01$ (when free information disclosure is allowed) the searcher’s expected benefit is 0.196, whereas when free information disclosure is not allowed the expert uses $c_e = 0.05$ and the searcher’s expected benefit is 0.247.

3.2 Model Robustness

One fear for an expert or a platform when considering switching to offering a service for free is that some parts of the population that were not using the service up until then because of its price, could start using it extensively once it is offered for free, causing a substantial unexpected expense for the expert, who may not previously have been aware of their existence. In this section we illustrate numerically that even with a relatively large population of such “free riders” the expert can still benefit from offering the service for free for some signals. For this purpose we consider two populations of searchers. The first is of searchers characterized by a relatively small search cost, hence with a smaller incentive to use the expert services (as they can potentially repeat their search process until running into an opportunity associated with a very high signal, and choose to terminate the search without ever querying the expert). This population will, however, use the expert’s services whenever offered for free, since this is a dominant strategy when available. The second

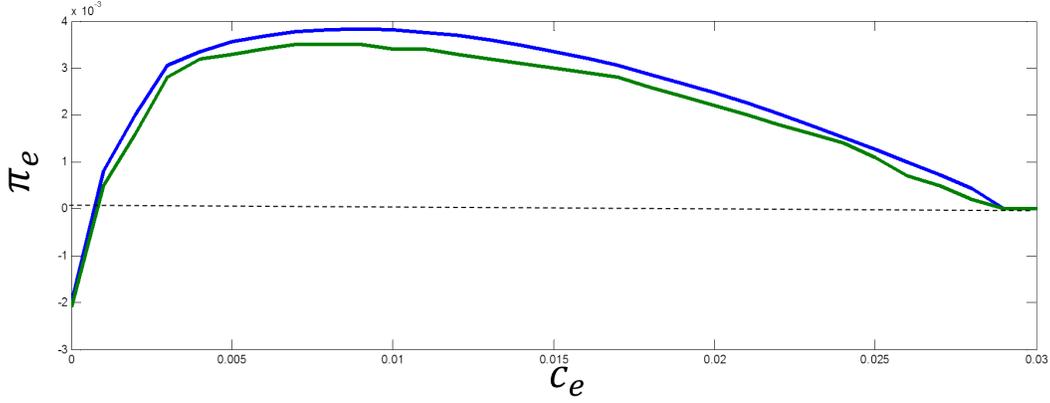


Figure 4: Expert’s expected profit with and without free information disclosure (upper and lower curve, respectively) as a function of c_e for a setting where: $c_s = 0.17$ and $d_e = 1.9 \cdot 10^{-4}$.

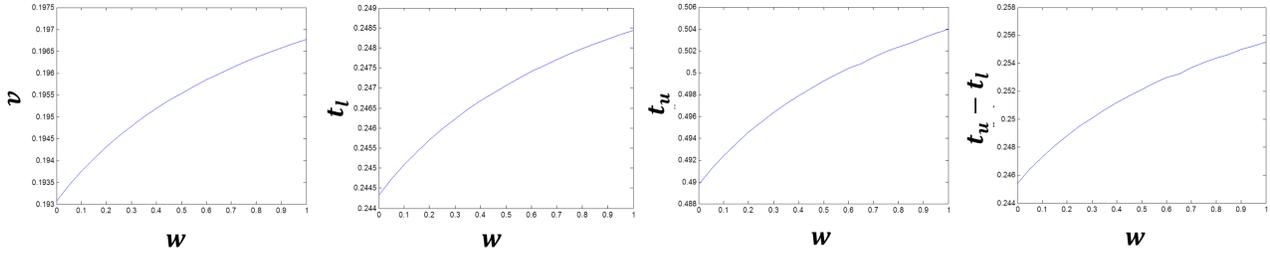


Figure 5: V, t_l, t_u and the size of the interval $[t_l, t_u]$ as a function of the parameter w (the percentage of the interval of signals $(t_u, 1)$ for which information is provided for free). The setting is $c_s = 0.17$, $d_e = 0.00019$ and $c_e = 0.01$.

population is characterized by a higher search cost, and uses the expert’s services for some signals even when offered for a fee $c_e \gg 0$. Both populations receive signals from the same distribution $f_s(y)$ and similarly share the same function $f_v(v|s)$ according to (3.1). The search costs of the two populations are $c_s^l = 0.0292$ for the low search cost population and $c_s^h = 0.17$ for the high search cost population. The expert’s marginal cost for providing the service is $d_e = 1.9 \cdot 10^{-4}$ for both populations.

Based on the parameters above there is no query fee $c_e \geq d_e = 1.9 \cdot 10^{-4}$ that results in the use of the expert’s services by the low search cost searchers (i.e., $t_l = t_u$ for this population). Therefore, the expert maximizes her expected-profit based on the second population only, resulting in the following equilibrium: $t_l = 0.246$, $t_u = 0.496$, $V = 0.0195$ and $c_e = 0.01$. When offering information for free for some of the signals, the expert, who cannot distinguish between searchers of the two populations, needs to take into consideration the loss due to the use of her services by searchers of the low search cost population.

Taking α to be the portion of the high search cost searchers in the general population, the expert’s expected profit is given by

$$(1 - \alpha)(-d_e) \frac{F(t_k) - F(t_u)}{A} + \alpha \left((-d_e) \frac{F(t_k) - F(t_u)}{A} + (c_e - d_e) \frac{F(t_u) - F(t_l)}{A} \right),$$

where A is the probability the search is terminated (calculated according to (5)). The first term corresponds to the loss due to the free

usage of the expert’s services by the low search cost searchers. The second term corresponds to the expected profit from the high search cost searchers and includes both the loss due to free service and the gain from the paid service. Both terms are weighted according to the proportion of the different searchers’ types in the population.

Figure 8 depicts the expert’s expected profit for the setting described above as a function of α , when free information disclosure is allowed and when it is not allowed. The figure demonstrates that, indeed, even for cases where the population of “free-rider” searchers is substantial (99% in this case), the expert can still benefit from free information disclosure.

4. RELATED WORK

Information platforms of the kind we discuss here (often referred to as middlemen or middle agents, brokers and matchmakers [7, 22, 16, 36]) are ubiquitous, especially in distributed multi-agent system environments where immediate reliable information about the different opportunities available to the agents is not public. As such, much recent work has focused on studying the dynamics associated with information search in such platforms [15, 26, 34, 5] and emergent behavior in two-sided markets [1, 29, 33, 14]. One of the main questions investigated within this context is that of how platforms should price their information services, i.e., who pays, and what fees to charge [4, 12, 8, 33]. This paper is among the first to consider a richer space of strategic choices for platforms, such as the option to partially disclose information for free [28]. To date, work that considers providing information for free has been limited to providing the information completely free to some users. For example, it

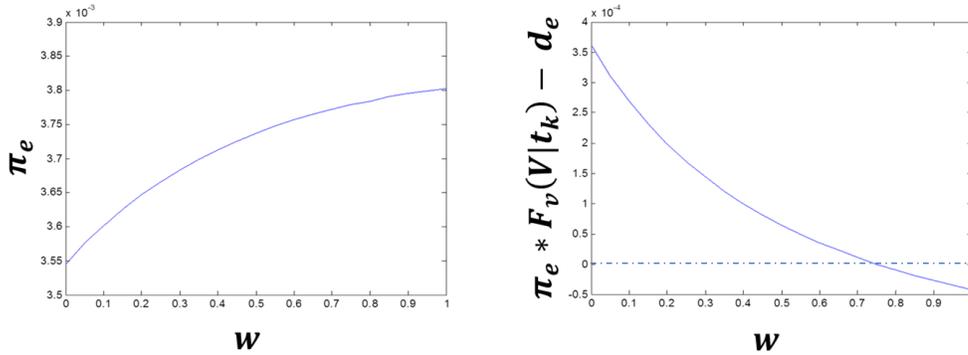


Figure 6: (a) The expert’s expected profit from having the searcher resume search as a function of the parameter w ; (b) The expected net benefit from providing the true value for free when the signal is t_k . The setting is identical to the one used for Figure 5. The expected net benefit in this example becomes zero for $w = 0.75$.

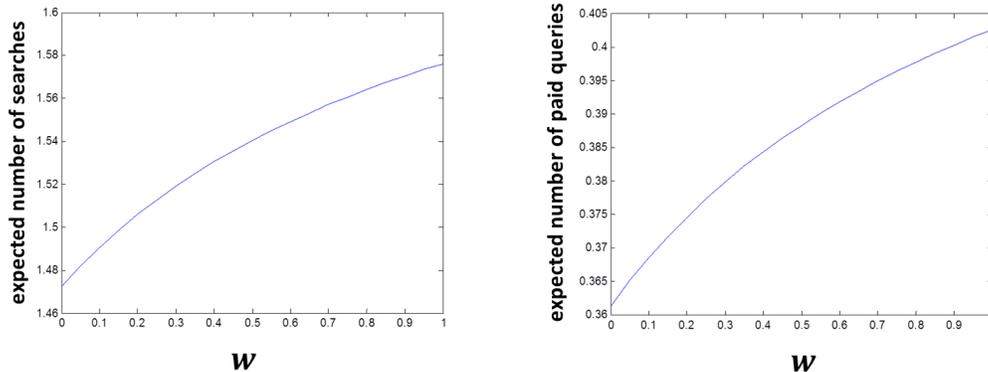


Figure 7: The expected number of search rounds carried out by the searcher (left) and the expected number of paid queries made by the searcher to the expert (right), as a function of w . The setting is identical to the one used for Figure 5.

has been suggested that platforms could charge only one side in a two-sided market while the other group is allowed to use the platform for free [3]. These models are also different from ours in the motivation for free service provision. Typically, the motivation in these models is intense competition among the players of one group (e.g., directories such as “yellow pages” that are supplied to readers for free) [1] or how platforms can attract elastic consumers and, as a result, obtain higher prices or more participation on the other side [29]. Our work analyzes partial free disclosure of information at the single user level, with the potential benefit that it may induce further consumption of the paid service.

Much recent work has been dedicated to applying search-theoretic principles in novel domains, e.g., in comparison shopping [31, 17]. The assumption in this line of work is that the provider’s sole purpose is to serve the user’s needs [23]. This assumption leads to the design or modeling of information providers which favor the user (e.g., buyers, in comparison shopping applications) [11, 27]. Existing work where information providers are modeled as self-interested autonomous entities [18, 19] 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 information [32]. In contrast, our work deals with an informa-

tion provider that is interested in maximizing its expected revenue from the process. Finally, there is a rich literature on variations of the secretary problem [9], a classical optimal-stopping online problem. Our setting is different in that it involves search costs rather than a limited list of possibilities, and the goal is to maximize expected utility rather than the probability of hiring the best candidate (for more on these differences and models that share some features of both types of problems, see Gilbert and Mosteller [10] and Das and Tsitsiklis [6]).

To the best of our knowledge, none of the one-sided search literature in either search theory or multi-agent systems has considered the market dynamics that result in cases where a self-interested expert can sometimes choose to disclose her information for free.

5. CONCLUSIONS

Our main contribution is to analyze a subtle strategic complexity (free information disclosure) in a common multi-agent environment (one-sided search with a self-interested information provider or platform). We find that allowing the information provider to choose to disclose her information or provide her services for free can be beneficial. The channel of operation is complex: when the expert sometimes provides her services for free, she changes the searchers’

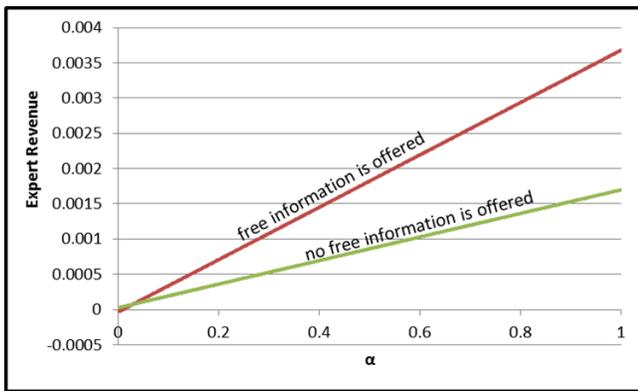


Figure 8: The expert’s expected benefit when free information disclosure is allowed and when it is not allowed, as a function of the percentage of the high search cost searchers in the general population, for the example described in the text.

optimal strategies, expanding the range at which users choose to use her non-free services.

One natural fear in using free disclosure strategies would be model robustness – suppose the expected higher profits were driven by a misestimation of the population? For example, it could be that only those with high search costs were using expert services earlier, so the expert assumes the population in general has high search costs – however, by offering her services for free, she suddenly draws out the population with low search costs that she was unaware of previously since they never used her services. We show that our result is robust to even a significant proportion of such “free riders” in the searching population. As such, the idea of free information disclosure could have significant practical value in search-based markets and systems. We note that the information-provider in our model is working within a somewhat restricted strategy space and could have incorporated different prices (including zero, i.e. free) for each signal. Yet, one of our major results is that, even with the restricted strategy space, there is a benefit to the information provider of providing some services for free. We also note that it is important that the expert will also observe the signal s , as otherwise the searcher could lie about the signal and always get the service for free. In various domains the signal can be verified (e.g., in the used car domain, the expert (e.g., mechanic) can verify the signal by checking the Carfax report herself). The non-verifiable signals domain is interesting for future work.

Other important avenues for future work include analysis of information provision and the incentives for free reporting in two-sided search markets (for example, matchmakers in a dating service), and analysis of searchers’ incentives for truthfulness. Here, we assume that searchers truthfully disclose their signals, which makes sense in verifiable settings like presenting a Carfax report to a mechanic; however, in more subjective settings like dating or travel preferences, how can the expert guarantee that the user is revealing his signal truthfully?

Acknowledgments

This work was supported in part by ISF grant 1083/13 and in part by an NSF CAREER award (1414452).

REFERENCES

[1] M. Armstrong. Competition in two-sided markets. *RAND Journal of Economics*, 37(3):668–691, 2006.

[2] J. Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Management science*, 43(12):1676–1692, 1997.

[3] B. Caillaud and B. Jullien. Chicken & egg: Competition among intermediation service providers. *RAND Journal of Economics*, 34:309–328, 2003.

[4] H. K. Cheng and G. J. Koehler. Optimal pricing policies of web-enabled application services. *Decision Support Systems*, 35(3):259–272, 2003.

[5] M. Chhabra, S. Das, and D. Sarne. Expert-mediated search. In *Proc. International Conference on Autonomous Agents and Multiagent Systems*, pages 415–422, 2011.

[6] S. Das and J. N. Tsitsiklis. When is it important to know you’ve been rejected? A search problem with probabilistic appearance of offers. *Journal of Economic Behavior and Organization*, 74:104–122, 2010.

[7] K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proc. International Joint Conference on Artificial Intelligence*, pages 578–583, 1997.

[8] D. S. Evans and R. Schmalensee. The industrial organization of markets with two-sided platforms. Technical report, National Bureau of Economic Research, 2005.

[9] T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.

[10] J. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73, 1966.

[11] M. F. Gorman, W. 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.

[12] A. Hagiu. Two-sided platforms: Pricing and social efficiency. Technical report, Princeton University, 2004.

[13] A. Hagiu and R. S. Lee. Exclusivity and control. *Journal of Economics & Management Strategy*, 20(3):679–708, 2011.

[14] A. Hagiu and J. Wright. *Multi-sided platforms*. Harvard Business School, 2011.

[15] C. Hajaj and D. Sarne. Strategic information platforms: Selective disclosure and the price of free. In *Proceedings of the ACM Conference on Economics and Computation*, pages 839–856. ACM, 2014.

[16] M. He, N. R. Jennings, and H. Leung. On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 15(4):985–1003, 2003.

[17] M. C. Janssen and J. L. Moraga-González. Strategic pricing, consumer search and the number of firms. *The Review of Economic Studies*, 71(4):1089–1118, 2004.

[18] J. O. Kephart and A. R. Greenwald. Shopbot economics. In *Game Theory and Decision Theory in Agent-Based Systems*, pages 119–158. Springer, 2002.

[19] J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, 2000.

[20] K. P. Lai, S. C. Chong, H. B. Ismail, and D. Y. K. Tong. An explorative study of shopper-based salient e-servicescape attributes: A Means-End Chain approach. *International Journal of Information Management*, 34(4):517–532, 2014.

[21] J. B. MacQueen. Optimal policies for a class of search and evaluation problems. *Management Science*, 10(4):746–759, 1964.

[22] P. Maes, R. H. Guttman, and A. G. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.

- [23] P. M. Markopoulos and L. H. Ungar. Shopbots and pricebots in electronic service markets. In *Game Theory and Decision Theory in Agent-Based Systems*, pages 177–195. Springer, 2002.
- [24] E. Maskin and J. Riley. Monopoly with incomplete information. *RAND Journal of Economics*, 15:171–196, 1984.
- [25] P. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12:380–391, 1981.
- [26] Y. Nahum, D. Sarne, S. Das, and O. Shehory. Two-sided search with experts. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 754–771. ACM, 2012.
- [27] M. Nermuth, G. Pasini, P. Pin, and S. Weidenholzer. The informational divide. *Games and Economic Behavior*, 78:21–30, 2013.
- [28] N. Peled, Y. Gal, and S. Kraus. A study of computational and human strategies in revelation games. *Autonomous Agents and Multi-Agent Systems*, 29(1):73–97, 2015.
- [29] M. Rysman. The economics of two-sided markets. *Journal of Economic Perspectives*, 23(3):125–43, 2009.
- [30] D. Sarne and S. Kraus. The search for coalition formation in costly environments. In *Cooperative Information Agents VII*, pages 117–136. Springer, 2003.
- [31] R. Waldeck. Search and price competition. *Journal of Economic Behavior & Organization*, 66(2):347–357, 2008.
- [32] Y. Wan and G. Peng. What’s next for shopbots? *IEEE Computer*, 43(5):20–26, 2010.
- [33] E. G. Weyl. A price theory of multi-sided platforms. *The American Economic Review*, pages 1642–1672, 2010.
- [34] 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.
- [35] R. Wilson. *Nonlinear pricing*. Oxford University Press, 1993.
- [36] M. Wooldridge. *An Introduction to Multiagent Systems*. Wiley, 2009.