Truthful stochastic and deterministic auctions for sponsored search*

Esteban Feuerstein † Pablo Ariel Heiber † Marcelo Mydlarz ‡

† *Depto. de Computación, FCEyN*
Univ. de Buenos Aires, Argentina
{efeuerst,pheiber}@dc.uba.ar

‡ *Yahoo! Research, Santiago, Chile*
marcem@yahoo-inc.com

Abstract

Incentive compatibility is a central concept in auction theory, and a desirable property of auction mechanisms. In a celebrated result, Aggarwal, Goel and Motwani [2] presented the first truthful deterministic auction for sponsored search (i.e., in a setting where multiple distinct slots are auctioned).

Stochastic auctions present several advantages over deterministic ones, as they are less prone to strategic bidding, and increase the diversity of the winning bidders. Meek, Chickering and Wilson [10] presented a family of truthful stochastic auctions for multiple identical items.

We present the first class of incentive compatible stochastic auctions for the sponsored search setting. This class subsumes as special cases the laddered *auctions of [2] and the stochastic auctions with the* condex *pricing rule of [10], consolidating these two seemingly disconnected mechanisms in a single framework. Moreover, when the price per click depends deterministically on the bids the auctions in this class are unique. Accordingly, we give a precise characterization of* all *truthful auctions for sponsored search, in terms of the expected price that each bidder will pay per click.*

We also introduce randomized algorithms and pricing rules to derive, given an allocation mechanism for the single- or multiple-identical-slots scenarios, a new mechanism for the multislot framework with distinct slots. These extensions have direct practical applications.

1. Introduction

A stochastic auction is a special kind of auction in which the winner is determined in a stochastic or randomized way. Naturally, by raising her bid a player should increase her chances of winning the auctioned good.

In this paper we consider a setting where many distinct items are simultaneously auctioned, and each participant bids a single amount that is interpreted as her bid for *any* of the auctioned items. However, no bidder can win more than one item. Finally, we assume that the relative values of the items are shared by all the players.

The previous assumptions are motivated by the framework of sponsored search and contextual advertisement. Given a query, search engines (like Yahoo!, Google or MSN) respond by presenting (ideally) the most relevant results, together with a set of ads. Similarly, in contextual advertisement, the ads are displayed within a web page that contains some specified terms. Usually, many advertisers compete for a limited number of slots available for these ads. This kind of advertising is continuously growing and has become the main source of revenue for many of the participants in this market.

Each advertiser places one bid, and the auctioneer decides, based on the bids and other public or private parameters, which ads will be published in which slot. The winning advertisers will pay a price established by the auctioneer each time a user clicks on their ads. Note that each of the slots may have different value for the advertisers. Indeed, there is an *attention decay* model which reflects the usual pattern of the users' behavior: people tend to click more on ads positioned at higher slots. Consequently, advertisers prefer slots with higher potential click-through-rate (CTR). Nevertheless, the position has usually no direct influence on the price paid (since all clicks are assumed to have the same expected revenue for the advertiser, independently of the slot where they originate).

The use of stochastic auctions for sponsored search has been recently addressed in [10], [6] and [3]. Stochastic auctions for the similar environment of multiple items with unlimited supply have been also considered in [7]. We summarize the main motivations for using them in place of deterministic ones:

- Stochastic auctions are less prone to vindictive and/or strategic bidding.

*This research was funded by a Yahoo! Research Alliance Grant.

978-0-7695-3397-1/08 $25.00 © 2008 IEEE
DOI 10.1109/LA-WEB.2008.17

- The fact that *anyone* can win the auction contributes to have a wider advertisers base, and therefore higher revenue in the medium term (e.g., see [9], section 8.2, and [4]).

- The increase in variety of the ads published brings about advantages, such as improved user experience and greater aggregate click through rates, due to better coverage of the possible intent and needs of the user.

- In order to avoid leaving out ads with high potential revenue –due to estimation errors, or simply lack of information–, there is a need to alternate among ads with high, small and unknown revenue expectation. This is known as the explore/exploit trade-off [12]. Stochastic auctions, in which each bidder has some probability of winning, provide an implicit way to implement this trade-off [6].

- Stochastic mechanisms are in general less vulnerable to fraudulent behavior (click fraud, [13]).

In [10], Meek, Chickering and Wilson propose a pricing rule to make stochastic auctions incentive compatible (truthful), and generalize Vickrey auctions [14], showing that the advantages of stochastic auctions can coexist with a pricing mechanism in which bidders have an incentive to bid truthfully their respective values. Their results, however, only apply to auctions of single or multiple identical items. That setting is not general enough to cover the typical sponsored search framework in which, as we mentioned before, it is normally assumed that many *distinct* slots are auctioned simultaneously, each of them having its own "position-CTR", that is, a factor that reflects the decay of users' attention. This position-CTR is modeled by a weight that is associated to each slot. Slots with higher weights are preferable: any given ad will be more likely to receive a click there.

One of the contributions of this paper is the first family of truthful stochastic auctions for sponsored search, together with a significant (although not immediate) implication: a unified mechanism generalizing the stochastic auctions of [10] to multiple-distinct-slots auctions, and the deterministic *laddered* auctions of [2, 5]. Laddered auctions are defined for certain deterministic allocation mechanisms (based on *ranking functions*), that we extend to a broader class. We also present a procedure which transforms any stochastic auction into an equivalent auction (in terms of expected revenue of the auctioneer and each of the bidders) where the price charged to each bidder is a deterministic function on the bids. In this way, a family of representative or canonical auctions is defined, which we call deterministic-price auctions. We then prove that, in the same way that laddered auctions determine the only possible truthful pricing scheme for deterministic ranking allocations, our mechanism is unique within the family of deterministic-price auctions. This leads to a general and purely arithmetic characterization of all truthful auctions. These contributions, as well as those in [10] and [2], may be seen as consequences of traditional results in microeconomics [8, 11] on how to design a truthful pricing rule given a certain allocation rule. Concisely, those very general results state that under certain restrictions of the type space (the space of the offers that bidders can make), given a fixed allocation rule and a payoff to the "last" bidder, there is a unique pricing rule that gives a truthful mechanism. Our results, however, are derived specifically and explicitly define techniques for the framework of sponsored search, which has its own peculiarities (multiple different slots auctioned simultaneously, payments subject to the occurrence of some contingency, etc.). We provide constructive proofs and define all technical details needed in order to use them in practice, and instantiate them in that particular world.

Another contribution of our work is a framework that, given any allocation mechanism, derives a truthful auction for sponsored search. We explore and evaluate some techniques to derive auctions for sponsored search given some simpler mechanism. This provides some insight and tools for further research in this direction. In particular, we show how to obtain a truthful auction for multiple distinct slots starting from a stochastic allocation scheme for either a single slot or multiple identical slots. In contrast with previously known truthful auctions, some of the auctions we propose do not need the values of position-CTRs in order to be implemented. In fact, while there is some consensus on the decay model followed by sponsored search, the actual parameters may be unknown, or vary across search terms. The resulting scheme, however, may achieve lower social welfare compared to an hypothetical one with knowledge of those values.

Finally, we present a drawing algorithm that provides an assignment of slots to bidders according to any set of rational probabilities.

In order to round up the introduction, we briefly mention other relevant recent work on stochastic auctions. Goldberg, Hartline, Karlin, Saks and Wright [7] address the question of designing stochastic competitive truthful auctions (i.e. auctions with provable revenue guarantees) for identical items in unlimited supply. Abrams and Gosh [1] follow this line, extending those ideas to sponsored search, proving that it is necessary to renounce to truthfulness to achieve competitiveness. A bidding heuristic in which small random perturbations are introduced to the bids to avoid cycling is considered in [3]. The result is that the heuristic converges in first- and second-price auctions to interesting equilibria in which bidders "share" items in a certain way. Finally, [6] introduces a family of stochastic allocation mechanisms whose aim is to increase diversity and avoid some classes of fraud, while not losing revenue significantly.

2. Assumptions and notation

The setting we consider involves a finite number, n, of risk-neutral bidders that compete for k slots. No bidder is allowed to win more than one slot. We assume that the number of slots is not greater than the number of bidders, that is, $k \leq n$.

We assume, as it is generally the case in the literature [2, 3, 5, 6], that the CTR can be separated into two factors, one advertisement-specific, the *ad-CTR* and the other position-specific, the *position-CTR*. This is called the "separability" of the CTR [2]. The ad-CTR of advertiser i will be denoted by c_i. Each slot j has an associated weight w_j, which may be interpreted as the click probability associated with the slot, or equivalently, the position-CTR. Thus, the expected number of clicks that ad i appearing in slot j will receive is exactly $w_j c_i$. For convenience, we assume without loss of generality that the weights are normalized in such a way that $1 = w_1 \geq w_2 \geq \ldots \geq w_k$.

We mainly follow the notation of [10]. In that article, all the auctioned items are equal, which corresponds to the case where $w_1 = w_2 = \ldots = w_k = 1$. The non-negative real-valued bids are denoted by $b_1, b_2, \ldots, b_n$, which we shorten by $\mathbf{b}$. Finally, $p_i(x)$ denotes the probability that bidder i wins exactly one item when bidding x; while this function clearly depends on $b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n$, we omit these bids in the notation, hoping to make our presentation easier to follow.

For an allocation rule following probability functions p_i, [10] defines the *condex* pricing rule for the auction as follows. The per-click price for bidder i is $\mu_i(\mathbf{b}) = \frac{\int_0^{b_i} x \; dp_i(x)}{p_i(b_i)}$. The name condex is short for conditional expectation: the price charged to bidder i is the expected value of her minimum winning bid given that she won the auction by bidding b_i.

As with the probability functions p_i, we will abuse notation and denote by $\mu_i(b_i)$ the price per click charged to bidder i when bidding b_i, omitting the bids of the other participants whenever they are fixed in the context. It is shown in [10] that the auction just described is incentive compatible (resp. strictly incentive compatible) if the functions p_i are non-decreasing (resp. strictly increasing). For the remainder of this paper we refer to this result as the "MCW Theorem". Finally, each bidder i has a private value v_i, reflecting how much a click on her ad is worth for her.

Due to space limitations, some of the proofs of the Theorems are presented in Section C of the Appendix.

3. Deterministic-price auctions

In a stochastic auction, the allocation rule, the price or both may be (independent or correlated) random variables. This implies, for example, that any given bidder may be charged different prices per click depending on which slot is assigned to her, but also the price she is charged for a click in the same slot could be a random variable. Examples of such auctions are given in [7], where even if all the auctioned items are identical, the price to be paid is determined in each opportunity as a result of some coin tosses.

We call the auctions where the price charged to each bidder follows deterministically from the set of bids *deterministic-price* auctions. Note that in a deterministic-price auction the slot allocated to any bidder is still a random variable.

Deterministic-price auctions are interesting because, apart from being easier to understand by the users, they are more predictable and auditable. Indeed, even if it may be unknown where an ad will be displayed, or whether it will be displayed at all, the price that will eventually be paid for a click is always known in advance. As mentioned in the introduction, sponsored search auctions used nowadays are deterministic-price.

We now show how to transform *any* auction into a deterministic-price auction without changing its allocation rule or its expected revenue, neither for the advertisers nor the auctioneer. This in turn implies that if this transformation is applied to a stochastic truthful auction it gives an "equivalent" stochastic truthful auction that is also deterministic-price. Instead of charging random prices, we can charge deterministically the expected price of each bidder, which yields the same result in expectation. Later on we make use of this procedure in the characterization of all truthful auctions for sponsored search.

Let A be a stochastic auction and x be bidder i's bid. We denote by $M_i(x)$ and $W_i(x)$ the random variables representing the price that bidder i will pay if her ad is clicked, and the weight of the slot allocated to her, respectively.

The expected amount that bidder i will pay *per impression* is given by the expression $E[M_i(x)W_i(x)c_i]$ (recall that c_i is i's ad-CTR). Since ad i's probability of a click in any impression is $E[W_i(x)c_i]$, it follows that bidder i's expected price per click is

$$\overline{\mu}_i(x) = \begin{cases} 0 & \text{if } E[W_i(x)] = 0 \\ \frac{E[M_i(x)W_i(x)]}{E[W_i(x)]} & \text{otherwise.} \end{cases}$$

Now we can define $D(A)$ as the auction with the same allocation rule as A, and pricing rule $\overline{\mu}_i$. It is clear that $D(A)$ is a deterministic-price auction.

Since the expected revenue of any auction is the sum over all bidders of their expected price per click times their probabilities of a click ($\sum_{i=1}^{n} \overline{\mu}_i E[W_i(x)]c_i$), $D(A)$ preserves the auctioneer's expected revenue.

Likewise, the expected revenue per click of bidder i under any auction is given by her valuation minus her expected price per click, that is, $v_i - \overline{\mu}_i E[W_i(x)]$. Thus, $D(A)$ preserves the expected revenue of each bidder. We record these two facts in the following lemma.

Lemma 1 *The expected revenue of the auctioneer under auctions A and $D(A)$ coincide. The expected revenue of each of the bidders under auctions A and $D(A)$ coincide as well.*

Lemma 1 implies that $D(A)$ preserves desirable properties; e.g., truthfulness.

4. Truthful stochastic auctions for distinct slots

We turn now into the design of truthful stochastic auctions for sponsored search. We will show necessary and sufficient conditions on the allocation rule for the existence of a truthful auction, and provide a mechanism to compute its pricing rule.

When the weights of the auctioned slots are equal, an allocation rule is given by the set of probability functions $p_i(x)$. Recall that $p_i(x)$ is the probability that bidder i is allocated one (any) slot. Accordingly, when slots have potentially different weights, an allocation mechanism follows implicitly from a set of functions $p_i^j(x)$, each denoting the probability that bidder i wins slot j when bidding x. The mechanisms we consider must ensure that each slot is assigned to some bidder, that is, $\sum_{i=1}^{n} p_i^j(b_i) = 1$ for all j.

Given an allocation rule as a set of probability functions p_i^j, we define $q_i(x) = \sum_{j=1}^{k} w_j p_i^j(x)$, the expected position-CTR bought by bidder i when bidding x (in terms of the notation introduced in the previous section, the random variable $W_i(x)$ has value w_j with probability $p_i^j(x)$, and $q_i(x) = E[W_i(x)]$). In this context, we define the *condex* pricing for the auction by $\mu_i(\mathbf{b}) = \frac{\int_0^{b_i} x \; dq_i(x)}{q_i(b_i)}$.

Theorem 1 *If an auction for multiple distinct slots has a non-decreasing (resp. strictly increasing) expected position-CTR and the corresponding condex pricing rule, then the auction is incentive compatible (resp. strictly incentive compatible).*

In the next theorem we show that requiring a non-decreasing (resp. strictly increasing) expected position-CTR is a necessary condition in any incentive-compatible auction. This justifies the preconditions required by Theorem 1. Something similar holds for identical items in unlimited supply [7], and for the premise of the MCW Theorem. All these results may be seen as variants of Myerson's results [11], as they establish some form of monotonicity of the allocation function as a necessary and sufficient condition for any incentive compatible mechanism.

Theorem 2 *Any incentive compatible (resp. strictly incentive compatible) auction requires a non-decreasing (resp. strictly increasing) expected position-CTR function.*

Table 1 in Appendix B shows an example of an application of the condex pricing rule to one allocation function.

Theorem 1 generalizes two well-known pricing mechanisms which create truthful auctions when coupled with their corresponding allocation mechanisms.

The first one is the MCW Theorem, which considers the special case where all auctioned items are equal, that is, $1 = w_1 = w_2 = \cdots = w_k$. Of course this yields that $q_i(x) = \sum_{j=1}^{k} w_j p_i^j(x) = \sum_{j=1}^{k} p_i^j(x)$, which may be regarded as bidder i's probability of winning exactly one item (since they all have the same value, there is no need to distinguish them). It is clear that this is exactly $p_i(x)$, and since the price expression in both cases coincides, it follows that Theorem 1 is in fact a generalization of that result.

The second result we generalize refers to deterministic auctions.

4.1. Deterministic auctions

The most frequently used family of deterministic auctions ranks the ads according to *ranking functions* $f_i(x)$, which assign "ranking points" to bidder i, considering not only her bid x, but also some potentially relevant properties of i (e.g., i's ad-CTR). The top-ranked ads are then assigned to the slots, higher ranked ads to slots with higher position-CTR. We call these auctions *deterministic ranking auctions*.

We assume that each f_i is non-decreasing and that, without loss of generality, bidders are ordered in such a way that $f_h(b_h) \geq f_{h+1}(b_{h+1})$. Now we define an inverse-like function $f_i^{-1}(y)$ as the minimum amount bidder i needs to bid in order to obtain at least y ranking points, that is,

$$f_i^{-1}(y) = \inf(\{x | f_i(x) \geq y\}).$$

Note that when f_i is bijective, f_i^{-1} is effectively its inverse. Given these ranking functions f_i, and setting $w_j = 0$ for $j > k$, we define the *extended laddered* pricing for bidder i:

$$\mu_i(b_i) = \sum_{j=i}^{k} \frac{w_j - w_{j+1}}{w_i} f_i^{-1}(f_{j+1}(b_{j+1})).$$

Theorem 3 *If a deterministic ranking auction for multiple distinct slots has a non-decreasing set of ranking functions and the corresponding extended laddered pricing rule, then the auction is incentive compatible.*

Proof: Deterministic auctions may be seen as stochastic auctions in which probabilities are either 0 or 1. First, for each bidder i we define the other bidders' scores $s_1, \ldots, s_{n-1}$ by

$$s_h = \begin{cases} f_h(b_h) & \text{for } h = 1, \ldots, i-1 \\ f_{h+1}(b_{h+1}) & \text{for } h = i, \ldots, n-1. \end{cases}$$

Now we can define p_i^j as needed in Theorem 1: [1]

$$p_i^j(x) = \begin{cases} 1 & \text{if } j = 1 \text{ and } f_i(x) > s_1 \\ 1 & \text{if } j > 1 \text{ and } s_{j-1} \geq f_i(x) > s_j \\ 0 & \text{otherwise.} \end{cases}$$

Then, the expected position-CTR defined in Theorem 1 becomes

$$q_i(x) = \begin{cases} w_1 & \text{if } f_i(x) > s_1 \\ w_j & \text{if } s_{j-1} \geq f_i(x) > s_j \\ 0 & \text{otherwise.} \end{cases}$$

Note that if f_i is non-decreasing, then q_i is non-decreasing as well. Moreover, since q_i is constant in all but a finite set of points (the set $\{f_i^{-1}(s_1), ..., f_i^{-1}(s_{n-1})\}$) and f_i^{-1} is non-decreasing by definition, μ_i defined in Theorem 1 becomes

$$\begin{aligned} \mu_i(b_i) &= \frac{1}{q_i(b_i)} \int_0^{b_i} x \, dq_i(x) \\ &= \frac{1}{w_i} \sum_{\{j | s_j < f_i(b_i)\}} f_i^{-1}(s_j) \, (w_j - w_{j+1}) \\ &= \sum_{j=i}^{n-1} \frac{w_j - w_{j+1}}{w_i} f_i^{-1}(s_j) \\ &= \sum_{j=i}^{k} \frac{w_j - w_{j+1}}{w_i} f_i^{-1}(f_{j+1}(b_{j+1})), \end{aligned}$$

where the first equality arises directly from the definition of the Riemann-Stieltjes integral. By Theorem 1, the resulting auction is incentive compatible, which is what we wanted to prove. □

A practical consequence of this result is that the pricing for deterministic cases has a simple and easy to compute formula (no integrals involved), provided that we are able to compute the inverse-like function, which is usually the case.

Theorem 3 generalizes the result in [2], which gives a pricing mechanism for deterministic auctions for potentially different slots. Concretely, [2] shows that a deterministic auction that ranks bidders according to scores of the form $u_i b_i$ is indeed truthful if the price paid for a click on ad i is

$$\sum_{j=i}^{k} \left(\frac{CTR_{i,j} - CTR_{i,j+1}}{CTR_{i,i}} \right) \frac{u_{j+1}}{u_i} b_{j+1},$$

where $CTR_{i,j}$ is the CTR of ad i when placed in slot j. The same can be proved using Theorem 3 with $f_i(b_i) = u_i b_i$, because under the separability assumption $CTR_{i,j} = c_i w_j$, where c_i is the ad-CTR of ad i. It follows immediately that the expressions for the price coincide.

5. Characterizing truthful auctions

Given the importance of truthfulness in auction theory, having a characterization of truthful auctions is highly desirable. It may be useful for both testing existing mechanisms or developing new ones, as well as for deriving properties of truthful auctions. In this section, building upon our results on sections 3 and 4, we provide such a characterization.

In Theorems 1 and 2 we have shown that monotonicity of the allocation rule is a necessary and sufficient condition for the existence of a truthful auction. We now show that, given such an allocation rule, the condex pricing rule is the only rule that makes the auction truthful and deterministic-price. Moreover, we show that the condex auction is particular in a stronger sense: for any truthful auction, the expected price per click must coincide with the condex price. This completes a strong and purely arithmetical characterization of truthful auctions for sponsored search in terms of the condex pricing rule.

Theorem 4 *Given an allocation rule with non-decreasing expected position-CTR, there exists exactly one pricing such that the auction is truthful and deterministic-price.*

Proof: Note that, whenever the expected position-CTR is 0 the price is irrelevant, so we set the price to 0 in those cases. As we saw in Theorem 1, the condex pricing makes the auction truthful and deterministic-price, so, it only remains to prove uniqueness. Note first that any pricing function μ corresponding to a truthful auction must satisfy

$$\mu(x) \leq x \text{ for all } x > 0. \quad (1)$$

Let i be an arbitrary bidder and assume that the bids of the others are fixed. For a bid x of bidder i, let $q_i(x)$ be her expected position-CTR; let $\mu_i(x)$ and $\nu_i(x)$ be two pricing rules, both yielding truthful auctions. By hypothesis the auction is truthful, so for all $v > 0$

$$\begin{aligned} v &\in \operatorname*{argmax}_x (v - \mu_i(x)) q_i(x), \text{ and} \\ v &\in \operatorname*{argmax}_x (v - \nu_i(x)) q_i(x) \end{aligned} \quad (2)$$

By way of contradiction, let as assume that μ_i and ν_i are different; without loss of generality we may assume that there exists a positive t such that

$$\mu_i(t) > \nu_i(t). \quad (3)$$

[1] This formula is not absolutely precise when ties are possible (although the theorem still holds true for those cases), yet taking this small liberty yields a much cleaner proof.

Let c be such that

$$\mu_i(t) - \nu_i(t) > c > 0 \text{ and } t/c > \lfloor t/c \rfloor .$$

Define k by

$$k = \max\{h \in \{0, \ldots, \lfloor t/c \rfloor\} : \mu_i(t-hc) - \nu_i(t-hc) > c\}$$

Now, $k = \lfloor t/c \rfloor$ implies $\mu_i(t - \lfloor t/c \rfloor c) > c$, but this is not feasible due to (1) and $t - \lfloor t/c \rfloor c < c$. Thus, $k < \lfloor t/c \rfloor$, implying that

$$\mu_i(t - kc) - \nu_i(t - kc) > c \quad (4)$$

$$\mu_i(t - (k+1)c) - \nu_i(t - (k+1)c) \le c \quad (5)$$

Define $a = t - kc$. Note that $q_i(a) > 0$ since by (4) $\mu_i(a) > 0$. Now,

$$\begin{aligned} &(a - \mu_i(a-c))q_i(a-c) \\ &\quad \le \quad (a - \mu_i(a))q_i(a) \\ &\quad < \quad (a - c - \nu_i(a))q_i(a) \\ &\quad \le \quad (a - c - \nu_i(a-c))q_i(a-c), \end{aligned}$$

where the strict inequality is due to (4), and the other inequalities are due to (2). The strict inequality implies that $q_i(a-c)$ must be non-zero, and since no value of q_i can be negative, we may cancel that factor from both ends:

$$a - \mu_i(a-c) < a - c - \nu_i(a-c).$$

However, this contradicts (5), and makes (3) false. □

Corollary 1 *Any truthful deterministic-price auction is a condex auction.*

The combination of previous results yields the following characterization of truthful auctions for sponsored search.

Theorem 5 *An auction is truthful if, and only if, for every set of bids the expected price per click of each bidder is the condex price.*

Proof: Let A be an auction. From lemma 1 A is truthful if, and only if, $D(A)$ is truthful. Since $D(A)$ is deterministic-price, from Corollary 1 $D(A)$ is a condex auction and by definition its pricing rule is the expected price per click of A. Since the allocation rule of A and $D(A)$ coincide, their condex prices also coincide. □

6. From one slot to distinct slots

An allocation mechanism reflects an intention on how a prize should be distributed among the bidders. However, when the prize consists of multiple distinct objects, and no bidder can be assigned more than one of them, not every intention may be implemented directly. For instance, if two identical slots are being auctioned, no bidder can be assigned more than 50% of the prize, no matter how much she bids. It is useful then to implement a way to emulate any single-slot distribution philosophy under the multiple distinct slots case. Note that although [10] deals with multiple identical slots, the only examples of allocation rules given therein are for the single-slot case, that is, they do not provide such an implementation.

We present here a family of allocation mechanisms for distinct slots that satisfy the premises of Theorem 1 and thus, together with the condex pricing define incentive compatible auctions for this scenario.

These mechanisms assign the first slot according to the basic single-slot probability, and each of the following slots according to the re-scaled probabilities of the remaining ads. Formally, if the single-slot allocation mechanism is represented by the probability functions $p_1(x), \ldots, p_n(x)$ (which give the probability of bidder i winning the auction when bidding x, assuming all other bids are fixed), then the multislot mechanism is given by the algorithm in Figure 1. The rescaling algorithm works for a large family of single-

set $S = \{1, \ldots, n\}$
for $j = 1$ to k
  randomly pick $i \in S$ with prob. $p_i(b_i)/\sum_{j \in S} p_j(b_j)$
  assign ad i to slot j
  set $S = S \setminus \{i\}$.

Figure 1. The rescaling algorithm

slot allocation mechanisms, wide enough to cover the ones used in practice and others present in the literature. We will call a single-slot auction *consistently-monotone* if, and only if, whenever a bidder raises her bid by any amount while the rest of the bids are fixed, her probability of winning the auction does not decrease, and the probability of any other bidder of winning the auction does not increase. In particular, note that this holds for all deterministic and stochastic auctions described earlier in this work and the cited bibliography.

Theorem 6 *If an allocation mechanism for the single-slot case is consistently-monotone, then its extension to a multislot framework via the rescaling algorithm gives an allocation system for which the expected position-CTR is non-decreasing.*

Corollary 2 *If an allocation mechanism for a single slot is consistently-monotone and it is extended with the rescaling algorithm and the corresponding condex pricing rule, then the resulting auction is incentive compatible.*

In Appendix B we give an example of an application of the rescaling algorithm and Corollary 2.

7. Linear extensions from single-slot to uniform multislot

The results in the previous sections provide a very general way of defining truthful auctions for multiple distinct items given a single-item probabilistic allocation rule. However the pricing rule arising from it is usually cumbersome and difficult to calculate, as it requires the calculation of potentially complex conditional probability functions, and their corresponding derivatives and integrals (following the definition of the condex pricing rule). We explore now other ways of extending a single-slot allocation rule whose pricing mechanism may be handled in a much simpler way.

We will first see how to use a single-slot allocation rule to generate multislot allocation and pricing rules for *identical* slots. These in turn can be used as input for a procedure described in the next section, which allows to jump from a multiple-identical-slots auction to one with multiple distinct slots. These two tools together conform a very simple way of extending a single-slot allocation mechanism to get a truthful auction for multiple distinct slots. As we will see, the goodness of this extension in terms of the revenue attainable by the auctioneer will depend on some properties of the original probability functions.

Let $\overline{p}_i$ be a probability distribution for one slot, that is, $\overline{p}_i(x)$ is the probability that bidder i wins the auctioned slot when bidding x. Recall that there are n bidders and k slots, and that we assume that $k \leq n$. It is clear that $\sum_{i=1}^{n} \overline{p}_i(b_i) = 1$.

Now we need to define a probability function p_i for the identical multislot case such that

$$\sum_{i=1}^{n} p_i(b_i) = k \text{ and } 0 \leq p_i(b_i) \leq 1 \forall i \tag{6}$$

Recall that the pricing rule may be hard to calculate for some underlying probability functions, due to the integrals involved. Thus, we would like to be able to use "nice" probability functions of our own choosing. In Appendix A, we present an algorithm that selects the ads to be assigned to the slots according to any set of rational probabilities satisfying (6). Consequently, we focus now on designing such probability functions.

We consider linear extensions, that is, functions of the form $p_i(x) = a\overline{p}_i(x) + c$, where a and c are constants to be defined. From (6) follow necessary and sufficient conditions on a and c for the feasibility of p_i. First, it follows that $k = \sum_{i=1}^{n} p_i(b_i) = \sum_{i=1}^{n} (a\overline{p}_i(b_i) + c) = a + nc$, so (*constant-coefficient bound*)

$$c = (k - a)/n. \tag{7}$$

The fact that $p_i(b_i)$ has to be a probability allows us to derive another bound, which we call *linear-coefficient bound*:

$$a \leq \min\left(\frac{n-k}{n\sup(\overline{p}_i) - 1}, \frac{k}{1 - n\inf(\overline{p}_i)}\right). \tag{8}$$

Since $a < 0$ yields a non-increasing probability (as opposed to what is needed, a non-decreasing probability), we assume that $a \geq 0$. Now, the condex pricing of the resulting auction becomes

$$\begin{aligned}\mu_i(b_i) &= \frac{\int_0^{b_i} x\, dp_i(x)}{p_i(b_i)} = \frac{\int_0^{b_i} x\, d(a\overline{p}_i(x) + c)}{a\overline{p}_i(b_i) + c} \\ &= \frac{a\int_0^{b_i} x\, d\overline{p}_i(x)}{a\overline{p}_i(b_i) + c} = \frac{a\overline{\mu}_i(b_i)\overline{p}_i(b_i)}{a\overline{p}_i(b_i) + c},\end{aligned}$$

and the expected revenue E for the auctioneer is

$$\begin{aligned}E &= \sum_{i=0}^{n} c_i p_i(b_i)\mu_i(b_i) \\ &= \sum_{i=0}^{n} c_i(a\overline{p}_i(b_i) + c)\frac{a\overline{\mu}_i(b_i)\overline{p}_i(b_i)}{a\overline{p}_i(b_i) + c} \\ &= a\sum_{i=0}^{n} c_i\overline{p}_i(b_i)\overline{\mu}_i(b_i) = a\overline{E},\end{aligned} \tag{9}$$

where c_i is the ad-CTR of ad i, and $\overline{\mu}_i$ and $\overline{E}$ are the pricing rule and revenue of the original auction, respectively. We note that we have assumed in (9) that the position-CTR for each of the slots in the multislot framework is the same as the one in the single-slot setting. Were this assumption not realistic, we could easily adjust the previous result since, under the separability assumption, the difference is only an appropriate constant.

From (7), (8) and (9) immediately follow necessary and sufficient conditions that let $\overline{p}_i$ be extended to a multislot environment generating greater income, maintaining the truthfulness and with no need of complex calculations. We summarize these results in the following theorem.

Theorem 7 *If a stochastic auction of a single slot uses a condex pricing rule and is extended linearly to a stochastic auction for multiple identical slots, respecting both linear-coefficient and constant-coefficient bounds then the following holds.*

- *The condex pricing rule of the extension is* $1 - \frac{c}{a\overline{p}_i(b_i)+c}$ *times the price in the original single-slot auction, where* a *and* c *are the linear and constant coefficients, respectively.*
- *The expected revenue for the publisher is* a *times the one in the single-slot auction.*

- *The resulting auction is incentive compatible.*

From the previous results also follows an upper bound on the publisher's revenue.

Corollary 3 *If a single-slot auction is extended with a linear extension, then the revenue of the resulting auction is not more than* $\frac{n-k}{n \sup(\overline{p}_i)-1}$ *times the revenue of the original auction.*

As for unbounded single-slot probability functions, we establish a negative result:

Corollary 4 *If an auction for a single slot with no upper bound on the probabilities of one bidder is extended via a linear extension, then the auctioneer's revenue will decrease.*

8. From uniform to weighted multislot auctions

As we mentioned in the previous section, here we present a procedure for extending, in a simple way, a multiple-identical-slot auction to a multiple-distinct-slot auction. The extension can be applied even without knowledge of the decay model, potentially at the cost of suboptimal social welfare, but the question is whether the optimal social welfare is achievable at all since, despite extensive research done in the area of sponsored search, it remains challenging to establish the "right" decay model (i.e., the correct values of the w_j's). Moreover, it may vary from one search term to another, or even change in time for the same keyword.

A uniform multislot auction consists of an allocation mechanism and a pricing rule. The pricing rule may be derived from the probabilistic allocation functions, using the MCW Theorem. Let us denote by A the existent allocation mechanism that selects a subset of k ads to be placed in the available slots, and p_i the probability functions that represent that mechanism.

We develop a simple extension of the multiple-identical-slots mechanism for the multiple-distinct-slots scenario: select a set S of k ads using the allocation mechanism A, and then assign them uniformly at random to the k slots. The condex pricing for the resulting auction is the same as the condex pricing for the original auction, as the following theorem states:

Theorem 8 *If an auction of identical slots has a non-decreasing (resp. strictly increasing) allocation rule and its corresponding condex pricing rule, then the auction of distinct slots resulting from applying the random-permutation algorithm to the identical-slot auction and charging the same price is incentive compatible (resp. strictly incentive compatible).*

This algorithm can be applied to multiple-identical-slots auctions that were derived from single-slot auctions (as shown in the previous section). Summarizing, we have devised a method for creating an auction for multiple distinct slots given one for a single slot.

Acknowledgments

We are grateful to anonymous referees for valuable comments and suggestions.

References

[1] Z. Abrams and A. Ghosh. Auctions with revenue guarantees for sponsored search. In *16th International World Wide Web Conference (WWW2007)*, 2007.

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

[3] C. Borgs, J. Chayes, O. Etesami, N. Immorlica, K. Jain, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *16th International World Wide Web Conference (WWW2007)*, 2007.

[4] J. Bullow and P. Klemperer. Auctions vs. negotiations. *American Economic Review*, 86:180–94, 1996.

[5] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Stanford Graduate School of Business Research Paper No. 1917 Available at SSRN: http://ssrn.com/abstract=861164, 2005.

[6] E. Feuerstein, P. Heiber, J. Martínez-Viademonte, and R. Baeza-Yates. New stochastic algorithms for placing ads in sponsored search. In *Proc. 5th Latin American Web Congress (LA-WEB 2007)*, 2007.

[7] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 2006.

[8] B. Holmstrom. Groves' scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.

[9] P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, 1999.

[10] C. Meek, D. Chickering, and D. Wilson. Stochastic and contingent-payment auctions. In *Workshop on Sponsored Search Auctions - ACM Conference on Electronic Commerce (EC'05)*, 2005.

[11] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[12] S. Pandey and C. Olston. Handling advertisements of unknown quality in search advertising. In *Proc. Twentieth Annual Conference on Neural Information Processing Systems (NIPS), Vancouver, Canada*, 2006.

[13] A. Penemberg. Click fraud threatens web. *Wired news*, October 13, 2004.

[14] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance, Vol. 16, No. 1*, pages 8–37, 1961.

A. Efficient drawing of sets with rational probabilities

Here we present a drawing algorithm that will distribute k slots among n bidders without repetition and following rational probabilities p_i such that $\sum_{i=1}^{n} p_i = k$. Let $\frac{n_i}{d_i} = p_i$ where n_i and d_i are coprime. Let $M = MCM(d_1, \ldots, d_n)$ and $a_i = \frac{n_i}{d_i} M$, which is clearly an integer. Note that $\sum_{i=1}^{n} a_i = Mk$.

Let A be a matrix of M rows and k columns. Fill A with a_1 1s, a_2 2s, etc, starting from the upper left cell and going downwards. When a column is full, go to the uppermost cell of the column to the left. Now, choose a random row of the matrix. The numbers in it are the indexes of the chosen ads. Since each a_i is no greater than M, no number appears twice in the same row, so this is a valid choice. Also, ad i appears in a_i out of M rows, so its probability of being chosen is $\frac{a_i}{M} = \frac{n_i}{d_i}$.

This idea can be implemented efficiently by the algorithm given in Figure 2, which is a symbolic calculation of the explained idea.

set c = a random integer between 0 and $M - 1$
set $R = \emptyset$, $d = 0$
for $i = 1$ to k
  set $d' = d + a_i$
  if $d \leq c < d'$ or $d \leq c + M < d'$ set $R = R \cup \{i\}$
  set $d = d' \mod M$
return R

Figure 2. A drawing algorithm

B. An example

The following is an example of an application of the rescaling algorithm and Corollary 2. Let the initial probabilities for a single slot be proportional to the bids, that is, $p_i(b_i) = \frac{b_i}{\sum_h b_h}$. We extend those probabilities to a scenario with $k = 2$ slots, weights $w_1 = 1$ and $w_2 = 1/2$, and $n = 3$ bidders. We only develop the formula corresponding to bidder $i = 1$, the formulas for the other bidders are symmetric. The original condex pricing for the single-slot case is given in [10] by

$$\mu_1(b_1) = (b_2 + b_3)\left[\left(1 + \frac{b_2 + b_3}{b_1}\right)\ln\left(1 + \frac{b_1}{b_2 + b_3}\right) - 1\right].$$

When we apply the rescaling algorithm we get the probability function

$$q_1(b_1) = \frac{b_1}{b_1 + b_2 + b_3}\left(1 + \frac{b_2}{2(b_1 + b_3)} + \frac{b_3}{2(b_1 + b_2)}\right).$$

We are then able to compute the corresponding price using the condex rule, obtaining

$$\mu_1(b_1) = -\frac{\frac{{b_2}^2}{b_2 + b_1} + \frac{{b_3}^2}{b_3 + b_1} + b_2 \ln(b_2 + b_1) + b_3 \ln(b_3 + b_1)}{2q_1(b_1)}$$

Table 1 presents the resulting probabilities, prices and expected revenue for the auctioneer, for different combinations of bids in both scenarios. We observe that the expected revenue with two slots is higher in the cases in which the bids are similar, decreasing as they set apart.

C. Proofs

C.1. Proof of Theorem 1

Proof: Let $\omega = \sum_{j=1}^{k} w_j$ and $\tilde{q}_i(x) = q_i(x)/\omega$. Clearly, $\tilde{q}_i$ is non-decreasing (resp. strictly increasing) if $q_i(x)$ is non-decreasing (resp. strictly increasing) as well. Moreover, it is easy to see that:

$$\sum_{i=1}^{n} \tilde{q}_i(b_i) = \frac{1}{\omega}\sum_{i=1}^{n}\sum_{j=1}^{k} w_j p_i^j(b_i) = \frac{1}{\omega}\sum_{j=1}^{k} w_j \sum_{i=1}^{n} p_i^j(b_i) = 1.$$

Thus, since $\tilde{q}_i$ is a probability function that fulfills the premises of the MCW Theorem, a condex price $\tilde{\mu}_i(x)$ can be defined, implying that

$$(v_i - \tilde{\mu}_i(x))\tilde{q}_i(x) \tag{10}$$

is maximized when $x = v_i$.

Simple calculations show that $\mu_i = \tilde{\mu}_i$. Recall that the revenue of the advertiser is computed as the per-click utility times the expected CTR, while the expected CTR is the product of the ad-CTR and the position-CTR. Thus, by letting c_i be the ad-CTR of ad i, this revenue is

$$\begin{aligned}(v_i - \mu_i(x))\sum_{j=1}^{k} c_i w_j p_i^j(x) &= (v_i - \mu_i(x))q_i(x)c_i \\ &= (v_i - \tilde{\mu}_i(x))\tilde{q}_i(x)\omega c_i, \quad (11)\end{aligned}$$

and since (11) is a constant multiplied by (10), it is also maximized when $x = v_i$. □

C.2. Proof of Theorem 2

Proof: Let us assume that a given auction has q_i as its expected position-CTR function and is in fact incentive compatible. By definition of incentive compatibility, letting c_i and v_i be the ad-CTR and private value of ad i, respectively, for any valid bid x:

$$\begin{aligned} q_i(x)c_i(v_i - \mu_i(x)) &\leq q_i(v_i)c_i(v_i - \mu_i(v_i)) \\ q_i(x)(v_i - \mu_i(x)) &\leq q_i(v_i)(v_i - \mu_i(v_i)) \end{aligned}$$

bid			exp. CTR 2 slots			prob. 1 slot			price 2 slots			price 1 slot			rev. 2 slots	rev. 1 slot
b_1	b_2	b_3	q_1	q_2	q_3	p_1	p_2	p_3	μ_1	μ_2	μ_3	μ_1	μ_2	μ_3		
1.00	1.00	1.00	0.50	0.50	0.50	0.33	0.33	0.33	0.39	0.39	0.39	0.43	0.43	0.43	0.58	0.43
1.00	1.00	2.00	0.42	0.42	0.67	0.25	0.25	0.50	0.40	0.40	0.65	0.45	0.45	0.77	0.77	0.61
1.00	2.00	2.00	0.33	0.58	0.58	0.20	0.40	0.40	0.43	0.70	0.70	0.46	0.83	0.83	0.96	0.76
1.00	10.00	10.00	0.09	0.70	0.70	0.05	0.48	0.48	0.48	2.43	2.43	0.49	3.94	3.94	3.46	3.77
0.10	0.10	1.00	0.30	0.30	0.91	0.08	0.08	0.83	0.04	0.04	0.16	0.05	0.05	0.23	0.17	0.20
0.10	0.10	10.00	0.25	0.25	0.99	0.01	0.01	0.98	0.04	0.04	0.37	0.05	0.05	0.60	0.38	0.59

Table 1. Example for different bid combinations for $n = 3$ bidders, $k = 1$ vs. 2 slots with $w_1 = 1$ and $w_2 = 1/2$

So, if v_i is any given value v,

$$q_i(v+\epsilon)(v - \mu_i(v+\epsilon)) \leq q_i(v)(v - \mu_i(v)), \tag{12}$$

but also, if v_i is $v + \epsilon$, then

$$q_i(v)(v+\epsilon - \mu_i(v)) \leq q_i(v+\epsilon)(v+\epsilon - \mu_i(v+\epsilon)). \tag{13}$$

Finally, adding up (12) and (13), and assuming $\epsilon \geq 0$ yields $q_i(v) \leq q_i(v+\epsilon)$.

The proof in the strictly incentive compatible case is analogous. □

C.3. Proof of Theorem 6

Proof: Let us consider a fixed bidder i for the entire proof. Given a subset C of $\{1,\ldots,n\}$ and a bid x of bidder i, let us denote by $R(C,x)$ the probability that the elements of C are assigned to the first $|C|$ slots in any order by the rescaling algorithm (while all other bids are fixed). By definition, for each non-negative integer m,

$$\sum_{C:|C|=m} R(C,x) = 1 \tag{14}$$

(the sum of the probabilities of all subsets of a given size is 1, because exactly one subset is always chosen and they are all disjoint choices).

First, we show that for any subset C of $\{1,\ldots,n\} \setminus \{i\}$, when C is fixed, $R(C,x)$ is non-increasing. We prove it by induction on $|C|$.

For the base case with $|C| = 0$, there is only one set $C = \emptyset$ and by definition $R(\emptyset, x) = 1$ is constant and therefore non-decreasing. In the inductive step, assuming $R(C,x)$ is non-increasing for sets of size $m-1$, we will show the same holds for sets of size m. Let $r_h(x)$ be the basic single-slot probability of bidder h when bidder i bids x (i.e., the probability that bidder h wins the first slot). By hypothesis r_h is non-increasing if $h \neq i$. Also, by definition of R:

$$R(C,x) = \sum_{h\in C} R(C\setminus\{h\},x)\frac{r_h(x)}{1-\sum_{c\in C\setminus\{h\}} r_c(x)} \tag{15}$$

where the fraction represents the probability of h being selected in the m^{th} iteration, given that the bidders in $C \setminus \{h\}$ were selected in the first $m-1$ iterations.

$R(C\setminus\{h\},x)$ is non-increasing by inductive hypothesis; $r_h(x)$ is non-increasing because $h \in C$ and by hypothesis that means $h \neq i$; and $1/(1-\sum_{c\in C\setminus\{h\}} r_c(x))$ is non-increasing because as each r_c is non-increasing, $\sum_{c\in C\setminus\{h\}} r_c(x)$ is non-increasing as well. Therefore, the sum in (15) is non-increasing.

Now, let $S_j(x)$ be the probability that bidder i gets any of the first j slots when bidding x. By (14) it follows that

$$S_j(x) = \sum_{C:|C|=j\wedge i\in C} R(C,x) = 1 - \sum_{C:|C|=j\wedge i\notin C} R(C,x)$$

Since all terms in the sum are non-increasing, $S_j(x)$ is non-decreasing.

Also, by definition $S_j(x) = S_{j-1}(x) + p_i^j(x)$, so

$$\begin{aligned} q_i(x) &= \sum_{j=1}^{k} w_j p_i^j(x) = \sum_{j=1}^{k} w_j(S_j(x) - S_{j-1}(x)) \\ &= \sum_{j=1}^{k} w_j S_j(x) - \sum_{j=1}^{k} w_j S_{j-1}(x) \\ &= \sum_{j=1}^{k-1} S_j(x)(w_j - w_{j+1}) + w_k S_k(x) - w_1 S_0(x) \end{aligned}$$

and all terms in the last expression are non-decreasing (since $S_0(x) = 0$, the last one is constant), so it is proved that $q_i(x)$ is non-decreasing. □

C.4. Proof of theorem 8

Proof: It is easy to see that with this allocation mechanism $p_i^j(x) = p_i(x)/k$ for all j. The premises of Theorem 1 hold because

$$q_i(x) = \sum_{j=1}^{k} w_j p_i^j(x) = \sum_{j=1}^{k} w_j \frac{p_i(x)}{k} = p_i(x)\frac{\sum_{j=1}^{k} w_j}{k}$$

is clearly non-decreasing (resp. strictly increasing).

If we apply the pricing of Theorem 1 we obtain an incentive compatible (resp. strictly incentive compatible) auction. Moreover, since $\frac{\sum_{j=1}^{k} w_j}{k}$ is constant, the resulting pricing rule equals the one obtained by applying the condex pricing rule to the function p_i. □