

Methodology for Designing Reasonably Expressive Mechanisms with Application to Ad Auctions

Michael Benisch

School of Computer Science
Carnegie Mellon University
mbenisch@cs.cmu.edu

Norman Sadeh

School of Computer Science
Carnegie Mellon University
sadeh@cs.cmu.edu

Tuomas Sandholm

Computer Science Department
Carnegie Mellon University
sandholm@cs.cmu.edu

Abstract

Mechanisms (especially on the Internet) have begun allowing people or organizations to express richer preferences in order to provide for greater levels of overall satisfaction. In this paper, we develop an operational methodology for quantifying the expected gains in economic efficiency associated with different forms of expressiveness. We begin by proving that the sponsored search mechanism (GSP) used by Google, Yahoo!, MSN, etc. can be arbitrarily inefficient. We then experimentally compare its efficiency to a slightly more expressive variant (PGSP), which solicits an extra bid for a premium class of positions. We generate random preference distributions based on published industry knowledge. We determine ideal strategies for the agents using a custom tree search technique, and we also benchmark using straightforward heuristic bidding strategies. The GSP's efficiency loss is greatest in the practical case where some advertisers ("brand advertisers") prefer top positions while others ("value advertisers") prefer middle positions, and that loss can be dramatic. It is also worst when agents have small profit margins. While the PGSP is only slightly more expressive (and thus not much more cumbersome), it removes almost all of the efficiency loss in all of the settings we study.

1 Introduction

A key trend on the Internet is a move toward more expressiveness in the mechanisms that mediate interactions such as the allocation of resources, matching of peers, and elicitation of opinions. Intuitively, one would think that more expressiveness leads to higher efficiency (sum of the agents' utilities) of the outcome (e.g., due to better matching of supply and demand). Efficiency improvements have indeed been reported from combinatorial and multi-attribute auctions (e.g., [Sandholm, 2007]), as well as expressive ad auction variants (e.g., [Lahaie *et al.*, 2008; Parkes and Sandholm, 2005; Even-Dar *et al.*, 2007; Boutilier *et al.*, 2008]). Yet, adding expressiveness does not always improve the outcome of a mechanism in practice. It generally increases the overhead

associated with running the mechanism (e.g., time and effort to collect data, computation requirements) [Martin *et al.*, 2008]. For some populations, increased expressiveness can be unnecessary [Abrams *et al.*, 2007] and can give rise to additional equilibria of poor efficiency [Milgrom, 2007]. It can even confuse (e.g., Herb Simon's "bounded rationality" [1955]), or aggravate (e.g., Barry Schwartz's "tyranny of choice" [2004]) people. What is missing is an operational methodology to quantify the expected efficiency gains from increasing expressiveness.

Work on information complexity of mechanisms has a long history (e.g., [Mount and Reiter, 1974; Hurwicz, 1972]). Recent theoretical work based on the notion of *shattering* from computational learning theory provided the foundations for studying expressiveness in a domain-independent manner [Benisch *et al.*, 2008]. In this paper we operationalize that theory with a methodology for comparing mechanisms with different degrees and forms of expressiveness. We apply it to *sponsored search*, or ad auctions conducted online for placement alongside search results.

The sponsored search industry accounts for tens of billions of dollars in revenue annually. The most frequent variant of these auctions, the *generalized second price (GSP)* mechanism used by Google, Yahoo!, MSN, etc. solicits a single bid from each advertiser (i.e., agent) for a keyword and assigns the advertisers to positions on a search result page according to the bids (roughly speaking, with the first position going to the highest bidder, the second position to the second highest, etc.). Since agents cannot offer a separate bid price for each ad position, the GSP mechanism is fundamentally inexpressive. We will attempt to characterize the loss of economic efficiency caused by this inexpressiveness, and to explore the conditions that affect that loss.

We begin by adapting our recent theoretical framework for studying expressiveness [Benisch *et al.*, 2008] to analyze the GSP. We find that for some preference distributions the GSP is arbitrarily inefficient.

In order to measure the inefficiency in practice we must be able to predict the outcome of the mechanism. The equilibrium of the GSP is known when it is assumed that agents have complete information (i.e., no private information about valuations) and monotonic preferences over positions (i.e., higher positions are always preferred) [Varian, 2007]; however when we relax these somewhat restrictive assumptions, the equi-

librium behavior is unknown. In fact, it is often difficult to characterize equilibrium behavior in less than fully expressive mechanisms when agents have complex preferences [Rosenthal and Wang, 1996; Wilenius and Andersson, 2007; Szentes and Rosenthal, 2003]. For that reason, we develop a general tree search technique for computing an upper bound on a mechanism’s expected efficiency, that involves finding social welfare maximizing strategies for the agents. In the worst case our search algorithm takes time that is exponential in the number of agents and types, but it can be applied to any preference distribution and provides an upper bound that tightens in an anytime manner.

We conclude with a series of experiments comparing the GSP to our slightly more expressive mechanism, which solicits an extra bid for premium ad positions, which we coin *Premium GSP (PGSP)*. We generate a range of realistic synthetic preference distributions based on published industry knowledge, and apply our search technique to compare the efficiency bounds achieved by social welfare maximizing strategies in the two mechanisms. We also examine the performance of the two mechanisms when agents use a straightforward heuristic bidding strategy.

While we must be careful not to read too much into experiments on synthetic data, they suggest that the GSP’s efficiency loss can be dramatic. It is greatest in the practical case where some agents (“brand advertisers”) prefer top positions while others (“value advertisers”) prefer middle positions (since customers who click on ads in middle positions are more likely to take action, resulting in revenue). The loss is also worst when agents have small profit margins. Despite the fact that our PGSP mechanism is only slightly more expressive (and thus not much more cumbersome), it removes almost all of the efficiency loss in all of the settings we study.

2 Setting and background results

The setting we study is a one-shot auction for a set of k advertising positions that are ranked from 1 to k (rank 1 is the highest rank). In the model there are n agents. Each agent i has some private information (not known by the mechanism or any other agent) denoted by a type, t_i , (e.g., a vector of valuations, one for each of the k positions) from the space of the agent’s possible types, T_i .

Settings where each agent has a utility function, $u_i(t_i, O)$, that depends only on its own type and the outcome (matching of agents to positions), $O \in \mathcal{O}$, chosen by the mechanism are called *private values* settings. We also discuss more general *interdependent values* settings, where $u_i = u_i(t^n, O)$, i.e., an agent’s utility depends on the others’ private signals as well (for example, if one agent’s value for a position depends on market estimates of the other agents). In both settings, agents report expressions to the mechanism, denoted θ_i , based only on their own types. In the GSP mechanism each agent can report a single real value indicating his/her bid. A mapping from types to expressions is called a *pure strategy*.

Based on these expressions the mechanism computes the value of an outcome function, $f(\theta^n)$, which chooses an outcome. In the GSP mechanism the outcome function maps agents to positions based on the order of their bids (the high-

est bidder is assigned the first position, the second highest bidder is assigned the second, etc.).¹ The mechanism may also compute the value of a payment function, $\pi(\theta^n)$, which determines how much each agent must pay or get paid. In this paper, we ignore the mechanism’s payment function because expressiveness is tied directly to a mechanism’s outcome function.²

We denote by $W(t^n, o)$ the social welfare of outcome o when agents have private types t^n , i.e., $W(t^n, o) = \sum_i u_i(t^n, o)$. Assuming that the expression of each agent in the mechanism’s most efficient Nash equilibrium is given by a function $m_i(t_i)$, we can describe the mechanism’s expected efficiency under that equilibrium, $\mathcal{E}(f, \pi)$, with the following equation (expectation is taken over the types of the agents, and their randomized equilibrium expressions).

$$E[\mathcal{E}(f, \pi)] = \int_{t^n} P(t^n) \int_{\theta^n} P(m(t^n) = \theta^n) W(t^n, f(\theta^n)) \quad (1)$$

2.1 A framework for characterizing expressiveness

The theoretical framework that we developed in our earlier work [Benisch *et al.*, 2008], provides the foundations for understanding the impact of making mechanisms more or less expressive, by providing meaningful, general definitions of a mechanism’s expressiveness.

In that work, we defined an *impact vector* to capture the impact of a particular expression by an agent under the different possible types of the other agents, and an expressiveness concept based on a notion called *shattering*, which we adapted from the field of computational learning theory [Vapnik and Chervonenkis, 1971]. The adapted notion captures an agent’s ability to distinguish among each of the impact vectors involving a subset of outcomes.

We also introduced a slightly weaker adaptation of shattering, called *semi-shattering*, for analyzing the more restricted setting where agents have private values. It captures an agent’s ability to cause each of the unordered pairs of outcomes (with replacement) to be chosen for every pair of types of the other agents, but without being able to control the *order* of the outcomes (i.e., which outcome happens for which type). We defined a measure of expressiveness based on the size of the largest outcome space that an agent can shatter or semi-shatter. It is called the *(semi-)shatterable outcome dimension*.

In addition to defining the expressiveness notions, we tied those notions to an upper bound on the expected efficiency of a mechanism’s *most efficient* equilibrium. We derived the bound by making the optimistic assumption that the agents play strategies which, taken together, attempt to maximize social welfare. The bound is given by the following equation

¹In practice the bids are adjusted by predicted click-through rates (CTR) before conducting the ranking. For simplicity, we do not weight by CTR. However, our formulation can be easily extended to account for this by multiplying each agent’s original bid by its CTR.

²Since the efficiency bound that we study does not directly depend on equilibrium behavior, this is without loss of generality, as long as agents do not care about *each others’* payments.

(the max is taken over all possible joint pure strategies).

$$E[\mathcal{E}(f)]^+ = \max_{\hat{h}(\cdot)} \int_{t^n} P(t^n) W(t^n, f(\hat{h}(t^n))) \quad (2)$$

Our earlier work provided several results relating this bound to a mechanism’s expressiveness. For the purposes of this paper the following result will prove useful.

Theorem 1 (reworded from [Benisch et al., 2008]). *For any setting, there exists a distribution over agent preferences such that the upper bound on expected efficiency for the best outcome function where agent i has semi-shatterable outcome dimension $d_i < |O|$ is arbitrarily lower than that of the best outcome function where agent i has semi-shatterable outcome dimension $d_i + 1$.*

3 Expressiveness in ad auctions

In order to study the expressiveness properties of the GSP’s outcome function, we first derive a mathematical representation of the function. Let $R(i, o)$ be the rank of the position given to the i ’th agent in the matching of agents to positions denoted by outcome o . For analysis purposes we will assume, without loss of generality, that each agent’s bid, θ_i , is restricted to be between 0 and 1 (this is not a limiting assumption due to the fact that we can losslessly map from any real valued space to this interval). Under this assumption, the following is functionally equivalent to the GSP’s outcome function.

$$f(\theta^n) = \arg \max_{o \in O} \sum_{i=1}^n \left(\theta_i 10^{-R(i,o)} \right) \quad (3)$$

Each bid in the sum is weighted by 10 raised to the negative power of the corresponding agent’s rank under the chosen outcome. Thus, agents with higher bids will contribute significantly more to the overall sum when they are placed in the first position.

We will now show that the outcome function of the GSP mechanism is inexpressive according to the notion of outcome semi-shattering introduced in the previous section.³

Theorem 2. *Consider a set of outcomes, $\{A, B, C, D\}$, under which agent i is assigned different positions. In the GSP mechanism, agent i cannot semi-shatter both pairs of outcomes $\{A, B\}$ and $\{C, D\}$ if the other agents have more than one joint type and the ranks satisfy $R(i, A) < R(i, C) < R(i, D) < R(i, B)$.*

This result, in conjunction with Theorem 1, implies that under some preference distributions the efficiency bound for the GSP is arbitrarily inefficient, and since it is an upper bound, the inefficiency exists under any strategy profile.

Corollary 1. *For any setting there exists a distribution over agent preferences such that the upper bound on expected efficiency (Equation 2) for the GSP mechanism’s outcome function is arbitrarily less than fully efficient.*

³Proof of Theorem 2 can be found in this paper’s appendix.

4 Premium GSP mechanism

To address GSP’s inexpressiveness without making the mechanism much more cumbersome, we introduce a new mechanism that only slightly increases the expressiveness. Later we show that this slight increase is extremely important in that it removes most of the efficiency loss entailed by GSP’s inexpressiveness.

The new mechanism separates the positions into two classes: premium and standard, and each agent can submit a separate bid for each class. We call this the *premium generalized second price (PGSP)* mechanism. The premium class might contain, for example, only the top position—as in our experiments.

The premium position(s) are assigned as if a traditional GSP were run on the premium bids (the top premium position goes to the agent with the highest premium bid, etc.). The standard positions are then assigned among the remaining agents according to the traditional GSP mechanism run on their standard bids.

5 Computing the efficiency bound

The results in Section 3 prove that there exist distributions over agent preferences for which the GSP is arbitrarily inefficient. However, in order to measure the inefficiency in practice we must be able to compute the value of the efficiency bound for any particular distribution over agent preferences. In this section we describe two general techniques for doing that. They take as input a distribution over agent preferences with a finite number of types (this distribution could be learned from data or approximated by a domain expert) and provide the value of the upper bound on the mechanism’s most efficient equilibrium. Although we present our techniques in the context of ad auctions, they can easily be generalized for use in other domains.

5.1 Integer programming formulation

First we will describe an integer programming formulation for computing the bound. The program includes a binary decision variable, z_o^t , for each outcome and each joint type of the agents. A value of 1 for z_o^t denotes that outcome o will be chosen by the mechanism when the agents have the joint type t , a value of 0 indicates that the outcome will not be chosen under t . The program also includes continuous variables representing the agents’ expressions (bids in the context of sponsored search) under each of their types, $\theta_i^{t_i}$. (We limit these expressions to be between 0 and 1, without loss of generality.) The following objective function is used to maximize the expected efficiency of the mechanism.

$$\max_{z_o^t, \theta_i^{t_i}} \sum_{t \in T^n} P(t) \sum_{o \in O} z_o^t W(t, o) \quad (4)$$

The first set of constraints enforces that exactly one outcome is chosen for each joint type. There are $|T^n|$ such constraints.

$$\text{s.t. } (\forall t \in T^n) \quad \sum_{o \in O} z_o^t = 1 \quad (5)$$

The next set of constraints ensures that for each z_o^t variable that is set to 1, the agents’ expressions under type t do indeed

cause the outcome function to choose outcome o . This set includes one constraint for each joint type and each pair of distinct outcomes. Thus there are $|T^n| \times (|O|^2 - |O|)$ such inequality constraints.⁴ These constraints depend on the outcome function of the mechanism we are studying. For GSP's outcome function, the constraints are as follows (we use M to denote a sufficiently large number such that the sum of all the agents' expressions cannot exceed it):

$$(\forall t, \forall o, \forall o' \neq o) \quad \sum_i \left(\theta_i^{t_i} 10^{-R(i,o)} \right) > \sum_i \left(\theta_i^{t_i} 10^{-R(i,o')} \right) - (1 - z_o^t) M \quad (6)$$

Finally, we have constraints on the decision variables:

$$(\forall t, \forall o) z_o^t \in \{0, 1\}, (\forall i, \forall t_i) 0 \leq \theta_i^{t_i} \leq 1 \quad (7)$$

An ad auction with k positions and n agents with two types each has $\frac{n!}{(n-k)!}$ distinct outcomes and 2^n joint types. The integer program has $|O| \times |T^n|$ binary decision variables, making it prohibitively large for general purpose integer program solvers, such as CPLEX, for mechanisms with more than 3 agents. These solvers do not explicitly take advantage of certain aspects of the problem structure, for example the fact that only one outcome can be chosen for each joint type.

5.2 Tree search for computing the bound

To address this problem, we developed a general tree search technique based on A* for computing the bound. We have applied the technique to GSP and PGSP on instances with up to five agents to find provable inefficiency. (In this paper we only report results with four agents in order to provide a larger number of experiments.)

Each level of the search tree corresponds to a different joint type. Each branch corresponds to the assignment of an outcome to the joint type. The tree has maximum depth $|T^n|$ and branching factor $|O|$. Figure 1 illustrates the search tree.

At any node j a partial assignment of outcomes to joint types can be constructed by traversing the edges from j to the root. We will denote the set of all joint types in the partial assignment at node j as T_j^n . For each type $t_j^n \in T_j^n$ we will denote the outcome it is assigned under the partial assignment at node j as o_{t_j} . In addition, for each joint type t^n we will denote any one of the outcomes that maximize social welfare as o_t^* (i.e., $o_t^* = \arg \max_o W(t^n, o)$).

As usual, our search orders the nodes in its open queue according to an admissible (i.e., optimistic) heuristic. The heuristic approximates the expected efficiency of the best assignment originating from a particular node under the assumption that any unassigned types will be assigned optimally.⁵ The priority of a node j , $\tilde{f}(j)$, is given by the expected welfare of its current partial assignment plus the expected welfare of the optimal assignment for any unassigned types:

$$\tilde{f}(j) = \sum_{t_j \in T_j^n} P(t_j) W(t_j, o_{t_j}) + \sum_{t \notin T_j^n} P(t) W(t, o_t^*) \quad (8)$$

⁴In practice we ensure that these inequality constraints are strict by adding a small ϵ term to one side.

⁵We need only calculate o^* once at the beginning of the search. It can be reused later by removing outcomes that are assigned.




Figure 1: Part of the search tree for a distribution with 3 types, $[t_1^n, t_2^n, t_3^n]$, and 4 outcomes $[A, B, C, D]$. Circles represent internal nodes and squares represent leaf nodes. The dashed nodes are not expanded, but they would be considered by the algorithm. The expanded path corresponds to the assignment of $[A, C, C]$ to types t_1^n , t_2^n , and t_3^n , respectively.

The $\tilde{f}(j)$ approximation is guaranteed to be greater than or equal to the true optimal value of any feasible assignment that descends from node j . It may overestimate this value if the optimal assignment is not achievable due to inexpressiveness, but it has the benefit of serving as a valid upper bound on the expected efficiency achievable by the mechanism. By using the A* node selection strategy, our search ensures that any node that it visits has a lower (or equal) $\tilde{f}$ value than any previously visited node. Thus, the $\tilde{f}$ value of the current node is a continually tightening upper bound on the mechanism's expected efficiency, and it can be provided at any time during the search.

Whenever a node is popped off the front of the open queue, its feasibility is checked. In both types of ad auction mechanisms we study, this check involves solving a linear feasibility problem (LFP). The LFP involves a set of constraints similar to those described in Equation 6, however the assignment of outcomes to types is fixed and there are no binary decision variables. If the node is not feasible, its children are not placed on the open queue.

6 Experiments

In this section we discuss the results of experiments using our search technique to compute the upper bound for the GSP mechanism and the slightly more expressive PGSP mechanism.

In order to gain additional insight, we also discuss the performance of the two mechanisms when agents use the straightforward strategy of always bidding their valuation for the top position (in the PGSP they bid their valuation for the top premium position and the top non-premium position as their two bids). We call the resulting efficiency *GSP heuristic* and *PGSP heuristic*, respectively.

Our experiments consist of collections of runs, each involving randomly generated instances with different parameter settings. The parameters are chosen to investigate circumstances under which the inexpressiveness of the GSP mechanism is costly (i.e., when the upper bound is low) and when

it is not. Each instance in one of our experiments represents a single auction for a single keyword with four agents.

Based on recent work examining different advertising attitudes on the Internet, in our experiments each agent is either a *brand advertiser* (with probability p_B) or a *value advertiser* (with probability $1 - p_B$) [Baye and Morgan, 2005]. Brand advertisers always prefer higher positions over lower ones. A value advertiser generally does not prefer the highest positions because middle positions tend to have higher *conversion rates* (e.g., the user’s probability of buying something conditional on having clicked is higher). Figure 2 illustrates prototypical brand and value preferences over different positions based on their rank.




Figure 2: Example of prototypical valuations for brand and value advertisers. The brand advertiser shown has $\mu = 1$ and the value advertiser has $\mu = 0.5$. Valuations are shown in expectation, not per-click. Rank 0% means the bottom position and Rank 100% means the top position.

We now describe how we generate preferences for brand and value advertisers. Let “clk” denote the event that the ad is clicked, and “cnv” denote that the click results in a conversion (e.g., a sale or user registration). Let C_i denote the amortized cost per click of running agent i ’s web site, and $V_i(\text{cnv})$ be the expected value of a conversion to agent i . Then, the expected value to agent i of having an ad in position ranked R is

$$E[V_i(R)] = P(\text{clk}|R, i) [P(\text{cnv}|\text{clk}, R, i)V_i(\text{cnv}) - C_i] \quad (9)$$

In order to keep the experiments simple and to focus on the impact of expressiveness, we assume that agents in the same instance are relatively similar. For one, we assume that the marginal cost of a click $C_i = C = \$1$ for all agents. Unless otherwise specified, we assume that $V_i(\text{cnv}) = V(\text{cnv}) = \50 for all agents. We assume that $P(\text{cnv}|\text{clk}, i) = P(\text{cnv}|\text{clk}) = 10\%$ for all agents. We also assume that click-through rates conditional on the rank of an ad’s position are the same for all agents. The specific rates are given in Table 1, along with the default values for all parameters. These click-through rates are from an Atlas Institute Digital Marketing publication [Brooks, 2007]. They were also used by Even-Dar *et al.* in their experiments [Even-Dar *et al.*, 2007].

Rather than generating arbitrary values of $P(\text{cnv}|\text{clk}, R, i)$, we assume that the probability of a conversion coming from a particular rank, $P(R|\text{cnv}, i)$, is normally distributed. The mean, μ , of this distribution is randomly chosen from $[0, 1]$

Parameter	Value	Parameter	Value
$P(\text{clk} R = 1)$	10%	$C(\text{clk})$	\$1
$P(\text{clk} R = 2)$	7.74%	Brand μ	$\sim \text{Uniform} [.8, 1]$
$P(\text{clk} R = 3)$	6.66%	Brand σ	25% of μ
$P(\text{clk} R = 4)$	5.74%	Value μ	$\sim \text{Uniform} [.4, .6]$
$P(\text{cnv} \text{clk})$	10%	Value σ	25% of μ
p_B	50%	$V_i(\text{cnv})$	\$35 to \$150

Table 1: Default settings for each parameter in our instance generation model.

for each agent, once for the case where she is a brand advertiser and once for the case where she is a value advertiser. (We also normalize the value of R to be between 0 and 1, so that, for example, the third position out of four has rank 0.25.) Values of μ closer to 1 indicate that the agent’s conversions tend to come from higher ranked ads, those closer to 0 indicate that conversions tend to come from lower ranked ads. The values of μ for the brand and value advertisers are given in Table 1, unless otherwise specified.

We transform $P(R|\text{cnv}, i)$ into $P(\text{cnv}|\text{clk}, R, i)$ using Bayes’ rule (and the observation that the cnv event implies the clk event):

$$P(\text{cnv}|\text{clk}, R, i) \propto P(R|\text{cnv}, i)P(\text{cnv}|\text{clk}, i) \quad (10)$$

Each data point in each figure below is the average over 50 instances.⁶ The confidence intervals represent standard error. (They are often so tight that they are barely visible.)

6.1 Experiment 1: Varying agents’ profit margin

In our first set of results we vary the expected value of a conversion, $V(\text{cnv})$, between \$35 and \$150 (i.e., 35 to 150 times the cost per click of running the site), Figure 3.




Figure 3: The value of the upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms.

These results demonstrate that when conversions generate relatively low profits, the efficiency loss due to inexpressiveness in the GSP mechanism, as measured by the upper bound, is more than 30%. As the profit margin of the agents increases, this loss decreases to around 10%.

⁶On occasion the search does not find the optimal value within our time limit of 20 minutes; this occurred on approximately 25% of the instances evenly distributed between mechanisms. For those instances we use the lowest $\tilde{f}$ value discovered prior to termination.

Additionally, the results show that the efficiency bound for the slightly more expressive PGSP mechanism is nearly 100% in all cases. This suggests that the added expressiveness in the PGSP is well suited to capture all the different types of preferences we generated.

We also see that the efficiency of the heuristic bidding strategy follows a similar qualitative pattern to the upper bound, which lends additional support to our findings. Specifically, this suggests that 1) the bound is meaningful in describing the efficiency of the mechanism, and 2) the conclusions apply more broadly than for fully rational game-theoretic agents.

6.2 Experiment 2: Varying agent diversity

The second experiment examines how the loss due to inexpressiveness depends on how similar value advertisers are to brand advertisers. Specifically, we vary the position that generates the most value for value advertisers. (Brand advertisers still always prefer the highest position the most.) In each run the mean of $P(R|cnv, i)$ for each value advertiser is drawn uniformly from an interval of size 0.2 (i.e., $\mu \sim \text{Uniform}[a, a + 0.2]$). The results are shown in Figure 4. The x-axis indicates the mid-point of the interval used in each run, which is also the expected value of μ for each value advertiser.




Figure 4: The value of our upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms. Large values of $E[\mu]$ correspond to runs in which higher ranking positions are more valuable for the value advertisers and vice versa.

These results demonstrate that the need for additional expressiveness is greatest when the value advertisers prefer middle ranking positions, as is typically the case in practice. For example, when those agents prefer the middle rank, the GSP can achieve at most 85% efficiency (with the heuristic bidding strategy achieving less than 75%) on average, whereas the PGSP can achieve over 95% (with the heuristic bidding strategy achieving about 85%). The expressiveness is less necessary when the value advertisers become more like the brand advertisers (i.e., large $E[\mu]$) or when they are drastically different than the brand advertisers (i.e., small $E[\mu]$).

Again the efficiency of the heuristic bidding strategy follows a similar qualitative pattern to the upper bound.

7 Conclusions and future research

A key trend on the Internet is a move toward more expressiveness in mechanisms. Yet, adding expressiveness does not

always improve the outcome of a mechanism in practice. In this paper we operationalized a recent theoretical framework for studying expressiveness with a methodology for comparing mechanisms with different degrees and forms of expressiveness, and applied it to sponsored search.

We began by proving that for some preference distributions the most commonly used sponsored search mechanism, GSP, is arbitrarily inefficient. In order to measure the inefficiency in practice we developed a general tree search technique for computing an upper bound on a mechanism’s expected efficiency. We concluded with a series of experiments comparing the GSP to our slightly more expressive mechanism, PGSP, which solicits an extra bid for premium ad positions. We generated a range of realistic preference distributions based on published industry knowledge, and applied our search technique to compare the efficiency bounds in the two mechanisms. We also examined the performance of the mechanisms when agents use a straightforward heuristic bidding strategy.

Our results suggest that the GSP’s efficiency loss due to inexpressiveness can be dramatic. It is greatest in the practical case where some agents (“brand advertisers”) prefer top positions while others (“value advertisers”) prefer middle positions. The loss is also worst when agents have small profit margins. Despite the fact that our PGSP mechanism is only slightly more expressive (and thus not much more cumbersome), it removes almost all of the efficiency loss in all of the settings we study.

Future research includes using our methodology to study efficiency in sponsored search with real data. We also plan to apply our methodology to other domains in order to design mechanisms that are not unnecessarily expressive yet remove most of the inefficiency of today’s inexpressive mechanisms.

8 Acknowledgement

This work has been supported by the National Science Foundation under grants IGERT-9972762, ITR-0205435, IIS-0427858, and CNS-0627513.

References

- [Abrams *et al.*, 2007] Zoe Abrams, Arpita Ghosh, and Erik Vee. Cost of conciseness in sponsored search auctions. In *Proceedings of Workshop on Internet Economics (WINE)*, 2007.
- [Baye and Morgan, 2005] Michael Baye and John Morgan. Brand and price advertising in online markets. Technical Report 0504009, EconWPA, 2005.
- [Benisch *et al.*, 2008] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. Theory of expressiveness in mechanisms. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 2008. Extended version: CMU tech report CMU-CS-07-178, 2007.
- [Boutilier *et al.*, 2008] Craig Boutilier, David Parkes, Tuomas Sandholm, and William Walsh. Expressive banner ad auctions and model-based online optimization for clearing. In *AAAI*, 2008.

- [Brooks, 2007] N. Brooks. The Atlas rank report: How search engine rank impacts traffic., 2007. Available at <http://atlassolutions.com/> in May 2008.
- [Even-Dar *et al.*, 2007] Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Sponsored search with contexts. In *WINE*, 2007.
- [Hurwicz, 1972] Leonid Hurwicz. On informationally decentralized systems. In C.B McGuire and R. Radner, editors, *Decision and Organization*. Amsterdam: North Holland, 1972.
- [Lahaie *et al.*, 2008] Sébastien Lahaie, David C. Parkes, and David M. Pennock. An expressive auction design for on-line display advertising. In *AAAI*, 2008.
- [Martin *et al.*, 2008] David J. Martin, Johannes Gehrke, and Joseph Y. Halpern. Toward expressive and scalable sponsored search auctions. In *Proceedings of Conference on Data Engineering*, 2008.
- [Milgrom, 2007] Paul Milgrom. Simplified mechanisms with applications to sponsored search and package auctions. Mimeo, 2007.
- [Mount and Reiter, 1974] Kenneth Mount and Stanley Reiter. The informational size of message spaces. *J. of Economic Theory*, 8(2):161–192, 1974.
- [Parkes and Sandholm, 2005] David Parkes and Tuomas Sandholm. Optimize-and-dispatch architecture for expressive ad auctions. In *Proceedings of Workshop on Sponsored Search Auctions*, 2005.
- [Rosenthal and Wang, 1996] Robert W. Rosenthal and Ruqu Wang. Simultaneous auctions with synergies and common values. *Games and Economic Behavior*, 17:32–55, 1996.
- [Sandholm, 2007] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.
- [Schwartz, 2004] Barry Schwartz. *The Paradox of Choice: Why More is Less*. Ecco/Harper Collins Publishers, 2004.
- [Simon, 1955] Herbert A Simon. A behaviorial model of rational choice. *Quarterly Journal of Economics*, 69:99–118, 1955.
- [Szentes and Rosenthal, 2003] Balazs Szentes and Robert W. Rosenthal. Beyond chopsticks: Symmetric equilibria in majority auction games. *Games and Economic Behavior*, 45:278–295, 2003.
- [Vapnik and Chervonenkis, 1971] Vladimir Vapnik and Alexey Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [Varian, 2007] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, pages 1163–1178, 2007.
- [Wilenius and Andersson, 2007] Jim Wilenius and Arne Andersson. Discovering equilibrium strategies for a combinatorial first price auction. In *Proceedings of E-Commerce Technology (CEC’07)*, 2007.

9 Appendix

Proof of Theorem 2. We will assume for contradiction that agent i can semi-shatter both pairs of outcomes, $\{A, B\}$ and $\{C, D\}$. First we restate the following lemma from our earlier work.

Lemma 1 (reworded from [Benisch *et al.*, 2008]). *Agent i can (semi-)shatter an outcome space $\mathcal{O}'$ when the agents other than i have more than one joint type only if there exists at least one pair of expressions by the other agents, $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$, which allows i to (semi-)shatter $\mathcal{O}'$.*

This lemma implies that there must be at least one pair of bids by the agents other than i , $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$, such that agent i can cause all four outcomes to happen by changing its own bid alone.

Let the weighted sum of the bids of the agents other than i for the first (second) profile under outcome A be a_1 (a_2), under outcome B be b_1 (b_2), and so on. Also, let the weights on agent i ’s bid under outcomes A through D in the the GSP outcome function be α_A through α_D . (Note that the predicate of our theorem implies that $\alpha_A > \alpha_C > \alpha_D > \alpha_B$.)

Let us assume (without loss of generality) that $b_1 - a_1 < b_2 - a_2$ and that A will happen against $\theta_{-i}^{(1)}$ and B will happen against $\theta_{-i}^{(2)}$. In order to cause A to happen against the first opponent profile and B against the second the following inequalities must hold (we assume that ties are broken consistently so that an agent cannot use them to semi-shatter:

$$A \text{ happens against } 1 \begin{cases} \alpha_A \theta_i + a_1 > \alpha_B \theta_i + b_1 \\ \alpha_A \theta_i + a_1 > \alpha_C \theta_i + c_1 \\ \alpha_A \theta_i + a_1 > \alpha_D \theta_i + d_1 \end{cases}$$

$$B \text{ happens against } 2 \begin{cases} \alpha_B \theta_i + b_2 > \alpha_A \theta_i + a_2 \\ \alpha_B \theta_i + b_2 > \alpha_C \theta_i + c_2 \\ \alpha_B \theta_i + b_2 > \alpha_D \theta_i + d_2 \end{cases}$$

By simplifying the above equations we derive the following set of constraints.

$$\begin{aligned} \frac{c_1 - a_1}{\alpha_A - \alpha_C} &< \theta_i < \frac{b_2 - d_2}{\alpha_D - \alpha_B} \\ \frac{d_1 - a_1}{\alpha_A - \alpha_D} &< \theta_i < \frac{b_2 - c_2}{\alpha_C - \alpha_B} \end{aligned}$$

In order to semi-shatter C and D we have the following inequalities generated in the same fashion,

$$\begin{aligned} \frac{b_2 - d_2}{\alpha_D - \alpha_B} &< \theta_i < \frac{c_1 - a_1}{\alpha_A - \alpha_C} \\ \frac{b_2 - c_2}{\alpha_C - \alpha_B} &< \theta_i < \frac{d_1 - a_1}{\alpha_A - \alpha_D} \end{aligned}$$

Now we can see that our assumption that agent i could semi-shatter both sets of outcomes when the other agents have more than a single type leads to a contradiction. $\square$