

Optimal Auctions Capturing Constraints in Sponsored Search*

Esteban Feuerstein¹, Pablo Ariel Heiber¹, Matías Lopez-Rosenfeld¹,
and Marcelo Mydlarz²

¹ Departamento de Computación, FCEyN, Universidad de Buenos Aires, Argentina

{efeuerst,pheiber,mlopez}@dc.uba.ar

² Yahoo! Research, Santiago, Chile

marcem@yahoo.com

Abstract. Most sponsored search auctions use the *Generalized Second Price* (GSP) rule. Given the GSP rule, they try to give an *optimal* allocation, an easy task when the only need is to allocate ads to slots. However, when other practical conditions must be fulfilled –such as budget constraints, exploration of the performance of new ads, etc.– optimal allocations are hard to obtain. We provide a method to optimally allocate ads to slots under the practical conditions mentioned above. Our auctions are stochastic, and can be applied in tandem with different pricing rules, among which we highlight two: an intuitive generalization of GSP and VCG payments.

1 Introduction

In the framework of sponsored search, advertisers compete in an auction to place their ads on a web page. Each advertiser places a bid, and the search engine decides, based on the bids and other public or private parameters, which ads will be published and where. In the widely used pay-per-click model, each advertiser is charged only when her ad receives a click. The position in which an ad is displayed has an impact on its likelihood of being clicked, and advertisers' preferences follow accordingly. Nevertheless, it is generally assumed that all clicks (independently of the ads' position) have the same value for an advertiser; consequently, search engines establish a price for a click that is not conditioned on the position where the ad is presented.

The most widely used mechanism for sponsored search auctions combines the sort-by-revenue allocation and Generalized Second Price (GSP) rules: bidders are ranked according to the revenue the auctioneer expects to obtain from them, while the price associated to each of the winning bidders –which are precisely the top-ranked bidders– is the minimum amount each of them needs to bid in order to maintain their position in the ranking [3,10].

* This research was funded by a Yahoo! Research Alliance Grant, and in part by UBACYT project X436 “Algoritmos de selección y asignación y mecanismos de pago para la publicidad online en Internet”.

Although in a basic setting the aforementioned tandem sort-by-revenue/GSP yields good revenue for the auctioneer, actual implementations must address many extra features, for instance: (a) Click-through-rates and expected revenues are not really known a priori, so they must be learned somehow by the auctioneer. In order to avoid leaving out ads with high potential revenue, there is a need to alternate among ads with high, small and unknown revenue expectation. This is known as the explore/exploit trade-off [25]. In terms of the allocation mechanism, the implementation of that trade-off may be seen as adding extra constraints to the problem, for example stating that each ad must receive at least a certain fraction of the impressions. (b) Users may aim at different meanings of a same query. Hence, the overall clickability of the published ads, and likewise the auctioneer's revenue may increase when the set of published ads covers a wide range of meanings, independently of the revenue expectation of each ad considered separately [14,30]. (c) The publication of certain ads may have a (possibly negative) influence on the click-probability of other ads. Therefore, the set of ads with the highest aggregate click probability is not necessarily the set of the best individual ads. These are called "contextual effects" [14,30]. (d) In the process of optimizing the performance of an auction, we must decide how many ads to display. This number influences the revenue in several conflicting ways. On the one hand, the more ads published the higher the probability that a user finds one that suits her needs. On the other hand, as the number of ads increases the fraction of the user's attention that each ad attracts decreases; moreover, there is also evidence that the user experience suffers [4]. (e) Bidders usually set budget constraints, i.e., upper bounds on the amount they are willing to spend for a keyword or set of keywords over a time period. These kind of constraints have been studied in [21,1,2]. (f) Advertisers may be allowed to place special requests such as being displayed only in certain positions.

Instead of performing ad-hoc modifications to allocations and pricing rules to model each of these extra features, it is useful to have an auction mechanism general enough to easily adapt to a changing environment. We provide such a mechanism, which simplifies and improves the usability of sponsored search auctions, by means of stochastic auctions. Stochastic auctions are auctions in which the allocation or the pricing rule (or both) are random variables. These auctions may be preferable to deterministic ones for several reasons: (1) they are less prone to vindictive and/or strategic bidding, since strategic behavior is impaired by the non-deterministic nature of the output [20]; (2) the fact that *anyone* can eventually win the auction contributes to have a wider advertisers base and therefore higher revenue in the medium term [16]; (3) they bear higher diversity of ads, which improves user experience and increases aggregate click through rates [14,30]; (4) they provide an implicit mechanism to implement an explore/exploit trade-off [12]; (5) they are in general less vulnerable to fraudulent behavior [26].

The method we propose, based on mathematical programming, creates a stochastic auction that achieves the best allocation with respect to some objective, that satisfies the constraints of the problem. Concretely, we provide an

algorithm \mathcal{M} , which can be subdivided into two parts: (1) a template algorithm \mathcal{A} that, given as input typical parameters (the advertisers' bids, estimations of the ad- and/or position-CTRs), possibly a set of constraints (e.g., budget or variety restrictions), and an objective function O , produces an *equivalence class of stochastic allocations* that satisfy the constraints, and are best possible according to O , and (2) a drawing algorithm that allocates ads according to the probabilities of some stochastic allocation in the class obtained through \mathcal{A} .

The method just described can be combined with several pricing rules, as we examine in Section 4. This combination yields, for a class of pricing rules which we will call *a priori*, optimal auctions –according to different objectives. A priori pricing rules include, among others, First-price and GSP. Other pricing rules we can use include VCG payments and Myerson's optimal truthful mechanism, which requires an assumption about the bidders' valuations (see Section 4).

We also introduce a natural extension of GSP for stochastic allocations, the *Extended Generalized Second Price* (EGSP) rule: prices are not only associated to the top-ranked bidders, but to *all* the bidders with a positive probability of being allocated a slot. Prices are computed in the same way as in GSP, therefore the prices associated to the top-ranked bidders coincide under both pricing rules.

EGSP is an a-priori pricing and can be coupled with different stochastic allocation rules, in particular those obtained using algorithm \mathcal{M} . This combination becomes a way of extending mechanisms currently in use towards a framework where a rich set of constraints can be explicitly included.

The computational complexity of algorithm \mathcal{A} depends on the objective function and the type of restrictions. A key observation is that, thanks to the stochastic nature of the allocations, there is no need to impose integrality restrictions. In particular, there are many interesting objective functions and restrictions that are linear (see Section 3), and yield polynomial time algorithms¹.

As opposed to auctions currently in use, which (to the best of our knowledge) can only handle constraints and objective functions in an ad-hoc way, \mathcal{M} can handle many of those constraints and objectives seamlessly, providing an optimal allocation for many pricing rules, including EGSP. Indeed, \mathcal{M} combined with EGSP brings the same or better allocations than sort-by-revenue/GSP auctions. Another advantage of \mathcal{M} is related to the nature of the stochastic allocations involved: we are able to optimize over a (continuous) polytope, as opposed to a discrete lattice, where optimization is computationally inefficient.

In summary, we present an extension to the most popular pricing rule in sponsored search, and a method to derive best stochastic allocation rules based on mathematical programming under different pricings. Neither of these contributions is a break-through result, yet their combination provides a powerful way to obtain optimal auctions in some real-life settings of sponsored search. Another

¹ Even when the types of restrictions entail a non-polynomial running time algorithm, if the size of the problem is reasonably small, \mathcal{M} may still be used in practice; e.g., by combining our approach with techniques for subdividing query-bidder graphs into smaller instances [8].

application is their use as benchmarks to measure the impact of the introduction of constraints on the overall performance of an auction.

Related work. The subject of including budgets in the design of sponsored auctions has received a lot of attention recently, for example in [7] and [1]. Mehta, Saberi, Vazirani and Vazirani [21] explore the problem from a competitive analysis point of view: they aim at optimizing the total revenue for a set of queries in an on-line manner, by trying to consume the maximum amount of each bidder's budget through a sequence of queries, of which neither the total length nor the frequency of each are known in advance, obtaining an optimal $(1 - 1/e)$ -competitive algorithm. Mahdian, Nazerzadeh and Saberi [19] consider the same framework, and present an algorithm that takes advantage of good estimations on the frequencies of keywords, while maintaining a good worst-case competitive ratio in case that those estimates are incorrect.

Linear Programming and Stochastic Algorithms have been used before in the framework of mechanism design. Just to cite a recent example in the framework of truthful mechanism design for combinatorial problems, Lavi and Swamy [17] propose a way to convert LP-based approximation algorithms into stochastic mechanisms that give approximate solutions to the winner determination problem and are truthful in expectation (i.e., all players maximize their expected utility by revealing their true values). One of the consequences of that work is a truthful approximation algorithm for both multi-unit auctions and multi-unit combinatorial auctions, which are problems related to ours. The focus there is, however, on a different aspect of the problem, more related to computational complexity of one-time auctions.

Much related to our approach, Abrams, Mendeleevitch and Tomlin [2] use LP trying to optimize sponsored search auctions subject to budget constraints. The difference with our work is manifold. Firstly, they impose restrictions on allocations: no ad can appear in a worse position than another ad with lower ranking. Secondly, while our formulation supports diverse pricing rules, theirs gets restricted to one: the price associated to each bidder is the minimum price needed to beat the ad allocated to the next slot. Thirdly, their model is more involved than ours, and requires more elaborated LP techniques such as delayed column generation; our model is then computationally far more efficient. Finally, it lacks the extra flexibility (given by stochastic allocations) for *easily* including other kinds of restrictions.

LP has been used in the framework of on-line advertising under the more traditional pay-per-impression model, for example in [23,29]. Finally, the use of stochastic auctions for sponsored search has been recently considered in [20], [12], [13] and [6].

2 The Model

Assumptions and notation. The setting we consider involves n risk-neutral bidders that compete for slots, but no bidder can win more than one; the number of slots is not set in advance. Each bidder i has a private value v_i for each click

received, and for which she places a bid of b_i . Following each query, the auctioneer decides which ads will be published along with their order.

A *bid vector* is a vector $\mathbf{b} = (b_1, \dots, b_n) \in \mathbb{R}^n$, while \mathcal{B} denotes the set of all bid vectors. An *allocation* s is an ordered subset of the ads to be displayed on a particular occurrence of a query; its size is denoted by $|s|$ (the number of ads to be displayed).

Unless stated otherwise, we assume that the click probability of an ad is not influenced by the identity of the other published ads. We denote by $CTR_{i,j,k}$ the expected click probability of ad i when presented in the j -th position while a total of k ads are displayed. The expected click-through rate of an ad i in an allocation s , denoted by $CTR_{i,s}$, is then $CTR_{i,j,|s|}$, where j is the position of i in s . A *stochastic allocation* S is a random variable with some probability distribution over allocations. The expected click through rate of ad i in S , denoted by $CTR_{i,S}$, is the sum over all possible allocations s of $CTR_{i,s}P(S = s)$ (where $P(S = s)$ denotes the probability that allocation s is chosen).

Let \mathcal{S} be the set of all stochastic allocations. An *allocation rule* is a function $\sigma : \mathcal{B} \rightarrow \mathcal{S}$. A *pricing* $\mu = \langle \mu_1, \dots, \mu_n \rangle \in \mathbb{R}^n$ is a vector of prices per click for each bidder. A *stochastic pricing* is a probability distribution over pricings. Let \mathcal{M} be the set of all stochastic pricings. A *pricing rule* is a function $p : \mathcal{B} \rightarrow \mathcal{M}$. An *auction* is a pair (a, p) , where a is an allocation rule and p is a pricing rule.

Basic Model. Now we present a basic model that uses mathematical programming in the framework of sponsored search auctions. First we define an equivalence relation over stochastic allocations, along with a polyhedron whose feasible region is the set of those equivalence classes. Then we provide an efficient algorithm that, given a point x in the polyhedron, obtains a stochastic allocation S from the equivalence class of x , followed by an allocation according to the probability distribution of S . Finally, we show that many typical auction measures (such as social welfare or the auctioneer's expected revenue) may be described as linear functions over the polyhedron, being therefore possible to efficiently compute the optimal equivalence class of allocations for that measure using linear programming.

We say that two stochastic allocations S and S' are equivalent if, and only if, for each ad i , each position j and each k , the probability that ad i is assigned to slot j when displaying k ads is the same under S as under S' . That two equivalent stochastic allocations are not necessarily equal is shown in the following simple example. Consider two stochastic allocations S_1 and S_2 ; S_1 allocates three bidders in order $(1, 2, 3)$, $(2, 3, 1)$, and $(3, 1, 2)$ each allocation with probability $1/3$, and S_2 allocates them in order $(1, 3, 2)$, $(3, 2, 1)$, and $(2, 1, 3)$ each with probability $1/3$ as well. While S_1 and S_2 are different, both have the same probability of allocation for each combination of advertiser and position (there are always three ads displayed), and thus $CTR_{i,S_1} = CTR_{i,S_2}$ for each ad i .

If two stochastic allocations are equivalent, the expected CTR of each ad coincides in both stochastic allocations and consequently, (noting that the price per click charged to any bidder is independent of the slot assigned to her) both the expected revenue for each bidder and for the auctioneer are the same as well in both auctions (under equal pricings).

The Stochastic Allocations Polyhedron. By regarding y_k as the probability of having k ads displayed, and $x_{i,j,k}$ as the probability of ad i being displayed on position j when a total of k ads are displayed ($1 \leq i \leq n$, $1 \leq j \leq k \leq n$), we define the Stochastic Allocations Polyhedron (SAP) by

$$x_{i,j,k} \geq 0 \quad \text{for each } i, j, k \quad (1) \quad \sum_{k=1}^n y_k = 1 \quad (2)$$

$$\sum_{i=1}^n x_{i,j,k} \leq y_k \quad \text{for each } j, k \quad (3) \quad \sum_{j=1}^k x_{i,j,k} \leq y_k \quad \text{for each } i, k. \quad (4)$$

For convenience we also define an extension of SAP (SAP-e), by adding non-negative (slack) variables $x_{i,j,k}$ for $j > k$, and replacing inequalities (3) and (4) with

$$\sum_{i=1}^n x_{i,j,k} = y_k \quad \text{for each } j, k \quad (3a) \quad \sum_{j=1}^n x_{i,j,k} = y_k \quad \text{for each } i, k. \quad (4a)$$

It is easy to see that the feasible region of SAP-e may be partitioned into subsets such that each of these subsets is associated with one feasible solution of SAP (by dropping the slack variables). While each feasible solution of SAP represents a class of stochastic allocations, SAP-e will prove useful for technical purposes.

Note that if we replace y_k by 1, and remove the third coordinate of the variables $x_{i,j,k}$, then (1), (3a) and (4a) describe the *bipartite perfect matching polytope* [11,27,18]. As in the bipartite matching polytope, we state in the next lemma that SAP extremes are also integral.

Lemma 1. *Every extreme of SAP (SAP-e) is integral.*

A drawing algorithm. Each solution of SAP (or SAP-e) can be associated with a set of *equivalent* stochastic allocations. Given such solution $(x_{i,j,k}, y_k)_{1 \leq i,j,k \leq n}$ in SAP-e, we show next how to obtain a stochastic allocation in its equivalence class.

For each k such that $y_k > 0$ we define the $n \times n$ matrix $Z^{(k)} = (z_{i,j}^{(k)})_{1 \leq i,j \leq n}$ by $z_{i,j}^{(k)} = x_{i,j,k}/y_k$. From this definition and restrictions (3a) and (4a) follows that each row and column of $Z^{(k)}$ sums up to 1, that is, $Z^{(k)}$ is a doubly stochastic matrix. In consonance with the Birkhoff-von Neumann theorem [5], $Z^{(k)}$ is a convex combination of permutation matrices. Accordingly, we give the following probabilistic algorithm that produces an allocation given a point in SAP:

Choose k with probability y_k .
 Construct $Z^{(k)}$.
 Find permutation matrices P_l and positive numbers λ_l
 such that $\sum_l \lambda_l = 1$ and $Z^{(k)} = \sum_l \lambda_l P_l$.
 Choose a permutation matrix P_l with probability λ_l .
 For $j = 1$ to k
 Let i be such that $P[i, j] = 1$.
 Display ad i in position j .

We can see that with the preceding algorithm the probability of displaying exactly k ads is y_k , and the probability that ad i is displayed in position j while having k ads on display is $y_k z_{i,j}^{(k)} = x_{i,j,k}$. Therefore, the stochastic allocation that results from the application of the algorithm to a point in SAP belongs to the equivalence class of the point (the stochastic allocation selected from the equivalence class depends on the convex combination found, which is not necessarily unique).

Now our procedure is clear: apply an instance of the template algorithm \mathcal{A} in order to obtain a solution of SAP, and feed that solution to the drawing algorithm in order to obtain an allocation. Since the convex combination for any given matrix $Z^{(k)}$ can be attained in polynomial time, the drawing algorithm takes polynomial time as well. As long as the instance of \mathcal{A} also runs in polynomial time, so will our procedure.

Lemma 1 implies that we are modeling an assignment problem, which can be solved with faster methods than using linear programming plus the drawing algorithm. Nevertheless, as it will become clear in the next section, this model provides an extra flexibility that enables the inclusion of different extensions.

3 Optimizing over SAP and Extensions

In order to round up the description of our model, we note that we can optimize any function over SAP. In particular, we consider linear functions, that yield linear programs. A natural instance of such functions is the social welfare, which can be maximized if we have the bidders' private values (or an estimation): $\sum_{i=1}^n \sum_{k=1}^n \sum_{j=1}^k x_{i,j,k} CTR_{i,j,k} v_i$. Alternatively, we can maximize the expected revenue of the auctioneer, $\sum_{i=1}^n \sum_{k=1}^n \sum_{j=1}^k x_{i,j,k} CTR_{i,j,k} \mu_i$. We can also maximize any linear combination of measures, therefore being able to tweak the trade-off between different objectives. We give an example of this kind of objective functions at the end of the next subsection.

We note that the model depends on some parameters of the environment, mainly the pricing and the click probabilities. As for the pricing, the vector $\mu = \langle \mu_1, \dots, \mu_n \rangle$ that we use may be a function of the bids and eventually other variables (like the click probabilities themselves), but we restrict μ to not depend on the allocation rule². In other words, when the objective function depends on

² We could also use $\mu_{i,j,k}$, allowing the price to depend on the number of ads displayed and their positions.

the price (e.g., the revenue function), the model is general enough to represent any stochastic sponsored search auction in which the prices are calculated “a priori” of the assignment. This includes classical pricing rules such as first-price or appropriate variants of the second-price rule. In Section 4 we present a natural extension of the broadly used Generalized Second Price (GSP) rule [10], which is suitable for our auctions. Note that we make no equilibrium analysis in these cases.

Since the extremes of SAP are binary solutions, maximizing any continuous objective function may be seen as an assignment problem. It is only natural then that we can obtain the same solutions with less sophisticated methods than linear programming. Nevertheless, the power of the model presented in Section 2 lies on its flexibility: it may be combined with different objective functions, restricted by adding different types of constraints, and extended by adding new variables that represent other features of the auction. Next we show a few practical instances of these extensions.

We restrict our analysis to linear programming, therefore both the objective functions and constraints considered are linear. We also note that although we consider each extension separately, they can be combined according to the features being modeled.

3.1 Variety Constraints

The first extension to the basic model that we consider are *variety constraints*. They are introduced as a means of granting each bidder some minimum reward in terms of impressions or click probability.

We consider two types of variety constraints: 1. Each ad is granted some probability of appearance on each occurrence of the query, 2. Each ad is granted a minimum expected click probability on each occurrence of the query. Both of them can be modeled by linear constraints to restrict the solution set of SAP. The first one may be appealing to some advertisers who are interested in impressions rather than clicks. However, the usual model for sponsored search considers that bidders only get profit on clicks, so we focus on type 2 constraints. In order to define type 2 constraints we make use of the *separability assumption* [3]: the CTR may be separated into two factors, one advertisement-specific, the *ad-CTR*, and the other position-specific, the *position-CTR*. Formally, denoting by a_i the ad-CTR of ad i , and by $w_{j,k}$ the position-CTR of slot j when k ads are displayed, the separability assumption states that $CTR_{i,j,k} = a_i w_{j,k}$. For convenience, we assume without loss of generality that the weights are sorted in such a way that $CTR_{i,j,k} \geq CTR_{i,j+1,k}$ ($w_{j,k} \geq w_{j+1,k}$). We also assume that $CTR_{i,j,k} \geq CTR_{i,j,k+1}$ ($w_{j,k} \geq w_{j,k+1}$), since having additional ads displayed can only reduce the visibility of the others. If the maximum number of ads to be displayed is m , then we set $w_{j,k} = 0$ for all $k > m$.

By denoting with l_i^a and l_i^p the lower bounds on bidder i 's expected impression probability and position-CTR, respectively, constraints of type 1 and 2 can be respectively expressed by

$$\sum_{k=1}^n \sum_{j=1}^k x_{i,j,k} \geq l_i^a \quad \text{for each } i \quad \text{and} \quad \sum_{k=1}^n \sum_{j=1}^k w_{j,k} x_{i,j,k} \geq l_i^p \quad \text{for each } i.$$

Since the lower bounds are part of the input to \mathcal{M} , different alternatives are possible. They are beyond the scope of this article. Nevertheless, we observe that for constraints of type 1, no feasible solution exists unless $\sum_{i=1}^n l_i^a < k$, for the largest k satisfying $w_{k,k} > 0$. It is less evident that for constraints of type 2 and assuming $l_i^p > l_{i+1}^p$, no feasible solution exists unless there is k such that for each $1 \leq t \leq n$, $\sum_{i=1}^t l_i^p \leq \sum_{i=1}^t w_{i,k}$; we omit the proof due to space limitations.

Alternatively, the need of periodically publishing every ad can be expressed via another objective function such as $\sum_{i=1}^n \sum_{k=1}^n \sum_{j=1}^k x_{ijk} (w_{jk} a_i \mu_i + E_i)$, where E_i is a measure of the benefit of exploring ad i , provided by an external source devoted to manage the explore/exploit trade-off [25].

3.2 Budget Constraints

We introduce now *budget constraints*, describing how much bidders are willing to spend. In order to model these constraints, we slightly modify our polyhedron, increasing its dimension. Nevertheless, all the results presented thus far can be easily extended to support this change.

Instead of working with a single query, we consider now a set of queries Q ; each bidders may place bids on many (possibly different) queries; Thus, we need to add to the variables and constants presented in Section 2 a new subindex q ranging over Q . For instance, we will have variables $y_{k,q}$ and $x_{i,j,k,q}$ whenever bidder i bids on q ; a priori prices $\mu_{i,q}$ may be part of the input. The basic restrictions of Section 2 become

$$\begin{aligned} x_{i,j,k,q} &\geq 0 && \text{for each } q, i, j, k && \sum_{k=1}^n y_{k,q} &= 1 && \text{for each } q \\ \sum_{i=1}^n x_{i,j,k,q} &\leq y_{k,q} && \text{for each } q, j, k && \sum_{j=1}^n x_{i,j,k,q} &\leq y_{k,q} && \text{for each } q, i, k. \end{aligned}$$

Each query $q \in Q$ is expected to occur c_q times during a certain time window; each bidder i may set a maximum budget $B_i^{(q)}$ for q and/or a maximum overall budget B_i . Note that some of the budgets may be set to infinity by dropping the associated restriction. Bidders that do not participate in a given query can be modeled with a 0 price, as they should never be displayed for that query.

Naturally, the objective function must be modified accordingly; e.g., the revenue maximization goal would be $\sum_q \sum_{i=1}^n \sum_{k=1}^n \sum_{j=1}^k c_q x_{i,j,k,q} CT R_{i,j,k,q} \mu_{i,q}$.

We state now the restrictions that *in expectation* preclude bidders from going over their budgets: the expected payments of each bidder should not exceed her budget during a time window.

$$\sum_q \sum_{k=1}^n \sum_{j=1}^k c_q x_{i,j,k,q} CTR_{i,j,k,q} \mu_{i,q} \leq B_i \quad \text{for each } i$$

$$\sum_{k=1}^n \sum_{j=1}^k c_q x_{i,j,k,q} CTR_{i,j,k,q} \mu_{i,q} \leq B_i^{(q)} \quad \text{for each } i, q.$$

Note that in practice it would be possible to display an ad whose budget has been exhausted. In such cases we may choose to replace this ad by any other one, without significantly affecting the expected revenue. A possible way of reducing the incidence of such situations is to use $B_i - \varepsilon$ instead of B_i in the LP constraints. We note that artificially retaining bidders with exhausted budgets may result in some illegitimate price hikes for other bidders, depending on the pricing rule. This effect however can be controlled if needed.

3.3 Contextual Effects and Other Restrictions

We consider now situations where the click probability of an ad is influenced by the other ads displayed. These are called *externalities* or *contextual effects*, and have been considered recently in [14,15]. The latter argues through experimental evidence that contextual effects do exist in sponsored search, and quantifies them.

One possible way of modeling contextual effects in our framework is by grouping together ads that have negative effects on each other's clickability and introducing a new kind of restriction, that aims at avoiding the joint publication of ads in the same group. In order to establish incompatibilities among groups of similar ads, *ads are partitioned into incompatibility groups* (each ad in exactly one group). The restriction states that *at most one ad of each group can be displayed* at the same time. This approach may be useful, for instance, when a query has different meanings – so users that search for that query may have different intentions – and the auctioneer tries to cover all the range without increasing the total number of ads shown.

In this extension, apart from adding restrictions to SAP, we also need to refine the drawing algorithm given in Section 2 that produces the allocation. This can be done with minor adjustments. Let m be the number of incompatibility groups. Since we will assign m groups instead of n bidders to the slots, we set n to be m . We need to ensure that each group receives, for each particular k , an aggregated probability of exactly y_k of being assigned some position. Letting G be the set of groups, this is captured by

$$\sum_{i \in g} \sum_{j=1}^m x_{i,j,k} = y_k \quad \text{for each } g \in G, k.$$

We can maximize different objective functions over this new polyhedron, looking for good “group allocations”. Concrete ad allocations will be produced by a modified version of the drawing algorithm. For each k such that $y_k > 0$, we

construct an $m \times m$ matrix $Z^{(k)}$ that is used to choose group permutations instead of bidder permutations, and then for each selected group g , choose a bidder from g and assign it to the slot. $Z^{(k)}$ is defined by $z_{l,j}^{(k)} = \sum_{i \in G_l} x_{i,j,k} / y_k$ for $1 \leq l, j \leq m$. Note that again $Z^{(k)}$ is a doubly stochastic matrix, so we apply the Birkhoff-von Neumann theorem as in Section 2. The new drawing algorithm will then be:

Choose k with probability y_k .
 Construct $Z^{(k)}$.
 Find permutation matrices P_l and positive numbers λ_l
 such that $\sum_l \lambda_l = 1$ and $Z^{(k)} = \sum_l \lambda_l P_l$.
 Choose a permutation matrix P_l with probability λ_l .
 For $j = 1$ to k
 Let g be such that $P[g, j] = 1$.
 Choose ad i of group g with probability $\frac{x_{i,j,k}}{\sum_{i \in g} x_{i,j,k}}$.
 Display ad i in position j .

It follows immediately that the probability of ad i of group l being placed on position j when k ads are displayed is $y_k z_{l,j}^{(k)} x_{i,j,k} / \sum_{i \in G_l} x_{i,j,k} = x_{i,j,k}$.

Grouping can also be used in other ways, such as setting the minimum number of ads to display from each group. However, this would require a deeper modification in the drawing algorithm.

4 Pricing Rules

Auctions are constituted by two main components: the allocation rule and the pricing rule. So far we have focused on obtaining optimal allocations (according to some criteria) that satisfy a number of restrictions. In this section, we focus on the pricing rules that can be combined with those allocation rules.

We divide our analysis in two directions: first we analyze pricing rules that are best suited for our optimization framework, and then we consider incentive-compatible (truthful) pricing rules. In the first case we do not make a game-theoretic analysis, we assume that bids do not change in response to the allocation rule. With truthful pricings, however, we can assume that the bids are in equilibrium since by definition every bidder maximizes her expected revenue by bidding her true value.

A priori pricing rules and EGSP. When the prices associated to the ads may depend on the ranking, but not on the allocation, these prices may be part of the input of \mathcal{M} , so as to find the best allocation with respect to some objective that depends on them³. We call these pricing rules *a priori*. Several well-known and

³ In the general case, we do allow a *light* dependence on the allocation: we only associate a non-zero price when an item is allocated; nevertheless, in the pay-per-click model this distinction disappears since an ad that is not displayed cannot be clicked and therefore will not be charged.

widespread pricing rules are indeed a priori; most notorious examples of this class are First-Price and the Generalized Second Price (GSP) [3,10]. While the former may be applied in our framework, it has been dropped from sponsored-search settings due to its instability. On the other hand, GSP has become the most widely used rule in that framework, enjoying good properties such as *envy-free* equilibria (see [10] for details).

We introduce a natural extension of GSP for stochastic allocations, the *Extended Generalized Second Price* (EGSP) rule. Like in GSP, EGSP assumes that the auctioneer ranks bidders according to some function on their bids while each winning bidder pays (for a click) the minimum price needed to retain her position in the ranking. However, with EGSP prices are not only associated to the top-ranked bidders, but to *all* the bidders with a positive probability of being allocated a slot. Since prices are computed in the same way as in GSP, the prices associated to the top-ranked bidders coincide under both pricing rules. Note that, as we are dealing with stochastic allocations, the ranking order is not necessarily the order in which ads are displayed each time; nevertheless, the resulting stochastic allocation rule will tend to allocate more/better slots to ads with a higher ranking.

Another extension of GSP to an allocation rule different than the simple sort-by-revenue rule has been proposed in [2]: given allocations that are subsets of the ads ordered by ranking, the price associated to each bidder is the minimum price needed to beat the ad allocated to the following slot (the price associated to the last ad is the reserve price). It is easy to see that, given any set of bids, EGSP charges strictly more than the pricing rule in [2].

Incentive Compatibility. Although the variations over the GSP rule currently in use in sponsored search auctions are not truthful, there are many reasons that make truthfulness a desirable property, which can be summarized in the fact that advertisers can define their optimal bids by themselves, without the need of invoking consultants or gurus, driving more resources to the sponsored search business, for the benefit of advertisers, auctioneer and users.

A natural way to incorporate truthfulness into our framework consists of the classical VCG approach [28,9], that is, the incentive-compatible pricing rule corresponding to the allocation obtained through \mathcal{M} that is individually rational and makes no positive transfers [24]. For instance, when bidder i 's bid is b_i (interpreted as values v_i , since in a truthful auction we can assume that each bidder bids her own private valuation), and the objective function is the social welfare $\sum_{i=1}^n \sum_{k=1}^n \sum_{j=1}^k x_{i,j,k} CTR_{i,j,k} b_i$, we associate to each bidder a price equal to the difference between social welfare value obtained with and without her participation. In this way, in order to compute the price for the n bidders we run the mathematical program $n + 1$ times. In practice the computation can be done "on demand", that is, only when an ad indeed receives a click.

Another approach to truthfulness within our framework is the following: given a distribution F_i on the valuation of each bidder i (known or inferred through historical data), we apply Myerson's incentive-compatible mechanism [22,16],

which maximizes the (expected) revenue of the auctioneer⁴. For an explanation on how to apply this mechanism, see for example [24].

Further Research. We are currently studying “a priori” pricing methods and their consequences, in particular the existence of equilibria for stochastic auctions under EGSP.

Another interesting research subject are auctions in which the expected position-CTR of the i -th ranked bidder is set to a value p_i . Such auctions, though stochastic in nature, behave like deterministic ones, so paired with EGSP will have interesting properties, such as the existence of envy-free equilibria [10].

Acknowledgments. We thank Ofer Mendelevitch and John Tomlin for providing us with data for the simulations.

References

1. Abrams, Z.: Revenue maximization when bidders have budgets. In: SODA, pp. 1074–1082 (2006)
2. Abrams, Z., Mendelevitch, O., Tomlin, J.: Optimal delivery of sponsored search advertisements subject to budget constraints. In: ACM Conference on Electronic Commerce, pp. 272–278 (2007)
3. Aggarwal, G., Goel, A., Motwani, R.: Truthful auctions for pricing search keywords. In: ACM Conference on Electronic Commerce, pp. 1–7 (2006)
4. Bhargava, H.K., Feng, J.: Paid placement strategies for internet search engines. In: WWW 2002: Proceedings of the 11th international conference on World Wide Web (2002)
5. Birkhoff, G.: Tres observaciones sobre el algebra lineal. Univ. Nac. Tucumán. Revista, 147–151 (1946)
6. Borgs, C., Chayes, J., Etesami, O., Immorlica, N., Jain, K., Mahdian, M.: Dynamics of bid optimization in online advertisement auctions. In: 16th International World Wide Web Conference (WWW 2007) (2007)
7. Borgs, C., Chayes, J.T., Immorlica, N., Mahdian, M., Saberi, A.: Multi-unit auctions with budget-constrained bidders. In: ACM Conference on Electronic Commerce, pp. 44–51 (2005)
8. Carrasco, J.J., Fain, D.C., Lang, K.J., Zhukov, L.: Clustering of bipartite advertiser-keyword graph. In: Workshop on Large Scale Clustering at IEEE International Conference on Data Mining (2003)
9. Clarke, E.H.: Multipart pricing of public goods. *Public Choice*, 17–33 (1971)
10. Edelman, B.G., Ostrovsky, M., Schwarz, M.: 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 (2005), <http://ssrn.com/abstract=861164>
11. Edmonds, J.: Paths, trees, and flowers. *Canad. J. Math.* 17, 449–467 (1965)
12. Feuerstein, E., Heiber, P., Martínez-Viademonte, J., Baeza-Yates, R.: New stochastic algorithms for placing ads in sponsored search. In: Proc. 5th Latin American Web Congress (LA-WEB 2007) (2007)

⁴ This may affect the incentives of the advertisers in previous stages of the mechanism.

13. Feuerstein, E., Heiber, P., Mydlarz, M.: Truthful stochastic and deterministic auctions for sponsored search. In: Proc. 6th Latin American Web Congress (LA-WEB 2008) (2008)
14. Ghosh, A., Mahdian, M.: Externalities in online advertising. In: 17th International World Wide Web Conference (WWW 2008) (2008)
15. Gunawardana, A., Meek, C.: Aggregators and contextual effects in search ad markets. In: WWW Workshop on Targeting and Ranking for Online Advertising (2008)
16. Klemperer, P.: Auction theory: A guide to the literature. *Journal of Economic Surveys* 13(3), 227–286 (1999)
17. Lavi, R., Swamy, C.: Truthful and near-optimal mechanism design via linear programming. In: FOCS, pp. 595–604 (2005)
18. Lovász, L., Plummer, M.D.: *Matching Theory*. In: *Annals of Discrete Mathematics*. North-Holland, Amsterdam (1986)
19. Mahdian, M., Nazerzadeh, H., Saberi, A.: Allocating online advertisement space with unreliable estimates. In: ACM Conference on Electronic Commerce, pp. 288–294 (2007)
20. Meek, C., Chickering, D.M., Wilson, D.B.: Stochastic and contingent-payment auctions. In: Workshop on Sponsored Search Auctions - ACM Conference on Electronic Commerce (EC 2005) (2005)
21. Mehta, A., Saberi, A., Vazirani, U., Vazirani, V.: Adwords and generalized on-line matching. In: Proc. Symposium on Foundations of Computer Science (2005)
22. Myerson, R.: Optimal auction design. *Mathematics of Operations Research* 6, 58–73 (1981)
23. Nakamura, A., Abe, N.: Improvements to the linear programming based scheduling of web advertisements: World wide web electronic commerce, security and privacy. In: zurko, M.e., greenwald, a. (guest eds.) *Electronic Commerce Research*, vol. 5, pp. 75–98 (2005)
24. Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V. (eds.): *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
25. Pandey, S., Olston, C.: Handling advertisements of unknown quality in search advertising. In: Proc. Twentieth Annual Conference on Neural Information Processing Systems (NIPS), Vancouver, Canada (2006)
26. Penenberg, A.: Click fraud threatens web. *Wired news*, October 13 (2004)
27. Schrijver, A.: *Theory of Linear and Integer Programming*. John Wiley & Sons, Chichester (1986)
28. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16(1), 8–37 (1961)
29. Wiggins, R., Tomlin, J.: Bid optimization for internet graphical ad auction systems via special ordered sets. Yahoo! Research Technical Report YR-2007-004 (2007)
30. Zhu, X., Goldberg, A., Van Gael, J., Andrzejewski, D.: Improving diversity in ranking using absorbing random walks. In: *Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAAC-HLT)* (2007)