Strategic Information Platforms - Selective Disclosure and The Price of "Free"

CHEN HAJAJ, Bar Ilan University
DAVID SARNE, Bar Ilan University

This paper deals with platforms that provide agents easier access to the type of opportunities in which they are interested (e.g., eCommerce platforms, used cars bulletins and dating web-sites). We show that under various common service schemes, a platform can benefit from not necessarily listing all the opportunities with which it is familiar, even if there is no marginal cost for listing any additional opportunity. The main implication of this result is that platforms should extract their expected-profit-maximizing service terms not based solely on the fees charged from users, but they should also use the subset that will be listed as the decision variable in the optimization problem. The analysis applies to four well-known service schemes that a platform may use to price its services. We show that neither of these schemes generally dominates the others or is dominated by any of the others. For the common case of homogeneous preferences, however, several dominance relationships can be proved, enabling the platform to identify the schemes that should be used as a default. Furthermore, the analysis provides a game-theoretic search-based explanation for a possible preference of buyers to pay for the service rather than receive it for free (e.g., when the service is sponsored by ads), a phenomena that has been justified in prior literature typically with the argument of willingness to pay a premium for an ad-free experience or more reliable platforms. The paper shows that this preference can hold both for the users and the platform in a given setting, even if both sides are fully strategic.

Categories and Subject Descriptors: I.2.11 [**ARTIFICIAL INTELLIGENCE**]: Distributed Artificial Intelligence - Intelligent agents, Multiagent systems

Additional Key Words and Phrases: Platforms and Services; Economics of Information; Two-Sided Markets; Price of Free; Service Schemes

1. INTRODUCTION

In many two-sided multi-agent markets the number of opportunities potentially available to agents is substantial. Nonetheless the agent needs to spend time and resources in locating such opportunities [He et al. 2003]. For example, a buyer who is interested in buying a specific product over the Internet can potentially find it in literally hundreds of online stores, most of which are unknown to her. Similarly, when looking for a mate over the Internet, there are potentially thousands of unknown candidates who are interested in matching as well. This plethora of opportunities has been a catalyst for the emergence of platforms that serve as mediators and are used primarily as a point of contact for users and opportunities, thereby saving users the need to invest their valuable resources in service or opportunity discovery [Bakos 1997; Rysman 2009]. For example, in eCommerce, B2B and B2C platforms, such as kompass.com and alibaba.com, are used to connect buyers and sellers from all over the world. Suppliers who register on these websites can upload their corporate profiles and product cata-

---

This work was partially supported by ISF grant 1083/13.
Authors' addresses: C. Hajaj and D. Sarne, Computer Science Department, Bar-Ilan University; email: chen.hajaj@biu.ac.il,sarned@cs.biu.ac.il.

EC'14, June 8–12, 2014, Stanford, CA, USA.
http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/2600057.2602864

logs and registered buyers can request quotations from them. Similarly, online dating platforms such as match.com and okcupid.com provide unmoderated matchmaking by allowing individuals to contact and communicate with each other over the Internet. Other examples include autotrader.com for connecting car shoppers with dealers and private owners and Yet2.com for connecting industrial firms with entrepreneurial ventures, research universities and individual inventors as a market for ideas.

The use of these platforms is beneficial in many ways. These intermediaries make many exchanges and matches that would not otherwise occur possible and increase the benefit of both sides in the market. Their main premise is that, with the use of the platform, users can save substantial time in locating applicable opportunities, avoid duplicate ones and have all of the information required for their decision in a compact, convenient and organized format [Bakos 1997; Kephart et al. 2000]. Recent research suggests many applications in which platforms can be used in order to facilitate the buyer's shopping experience [Rysman 2009; Hagiu and Lee 2011]. Still, even with such platforms, the users need to invest some of their resources in order to reason about the actual value of each opportunity listed (e.g., contacting a seller in order to receive a quote, chat with or actually date a person to reason about her qualities, contact a car owner or actually take the car to the mechanic to get a better estimate of its worth).

One of the most prominent questions in the research on platforms is how a platform should price its information-providing services [Nahum et al. 2012]. Being a self-interested agent, the platform seeks to maximize its profit, taking into consideration its own information providing costs and the payments it receives. The payments can be received from any of the sides using the market (e.g., buyers and sellers), taking over some of their surplus from using the platform. Alternatively, payments can be received from external entities that may benefit from the activity taking place in the market (e.g., advertisers or repositories willing to buy user information collected by the platform). As such, the service pricing question has become a prevalent theme in electronic commerce literature [Markopoulos and Kephart 2002; Waldeck 2008; Weyl 2010], and various schemes have been proposed to date (e.g., bundling, pay-per click, subscription) [Geng et al. 2005; MacQueen 1964; Sundararajan 2004]. Nevertheless, prior work which considered the service-pricing question took the set of opportunities listed by the platform to be given. In real-life, however, the platform is not limited to price-setting only but can control what opportunities will be included in its listings.

In this paper we analyze richer service-terms-setting strategies for the platform, ones that do not take the set of opportunities to be presented to be fixed, but rather allow the platform to control what subset of these will be presented. For example, in cases where the platform's profit is fully based on advertisements scattered throughout the different pages it presents to users the platform can potentially benefit from omitting some of the more attractive opportunities from its listings. This will increase the number of opportunities that users will check (and consequently the number of pages visited) as more exploration will now be required in order to distinguish the "good" opportunities from the "bad" ones.[1]

The analysis given in the paper encompasses four common service fee-setting schemes: (a) "subscription", where the user pays a one-time fee and gains access to all the opportunities listed in the platform's database; (b) "sponsored", where the service is free for the user and the platform gains from ads presented to the user along the process; (c) "per-click", where the user pays the platform for each opportunity it wants to explore further; and (d) "full service" where the user pays a one time fee to the platform and obtains the opportunity with the best value of those listed in the plat-

[1]This is of course not a general guideline, since it is possible that by removing the best opportunities some users will opt not to use the platform in the first place.

form's database. For each scheme, we show how the platform sets its expected-profit-maximizing service terms, based on the Stackleberg game (in which the platform is the leader and the users are the followers). We manage to demonstrate that the platform's preference of including in its listings only a subset of the opportunities it can potentially present holds even when the operational cost for the platform is fixed (i.e., when including more opportunities does not incur additional costs). This finding suggests that even in online markets, where one may consider the marginal cost of storing and maintaining information to be negligible, it is important for the platform to consider the listed subset as a parameter in its optimization considerations.

We demonstrate that for the general case of heterogeneous users, when given the option to also control the set of opportunities to be listed, none of the service schemes generally dominates any of the others from the platform's point of view. The main implication of this finding is that platforms should not use any of these schemes as a default but rather calculate their expected profit for each scheme individually, as presented in the paper. We also refer to the scenario, which is particularly common in eCommerece platforms, where users are homogeneous. For this case we prove several dominance relationships related to the choice of the service scheme, both from the users' and the platform's point of view. These dominance relationships are important for platforms when the option to choose their service scheme exists, and also for market designers or regulatory entities that might be interested in users' profit or social welfare, whenever they can influence the choice of the service scheme.

Another important finding, of an existential nature, is that from the users' point of view, the free use of platforms, e.g., those that are sponsored by ads, is not necessarily the preferred choice. Such preference towards "costly" usage may hold only when the platform optimizes based both on service terms and the subset to be used, as in the model analyzed in this paper. This preference is explained in our case by a more favorable listing. Furthermore, we demonstrate that for some settings the costly option is the dominant one both for the platform and for all users individually.

In the following section we formally present the model and the four service schemes used. In Section 3, we provide an equilibrium analysis for the resulting Stackleberg game, demonstrating the advantage of partial information disclosure. Later, in Sections 4 and 5, we prove some dominance relationships between the different schemes for the homogeneous case and demonstrate that there is no general dominance for the heterogeneous case. Related work is reviewed in Section 6. Finally, we conclude with a discussion and directions for future research in Section 7.

2. MODEL

The model slightly extends the standard one-sided platform-based search model of the kind commonly used in prior literature [Hagiu and Wright 2011; Sarne 2013]. For illustration purposes, we use the procurement application domain terminology, though as discussed towards the end of this section the model applies to a variety of domains. The procurement setting considers platforms such as alibaba.com, made-in-china.com, and gobizkorea.com bringing together procurement agents (hereafter denoted "buyers") and sellers (typically small businesses worldwide).

Sellers in such platforms offer well-defined products for sale, however the exact price of a product, as well as any other terms and conditions, are a priori uncertain and can be obtained only by contacting the seller directly (either by e-mail, a designated online inquiry form or through a chat service offered by the platform). It is assumed that sellers' offers are set exogenously, unaffected by the search strategies used by the buyers visiting the specific platform which are the focus of this paper. Many justifications for this has been presented in prior work [Xing et al. 2006; Sarne 2013], e.g., sellers usually operate in various parallel markets (thus prices are set by the aggregate of demands) and respond to a rich set of external factors (e.g., shipment and insurance constraints).

Buyers are assumed to be self-interested fully-rational agents. In order to successfully complete their task, buyers need to locate sellers that sell the product in which they are interested and check their offers. In its most basic form, the model assumes that buyers are fully price-sensitive, hence the benefit in a seller's offer is fully captured by its price attribute. Later we show that the analysis provided also supports the case where sellers' offers are evaluated according to a richer set of considerations (e.g., warranty and shipping time) by the value probability distribution function associated to each of them. The platform saves the buyers the trouble of locating sellers, however, as explained above, the price charged by each seller remains uncertain in the platform's level. This uncertainty about the price, q, requested by seller s_i is represented by the probability distribution function $f_i(q)$, which is assumed to be known to the buyer, e.g., based on the accumulated history of the seller's prices, characteristics and reputation (information that is potentially supplied by the platform).[2]

In order to obtain the price requested by seller s_i listed in the platform, the buyer needs to communicate with that seller, as explained above, which incurs a cost c_i^t (e.g., cost of time)[3], where t is the buyer's type (buyers can be heterogeneous in their characteristics, as explained below). The buyer can query (or ask the platform to query) as many sellers as it wants of those listed in the platform. Since the focus of the paper is the individual platform, the other alternatives to the platform that are available to a user of type t (e.g., competing platforms or locating and querying sellers without the use of a platform) are modeled through the use of a private value, denoted R_b^t, expressed in terms of the alternative expense for purchasing the product not through the platform, as provided in prior works [Niu et al. 2008].

The buyer's problem can be expressed as follows: Given the $\hat{N}$ sellers $S = \{s_1, ..., s_{\hat{N}}\}$ listed in the platform, the appropriate probability distribution function of prices $f_i(q)$ and the querying cost c_i^t $\forall i$, and given the platform's service terms, should the buyer use the platform, and if so, how should the different sellers to be queried (namely, who should be queried next and when to terminate the process, based on the results obtained so far). The goal of the buyer is to minimize its expected total expense for buying the product, thus the optimal strategy is the expected-expense minimizing one. The population of buyers is not necessarily homogeneous — buyers can be of diverse types, differing in their private value, their cost of querying the different sellers and the way in which they evaluate the different sellers' opportunities (e.g., when the valuation is not solely based on the price). We denote the portion of the buyers population that belongs to type t as $g(t)$. In this case, we use $f_i^t(q)$ to denote the value probability distribution function associated with seller s_i from buyers with type t point of view.

The platform is assumed to be a self-interested agent, thus driven by expected-profit-maximization considerations. There are many possible service schemes that the platform can potentially use, and obviously attempting to cover them all is implausible. In this paper, we focus on four common service schemes, each offering different advantages and disadvantages to both the platform and the buyer:

[2]The existence of such a probability distribution is commonly assumed in e-commerce research [Janssen et al. 2005; Waldeck 2008; Tang et al. 2010] and is also supported by empirical research in well-established online markets [Clay et al. 2002; Brynjolfsson et al. 2003; Baye et al. 2004; Baye and Morgan 2006]. It is commonly explained by frequent price changes due to dynamic pricing (that go beyond the platform level) [Kephart et al. 2000; Jumadinova and Dasgupta 2008] and the use of different sales strategies (e.g., "hit and run" [Baye et al. 2004; Jumadinova and Dasgupta 2008]).

[3]The time cost for a platform to query a seller is denoted c_i^p.

Subscription. The buyer pays a one-time subscription fee, c_{sub}, to use the platform, thereby gaining access to the set S' of listed sellers. This allows the buyer to query any seller $s_i \in S'$, wich incurs a cost c_i^t, and receive its actual price.

Sponsored. The platform is sponsored by advertisers and gains a payment ad_{rev} for each ad presented to the buyer. The buyer uses the platform for free, i.e., from its point of view this scheme is a specific variant of the Subscription scheme with $c_{sub} = 0$. The number of ads presented to a buyer is linear in the number of sellers queried by that buyer, as the latter measure is correlated with the amount of time the buyer spends using the platform and the number of pages it visits.

Per-click. The buyer can see the properties of all listed sellers (price distribution and querying cost), however in order to query seller $s_i \in S'$ (e.g., in order to reveal its identity and contact it) the buyer needs to pay the platform a payment c_c (in addition to incurring the appropriate cost c_i^t).

Full Service. The buyer pays the platform a fixed payment, c_{usage}, and in exchange the platform queries all listed sellers (incurring a cost c_i^p for each seller s_i queried), returning a list of actual prices to the buyer.

These four service schemes are not limited to eCommerce only. Taking dating platforms as an example, the Subscription scheme is the commonly used one. Still, there are also dating platforms that use the Per-click scheme, e.g., enabling one to see all profiles, however providing the communication details only for a fee. Some dating platforms offer the services for free however use ads. Finally, the Full Service scheme in that domain can represent the case of a “premium” subscription, where the user can get far more details, based on, for example, interviews the platforms carry out.

As demonstrated in the following sections, in all of the above mentioned service schemes the platform can potentially benefit from including in its listings only a subset of the full set of sellers S who can potentially be included in the platform. Therefore, the optimal (expected-profit-maximizing) strategy for each of the above service schemes should specify both the subset $S' \subseteq S$ of sellers to be included in the platform's listings and the charge to the buyers (whenever applicable). To further generalize the model, we introduce the operational expense ($d(S') \geq 0$) that the platform incurs for including and maintaining a set S' of sellers in its listing.

The model, which is described in the context of the procurement application domain, applies to any domain where platforms connect agents with opportunities, where each agent's goal can potentially be satisfied by any of the opportunities the platform lists and the value of each opportunity is a priori uncertain. Table I provides mappings of several other applications to the model.

3. ANALYSIS

The problem of choosing a service scheme and its parameters can be considered a Stackelberg game [Fudenberg and Tirole 1991; Conitzer and Sandholm 2006] where the platform is the first mover, that wants to maximize its expected profit with respect to the subset of sellers it lists, S', and the service terms it sets. We first analyze the optimal (expected-expense-minimizing) search strategy for individual agents of any type t, given S' and the platform's service scheme, and then show how the platform sets these two variables.

For illustration purposes, this section uses a synthetic toy environment with homogeneous buyers. Buyers' homogeneity suggests that they are all of the same type t and share the same private value R_b^t, distribution function of prices $f_i^t(q)$ and querying cost per seller c_i^t. The environment also includes a set of six sellers $S = \{s_1, ..., s_6\}$ that the platform may include in its listings. Each seller's price distribution is normal, with a

Table I. The mapping of different applications to the underlying model.

Domain	Platforms	Agent's Goal	Opportunities	Value	Query cost
Match-making	match.com, okcupid.com	Find a partner	Potential partners' profiles (partial and lack face-to-face impression)	Happiness and long term benefits	Time (and possibly cost) of contacting and dating the candidate
Car buying	autotrader. com	Buy a car	Dealers' and private owners' ads (partial impression and many details missing)	Value of the car minus its asking price	Contacting the owner for getting more details; test drive; mechanic check
Mortgage	mortgage-mediator. com	Get a good mortgage	Bank offers (missing/confusing terms and conditions)	Mortgage real value	Communicating to receive more info; expert advice
Real-estate	forrent.com	Rent an apartment	Apartments available on the market, cannot estimate the real value from an ad	Rent price	Time to estimate each apartment

different mean and standard deviation, defined over the interval [0,1] (trimmed and normalized). Sellers $s_2 - s_6$ share the same price probability distribution function (i.e., $f_2^t(q) = \ldots = f_6^t(q)$), while seller s_1 is associated with a different price distribution. We use S_i to denote the subset which includes s_1 and i additional sellers out of the 5 identical ones (hence, S_0 represents $\{s_1\}$ and $S_5 = S$). Similarly, $\overline{S}_i$ is used to denote the subset which includes only i sellers of the 5 identical ones . Therefore, the number of sets that the platform should consider (including S) is 11. The platform in this example is assumed to have no operational expenses, i.e., $d(S') = 0 \; \forall S' \subseteq S$.

3.1. Buyer's Strategy

The optimal strategy for the basic underlying search problem in our model can be found in classic economic search theory [Weitzman 1979]. Here, a searcher faces N opportunities, where each opportunity s_i is associated with a value q_i which is a priori unknown to her, i.e., only the distribution from which this value is drawn is known. The searcher's goal is to maximize her expected revenue (or, as in our case, minimize her expected expense) while the true value of any opportunity i can be revealed for a cost c_i. The solution for the problem in its expected-expense-minimization form is as follows: The searcher should assign each opportunity i a reservation value (i.e., a threshold) R_i, calculated as the solution for the following equation which represents the indifference of querying for the value of opportunity i or exploiting the best value revealed so far:

$$c_i = \int_{q_i=0}^{R_i} (R_i - q_i) f_i(q_i)\, dq_i \tag{1}$$

The searcher should query the opportunities according to their reservation value, in ascending order, and terminate once the best (lowest) value found so far is lower than the reservation value of the next opportunity according to that order. In the following paragraphs we assume, WLOG, that the sellers $\{s_1, \ldots, s_n\}$ are ordered according to their reservation value as calculated in (1). Based on the above principles we can extract the buyer's best-response to the platform's service terms with any of the methods and consequently the terms that will be set by the platform.

3.2. Subscription

From the buyer's point of view, upon paying a cost c_{sub}, it becomes a sunk cost, hence the buyer's optimal search strategy is reservation-value based as explained above. We use $E^t_{sub}(S', i, v)$ hereafter to denote the expected expense of a buyer of type t, upon approaching the ith seller in the available set S' (according to the optimal search sequence for agents of this type) and the best price found so far (when querying the previous $i-1$ sellers) is v. $E^t_{sub}(S', i, v)$ is calculated as the sum of the price that the buyer will eventually end up being charged and the accumulated costs incurred from this point on, given by:

$$E^t_{sub}(S', i, v) = \begin{cases} c^t_i + \int_{y=-\infty}^{\infty} E^t_{sub}(S', i+1, \min(v, y)) f^t_i(y) dy & v \geq r^t_i \bigwedge i < |S'| \\ c^t_i + \int_{y=-\infty}^{\infty} \min(v, y) f^t_i(y) dy & v \geq r^t_i \bigwedge i = |S'| \\ v & v < r^t_i \end{cases} \tag{2}$$

If the platform is used the buyer's overall expected expense is thus $E^t_{sub}(S', 1, R^t_b) + c_{sub}$. The buyer will use the platform only if $E^t_{sub}(S', 1, R^t_b) + c_{sub} \leq R^t_b$. Therefore, the platform's profit-maximizing fee for that buyer, if S' is used, denoted $c^t_{sub}(S')$, is $c^t_{sub}(S') = R^t_b - E^t_{sub}(S', 1, R^t_b)$, i.e., the one resulting in taking over the buyer's entire surplus, leaving it with an expected expense of R^t_b.

Based on the value of $c^t_{sub}(S')$, the platform's optimal fee can be calculated as follows. A buyer of type t will use the platform given a fee c_{sub} and a subset S' only if $c^t_{sub}(S') \geq c_{sub}$. Therefore the platform's expected profit if (S',c_{sub}) is used, denoted $V_{Subscription}(S', c_{sub})$, is:

$$V_{Subscription}(S', c_{sub}) = \sum_{c^t_{sub}(S') \geq c_{sub}} g(t) \cdot (c_{sub} - d(S')) \tag{3}$$

thus the platform needs to extract the pair (S',c_{sub}) that maximizes (3), denoted (S^*,c^*_{sub}).[4]

PROPOSITION 1. *When extracting the pair (S^*,c^*_{sub}) that maximizes the platform's expected profit, c^*_{sub} is necessarily one of the platform's maximizing fees, i.e., $c^*_{sub} \in \{c^t_{sub}(S') | S' \subseteq S\}$.*

PROOF. Assume otherwise, i.e., the use of some $(S', c_{sub} \notin \{c^t_{sub}(S') | S' \subseteq S\})$ maximizes the platform's expected profit. Incrementing c_{sub} to the closest higher value $c'_{sub} = \min\{c^t_{sub}(S') | c^t_{sub}(S') > c_{sub}, S' \subset S\}$ will not change the willingness of any of the buyers to use the platform — with c'_{sub} any buyer that used the platform's services when c_{sub} is charged will still use its service (to the same extent) and all others will still avoid using it. Therefore $V_{Subscription}(S', c'_{sub}) > V_{Subscription}(S', c_{sub})$ hence (S', c_{sub}) cannot be the expected-profit-maximizing strategy, contradicting the proof's assumption. □

The fact that the value of c^*_{sub} is necessarily one of the platform's maximizing fees $\{c^t_{sub}(S') | S' \subseteq S\}$ substantially simplifies the process of finding the optimal pair that will maximize (3) — for each subset $S' \subseteq S$ the platform needs to extract the expected-profit-maximizing fee, denoted $c_{sub}(S')$, as the one maximizing (3). Still, the number of subsets for which $c_{sub}(S')$ will need to be extracted is combinatorial. Since the focus of the paper is on the equilibrium analysis of the different schemes and the resulting preference of the buyers and platform, we place less importance on the computational aspect and the task of dealing with this complexity is left for future research.

Table II depicts the platform's profit when it uses its optimal fee according to the calculation procedure described above for the different subsets of the illustrative setting used in this section. The setting uses normal probability distribution functions with

[4]If the expected profit that uses (S^*,c^*_{sub}) is negative, the platform will opt not to offer its services in the first place under that service scheme.

Table II. Subscription: Subset's effect on c^*_{sub}

Subset	S_0	S_1	S_2	S_3	S_4	S_5	$\overline{S}_1$	$\overline{S}_2$	$\overline{S}_3$	$\overline{S}_4$	$\overline{S}_5$
c^*_{sub}	0.19	0.29	0.35	0.39	0.41	0.42	0.19	0.29	0.35	0.39	0.40

parameters ($\mu = 0.40, \sigma = 0.59$) for s_1 and ($\mu = 0.26, \sigma = 0.60$) for the remaining sellers. The buyers' private value is $R^t_b = 0.6$ and the cost of querying a seller is $c^t_i = 0.01\ \forall i$. From the table we observe that the platform's profit is maximized by disclosing all of the sellers known to the platform. This result is not surprising, since in our example the platform does not incur additional cost for providing more sellers. In fact Proposition 2 states that when using the Subscription scheme, if $d(S') = const$ then it is essentially optimal for the platform to disclose all sellers.

PROPOSITION 2. *If the platform's operational expense does not depend on the disclosed set, $d(S') = const\ \forall S'$ (e.g., if it uses a database that has already been generated), it is inevitably optimal for the platform to disclose all sellers, i.e., $V_{Subscription}(S, c_{sub}(S)) \geq V_{Subscription}(S', c_{sub}(S'))\ \forall S'$.*

PROOF. Assume otherwise, i.e., $\exists S'$, $V_{Subscription}(S', c_{sub}(S')) > V_{Subscription}(S, c_{sub}(S))$. Offering the terms $(S, c_{sub}(S'))$ instead of $(S', c_{sub}(S'))$, i.e., offering all sellers S while keeping the same value $c_{sub}(S')$ used for S', will not result in any additional cost for the platform, as $d(S') = const\ \forall S'$. Since $S' \subseteq S$ and the platform uses the same $c_{sub}(S')$, buyers' expected expense when offered S essentially does not increase, regardless of their types: $E^t_{sub}(S', 1, R^t_b) + c_{sub}(S') \geq E^t_{sub}(S, 1, R^t_b) + c_{sub}(S')$ $\forall t$. Therefore, with the new terms $(S, c_{sub}(S'))$ at least the same expected number of buyers will use the platform, regardless of their types, paying $c_{sub}(S')$ each, leading to $V_{Subscription}(S, c_{sub}(S')) \geq V_{Subscription}(S', c_{sub}(S'))$. Finally: $V_{Subscription}(S, c_{sub}(S)) \geq V_{Subscription}(S, c_{sub}(S')) \geq V_{Subscription}(S', c_{sub}(S'))$ hence contradicting the proof's initial assumption. $\square$

3.3. Sponsored

From the buyer's point of view the Sponsored case is a specific case of the Subscription scheme where $c_{sub} = 0$. Therefore its optimal strategy is inevitably reservation-value based where reservation values are calculated according to (1) and the expected expense is calculated according to (2). Using this scheme, the platform's expected profit is based on an external parameter (ad_{rev}) that does not affect the buyers' decision-making process. In order to extract the platform's profit in this case, if a subset S' is chosen, we need to calculate the expected number of sellers the buyers will query, denoted $\varphi(S')$. A type t buyer will end up querying the ith seller in its sequence (though not necessarily terminate its search process immediately) only if the minimum between the best value found so far (when querying the $i-1$ sellers characterized with lower reservation values according to (1)) and the value R^t_b is higher than its reservation value r^t_i, i.e., a probability of $\prod_{j=1}^{i-1}(1 - F^t_j(r^t_i))$ for $r^t_i < R^t_b$ and otherwise zero. Therefore, the expected number of sellers a type t buyer will query is $\sum_{r^t_i < R^t_b} \prod_{j=1}^{i-1}(1 - F^t_j(r^t_i))$, and the expected number of sellers a random buyer will query (and consequently the number of ads that will be presented, per buyer), is given by:

$$\varphi(S') = \sum_t g(t) \cdot \sum_{r^t_i < R^t_b} \prod_{j=1}^{i-1}(1 - F^t_j(r^t_i)) \tag{4}$$

The platform's expected profit when using a subset S' and receiving ad_{rev} for each ad displayed, denoted $V_{Sponsored}(S')$, is thus:

$$V_{Sponsored}(S') = \varphi(S') \cdot ad_{rev} - d(S') \tag{5}$$

and the optimal subset to be used can be calculated accordingly (i.e., $argmax_{S'}(V_{Sponsored}(S'))$).

Table III. Sponsored: Subset's effect on $\varphi(S')$.

Subset	S_0	S_1	S_2	S_3	S_4	S_5	$\overline{S}_1$	$\overline{S}_2$	$\overline{S}_3$	$\overline{S}_4$	$\overline{S}_5$
$\varphi(S')$	1.00	1.58	2.09	2.55	2.96	3.33	1.00	1.88	2.68	3.38	4.01

If $d(S') = const$, the platform's expected profit will be maximized by choosing the subset S' that results in the maximum expected number of queries. Table III illustrates that indeed the expected-number-of-queries-maximizing set is not necessarily the full set of sellers. It depicts the expected number of queries made by buyers for a setting with normal probability distribution functions with parameters $(\mu = 0.56, \sigma = 0.88)$ for s_1 and $(\mu = 0.66, \sigma = 0.19)$ for the remaining sellers. The buyers' private value is $R_b^t = 0.5$ and the query cost for each of the sellers is $c_i^t = 0.01 \forall i$. In this example the optimal subset the platform should disclose is $S^*=\overline{S}_5 = \{s_2 - s_6\}$.

3.4. Per-click

From the buyer's point of view, the Per-click method is a specific case of the Subscription scheme where $c_{sub} = 0$ and the cost incurred when querying seller s_i is $c_i' = c_i + c_c$ $\forall i$. Therefore, its optimal strategy is once again reservation-value based according to the principles of (1) and (2), using c_i' instead of c_i. The platform's expected profit in this case, denoted $V_{PerClick}(S', c_c)$, is the same as in (5), replacing ad_{rev} with c_c:

$$V_{PerClick}(S', c_c) = \varphi(S') \cdot c_c - d(S') \tag{6}$$

In order for the platform to find the optimal service terms (S', c_c) to be offered to the buyers, it needs to calculate for each subset S' the optimal c_c value and then pick the subset for which the maximum expected profit will be obtained (and its corresponding c_c value), i.e., $argmax_{(S',c_c)}\ V_{PerClick}(S', c_c)$.

Table IV. Per-click: Subset's effect on $\varphi(S')$ and $V_{PerClick}$(S',c_c)

Subset (S')	S_0	S_1	S_2	S_3	S_4	S_5	$\overline{S}_1$	$\overline{S}_2$	$\overline{S}_3$	$\overline{S}_4$	$\overline{S}_5$
c_c	0.16	0.09	0.08	0.08	0.07	0.07	0.05	0.05	0.05	0.05	0.05
$\varphi(S')$	1.00	1.28	1.49	1.63	1.75	1.80	1.00	1.71	2.21	2.57	2.82
$V_{PerClick}$(S',c_c)	0.16	0.12	0.12	0.13	0.13	0.13	0.05	0.08	0.11	0.12	0.14

Table IV depicts the platform's expected profit when disclosing different possible subsets to the buyers, using the optimal c_c value for each subset. The setting uses normal probability distribution functions with parameters $(\mu = 0.18, \sigma = 0.90)$ for s_1 and $(\mu = 0.97, \sigma = 0.43)$ for the remaining sellers. The buyers' private value is $R_b^t = 0.55$ and the cost of querying a seller is $c_i^t = 0.005\ \forall i$. In this case, the maximal platform's expected profit is achieved by disclosing the subset S_0 (i.e., seller s_1 only) and setting the per-click fee to $c_c = 0.16$. This is despite the fact that subset $\underline{S}_0$ is not the one that will maximize the expected number of queries (1 versus 2.82 with $\overline{S}_5$). As with the other service schemes analyzed above, when the population of buyers is heterogeneous, the extraction of the maximizing (S^*, c_c^*) pair requires weighing the expected profit from any buyer type according to its weight in the population for any possible pair (S', c_c).

3.5. Full Service

In this case, if a buyer of type t uses the platform, its expected expense, denoted $E^t_{FullService}$, is given by:

$$E^t_{FullService}(S') = c^t_{usage}(S') + \int_{q=0}^{R^t_b} q f_{S'}(q)dq + (1 - F_{S'}(R^t_b))R^t_b \tag{7}$$

where $f_{S'}(q)$ is the probability distribution function of the minimum price among those returned by the sellers in the set S'. The first term in (7) is the payment to the platform and the other two relate to the expected expense for the actual purchase (differentiating between the case where the price returned by the platform is lower than R^t_b and when it is greater (in which case buying elsewhere will incur a total expense of R^t_b)). The function $f_{S'}(q)$ can be calculated as the derivative of $F_{S'}(q)$, the probability that the minimum price returned will be equal to or lower than q:

$$f_{S'}(q) = \frac{\partial F_{S'}(q)}{\partial q} = \frac{\partial[1 - \prod_{s_i \in S'}(1 - F_i(q))]}{\partial q} \tag{8}$$

By substituting (8) in (7), and using integration by parts, we will obtain: [5]

$$E^t_{FullService}(S') = c^t_{usage}(S') + R^t_b - \int_{q=0}^{R^t_b} \big(1 - \prod_{s_i \in S'} (1 - F_i(q))\big) dq \tag{9}$$

The platform has to query each seller in the subset hence its expected profit, denoted $V_{FullService}(S', c^t_{usage})$, is the payment it charges the buyer for its service minus the operational expense it incurs for maintaining the database of sellers and querying them upon demand, i.e.:

$$V_{FullService}(S', c^t_{usage}(S')) = c^t_{usage}(S') - \sum_{s_i \in S'} c^p_i - d(S') \tag{10}$$

Similar to the Subscription case, the platform will set $c^t_{usage}(S')$ such that it takes over the entire buyer's surplus, i.e., $E^t_{FullService}(S') = R^t_b$, resulting in:

$$c^t_{usage}(S') = \int_{q=0}^{R^t_b} \big(1 - \prod_{i=1}^{S'}(1 - F_i(q))\big) dq \tag{11}$$

If the market is populated with heterogeneous buyers, then the platform's expected profit from using a fee c_{usage} is calculated according to the different agent types, i.e.:

$$V_{FullService}(S', c_{usage}) = \sum_{c^t_{usage}(S') \geq c_{usage}} g(t) \cdot (c_{usage} - \sum_{s_i \in S'} c_i - d(S')) \tag{12}$$

and the optimal service terms S^* and c^*_{usage} can be extracted by calculating the platform's expected profit for all combinations S'⊆S and their associated fees $c^t_{usage}(S')$. [6] Table V depicts the platform's profit when it uses its optimal fee according to the calculation procedure described above for the different subsets of the illustrative setting used in this section. The setting uses normal probability distribution functions with parameters $(\mu = 0.27, \sigma = 0.02)$ for s_1 and $(\mu = 0.18, \sigma = 0.43)$ for the remaining sellers.

---

[5] $\int_{q=0}^{R^t_b} q f_{S'}(q)dq = R^t_b F_{S'}(R^t_b) - \int_{q=0}^{R^t_b} (F_{S'}(q))dq$

[6] If the expected profit using (S^*, c^*_{usage}) is negative, the platform will opt not to offer its services in the first place using that service scheme.

Table V. Subscription: Subset's effect on $V_{FullService}(S', c_{usage})$

Subset	S_0	S_1	S_2	S_3	S_4	S_5	$\overline{S}_1$	$\overline{S}_2$	$\overline{S}_3$	$\overline{S}_4$	$\overline{S}_5$
$V_{FullService}(S', c_{usage})$	0.11	0.14	0.16	0.17	0.17	0.17	0.09	0.14	0.17	0.18	0.18

The buyers' private value is $R_b^t = 0.4$ and the cost of querying a seller is $c_i^p = 0.02\ \forall i$. In this example the optimal subset the platform will disclose is S^*=$\overline{S}_5 = \{s_2 - s_6\}$.

4. HOMOGENEOUS BUYERS

In many platforms, the population of buyers is quite homogeneous in terms of their perceived utility distribution, their opportunity-querying costs and their private value. For example, buyer agents in eCommerce platforms typically attempt to minimize their expected expense, hence for them the value distribution is actually the price distribution which characterizes each seller. Similarly, in this domain, buyers are typically associated with the same seller-querying cost, as their architecture and the communication protocols they use with sellers are identical. Finally, they all share the same private value, as they are all capable of using all of the other competing platforms and exploit the other opportunities one may find there. In this case, in contrast to the heterogeneous case, as will be demonstrated in the following section, several dominance relationships can be proven between the different service schemes, both from the buyers' and the platform's points of view.

THEOREM 1. *If buyers are fully homogeneous then: (a) from the buyer's point of view, the Per-click and Sponsored schemes (weakly) dominate (e.g., result in an equal or lower overall expected expense) the Full Service and Subscription schemes; and (b) from the platform's point of view: the Subscription scheme (weakly) dominate (e.g., result in an equal or greater expected profit) the Per-click scheme; if $c_i^p = c_i\ \forall i$ (i.e., if the platform's querying cost is similar to the buyers' querying cost) then the Subscription scheme (weakly) dominates the Full Service scheme; and if $c_i^p = 0\ \forall i$ (e.g., if the platform incurs a one-time setup fee and all subsequent queries are free), the Full Service scheme dominates the Subscription and Per-click schemes.*

PROOF. (a) From 3.2 and 3.5 we know that in both the Full Service and Subscription schemes the platform sets a usage fee such that the buyers' expected expense is exactly R_b, whereas with the other two schemes the buyers will end up with an expected expense equal to or lower than R_b (see 3.3 and 3.4).

(b) For the dominance of Subscription over Per-click, assume otherwise, i.e., by setting some terms (S', c_c) the platform can make a better profit compared to the optimal Subscription usage fee $(V_{PerClick}(S', c_c) > V_{Subscription}(S^*, c^*_{sub}))$. In this case, the platform can offer buyers the Subscription scheme with the set S' and its appropriate profit maximizing fee, i.e., $(S', c_{sub}(S'))$. The platform will now gain the difference between the buyers' private value and their expected expense. With the per-click fee, the platform will gain at most the difference between buyers' private value and their expected expense. Since in both cases the same set S' is used, and the buyers' "effective" querying cost is greater in the Per-click case $(c_c + c_i > c_i)$, and the buyers expected expense excluding the payments to the platform is greater in the Per-click case. Therefore the platform can gain the same or more (compared to Per-click) by offering the Subscription scheme with the same S' used with Per-click; hence the optimal subset used with Subscription, S^*, results in at least the same benefit. Therefore the Subscription scheme (weakly) dominates Per-click.

As for the dominance of Subscription over Full Service in cases where $c_i^p = c_i \forall i$ (i.e., when the same querying costs apply), here, once again, the platform can offer the Subscription scheme with the same set S^* which was found to be optimal for the Full Service scheme, $(S^*, c^*_{usage}) = \arg\max_{S', c_{usage}} (V_{FullService}(S', c_{usage}) | S' \subseteq S \wedge c_{usage} \in \mathbb{R})$. In this case, the buyer will be facing the exact same setting of sellers S' and their associated querying costs, except that with the Subscription scheme it can search sequentially, whereas with the Full Service scheme all opportunities are queried in parallel. Given similar conditions, sequential search is known to dominate parallel search, resulting in a lesser expected overall cost for the searcher [Morgan and Manning 1985]). Also, recall that with both the Subscription and Full service the platform takes over the buyer's surplus, which is the difference between R_b and the expected expense on the search. Therefore, the Subscription scheme will result in at least the same profit compared to the Full Service scheme.

The dominance of the Full Service scheme over Subscription and Per-click schemes is proven in a similar manner. In all three methods the profit of the platform is bounded by $R_b - E_S(q)$, where $E_S(q)$ is the expected minimum price when checking all possible sellers (calculated as $E_S(q) = \int min(R_b, q) f_S(q) dq$). $R_b - E_S(q)$ is in fact the maximum possible benefit (surplus) of the buyer from using the platform. As in this case its expected expense is minimized, since it uses the platform fully free and it can access all possible sellers. Therefore in the case of the Full Service scheme the platform, which does not incur any cost when checking prices, and can provide all prices in S and charge exactly $R_b - E_S(q)$ (as it can take over the entire buyer's surplus). In both other schemes the buyer's surplus is necessarily lower, regardless of the terms used, since the buyer incurs additional costs when using the platform. Hence the platform's expected profit in this case when using Full Service is greater than the one obtained with either of the two other methods. $\square$

We note that the above results do not depend on the function $d(S')$ as our proofs are based on evaluating the same subset (S') for the different schemes.

Based on Theorem 1, the buyer's decision, if given the option to choose between the different schemes when buyers are homogeneous, should be either Per-click or Sponsored. Between these two there is no dominance relationship from the buyer's point of view. For example, consider a setting where there is one seller which, with a probability of 0.25, will offer the product for free (e.g., via a promotion or by full rebate) and with a probability of 0.75 will sell the product at the full price of 1. All other sellers are characterized by a uniform distribution of prices between 0 and 1, and there is no limit as to the quantity of sellers that can be included in the platform. Also assume $c_i = 0.02\ \forall i$, $ad_{rev} = 0.01$, $R_b = 0.3$ and $d(S') = 0\ \forall S'$. In this case, if the Sponsored scheme is used, the platform will choose to exclude the first seller from its listings and keep only sellers of the other type. This is due to the fact that when only sellers of the second type are on the list, the reservation value used by the buyers is 0.2 (based on (1)); hence the expected number of sellers queried is 5 (resulting in an expected profit of 0.05 for the platform) and the buyer's expected expense is 0.2. If the first seller is included in the list, then that seller will be queried first (as its reservation value is 0.08) and the expected number of sellers queried is 4.75. With the Per-click scheme, the buyer's reservation value is $0.08 + 4c_c$ for the first seller and $\sqrt{2(0.02 + c_c)}$ for the others (according to (1)). The platform's expected profit with the Per-click scheme will thus be maximized (taking into consideration the buyer's private value) when the first seller is included and $c_c = 0.025$ is charged, resulting in a total expected profit of 0.0875. The alternative is to exclude the first seller and charge $c_c = 0.025$, resulting in expected profit of 0.0833. From the buyer's point of view, incurring an additional cost of $c_c = 0.025$ for each query, however with the option to query the more favorable seller, results in an expected expense of 0.18 which is better than its expected expense with the Sponsored scheme (0.2). Interestingly, in this case, the platform's expected profit when $c_c = 0.025$ and all sellers are included in the list is 0.0875, which is greater than its expected profit when the optimal subset is used in the Sponsored scheme (0.05).

The above example illustrates a scenario where both the platform and the buyer prefer a non-sponsored (i.e., a non-free) service scheme. This result and its implications are discussed in detail in the discussion and conclusions section.

The reverse preference is also possible. For example, if we simply remove the more favorable seller from the above example, the buyer's preference will be the Sponsored scheme. This can be further generalized to the following proposition.

PROPOSITION 3. *When sellers are fully homogeneous, i.e., characterized by the same querying cost c_i, distribution of values $f_i(q)$, and $d(S') = const$, the Sponsored scheme dominates all the other schemes from the buyers point of view.*

PROOF. From Theorem 1 we know that the buyers will prefer the Sponsored or Per-click schemes. Hence, we can prove that in this case the Sponsored scheme is preferred over the Per-click scheme. Since all sellers are identical, the platform's expected profit is necessarily maximized by listing the full set S, in both schemes. Therefore buyers encounter the same search problem, except that with the per-click scheme they also incur the per-click cost c_c; consequently they will prefer the Sponsored scheme where the search is free. $\square$

From the platform's point of view, no dominance relations between the Subscription and Sponsored schemes are proven in Theorem 1. Still, it is easy to prove that settings where one dominates the other exist.

PROPOSITION 4. *From the platform's point of view, a threshold exists such that for values exceeding ad_{rev} the Sponsored scheme dominates the Subscription scheme and for values below ad_{rev} the Subscription scheme dominates the Sponsored scheme.*

PROOF. The proof is based on the fact that the change in the value of ad_{rev} affects only the expected profit achieved with the Sponsored scheme, and that the latter monotonically increases as ad_{rev} increases. Now notice that for $ad_{rev} \rightarrow \infty$ the expected profit with the Sponsored scheme is greater (than with Subscription) and for $ad_{rev} \rightarrow 0$ vice-versa. Therefore undoubtedly there is a threshold value $\hat{ad_{rev}}$ above which the Sponsored scheme is preferred and below which the Subscription scheme is preferred. $\square$

5. HETEROGENEOUS BUYERS

The proofs given in the former section do not necessarily carry over to markets where buyers are heterogeneous, i.e., of different types. In fact, we can demonstrate that in such markets none of the schemes generally dominates any of the others from the platform's point of view.[7] In this section we demonstrate this phenomenon, showing that for each scheme we can find a setting in which that scheme performs best (i.e., yields the maximum expected profit for the platform). For this purpose we use two types of buyers that differ solely in their private value.

5.1. Subscription

We consider a setting with six possible sellers, $\{s_1, ..., s_6\}$, each associated with a normal distribution of prices (truncated and normalized between 0 and 1) according to Table VI. The revenue from presenting an ad is $ad_{rev} = 0.05$, the operational expense of the platform is $d(S') = 0 \; \forall S'$ and the querying cost is $c_i = c_i^p = 0.04 \; \forall i$. There are two

[7]Unlike with platforms, buyers' preference dominance is not well defined in the heterogeneous case, as it is possible that buyers of some types will prefer one scheme while buyers of other types will prefer another.

Table VI. Subscription: Sellers' population.

Seller	s_1	s_2	s_3	s_4	s_5	s_6
μ	0.27	0.18	0.09	0.06	0.21	0.03
σ	1.49	0.16	1.82	1.65	1.99	0.88

types of buyers: in the first (75% of the agents) the buyers are characterized by a private value $R_b^1 = 0.4$ and in the latter (25% of the agents) the buyers are characterized by $R_b^2 = 0.5$. In this case, the optimal pricing is: (a) to list sellers $\{s_1, s_2, s_3, s_4, s_5, s_6\}$ and charge $c_{sub} = 0.17$ in the Subscription scheme; (b) to list $\{s_1, s_3, s_5, s_6\}$ in the Sponsored scheme; (c) to list only s_2 and charge $c_c = 0.14$ in the Per-click scheme; and (d) to set $c_{usage} = 0.18$ for querying seller s_2 only in the Full Service scheme. The platform's expected profit in this case is maximized with the Subscription scheme (0.17, compared to 0.13, 0.14, and 0.13 with the Sponsored, Per-click and Full Service schemes, respectively).

5.2. Sponsored

When slightly changing the above example by increasing the revenue from each ad, namely from $ad_{rev} = 0.05$ to $ad_{rev} = 0.07$, we will obtain a different preference — the platform's expected profit with the Sponsored scheme will increase to 0.19 and the expected profit with all other schemes will remain the same, hence the Sponsored scheme will become the preferred method.

5.3. Per-click

Table VII. Per-click: Sellers' population.

Seller	s_1	s_2	s_3	s_4	s_5
μ	0.18	0.11	0.2	0.19	0.17
σ	0.01	0.09	0.37	1.24	0.98

We consider a setting with five possible sellers, $\{s_1, ..., s_5\}$, each associated with a normal distribution of prices (truncated and normalized between 0 and 1) according to Table VII. The revenue from presenting an ad is $ad_{rev} = 0.05$, the operational expense of the platform is $d(S') = 0\ \forall S'$ and the querying cost is $c_i^t = c_i^p = 0.3\ \forall i, t$. There are two types of buyers: in the first (60% of the agents) the buyers are characterized by a private value $R_b^1 = 0.3$ and in the second (40% of the agents) the buyers are characterized by $R_b^2 = 0.5$. In this case, the optimal pricing is: (a) to list sellers $\{s_1, s_2, s_3, s_4, s_5\}$ and charge $c_{sub} = 0.09$ in the Subscription scheme; (b) to list $\{s_1, s_2, s_3\}$ in the Sponsored scheme; (c) to list only s_4 and charge $c_c = 0.26$ in the Per-click scheme; and (d) to set $c_{usage} = 0.11$ for querying seller s_4 only in the Full Service scheme. The platform's expected profit in this case is maximized with the Per-click scheme (0.1, compared to 0.09, 0.09, and 0.03 with the Subscription, Sponsored and Full Service schemes, respectively).

5.4. Full Service

By slightly changing the example used for the Subscription and Sponsored schemes, setting $c_i^p = 0$ rather than $c_i^p = c_i^t\ \forall i, t$, a different preference is revealed. In this case the platform's expected profit with the Full Service scheme increases to 0.21 and does not change for the other schemes, hence the Full Service scheme is the preferred method.

6. RELATED WORK

The role of platforms as mediators (or middle agents, brokers and matchmakers [Klusch and Sycara 2001]) in two-sided markets has been extensively investigated in recent years both in multi-agent [Decker et al. 1997; Sarne et al. 2007; Evans 2012] and in general [Armstrong 2006; Rysman 2009; Weyl 2010; Hagiu and Wright 2011] literature. The platform offers many advantages, especially in distributed MAS environments where immediate reliable information about the different opportunities available to the agents is not publicly available. As such, these platforms have proved to fulfill a major role in MAS (e.g., for comparison shopping, dating, cars trading, mortgage and negotiation) and 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 consumers buying experience [Ripper et al. 2000].

One of the main questions investigated within this context was how should platforms price their information services, i.e., who and what fees to charge [Evans and Schmalensee 2007; Weyl 2010]. As such, many service pricing models have been introduced and supported, e.g., charging buyers, using sponsored ads, and charging sellers [Sarne 2013; Chhabra et al. 2014]. The idea that a platform will charge only one side in a two-sided market while the other group is allowed to use the platform for free can be explained, in general, by intense competition among the players of that group (e.g., directories such as "yellow pages" that are supplied to readers for free) [Armstrong 2006] and as an attempt to attract elastic consumers and, as a result, obtain higher prices or more participation on the other side [Rysman 2009]. Unlike our work, prior research has concentrated on how platforms should price their services without considering the nature of their listings, and thus without taking into consideration the option of limiting or selectively disclosing the information available to them.

The idea of selective disclosure, which is one of the cornerstones of our analysis, where only a partial set of the information/opportunities known to an agent is revealed when interacting with other agents, has been widely used in psychology and behavioral economics [Thaler and Sunstein 2008; Iyengar 2010] and in multi-agent literature [Sarne et al. 2011; Peled et al. 2014]. For example, selective disclosure of information has been shown to be beneficial in adversarial settings, such as in the case of comparison shopping agents (CSAs). Through the use of selective disclosure, CSAs can influence their human users not to query additional, competing, CSAs [Hajaj et al. 2013, 2014].[8] Similarly, it has been shown that for advice-generating agents in settings where the advisor's interests may conflict with the advisee's interests, selective disclosure of information can improve the advising-agent's expected benefit [Azaria et al. 2012]. Yet, all these works consider users to be human, rather than fully rational and computationally unbounded agents. Therefore they predominantly use people's inherent irrationality as a basis for their proposed methods. In this paper, on the other hand, users are assumed to be fully rational, and the benefit in selective information disclosure derives from the equilibrium analysis.

Finally, we note that the user's strategy in this paper is fully based on economic search theories [Morgan and Manning 1985; McMillan and Rothschild 1994; Sarne and Kraus 2003; Kang 2005; Sarne and Kraus 2008; Grosfeld-Nir et al. 2009]. Nevertheless, despite the rich literature in this field, no previous work has addressed adversarial settings of the types analyzed in this paper.

[8]Within this context a CSA can be seen as a platform that applies the Full Service scheme.

7. DISCUSSION AND CONCLUSIONS

The analysis and numerical examples given in the paper support the claim that platforms can gain from not necessarily including all of the opportunities that they can potentially provide to their users. This was shown for all of the service schemes analyzed and holds even if the platform's operational expense does not depend whatsoever on the set of opportunities listed (except for the Subscription scheme for which we prove otherwise). This result is especially important in light of the many works in the area of two-sided markets which advocate that platforms should try to increase the number and richness of opportunities they list [Rochet and Tirole 2003]. The difference is of course that works of the latter type do not take into consideration the effect of the users' search behaviors on the platform's gain (e.g., in the form of the number of opportunities explored).

The implications of potentially benefiting from not listing all possible opportunities are twofold. First, platforms should not focus solely on price settings, i.e., taking the information to be listed as given, but rather optimize both based on the fees charged and the set of opportunities to be listed. The analysis given in the paper supports such optimization as it supplies the appropriate equations based on which platforms can calculate their expected-maximizing strategy (set of opportunities to list and the corresponding usage fees) for four highly common service schemes that have been widely discussed in literature. For the homogeneous case, which, as contended in Section 4, is highly common in domains such as eCommerce, several dominance relationships between the different schemes are proven, directing the platform to a default scheme in various cases. For the heterogeneous case, the platform must calculate the optimal service scheme terms for each specific scheme or on a case-by-case basis, since, as demonstrated in Section 5, none of the schemes generally dominates the others.

Second, it is possible, as demonstrated in the paper, that users will prefer costly platforms over "free" (sponsored) platforms. This preference has been justified in prior literature by arguments such as consumers' willingness to pay a premium for an ad-free experience [Cho and Cheon 2004], the benefit of having more reliable (i.e., not obligated to their sponsors) platforms [Belleflamme and Peitz 2010], and using a cooperative platform rather than a self-interested one [Sarne 2013]. Our paper, on the other hand, shows that such preference derives from the differences in the set of opportunities that the platform will choose to list in each of the cases, and the resulting Stackelberg equilibrium. Furthermore, we show that the preference of the costly alternative may hold, for a given setting, both for the platform and its users, even when the two are fully strategic. This latter result is very significant as it supplies motivation for the transition of platforms from the commonly used Sponsored model to user-charging business models without risking a decline in users' satisfaction.

We envision several important extensions of the model analyzed in the paper that can further contribute to the understanding of platform-based MAS dynamics. First our model assumes that opportunities values are exogenously set. In this sense, an important extension would be the integration of strategic behavior in the equilibrium analysis from the opportunities' side (e.g., where opportunities represent sellers). Second, we would like to consider a richer equilibrium analysis in which the platform can also choose to charge the owners of opportunities for being listed (or per transaction).

REFERENCES

M. Armstrong. 2006. Competition in Two-Sided Markets. *RAND Journal of Economics* 37, 3 (2006), 668–691.

A. Azaria, Z. Rabinovich, S. Kraus, C. V. Goldman, and O. Tsimhoni. 2012. Giving advice to people in path selection problems. In *Proc. of AAMAS*. 459–466.

Y. Bakos. 1997. Reducing Buyer Search Costs: Implications for Electronic Marketplaces. *Management Science* 42(12) (1997), 1676–1692.

M. R. Baye and J. Morgan. 2006. Persistent Price Dispersion in Online Markets. In *The New Economy and Beyond*, D. Jansen (Ed.). Edward Elgar Press.

M. R. Baye, J. Morgan, and P. Scholten. 2004. Temporal Price Dispersion: Evidence from an Online Consumer Electronics Market. *Interactive Marketing* 18(4) (2004), 101–115.

P. Belleflamme and M. Peitz. 2010. Platform competition and seller investment incentives. *European Economic Review* 54, 8 (2010), 1059–1076.

E. Brynjolfsson, E. Hu, and M. Smith. 2003. Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Bookseller. *Management Science* 49(11) (2003), 1580–1596.

M. Chhabra, S. Das, and D. Sarne. 2014. Expert-mediated sequential search. *European Journal of Operational Research* 234, 3 (2014), 861–873.

C. Cho and J. Cheon. 2004. Why do people avoid advertising on the internet? *Journal of advertising* 33, 4 (2004), 89–97.

K. Clay, R. Krishnan, E. Wolff, and D. Fernandes. 2002. Retail Strategies on the Web: Price and Non-price Competition in the Online Book Industry. *Journal of Industrial Economics* 50 (2002), 351–367.

V. Conitzer and T. Sandholm. 2006. Computing the optimal strategy to commit to. In *Proc. of ACM EC*. 82–90.

K. Decker, K. Sycara, and M. Williamson. 1997. Middle-agents for the internet. In *IJCAI (1)*. 578–583.

D. S. Evans. 2012. Governing Bad Behavior by Users of Multi-Sided Platforms. *Berkeley Technology Law Journal* 2, 27 (2012).

D. S. Evans and R. Schmalensee. 2007. The industrial organization of markets with two-sided platforms. (2007).

D. Fudenberg and J. Tirole. 1991. Game theory. (1991).

X. Geng, M. B. Stinchcombe, and A. B. Whinston. 2005. Bundling information goods of decreasing value. *Management Science* 51, 4 (2005), 662–667.

A. Grosfeld-Nir, D. Sarne, and I. Spiegler. 2009. Modeling the search for the least costly opportunity. *EJOR* 197 (2009), 667–674.

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

A. Hagiu and J. Wright. 2011. *Multi-sided platforms*. Harvard Business School.

C. Hajaj, N. Hazon, and D. Sarne. 2014. Ordering Effects and Belief Adjustment in the Use of Comparison Shopping Agents. In *Proc. of AAAI*.

C. Hajaj, N. Hazon, D. Sarne, and A. Elmalech. 2013. Search More, Disclose Less. In *Proc. of AAAI*.

M. He, N. R. Jennings, and H. Leung. 2003. On agent-mediated electronic commerce. *TKDE* 15(4) (2003), 985–1003.

S. Iyengar. 2010. *The Art of Choosing*. Twelve.

M. Janssen, J. L. Moraga-Gonzalez, and M. R. Wildenbeest. 2005. Truly costly sequential search and oligopolistic pricing. *IJIO* 23, 5-6 (2005), 451–466.

J. Jumadinova and P. Dasgupta. 2008. Firefly-Inspired Synchronization for Improved Dynamic Pricing in Online Markets. In *Proc. of SASO*. 403–412.

B.-K. Kang. 2005. Optimal stopping problem with double reservation value property. *European Journal of Operational Research* 165, 3 (2005), 765–785.

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

M. Klusch and K. Sycara. 2001. Brokering and Matchmaking for Coordination of Agent Societies: A Survey. (2001).

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

P. M. Markopoulos and J. O. Kephart. 2002. How valuable are shopbots?. In *Proc. of AAMAS*. 1009–1016.

J. McMillan and M. Rothschild. 1994. Search. In *Handbook of Game Theory with Economic Applications*, R. Aumann and S. Hart (Eds.). 905–927.

P. Morgan and R. Manning. 1985. Optimal search. *Econometrica* 53, 4 (1985), 923–944.

Y. Nahum, D. Sarne, S. Das, and O. Shehory. 2012. Two-sided search with experts. In *Proc. of EC*. 754–771.

J. Niu, K. Cai, S. Parsons, E. Gerding, P. McBurney, T. Moyaux, S. Phelps, and D. Shield. 2008. JCAT: A platform for the TAC Market Design Competition. In *Proc. of AAMAS*. 1649–1650.

N. Peled, Y. Gal, and S. Kraus. 2014. A Study of Computational and Human Strategies in Revelation Games. *Proc. of AAMS* (2014).

P. S. Ripper, M. F. Fontouraand A. M. Neto, and C. J. P. De Lucena. 2000. V-Market: A framework for agent e-commerce systems. *World Wide Web* 3, 1 (2000), 43–52.

J. C. Rochet and J. Tirole. 2003. Platform competition in two-sided markets. *Journal of the European Economic Association* 1, 4 (2003), 990–1029.

M. Rysman. 2009. The Economics of Two-Sided Markets. *Journal of Economic Perspectives* 23, 3 (2009), 125–43.

D. Sarne. 2013. Competitive Shopbots-Mediated Markets. *ACM Trans. Econ. Comput.* 1, 3, Article 17 (2013), 41 pages.

D. Sarne, A. Elmalech, B. J. Grosz, and M. Geva. 2011. Less is more: restructuring decisions to improve agent search. In *Proc. of AAMAS*. 431–438.

D. Sarne and S. Kraus. 2003. The search for coalition formation in costly environments. In *Cooperative Information Agents VII*. Springer, 117–136.

D. Sarne and S. Kraus. 2008. Managing parallel inquiries in agents' two-sided search. *Artificial Intelligence* 172, 4-5 (2008), 541–569.

D. Sarne, S. Kraus, and T. Ito. 2007. Scaling-up shopbots: a dynamic allocation-based approach. In *Proc. of AAMAS*. 256.

A. Sundararajan. 2004. Nonlinear pricing of information goods. *Management Science* 50, 12 (2004), 1660–1673.

Z. Tang, M. D. Smith, and A. Montgomery. 2010. The impact of shopbot use on prices and price dispersion: Evidence from online book retailing. *International Journal of Industrial Organization* 28, 6 (2010), 579–590.

R. H. Thaler and C. R. Sunstein. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

R. Waldeck. 2008. Search and price competition. *JEBO* 66, 2 (2008), 347–357.

M. L Weitzman. 1979. Optimal Search for the Best Alternative. *Econometrica* 47, 3 (1979), 641–54.

E Glen Weyl. 2010. A price theory of multi-sided platforms. *The American Economic Review* (2010), 1642–1672.

X. Xing, Z. Yang, and F. Tang. 2006. A comparison of time-varying online price and price dispersion between multichannel and dotcom DVD retailers. *Journal of Interactive Marketing* 20, 2 (2006), 3–20.