Competitive Comparison-Shopping Mediated Markets

David Sarne
Computer Science Department
Bar-Ilan University,
Ramat-Gan, 52900 Israel
Email: sarned@cs.biu.ac.il

***Abstract*—This paper considers markets mediated by self-interested comparison shopping agents. The comparative search conducted by the agents is driven by incentives offered by sellers, the cost incurred by the search, and competition dynamics that arise in the multi-agent setting. Based on models of economic search theory, the paper provides a formal analysis of the strategies used by the agents and the corresponding expected buyers' expense and sellers' net revenue. Equilibrium analysis is given for homogeneous environments in which all agents share the same search characteristics. Using this latter environment, it is demonstrated how the transition to competitive comparison-shopping mediated market can in some cases result both with lower expected expense to buyers and higher expected net revenue to sellers.**

***Keywords*-Agent-Mediated Markets; Multiagent Systems;**

I. Introduction

In an effort to fully exploit the plethora of retailers and virtual stores over the internet in electronic marketplaces, small-volume buyers and individuals adopt the use of autonomous agents for enhancing their buying experience. Recent research has suggested many applications in which agents can be used in order to facilitate consumer-related activities over the different stages of the consumer's buying experience [10]. In particular, emphasis has been placed on the integration of software agents into the *Merchant brokering* stage. In this stage the buyer searches for sellers who offer a specific desired product, and learns their posted price [11]. This process can be facilitated by many commercial comparison-shopping agents (CSAs) that can be found over the Web (e.g., PriceScan.com, Shopping.com, MySimon.com). The main advantage of CSAs, in this context, is in their capability to automatically query multiple vendors, searching for price information on desired goods and present it in a consolidated and compact format [1], [14].

While the implementation of CSAs can be based on maintaining a database of prices (i.e., storing the price in which each seller is selling a given product) which updates on a timely basis, it is definitely not the preferred option. This is due to the ever-increasing frequency of price updates in electronic markets nowadays. This phenomena has empirical evidence in literature [2] and is also well supported in theory. For example, dynamic pricing theories suggest that sellers can benefit from re-pricing their goods as often as possible based on their observations of the competitors' prices [14]. Alternatively, E-retail managers may use "hit and run" sales strategies undertaking short-term price promotions at unpredictable intervals - a method shown to be effective and widely used [2]. Therefore a reliable CSA is expected to use real-time querying of electronic merchants (rather than retrieve price data from a formerly collected price database) upon the arrival of price comparison requests from its users.

The real-time querying process is well recognized to be costly [1], [6], [13], [14] in a sense that the CSA needs to invest/consume some of its resources (e.g., computation, communication) for opening a connection with the remote server, extracting and filtering the relevant information and comparing it with the other results obtained. Some authors have argued that recent advances in communication technologies reduce search costs and other environmental inefficiencies in multi-agent environments [4]. However the general agreement is that these cannot be ignored completely [1]. The growing interoperability between different systems and environments in the internet age, followed by a phenomenal increase in the number and complexity of opportunities available, make the overall cost of acquiring information an important factor that affect performance [6], [13]. While current CSAs query all the sellers they are familiar with, upon the arrival of a new request, and return a list of price quotes, the expected increase both in the number of electronic merchants offering any specific product and in the demand for CSAs services suggest that CSAs are continuously proceeding to a point where their querying strategies will need to take into consideration search costs. These can either be direct costs associated with the search or the alternative gain that could have been obtained if the resources required for querying would have been allocated for the sake of other incoming requests [17].

While economic search theory supplies a framework for operating CSAs for the purpose of searching efficiently for the minimum price [15], [16], [20], [8], one must keep in mind that CSAs are, in most cases, self interested entities that attempt to maximize their own revenues. Today, CSAs business models are commonly based on advertisements or payments coming from vendors either as commissions for directing buyers or payments for being listed. Operating as a self-interested agent, the CSA is not necessarily driven by price minimization considerations, but rather by the expected gain from introducing different information to the user. This has many equivalents in physical markets. For example, consider real-estate brokers and travel agents. These mediators will not necessarily introduce to the customer the flight/hotel option or the apartment that best fit her needs but might often prefer to introduce first those options from which they gain the highest commission overall.

This paper considers buyer-seller markets mediated by several self-interested CSAs whose (costly) search is motivated by commissions offered by sellers for purchases made by buyers they direct to their web-sites. The CSAs thus need to consider, when constructing their search strategy, not only the benefits in extending their search versus the costs associated with it but also the competition dynamics that arise from the influence their search strategy has on other CSA strategies and vice versa. For comparison purposes a model where the CSAs are buyer-operated, i.e., the buyer agent controls the CSAs' search however also needs to cover the search cost incurred, is used. The main contributions of this paper are twofold. First, it supplies a formal modeling and analysis of a multi-CSA based market, where CSAs are self-interested competing agents. While the model provides a more realistic representation of future markets, the sequencing aspect associated with the CSAs' search precludes direct extraction of the equilibrium set of CSAs' strategies for general environments. Therefore, a second contribution of the paper is in supplying equilibrium analysis of multi-CSA mediated homogeneous markets, i.e., where CSAs share the same search cost and all seller agents are offering a similar fixed commission. The importance of this latter analysis is that it is then used for illustrating several important non-intuitive performance characteristics of the model. Particularly, it is shown that there is no guarantee that the beneficiaries of the transition to self-interested commission-driven CSAs (in comparison to buyer-operated CSAs) are the buyer agents on the expense of the seller agents. This, despite the buyer's direct saving of the costs incurred by the search process and the commission expenses of sellers. In fact, the paper presents possible scenarios where both agent types can benefit from the new model.

II. Model and Analysis

The new model considers an electronic marketplace populated by N_s seller agents, and some buyer agents. To facilitate their search for the products they seek to buy, buyer agents use the services of CSAs. A CSA can query any seller agent s_i for a price quote while the querying process itself incurs cost c_i. The price quote received from specific seller agent s_i, denoted q_{s_i}, is assumed to be drawn from a probability distribution function $f_i(q)$ that remains constant along time. This latter assumption is supported by recent empirical research in well-established online retail markets, presenting evidence for the persistence of price dispersion [2], [3], [5], [7].[1] Furthermore, the model assumes that while the CSAs and the buyer agents are not familiar with the specific prices offered by the seller agents at any given time they are acquainted with (or can learn over time) the distribution functions, from which these derive, $\{f_i(q)|0 \leq i \leq N_s\}$, and the cost of querying each seller agent in the market.

For exposition purposes, during any stage of a CSA's search, S_{sampled} is used to denote the set of sellers that have already been sampled by the CSA (and thus their posted price has already been revealed) and $S_{\neg\text{sampled}}$ is used to denote the set of seller agents whose prices are still unknown. The CSA's strategy is thus the mapping from the price quotes received so far $\{q_{s_i}|s_i \in S_{\text{sampled}}\}$ to a seller $s_j \in \{\emptyset, S_{\neg\text{sampled}}\}$ to be queried next, where $s_j = \emptyset$ if the CSA decides to terminate the search (or if $S_{\neg\text{sampled}} = \emptyset$), in which case it returns to the buyer agent a price quote $q \in \{q_i|s_i \in S_{\text{sampled}}\}$.

A. Buyer-Operated CSA

In the buyer-operated scenario, it is assumed that whenever a buyer agent requests the services of a CSA it gains full control over that agent, however needs to account for its cost of search. Therefore the optimal search strategy is the one that minimizes the expected sum of the accumulated cost of search and the minimal product price found. In this case, determining the optimal search sequence and stopping rule could be mapped to "Pandora's problem" [20] and consequently the optimal search strategy is as follows: The CSA initially assigns a reservation value to each seller s_i, denoted by R_i, satisfying: $c_i = \int_{q \leq R_i}(R_i - q)f(q)dy$. The

[1] In particular, we refer to findings of a considerable turnover in firms' relative positions in the distribution of prices over time and a significant variation in the identity of the low-price firm for the same product over time [2]. These findings contrast with the classic Law of One Price in spite of the fact that the underlying products being compared are homogeneous and the marginal costs of the products are essentially identical across retailers [5].

reservation value in this case is used as a reservation price. At each stage of its search process, the CSA picks the seller from the set $S_{\neg\text{sampled}}$ with the lowest reservation price. If this seller's reservation price is lower than the minimum price quote received from sellers in S_{sampled}, then the CSA queries this seller (formally, query seller $s_i \in S_{\neg\text{sampled}}$ for which $R_i \leq min(q_{s_j}|s_j \in S_{\text{sampled}})$ and $R_i \leq R_k$ $\forall R_k \in S_{\neg\text{sampled}}$, if one exists). If no such seller agent exists then the CSA terminates its search and returns the minimum price found so far and the seller associated with that price.

B. Self-Interested CSA

An alternative mode of operation for the CSA is as an autonomous self-interested agent with the goal of maximizing its own net revenue. Here, the incentive for search relates to commissions offered by the seller agents. In addition, the CSA faces competition, resulting from the existence of other CSAs in the market, thus it needs to include in its decision making process also its estimate of the buyer's buying probability given the quotes it may obtain from each seller. Let $P(q)$ denote the probability that the buyer agent will actually buy the product if offered at price q, and let $M_i(q)$ denote the commission payed to the CSA by seller agent s_i if a CSA-directed buyer agent buys the product at price q.

The optimal search strategy of the the CSA in this case can be found, once again, by mapping the search problem to Pandora's problem, taking into account buying probability and offered commissions. Here, the CSA will assign a reservation value to each seller, however this time the reservation value will be expressed in terms of revenue, and correlated with the expected commission from sampling the seller agent (as opposed to a reservation price). The reservation value in this case will be referred to as reservation expected commission. The reservation expected commission of seller s_i, denoted R_i, can be extracted from:

$$c_i = \int_{M_i(q)P(q) \geq R_i} (M_i(q)P(q) - R_i) f_i(q) dq \qquad (1)$$

At each stage of its search process, the CSA picks the seller from the set $S_{\neg\text{sampled}}$ with the highest reservation expected commission. If this seller's reservation expected commission is greater than the expected commission associated with each of the price quotes found so far, then the CSA queries this seller (formally, query seller $s_i \in S_{\neg\text{sampled}}$ for which $R_i \geq max(q_{s_j}P(q_{s_j})|s_j \in S_{\text{sampled}})$ and $R_i \geq R_k$ $\forall R_k \in S_{\neg\text{sampled}}$, if one exists). If no such seller agent exists then the CSA terminates its search and directs the buyer agent to the seller agent whose price quote maximizes the CSA's expected revenue, i.e., to seller agent $s_j \in S_{\text{sampled}}$ offering price quote q_{s_j}, satisfying: $M_j(q_{s_j})P(q_{s_j}) \geq M_k(q_{s_k})P(q_{s_k}), \forall s_k \in S_{\text{sampled}}$. For convenience, s_i is re-defined as the seller agent associated with the i-th highest reservation expected commission. The expected net benefit of the CSA when using the set of reservation values $(R_1, ..., R_{N_s})$, denoted $V(R_1, ..., R_{N_s})$ is thus:

$$V(R_1, ..., R_{N_s}) = \int_{(M_1P(q_1) \geq R_2)} (M_1P(q_1) - c_1) f_1(q_1) dq_1 + \qquad (2)$$
$$\sum_{i=1}^{N_s-1} \int_{M_1P(q_1)<R_i} f_1(q_1) \int_{max_2<R_i} f_2(q_2) \quad ...$$
$$\int_{(max_i > R_{i+1})} \Big(max_i - \sum_{k=1}^{i} c_k \Big) f_i(q_i) dq_i dq_{i-1} ... dq_1 +$$
$$\int_{M_1P(q_1)<R_{N_s}} f_1(q_1) \int_{max_2<R_{N_s}} f_2(q_2) \; ... \int_{max_{N_s-1}<R_{N_s}} f_{N_s-1}(q_{N_s-1}) \int_{q_{N_s}} \Big(max_{N_s} -$$
$$\sum_{k=1}^{N_s} c_k \Big) f_{N_s}(q_{N_s}) dq_{N_s} dq_{N_s-1} ... dq_1$$

where $M_iP(q) = M_i(q)P(q)$ is the expected commission received from seller s_i if choosing its quote q and $max_k = max(M_1P(q_1), M_2P(q_2), ..., M_kP(q_k))$. The first component on the right hand side of the equation deals with the case where the expected commission from the first quote received is greater than the next highest reservation expected commission. In this case, the CSA terminates its search, incurring the cost c_1 only. The following component deals with the case where the maximum of the expected commissions associated with the price quotes received up to querying seller s_i is greater than the reservation expected commission R_{i+1}. In this case the CSA terminates the search and directs the buyer agent to the seller associated with the maximum expected commission. The accumulated cost incurred in this case is $\sum_{k=1}^{i} c_k$. The probability that seller agent s_i is actually queried equals the probability that none of the quotes received from sellers $\{s_1, ..., s_{i-1}\}$ is associated with expected commission greater than R_i. Finally, the last component relates to the case where the CSA queries all N_s seller agents. In this case it simply directs the buyer agent to the seller associated with the maximum expected commission.

From the buyer agent's viewpoint, the above search strategy results with a specific distribution of the price q, returned eventually by the CSA. The probability that the CSA returns a price quote smaller or equal to q is denoted by $G(q)$. Given a reservation values set $\{R_1, ..., R_{N_s}\}$ used by the CSA, the probability $G(q)$ can be calculated as:

$$G(q)=\int_{(M_1P(q_1)\geq R_2)\wedge(q_1<q)} f_1(q_1)dq_1 + \sum_{i=1}^{N_s-1}\int_{M_1P(q_1)<R_i} f_1(q_1) \int_{max_2<R_i} f_2(q_2)... \int_{(max_i>R_{i+1})\wedge(argmax_{q_j}(M_jP(q_j)|1\leq j\leq i)\leq q)} f_i(q_i)dq_idq_{i-1}...dq_1 + \int_{M_1P(q_1)<R_{N_s}} f_1(q_1) \int_{M_2P(q_2)<R_{N_s}} f_2(q_2) ... \int_{M_{N_s-1}P(q_{N_s-1})<R_{N_s}} f_{N_s-1}(q_{N_s-1}) \int_{argmax_{q_j}(M_jP(q_j)|1\leq j\leq N_s)\leq q} f_{N_s}(q_{N_s})dq_{N_s}dq_{N_s-1}...dq_1 \quad (3)$$

The first derivative of $G(q)$, $g(q) = \frac{dG(q)}{dq}$, is thus the probability distribution function of the price quote received from the CSA and the expected returned quote is given by:

$$E(q_{received}) = \int_{q=0}^{\infty} qg(q)dq \quad (4)$$

C. Multi-CSA Environment

The probability that the buyer agent will actually buy the product based on the CSA's returned price quote q, $P(q)$, depends on the search strategies used by other CSAs the user can contact. Since the user is merely interested in minimizing its expense, the probability it will buy the product, if given a quote q, is the probability that none of the other CSAs contacted in parallel returns a better (i.e., lower) price quote. Extending the analysis to the multi-CSA case, the model uses N_c to denote the number of CSAs in the market and $P_i(q)$ to denote the probability that given a price quote q returned by the i-th CSA, denoted CSA_i, the buyer agent will buy the product from the seller associated with that price. Notice that different CSAs may incur different costs and be suggested different commissions when querying the same seller, thus the buying probability function is CSA-specific. The probability $P_i(q)$ can be calculated as:

$$P_i(q) = \Pi_{j=1,j\neq i}^{N_c}(1 - G_j(q)) \quad (5)$$

where $G_j(q)$ is the probability that CSA_j returns a price quote smaller or equal to q. Substituting $P_i(q)$ in Equation 1, an equation that describes each CSA's optimal search strategy as a function of the search strategies used by the other CSAs in the market can now be constructed. The equilibrium in this case is a set of search strategies from which none of the CSAs have an incentive to deviate, i.e., a set of reservation-expected-commissions sets where none of the agents can benefit (according to Equation 2) from changing any subset of reservation-expected-commissions it uses. Given a set of strategies, one can determine if it is an equilibrium set by constructing the set of equations describing the optimal reservation-expected-commission assigned to each seller agent by each CSA, based on Equation 1, and checking if each CSA's reservation values, given those used by the other agents, are in fact the ones assigned to it. As for extracting a set of strategies which is in equilibrium for the general case, this can become extremely complex. The complexity derives both from the inclusion of probability distribution functions in Equation 3 and the fact that the CSAs' search strategies are defined as sequences. The first problem implies that for some distribution function (e.g., normal distribution function) it will be difficult to extract the appropriate reservation values using direct immediate calculation. Here, one will need to employ specific function-dependent approximation techniques (e.g., solving Taylor series expansion, or using the Trapezoidal Rule and Simpson's Rule). The second and more acute problem, relates to the fact that Equation 3 describes a sequential process and the order of the seller agents queried by a given CSA changes when other CSAs' strategies change. These sequence changes require re-constructing Equations 1 and 3 for all agents whenever one of the agents changes its search strategy. Nevertheless, as illustrated in the following section, several interesting results of multi-CSA mediated markets can be illustrated even by using simpler variants of the model.

III. Homogeneous Environment Equilibrium

To demonstrate equilibrium dynamics in a multi-CSA mediated environment a simpler homogeneous variant of the model is used, where all CSAs share the same search cost, denoted c, all seller agents are offering identical fixed commission, M, and CSAs are not limited by a finite decision horizon (i.e., can obtain as many price quotes, each incurs a cost c, as requested).[2] These assumptions substantially simplify the analysis and yet enable illustrating some of the important effects of the new model.

A. Buyer-Operated CSA

The buyer's problem of operating a single self-owned CSA in homogeneous environment is equivalent to the classical one-sided sequential search model [9], [15]. In this case, the CSA uses a fixed reservation price R, and sequentially queries (in a random order) seller agents, terminating its search upon finding a price quote smaller or equal to R. The optimal reservation price (i.e., the one that minimizes

[2]This latter assumption can be justified by the high entrance and leave rates of seller agents to/from the market. It is also supported by the emergence of dynamic pricing as a recurring later query to the same seller may yield a different price quote.

the expected cost) in this case can be extracted from:

$$c = \int_{q=0}^{R} F(q)dq \tag{6}$$

The expected expense of the CSA (and thus, the expected expense of the buyer agent) in this model equals the optimal reservation price. Furthermore, since the cost of search is linear in the number of queries made, it is straightforward that the optimal search strategy is to use a single CSA rather than owning and operating several of them.[3] The expected overall number of queries made in this case is given by $\frac{1}{F(R)}$, thus the expected expense for the buyer can be divided into the expected cost of the search, $\frac{c}{F(R)}$, and the expected payment for the product, which is also the expected revenue of the seller, $(\int_{q=0}^{R} \frac{qf(q)}{F(R)}dq)$ [9].

B. Equilibrium Characteristics

Since the sellers are identical, they are all associated, from the CSA's point of view, with the same reservation value, R. Obviously, the smaller the price the higher the probability it is the minimum out of the quotes returned by the CSAs. Since the commission is fixed, the reservation value can be expressed in terms of reservation prices, i.e., A CSA using a reservation value R will never return a price quote greater than R but rather will keep on searching until a price quote smaller or equal to R is found. Given $P_i(q)$, the optimal reservation value used by CSA_i can be extracted from the following equation:

$$c = M \int_{q=0}^{R} (P(q) - P(R))f(q)dy \tag{7}$$

The probability that the CSA will return a price quote smaller or equal to q, $G(q)$, is thus given by $G(q) = \frac{F(q)}{F(R)}$ for $q \leq R$ and $G(q) = 1$ otherwise. Consequently, the probability that the buyer agent will buy the product at a price q found by a given CSA equals the probability that none of the other $N_c - 1$ CSAs returns a lower quote, thus $P(q) = (1 - \frac{F(q)}{F(R)})^{N_c-1}$, $q \leq R$ (and $P(q) = 0$ otherwise). As expected, $P(q)$ decreases as q increases and as N_c increases. When there is no competition at all (i.e., $N_c = 1$), $P(q) = 1$ is obtained. Substituting $P(q) = (1 - \frac{F(q)}{F(R)})^{N_c-1}$ in Equation 7, and setting $R' = R$, the equilibrium reservation price R used by the CSAs can be extracted from:

$$c = \frac{MF(R)}{N_c} \tag{8}$$

[3]A pure sequential search which samples a single observation at a time dominates in this case a parallel search, according to economic search theory.

Proposition 1: When using the equilibrium strategy, the expected net benefit of each CSA is zero.

Proof: The proof is obtained by a simple reduction to a Bertrand competition model. The CSAs can be seen as non-cooperating firms producing homogeneous products, having the same (constant) marginal cost. The demand is linear and the firms compete in price, choosing their respective prices simultaneously. The equilibrium result in this case is marginal cost pricing, thus the expected net benefit of each CSA is zero.

This can also be seen directly from Equation 2. For the homogeneous environment, the modification of Equation 2 is:[4]

$$V(R) = \frac{-c + \int_{q=0}^{R} MP(q)f(q)dy}{F(R)} \tag{9}$$

which transforms, after substituting $P(q) = (1 - \frac{F(q)}{F(R)})^{N_c-1}$, into $V(R) = \frac{M}{N_c} - \frac{c}{F(R)}$. Finally, substituting $F(R) = \frac{cN_c}{M}$ (Equation 8), $V(R) = 0$ is obtained. ■

One may wonder, given the last result, what is the incentive for the CSAs to operate in homogeneous competitive environments if their search net-benefit is zero. As shown in the following subsection, having the CSAs operate in homogeneous environment can improve both buyers' expense and sellers' net benefit. Therefore, market makers and market designers may suggest fixed incentive to CSAs for taking an active part in the market despite not making any profit from the search itself.[5]

The expected price in which the buyer agent buys the product eventually, denoted q_{min} is the expected minimum of the price quotes received from the different CSAs. The probability distribution function and commutative distribution function of q_{min} when using N_c identical CSAs, denoted $f^{N_c}(q)$ and $F^{N_c}(q)$, respectively, can be calculated using:

$$F^{N_c}(q) = 1 - (1 - G(q))^{N_c} \tag{10}$$

$$f^{N_c} = \frac{dG_{N_c}(q)}{dq} = N_c f(q)(1 - G(q))^{N_c-1} \tag{11}$$

Buyer agents' expected minimum expense is thus $E[q_{min}] = \int_{q=0}^{R} qf^{N_c}(q)dq$. Using integration by parts this transforms

[4]Since the CSA is not concerned with a finite decision horizon, it is using a stationary reservation value. Equation 9 is a modification of the expected overall cost of a searcher engaged with economic one-sided search [15].

[5]Notice that a sufficient condition that guarantees an active participation of the CSA in the marketplace is requiring that the CSA will supply at least one price quote whenever addressed. Given this requirement, the CSA will necessarily search according to the equilibrium reservation value because any other search policy will result with a loss.

into:

$$E[q_{min}] = \int_{q=0}^{R} (1 - \frac{F(q)}{F(R)})^{N_c} dq \qquad (12)$$

Similarly, seller agents' net revenue is $E[q_{min}] - M$.

Proposition 2: As the number of competing CSAs increases, the expected minimum quote received by the buyer agent (and consequently sellers' net revenue) increases.

Proof: Substituting $F(R) = \frac{cN_c}{M}$ (Equation 8) in Equation 12, $E[q_{min}] = \int_{q=0}^{R} (1 - \frac{MF(q)}{cN_c})^{N_c} dq$ is obtained. Since the integrated term increases as N_c increases, $E[q_{min}]$ also necessarily increases as N_c increases. ∎

Proposition 2 above suggests a non-intuitive market behavior. While one would expect the increase in the number of competing CSAs to induce competition, resulting in further search, the expected expense of buyer agents actually increases. This surprising result can be explained using Proposition 1. Since the CSAs have a zero net revenue from their search, the increased competition results in an increase in the reservation price they use. The expected minimum quote is now affected positively by the additional quotes (from the added CSAs) and negatively by the increased reservation price.

While the focus of this paper is environments mediated by self-interested competing CSAs, several results that can be obtained for the homogeneous environment, when mediated by non-competing CSAs, are highlighted. First, if the agents were allowed to join forces and form a monopoly, the optimal strategy for them would be to randomly select a single CSA that will be willing to serve the buyer agent upon receiving a request for price comparison. The selected CSA would sample a single seller and receive the commission (as there is no competition). The net revenue of the CSAs in this case is $\frac{M-c}{N_c}$, the expected quote received (i.e., expected buyer's expense) is $E[q]$ and the expected net revenue of the seller agent from which the product is eventually purchased is $E[q] - M$.

Second, if the CSAs were allowed to create a cartel (and had the means for ensuring none of them deviates from it) the optimal strategy for them would be to obtain a single quote (each). The net revenue of the CSAs in this case is $\frac{M}{N_c} - c$, the expected buyer expense is the minimum of a sample of size N_c, $E[q_{min}] = \int_{q=0}^{\infty} (1 - F(q))^{N_c} dq$, and the expected net revenue of the seller agent from which the product is eventually purchased is $E[q_{min}] - M$. It is worth noticing that in the cartel case the expected minimum quote returned to the buyer agent, q_{min}, actually decreases as N_c increases. This, in comparison to an increase, in the competitive model according to Proposition 2.

Finally, if the CSAs are operated by the buyer agents (i.e., non-seller-incentive environment), their optimal reservation value can be extracted from $c = \int_{q=0}^{R} F(q)dq$ (according section III-A). In this case the expected expense for buyer agents and the expected net revenue of the sellers are R and $\int_{q=0}^{R} \frac{qf(q)}{F(R)} dq$, respectively.

C. Results from Specific Distributions of Prices

For illustration purposes a synthetic environment based on the uniform distribution function ($f(q) = 1$, $F(q) = q$, $0 \le q \le 1$) is used. Table I summarizes the appropriate formulations in this case, based on the analysis given in the former paragraphs.

Figure 1 depicts the expected expense of buyers and the expected net-revenue for sellers as a function of the number of competing CSAs, N_c, using the analysis given above. The commission offered by seller agents in the environment used is $M = 0.01$, the cost of search is $c = 0.0003$ and prices are distributed uniformly in the interval $(0, 1)$. The vertical axis represents monetary values thus the higher the value the greater is the expected expense for the buyer agent (and thus the worse is its performance) and the greater is the expected net revenue of the seller agent. When the buyer agent is the one operating the CSA, both its own and sellers' performance do not depend on the number of competing CSAs as the buyer's optimal strategy is to search sequentially. However, in the scenario where CSAs are self-interested agents, different levels of competition (resulting from the number of CSAs) yield different performance for the buyers and sellers. As expected from the theoretical analysis given in the former section, both buyer and seller performance curves increase as the number of competing CSAs increase. As observed from Figure 1, for some N_c values (e.g., for $N_c \ge 5$) the buyers are better off if using self-operated CSAs, and for others will benefit from CSAs competition. Similarly, for some N_c values (e.g., $N_c \ge 3$) seller agents get greater expected net revenue when facing competing self-interested CSAs (despite the commission paid). The most striking result in the figure is that for some N_c values ($N_c = 3$ and $N_c = 4$) both buyer and seller agents benefit from having self-interested CSAs competition (i.e., buyers spend less while sellers' net revenue rises) while none of the CSAs actually lose from operating in the market. This win-win situation can be explained as follows. By offering a commission to the CSAs, the seller agents fully subsidize the costs associated with search. If this subsidy could have been transferred completely to the buyer agents, the latter ones would improve their performance and the seller would worsen theirs. Nevertheless, the multi-CSA scenario suggests several agents searching in parallel instead of one agent searching sequentially as in the self-operated case, which makes the overall search process less efficient. Thus despite the spending on subsidizing the search process,

Measure	Uniform			
	Buyer-Operated	Competing	Monopoly	Cartel
Reservation Value	$R=\sqrt{2c}$	$R=\frac{cN_c}{M}$	N/A	N/A
E[total queries]	$\frac{1}{\sqrt{2c}}$	$\frac{M}{c}$	1	N_c
CSA's revenue	N/A	0	$M-c$	$\frac{M}{N_c}-c$
Buyer's expense	$\sqrt{2c}$	$\frac{cN_c}{M(N_c+1)}$	0.5	$\frac{1}{N_c+1}$
Seller's revenue	$\frac{\sqrt{2c}}{2}$	$\frac{cN_c}{M(N_c+1)}-M$	$0.5-M$	$\frac{1}{N_c+1}$-M

Table I
SUMMARY OF RESULTS FOR THE UNIFORM DISTRIBUTION HOMOGENEOUS ENVIRONMENT.

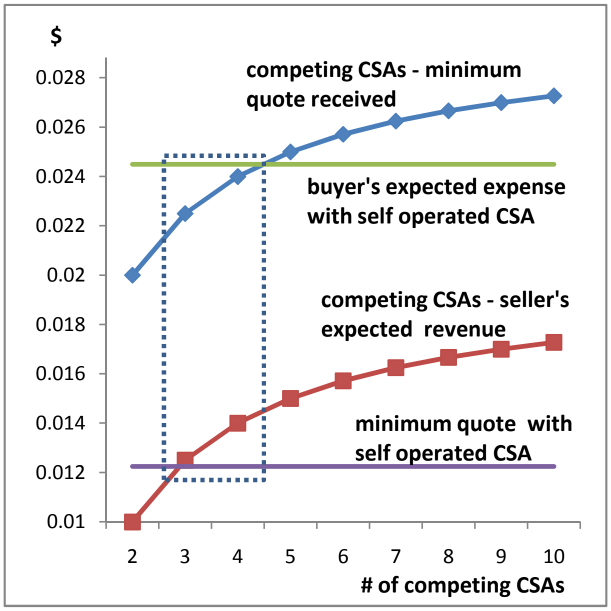

Figure 1. Buyer agents' expected expense and seller agents' expected net revenue.

the seller agents benefit from its inherent inefficiencies. In a similar manner, despite the inefficiencies of the search, the buyer agents benefit from having CSAs perform the search for them for free.

It is notable that in this example, colluding behavior is highly beneficial for the CSAs as well as for the seller agents. The seller agents' expected revenue rises to 0.49 in the monopoly case and in the cartel case starts at 0.323 (for $N_c = 2$) and gradually decreases until converging to 0.019 (to which the expected net revenue in the competing CSAs scenario converges).

IV. Related Work

The comparison-shopping domain has attracted the attention of researchers and market designers for the past 10 years [1], [10], [11]. The majority of the analysis in this area concern the influence of CSAs on retailers' and consumers' behavior [7], [12], [18], given the premise that shopbots can significantly reduce search costs (associated with obtaining price information).

The investigation of CSAs' search strategies [6], [13] builds on economic search theory, which considers an individual interested in locating an opportunity which will minimize its expected cost (or maximize its expected utility), while the search process is associated with a search cost (see [15] for literature review of search theory). The three main search models found in literature are the fixed sample size model, the sequential model and the variable sample size model. In the fixed sample size model [19] the searcher draws a single sample where all observations are taken simultaneously. In the sequential search strategy [16], which for the general finite decision horizon case is also known as "Pandora's Problem" [20], the searcher draws one observation at a time, allowing multiple search stages. The last search method [8] suggests a combined approach in which several observations may be obtained in any period. Nevertheless, despite recent advances in investigating buyer-seller pricing and search dynamics when self-interested CSAs are used [13], [14], none of the previous works includes in its framework and analysis the dynamics that arise from multi-CSA competition.

V. Conclusions

The formulation supplied in this paper extends the single CSA model into a multi-CSA framework by including the unique dynamics that arise when one CSA's search strategy affects the search strategies used by other CSAs (and vice versa). In the model presented, these dynamics are directly correlated with the changes in the buyers' purchase probability given the price quote returned by the CSA. The analysis methodology given for the general case is used in this paper for supplying an equilibrium analysis for the homogeneous case, however can also be used as an infrastructure for nurturing further buyer and seller oriented analysis in multi-CSA mediated environments. For example, investigating how seller agents should set the commissions they offer and finding the set of commissions used in equilibrium.

The homogeneous environment equilibrium analysis suggests several interesting (and often non-intuitive) results that apply to multi-CSA-mediated markets. For example, when considering the transition from a buyer-operated CSA to self-interested competing CSAs, one would expect the result to be a reduction in the buyer agents' expected expense (as they do not incur any costs associated with the search anymore) and a reduction in the seller agents' net revenue (as they now need to pay commissions). This is indeed the result in some settings, however in many others an increase in buyer agent's expected expense can be observed. Similarly, the result to seller agents can be either an increases or a decreased in the net-revenue. The most interesting result, though, is that there are scenarios where both buyer and seller agents benefit from the transition to self-interested competing CSA-mediated environment. Another non-intuitive result relates to the effect of CSAs' competition over buyer agents' expected expense. One would expect a decrease in the expected expense as the number of competing CSAs in the market increases (i.e., when the competition rises). Alas, as proved for the homogeneous environment, an increase in the number of competing CSAs results with an increase in the buyer agents' expected expense.

Finally, while equilibrium analysis often fails to predict the strategies used in practice by people, in this case the analysis directly addresses a reality in which artificial agents are the main players (as seen in CSA-based markets today). These agents are inherently more rational and less computationally bounded than people, thus there is substantial potential that the theoretical analysis and computational techniques developed in this paper will be implemented in the real world.

ACKNOWLEDGMENT

This work was partially supported by BSF grant. Special thanks go to Sanmay Das for his help in motivating the multi-CSA model and constructing its initial formulation.

REFERENCES

[1] Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, 42(12):1676–1692, June 1997.

[2] M. Baye, J. Morgan, and P. Scholten. Temporal price dispersion: Evidence from an online consumer electronics market. *Journal of Interactive Marketing*, 18(4):101–115, 2004.

[3] M. Baye, J. Morgan, and P. Scholten. Persistent price dispersion in online markets. In D. Jansen, editor, *The New Economy and Beyond*. Edward Elgar Press, Northampton, MA, 2006.

[4] S. Biswas and Y. Narahari. Analysis of supplier competition in electronic marketplaces. In *Proceedings of the International Conference on Automation, Energy, and Information Technology, EAIT-2001*, Indian Institute of Technology, Kharagpur, December 2001.

[5] E. Brynjolfsson, E. Hu, and M. Smith. Consumer surplus in the digital economy: Estimating the value of increased product variety at online bookseller. *Management Science*, 49(11):1580–1596, 2003.

[6] S. Choi and J. Liu. Optimal time-constrained trading strategies for autonomous agents. In *Proc. of MAMA'2000*, 2000.

[7] K. Clay, R. Krishnan, E. Wolff, and D. Fernandes. Retail strategies on the web: Price and non-price competition in the online book industry. *Journal of Industrial Economics*, 50:351–367, 2002.

[8] S. Gal, M. Landsberger, and B. Levykson. A compound strategy for search in the labor market. *International Economic Review*, 22(3):597–608, 1981.

[9] A. Grosfeld, D. Sarne, and I. Spiegler. Modeling the search for the least costly opportunity. *European Journal of Operational Research*, 197(2):667-674, 2009.

[10] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review*, 13(2):147–159, June 1998.

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

[12] E. Johnson, W. Moe, P. Fader, S. Bellman, and G. Lohse. On the depth and dynamics of online search behavior. *Management Science*, 50(3):299–308, 2004.

[13] J. Kephart and A. Greenwald. Shopbot economics. *JAAMAS*, 5(3):255–287, 2002.

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

[15] J. McMillan and M. Rothschild. Search. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. 1994.

[16] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82:689–711, 1974.

[17] D. Sarne, S. Kraus, and T. Ito. Scaling-up shopbots: a dynamic allocation-based approach. In *AAMAS '07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 256–263, 2007. ACM.

[18] M. Smith. The impact of shopbots on electronic markets. *Journal of the Academy of Marketing Science*, 30(4):442–450, 2002.

[19] G. Stigler. The economics of information. *Journal of Political Economy*, 69(3):213–225, 1961.

[20] M. L Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–54, May 1979.