

On the Efficiency and Equilibria of Rich Ads

MohammadAmin Ghiasi¹, MohammadTaghi Hajiaghayi¹, Sébastien Lahaie² and Hadi Yami¹

¹University of Maryland

²Google Research

{amin, hajiaghayi}@cs.umd.edu, slahaie@google.com, hadiyami@cs.umd.edu

Abstract

Search ads have evolved in recent years from simple text formats to rich ads that allow deep site links, ratings, images and videos. In this paper, we consider a model where several slots are available on the search results page, as in the classic generalized second-price auction (GSP), but now a bidder can be allocated several consecutive slots, which are interpreted as a rich ad. As in the GSP, each bidder submits a bid-per-click, but the click-through rate (CTR) function is generalized from a simple CTR for each slot to a general CTR function over sets of consecutive slots. We study allocation and pricing in this model under subadditive and fractionally subadditive CTRs. We design and analyze a constant-factor approximation algorithm for the efficient allocation problem under fractionally subadditive CTRs, and a log-approximation algorithm for the subadditive case. Building on these results, we show that approximate competitive equilibrium prices exist and can be computed for subadditive and fractionally subadditive CTRs, with the same guarantees as for allocation.

1 Introduction

As search results pages have evolved from ‘ten blue links’ to rich informational layouts including images, video, product listings and the like, so have the advertising units on the page. Search ads today go beyond the standard template of title-link-description and can admit various enhancements known as *rich ads* or *ad extensions*. For instance, a location extension shows a business’ address and distance; a seller’s rating extension shows star ratings and review snippets. A common and important variant is the site links extension, which adds additional lines to the ad with more links to the business (e.g., to a hotel restaurant or reservations). Rich ads offer better information to the user and higher click-through rates for advertisers, and are now commonplace on Google, Bing, and Yahoo.

Rich ads pose a market design challenge because the standard auction for text ads, namely the Generalized Second Price auction (GSP), does not directly extend to ads of various sizes. The GSP proceeds by ranking text ads according to

the product of the advertiser’s bid and the ad’s click-through rate. An advertiser is charged, per click, the lowest bid that would maintain its ad’s rank. The GSP is therefore intimately connected with the notion of *ranking*, but with rich ads the problem moves from ranking to general *allocation*. With extensions allowed, it is conceivable that the highest bidder may be placed below the top slot, but with a larger ad including extensions like site links. The GSP pricing rule is now ill-defined: because ad extensions lead to higher click-through rates (CTRs), GSP pricing taken verbatim may counterintuitively lead to a lower price-per-click for extensions, even though they occupy more space on the page. These kinds of complications indicate the need to rethink allocation and pricing together in the presence of rich ads.

In this paper, we cast the problem of allocating and pricing ads and ad extensions on the page as a generalization of the standard slot-allocation problem handled by GSP, where each agent can now obtain multiple consecutive slots. When an agent obtains several slots, the interpretation is that its ad is shown with extensions (i.e., as a rich ad). Since the number of slots on the page is constrained, the allocation problem remains interesting and does not simply call for every ad to be shown with an extension. The bidding language remains the same as GSP: each agent submits a single bid per click (agents do not directly express any preferences for extensions). The crucial generalization to GSP is that the CTR function is now defined over sets of slots, beyond just individual slots. There are two structural extremes for CTR functions: unit-demand and additive. When CTRs are unit-demand, we recover the standard GSP model. When CTRs are additive, the allocation problem becomes trivial and unrealistic, because all slots go to the highest bidder. We therefore, consider families of *subadditive* CTR functions in this work.

We first consider the question of efficiently *allocating* rich ads, just as GSP ranks ads efficiently in the single-slot case. We initially consider allocation separately from questions of pricing and equilibrium. Because search results pages need to be rendered within milliseconds, we eschew general approaches based on integer programming and seek poly-time allocation schemes. Our first result is a randomized allocation scheme that achieves a $1/4$ approximation to efficiency under *fractionally subadditive* CTRs (also known as XOS, which includes submodular and substitutes CTRs as special cases).

The scheme is based on rounding of the solution to the associated configuration LP, and the constant factor approximation relies on the consecutive-slot structure. We provide a closely related scheme that achieves an $\mathcal{O}(\log \log m / \log m)$ approximation for general subadditive valuations, with m slots.

We next turn to the question of *pricing* for rich ads to achieve desirable equilibrium properties. The extension of GSP pricing to multi-slot allocation is ill-defined because the GSP auction is intimately linked to ad ranking rather than ad allocation. We consider two natural generalizations, *lowest bid* and *incremental* pricing, and show that they lack a crucial property of the original GSP: for both variants, an efficient pure-strategy Nash equilibrium (NE) may not exist. Recall that a pure NE is efficient if the slots are allocated efficiently with respect to the agents’ true underlying values-per-click, which is not equivalent to efficient allocation with respect to bids (because, like GSP, neither pricing rule is truthful). Under the lowest bid pricing rule, an agent is charged (per click) the lowest bid it could have placed and still obtained the same set of slots, holding other agents’ bids fixed. We show that a family of pricing rules that includes lowest-bid pricing fails to guarantee an efficient pure NE. Under incremental pricing, an agent is charged (per impression) for each slot the smallest impression price that clinches the slot; the total charge is then converted to a price per click according to the aggregate CTR of all slots received (i.e., the final rich ad). Again, we show that this pricing rule fails to guarantee an efficient pure NE.

With these results in mind, we turn to pricing rules specifically suited to combinatorial allocation. A natural first choice is the VCG pricing rule since it supports efficient allocation not just in pure NE, but in dominant strategies. However, VCG has been considered and rejected for rich ads in several papers for the same consistent reason: it is not a generalization of GSP, and the switch from GSP to VCG (in the standard setting with just text ads) would lead to significant revenue losses [Varian and Harris, 2014; Cavallo *et al.*, 2017; Hartline *et al.*, 2018]. Since rich ads allocation must handle text ads as a special case, this is a fatal drawback. An alternative known as *core pricing*, a standard scheme used in practice for combinatorial auctions, was recently proposed for rich ads [Hartline *et al.*, 2018]. Core pricing is attractive for rich ads because it admits natural efficient NE and guarantees a better revenue standard than VCG [Hartline *et al.*, 2018]. We consider *competitive equilibrium* pricing, a special case of core pricing where prices are defined over sets of slots such that each agent obtains its preferred (i.e., utility-maximizing) set of slots as bid. This is closely related to the original envy-freeness property of certain GSP equilibria in the original text ads setting [Edelman *et al.*, 2007].

Our main pricing result shows how to construct prices over bundles of slots such that the approximately-efficient allocations constructed via our earlier algorithms are individually-rational and envy-free. An allocation is individually-rational if each agent’s allocated bundle has non-negative utility at the given prices, and it is envy-free if each agent prefers its own bundle to that of any other. To complete the result, we show that the prices can be extended so that each agent obtains a utility-maximizing bundle (over all possible bundles) at the given prices, which forms a competitive equilibrium.

1.1 Related Work

The allocation, pricing, and incentive properties of the basic GSP were studied in a series of celebrated early papers [Aggarwal *et al.*, 2006; Edelman *et al.*, 2007; Varian, 2007]. Our aim in this paper is to stay true to this basic model and depart from it as slightly as possible. Our model has a fixed set of slots on the page, and CTR is separable into agent- and slot-dependent effects. This means that agents only place a single bid-per-click as in the original GSP. More general non-separable CTRs quickly lead to complications like inexistence of efficient equilibria [Cavallo and Wilkens, 2014]. Varian and Harris (2014) report that Google considered a switch to VCG to accommodate rich ads in search results, but this failed to launch because the revenue consequences were too severe; on the other hand, they report that VCG is used for contextual display advertising.

The problem of auction design for rich ads has attracted increased attention in recent years. Cavallo *et al.* (2017) report on a redesign of Yahoo’s sponsored search auction to handle rich ads, using a local search algorithm for allocation and pricing. Deng *et al.* (2010) and Bachrach *et al.* (2014) consider intuitive generalizations of the GSP pricing rule for rich ad allocation, different from the ones we examine but with similar conclusions that efficient equilibria may not exist.

Recent work on rich ads pricing has examined core pricing, drawing on ideas from the literature on combinatorial auctions [Day and Raghavan, 2007; Day and Milgrom, 2008]. Goel *et al.* (2015) study the revenue performance of core pricing for the problem of allocating text ads vs. larger image ads. Hartline *et al.* (2018) give efficient algorithms for bidder-optimal core pricing, specifically for rich ads. In contrast, our work specifically focuses on competitive equilibrium pricing, which connects with the original GSP papers given the close connection between envy-freeness and competitive equilibrium [Edelman *et al.*, 2007].

1.2 Preliminaries

Throughout this paper, we denote the set of advertisers by \mathcal{N} and the set of slots by \mathcal{M} . Let $n = |\mathcal{N}|$ and $m = |\mathcal{M}|$. We refer to the advertisers by a_i and to the slots by s_i , i.e., $\mathcal{N} = \{a_1, a_2, \dots, a_n\}$ and $\mathcal{M} = \{s_1, s_2, \dots, s_m\}$. Let b_i be the bid of agent a_i in the auction. Without loss of generality we assume that $b_i \geq b_j$ for each $i < j$. We also have a monotone set function for the click-through rate (CTR) of the slots. We denote by θ_S the CTR for set S of slots, and we consider the following set functions for the CTR function:

- **Additive:** A set function $V(\cdot)$ is additive if $V(S_1) + V(S_2) = V(S_1 \cup S_2) + V(S_1 \cap S_2)$ for every two sets $S_1, S_2 \in \text{ground}(V)$.
- **Submodular:** A set function $V(\cdot)$ is submodular if $V(S_1) + V(S_2) \geq V(S_1 \cup S_2) + V(S_1 \cap S_2)$ for every two sets $S_1, S_2 \in \text{ground}(V)$.
- **Fractionally Subadditive (XOS):** An XOS set function $V(\cdot)$ can be shown via a finite set of additive functions $\{V_1, V_2, \dots, V_\alpha\}$ where $V(S) = \max_{i=1}^\alpha V_i(S)$ for any set $S \subseteq \text{ground}(V)$.

- **Subadditive:** A set function $V(\cdot)$ is subadditive if $V(S_1) + V(S_2) \geq V(S_1 \cup S_2)$ for every two sets $S_1, S_2 \subseteq \text{ground}(V)$.

In this paper, the above definitions work for all the possible sets of slots (consecutive or non-consecutive). However, in the setting of consecutive slots we assume that we cannot allocate a set of slots with gaps to any advertiser. Hence, this is a constraint that we have in the allocation, not in the CTR functions.

We define $w_i(S) = b_i \theta_S$ as the welfare of bundle S for agent a_i . Our CTR function might be additive, gross substitutes, submodular, XOS, or subadditive. We say that an allocation of the slots to the agents is *efficient* if $\sum_i w_i(S_i)$ is maximized, where S_i is the set of allocated slots to agent a_i .

A *competitive equilibrium* (CE) consists of an allocation of slots to the advertisers with a price for each allocated bundle such that our allocation and pricing satisfy envy-freeness, market clearance, individual rationality, and budget balance.

1.3 Our Results

In this work, first we consider the efficient allocation problem under the rich ads model in Section 2. We give some computational results for approximating the efficiency for different levels of CTR functions—additive, submodular, XOS, etc. We give approximation results for XOS and subadditive CTR functions, which generalize submodular CTR functions. For the XOS CTR function over consecutive slots, we approximate the efficiency by a constant factor of $1/4$, and for the subadditive CTR function over consecutive slots, we approximate the efficiency by a factor of $\mathcal{O}(\frac{\log \log m}{\log m})$.

In the classic sponsored search auction, when we are allowed to allocate at most one slot to each advertiser, the GSP mechanism guarantees the existence of Nash Equilibrium. For the rich ads model, we consider two natural generalizations of GSP, Lowest Bids (LB) and Incremental Pricing (IP), and show that none of them guarantees the existence of an efficient Nash Equilibrium. We elaborate on these generalizations briefly in Section 3.

In Section 4, we obtain pricing rules that guarantee competitive equilibrium along with efficiency in the rich ads model, and we provide an algorithm which gives an allocation and pricing rule guaranteeing competitive equilibrium and $1/4$ approximation of the efficient allocation for XOS valuations and $\mathcal{O}(\frac{\log \log m}{\log m})$ approximation for subadditive valuations.

2 The Efficient Allocation

In this section, we study the efficient allocation of the slots to the advertisers for different CTR functions. As mentioned previously, an efficient allocation in the rich ads setting is an allocation maximizing $\sum_i b_i \theta_{S_i}$ where S_i is the set of allocated slots to bidder i . We study the problem of finding the efficient allocation for different CTR functions. First, we start with *additive* CTR functions.

Observation 2.1. *When the CTR function is additive, the efficient allocation allocates all the slots to the advertiser with the highest bid, b_1 .*

The problem of finding the efficient allocation becomes more challenging when we have more general CTR functions. The efficiency problem is closely related to the *Welfare Maximization Problem*, which is formulated as follows:

The Welfare Maximization Problem: Given m items and n players with monotone utility functions $w_i : 2^{[m]} \rightarrow R_+$, we seek a partition of the items into disjoint sets S_1, S_2, \dots, S_n in order to maximize $\sum_{i=1}^n w_i(S_i)$.

The welfare maximization problem has been extensively studied for different valuation functions. The Submodular Welfare Problem was first studied by Lehmann *et al.* (2006). They showed that a simple online greedy algorithm gives a $1/2$ -approximation for this problem. Khot *et al.* (2008) proved that the Submodular Welfare Problem cannot be approximated to a factor better than $1 - 1/e$, unless $P = NP$. Mirrokni *et al.* (2008) showed that a better than $(1 - 1/e)$ -approximation would require exponentially many value queries, regardless of $P = NP$.

[Feige, 2006] considers more general valuation functions. He gives a $(1 - 1/e)$ -approximation ratio for fractionally subadditive and a $1/2$ -approximation for subadditive valuation functions.

As the first step, we consider the problem of finding the efficient allocation when we are allowed to allocate non-consecutive slots to the advertisers. This problem is NP-hard for gross substitutes valuations (and clearly for more general valuations) by a reduction from Knapsack Problem [Lehmann *et al.*, 2006].

Theorem 2.1. *Using the results by [Feige, 2006], We can guarantee $(1 - 1/e)$ and $1/2$ approximation ratio of the efficient allocation for respectively fractionally subadditive and subadditive CTR functions.*

The above theorem gives theoretical results for the efficient allocation when the allocations can be non-consecutive. However, in practice, we usually need to allocate consecutive slots to each advertiser. With this constraint, unfortunately these results do not apply anymore.

In sponsored search auctions, higher slots typically have higher CTRs. In the classic version of the sponsored search auction, we allocate at most one slot to each advertiser. In such cases, GSP allocation guarantees the efficient allocation. It allocates the first slot to the highest bidder, the second slot to the second highest bidder, and so on. If we want to apply such a method of allocation in the consecutive rich ads setting, we should allocate the first consecutive bundle of slots to the highest bidder, the second consecutive bundle of slots to the second highest bidder, and so on. However, this method, in the consecutive rich ads setting does not always guarantee the efficiency.

2.1 Welfare Maximization in Consecutive Allocations

Now, in this part, we give an algorithm with $1/4$ approximation ratio for the welfare maximization problem when the allocation is consecutive and the CTR functions are fractionally subadditive. We also give an algorithm for the same

problem when the CTR functions are subadditive. Both of these algorithms are based on the randomized rounding of the configuration LP of the problem. For the welfare maximization problem, [Dobzinski *et al.*, 2010] considered the following configuration LP of the problem, that we call *welfare maximizing* LP. By rounding this configuration LP, they gave the first approximation ratios for general valuation settings. Later, [Feige, 2006] used a better rounding method of this configuration LP and improved the approximation ratios for subadditive and fractionally subadditive valuation functions. The LP is as follows:

$$\begin{aligned} \max_{x_{i,S} \geq 0} \quad & \sum_{i,S} x_{i,S} w_i(S) \\ \text{s.t.} \quad & \sum_{i,S|j \in S} x_{i,S} \leq 1 \quad (j \in \mathcal{M}) \quad (1) \\ & \sum_S x_{i,S} \leq 1 \quad (i \in \mathcal{N}) \quad (2) \end{aligned}$$

In the previous works, subset S could be any subset of items, but here, we only consider the consecutive bundle of items as feasible subsets in our LP. Here, $x_{i,S}$ is an indicator variable specifying whether agent i gets set S . The constraints in Inequality (1) are the *item constraints*, and the constraints in Inequality (2) are the *agent constraints*. In each item constraint, we make sure each item is not allocated in more than one allocation, and in agent constraints we make sure that we have not allocated more than one set to each agent. Note that we mean the summations of fractional allocations should not be more than one. As mentioned in [Feige, 2006], we can assume the solution given to the LP is such that all item and agent constraints are satisfied with equality. This convention can easily be made without loss of generality as follows: If the item constraint for item j is not satisfied with equality, we can add a dummy agent a_j with a utility function that is identically zero, and set the value of variable $x_{a_j,j} = 1 - \sum_{i,S|j \in S} x_{i,S}$. Similarly, if the agent constraint for agent i is not satisfied with equality, we can set the value of variable $x_{i,\emptyset}$ to $1 - \sum_{S \neq \emptyset} x_{i,S}$.

The previous works considered this configuration LP for any subset S of items, but here, in our application, we only consider the consecutive subsets as feasible subsets. Unfortunately, here, in the consecutive allocation setting, we cannot directly use their rounding methods. The reason is that in all these methods, in a part of the algorithm, they allocate every single item randomized. Here, since the allocation for each agent should be consecutive, we cannot have such randomizations in our solution. Hence, we should round our configuration LP in another way or somehow modify their solution. To this end, we devise Algorithm 1 which provides a constant approximation solution for the fractionally subadditive valuations. In the first step of this algorithm, we find the fractional solution of the LP. Then, for each agent, with probability $1/2$ we do not allocate her any bundle, and with probability $1/2$ we allocate her exactly one bundle. If we allocate one bundle to agent i , the bundle (say bundle S) is chosen with probability $x_{i,S}$. Since we do not allocate any bundle to each agent with probability $1/2$ the actual probability for agent i and bundle S is $x_{i,S}/2$. Now, up to this point of the algorithm,

Algorithm 1: Consecutive Welfare Maximization with XOS Valuations

Input: Set of agents \mathcal{N} and set of items \mathcal{M} (with a specific order)

- 1 Find the fractional solution for the configuration LP;
- 2 **for each agent** $i \in \mathcal{N}$ **do**
- 3 With probability $1/2$ do not allocate any item to her, and with the remaining probability allocate her a single consecutive subset S which is chosen with probability $x_{i,S}$;
- 4 **end**
- 5 **for each item** $j \in \mathcal{M}$ **do**
- 6 **if** item j was allocated to more than one agent **then**
- 7 Deallocate it from all the agents;
- 8 **end**
- 9 **end**
- 10 **for each agent** $i \in \mathcal{N}$ **do**
- 11 Allocate her all the items from x to y , if x was the leftmost allocated item to i , and y was the rightmost allocated item to her;
- 12 **end**
- 13 **return** the allocations;

each agent has at most one allocated bundle, but still, each item may be allocated to more than one agent. To avoid this, first, we deallocate all the items allocated to more than one agent. However, this may make our allocations disconnected. To handle this, if we had a disconnection in the allocation of an agent, we allocate all the items in the gaps to the agent to make it a connected allocation. In Theorem 2.3, we prove this algorithm guarantees $1/4$ approximation ratio of the optimal solution. The following proposition is from [Feige, 2006] which is useful in the proof of the theorem.

Proposition 2.2. *Let $k \geq 1$ be an integer and let w be an arbitrary fractionally subadditive utility function. For a set S , consider a distribution over subsets $S' \subset S$ such that each item of S is included in S' with probability at least $1/k$. Then $E[w(S')] \geq w(S)/k$.*

Theorem 2.3. *Algorithm 1 guarantees a $1/4$ approximation ratio for the welfare maximization problem, when the allocation is consecutive and valuation functions are fractionally subadditive.*

Proof. Assume that $w(LP)$ is the welfare of fractional solution of our configuration LP, and $w(IP)$ is the welfare of our integral solution. As the first step, Algorithm 1 finds the fractional solution of the LP. We know $w(LP) \geq w(IP)$, thus any approximation guarantee for $w(LP)$ gives a guarantee for $w(IP)$. In the first rounding step of the algorithm, we randomly round the LP and assign each bundle S with probability $x_{i,S}/2$ to agent i . Hence, after this rounding, each agent with probability $1/2$ does not receive any bundle, and with probability $1/2$ exactly receives one bundle of consecutive items. Hence, we can argue the current allocation in expectation has welfare $w(LP)/2$, but unfortunately, the allocation may assign each item to multiple agents.

Now, we prove for each agent i , each item j assigned to i is not assigned to any other agent with probability at least $1/2$. The idea of the proof comes from the item constraints and the way that we rounded the LP. The probability that we allocate item j to any other agent i' is $p_{j,i'} = \sum_{S|j \in S} x_{i',S}/2$. Hence, with probability $1 - p_{j,i'}$, item j is not allocated to i' , and we can conclude with probability $\prod_{i' \neq i} (1 - p_{j,i'})$, item j is not allocated to any agent $i' \neq i$. Since $\sum_{i' \neq i} p_{j,i'} \leq 1/2$, we can argue that $\prod_{i' \neq i} (1 - p_{j,i'}) \geq 1/2$. With this in hand, and using Proposition 2.2, we can argue that if we did not need to allocate the items consecutively, the solution by ignoring the items allocated to more than one agent is at least $w(LP)/4$. In the last step of the algorithm, to make the allocated bundles consecutive we may only allocate some items to each agent. Hence after this step, our social welfare does not decrease, and it is still at least $w(LP)/4$.

Note that Algorithm 1 does not necessarily allocate all the items. If we want to make sure all the items are allocated, we should allocate all the items in each gap to the agent who has received the next or previous consecutive items. \square

Note that the above theorem is a result in expectation. To have a result in the worst case, we should run the same algorithm several times and take the best answer.

For subadditive valuations, similar to Algorithm 1, first, we find the fractional solution for the configuration LP. However, here, we use a rounding technique which is also used in [Feige, 2006]. We allocate consecutive bundle S to agent i with probability $x_{i,S}$. Please note that we allocate exactly one bundle to each agent. This is doable because according to the agent constraints, for each agent i , we have $\sum_S x_{i,S} = 1$. However, each item may be allocated to more than one agent. According to [Feige, 2006], with probability at least $1 - 1/m$, no item belongs to more than $k = \mathcal{O}(\log m / \log \log m)$. Now, we want to give a color to each agent such that two agents do not have the same color if they share an item. To this end, we define the *intersection graph* of the agents as follows:

Definition 2.4. *In the intersection graph of the agents, we have a node for each agent, and we connect two agents with an edge if and only if there exists a common item in their allocations.*

We color the vertices of the intersection graph. This is doable with at most $\Delta + 1$ colors where Δ is the maximum degree of the graph. Since with high probability, the degree of each vertex is k , with high probability the coloring is doable with $k + 1$ colors. As the last step of the algorithm, we choose one of the colors uniformly at random and deallocate the items from the agents with unchosen colors.

Theorem 2.5. *Algorithm 2 guarantees an $\mathcal{O}(\log \log m / \log m)$ approximation ratio for the welfare maximization problem, when the allocation is consecutive and valuation functions are subadditive.*

3 On Generalizations of GSP

In sponsored search auctions, Generalized Second Price (GSP) auction is used to guarantee the existence of a Nash

Algorithm 2: Consecutive Welfare Maximization with Subadditive Valuations

Input: Set of agents \mathcal{N} and set of items \mathcal{M} (with a specific order)

- 1 Find the fractional solution for the configuration LP;
- 2 **for each agent** $i \in \mathcal{N}$ **do**
- 3 With probability $x_{i,S}$, allocate a unique consecutive subset S to agent i ;
- 4 **end**
- 5 Color the intersection graph of the agents;
- 6 Choose one of the colors of the graph uniformly at random;
- 7 Deallocate all the items from the agents with unchosen colors;
- 8 **return** the allocations;

Equilibrium and efficiency. In order to generalize GSP, we study two models that we call *Lowest Bids Method* and *Incremental Pricing Method*. Although these are intuitive generalizations, we provide counterexamples with two or three advertisers and two slots to show that these methods do not guarantee the existence of Nash Equilibrium.

The next theorem indicates that the efficient allocation with additive CTR functions is to allocate all slots to the advertiser with the highest bid.

Theorem 3.1. *When the CTR function is additive the following assignment and pricing rule is efficient and guarantees Nash equilibrium: Allocate all the slots to the advertiser with the highest bid with price b_2 per click. The total payment is $b_2 \theta_{\mathcal{M}}$.*

Consequently, in this section, we focus on the generalizations of GSP for more general CTR functions.

Definition 3.2 (Lowest Bids Rule). *Assume that p_i is the price of the i^{th} slot when we run regular GSP (in the regular GSP we allocate at most one slot to each agent). In Lowest Bid rule, first we allocate the slots in an efficient way, then, we charge each bidder the lowest bid with which the allocation does not change.*

Theorem 3.3. *The Lowest Bid method does not guarantee Nash Equilibrium.*

Now, we consider another generalization of GSP that we call *Incremental Pricing Rule*.

Definition 3.4 (Incremental Pricing Rule). *Under incremental pricing, an agent is charged (per impression) for each slot the smallest impression price that clinches the slot; the total charge is then converted to a price per click according to the aggregate CTR of all slots received. We note that IP does not necessarily allocate consecutive slots.*

Theorem 3.5. *The Incremental Pricing method does not guarantee the existence of pure Nash equilibrium even if the non-consecutive allocation is allowed.*

4 Competitive Equilibria

In this section, we study the pricing rules guaranteeing competitive equilibrium along with efficiency in the rich ads setting. First, consider additive CTR functions. In this setting,

Algorithm 3: Efficient Envy-free Pricing Method

Input: Set of agents \mathcal{N} and set of items \mathcal{M}

- 1 Find the approximated efficient allocation;
 - 2 Construct bipartite graph G based on the advertisers and their allocated bundles of the slots;
 - 3 Find the minimum weighted vertex cover of G ;
 - 4 **for each bundle S_i do**
 - 5 set $\frac{c_i}{\theta_{S_i}}$ as its price, and c_i is the assigned number by the minimum weighted vertex cover to S_i
 - 6 **end**
 - 7 Return the efficient allocation with the prices assigned to the bundles;
-

if we allocate all the slots to advertiser a_1 with the price per click b_2 , clearly it is efficient and guarantees the competitive equilibrium. Hence, we consider the problem for more general CTR functions.

First we normalize the CTR functions such that $\theta_M = 1$ and provide Algorithm 3. As the first step of the algorithm, we find the approximately efficient allocation (using the methods that we discussed in Section 2) and assign the slots to the agents based on the allocation. This is our allocation, and now our goal is to price the assigned bundles of slots such that we have an envy-free allocation.

As the second step of the algorithm, we construct a bipartite graph G as follows: The nodes of the first partition of the graph represent our advertisers and the second partition of the graph represent the assigned bundles. We connect each advertiser a_i to a bundle S_j using an edge with weight $b_i\theta_{S_j}$, which is the value of bundle S for advertiser a_i . The maximum weighted matching in graph G finds our (approximately-) efficient assignment. Now, as the next step, we find the minimum weighted vertex cover of G . This algorithm assigns a weight to each vertex of G . Assume that c_i is the assigned number to the vertex of bundle S_i . We set $p_i = c_i/\theta_{S_i}$ as the price of S_i . We also assign infinite price to a bundle of items S if there exists no node for S in G .

In the next lemma, we prove that our pricing rule guarantees envy-freeness.

Lemma 4.1. *Algorithm 3 guarantees envy-freeness and individual rationality.*

Proof. First, we prove the profit of each advertiser is non-negative. In the minimum weighted vertex cover of a bipartite graph, we know the summation of the assigned weights to the endpoints of an edge is equal to the weight of the edge if the edge is in the maximum weighted matching. Hence, we argue that the assigned number to any bundle S_i is not more than the utility of the bundle for the advertiser receiving the bundle in the efficient allocation. Assume that in the efficient allocation S_i is assigned to advertiser a_j , and c_i is the assigned number to S_i by the minimum weighted vertex cover algorithm. As we mentioned, $b_j\theta_{S_i} \geq c_i$. Since $p_i = c_i/\theta_{S_i}$, we have $b_j \geq p_i$. Hence, the profit of each advertiser is non-negative.

Now, for the sake of contradiction, assume that our auction is not envy-free, and advertiser a_i envies advertiser a_j . Assume that $S_{i'}$ and $S_{j'}$ are the assigned bundles to a_i and a_j in

the efficient allocation respectively. Now, we have:

$$b_i\theta_{S_{i'}} - c_{i'} < b_i\theta_{S_{j'}} - c_{j'}.$$

Hence, we have the following inequality:

$$b_i(\theta_{S_{i'}} - \theta_{S_{j'}}) < c_{i'} - c_{j'} \quad (3)$$

Assume that c_i , $c_{i'}$, and $c_{j'}$ are the assigned numbers by the minimum weighted vertex cover to a_i , $S_{i'}$, and $S_{j'}$ respectively. According to the properties of the minimum weighted vertex cover, we have the following:

$$c_i + c_{i'} = b_i\theta_{S_{i'}} \quad (4)$$

$$c_i + c_{j'} \geq b_i\theta_{S_{j'}} \quad (5)$$

Hence, using Equality [4] and Inequality [5], we have:

$$b_i(\theta_{S_{i'}} - \theta_{S_{j'}}) \geq c_{i'} - c_{j'} \quad (6)$$

As we see, Inequalities [3] and [6] clearly contradict each other. Hence, our assignment is envy-free. \square

Algorithm 3 and Lemma 4.1 imply the following theorem:

Theorem 4.2. *Algorithm 3 gives an allocation and non-linear pricing rule guaranteeing competitive equilibrium and 1/4 approximation of the efficient allocation for XOS valuations and $\mathcal{O}(\frac{\log \log m}{\log m})$ approximation for subadditive valuations.*

5 Conclusions

This paper cast the problem of auctioning search ads with extensions as a generalization of the standard GSP model where advertisers can now obtain multiple consecutive slots, interpreted as a rich ad. We obtained a 1/4 approximation algorithm for the efficient allocation problem when CTR functions are fractionally subadditive, and a $\mathcal{O}(\log \log(m)/\log(m))$ approximation ratio when the CTR functions are subadditive. Our final pricing result shows how to support these approximately efficient allocations using bundle prices such that the allocation-price pairs are individually-rational and envy-free.

A key open question concerns generalizations of GSP that admit pure strategy Nash Equilibria (as opposed to competitive equilibria). We showed that two natural generalizations of the GSP pricing rule fail to guarantee Nash Equilibria even under two slots, but with appropriate restrictions on the CTR function (e.g., gross substitutes), we believe that a generalization of all the GSP properties (efficiency, envy-freeness, Nash Equilibrium) may still be possible. Another interesting direction would be to generalize the model and results to the problem of whole page optimization in display advertising, where ads need to be allocated in a two-dimensional layout on the page.

Acknowledgments

This paper is supported in part by Guggenheim Fellowship, NSF grants CCF:SPX 1822738 from , IIS:BIGDATA 1546108, DARPA grant SI3CMD, UMD Year of Data Science Program Grant, and Northrop Grumman Faculty Award.

References

- [Aggarwal *et al.*, 2006] Gagan Aggarwal, Ashish Goel, and Rajeew Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 1–7, 2006.
- [Bachrach *et al.*, 2014] Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, and David Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the fifteenth ACM Conference on Economics and Computation*, pages 75–92, 2014.
- [Cavallo and Wilkens, 2014] Ruggiero Cavallo and Christopher A. Wilkens. GSP with general independent click-through-rates. In *International Conference on Web and Internet Economics*, pages 400–416. Springer, 2014.
- [Cavallo *et al.*, 2017] Ruggiero Cavallo, Prabhakar Krishnamurthy, Maxim Sviridenko, and Christopher A. Wilkens. Sponsored search auctions with rich ads. In *Proceedings of the 26th International Conference on World Wide Web*, pages 43–51, 2017.
- [Day and Milgrom, 2008] Robert Day and Paul Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, 36(3-4):393–407, 2008.
- [Day and Raghavan, 2007] Robert Day and Subramanian Raghavan. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53(9):1389–1406, 2007.
- [Deng *et al.*, 2010] Xiaotie Deng, Yang Sun, Ming Yin, and Yunhong Zhou. Mechanism design for multi-slot ads auction in sponsored search markets. In *International Workshop on Frontiers in Algorithmics*, pages 11–22. Springer, 2010.
- [Dobzinski *et al.*, 2010] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Mathematics of Operations Research*, 35:1–13, 2010.
- [Edelman *et al.*, 2007] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the Generalized Second-Price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [Feige, 2006] Uriel Feige. On maximizing welfare when utility functions are subadditive. In *Symposium on the Theory of Computing*, pages 41–50, 2006.
- [Goel *et al.*, 2015] Gagan Goel, Mohammad Reza Khani, and Renato Paes Leme. Core-competitive auctions. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 149–166, 2015.
- [Hartline *et al.*, 2018] Jason Hartline, Nicole Immorlica, Mohammad Reza Khani, Brendan Lucier, and Rad Niazadeh. Fast core pricing for rich advertising auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 111–112, 2018.
- [Khot *et al.*, 2008] Subhash Khot, Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica* 52(1), pages 3–18, 2008.
- [Lehmann *et al.*, 2006] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55(2), pages 270–296, 2006.
- [Mirrokni *et al.*, 2008] Vahab Mirrokni, Michael Schapira, and Jan Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 70–77, 2008.
- [Varian and Harris, 2014] Hal R. Varian and Christopher Harris. The VCG auction in theory and practice. *American Economic Review*, 104(5):442–45, 2014.
- [Varian, 2007] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.