

Competitive Information Provision in Sequential Search Markets

Meenal Chhabra
Virginia Tech
Blacksburg, VA, USA
meenal@vt.edu

Sanmay Das
Washington University in St.
Louis
St. Louis, MO, USA
sanmay@seas.wustl.edu

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

ABSTRACT

We study competitive information provision in search markets. Consider the used car market: as a consumer searches, she receives noisy signals of the values of cars. She can consult an expert (say Carfax or a mechanic) to find out more about the true value before deciding whether to purchase a particular car or keep searching. Prior research has studied the pricing problem faced by a monopolistic expert who provides searchers with perfect information. Here, we study a richer model that augments prior work in two important respects. First, we analyze expert duopolies; thus each expert must now reason about the influence of her strategy on the other. Second, we consider experts who provide uncertain information, with higher quality experts providing more certainty; experts can compete on both price and quality. We show that, in equilibrium, prices will be set such that the searcher consults the worse quality expert for low or high signals, and the higher quality expert for intermediate signals. Surprisingly, we find cases where an improvement in the quality of the higher-quality expert (holding everything else constant) can be pareto-improving: not only does that expert's profit increase, so does the other expert's profit and the searcher's utility.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

Keywords

Economics of information; Sequential search; Duopoly

1. INTRODUCTION

Consider a consumer looking to buy a used car. She will examine cars ads sequentially until she finds a car she likes and buy it. When she looks at an advertisement, she may not interpret the quality of the car correctly, as the seller may not reveal the true condition of the car. The consumer may choose to consult an agency like Carfax, which provides information about the car's history. Prior research has modeled the impact of the presence of monopolistic information brokers or experts like Carfax on search markets

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

like `autotrader.com` which serve as platforms for such consumer search [3].

While it is convenient to model them as perfectly informative, in reality, experts' signals are themselves likely to be noisy. MacQueen has studied the searcher's problem when the information provided by the expert is noisy [17]; however, MacQueen's expert is not strategic, or even an optimizer. In contrast, Chhabra et. al.'s expert is an optimizing monopolist, but she has access to perfect information [3, 4]. The noise in the expert's signal is a direct measure of the quality of the expert: the less noisy, the higher quality. Recognizing that quality may be determined exogenously or it might be a choice made by the experts themselves, in this paper we model the dynamics of competition between experts who can compete on both price and quality. For example, the used car information services Carfax and AutoCheck compete in both the price space (as of this writing, AutoCheck charges \$29.99 for a single report and Carfax charges \$39.99) and anecdotally, at least, are of different qualities. Our main contribution is to develop a model that allows us to study competition between experts of different quality and characterize optimal search and market dynamics at equilibrium in a duopoly setting.

Related Work.

Our work relates to the literature on noisy one-sided search, the literature on third-party certifiers of quality, and the literature on equilibrium analysis of firms providing differentiated quality goods. In the one-sided search literature, we extend the model of Chhabra et. al., which considers noisy search with a perfectly informed expert and derives a double threshold strategy [3].

Costly search [13] and in particular sequential costly search [22, 15, 14] is a prevalent theme in MAS. It is of great importance whenever there is no central source that can supply an agent full immediate reliable information on the environment thus an agent needs to consume some of its resources obtaining this information [11]. Most sequential search problems assume that the true value of the opportunity is observed and the optimal solution is usually in the form of a reservation-strategy [23, 19, 18, 24]. Work on multi-attribute sequential search must typically incorporate noise into the models [16, 25]. Most closely aligned with our work is that of MacQueen, whose basic set-up is used here [17]. However, in all the work on noisy search (with the exception of Chhabra et al), the focus is solely on optimal search, when the quality and price of the expert is exogenous. In contrast, we use the optimal search strategy to

analyze the market for experts. We model competitive markets and derive strategies for both searchers and experts.

The literature on certification is relevant because the expert can be thought of as a third party certifier providing quality ratings for a product. Dranove and Jin summarize the literature on third party certifiers [7]. Most prior work either focuses on the seller’s incentives to disclose complete information versus hiding their quality, or on the strategic behavior of certifiers [2]. In this paper, the incentives for experts to acquire more information are dependent on the costs of such information acquisition, but there is no value to hiding information once acquired. We also study the strategic issues arising from the competition dynamics of experts with different levels of information. The major difference between a search setting such as ours and the literature on ratings (for example, credit ratings) is that a central theme of the ratings world is that organizations (or sellers) pay certifiers in order to get rated (the issuer-pays model). So competition can create completely different dynamics than in a search setting; these certifiers are struggling to stay in the market so they may have incentives to provide generous ratings [9]. The business model is different for our expert, since she provides services to buyers directly. She may have to make a decision on the quality of information to acquire in order to compete, but there are no incentives for misrepresentation.

Another line of related research is on equilibrium analysis in duopoly when firms produce similar goods with different qualities. In this paper, experts compete on price and quality, so we focus on literature geared towards finding Bertrand solutions. Motta studies a two stage game in which the two participating firms first choose their quality and then select price or quantity when the marginal cost varies with quality in order to compare Cournot and Bertrand equilibrium outputs [21]. Crampes and Hollander study how a duopoly equilibrium can be perturbed if there is a minimum quality requirement [5]. Economides studies oligopolistic competition with an infinite number of firms who can select quality from a given range, and he studies equilibrium characteristics with both fixed and variable costs of production [8]. However, in all of these papers, the model used are simple enough to be solved analytically. As far as we are aware of, we are the first to study competition in information provision in search markets.

Finally, we note that there is also related work on providing information to people in adversarial settings [1, 10]; this literature tends to focus, however, on factors related to irrationality and the computationally bounded nature of human reasoning.

Contributions.

- We introduce a new model of competitive information provision in search markets. Within this model, we characterize the searcher’s strategy when there are two competitive experts with different levels of information quality, and show that, under certain conditions, the optimal strategy of the searcher with imperfect experts in a duopoly is reservation-based: there are lower and upper thresholds between which an expert is consulted, as in the monopolistic expert case. Now, the region between the lower and upper thresholds is itself divided into three, with the higher quality expert being con-

sulted in the middle of the region and the lower quality expert in the upper and lower parts of that region.

- We show how to compute the optimal price for experts as a function of information quality and cost, both when information quality is controlled by the expert and when it is exogenous. We use this to characterize equilibrium in the duopoly setting.
- We examine in detail the dynamics of competition between a higher quality expert and a lower quality expert (where quality is characterized by the variance of white noise that affects an expert’s signal), and the effects of competition on the profits of the experts and the welfare of searchers. Surprisingly, we find cases where an improvement in the quality of the higher quality expert (holding everything else constant) can be pareto-improving: not only does that expert’s profit increase, so does the other expert’s profit and the searcher’s utility.

2. THE MODEL

We consider a one-sided noisy search environment where a searcher observes a stream of opportunities sequentially, from which she eventually needs to choose one. The value v of each opportunity is a priori unknown, however the searcher receives, upon paying a cost c_s , a noisy signal s which is correlated with the true value v according to a known probability density function $f_s(s|v)$. Let $f_v(x)$ ($x \in \mathbb{R}$) denote the (stationary) probability density function for the true values (we assume the searcher knows $f_v(x)$). Even with the signal, the true value of the opportunity remains unknown and revealed only if it is accepted. The searcher thus has to decide whether to accept the opportunity and terminate the search or reject and pay to receive the signal for the next opportunity. The searcher also has a third option. She can consult an expert for some fee in order to get more information about the opportunity. An expert provides a noisy signal s_e . The conditional density function of s_e given the true value v is given by $f_e(s_e|v)$. The searcher’s signal and the expert’s signal are assumed to be conditionally independent given the true value, i.e., for each expert $f(s, s_e|v) = f_s(s|v)f_e(s_e|v)$; thus one can easily calculate the joint density $f(s, s_e, v)$ and other conditional densities using Bayes’ rule. The searcher is a risk-neutral utility maximizer.

The model above is standard in the literature and various motivations and justifications for the above assumptions can be found in the literature cited in this paper. In particular the above model is similar to the one proposed by MacQueen [17] and used as a basis by many others, mostly under the assumption that expert signal is certain [3, 16].

There are two competing experts the searcher can turn to. Each expert offers a different signal quality, captured by its unique conditional density $f_e(s_e|v)$, and the searcher, who is aware of the qualities of the two experts, can only use one of them. Specifically, we assume that signals are corrupted by zero-mean noise, ϵ , therefore $s_e = v + \epsilon$, (i.e., the mean of $f_e(s_e|v)$ is v). The variance σ_e^2 of the distribution is thus a measure for the expert’s signal quality — the greater the variance, the lower is the quality and vice versa. Signals received from a zero variance expert represent perfect information.

The two experts are self-interested, risk-neutral, utility maximizers and act strategically (the expected utility of an

expert is the expected payments received from searchers minus the cost of producing its signals in response to queries). We consider both exogenous and endogenous quality. In the first case each expert needs to set merely its fee, denoted by c_{e_i} for Expert $i \in \{1, 2\}$ while in the second case the decision involves both the quality and the fee to be charged. We assume the marginal cost of provision of services for each expert, denoted d_{e_i} , may vary for different qualities of information, and therefore represent it as a function of its variance $\sigma_{e_i}^2$ respectively for Expert $i \in \{1, 2\}$.

We assume that the searcher can only query once per signal, so if she chooses to consult an expert she also has to decide which expert to select. This restriction, where only one expert can be queried, is applicable whenever the nature of the opportunity, or constraints imposed by the provider of the opportunity, preclude obtaining a second opinion. For example, a car owner might not be willing to let a potential buyer take his car to be inspected by more than one mechanic.

3. OPTIMAL POLICIES

In this section, we derive the optimal strategy for the searcher and then use it to characterize the strategies for both experts. We first summarize the searcher's strategy in case of a monopoly, drawing from existing literature. We then turn to deriving the form of the optimal strategy when multiple experts are available. Using the searcher's optimal strategy, we can calculate the expected demand for each expert's services, and thus their profit. This enables us to characterize equilibrium conditions and to study the phenomenological properties of the market in equilibrium.

3.1 The searcher's strategy

3.1.1 Monopoly

If signals and values are not meaningfully correlated, the optimal strategy for the searcher is hard to characterize. For any arbitrary conditional distribution of signals and values, there is no guarantee that the optimal strategy is reservation value based. However, it is often the case that if we assume some natural relationship between values and signals, the optimal strategy can be shown to have a simple structure. In our case, we assume the *monotone likelihood ratio property*.

DEFINITION 1. Monotone likelihood ratio property (MLRP): A distribution $f(y|s)$ satisfies MLRP if the ratio $\frac{f(y|s')}{f(y|s)}$ is non-decreasing in y for $s' > s$. This also implies that $f(y|s')$ first-order, second-order and third-order stochastically dominates $f(y|s)$ for $s' > s$.

The MLRP implies that a searcher is more likely to get a higher value if she sees a higher signal than if she sees a lower signal [20, 26]. If both the conditional distribution of the true value v given the searcher's signal, s , ($f_v(v|s)$) and the conditional distribution of the true value v given the expert's signal s_e ($f_v(v|s_e)$) satisfy the MLRP as given in Definition 1 then the optimal policy is reservation based [17]. The optimal search strategy in this case can be described by the tuple (V, t_l, t_u) , 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 current opportunity should be accepted and the search should be terminated; and (c) V is the expected utility of the searcher; the searcher should query the

expert if the signal s , she observes, lies in between t_l and t_u ; On receiving the signal, s_e , from the expert, the searcher should terminate the search if $\mathbb{E}(v|s, s_e) \geq V$ otherwise resume search. The values t_l , t_u , and V can be calculated using Equations 1-3:

$$V = -c_s + VF_s(t_l) + \int_{t_u}^{\infty} f_s(s)\mathbb{E}(v|s) ds + \int_{t_l}^{t_u} f_s(s) \left(VF_Z(V|s) + \int_V^{\infty} z f_z(z|s) dz - c_e \right) ds \quad (1)$$

$$V = VF_Z(V|t_l) + \int_V^{\infty} z f_z(z|t_l) dz - c_e \quad (2)$$

$$\mathbb{E}(v|t_u) = VF_Z(V|t_u) + \int_V^{\infty} z f_z(z|t_u) dz - c_e \quad (3)$$

where $Z = \mathbb{E}_v(v|s, s_e)$ and the density function of Z conditional on the searcher's signal s is calculated as: $F_Z(z|s) = \Pr(\mathbb{E}_v(v|s, s_e) < z|s)$.

3.1.2 Duopoly

We have two competing experts E_1 and E_2 , each providing a different degree of expertise such that their conditional distributions of the signals given the values are $f_{e_1}(s_{e_1}|v)$ and $f_{e_2}(s_{e_2}|v)$ respectively, and they charge c_{e_1} and c_{e_2} per query respectively. Without loss of generality, let E_1 be the lower quality (higher variance) expert and E_2 be the higher quality (lower variance) expert. For convenience, we refer to E_1 and E_2 as E_{low} and E_{high} , and their corresponding signals as $s_{e_{low}}$ and $s_{e_{high}}$, respectively. If the searcher decides to query and get more information about the current opportunity, she also has to decide which expert to query. Theorem 1 characterizes the structure of the searcher's optimal strategy for the duopoly problem under the MLRP.

THEOREM 1. *If the conditional distribution of the true value v given the searcher's signal, s , ($f_v(v|s)$) and the conditional distribution of the true value v given either expert's signal $s_{e_{low}}$ or $s_{e_{high}}$ (i.e., $f_v(v|s_{e_{low}})$ and $f_v(v|s_{e_{high}})$) satisfy MLRP, the searcher's signal s and the experts' signal $s_{e_{low}}$ and $s_{e_{high}}$ are conditionally independent given the true value v , and the conditional distribution of Z given s is differentiable, then the optimal search strategy can be described by the tuple (V, t_l, t_u) , 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 current opportunity should be accepted and the search should be terminated; (c) V is the expected utility of the searcher; if the signal s , she observes, lies in between t_l and t_u then the searcher should query one of the experts depending on whichever yields higher expected value from querying. The expected value of querying the experts E_{low} and E_{high} is given by Equations 4 and 5 respectively given the searcher's signal s :*

$$U_{low}(s) = -c_{e_{low}} + VF_{z_{low}}(V|s) + \int_V^{\infty} z f_{z_{low}}(z|s) dz \quad (4)$$

$$U_{high}(s) = -c_{e_{high}} + VF_{z_{high}}(V|s) + \int_V^{\infty} y f_{z_{high}}(z|s) dz \quad (5)$$

On receiving the signal s_e , from the expert she chooses, she should terminate search if $\mathbb{E}(v|s, s_e) \geq V$ otherwise resume search. The values t_l , t_u , and V can be calculated using

Equations 6-8:

$$V = -c_s + V F_s(t_l) + \int_{t_u}^{\infty} \mathbb{E}(v|s) f_s(s) ds + \int_{t_l}^{t_u} \max(U_{high}(s), U_{low}(s)) ds \quad (6)$$

$$V = \max(U_{high}(t_l), U_{low}(t_l)) \quad (7)$$

$$\mathbb{E}(v|t_u) = \max(U_{high}(t_u), U_{low}(t_u)) \quad (8)$$

Proof sketch: This is a simple corollary to MacQueen's result showing the optimality of a double reservation strategy for the searcher when there is only one expert under a stochastic dominance assumption on the distribution $F_Z(v|s)$ [17]. Since the MLRP implies stochastic dominance, MacQueen's proof can be extended by replacing the terms related to the utility of querying an expert in a monopolistic world with the maximum of the utilities that can be achieved by querying each expert separately. \square

Equation 6 is the main Bellman equation which summarizes the searcher's expected utility from continuing search, in which case she incurs a cost c_s for obtaining the signal of one additional opportunity. If the signal obtained is below t_l (with probability $F_s(t_l)$) then the search resumes, yielding expected utility V . Otherwise, if the signal obtained is above t_u , the search terminates and the agent's utility is $\mathbb{E}(v|s)$. Finally, if the signal is between t_l and t_u , the expert who provides higher expected utility to the searcher is used. Equations 7 and 8 can be obtained by setting the first derivative of Equation 6 w.r.t t_l and t_u respectively equal to 0.; they can also be seen as indifference conditions at signal t_l , and t_u ; for example, Equation 7 represents the fact that when the signal value is equal to t_l , the expert is indifferent between querying an expert (whichever one would be more useful given signal value t_l), and rejecting the opportunity, thus resuming search. There is also a degenerate case when neither expert is being queried. In that case, a single threshold serves as an optimal strategy for the searcher.

What happens between t_l and t_u ? Although, Theorem 1 tells us that there is a region (between t_l and t_u), where one of the experts will be queried, it is in general non-trivial to determine how the searcher's strategy behaves in $[t_l, t_u]$. However, for one of the most common and well-known value-signal structures in the literature [6, 12], we can show that it is optimal for the searcher to use a double reservation-value strategy to partition the region $[t_l, t_u]$ itself.

THEOREM 2. *Assume the true value is normally distributed ($v \sim \mathcal{N}(\mu, \sigma_v^2)$) and signals are corrupted by additive white Gaussian noise (AWGN), i.e., $s = v + \epsilon_s$, $s_{e_{high}} = v + \epsilon_{e_{high}}$ and $s_{e_{low}} = v + \epsilon_{e_{low}}$ where ϵ_s , $\epsilon_{e_{high}}$ and $\epsilon_{e_{low}}$ represent independent draws from zero-mean Gaussians. Then the optimal strategy of the searcher can be characterized by a tuple (V, t_1, t_2, t_u) such that a rational searcher should: (1) reject all the signals $s < t_1$; (2) accept all the signals $s > t_u$ without querying the expert; (3) query the low quality expert if $t_1 \leq s \leq t_2$ and $t_2 \leq s \leq t_u$. On consulting the expert, if $Z_{low} = \mathbb{E}(v|s, s_{e_{low}}) \geq V$, then accept the opportunity and terminate the search, otherwise resume; (4) query the high quality expert if the signal s lies in between t_1 and t_2 . On consulting the expert if $Z_{high} = \mathbb{E}(v|s, s_{e_{high}}) \geq V$, then accept the opportunity and terminate the search, otherwise resume;*

PROOF. Since $v \sim \mathcal{N}(\mu_v, \sigma_v^2)$, $f_{e_{high}}(s_e|v) \sim \mathcal{N}(v, \sigma_{e_{high}}^2)$ and $f_{e_{low}}(s_e|v) \sim \mathcal{N}(v, \sigma_{e_{low}}^2)$, the following holds:

$$f_{Z_{high}}(z|s) \sim \mathcal{N}\left(\frac{s\sigma_v^2 + \mu\sigma_s^2}{\sigma_v^2 + \sigma_s^2}, \frac{\sigma_s^4\sigma_v^4(\sigma_v^2 + \sigma_s^2)^{-1}}{\sigma_s^2\sigma_v^2 + \sigma_s^2\sigma_{e_{high}}^2 + \sigma_v^2\sigma_{e_{high}}^2}\right)$$

$$f_{Z_{low}}(z|s) \sim \mathcal{N}\left(\frac{s\sigma_v^2 + \mu\sigma_s^2}{\sigma_v^2 + \sigma_s^2}, \frac{\sigma_s^4\sigma_v^4(\sigma_v^2 + \sigma_s^2)^{-1}}{\sigma_s^2\sigma_v^2 + \sigma_s^2\sigma_{e_{low}}^2 + \sigma_v^2\sigma_{e_{low}}^2}\right)$$

This enables the calculation of $F_{Z_{high}}(z|s)$ and $F_{Z_{low}}(z|s)$. Since the mean is equal for the two conditional distributions above, and is expressed as a function of the variable s , it can be represented as $\mu(s)$. Let $\sigma_{Z_{high}}$ and $\sigma_{Z_{low}}$ represent the respective standard deviations. The difference in the utility from querying either expert is:

$$U_{low}(s) - U_{high}(s) = -c_{e_{low}} + c_{e_{high}} + \int_V^{\infty} (\Phi(z, \mu(s), \sigma_{Z_{high}}^2) - \Phi(z, \mu(s), \sigma_{Z_{low}}^2)) dz$$

where $\Phi(x; \mu, \sigma^2)$ represents the normal CDF with mean μ and variance σ^2 evaluated at x . We want to find the region in which each expert is queried. In order to do so, we check how many times the sign of the difference of utilities received by querying the two experts changes. Note, that we only need to consider cases where $c_{e_{high}} - c_{e_{low}} > 0$, as otherwise the strategy is trivial — consult the high quality expert. We claim that $U_{low}(s) - U_{high}(s)$ is single peaked function, and therefore can change sign twice (have at most two roots). In order to do so we first find the points at which the first derivative is 0, and then perform higher order derivative test.

$$\frac{d(U_{low}(s) - U_{high}(s))}{ds} = \frac{d(U_{low}(s))}{ds} - \frac{d(U_{high}(s))}{ds} = -\frac{\sigma_v^2}{\sigma_v^2 + \sigma_s^2} \left(\Phi(V; \mu(s), \sigma_{Z_{low}}^2) - \Phi(V; \mu(s), \sigma_{Z_{high}}^2) \right)$$

Putting the derivative of $U_{low}(s) - U_{high}(s)$ equal to zero, we get: $\Phi(V; \mu(s), \sigma_{Z_{low}}^2) = \Phi(V; \mu(s), \sigma_{Z_{high}}^2)$. The CDF for two normal distribution with same mean and different variance are equal at the extremes $\{-\infty, \infty\}$ or at the mean because CDF of any normal distribution at the mean is 0.5. We know V is finite, therefore, $\mu(s) = V \Rightarrow s = \frac{V(\sigma_v^2 + \sigma_s^2)}{\sigma_v^2}$. To check if this is an extremum, we calculate the second derivative of $U_{low}(s) - U_{high}(s)$:

$$\frac{d^2(U_{low}(s) - U_{high}(s))}{ds^2} = \frac{d}{ds} \left(\frac{d(U_{low}(s))}{ds} - \frac{d(U_{high}(s))}{ds} \right) = \left(\frac{\sigma_v^2}{\sigma_v^2 + \sigma_s^2} \right)^2 \left(\phi(V; \mu(s), \sigma_{Z_{low}}^2) - \phi(V; \mu(s), \sigma_{Z_{high}}^2) \right)$$

$$\frac{d^2(U_{low}(s) - U_{high}(s))}{ds^2} \Big|_{s = \frac{V(\sigma_v^2 + \sigma_s^2)}{\sigma_v^2}} = \left(\frac{\sigma_v^2}{\sigma_v^2 + \sigma_s^2} \right)^2 \frac{1}{\sqrt{2\pi}} \left(\frac{1}{\sigma_{Z_{low}}} - \frac{1}{\sigma_{Z_{high}}} \right) > 0$$

$$\text{as } \sigma_{e_{high}} < \sigma_{e_{low}} \Rightarrow \sigma_{Z_{high}} > \sigma_{Z_{low}} \Rightarrow \frac{1}{\sigma_{Z_{high}}} < \frac{1}{\sigma_{Z_{low}}}$$

where $\phi(x; \mu, \sigma^2)$ represent the normal probability density function with mean μ and variance σ^2 . As the function is single peaked it can at most change sign twice (intersect with x-axis). Note, if the maximum of $(\int_V^{\infty} (\Phi(z; \mu(s), \sigma_{Z_{high}}^2) - \Phi(z; \mu(s), \sigma_{Z_{low}}^2)) dz)$ is less than $c_{e_{low}} - c_{e_{high}}$ then the high quality expert will never be queried. If there are two roots

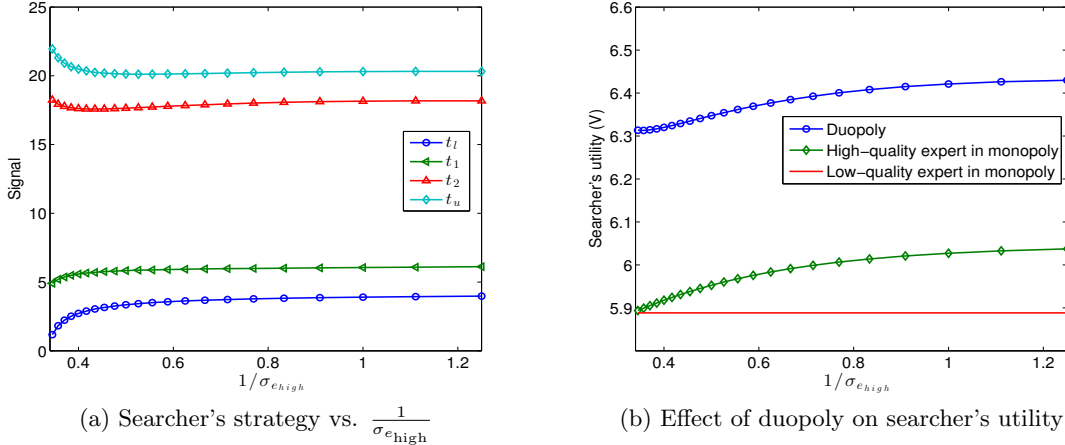


Figure 1: (a) The optimal strategy for the searcher (b) Searcher's utility as a function of the quality of the high-quality expert ($\frac{1}{\sigma_{e_{\text{high}}}}$). For comparison, the green and the red line show the searcher's utility when the high-quality expert and the low-quality expert are the monopolist respectively. For both graphs, we keep the low-quality expert's standard deviation, $\sigma_{e_{\text{low}}}$, constant at 3, $d_{e_{\text{low}}} = \frac{0.01}{\sigma_{e_{\text{low}}}^2}$, $d_{e_{\text{high}}} = \frac{0.01}{\sigma_{e_{\text{high}}}^2}$, $c_s = 0.1$, $\mu_v = 5$, $\sigma_s = 4$, $\sigma_v = 2$.

of $U_{\text{low}}(s) - U_{\text{high}}(s) = 0$, say t_1 and t_2 , then as the point of extremum calculated is a minima (the second derivative is positive and the first derivative is zero) then for signals less than t_1 or greater than t_2 , the low-quality expert will be preferred, and in between the high-quality expert. If $t_1 < t_l$ and $t_2 > t_u$ then the low quality expert is never queried in practice because then resuming the search or accepting without querying are better alternatives for signals $s < t_l$ and $s > t_u$ respectively. \square

3.2 The experts' strategies

While the strategies described above are optimal for the searcher no matter what prices the experts set, the non-degenerate cases only hold when the experts pursue sensible pricing strategies, given their respective qualities. In a duopoly, for both E_{low} and E_{high} to survive, it must be the case that the lower quality expert charges less in equilibrium ($c_{e_{\text{low}}} < c_{e_{\text{high}}}$), otherwise it will never be queried.

Under the white Gaussian noise assumption, the query region of the two experts is nicely partitioned, therefore, we can calculate each expert's profit analytically (similar to the monopolist profit calculation discussed in [3]) and we use that to solve for equilibrium using best-response dynamics. Let A , B , C , D , E and F represent the probability of (1) rejecting and resuming the search; (2) accepting the search without querying the expert; (3) querying the low quality expert and terminating the search; (4) querying the low-quality expert and resuming the search; (5) querying the high-quality expert and terminating the search; (6) querying the high-quality expert and resuming the search respectively.

The expected number of opportunities examined $\eta_s = \frac{1}{B+C+E}$ (because termination is a Bernoulli trial with probability $B+C+E$). The expected number of times an expert is queried, denoted by η_{c_e} , is given by:

$$\eta_{c_e} = \Pr(\text{Expert is queried})\eta_s$$

Therefore, the expected number of times E_{low} and E_{high} are queried is $\eta_{c_{e_{\text{low}}}} = \frac{C+D}{B+C+E}$ and $\eta_{c_{e_{\text{high}}}} = \frac{E+F}{B+C+E}$ respec-

tively. We now turn to a specific example to understand the properties of equilibrium.

3.3 An example

Suppose we have two experts E_{high} and E_{low} such that the conditional distributions of their signals, $s_{e_{\text{high}}}$ and $s_{e_{\text{low}}}$ respectively, given the true value v are both normally distributed with mean v and variance $\sigma_{e_{\text{high}}}^2$ and $\sigma_{e_{\text{low}}}^2$ (i.e., $f_{e_{\text{high}}}(s_e|v) \sim \mathcal{N}(v, \sigma_{e_{\text{high}}}^2)$ and $f_{e_{\text{low}}}(s_e|v) \sim \mathcal{N}(v, \sigma_{e_{\text{low}}}^2)$, and $\sigma_{e_{\text{high}}} < \sigma_{e_{\text{low}}}$). The distribution of the true value $v \sim \mathcal{N}(\mu_v, \sigma_v^2)$. This example satisfies the assumptions for Theorem 2.

3.3.1 Exogenous expert quality

First, we assume that the quality of both experts is fixed and different, but they compete on price. We use a best-response dynamic to find equilibrium prices for the experts to set, as mentioned above.

Searcher strategy and utilities: The searcher's optimal strategy is characterized by a 5-tuple (V, t_l, t_1, t_2, t_u) as per Theorem 2. Figure 1(a) shows the searcher's strategy as a function of the quality of the higher-quality expert ($1/\sigma_{e_{\text{high}}}$), holding the quality of the lower-quality expert ($1/\sigma_{e_{\text{low}}}$) constant. Therefore, the quality of the higher-quality expert increases as we go right. We can see that the lower-quality expert is only utilized for signals that are "close to the edge" in terms of whether or not the searcher wishes to consult an expert at all. The higher quality expert is consulted far more, since it is consulted for all signals that fall in the intermediate range. Figure 1(b) shows the expected utility of the searcher as the quality of the better expert increases; the rate of increase is similar to what would occur if the better expert were a monopolist, and, as expected, the searcher is much better off in a world with competition than a world where the expert is a monopolist. The extent to which this effect holds is interesting. Note that the searcher potentially has access to much better information at the rightmost point of the green line (a monopolist with $1/\sigma_{e_{\text{high}}} \approx 1.25 \Rightarrow \sigma_{e_{\text{high}}} \approx 0.8$), than at the leftmost

point of the blue line (duopolistic experts with $\sigma_{e_{low}} = 3.0$ and $1/\sigma_{e_{high}} \approx 0.35 \Rightarrow \sigma_{e_{high}} \approx 2.85$ respectively). However, the searcher's expected utility is much higher for the leftmost point of the blue line than the rightmost point of the green line, showing that the high expert prices charged in monopoly are much worse for the searcher than having potentially worse information on the basis of which to make a decision.

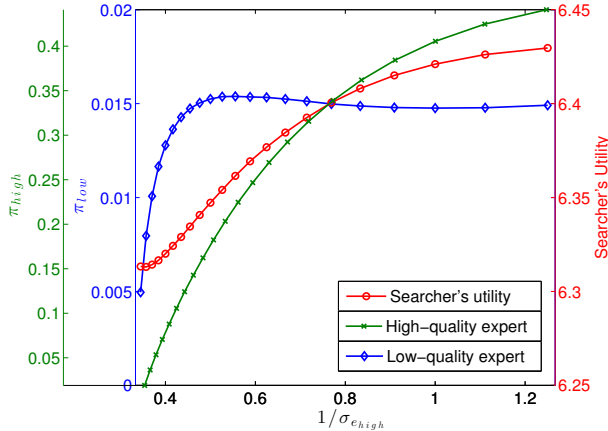


Figure 2: Equilibrium profits of the two experts and the searcher's utility as a function of quality of the high-quality expert ($\frac{1}{\sigma_{e_{high}}}$), keeping the quality of the low-quality expert constant ($\sigma_{e_{low}} = 3$). For $\frac{1}{\sigma_{e_{high}}} < 0.5 \Rightarrow \sigma_{e_{high}} > 2$, we see that the profit of both high-quality and low-quality expert and the utility of the searcher increases as the quality of the high-quality expert increases, therefore, the overall social welfare also increases. The settings are same as in Figure 1.

Expert behavior and profits: Figure 2 shows the profits accruing to each expert (along with the searcher's utility) as a function of the quality of the high-quality expert ($1/\sigma_{e_{high}}$), holding $\sigma_{e_{low}}$ constant ($= 3$). For this figure, the marginal cost of production of expert reports is a function of the quality provided (specifically, $d_e = 0.01/\sigma_e^2$) for both experts. Several interesting things jump out from this picture. First, the profit of the higher quality expert is an order of magnitude higher than the profit of the lower quality expert (the Y axes for the two curves are different, the higher quality expert always makes more profit than the lower quality one). The profit of the higher quality expert increases significantly as her quality increases. Despite the probability of the higher quality expert getting queried in each round ($\Pr(t_1 \leq s \leq t_2)$) decreases the amount it is able to charge increases at a much faster rate (see Figure 3), and hence the profit. The behavior of the profit of low quality expert is more complex — it initially increases rapidly when the difference in the quality of the high-quality expert increases and then starts decreasing very slowly. This initial increase is particularly surprising because one would imagine that enhancing the quality of the better expert would hurt the low quality expert. This effect can be teased out by again examining the probabilities of each expert being consulted and their fees in Figure 3. Interestingly, although the probability of consulting the low quality expert decreases as the

quality of the high quality expert increases, the price she charges also increases. The higher quality expert charges more when its quality is improved in order to offset its increasing marginal cost associated with increasing quality. It is possible that in this situation, the low-quality expert is also able to take advantage of the higher price charged by high-quality expert by increasing its own fee and thus gaining more profit. However, eventually the rate of increase of the query price at equilibrium for the low quality expert decreases, and the declining probability that she is used becomes more important.

It is worth noting that the searcher's utility (the red line in Figure 2) also increases with the increase in the quality of the high-quality expert, mainly because better information is available to the searcher. Clearly, the searcher and the high-quality expert are doing better as the quality of the high-quality expert improves. The fact that the low-quality expert's profit also increases until $\sigma_{e_{high}} > 2$ (or $1/\sigma_{e_{high}} < 0.5$) therefore demonstrates the surprising result that improving the quality of the high-quality expert can be pareto-improving.

3.3.2 Endogenous expert quality


When the expert not only has control over the price it charges but also over the quality it offers, additional equilibrium dynamics must be considered. One option in this case is that the experts have already done enough research on their part and are able to provide any quality and choose the one which is in their best interest, i.e., in order to maximize their expected profit. Alternatively, it is possible that the experts incur some switching cost whenever changing their quality. We analyze equilibrium conditions for these two situations for two experts with identical marginal costs of producing expert reports.

No switching costs.


It is a fairly standard result on differentiated Bertrand competition that the following conditions must hold in equilibrium (cf. [8]):

1. The two experts should make identical profits, otherwise one expert can slightly undercut the other expert in price at the same quality and increase her profit, taking over the entire market share.
2. At equilibrium the price charged by each expert should be equal to the marginal cost, otherwise the experts will tend to undercut each other and capture the complete market. Hence, their profit at equilibrium is zero. However, not all marginal cost pricing strategies are in equilibrium.

In order to find an equilibrium in this case one needs to find a quality (σ_e) (the price will be equal to the marginal cost because of the condition discussed above), such that the best response of the second expert is to charge a price equal to the marginal cost for any quality. For example, if the marginal cost of production is not dependent on the quality and is a fixed identical constant for both experts, then in equilibrium both experts provide perfect information at a price that is equal to their marginal cost (e.g., in the case of informational goods, that are characterized with zero marginal cost of production, the service should be provided for free at equilibrium).



(a) Probability of querying either expert



(b) Prices charged by the two experts

Figure 3: Probability that an expert will be queried (left) and price set by each expert (right) as a function of the quality of the high-quality expert ($\frac{1}{\sigma_{e_{high}}}$). The definitions and settings are the same as in Figure 1

In the presence of switching costs.

Consider two experts with quality and price combinations (σ_{e_1}, c_{e_1}) and (σ_{e_2}, c_{e_2}) respectively such that c_{e_1} and c_{e_2} are equilibrium prices given that qualities are fixed. If the experts are also allowed to change their quality and either of them find a strategy (σ'_e, c'_e) to be more beneficial, taking the strategy of the other expert to be fixed, then the solution no longer complies with the equilibrium stability condition. However, if changing one's quality incurs a cost, which is greater than the benefit achieved by deviating to any new quality (and its corresponding price), the equilibrium conditions will still hold. It is natural to assume zero costs for decreasing quality but positive costs for improving quality, so we model switching costs using a hinge function. In this case, any quality pairs $(\sigma_{e_1}, \sigma_{e_2})$ for which neither of the agents finds a benefit in switching to a lower quality can be stabilized, and thus remain in equilibrium, using an appropriate switching cost (that applies only to increasing quality).

In order to understand how the switching cost affects the nature of equilibria, consider an example where the switching cost depends linearly upon the amount of improvement required to improve the profit, $S = \max(A(\frac{1}{\sigma'_{e_i}} - \frac{1}{\sigma_{e_i}}), 0)$ where $A > 0$ and σ'_{e_i} is the new standard deviation for Expert $i \in \{1, 2\}$ (the higher the standard deviation of the noise, the lower the quality). In order for a pair of standard deviations to be in equilibrium, the proportionality constant A must be sufficiently high to incentivize the two experts to not deviate from the qualities they are offering. Figure 4 shows how the minimum value of the proportionality constant A required to keep the two experts in equilibrium varies as a function of the quality of the high-quality expert ($1/\sigma_{e_{high}}$), keeping the quality of the low quality expert constant. While the figure shows the constant A needed for both experts separately, under the assumption that there is a universal switching cost function for both experts, A would effectively be the max of the two at each point. Note that in this example the proportionality constant for the low-quality expert is always higher than that of the high-quality expert. One reason for this could be that the low-quality expert can gain more by improving her quality. Also, the minimum

value needed for the proportionality constant decreases as the quality of the high-quality expert increases. Therefore, if we choose any value of A (say $A = 1$), then only the points lying on the right of $1/\sigma_{e_{high}} \cong 0.7$ will be in equilibrium.

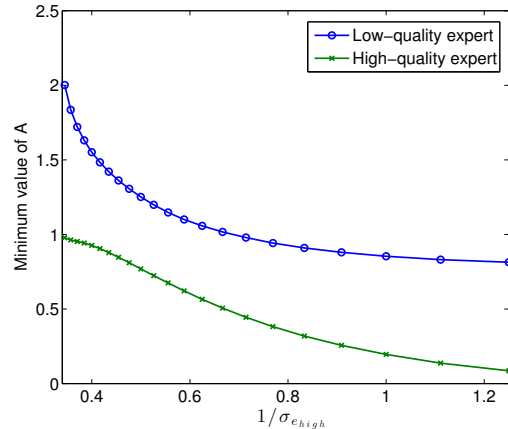


Figure 4: Minimum value of A , the proportionality constant for the switching cost, required for the two experts to be in equilibrium as a function of the quality of the high-quality expert ($\frac{1}{\sigma_{e_{high}}}$) keeping $\sigma_{e_{low}}$ constant. The switching cost $S = \max(A(\frac{1}{\sigma'_{e_i}} - \frac{1}{\sigma_{e_i}}), 0)$ where σ'_{e_i} and σ_{e_i} are the new and the current standard deviations respectively for Expert $i \in \{1, 2\}$. We find that the minimum value of this constant decreases as the quality of the high quality expert increases further, and the value is always higher for the low quality expert. Therefore, if A is equal to any point on the blue curve, then all the points to its right will be in equilibrium but none of those to its left. The settings are same as in Figure 1.

4. DISCUSSION

We have introduced a model for studying competitive information provision in search markets, focusing especially on

the duopolistic case. Our model allows us to study the impact of experts (information providers) of different quality, and strategic issues in how these experts should price their services. In this paper we demonstrate some basic results – e.g, with normally distributed noise for both experts, the higher quality expert will be consulted for “more ambiguous” signals, while the lower quality expert will be consulted for ones that are closer to the searcher’s original decision thresholds. Using those results, we discuss the characteristics of equilibria under both exogenous and endogenous assumptions on expert quality. Surprisingly, we show that improvements in the quality of the higher quality expert can be pareto-improving, increasing not just her profit and the utility of the searcher, but also the profit of the lower quality expert (whose quality does not change). We also characterize the level of switching cost needed in order to “stabilize” the (price, quality) pairs offered by the two experts so that they are in equilibrium.

These results are of importance to MAS designers, platform owners and regulators that often have control of, or can impose constraints on, experts and the quality they offer. For example, if the quality of the expert relies on the amount of information it has or its access to databases that are under the control of the MAS designer, then the latter can dictate the set of expert qualities that maximizes user’s benefit. Alternatively, the higher quality expert can be paid to improve its quality, in a way that benefits all players.

The plausibility and tractability of our model, with normally distributed noise of different variances as the defining features of experts of different qualities, opens up many interesting questions for further study. For example, extending the model to multiple types of consumers/searchers may lead to some interesting insights. Suppose there were some consumers with noisier signals than others, or some who were pickier than others. Would the dynamics of competition lead to market segmentation where different qualities of information were provided to different types of searchers? We believe our model can provide a useful foundation for studying such questions.

5. ACKNOWLEDGMENTS

This work was supported by a US-Israel Binational Science Foundation grant (2008404), a US National Science Foundation CAREER award (IIS-1414452), and an Israel Science Foundation grant (1083/13).

6. REFERENCES

- [1] A. Azaria, Z. Rabinovich, S. Kraus, and C. Goldman. Strategic information disclosure to people with multiple alternatives. In *Proc. of AAAI*, pages 594–600, 2011.
- [2] B. Becker and T. Milbourn. How did increased competition affect credit ratings? *Journal of Financial Economics*, 101(3):493–514, 2011.
- [3] M. Chhabra, S. Das, and D. Sarne. Expert mediated search. In *Proc. of AAMAS*, pages 415–422, 2011.
- [4] M. Chhabra, S. Das, and D. Sarne. Expert-mediated sequential search. *European Journal of Operational Research*, 234(3):861–873, 2014.
- [5] C. Cramps and A. Hollander. Duopoly and quality standards. *European Economic Review*, 39(1):71–82, 1995.
- [6] S. Das. The effects of market-making on price dynamics. In *Proc. of AAMAS*, pages 887–894, 2008.
- [7] D. Dranove and G. Jin. Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–963, 2010.
- [8] N. S. Economides. A note on equilibrium in price-quality competition. *Greek Economic Review*, 7(2):179–186, 1985.
- [9] E. Farhi, J. Lerner, and J. Tirole. Fear of rejection? Tiered certification and transparency. Technical report, National Bureau of Economic Research, 2008.
- [10] C. Hajaj, N. Hazon, D. Sarne, and A. Elmalech. Search more, disclose less. In *Proc. of AAAI*, pages 401–408, 2013.
- [11] N. Hazon, Y. Aumann, S. Kraus, and D. Sarne. Physical search problems with probabilistic knowledge. *AI*, 196:26–52, 2013.
- [12] D. Hennes, D. Claes, W. Meeussen, and K. Tuyls. Multi-robot collision avoidance with localization uncertainty. In *Proc. of AAMAS*, pages 147–154, 2012.
- [13] J. Huang and A. Kazeykina. Optimal strategies for reviewing search results. In *Proc. of AAAI*, pages 1321–1326, 2010.
- [14] J. Kephart and A. Greenwald. Shopbot economics. *JAAMAS*, 5(3):255–287, 2002.
- [15] J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32:731–752, 2000.
- [16] C. Lim, J. Bearden, and J. Smith. Sequential search with multiattribute options. *Decision Analysis*, 3(1):3–15, 2006.
- [17] J. MacQueen. Optimal policies for a class of search and evaluation problems. *Management Science*, 10(4):746–759, 1964.
- [18] E. Manisterski, D. Sarne, and S. Kraus. Cooperative search with concurrent interactions. *JAIR*, 32:1–36, 2008.
- [19] J. McMillan and M. Rothschild. Search. In *Handbook of Game Theory with Economic Applications*, pages 905–927, 1994.
- [20] P. Milgrom. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12:380–391, 1981.
- [21] M. Motta. Endogenous quality choice: price vs. quantity competition. *Journal of Industrial Economics*, pages 113–131, 1993.
- [22] D. Sarne. Competitive shopbots-mediated markets. *ACM Transactions on Economics and Computation*, 1(3):17:1–17:41, 2013.
- [23] D. Sarne and S. Kraus. Managing parallel inquiries in agents’ two-sided search. *AI*, 172(4-5):541–569, 2008.
- [24] M. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.
- [25] 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.
- [26] R. Wright. Job search and cyclical unemployment. *Journal of Political Economy*, 94(1):38–55, 1986.