

Multi-Keyword Sponsored Search

Peerapong Dhangwatnotai*
Department of Computer Science, Stanford University
460 Gates Building, 353 Serra Mall, Stanford, CA
pdh@cs.stanford.edu

ABSTRACT

We investigate search engines' mechanism for allocating impressions generated from different search terms. This mechanism is equivalent to running an independent GSP auction for each search term *only* when the number of search terms is small. In practice, the number of search terms is so large that an advertiser cannot possibly communicate to the search engine all the GSP auctions that he wishes to participate in. For example, a travel agency is interested in all search terms pertaining to flight, including "flight to boston", "ticket to SFO", "cheap airfare", etc. Therefore, the search engine introduces broad match keywords as a bidding language that allows an advertiser to submit a bid for multiple GSP auctions at once. However, with broad match keywords, the GSP auctions are no longer independent, i.e. an advertiser's bid in one auction may depend on his bid in another auction.

We propose the broad match mechanism as a model that captures this aspect of the multi-keyword sponsored search mechanism. We study the performance of this mechanism under the price of anarchy (POA) framework. We identify two properties of broad match keywords, namely expressiveness and homogeneity, that characterize the POA and we prove almost tight bounds on the POA. The bounds allow us to explore trade-offs between the two properties. We introduce the exact-match-only mechanism whose performance, when compared to that of broad match mechanisms, gives us an insight into the net benefit of broad match keywords. The broad match mechanism can also be viewed as a mechanism that copes with severe communication constraint i.e. the valuation of an advertiser is described by many more numbers than the search engine can solicit.

Categories and Subject Descriptors

F.0 [Theory of Computation]: General; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

*Supported in part by the ONR Young Investigator Award of Tim Roughgarden.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'11, June 5–9, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0261-6/11/06 ...\$10.00.

General Terms

Economics, Theory

Keywords

Sponsored search, mechanism design, auctions

1. INTRODUCTION

It is important to distinguish *search terms* from *keywords* when we discuss search advertising with multiple search terms. When a user enters a *search term* into a search engine, the search engine has an opportunity to show an ad to the user by putting it alongside the user's search result. This opportunity is called an *impression*. Advertisers want to buy impressions because they are a source of traffic to their websites. Therefore, the commodity that the search engine sells to the advertisers is impression.¹ However, each advertiser values the impressions generated from certain search terms more than others.² Therefore, the search engine sells impressions by running a separate auction for each search term. We call the mechanism that allocates impressions generated from *all* different search terms the *multi-keyword sponsored search mechanism* (MKSS), for reasons that will become clear later. The auction that the MKSS mechanism uses to allocate impressions generated from a *specific* search term is called the *GSP auction*. The big problem with this approach is that the number of search terms that an advertiser wants is enormous e.g. every phrase that contains "android". It will take practically forever for an advertiser to communicate all the desired search terms to the search engine.

The search engine introduces keywords to deal with the large number of search terms. A keyword can be associated with multiple search terms. When an advertiser bids using a keyword, he automatically bids on all search terms associated with it. We can think of the keywords as a bidding language that allows an advertiser to bid on many search terms at once. There are two important types of keywords. A *broad match* keyword "b" is associated with all search terms deemed relevant to b by the search engine. For example, broad match keyword "apple" is associated with search

¹Search engines may sell impressions by clicks. As long as the number of clicks per impression (CTR) for each advertiser is fixed, clicks and impressions are just different units of the same commodity.

²For example, Disney values impressions generated from the search term "disney" much more than those generated from "grocery" (because users who search for "disney" are more likely to follow the ad to Disney's website).

terms “apple”, “apple cider”, “apple iphone”, “fruit”, etc. In contrast, an *exact match* keyword “b” is associated with the only search term “b”. One benefit of a broad match keyword is that it allows advertisers to bid on rare search terms e.g. search terms with more than 3 words. Since there are many of them, the set of such search terms cannot be communicated efficiently without broad match keywords and some would have gone wasted because of the lack of bidders.

Our understanding of the GSP auction does not extend easily to the multi-keyword sponsored search mechanism. A number of papers studied the GSP auction [7, 4, 6, 5, 20, 14, 17, 10]. All of them assume the GSP auction is run in isolation or, in other words, there is only one search term in the world. Without broad match keywords, the sponsored search mechanism is equivalent to multiple of those GSP auctions running independently. With broad match keywords, such equivalence no longer holds.³ The performance metric that we use to study sponsored search is welfare. In this setting, it corresponds to the sum of each advertiser’s value for his allocated impressions.⁴ [4] showed that, under some assumptions, the welfare of every equilibrium of every GSP auction is close to the optimal. It would be nice to have an analogous theorem for the multi-keyword sponsored search mechanism.

Broad match keywords play an important role in the performance of the sponsored search mechanism. Broad match keywords can improve welfare because they allow advertisers to obtain impressions that would otherwise be wasted. Broad match keywords also have their drawbacks. For example, if we search for “blackberry preserve” on a popular search engine, we may find an ad from Research in Motion, the company that makes smartphones, at the top position. The welfare would have been greater if the search engine had allocated this impression to a jam company instead (replaced Research in Motion’s ad with the jam company’s ad).⁵ We suspect this phenomenon occurs because Research in Motion places a bid on the broad match keyword “blackberry” which is associated with the search term “blackberry preserve”. In short, broad match keywords can cause the mechanism to misallocate an impression, taking away an impression from the advertiser who values it the most and giving it to another advertiser who does not value it at all.

The communication constraint necessitates the use of broad match keywords. Since the GSP auction is not truthful, advertisers may need to continually adjust their bids to respond to others’ bids. An advertiser with limited resources may find the task of adjusting one million bids every hour daunting. On the search engine side, hundreds of thousands of advertisers collectively submitting 10000000000 new bids every hour can put too much load on the server. Search engines typically place a limit on the number of keywords per advertiser, which is the same as a limit on the number of bids. We can also view this as an implementation constraint. Here’s an excerpt from AdWords’ documentation:

³When an advertiser uses a broad match keyword, he is forced to use the same bid for all GSP auctions whose search terms are associated with that keyword.

⁴Welfare also correlates with user experience. Users tend to have good experience when the right impression is matched to the right advertiser.

⁵We have been observing the sponsored search result and notice that indeed jam companies lose some impressions with this search term to Research in Motion.

To ensure that accounts do not place unnecessary load on our advertising servers, we ask that you keep the overall number of keywords in your account to a manageable size (typically 50,000 keywords or less).

With this constraint, broad match keywords become necessary for the advertisers and for any model of the multi-keyword sponsored search mechanism.

Our model, called the broad match mechanism, is designed to capture the effects of broad match keywords on advertiser behavior and the performance of the MKSS mechanism. In our model, each advertiser can submit a bid for each predefined keyword but the total number of keyword bids must not exceed κ . Then each advertiser’s keyword bids are translated into bids on search terms. A GSP (Vickrey) auction⁶ is run for each search term to allocate the corresponding impressions. To better illustrate the effects of broad match keywords, we assume that the GSP auction has only 1 slot and, while keeping the allocation rule the same, the search engine charges payment per impression. We believe these two details are not important for understanding the effects of broad match keywords and can only serve as a distraction.

There is also a theoretical motivation for studying the broad match mechanism. It is an example of a mechanism with severely constrained communication. Each advertiser has a different valuation for each search term. Since it is impossible for advertisers to communicate their valuations for every search term, it is impossible for the mechanism to obtain enough information to maximize welfare. It is also a situation when the revelation principle does not apply. It is a situation where designing the bidding language is as important as designing the mechanism. We attempt to answer the question: What makes good broad match keywords? We identify two properties of broad match keywords that characterize the upper bound on the price of anarchy of the mechanism. The two properties are expressiveness and homogeneity. They are in conflict with each other and thus it is difficult for the mechanism designer to obtain both at the same time. Therefore, the best mechanism needs to strike a balance between the two.

Our Results .

- We propose a model of the multi-keyword sponsored search mechanism, called the broad match mechanism. The model captures the interactions between advertisers in the presence of broad match keywords. It is also an example of a mechanism that copes with severe communication constraint.
- We prove almost-tight worst-case bounds on the welfare of equilibria of broad match mechanisms. We assume that the advertisers have additive valuations. While it is impossible to prove any bound on every equilibrium, even when there is only one search term⁷, our bounds apply to “plausible” equilibria. Implausible equilibria are equilibria in which an advertiser

⁶The GSP auction is equivalent to the Vickrey auction if there is only 1 slot.

⁷Single-item Vickrey auction is a special case of the broad match mechanism. Then consider two bidders with valuations 1 and 0 who bid 0 and 1 respectively.

plays a weakly dominated strategy. An analogy for the Vickrey auction would be when a bidder bids higher than his valuation. We identify a particular set of undominated strategies for each bidder. The bounds are parameterized by expressiveness (α) and homogeneity (c). We prove that the welfare of every pure Nash equilibrium, in which bidders play undominated strategies, in fully-expressive broad match mechanisms is at least $\frac{1}{(c^2+c)}$ times that of the welfare-maximizing allocation. This almost matches our lower bound of $\frac{4}{c^2}$ in the worst case.

- We add a flexibility for the mechanism designer to trade off expressiveness and homogeneity of the broad match mechanism. We prove that the welfare of every pure Nash equilibrium, in which bidders play undominated strategies, in broad match mechanisms with parameters (α, c) is at least $\frac{\alpha}{(c^2+c)}$ times that of the welfare-maximizing allocation.
- To weigh the net benefit of broad match keywords, we introduce the exact-match-only (EMO) mechanism, a mechanism that does not have any broad match keyword. A comparison of the price of anarchy between a broad match mechanism and an EMO mechanism can be used as a quick test to show how much broad match keywords benefit the multi-keyword sponsored search mechanism.
- We extend our bound on the price of anarchy of the EMO mechanism to subadditive valuation. EMO mechanisms corresponds to combinatorial auctions with item-bidding in [3, 15]. However, our mechanisms have an extra constraint on communication complexity. Communication constraint is parameterized by γ , the number of items an advertiser is interested in divided by the number of items that κ keywords can cover. We prove that the upper bound on the price of anarchy of EMO mechanisms degrades gracefully with parameter γ .

Related Work.

Even-Dar et al. studied the algorithmic problem of computing the utility-maximizing bid for a single bidder when there are broad match keywords [12]. Mahdian and Wang observed that after the advertisers commit to a set of keywords (but not bid amounts), the sponsored search mechanism with overlapping broad match keywords may not have an equilibrium [13]. They also gave an algorithm for computing a set of non-overlapping broad match keywords that approximate the welfare of the optimal non-overlapping broad match keywords in a setting where advertisers are not allowed to bid on exact match keywords. The main difference from our work is they assume the algorithm knows the valuations of advertisers. Kumar et al. studied the problem of allocating a large number of distinct items (impressions) to bidders whose valuations are drawn from distributions [18]. The realized valuation is more complex than what the auctioneer can elicit but the distributions are known. Singh et al. studied the setting in which the auctioneer can withhold information on the presence of some search terms [19]. They don't consider any communication restriction as each bidder can submit a different bid for every search term.

Our approach of using the price of anarchy to evaluate mechanisms without dominant strategies is similar to [3, 15, 4]. [3, 15] studied combinatorial auctions (CA) with item-bidding, which run a separate auction for each item and allow bidders to bid in every auction. The mechanism uses less communication than the complexity of bidder valuation which is subadditive in [3]. However, the number of items in our setting is much larger and so large that bidders cannot participate in every auction that they are interested in. Our work also introduces broad match keywords into the bidding language and investigates the problem of designing the mechanism with the best equilibria.

As mentioned earlier, many works focus on single search-term GSP auctions. For example, [4, 7, 16, 20, 14, 2, 17] study the revenue or the efficiency of the equilibria of single search-term GSP auctions. [5, 6] designed truthful auctions for a single search term. It is impossible to have truthful auctions in our setting with many search terms because of the communication constraint. There are models and results for multi-search-term algorithms/mechanisms but they don't consider communication requirements and thus effectively don't have keywords [8, 9, 11]. Ghosh et al. studied the algorithmic problem of bundling search terms to maximize revenue and welfare [1].

2. THE MODEL

Assume all impressions generated from the same search term are identical. We denote the group of impressions generated from the same search term j by *query* j .⁸ We consider a one-period model i.e. the search engine has a set M of m queries available for sale in a given period (the number of impressions in each query is known to all parties). The number of queries m is finite but very large. There are n advertisers. Each advertiser i has a *valuation* $v_i : 2^m \rightarrow \mathbb{R}^+$ that describes his value for each subset of the queries.⁹ We also define $v_i(j) = v_i(\{j\})$. We assume $v_i(\emptyset) = 0$ and $v_i(S) \leq v_i(T)$ for all $S \subseteq T$. The advertisers' valuations are additive across queries i.e. $v_i(A) = \sum_{j \in A} v_i(j)$. We will point out when our results also apply to subadditive valuations. One can see that this model is a type of combinatorial auction, which may be of independent interest. We will now use "advertiser" and "bidder" interchangeably.

Bidder i 's *support* \mathcal{S}_i is defined by $\mathcal{S}_i = \{j | v_i(j) > 0\}$, i.e. the set of queries for which i has positive valuation. One of the main challenges of sponsored search is some $|\mathcal{S}_i|$'s are so large that it is impossible for the bidders to communicate all $v_i(j)$'s in their supports. An allocation $X = \{X_1, X_2, \dots, X_n\}$ is an assignment of queries to bidders (X_i to bidder i) such that each query is assigned to at most one bidder, i.e. $X_i \subseteq M$ and $X_i \cap X_j = \emptyset$ for all i, j . The *welfare* $SW(X)$ of an allocation X is defined as $\sum_i v_i(X_i)$. We will examine the welfare of a few mechanisms under reasonable communication requirements (which will be defined later on).

⁸Since we consider a setting with one slot and no budget, an advertiser who wins an auction for search term j will win all impressions with that search term.

⁹If we assume the number of impressions in query j (in the given period) is η , an advertiser whose value-per-click is c and click-through-rate is μ will have valuation $v_i(\{j\}) = \eta\mu c$.




Figure 1: An example of a broad match mechanism. The queries (a,b,c,d) are shown as circles on the bottom and the keywords (K,L,P) are shown as squares. The broad match keywords are K and L. The exact match keyword is P. An advertiser who places a bid of 2 on K and 3 on L automatically bids 2 on queries a and 3 on queries b,c,d.

2.1 The Broad Match Mechanism

We propose a model of the multi-keyword sponsored search mechanism employed by popular search engines. The broad match mechanism deals with the enormous number of queries by defining sets over them. Each set is called a *keyword*. For example, keyword “apple” contains queries “apple”, “apple cider”, “apple iphone”, etc. When an advertiser bids on a keyword, he automatically bids on all queries contained in that keyword (using the same bid amount). There are three types of keywords: exact match, phrase match and broad match. An exact match keyword b contains the only query “ b ”. A phrase match keyword b contains all queries containing b , e.g. “ $a b c$ ”. A broad match keyword b contains all queries deemed relevant to b by the search engine (they can be anything). Since phrase match is just a special case of broad match, we will lump them together and call them broad match from now on. Figure 1 shows an example of a mechanism with multiple keywords.

Using set notation, we denote the set of queries that keyword A contains by A . To control communication complexity, each bidder is allowed to bid on at most κ keywords (all types combined). Let $\mathbf{b} = (b_1, b_2, \dots, b_n)$ denote the bid profile where b_i is bidder i ’s bid vector. Denote bidder i ’s bid on keyword A by $b_i(A)$ and, overloading the notation, bidder i ’s bid on query j by $b_i(j)$, interpreting $b_i(A) = 0$ for keywords without a bid. We allow keywords to be overlapping. To resolve the bid on query j for bidder i , $b_i(j) = \max_{A:j \in A} b_i(A)$, the maximum of i ’s bids over all keywords containing the query j . For each query, the mechanism runs a Vickrey auction i.e. allocates query j to the highest bidder i at a price $p_i(j)$ equal to the second-highest bid $\max_{k \neq i} b_k(j)$ for that query. The following summarizes the broad match mechanism:

1. Each bidder i submits bid b_i on keywords upfront. The number of keywords A such that $b_i(A) > 0$ must be at most κ .
2. Each bidder’s bid is translated into bids on queries using the formula $b_i(j) = \max_{A:j \in A} b_i(A)$.
3. For each query j , allocate the query to the bidder with the highest $b_i(j)$ and charge him the price $p_i(j) = \max_{k \neq i} b_k(j)$.

Since this broad match mechanism charges a price $p_i(j)$ for the entire search term, it is called the per-search-term

model. One might want to consider a per-impression model that charges payment per impression. We present a per-impression model in section 7 and show that our analysis of this per-search-term model generalizes to the per-impression model.

Denote the queries allocated to bidder i by $X_i(\mathbf{b})$. Each bidder i wants to maximize his *utility*: $u_i(X_i(\mathbf{b}), b_{-i}) = \sum_{j \in X_i(\mathbf{b})} v_i(j) - p_i(j)$.

As mentioned in the introduction, no welfare guarantee is possible if a bidder uses an unreasonable bid. We refine the bidders’ strategy spaces using the concept of undominated strategy. We say bidder i ’s bid vector b_i *weakly dominates* b'_i if for all bid profiles b_{-i} of the other bidders, $u_i(X_i(b_i, b_{-i}), b_{-i}) \geq u_i(X_i(b'_i, b_{-i}), b_{-i})$, where (b'_i, b_{-i}) denotes the bid profile when bidder i bids b'_i and all other bidders bid according to \mathbf{b} . For the rest of this paper, we restrict the strategy space of each bidder to a set of undominated strategies. The strategy space D_i of bidder i is the set of bid vectors b_i such that 1) the number of keywords A with $b_i(A) > 0$ is at most κ and 2) for every keyword A , there exists a query $j \in A$ such that $b_i(A) \leq v_i(j)$. The bid vectors outside of D_i are unreasonable in the sense that they overbid on every single query associated with some keyword. Observe that for every bid vector $b'_i \notin D_i$, there exists a bid vector in D_i that weakly dominates it and this transition is polynomial-time computable.

2.2 The Bidding Game

Players generally have no dominant strategies in this mechanism. Therefore, we analyze the performance of this mechanism by looking at its equilibria. Let $\mathbf{v} = \{v_1, v_2, \dots, v_n\}$ denote the fixed valuation profile of the bidders. A bid profile $\mathbf{b} = \{b_1, b_2, \dots, b_n\}$ is a *pure Nash equilibrium* if $u_i(\mathbf{b}) \geq u_i(b'_i, b_{-i})$ for every player i and b'_i . Note that any Nash equilibrium in which bidders use strategies in $D = \{D_i\}$ is also a Nash equilibrium without any strategy-space restriction.

Let $OPT(\mathbf{v})$ denote the optimal allocation, i.e. $OPT(\mathbf{v}) = \text{argmax}_X SW(X)$. The *pure price of anarchy* (POA) is the ratio between the welfare of the optimal allocation and the welfare of the worst pure Nash equilibrium:

$$\text{pure POA} = \max_{\mathbf{b}: \text{a pure Nash eq.}} \frac{SW(OPT(\mathbf{v}))}{SW(\mathbf{b})}$$

where we define $SW(\mathbf{b}) = SW(X(\mathbf{b}))$.

As is standard in Nash equilibrium analysis, our setting is a full-information setting. That is every bidder knows all the information necessary to compute the payoff matrix, including every bidder’s valuation. The search engine (auctioneer) needs not know anything about bidder valuations.

3. HOMOGENEOUS KEYWORDS

We begin by showing that the performance of broad match mechanisms can be parameterized by a property of the keywords, called homogeneity. When a bidder uses broad match keywords, the mechanism loses some information about the bidder’s valuation. This loss can be minimized by designing only keywords that, in the bidders’ perspectives, group similar queries. We conceptualize this notion in the following definition.

Definition 3.1 (Homogeneity)

A keyword A is c -homogeneous if for every bidder i and every two queries $q_1, q_2 \in A$, $v_i(q_1) \leq c \cdot v_i(q_2)$.

By definition, exact match keywords are 1-homogeneous. Note that valuations of different bidders for the same query can still differ greatly.

Recall that the number of keyword bids is limited for each bidder. Therefore, in order to provide any welfare guarantee, the keywords must have good coverage of the bidders' valuations.

Definition 3.2 A set K of keywords is said to be *fully expressive* if for every bidder i , it is always possible to simultaneously bid on every query in \mathcal{S}_i using at most κ keywords from K .

Definition 3.3 A broad match mechanism is said to be *fully expressive* if the set of all keywords in the mechanism is fully expressive.

Now we will prove a bound of $c^2 + c$ on the price of anarchy for broad match mechanisms. The main technique is to use each bidder's equilibrium condition to bound the optimal welfare by prices, which, in turns, are bounded by the welfare of the equilibrium. The main challenge is that the limited number of keywords requires the bidder to drop one of the current keyword bids before he can bid on a new keyword. This adds to the equilibrium condition a switching cost, which works against our bound. We limit the total switching cost by using only κ equilibrium conditions per bidder. The limited number of equilibrium conditions means we need to carefully choose deviation bids to obtain the desirable bound.

Theorem 3.4 *For every fully-expressive broad match mechanism with c -homogeneous keywords and at least one pure Nash equilibrium, the pure POA is at most $c^2 + c$.*

PROOF. Let $SW(\mathbf{b})$ denote the welfare of a pure Nash equilibrium. Let $OPT(\mathbf{v}) = \{O_1, O_2, \dots, O_n\}$ denote the welfare-maximizing allocation. WLOG, we assume that $O_i \subseteq \mathcal{S}_i$. We have that $SW(OPT(\mathbf{v})) = \sum_i \sum_{j \in O_i} v_i(j)$. We will show that $SW(OPT(\mathbf{v})) \leq (c^2 + c)SW(\mathbf{b})$.

Fix a bidder i . Consider the set of queries $Q = O_i \setminus X_i(\mathbf{b})$. Let $h(q)$ denote the highest bidder other than i for query $q \in Q$, i.e. $b_{h(q)}(q) = p_i(q)$. Let K_i be the set of keywords on which bidder i is currently bidding. Since the mechanism is fully-expressive, there exists a set of at most κ keywords $K'_i = \{A_1, A_2, A_3, \dots\}$ such that $Q \subseteq \bigcup_{A \in K'_i} A$ (K'_i covers all queries in Q). Now we assign each query in Q to a keyword in K'_i . Let $A_k(Q)$ denote $Q \cap (A_k \setminus \bigcup_{i=1}^{k-1} A_i)$.

Next we consider a particular deviation for i . Consider a keyword $A \in K'_i$ and the set $A(Q)$, dropping the subscript for succinctness. Order the queries $q_1, q_2, \dots, q_t \in A(Q)$ so that $p_i(q_1) \leq p_i(q_2) \leq \dots \leq p_i(q_t)$. Let t be the largest integer such that $\frac{1}{c}v_i(q_t) > p_i(q_t)$. Let b_i^A be the bid vector obtained by 1) removing a keyword bid from b_i if b_i has κ keyword bids, and 2) setting the keyword bid $b_i^A(A)$ to $p_i(q_t)$. Observe that b_i^A is in i 's undominated strategies. Let $\sigma(A)$ denote the keyword bid that is removed when b_i^A is constructed from b_i ($\sigma(A)$ is a dummy keyword \emptyset if no keyword bid is removed). We choose the removed bids so that $\sigma(A)$

is unique for each A (this is possible because $|K'_i| \leq \kappa$). Let $W = X_i(b_i, b_{-i})$. Consider bidder i switching his bid vector from b_i to b_i^A . Let $W' = X_i(b_i^A, b_{-i})$. The equilibrium condition implies $\sum_{j \in W} v_i(j) - p_i(j) \geq \sum_{j \in W'} v_i(j) - p_i(j)$.

Let $U_i^A = \sum_{j \in W \setminus W'} v_i(j) - p_i(j)$ (the switching cost). Since the price that a bidder pays is independent of his bids, $p_i(j)$ remains the same for $j \in W \cap W'$. Canceling out the terms with $j \in W \cap W'$, we have

$$\begin{aligned} U_i^A &\geq \sum_{j \in W' \setminus W} v_i(j) - p_i(j) \\ &\geq \sum_{r=1}^t v_i(q_r) - p_i(q_r) + \sum_{j \in (W' \setminus W) \setminus Q} v_i(j) - p_i(j) \\ &\geq \sum_{r=1}^t v_i(q_r) - p_i(q_r) + \sum_{j \in (W' \setminus W) \setminus Q} \frac{1}{c} v_i(q_t) - p_i(q_t) \end{aligned}$$

We obtain the second inequality by separating the terms with queries in Q from the other terms and noting that i wins only queries with $p_i(j) \leq p_i(q_t)$. The last inequality comes from the assumption that keywords are homogeneous and the fact that $p_i(j)$ is at most the bid of $p_i(q_t)$. Next we drop the terms $\sum \frac{1}{c} v_i(q_t) - p_i(q_t)$, using the fact that they are positive.

$$U_i^A + \sum_{r=1}^t p_i(q_r) \geq \sum_{r=1}^t v_i(q_r) \quad (1)$$

$$U_i^A + \sum_{r=1}^t p_i(q_r) + \sum_{r=t+1}^l p_i(q_r) \geq \sum_{r=1}^t v_i(q_r) + \sum_{r=t+1}^l \frac{1}{c} v_i(q_r) \quad (2)$$

$$c \cdot U_i^A + c \cdot \sum_{r=1}^l p_i(q_r) \geq \sum_{r=1}^l v_i(q_r) \quad (3)$$

$$c \cdot U_i^A + c^2 \cdot \sum_{r=1}^l v_{h(q_r)}(q_r) \geq \sum_{r=1}^l v_i(q_r) \quad (4)$$

$$c \cdot U_i^A + c^2 \sum_{j \in A(Q)} v_{h(j)}(j) \geq \sum_{j \in A(Q)} v_i(j) \quad (5)$$

We obtain (2) using the fact that $\frac{1}{c} v_i(q_r) \leq p_i(q_r)$ for all $r > t$. Inequality (4) uses the assumption that $p_i(q_r) = b_{h(q_r)}(q_r)$ is in undominated strategies (of bidder $h(q_r)$), which implies there exists a keyword A and a query $j \in A$ such that $q_r \in A$ and $b_{h(q_r)}(q_r) = b_{h(q_r)}(A) \leq v_{h(q_r)}(j)$; then the homogeneity of keyword A implies $v_{h(q_r)}(j) \leq c \cdot v_{h(q_r)}(q_r)$.

Now we bound the total switching cost $\sum_{A \in K'_i} U_i^A$ by the total welfare of bidder i . Recall that $U_i^A = \sum_{j \in W \setminus W'} v_i(j) - p_i(j)$ which implies

$$U_i^A \leq \sum_{j \in W \setminus W'} v_i(j) = \sum_{j \in X_i(b_i, b_{-i}) \setminus X_i(b_i^A, b_{-i})} v_i(j).$$

Summing both sides over $A \in K'_i$, we have

$$\sum_{A \in K'_i} U_i^A \leq \sum_{A \in K'_i} \sum_{j \in X_i(b_i, b_{-i}) \setminus X_i(b_i^A, b_{-i})} v_i(j) \quad (6)$$

Let $\delta(A)$ denote the set of queries that bidder i would lose if he drops his bid on keyword A . The set $X_i(b_i, b_{-i}) \setminus X_i(b_i^A, b_{-i})$ is precisely $\delta(\sigma(A))$. Using the fact that $\sigma(A) \in K_i$ and the fact that $\sigma(A)$'s are distinct, we have

$$\sum_{A \in K'_i} U_i^A \leq \sum_{A \in K'_i} \sum_{j \in \delta(\sigma(A))} v_i(j) \quad (7)$$

$$\leq \sum_{A \in K_i} \sum_{j \in \delta(A)} v_i(j) \quad (8)$$

Since bidder i can lose a query j only if he drops his bid on the keyword $\text{argmax}_{A: j \in A} b_i(A)$, for every two distinct keywords $A, B \in K_i$, the sets $\delta(A)$ and $\delta(B)$ are disjoint. Therefore, $\sum_{A \in K_i} \sum_{j \in \delta(A)} v_i(j) \leq \sum_{j \in X_i(b_i, b_{-i})} v_i(j)$ and thus

$$\sum_{A \in K'_i} U_i^A \leq \sum_{j \in X_i(b_i, b_{-i})} v_i(j) \quad (9)$$

Summing both sides of inequality (5) over $A \in K'_i$ and applying inequality (9), we have

$$\begin{aligned} \sum_{j \in Q} v_i(j) &\leq c \cdot \sum_{j \in X_i(\mathbf{b})} v_i(j) + c^2 \sum_{j \in Q} v_{h(j)}(j) \\ \sum_{j \in O_i \setminus X_i(\mathbf{b})} v_i(j) &\leq c \cdot \sum_{j \in X_i(\mathbf{b})} v_i(j) \\ &\quad + c^2 \sum_{j \in O_i \setminus X_i(\mathbf{b})} v_{h(j)}(j) \end{aligned} \quad (10)$$

Then summing both sides over i , adding $\sum_i \sum_{j \in O_i \cap X_i(\mathbf{b})} v_i(j)$ to the left and adding $c^2 \sum_i \sum_{j \in O_i \cap X_i(\mathbf{b})} v_i(j)$ to the right yields

$$\begin{aligned} \sum_i \left(\sum_{j \in O_i \setminus X_i(\mathbf{b})} v_i(j) + \sum_{j \in O_i \cap X_i(\mathbf{b})} v_i(j) \right) &\leq c \cdot \sum_i \sum_{j \in X_i(\mathbf{b})} v_i(j) \\ &\quad + c^2 \sum_i \left(\sum_{j \in O_i \setminus X_i(\mathbf{b})} v_{h(j)}(j) + \sum_{j \in O_i \cap X_i(\mathbf{b})} v_i(j) \right) \end{aligned}$$

$$SW(OPT(\mathbf{v})) \leq c \cdot SW(\mathbf{b}) + c^2 \cdot SW(\mathbf{b})$$

Rearranging the terms yields the POA of $c^2 + c$. \square

4. THE LOWERBOUND

In this section, we show that the bound given by theorem 3.4 is almost tight. Consider the following broad match setting (see figure 2). There are c bidders, namely $0, 1, \dots, c-1$. There are $c-1$ queries, namely b_1, b_2, \dots, b_{c-1} . The only keyword is $A = \{b_1, b_2, \dots, b_{c-1}\}$. $v_0(b_1) = c, v_0(b_j) = 1, \forall j > 1$. For each $i > 0$, $v_i(b_i) = \frac{1}{2}c^2$ and $v_i(b_j) = \frac{1}{2}c$ for all $j \neq i$. There is a pure Nash equilibrium where bidder 0 bids c on keyword A and the other bidders bid nothing. The welfare of this equilibrium is $c+c-2$. The optimal welfare is $\frac{1}{2}c^2(c-1)$, which results from the allocation: $X_0 = \emptyset, X_i = \{b_i\}$. Hence the pure POA is at least $\frac{1}{4}c^2$.




Figure 2: An example with POA at least $\frac{1}{4}c^2$. Queries are b_1, b_2, \dots, b_{c-1} and the keyword is A .

Theorem 4.1 *There exists a fully-expressive broad match mechanism with c -homogeneous keywords such that its pure POA is at least $\frac{1}{4}c^2$.*

5. THE EXPRESSIVENESS-HOMOGENEITY TRADE-OFF

In this section, we tackle the mechanism design problem: How to design keywords for a broad match mechanism? As shown in the previous two sections, the designer wants the keywords to be homogeneous and fully-expressive. However, designing such keywords can be difficult or even impossible, given the constraint on κ , the number of keyword bids. Note that it is easy to design keywords that are fully-expressive but not necessarily homogeneous: create 1 keyword that contains all queries in the universe. It is also easy to design keywords that are 1-homogeneous but not necessarily expressive: for every query j , create a keyword $J = \{j\}$. That leads us to explore the trade-off between the two properties. The best mechanism probably lies between those two extremes. First, we define semi-expressiveness.

Definition 5.1 A set L of keywords is α -expressive if for every bidder i and every set $Q \subseteq \mathcal{S}_i$ of queries such that $|Q| \leq \lceil \alpha |\mathcal{S}_i| \rceil$, there exists a set $K \subseteq L$ such that $|K| \leq \kappa$ and $Q \subseteq \bigcup_{A \in K} A$.

We say a broad match mechanism is α -expressive if the set of all keywords in the mechanism is α -expressive. We can show that the upper bound on the POA of broad match mechanisms is a function of α and c .

Theorem 5.2 *For every α -expressive broad match mechanism with c -homogeneous keywords and at least one pure Nash equilibrium, the pure POA is at most $\frac{1}{\alpha}(c^2 + c)$.*

This bound advises us that when we introduce a set of new broad match keywords that improves expressiveness by a factor of x , the homogeneity constant of the keywords should not increase by more than a factor of \sqrt{x} .

Note that the mechanism designer does not need to know bidder valuation in order to design broad match keywords. For homogeneity, the designer only needs to know the relative valuations of a bidder for different queries. For example, it is enough to know that advertisers value “sneakers” and “running shoes” roughly the same to create a keyword for them. For expressiveness, the designer only needs to know the set of queries that each bidder is interested in (queries with positive valuation).

The proof of theorem 5.2 combines the arguments in the proof of theorem 3.4 with the next lemma.

Lemma 5.3 *For every constant $0 \leq \alpha \leq 1$, every additive valuation v and every set T , there exists a set $T' \subseteq T$ of size at most $\lceil \alpha|T| \rceil$ such that $v(T') \geq \alpha v(T)$.*

PROOF. Rename the elements in $T = \{1, 2, \dots\}$ so that $v(1) \geq v(2) \geq v(3) \geq \dots$. Let $T' = \{1, 2, \dots, \lceil \alpha|T| \rceil\}$. If $v(T') \geq \alpha v(T)$, we are done. Suppose not. That means $v(\lceil \alpha|T| \rceil) < \frac{\alpha v(T)}{\lceil \alpha|T| \rceil}$. Let $S = T \setminus T'$. Then $v(S) \leq \frac{\alpha v(T)}{\lceil \alpha|T| \rceil} (|T| - \lceil \alpha|T| \rceil)$. That implies $v(T') + v(S) < \alpha v(T) + \frac{\alpha v(T)}{\lceil \alpha|T| \rceil} (|T| - \lceil \alpha|T| \rceil) = \alpha v(T) \left(\frac{|T|}{\lceil \alpha|T| \rceil} \right) \leq v(T)$. Contradiction to the assumption that v is additive. \square

Now we modify the proof of theorem 3.4 to obtain the proof of theorem 5.2.

Proof of Theorem 5.2: Using the same notation, the welfare-maximizing allocation is $\{O_1, O_2, \dots, O_n\}$. WLOG, assume $O_i \subseteq S_i$. By lemma 5.3, there exists a set $O'_i \subseteq O_i$ such that $|O'_i| \leq \lceil \alpha|O_i| \rceil$ and $v_i(O'_i) \geq \alpha v_i(O_i)$. Then we redefine Q in the original proof as $Q = O'_i \setminus X_i(\mathbf{b})$. We redefine K'_i using the new definition of expressiveness. By the assumption that the mechanism is α -expressive, there exists a set of keywords K'_i such that $|K'_i| \leq \kappa$ and $Q \subseteq O'_i \subseteq \bigcup_{A \in K'_i} A$. Then we proceed as in the original proof. Inequality (10) becomes

$$\sum_{j \in O'_i \setminus X_i(\mathbf{b})} v_i(j) \leq c \cdot \sum_{j \in X_i(\mathbf{b})} v_i(j) + c^2 \sum_{j \in O'_i \setminus X_i(\mathbf{b})} v_{h(j)}(j)$$

Summing both sides over i , adding $\sum_i \sum_{j \in O'_i \cap X_i(\mathbf{b})} v_i(j)$ to the left and adding $c^2 \sum_i \sum_{j \in O'_i \cap X_i(\mathbf{b})} v_i(j)$ to the right yields

$$\begin{aligned} \sum_i \left(\sum_{j \in O'_i \setminus X_i(\mathbf{b})} v_i(j) + \sum_{j \in O'_i \cap X_i(\mathbf{b})} v_i(j) \right) &\leq c \cdot \sum_i \sum_{j \in X_i(\mathbf{b})} v_i(j) \\ + c^2 \sum_i \left(\sum_{j \in O'_i \setminus X_i(\mathbf{b})} v_{h(j)}(j) + \sum_{j \in O'_i \cap X_i(\mathbf{b})} v_i(j) \right) \end{aligned}$$

Combining the terms, we have $\sum_i v_i(O'_i) \leq c \cdot SW(\mathbf{b}) + c^2 \cdot SW(\mathbf{b})$. Using the fact that $v_i(O'_i) \geq \alpha v_i(O_i)$, we obtain $\alpha SW(OPT(\mathbf{v})) \leq c \cdot SW(\mathbf{b}) + c^2 \cdot SW(\mathbf{b})$. Therefore,

$$SW(OPT(\mathbf{v})) \leq \frac{1}{\alpha} (c^2 + c) SW(\mathbf{b}).$$

■

We note one special case, when the mechanism does not have any broad match keyword.

Definition 5.4 (EMO mechanism)

The *exact-match-only* mechanism is a broad match mechanism that has one exact match keyword $J = \{j\}$ for each query j and no other keywords.

All keywords in EMO mechanisms are 1-homogeneous. Let $\rho = \max_i |\mathcal{S}_i|$ and $\gamma = \rho/\kappa$. By observing that an EMO mechanism is $\frac{1}{\gamma}$ -expressive, we have the following corollary.

Corollary 5.5 *For every EMO mechanism that limits the number of keyword bids to κ , the pure POA is at most 2γ .*

For example, if every advertiser has a positive valuation for at most 500,000 queries and the EMO mechanism limits the number of keyword bids to 50,000, the pure POA of this mechanism is at most 20.

6. AN EXTENSION TO SUBADDITIVE VALUATION

In this section, we generalize our bound for EMO mechanisms to cover subadditive valuation. A valuation v is *subadditive* if for all subsets S_1, S_2, \dots, S_k of the queries, $v(S_1 \cup S_2 \cup \dots \cup S_k) \leq v(S_1) + v(S_2) + \dots + v(S_k)$. Define the strategy space D_i of each bidder i to be the set of bid vectors b_i such that 1) the number of keywords A with $b_i(A) > 0$ is at most κ and 2) $\sum_{j \in S} b_i(j) \leq v_i(S)$ for every set S of queries. In short, D_i corresponds to the set of bid vectors that don't overbid on any subset of queries. The next lemma, which is proven by [3], implies that for every bid vector $b'_i \notin D_i$, there exists a bid vector in D_i that weakly dominates it.

Lemma 6.1 *Let O be a set of at most κ queries and $p(j)$ be a price for the query $j \in O$. Let v_i be a subadditive valuation of bidder i . There is a bid vector $b_i \in D_i$ such that $b_i(j) = 0$ for all $j \notin O$ and*

$$v_i(S) - \sum_{j \in S} p(j) \geq v_i(O) - \sum_{j \in O} p(j)$$

where $S = \{j \in O : b_i(j) \geq p(j)\}$.

We can think of $v_i(S) - \sum_{j \in S} p(j)$ as the utility of bidder i from bidding b_i against prices p induced by the bids of other bidders.

Unfortunately, lemma 5.3 does not hold for subadditive valuation but we can show that it holds when the requirements are slightly relaxed.

Lemma 6.2 *For every constant $0 < \alpha \leq 1$, every subadditive valuation v and every set T , there exists a set $T' \subseteq T$ such that*

$$|T'| \leq \lceil \alpha|T| \rceil \quad \text{and} \quad v(T') \geq \left\lceil \frac{1}{\alpha} \right\rceil v(T).$$

PROOF. Let $a = \lceil \frac{1}{\alpha} \rceil$. Divide T arbitrarily into a sets of size at most $\lceil \alpha|T| \rceil$ and name them S_1, \dots, S_a . By the subadditivity of v , $\sum_{i=1}^a v(S_i) \geq v(T)$. That implies $\frac{1}{a} \sum_{i=1}^a v(S_i) \geq \frac{1}{a} v(T)$. There must be at least one set $S_i \in \{S_1, S_2, \dots, S_a\}$ such that $v(S_i) \geq \frac{1}{a} \sum_{i=1}^a v(S_i) \geq \frac{1}{a} v(T)$. \square

We keep the same definition for ρ and γ .

Theorem 6.3 *For every EMO mechanism that has only bidders with subadditive valuations and a limit of κ on the number of keyword bids, the pure POA is at most $2 \lceil \gamma \rceil$.*

PROOF. Denote the welfare-maximizing allocation by $O = \{O_1, O_2, \dots, O_n\}$. WLOG, assume $O_i \subseteq S_i$. Let \mathbf{b} denote an equilibrium bid profile. Let $\alpha = \frac{1}{\gamma}$. Consider bidder i . According to lemma 6.2, there exists a set of queries $Q_i \subseteq O_i$

such that $|Q_i| \leq \lceil \alpha |O_i| \rceil$ and $v_i(Q_i) \geq \frac{1}{\lceil \alpha \rceil} v_i(O_i)$. Observe that $|Q_i| \leq \kappa$. Now we will invoke lemma 6.1, by letting $p(j)$ be the price of query j for bidder i when the other bidders bid b_{-i} . By lemma 6.1, there exists a bid vector $b'_i \in D_i$ such that $b'_i(j) = 0$ for all $j \notin Q_i$ and

$$v_i(X_i(b'_i, b_{-i})) - \sum_{j \in X_i(b'_i, b_{-i})} p_i(j) \geq v_i(Q_i) - \sum_{j \in Q_i} p_i(j).$$

Consider bidder i switching his bid from b_i to b'_i . Denote bidder i 's utility in the equilibrium by u_i . The equilibrium condition gives

$$\begin{aligned} u_i &\geq v_i(X_i(b'_i, b_{-i})) - \sum_{j \in X_i(b'_i, b_{-i})} p_i(j) \\ &\geq v_i(Q_i) - \sum_{j \in Q_i} p_i(j) \\ &\geq \frac{1}{\lceil \alpha \rceil} v_i(O_i) - \sum_{j \in Q_i} p_i(j) \end{aligned}$$

Summing both sides over i , we have

$$\begin{aligned} \sum_i u_i &\geq \frac{1}{\lceil \alpha \rceil} \sum_i v_i(O_i) - \sum_i \sum_{j \in Q_i} p_i(j) \\ &\geq \frac{1}{\lceil \alpha \rceil} SW(OPT(\mathbf{v})) - \sum_i \sum_{j \in O_i} p_i(j) \\ &\geq \frac{1}{\lceil \alpha \rceil} SW(OPT(\mathbf{v})) - \sum_i \sum_{j \in X_i(\mathbf{b})} b_i(j) \\ &\geq \frac{1}{\lceil \alpha \rceil} SW(OPT(\mathbf{v})) - \sum_i v_i(X_i(\mathbf{b})) \end{aligned}$$

The third inequality comes from the fact that the price of an item is at most the bid of the winner. The last inequality uses the assumption that $b_i \in D_i$. Applying the fact that $SW(\mathbf{b}) \geq \sum_i u_i$, $SW(\mathbf{b}) \geq \sum_i v_i(X_i(\mathbf{b}))$, and rearranging the terms, we have

$$SW(\mathbf{b}) \geq \frac{1}{2\lceil \alpha \rceil} SW(OPT(\mathbf{v}))$$

□

7. THE PER-IMPRESSION BROAD MATCH MECHANISM

In this section, we analyze a per-impression variant of the broad match mechanism. The main difference between the model in section 2 and the per-impression model is that in the latter, the number of impressions for each search term is made explicit. As a result, we can explicitly model each bidder's valuation per impression. This leads to a more practical definition of homogeneity, which now operates on value-per-impression (rather than value-per-search-term in section 2). Despite this fundamental difference, their analyses are syntactically very similar and our bound on the POA for the per-impression mechanism mirrors that of the per-search-term mechanism.

7.1 The Per-impression Model

There are n bidders and m search terms. For each search term j , the search engine (auctioneer) has η_j impressions

for sale. Bidder i has valuation $v_i(j)$ for each impression of search term j . The bidders' valuations are additive across impressions. An allocation is an assignment of impressions to bidders such that each impression is assigned to at most one bidder. The per-impression broad match mechanism is defined as follows

1. Each bidder i submits bid b_i on keywords upfront. The number of keywords A such that $b_i(A) > 0$ must be at most κ .
2. Each bidder's bid is translated into bids on search terms using the formula $b_i(j) = \max_{A:j \in A} b_i(A)$.
3. For each search term j , for each impression generated from j , allocate the impression to the bidder with the highest $b_i(j)$ and charge him the price $p_i(j) = \max_{k \neq i} b_k(j)$.

Note that this mechanism always allocates all impressions from the same search term to the same bidder. The optimal (welfare-maximizing) allocation assigns each impression from search term j to the bidder i with the highest $v_i(j)$. Therefore, our analysis deals only with allocations that assign all impressions from the same search term to the same bidder. Denote an allocation by $X = \{X_1, X_2, \dots, X_n\}$ where X_i is the set of search terms that go to bidder i . Given a set of search terms T , we define $v_i(T) = \sum_{j \in T} \eta_j v_i(j)$. The *welfare* $SW(X)$ of an allocation X is defined as $\sum_i \sum_{j \in X_i} \eta_j v_i(j)$. The utility of bidder i is $\sum_{j \in X_i(\mathbf{b})} \eta_j (v_i(j) - p_i(j))$. We restrict the strategy space of each bidder i to D_i which is defined the same way as before.

One can view the previous per-search-term mechanism as a special case of the per-impression mechanism where $\eta_j = 1$ for all j .

7.2 The Upper Bound on the POA

In this model, we use value-per-impression in the definition of homogeneity.

Definition 7.1 A keyword B is c -homogeneous if for every bidder i and every two search terms $j_1, j_2 \in B$, $v_i(j_1) \leq c \cdot v_i(j_2)$.

The support of a bidder is defined as in section 2 except it is done in terms of search terms instead of queries. The definition of α -expressiveness remains the same. We have a lemma similar to lemma 5.3.

Lemma 7.2 For every constant $0 \leq \alpha \leq 1$, every additive valuation v and every set T of search terms, there exists a set $T' \subseteq T$ of size at most $\lceil \alpha T \rceil$ such that $v(T') \geq \alpha v(T)$.

We have an upper bound on the POA similar to theorem 5.2.

Theorem 7.3 For every α -expressive broad match mechanism with c -homogeneous keywords and at least one pure Nash equilibrium, the pure POA is at most $\frac{1}{\alpha}(c^2 + c)$.

Since the proof is very similar to the proof of theorem 5.2, we present here a proof sketch.

Proof (sketch): We use the same notations as in the proof of theorem 5.2 except the queries are now search terms. The equilibrium condition implies

$$U_i^A \geq \sum_{j \in W \setminus W} \eta_j(v_i(j) - p_i(j))$$

Then we proceed as before except we now have η_j in the terms involving $v_i(j)$ and $p_i(j)$. Whenever we use the fact that $\frac{1}{c}v_i(q_t) > p_i(q_t)$, we substitute it with $\frac{1}{c}\eta_j v_i(q_t) > \eta_j p_i(q_t)$. Then we have (compared to inequality (5))

$$c \cdot U_i^A + c^2 \sum_{j \in A(Q)} \eta_j v_{h(j)}(j) \geq \sum_{j \in A(Q)} \eta_j v_i(j)$$

We can bound the total switching cost and combine it with the above inequality as before. Then we have

$$\sum_{j \in O_i' \setminus X_i(\mathbf{b})} \eta_j v_i(j) \leq c \sum_{j \in X_i(\mathbf{b})} \eta_j v_i(j) + c^2 \sum_{j \in O_i' \setminus X_i(\mathbf{b})} \eta_j v_{h(j)}(j)$$

We conclude the proof by summing both sides over i and following similar syntactic manipulations as in the proof of theorem 5.2. ■

8. CONCLUSION

Broad match keywords are a crucial part of search advertising. Our tight bounds on the price of anarchy of broad match mechanisms suggest that the performance of the multi-keyword sponsored search mechanism depends heavily on the design of each broad match keyword (the set of queries it is associated with). Our parameterization of the upper bound on the price of anarchy by expressiveness and homogeneity gives the mechanism designer a flexible design space that ranges from no broad match keyword to a fully-expressive bidding language.

Another useful design choice is to personalize keywords for advertisers. That is to create keywords that are exclusive to each advertiser. We believe that this is happening in the industry. Our model can be easily modified to allow such keywords by changing the definition of homogeneity to bind only for the advertisers who can bid on them, and our bounds will remain the same.

Our work suggests a number of interesting open questions. First, is it possible to generalize theorem 3.4 to subadditive valuation? Second, broad match keywords are just one way to design the bidding language. What else can we do? Third, since we cannot make the assumption that a GSP auction for a search term is independent from GSP auctions for other search terms, we might as well not use them. What are other alternative mechanisms for allocating impressions for all different search terms?

Acknowledgments

We thank Tim Roughgarden and anonymous reviewers for helpful comments.

9. REFERENCES

[1] A. Ghosh, H. Nazerzadeh, and M. Sundararajan. Computing Optimal Bundles for Sponsored Search. In *WINE*, pages 576-583, 2007.

[2] A. Ghosh and A. Sayedi. Expressive auctions for externalities in online advertising. In *WWW*, pages 371-380, 2010.

[3] K. Bhawalkar and T. Roughgarden. Welfare Guarantees for Combinatorial Auctions with Item Bidding. In *SODA*, 2011.

[4] B. Lucier and Renato Paes Leme. Improved Social Welfare Bounds for GSP at Equilibrium. *CoRR abs*, 1011.3268, 2010.

[5] G. Aggarwal, S. Muthukrishnan, D. Pál, and M. Pál. General auction mechanism for search advertising. In *WWW*, pages 241-250, 2009.

[6] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. *ACM Conference on Electronic Commerce*, pages 1-7, 2006.

[7] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1), pages 242-259, 2007.

[8] G. Goel and A. Mehta: Online budgeted matching in random input models with applications to Adwords. In *SODA*, pages 982-991, 2008.

[9] C. Borgs, J. T. Chayes, N. Immorlica, K. Jain, O. Etesami, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *WWW*, pages 531-540, 2007.

[10] D. J. Martin, J. Gehrke, and J. Y. Halpern. Toward Expressive and Scalable Sponsored Search Auctions. In *ICDE*, pages 237-246, 2008.

[11] P. Rusmevichientong and D. P. Williamson: An adaptive algorithm for selecting profitable keywords for search-based advertising services. *ACM Conference on Electronic Commerce*, pages 260-269, 2006.

[12] E. Even-Dar, V. S. Mirrokni, S. Muthukrishnan, Y. Mansour, and U. Nadav. Bid optimization for broad match ad auctions. In *WWW*, 231-240, 2009.

[13] M. Mahdian and G. Wang. Clustering-Based Bidding Languages for Sponsored Search. In *ESA*, pages 167-178, 2009.

[14] I. Giotis and A. R. Karlin: On the Equilibria and Efficiency of the GSP Mechanism in Keyword Auctions with Externalities. In *WINE*, pages 629-638, 2008.

[15] G. Christodoulou, A. Kovács, and M. Schapira: Bayesian Combinatorial Auctions. In *ICALP*, pages 820-832, 2008.

[16] H. Varian, Position Auctions, *International Journal of Industrial Organization*, pages 1163-1178, 2007.

[17] Z. Abrams, A. Ghosh, and E. Vee. Cost of Conciseness in Sponsored Search Auctions. *WINE*, pages 326-334, 2007.

[18] R. Kumar, M. Mahdian, and A. Sayedi. Mechanism Design for Complexity-Constrained Bidders. In *WINE*, pages 513-520, 2009.

[19] S. K. Singh and V. P. Roychowdhury. To Broad-Match or Not to Broad-Match : An Auctioneer's Dilemma ?. *CoRR abs*, 0802.1957, 2008.

[20] S. K. Singh, V. P. Roychowdhury, M. Bradonjic, and B. A. Rezaei. Exploration via design and the cost of uncertainty in keyword auctions. *CoRR abs*, 0707.1053, 2007.