Externalities among Advertisers in Sponsored Search*

Dimitris Fotakis[1] Piotr Krysta[2] Orestis Telelis[2]

[1] School of Electrical and Computer Engineering,
National Technical University of Athens, 157 80 Athens, Greece.
Email: fotakis@cs.ntua.gr
[2] Department of Computer Science,
University of Liverpool, L69 3BX Liverpool, UK.
Email: {p.krysta,o.telelis}@liverpool.ac.uk

Abstract. We introduce a novel computational model for single-keyword auctions in sponsored search, which explicitly models externality effects among the advertisers, an aspect that has not been (fully) reflected in the existing models, and is known to be prevalent in the behavior of real advertisers. Our model takes into account both positive and negative correlations between any pair of advertisers, and appropriately reflects them in the way they perceive an outcome of a sponsored search auction. In our model, the click-through rate of an ad depends on the set of other ads appearing in the sponsored list, on their relative order, and on their distance in the list. In contrast to previous modeling attempts, we avoid modeling end-users' behavior, but only resort to a reasonable assumption that their browsing focuses on any bounded scope section of the list. We present a comprehensive collection of computational results concerning the winner determination problem in our model with an objective to maximize the social welfare, showing both hardness of approximation results and polynomial-time approximation algorithms. We also present an exact polynomial-time algorithm for the practically relevant cases of our model, where the length of the sponsored list is at most logarithmic in the number of advertisers. This exact algorithm can be used as a truthful mechanism by employing the VCG payments, thus showing that we can fully cope with selfish behavior of advertisers when the mechanism is fully aware of their correlations. We finally study the performance of the practically used Generalized Second Price auction mechanism in our model and show that, in presence of externalities, pure Nash equilibria may not exist for conservative bidders that do not outbid their valuation. Moreover, we exhibit instances where pure Nash equilibria do exist, but they carry an unbounded loss in social welfare as compared to the socially optimal solution assignment of slots, even for such conservative bidders.

* Research partially supported by an NTUA Basic Research Grant (PEBE 2009), EPSRC grant EP/F069502/1, and DFG grant Kr 2332/1-3 within Emmy Noether Program.

1 Introduction

Sponsored search advertising is nowadays a predominant and arguably most successful paradigm for advertising products in a market, facilitated by the Internet. It constitutes a major source of income for popular search engines like `Google`, `Yahoo!` or Microsoft `Bing`, who allocate up to $8-9$ advertisement slots in their sites, alongside the organic results of keyword searches performed by end users. Each time an end user makes a search for a keyword, slots are allocated to advertisers by means of an auction performed automatically; advertisers are ranked in non-increasing order of a *score*, defined as the product of their *bid* with a characteristic *relevance* quantity per advertiser. The relevance of each advertiser is interpreted as the probability that his ad will be clicked by an end user. The score corresponds then to the *declared expected revenue* of advertisers. Advertisers ranked higher are matched to higher slots. Each advertiser is charged an amount depending on the score (hence, the bid) of the one ranked below him, when his ad is clicked.

The described auction is used in varying flavors by search engines. Apart from the mentioned *Rank-By-Revenue* rule, a plainer *Rank-By-Bid* rule has also been used (e.g. by `Yahoo!`). This auction, known as the *Generalized Second Price* (GSP) auction, constitutes a generalization of the well-known strategyproof Vickrey auction [30]. The GSP auction is not strategyproof though; it encourages strategic bidding by the advertisers instead of eliciting their valuations truthfully. This induces a strategic game rich in Nash equilibria, and competitive behavior among advertisers that incurs significant revenue to the search engines. Pure Nash Equilibria of the GSP mechanism were first studied by Edelman, Ostrovsky, Schwartz [17] and Varian [29], under what came to be known as the *separable click-through rates* model. In this model every slot is associated with a *Click-Through Rate* (CTR), i.e. the probability that an ad displayed in this slot will be clicked. The joint probability that an ad is clicked is given by the product of the slot's CTR with the relevance of the ad. Since [17, 29], a rich literature on keyword auctions has been published, concerning algorithmic and game theoretic issues such as bidding strategies, social efficiency, revenue, see, e.g. [27] (chapter 28).

Recent experiments [22] show that the probability of a displayed ad being clicked (hence, the utility of the advertiser) is affected by its relative order and distance to other ads on the list. E.g., two competing ads displayed nearby each other may distract a user's attention from any of them. However, two advertisers may also profit from being displayed nearby (e.g. a cars manufacturer and a spare parts supplier). We introduce a model for expressing such *externalities* in keyword auctions. Externalities result in complicated strategic competition among advertisers, precluding stable outcomes and social welfare optimization. A way of alleviating these effects is by solving optimally the *Winner Determination* problem [2, 1, 23]; i.e., the problem of selecting winners and their assignment to slots, to maximize the *Social Welfare* (the sum of the advertisers' utilities and the auctioneer's revenue – defined formally in Section 2). Such a solution, paired with payments of the Vickrey-Clarke-Groves mechanism [14] (chapter 1), yields a truthful mechanism. On the other hand, the study of the GSP auction's performance in presence of externalities [20] yields insights for the current practice of sponsored search. In our model we analyze exact and approximation algorithms for *Winner Determination*, and show how externalities can harm the GSP auction's stability and social efficiency.

Several recent works concern theoretical and experimental study of externalities in keyword auctions [1, 7, 19, 20, 23, 21, 22]. These works associate the occurrence of externalities with a model of how end users search through the list of ads and how this affects the probability of an ad being clicked. This search is commonly modeled by a top-down ordered scan of the list [7, 1, 23, 20]. The popular *Cascade Model* [1, 23], associates a *continuation probability* with every ad, i.e. the likelihood of the user continuing his ordered scan after viewing the ad. Using real data from keyword auctions, Jeziorski and Segal [22] found that previously proposed models fail to express externalities. The *Cascade Model* is contradicted by the fact that about half of the users *do not* click on the ads sequentially, i.e., they return to higher slots after clicking on lower slots. Jeziorski and Segal arrived at a structural model of end user behavior advocating that, after scanning all the advertisements, users focus on a subset of consecutive slots. Within this focus *window*, they observed externalities due to proximity and relative order of the displayed ads. They highlight that an ad's CTR on the list depends crucially on the ads displayed in the other slots, above and below it. With our model we aim at quantifying rigorously these observations, while avoiding explicit modeling of end users' behavior.

Contribution and Techniques. Our main contribution is a novel and fairly expressive model for describing positive or negative influences between any two advertisers in keyword auctions. Our model is built by usage of a social context graph [6] and upon the assumption of separable click-through rates, which facilitates analytical study. (see Section 2 for precise definition). It provides for the description of influences among the advertisers' relevances, depending on: **(i)** The IDs of advertisers, **(ii)** Their relative order in the list of sponsored links and **(iii)** The distance of their ads in the list. Motivated by results of [22], we make the assumption that the end users' attention may be captured by any size c *window* of consecutive slots, within which externalities take effect among advertisers. Our model is the first to express inherently the practically relevant possibility that social welfare maximization may occur under partial allocation of slots (i.e., where some of the slots are left empty).

Under the proposed model we study the *Winner Determination* problem with the objective of maximizing the social welfare. Using reductions from the Longest Path problem and the Traveling Salesperson Problem with distances 1 and 2 [28], we show that the problem is $\mathbf{NP}$-hard and $\mathbf{APX}$-hard respectively, even for window size 1 and *positive-only* externalities. For positive-only externalities we study approximation algorithms. We develop two interesting approximation-preserving reductions of the Winner Determination problem to the Weighted m-Set Packing problem with sets of size $m = 3$ and $m = 2c + 1$ respectively. These result in a nice tradeoff between approximability and computational efficiency for the Winner Determination problem; for k slots and n advertisers, we obtain algorithms with approximation factors: **(i)** $6c$ in time $O(kn^2 \log n)$, **(ii)** $4c$ in time $O(k^2n^4)$ and **(iii)** $2(c + 1)$ in polynomial time for any constant c. These results settle almost tightly the approximability of the problem for positive-only externalities and $c = O(1)$.

On the positive side, we build on the color coding technique [3] to obtain an exact algorithm for Winner Determination in the full generality of our model. A derandomization of our color coding is possible, that yields a deterministic algorithm. The algorithm has running time $2^{O(k)}n^{2c+1} \log^2 n$, which shows that the class of practically interesting instances of the problem with $c = O(1)$ and even $k = O(\log n)$ is in $\mathbf{P}$. This algorithm, paired with VCG payments [14] yields a truthful mechanism when the advertisers' valuations are private information, that optimizes the social welfare in presence of externalities. Notice that the winner determination problem can be solved easily by complete enumeration of all $\binom{n}{k}$ tuples of ads, in $O(n^k)$ time. For any constant size window our algorithm is significantly faster, and even for larger than constant values of k. Thus the problem is *not* $\mathbf{NP}$-hard for $k = O(\text{poly}(\log n))$, unless $\mathbf{NP} \subseteq \mathbf{DTIME}(n^{\text{poly}(\log n)})$.

We conclude our study with an investigation of the GSP mechanism's behavior under externalities. We find that pure Nash equilibria do not exist in general, for *conservative bidders* that do not outbid their valuation. This is a standard assumption when studying non-truthful mechanisms [11, 26, 13], since "overbidding" results in unboundedly low social welfare. Even when such conservative equilibria exist, we show that their social welfare can be arbitrarily low compared to the social optimum. This should be contrasted with an upper bound of 1.282 shown recently by Caragiannis *et al.* [11] for the social inefficiency of GSP equilibria.

2 A Model for Externalities in Keyword Auctions

We consider a set $\mathcal{N} = [n] = \{1, 2, \ldots n\}$ of n advertisers (or players/bidders) and a set $\mathcal{K} = \{1, 2, \ldots, k\}$ of $k \leq n$ advertisement slots. Each advertiser $i \in N$ has valuation v_i for each click and is associated with a probability q_i that her ad receives a click, independently of slot assignment. q_i is often termed *relevance* and measures the intrinsic "quality" of ad i. Each slot j is associated with a probability λ_j that an ad displayed in slot j will be clicked; this is called the *Click-Through Rate* (CTR) of the slot. The (overall) Click-Through Rate *of an ad* $i \in \mathcal{N}$ occupying slot j is $\lambda_j \cdot q_i$ (separable CTRs). We assume $1 \geq \lambda_1 \geq \ldots \geq \lambda_k > 0$, i.e. that the top slot is slot 1, the second one is slot 2 and so on. Let $S \subseteq [n]$, $|S| \leq k$, be the set of winning ads and $\pi : S \to \mathcal{K}$ their assignment to slots, i.e. $\pi(i) \in \mathcal{K}$ is the slot of advertiser $i \in S$. Every advertiser $i \in S$ issues to the search engine (auctioneer) a payment p_i per received click, hence receives expected utility $u_i(S, \pi) = \lambda_{\pi(i)} \cdot q_i \cdot (v_i - p_i)$. The *Social Welfare* is then:

$$\mathrm{sw}(S, \pi) = \sum_{i \in S} \lambda_{\pi(i)} \cdot q_i \cdot v_i = \sum_{i \in S} u_i(S, \pi) + \sum_{i \in S} \lambda_{\pi(i)} \cdot q_i \cdot p_i, \tag{1}$$

where $\sum_{i \in S} \lambda_{\pi(i)} \cdot q_i \cdot p_i$ is the expected revenue of the auctioneer. Our modeling of externalities is built on top of the separable CTRs model and induces amplified or diminished actual relevance Q_i (compared to q_i) to the advertisers, depending on their relative position and distance of their ads on the list. We use a directed *social context graph* [6] $G(\mathcal{N}, E^+, E^-)$, defined upon the set of advertisers $\mathcal{N}$. Each edge $(i, j) \in E^+ \cup E^-$ is associated with a function $w_{ij} \in (0, 1)$. An edge (j, i) in E^+ or in E^- denotes potential positive or negative influence respectively of i by j. Namely, edges in E^+ and E^- model respectively *positive* and *negative* externalities among advertisers. For any edge $(j, i) \in E^+ \cup E^-$, the potential influence of j to i is quantified by a function $w_{ji} : \{-k+1, -k+2, \ldots, -1, 1, \ldots, k-2, k-1\} \mapsto [0, 1]$. If $d_\pi(j, i) = \pi(i) - \pi(j)$ is the distance of ad i from ad j in the list, $w_{ji}(d_\pi(j, i))$ is the probability that a user's interest in a displayed ad j may result in attraction or distraction of his attention to or from an ad i respectively, depending on whether $(j, i) \in E^+$ or $(j, i) \in E^-$ [3]. It is reasonable to assume that the closer i and j are in π, the stronger the influence of j on i is. Formally, for every ℓ, ℓ', with $|\ell|, |\ell'| \geq 1$, if $|\ell| < |\ell'|$, then $w_{ji}(\ell) \geq w_{ji}(\ell')$.

Let $S \subseteq \mathcal{N}$ be the set of winners and π be the permutation assigning them to the slots. Define the subgraph $G_S(S, E_S^+, E_S^-)$ of the context graph, induced by S. The probability $Q_i(S, \pi)$ that a user is attracted by ad i is expressed as product $Q_i^+(S, \pi) \times Q_i^-(S, \pi)$. $Q_i^+(S, \pi)$ is the probability that i attracts the user's attention either by its intrinsic relevance q_i or by receiving positive influence. $Q_i^-(S, \pi)$ is the probability that the user's attention is *not* distracted from i due to negative influence of others. For each $i \in S$ define $N_i^+(S) = \{j \in S : (j, i) \in E_S^+\}$ and $N_i^-(S) = \{j \in S : (j, i) \in E_S^-\}$ to be the set of neighbors of i in G_S with positive and negative influence respectively. Let us derive $Q_i^+(S, \pi)$ first. A user's attention is *not* attracted by the ad of i with probability $(1 - q_i)$ and – independently – if i is not positively influenced by any $j \in N_i^+(S)$. The latter occurs either because j itself does not attract the user (with probability $1 - q_j$), or the positive influence of j to i does not occur (with probability $q_j(1 - w_{ji}(d_\pi(j, i)))$. Then j does not influence i with probability $(1 - q_j) + q_j(1 - w_{ji}(d_\pi(j, i))) = (1 - q_j w_{ji}(d_\pi(j, i)))$ and:

$$Q_i^+(S, \pi) = 1 - (1 - q_i) \cdot \prod_{j \in N_i^+(S)} (\, 1 - q_j w_{ji}(d_\pi(j, i)) \,) \tag{2}$$

[3] We note that if $d_\pi(j, i) > 0$, i appears below j, and if $d_\pi(j, i) < 0$, j appears below i in π.

For $Q_i^-(S,\pi)$ we use similar reasoning. A user's attention is not distracted by an ad of i due to negative influence of $j \in N_i^-(S)$ if either his attention is not captured by j or, if it is, j fails to influece i negatively. This event occurs with probability $(1-q_j)+q_j(1-w_{ji}(d_\pi(j,i))) = 1-)q_j w_{ji}(d_\pi(j,i)$. Assuming independence of the events for all $j \in N_i^-(S)$, we have:

$$Q_i^-(S,\pi) = \prod_{j \in N_i^-(S)} (\, 1 - q_j w_{ji}(d_\pi(j,i))\,)$$

After deriving $Q_i(S,\pi) = Q_i^+(S,\pi) \times Q_i^-(S,\pi)$, we can restate the social welfare as:

$$\mathrm{sw}(S,\pi,G) = \sum_{i \in S} \lambda_{\pi(i)} \cdot Q_i(S,\pi) \cdot v_i\,. \tag{3}$$

Arguably, users' attention and memory when they process the advertisements list have a bounded scope. Therefore, we assume that there is an integer constant $c > 0$, called *window size*, so that each ad j can only affect other ads at a distance at most c in π. Formally, if window size is c, for all $(j,i) \in E^+ \cup E^-$ and all integers ℓ with $|\ell| > c$, $w_{ji}(\ell) = 0$. Then the relevance Q_i only depends on the ads in $N_i^+(S) \cup N_i^-(S)$ assigned to slots at distance at most c from i.

Related Work. Edelman, Ostrovsky, Schwartz [17] and Varian [29] first modeled the game induced by the GSP auction mechanism under the assumption of separable CTRs, i.e. that the probability of a specific ad being clicked when displayed in a certain slot is given by the product of the slot's CTR and the ad's relevance. They identified socially optimal pure Nash equilibria of the game, in which advertisers are ranked by non-increasing score computed by their valuations. Prior to these works, Aggarwal, Goel and Motwani [2] had already designed a truthful mechanism for non-separable CTRs (in this case the VCG auction is also not applicable). For separable CTRs they proved revenue equivalence of their mechanism to the VCG.

There has been a growing interest in modeling externalities in sponsored search and studying how externalities affect the advertisers' bidding strategies and the properties of GSP equilibria. The unpublished work of Das *et al.*[16] is probably the closest in spirit to ours. Das et al. consider externalities among advertisers based on their relative quality (i.e., relevance), and quantify the fact that for any sponsored list of ads, as the relevance of an ad increases, its CTR should increase and the CTRs of the other ads should decrease. However, the model of [16] treats advertisers as anonymous, in the sense that externalities do not depend on their bids, but just on their relative quality. Our model makes a step further, since in addition to the advertisers' relative quality as measured by their relevance, we also take into account their IDs, their dependencies in the context graph, and their relative distance in the sponsored list. Moreover, our model quantifies how as the relevance of an ad increases, its CTR and the CTRs of positively correlated ads appearing nearby in the list should increase, and the CTRs of negatively correlated ads appearing nearby in the list should decrease.

Athey and Ellison [7] gave one of the first models for externalities, where they assumed an ordered top-down scan of slots by end users. By assuming a certain cost incurred to the end users for clicking on an ad, they derived the equilibria of the GSP auction for their game. Along similar lines, Aggarwal *et al.* [1] and Kempe and Mahdian [23] studied *Cascade Models* involving Markovian end-users, where every ad is associated with its individual CTR and a *continuation probability*. The latter denotes the probability that an end-user will continue scanning the list of ads (in a top-down order) after viewing that particular ad. The authors in these works treated the algorithmic question of winner determination towards maximizing the social welfare. Equilibria of the GSP auction mechanism in the cascade model were studied by Giotis and Karlin in [20]. Kuminov and Tennenholtz [25] study equilibria of the GSP and VCG (cf. [27, Ch. 9]) auctions in a similar model. The cascade model was assessed experimentally by Craswell *et al.* [15]. Gomes, Immorlica and Markakis [21] were the first to document empirically externalities under a version of the cascade model, using real data from keyword auctions.

Externalities among the bidders have also been considered in other similar settings. For instance, Ghosh and Mahdian [19] present a model for externalities in online lead generation, study the complexity of the winner determination problem, and present computationally efficient incentive compatible mechanisms for several cases. Chen and Kempe [12] consider positive and negative social dependencies among the bidders for single item auctions, employing an idea similar to social context graph, and study the properties of equilibria for the first and second price auctions.

Winner Determination Problems. We denote by $\mathrm{MSW\text{-}E}$ the problem of *Maximum Social Welfare with Externalities*, i.e. selecting a subset of winning players S and a permutation π of S that maximize $\mathrm{sw}(S, \pi, G)$ in our model. The complexity and performance of our algorithms are parameterized by the window size c and we write $\mathrm{MSW\text{-}E}(c)$. An interesting special case of $\mathrm{MSW\text{-}E}(c)$ occurs for $E^- = \emptyset$. We refer to this *"positive-only"* externalities version by $\mathrm{MSW\text{-}PE}(c)$. Motivated by the Cascade Model, we also consider the case of *forward-only* positive-only externalities – denoted by $\mathrm{MSW\text{-}FPE}(c)$ – where an ad j may only influence an ad i iff $\pi(j) < \pi(i)$. For deriving our hardness results we consider a simplification of the model where the probability w_{ji} that ad j affects ad i only depends on the window size in π, and not on their ids, i.e. $w_{ji} = w$ for all $(j, i) \in E^+ \cup E^-$.

In presence of negative externalities, it may be profitable to select less than k ads and arrange them in a *broken list*, namely a list with empty slots appearing between slots occupied by negatively correlated ads. In practice however, this may not be feasible. Therefore, we restrict feasible solutions to so-called *unbroken lists*, namely lists where the selected ads occupy consecutive slots in the ad list starting from the first one[4]. Hence a feasible solution can select less than k adds, but it is not allowed to use empty slots and separate negatively correlated ads (i.e. the empty slots, if any, occupy the last $k - |S|$ slots in π). The analysis of our algorithms assumes the case of unbroken lists. However, it is not difficult to generalize our algorithmic results to the case of broken lists.

3 Computational Hardness and the GSP Mechanism

In this section, we restrict our attention to the simplest special case of $\mathrm{MSW\text{-}FPE}(1)$, where there is a single function w, and position multipliers λ_j, valuations v_i, and qualities q_i are uniform. We prove that even in this very restricted setting, $\mathrm{MSW\text{-}FPE}(1)$ is $\mathbf{APX}$-hard. We note that a simple and elegant transformation from the Longest Path problem can be used to prove $\mathbf{NP}$-hardness of $\mathrm{MSW\text{-}FPE}(1)$; this proof is described in the **Appendix**, Section **A.1**. Below we describe a PTAS reduction for $\mathbf{APX}$-hardness.

Theorem 1. $\mathrm{MSW\text{-}FPE}(1)$ *is* $\mathbf{APX}$*-hard even in the special case of uniform position multiplies, valuations, and qualities.*

Proof sketch. The proof is by a PTAS reduction from the Traveling Salesperson Problem with distances 1 and 2, aka $\mathrm{TSP}(1, 2)$, which is known to be $\mathbf{APX}$-hard [28]. An instance of $\mathrm{TSP}(1, 2)$ is defined by an undirected graph $G(V, E)$. Each edge $\{u, u'\} \in E$ has length 1 and each non-edge $\{u, u'\} \notin E$ has length 2. The goal is to find a tour that includes all vertices and has a minimum

[4] A reasonable assumption that may justify the restriction above is that there exist an adequate number of "neutral" ads ($k - 1$ of them suffice) that do not negatively affect any other ad.

total length, or equivalently contains a minimum number of non-edges. $\mathrm{TSP}(1,2)$ is inapproximable within 741/740 unless $\mathbf{P} = \mathbf{NP}$ [18], and approximable within 8/7 [9].

Given an undirected graph $G(V,E)$, $|V| = n$, we construct an instance I of $\mathrm{MSW\text{-}FPE}(1)$ with n advertisers, all of them with valuation 1 and quality $q \in (0,1)$, context graph $G(V,E,\emptyset)$, a single function w with $w(1) = \beta \in (0,1)$ and $w(\ell) = 0$ for all $\ell \neq 1$, n slots available in the sponsored list, and uniform position multipliers equal to 1. In the Appendix, Section A.2, we show that given a $(1+\varepsilon)$-approximate solution to I, we can efficiently construct a $(1+O(\varepsilon))$-approximate solution to the instance of $\mathrm{TSP}(1,2)$ defined by $G(V,E)$. □

4 An Exact Algorithm Based on Color Coding

In this section we develop and analyze an exact algorithm for $\mathrm{MSW\text{-}E}(c)$. We employ *color coding* [3] and dynamic programming, to prove the following result:

Theorem 2. $\mathrm{MSW\text{-}E}(c)$ *can be solved optimally in* $2^{O(k)}\, n^{2c+1} \log^2 n$ *time.*

For simplicity, we assume that the optimal solution consists of k ads. We can remove this assumption by running the algorithm for every sponsored list size up to k, and keep the best solution. Since the running time is exponential in k, this does not change the asymptotics of the running time. To apply the technique of color coding, we consider a fixed coloring $h : N \mapsto [k]$ of the ads with k colors. A list (S,π) of k ads is *colorful* if all ads in S are assigned different colors by h. In the following, we formulate a dynamic programming algorithm that computes the best *colorful* list.

For each $2c$-tuple of ads $(i_1,\dots,i_{2c}) \in N^{2c}$, with all ads assigned different colors by h, and each color set $C \subseteq [k]$, $|C| \le k-2c$, that does not include any of the colors of $i_1,\dots,i_{2c}$, we compute $\mathrm{sw}(i_1,\dots,i_{2c},C)$, namely the maximum social welfare if the last $|C|$ positions in the list are colored according to C, and on top of them, there are ads $i_1,\dots,i_{2c}$ in this order from top to bottom. More precisely, the solution corresponding to $\mathrm{sw}(i_1,\dots,i_{2c},C)$ assigns ad i_p to slot $k-(|C|+2c)+p$, $p = 1,\dots,2c$, and considers the best choice of ads colored according to C for the last $|C|$ slots. Clearly, there are at most $n^{2c}\,2^{k-2c}$ different sw values to compute, and the maximum sw value for all colorful tuples $(i_1,\dots,i_{2c},C)$, with $|C| = k-2c$, corresponds to the best colorful list of k ads. The proof of Theorem 2 follows:

Proof. For the basis of our dynamic programming, let $C = \emptyset$ and for all $2c$-tuples $(i_1,\dots,i_{2c}) \in N^{2c}$ with all ads assigned different colors by h, we have:

$$\begin{aligned}\mathrm{sw}(i_1,\dots,i_{2c},\emptyset) = &\sum_{p=1}^{c} \lambda_{k-2c+p} \cdot Q_{i_p}(i_1,\dots,i_p,\dots,i_{p+c}) \cdot v_{i_p} \\ &+ \sum_{p=c+1}^{2c} \lambda_{k-2c+p} \cdot Q_{i_p}(i_{p-c},\dots,i_p,\dots,i_{2c}) \cdot v_{i_p}\,,\end{aligned}$$

where $Q_{i_p}(i_1,\dots,i_p,\dots,i_{c+p})$ (resp. $Q_{i_p}(i_{p-c},\dots,i_p,\dots,i_{2c})$) is the CTR of ad i_p given that the (only) ads in the list at distance at most c from i_p are $i_1,\dots,i_{c+p}$ (resp. $i_{p-c},\dots,i_{2c}$) arranged in this order from top to bottom.

Given the values of sw for all $2c$-tuples of ads and all color sets of cardinality $s < k - 2c$, we compute the values of sw for all $2c$-tuples $(i_1, \dots, i_{2c})$ and all color sets C of cardinality $s + 1$:

$$\begin{aligned}\mathrm{sw}(i_1, \dots, i_{2c}, C) = \max_{i: h(i) \in C} \{ & \mathrm{sw}(i_2, \dots, i_{2c}, i, C - \{h(i)\}) + \\ + \; & \lambda_{k-(|C|+2c)+1} \cdot Q_{i_1}(i_1, \dots, i_{c+1}) \cdot v_{i_1} + \\ + \; & \sum_{p=2}^{c} \lambda_{k-(|C|+2c)+p} \cdot [Q_{i_p}(i_1, i_2, \dots, i_p, \dots, i_{p+c}) - Q_{i_p}(i_2, \dots, i_p, \dots, i_{p+c})] \cdot v_{i_p} + \\ + \; & \lambda_{k-|C|-c+1} \cdot [Q_{i_{c+1}}(i_1, i_2, \dots, i_{c+1}, \dots, i_{2c}, i) - Q_{i_p}(i_2, \dots, i_{c+1}, \dots, i_{2c}, i)] \cdot v_{i_{c+1}} \}\end{aligned}$$

In the recursion above, the second term accounts for the additional social welfare due to i_1, and the third and the fourth term account for the difference in the social welfare due to ads $i_2, \dots, i_{c+1}$, whose CTRs $Q_{i_2}, \dots, Q_{i_{c+1}}$ are affected by i_1. Ads $i_{c+2}, \dots, i_{2c}, i$ are used to calculate the difference in the CTRs $Q_{i_2}, \dots, Q_{i_{c+1}}$. The CTRs of ads $i_{c+2}, \dots, i_{2c}, i$ and of the ads at the bottom of the list with colors in $C - \{h(i)\}$ are not affected by i_1, since their distance to i_1 is greater than c.

Therefore, for any fixed coloring h, the best colorful list of k ads can be computed in time $O(n^{2c+1}\, 2^k)$. If we select a random coloring h, the probability that the optimal solution is colorful under h is $k!/k^k > e^{-k}$. If we run the algorithm for $e^k \ln n$ randomly and independently chosen colorings and keep the best solution, the probability that we fail to find the optimal solution is at most $1/n$. The approach can be derandomized using a k-perfect family of hash functions of size $2^{O(k)} \log^2 n$ (see [3, Section 4] for the details). □

Hence we obtain the following result for practically interesting sizes of the problem:

Corollary 1. *For $k = O(\log n)$ and $c = O(1)$,* MSW-E(c) *is in* $\mathbf{P}$.

Moreover, unless $\mathbf{NP} \subseteq \mathbf{DTIME}(n^{\mathrm{poly}(\log n)})$, MSW-E($c$) is not $\mathbf{NP}$-hard for $c = O(1)$ and $k = O(\mathrm{poly}(\log n))$. In a mechanism design setting, suppose that advertisers' click valuations are their private information. Having an exact algorithm for MSW-E(c), we can directly use it together with the VCG payments (see e.g. [27, Ch. 9]) to obtain a truthful mechanism for MSW-E. Computation of payments incurs an additional factor $O(k)$ to the running time of the algorithm given by Theorem 2.

5 $O(c)$-Approximation Algorithms for Positive Externalities

In this section, we show how to use polynomial-time approximation algorithms for the Weighted m-Set Packing problem and approximate MSW-PE(c) within a factor of $O(c)$ in polynomial time. In Weighted m-Set Packing, we are given a collection of sets, each with at most m elements and a positive weight, and seek a collection of disjoint sets of maximum total weight. The greedy algorithm for Weighted m-Set Packing achieves an approximation ratio of m, the algorithm of [10] achieves an approximation ratio of $(2/3)m$ in time quadratic in the number of sets, and the algorithm of [8] achieves an approximation ratio of $(m + 1)/2$ in polynomial time for any constant m.

Theorem 3. *Given an α-approximation $T(\nu)$-time algorithm for Weighted 3-Set Packing with ν sets, we obtain a $2\alpha c$-approximation $T(kn^2)$-time algorithm for* MSW-PE(c) *with n ads and k slots.*

Proof. We transform any instance of MSW-PE(c) to an instance of Weighted 3-Set Packing with $kn^2/4$ sets so that any α-approximation to the optimal set packing gives a $2\alpha c$-approximation to the optimal social welfare for the original $\mathrm{MSW\text{-}PE}(c)$ instance. To simplify the presentation, we assume that k is even. Our proof can be easily extended to the case where k is odd.

Given an instance of $\mathrm{MSW\text{-}PE}(c)$ with n ads and k slots, we partition the list into $k/2$ blocks of 2 consecutive slots each. The set packing instance consists of $\binom{n}{2}$ 3-element sets for each block. Namely, for every block $p = 1, 3, 5, \ldots, k-1$ and every subset $\{i_1, i_2\}$ of 2 ads, there is a set $\{i_1, i_2, p\}$ in the set packing instance[5]. The weight $W(i_1, i_2, p)$ of each set $\{i_1, i_2, p\}$ is the maximum social welfare if ads i_1 and i_2 are assigned to slots p and $p+1$, and there is no influence on i_1 and i_2 from any other ad in the list. Formally,

$$W(i_1, i_2, p) = \max\{\lambda_p Q_{i_1}(i_1, i_2) v_{i_1} + \lambda_{p+1} Q_{i_2}(i_1, i_2) v_{i_2}, \lambda_p Q_{i_2}(i_2, i_1) v_{i_2} + \lambda_{p+1} Q_{i_1}(i_2, i_1) v_{i_1}\}$$

where $Q_{i_1}(i_1, i_2) = 1-(1-q_{i_1})(1-q_{i_2} w_{i_2 i_1}(1))$ (resp. $Q_{i_2}(i_1, i_2) = 1-(1-q_{i_2})(1-q_{i_1} w_{i_1 i_2}(-1))$) denotes the relevance of ad i_1 (resp. i_2) given that the only ad in the list with an influence on i_1 (resp. i_2) is i_2 (resp. i_1) located just above (resp. below) i_1 (resp. i_2) in the list.

Given an instance of $\mathrm{MSW\text{-}PE}(c)$ with n advertisers and k slots, the corresponding instance of Weighted 3-Set Packing can be computed in $O(kn^2)$ time. To show that the transformation above is approximation preserving, we prove that **(i)** the optimal set packing has weight at least $1/(2c)$ of the maximum social welfare, and that **(ii)** given a set packing of weight W, we can efficiently compute a solution for the original instance of $\mathrm{MSW\text{-}PE}(c)$ with a social welfare of at least W.

To prove **(i)**, we assume (by renumbering the ads appropriately if needed) that the optimal list for the $\mathrm{MSW\text{-}PE}(c)$ instance is $(1, \ldots, k)$. We let Q_i^* be the relevance of ad i in $(1, \ldots, k)$, and let $W^* = \sum_{i=1}^{k} \lambda_i Q_i^* v_i$ be the optimal social welfare. We construct a collection of $2c$ feasible set packings of total weight at least W^*. Thus, at least one of them has a weight of at least $W^*/(2c)$. The construction is based on the following claim, which can be proven by induction on c (see **Appendix A.3**).

Claim 1 *Let c be any positive integer. Given a list $(1, \ldots, k)$, there is a collection of $2c$ feasible 3-set packings such that for each pair i_1, i_2 of ads in $(1, \ldots, k)$ with $|i_1 - i_2| \le c$, the union of these packings contains a set $\{i_1, i_2, p\}$ with $p \le \min\{i_1, i_2\}$.*

Intuitively, for each pair i_1, i_2 of ads located in $(1, \ldots, k)$ within a distance no more than the window size c, and thus possibly having a positive influence on each other, the collection of set packings constructed in the proof of Claim 1 includes a set $\{i_1, i_2, p\}$ whose weight accounts for the increase in i_1's and i_2's social welfare due to i_2's and i_1's positive influence, respectively. Summing up the weights of all those sets, we account for the positive influence between all pairs of ads in $(1, \ldots, k)$, and thus end up with a total weight of at least W^*.

Formally, let $W^{(j)}$ be the total weight of the j-th set packing constructed in the proof of Claim 1, and let i_1, i_2, with $i_1 < i_2$, be any pair of ads in $(1, \ldots, k)$ included in the same set $\{i_1, i_2, p\}$ of the j-th packing. Since each ad appears in each set packing at most once, we let $Q_{i_1}^{(j)} = Q_{i_1}(i_1, i_2)$ and $Q_{i_2}^{(j)} = Q_{i_2}(i_1, i_2)$ be the relevance of i_1 and i_2 in the calculation of $W(i_1, i_2, p)$. Since Claim 1 ensures that $p \le i_1$, and since slot CTRs are non-increasing, λ_{i_1} (resp. λ_{i_2}) is no greater than λ_p (resp. λ_{p+1}). Therefore, $W(i_1, i_2, p) \ge \lambda_{i_1} Q_{i_1}^{(j)} v_{i_1} + \lambda_{i_2} Q_{i_2}^{(j)} v_{i_2}$. Setting $Q_i^{(j)} = 0$ for all ads i in $(1, \ldots, k)$ which do not appear in any set of the j-th packing, we obtain that $W^{(j)} \ge \sum_{i=1}^{k} \lambda_i Q_i^{(j)} v_i$. We show that

$$\sum_{j=1}^{2c} W^{(j)} \ge \sum_{i=1}^{k} \lambda_i v_i \sum_{j=1}^{2c} Q_i^{(j)} \ge \sum_{j=1}^{k} \lambda_i Q_i^* v_i = W^*$$

[5] Throughout the proof, we implicitly adopt the simplifying (and easily implementable) assumption that the range of block descriptors $1, 3, 5, \ldots, k-1$ and the range of ad descriptors $1, \ldots, n$ are disjoint.

The first inequality follows from the discussion above and by changing the order of the summation. To establish the second inequality, we show that for every ad i in $(1, \ldots, k)$, $\sum_{j=1}^{2c} Q_i^{(j)} \geq Q_i^*$. To simplify the presentation, we focus on an ad i with $c < i \leq k - c$. Ads $1, \ldots, c$ and $k - c + 1, \ldots, k$ can be treated similarly.

We recall that $Q_i^* = 1 - (1 - q_i) \prod_{j=i-c, j \neq i}^{i+c} P_i(j)$, where for each ad j in the sublist $(i - c, \ldots, i - 1, i + 1, \ldots, i + c)$, $P_i(j) = (1 - q_j w_{ji}(j - i)) \in [0, 1]$ accounts for j's positive influence on i's relevance. Claim 1 ensures that for each j in $(i - c, \ldots, i - 1, i + 1, \ldots, i + c)$, ads i and j are included in the same set of some set packing. Since ad i appears in each set packing at most once, for simplicity, we can renumber the set packings of Claim 1, and say that i and j are included in the same set of the j-th set packing. Then, if $j < i$,

$$Q_i^{(j)} = 1 - (1 - q_i)(1 - q_j w_{ji}(-1)) \geq 1 - (1 - q_i) P_i(j)\,,$$

because w_{ji} is non-increasing with the distance of j and i in the list, and thus $w_{ji}(-1) \geq w_{ji}(j - i)$. The same holds if $j > i$. Therefore, $Q_i^{(j)} \geq 1 - (1 - q_i) P_i(j)$, for any j. To conclude the proof of **(i)**, we observe that:

$$\sum_{j=i-c, j \neq i}^{i+c} (1 - (1 - q_i) P_i(j)) \geq 1 - (1 - q_i) \prod_{j=i-c, j \neq i}^{i+c} P_i(j) = Q_i^* \tag{4}$$

To establish (4), we repeatedly apply that for every $x, y, z \in [0, 1]$, $(1 - xy) + (1 - xz) \geq 1 - xyz$.

We proceed to establish claim **(ii)**, namely that given a set packing of weight W, we can efficiently construct a sponsored list of social welfare at least W for the original instance of MSW-PE(c). By construction, we can restrict our attention to set packings of the form $\{\{i_p, i_{p+1}, p\}\}_{p=1,3,\ldots,k-1}$, where the weight of the packing is $W = \sum_p W(i_p, i_{p+1}, p)$, and where ads i_p and i_{p+1} are indexed according to their best order, with respect to which $W(i_p, i_{p+1}, p)$ is calculated. Since we consider positive externalities, the sponsored list $(i_1, i_2, \ldots, i_{k-1}, i_k)$ has a social welfare of at least W. □

Combining Theorem 3 with the greedy 3-approximation algorithm for Weighted 3-Set Packing and with the algorithm of [10] we obtain respectively:

Corollary 2. *For n ads and k slots, the* MSW-PE(c) *problem can be approximated within factor* $6c$ *in $O(kn^2 \log n)$ time and within factor $4c$ in $O(k^2 n^4)$ time.*

A similar approximation preserving reduction from MSW-PE(c) to Weighted $(2c + 1)$-Set Packing yields:

Theorem 4. *An $f(m)$-approximation $T(\nu, m)$-time algorithm for Weighted m-Set Packing with ν sets yields a $2f(2c+1)$-approximation $O(ckn^{2c} + T(kn^{2c}, 2c+1))$-time algorithm for* MSW-PE(c) *with n ads and k slots.*

The full proof appears in the **Appendix**, Section **A.4**; the list is partitioned into $k/(2c)$ blocks of $2c$ consecutive slots each, and the set packing instance consists of $\binom{n}{2c}$ sets of $2c + 1$ elements each for each block. Then we exhibit a pair of feasible set packings whose total weight is no less than the optimal social welfare. Combining Theorem 4 with the approximation algorithm of [8], we obtain a polynomial time $2(c + 1)$-approximation algorithm MSW-PE(c), for any constant c.

6 On the GSP Mechanism with Externalities

In examining the behavior of the GSP auction mechanism we make the reasonable assumption that only a snapshot of the players' intrinsic relevances q_i, $i \in \mathcal{N}$ is available to the mechanism. In practice, estimates of the players' relevances are deduced by software machinery of the sponsored search platform; it is therefore conceivable that the mechanism will eventually extract information indicative of externalities. By the time this occurs however, associations among advertisers may also be updated. This justifies the mechanism's unawareness of externalities. On the other hand, each advertiser is aware of the social context associations that may harm him or boost his relevance. Given a social context graph $G(\mathcal{N}, E^+, E^-)$, with functions w_{ji} for every $(j,i) \in E^+ \cup E^-$ and window size c, we assume that every advertiser $i \in \mathcal{N}$ is aware of: $E_i^+ = \{(i', i) \in E^+ | i' \in \mathcal{N}\}$, $E_i^- = \{(i', i) \in E^- | i' \in \mathcal{N}\}$ and of c and w_{ji} for every $j \in E_i^+ \cup E_i^-$.

In a keyword auction each advertiser i bids b_i for receiving a slot in the list. The GSP mechanism in its most common flavors uses the *Rank-By-Revenue* (RBR) rule to assign advertisers to slots. Under RBR, advertisers are ranked in order of non-increasing score $q_i \cdot b_i$ and higher scores are assigned higher CTR slots. The score of a bidder is his *declared expected revenue* for a click. The plainer RBB rule is obtained by taking $q_i = 1$ for all $i \in \mathcal{N}$. Given a bid vector $\boldsymbol{b}$, let $\phi_{\boldsymbol{b}} : \mathcal{K} \to \mathcal{N}$ denote the ranking of bidders in order of non-increasing expected revenue, i.e. $\phi_{\boldsymbol{b}}(j)$ is the bidder assigned to slot j. According to the previously used definition of π, we use ϕ as an extension of π^{-1}. Every slot winning player $\phi_{\boldsymbol{b}}(j)$, for $j = 1, \ldots, k$ pays *per click* a price equal to $(q_{\phi_{\boldsymbol{b}}(j+1)} \cdot b_{\phi_{\boldsymbol{b}}(j+1)})/q_{\phi_{\boldsymbol{b}}(j)}$; i.e., the score of the bidder occupying the next position under $\boldsymbol{b}$, divided by the intrinsic relevance of $\phi_{\boldsymbol{b}}(j)$. Using his knowledge of externalities that influence him, player $\phi_{\boldsymbol{b}}(j)$ experiences a relevance $Q_{\phi_{\boldsymbol{b}}(j)}(\boldsymbol{b})$ and estimates his expected profit (utility) as:

$$u_{\phi_{\boldsymbol{b}}(j)}(\boldsymbol{b}) = \lambda_j \cdot Q_{\phi_{\boldsymbol{b}}(j)}(\boldsymbol{b}) \times \left(v_{\phi_{\boldsymbol{b}}(j)} - \frac{q_{\phi_{\boldsymbol{b}}(j+1)}}{q_{\phi_{\boldsymbol{b}}(j)}} \cdot b_{\phi_{\boldsymbol{b}}(j+1)} \right) \tag{5}$$

We assume a complete information setting, as advertisers typically employ machine learning techniques to estimate how much they should outbid a competitor. Such techniques reveal the ranking information used by the GSP mechanism. Thus, in computing his best response under a bid vector $\boldsymbol{b}_{-i}$, every advertiser $i \in \mathcal{N}$ is assumed to know only $q_{i'}$ and $b_{i'}$ for each $i' \in \mathcal{N} \setminus \{i\}$, and *not* the actual relevance $Q_{i'}$ perceived by i'.

In studying pure Nash equilibria of the GSP mechanism we make the standard assumption of *conservative* (or *ex-post individually rational*) bidders. It states that in fear of receiving negative utility, no bidder ever outbids his valuation, i.e. $b_i \leq v_i$ holds for all $i \in \mathcal{N}$. This allows to prove meaningful bounds for the social welfare relative to the social optimum [26, 13] since, otherwise, the worst-case social welfare of equilibria may be arbitrarily low. For conservative bidders and under our definitions of the mechanism's and the bidders' awareness of externalities we show:

Proposition 1. *The strategic game induced by the Generalized Second Price Auction mechanism under the RBR rule and deterministic tie-breaking does not generally have pure Nash equilibria in presence of forward positive externalities, even for* 3 *conservative players and* 2 *slots.*

Proof. Consider 3 players, and 2 slots. The tie-breaking rule picks player 3 first if he issues the same bid with any of the other two players. Among the other 2 players ties may be resolved in any arbitrarily chosen fixed way. All players have intrinsic relevance q and valuations $v_1 = v_2 = V > \frac{2-q}{1-q} \cdot v$, where $v = v_3$. The slots have CTRs $\lambda_1 > \lambda_2$, such that $\gamma = \lambda_2/\lambda_1 > V/[(2-q) \cdot (V - v)]$. Notice that $V = (2-q) \cdot V - (1-q) \cdot V < (2-q) \cdot (V - v)$, because $V > \frac{2-q}{1-q} \cdot v$. Thus γ can be feasibly chosen within a non-empty range $(V/[(2-q) \cdot (V - v)], 1]$.

Set the window size to be $c = 1$ and consider a directed cycle $\{(1,2),(2,1)\}$ as the context graph. Let $w_{ji}(1) = 1$. When both players 1,2 are slot winners, if player i is ranked below $i' \neq i$, his relevance is amplified to $Q = q(2-q)$. We claim that none of the orderings $\langle 1,2 \rangle$ and $\langle 2,1 \rangle$ has a bid vector that is a pure Nash equilibrium. For every bid vector $\boldsymbol{b}$ that makes player 3 a looser (without ties), we show that player $\phi_{\boldsymbol{b}}(1)$ has incentive to aim for slot 2. $\boldsymbol{b}'$ will denote the bid vector after such a deviation of player $\phi_{\boldsymbol{b}}(1)$, given $\boldsymbol{b}_{-\phi_{\boldsymbol{b}}(1)}$. Then:

$$\begin{aligned}
u_{(1)}(\boldsymbol{b}) &= \lambda_1 \cdot q \cdot (v_{\phi_{\boldsymbol{b}}(1)} - b_{\phi_{\boldsymbol{b}}(2)}) \\
&= \gamma^{-1} \cdot \lambda_2 \cdot q \cdot (v_{\phi_{\boldsymbol{b}}(1)} - b_{\phi_{\boldsymbol{b}}(2)}) \\
&< \gamma^{-1} \cdot \lambda_2 \cdot q \cdot V \\
&< \tfrac{(2-q)\cdot(V-v)}{V} \lambda_2 \cdot q \cdot V \\
&= \lambda_2 \cdot Q \cdot (V - v) \\
&\leq \lambda_2 \cdot Q \cdot (v_{\phi_{\boldsymbol{b}}(1)} - b_3) = u_{\phi_{\boldsymbol{b}}(1)}(\boldsymbol{b}')
\end{aligned}$$

When at least one of players 1, 2 bids equally to b_3, at least one of 1,2 looses (i.e. does not win a slot), because of the tie-breaking rule. Under such a configuration the loosing player has incentive (and the capacity - because $V > v \geq b_3$) to outbid b_3 and get a slot. This proves that for particular deterministic tie-breaking rules, conservative pure Nash equilibria do not exist for the GSP under the RBR rule. □

The proof of the following proposition can be found in **Appendix A.5**.

Proposition 2. *The strategic game induced by the Generalized Second Price Auction under the RBR rule and deterministic tie-breaking does not generally have pure Nash equilibria in presence of bidirectional positive externalities, even for* 3 *conservative players and* 3 *slots.*

In proving these propositions we used equal intrinsic relevance factors among players. Thus, even the relevance-independent RBB rule cannot yield stability against the players' perceptions about their expected payoff. Even when pure Nash equilibria exist, externalities may cause the optimum social welfare to be arbitrarily larger than the *best* welfare achievable in equilibrium, i.e. an unbounded *Price of Stability* [4]:

Proposition 3. *There is an infinite family of instances of the stategic game induced by the Generalized Second Price auction mechanism with unbounded Price of Stability, even with conservative bidders.*

Proof. Consider first an instance with 2 slots of equal CTRs, $\lambda_1 = \lambda_2 = 1$, and 5 players. Take $q_1 = q_2 = q_5 = \delta$, $v_1 = v_2 = v_5 = 1$, and $q_3 = q_4 = 1$, $v_3 = v_4 = \delta'$ such that δ' is always very much smaller than δ. Let the social context graph have edge set $E^+ = \{(3,1),(4,2)\}$ and $E^- = \emptyset$. Set the window size to $c = 1$ and let $w(1) = 1$. The bid vector $\boldsymbol{b}$ where $b_i = v_i$ is a pure Nash equilibrium: under any tie-breaking rule, two players out of 1, 2, 5 receive a slot and have zero utility. Then, $\mathrm{sw}(\boldsymbol{b}) = 2\delta$. For the social optimum, assign the two slots to any of the pairs $\{3,1\}$ or $\{4,2\}$. The optimum social welfare is then $OPT = 1 + \delta'$, in any of these pairs one player has his relevance amplified to 1, due to positive influence by the other. $OPT/\mathrm{sw}(\boldsymbol{b}) = (1+\delta')/(2\delta)$, which becomes arbitrarily large as $\delta \to 0$. Notice that, because players are conservative, any bid configuration in which a player $i \in \{3,4\}$ receives a slot, cannot be an equilibrium. This is because at least two of players 1,2,5 would not receive a slot and any of them can deviate by outbidding i. Thus all equilibria of the game have equal social welfare. One can straightforwardly generalize this instance for any number of slots k, using $2k+1$ players. □

This result demonstrates how externalities may harm the performance of the GSP mechanism arbitrarily, in contrast to the recent upper bound of 1.282 shown recently [11] on the mechanism's *Price of Anarchy* [24] (the ratio of the optimum welfare over the *worst* equilibrium welfare), without externalities. Giotis and Karlin showed in [20] a lower bound of k in the Cascade Model of externalities with conservative bidders.

7 Conclusion and Further Directions

We presented a comprehensive list of computational results for alleviating the impact of externalities in sponsored search, under a novel fairly expressive model. We gave an exact algorithm for optimizing the social welfare under positive and/or negative externalities, which can be paired with VCG payments, to yield a truthful polynomial time mechanism for practically relevant input sizes. We settled $\mathbf{APX}$-completeness of the associated Winner Determinaton problem with positive-only externalities and constant-size window of their scope, by proving hardness and designing constant factor approximation algorithms. These algorithms do not yield truthful mechanisms in dominant strategies, because they are not *monotone* [6]. The dependence of the advertisers' CTRs depend on the CTRs of others caused even simple approximation algorithms that we tried not to be monotone. On the other hand we showed that the non-truthful GSP mechanism may not have pure Nash equilibria in presence of externalities; when it does, these may have unboundedly low social welfare. As our negative results rely on the generality of our model, many directions emerge regarding the design and performance analysis of mechanisms (including the GSP), for special practically relevant cases; these include e.g. simple cases of social context graphs and bounded influence among advertisers.

Acknowledgments: Dimitris Fotakis wishes to thank Vasileios Syrgkanis for many helpful discussions on the model of externalities among advertisers presented in this work. We would also like to thank Moshe Tennenholtz for many interesting discussions and suggestions concerning this project.

References

1. G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pal. Sponsored Search Auctions with Markovian Users. In C. H. Papadimitriou and S. Zhang, editors, *WINE 2008*, volume 5385 of *LNCS*, pages 621–628. Springer, Heidelberg, 2008.
2. G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proc. of 7th ACM Conference on Electronic Commerce (EC)*, pages 1–7. ACM, 2006.
3. N. Alon, R. Yuster, and U. Zwick. Color Coding. *Journal of the ACM*, 42:844–856, 1995.
4. E. Anshelevich, A. Dasgupta, J. M. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
5. A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 482–491. IEEE, 2001.
6. I. Ashlagi, P. Krysta, and M. Tennenholtz. Social Context Games. In C. H. Papadimitriou and S. Zhang, editors, *WINE 2008*, volume 5385 of *LNCS*, pages 675–683, 2008.
7. S. Athey and G. Ellison. Position Auctions with Consumer Search. *Quartetly Journal of Economis*, page (to appear), 2011.
8. P. Berman. A $d/2$ Approximation for Maximum Weight Independent Set in d-Claw Free Graphs. In M. M. Halldórsson, editor, *SWAT 2000*, volume 1851 of *LNCS*, pages 214–219, 2000.
9. P. Berman and M. Karpinski. $8/7$-approximation algorithm for $(1, 2)$-TSP. In *Proc. of the 17th Annual ACM Symposium on Discrete Algorithms (SODA)*, pages 641–648. ACM, 2006.
10. P. Berman and P. Krysta. Optimizing misdirection. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 192–201, 2003.

[6] A certain monotonicity property is required for an algorithm to imply a truthful mechanism in single parameter settings as is ours, see, e.g., [5].

11. I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On the Efficiency of Equilibria in Generalized Second Price Auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC), to appear*. ACM, 2011.
12. P. Chen and D. Kempe. Bayesian Auctions with Friends and Foes. In M. Mavronicolas and V. G. Papadopoulou, editors, *SAGT 2009*, volume 5814 of *LNCS*, pages 335–346. Springer, Heidelberg, 2009.
13. G. Christodoulou, A. Kovács, and M. Schapira. Bayesian Combinatorial Auctions. In L. Aceto, I. Damgaard, L. A. Goldberg, M. M. Halldórsson, A. Ingólfsdóttir, and I. Walukiewicz, editors, *ICALP 2008 (1)*, volume 5125 of *LNCS*, pages 820–832. Springer, Heidelberg, 2008.
14. P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. MIT Press, 2006.
15. N. Craswell, O. Zoeter, M. Taylor, and B. Ramsey. An experimental comparison of click position-bias models. In *WSDM '08: Proceedings of the international conference on Web search and web data mining*, pages 87–94. ACM, 2008.
16. A. Das, I. Giotis, A.R. Karlin, and C. Mathieu. On the Effects of Competing Advertisements in Keyword Auctions. In *Unpublished manuscript*, 2008.
17. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
18. L. Engebretsen and M. Karpinski. TSP with bounded metrics. *Journal of Comp. and Sys. Sciences*, 72(4):509–546, 2006.
19. A. Ghosh and M. Mahdian. Externalities in Online Advertising. In *Proc. of the 17th International Conference on World Wide Web Conference (WWW)*, pages 161–168. ACM, 2008.
20. I. Giotis and A. R. Karlin. On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. In C. H. Papadimitriou and S. Zhang, editors, *WINE 2008*, volume 5385 of *LNCS*, pages 629–638. Springer, Heidelberg, 2008.
21. R. Gomes, N. Immorlica, and E. Markakis. Externalities in Keyword Auctions: an empirical and theoretical assessment. In S. Leonardi, editor, *WINE 2009*, LNCS, pages 172–183. Springer, Heidelberg, 2009.
22. P. Jeziorski and I. Segal. What makes them click: Empirical analysis of consumer demand for search advertising. *Working paper*, Department of Economics, Stanford University, 2009.
23. D. Kempe and M. Mahdian. A Cascade Model for Externalities in Sponsored Search. In C. H. Papadimitriou and S. Zhang, editors, *WINE 2008*, volume 5385 of *LNCS*, pages 585–596. Springer, Heidelberg, 2008.
24. E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
25. D. Kuminov and M. Tennenholtz. User Modeling in Position Auctions: Re-Considering the GSP and VCG Mechanisms. In *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 273–280. ACM, 2009.
26. R. Paes Leme and E. Tardos. Pure and Bayes-Nash Price of Anarchy for Generalized Second Price Auction. In *Proc. of the IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 735–744. IEEE, 2010.
27. N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
28. C.H. Papadimitriou and M. Yannakakis. The Travelling Salesman Problem with Distances One and Two. *Mathematics of Operations Research*, 18(1):1–11, 1993.
29. H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.
30. W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37, 1961.

A Appendix

A.1 NP-hardness of $\text{MSW-FPE}(1)$

Proposition 4. $\text{MSW-FPE}(1)$ *is* $\mathbf{NP}$*-hard even in the special case of uniform position multipliers, valuations, and qualities.*

Proof. We use transformation from the Longest Path problem. Given a directed graph $G(V, E)$ and an integer $k \geq 2$, we construct an instance I of $\text{MSW-FPE}(1)$. In I, there are $|V|$ advertisers, all of them with valuation 1 and quality $1/2$, there are k slots in the sponsored list, all position multipliers are equal 1, the advertisers' context graph is $G(V, E, \emptyset)$, and there is a single externality function with $w(1) = 1/2$ and $w(\ell) = 0$ for all $\ell \neq 1$. The ad at slot 1 contributes $1/2$ to the social welfare. The ad at each slot ℓ, $\ell \in \{2, \ldots, k\}$, contributes $5/8$ to the social welfare if the ad on slot $\ell - 1$ is correlated to it, and $1/2$ otherwise. It is immediate that G has a simple path of length $k - 1$ iff I admits a solution of social welfare at least $(5k - 1)/8$. □

A.2 The Proof of Theorem 1

We show that for any constant $\varepsilon > 0$, an $(1 + \varepsilon)$-approximation algorithm for $\text{MSW-FPE}(1)$ implies a $(1 + O(\varepsilon))$-approximation algorithm for $\text{TSP}(1, 2)$.

Given an instance $G(V, E)$, $|V| = n$, of $\text{TSP}(1, 2)$, we construct an instance I of $\text{MSW-FPE}(1)$. In I, there are n advertisers, all of them with valuation 1 and quality $q \in (0, 1)$, the context graph is $G(V, E, \emptyset)$, there is a single function w with $w(1) = \beta \in (0, 1)$ and $w(\ell) = 0$ for all $\ell \neq 1$, the number of slots is n (i.e. all ads fit in the sponsored list), and all position multipliers are equal to 1.

A solution to I is a permutation π of the n advertisers corresponding to the vertices of G. The ad at slot 1 contributes q to the social welfare. The ad at each slot ℓ, $\ell \in \{2, \ldots, n\}$, contributes $q(1 + \beta(1 - q))$ to the social welfare if the ad at slot $\ell - 1$ is correlated to it (i.e. there is an edge in G between $\pi^{-1}(\ell - 1)$ and $\pi^{-1}(\ell)$), and q otherwise. To simplify the notation, we let $\zeta = \beta(1 - q)$. Therefore any solution to I that contains μ non-edges of G, i.e. assigns μ pairs of ads not connected by an edge in G to consecutive slots on the ad list, has a social welfare of $q(n(1 + \zeta) - (\mu + 1)\zeta)$.

Let $t^* = (u_1, \ldots, u_n)$ be an optimal tour for the $\text{TSP}(1, 2)$ instance, and let α be the number of non-edges in t^*, and $n + \alpha$ be t^*'s total length. For simplicity, we assume that if $\alpha > 0$, $\{u_n, u_1\} \notin E$. The solution to I that arranges the ads in the same order as in t^*, namely $(u_1, \ldots, u_n)$, has a social welfare of at least $q(n(1 + \zeta) - \zeta)$ if $\alpha = 0$, and $q(n(1 + \zeta) - \alpha\zeta)$ otherwise.

For any constant $\varepsilon > 0$, we consider a $(1 + \varepsilon)$-approximate solution π to I. Let π contain α' non-edges. If $\alpha > 0$, the social welfare of π is:

$$q(n(1 + \zeta) - (\alpha' + 1)\zeta) \geq \frac{q(n(1 + \zeta) - \alpha\zeta)}{1 + \varepsilon},$$

which implies that

$$\alpha' + 1 \leq \tfrac{\varepsilon}{1+\varepsilon} \tfrac{\zeta+1}{\zeta} n + \tfrac{1}{1+\varepsilon} \alpha$$

If $\alpha = 0$, we work similarly and obtain that

$$\alpha' + 1 \leq \tfrac{\varepsilon}{1+\varepsilon} \tfrac{\zeta+1}{\zeta} n + \tfrac{1}{1+\varepsilon}$$

Arranging the vertices of G according to π gives a tour with at most $\alpha' + 1$ non-edges and a length no greater than $n + \alpha' + 1$. If $\alpha > 0$, we obtain that

$$n + \alpha' + 1 \leq (n + \alpha) \left(1 + \varepsilon \tfrac{\zeta+1}{\zeta}\right)$$

If $\alpha = 0$, we obtain that

$$n + \alpha' + 1 \leq n \left(1 + \tfrac{\varepsilon}{1+\varepsilon}\,\tfrac{\zeta+1}{\zeta} + \tfrac{1}{n(1+\varepsilon)}\right) \leq n \left(1 + \varepsilon\,\frac{\zeta+1}{\zeta}\right) ,$$

where the last inequality holds for all $n \geq \frac{\zeta}{\varepsilon^2(1+\zeta)}$.

Therefore, for any constant $\varepsilon > 0$, any $(1+\varepsilon)$-approximate solution to $\mathrm{MSW\text{-}FPE}(1)$ instance I can be translated into a $(1+\varepsilon\frac{\zeta+1}{\zeta})$-approximate solution to the original instance of $\mathrm{TSP}(1,2)$. We recall that $\zeta = \beta(1-q)$. Thus $\frac{\zeta+1}{\zeta}$ can be arbitrarily close to 2 if q is selected sufficiently close to 0 and β is selected sufficiently close to 1. □

A.3 The Proof of Claim 1 in Theorem 3

The proof is by induction on c. For $c = 1$, the first set packing is $\{\{p, p+1, p\}\}_{p=1,3,5,\ldots,k-1}$ and the second set packing is $\{\{p+1, p+2, p\}\}_{p=1,3,5,\ldots,k-3}$ [7]. It is straightforward to verify that for each ad i in $(1, \ldots, k-1)$, there is a set $\{i, i+1, p\}$ with $p = i-1$ if i is even, and with $p = i$ if i is odd (see also the pair of rows corresponding to $c = 1$ in Fig. 1).

For some integer $c \geq 2$, we inductively assume that the claim holds for a window size up to $c-1$, and we show that the claim holds for a window size of c. Thus, by inductive hypothesis, we already have a collection of $2(c-1)$ feasible set packings such that for each pair i_1, i_2 of ads with $|i_1 - i_2| \leq c-1$, the union of these packings contains a set $\{i_1, i_2, p\}$ with $p \leq \min\{i_1, i_2\}$ (see e.g. the eight rows corresponding to $c = 1, 2, 3, 4$ in Fig. 1). To complete the construction, we add two new set packings such that for each ad $i \leq k - c$, they include a set $\{i, i+c, p\}$, with $p \leq i$.

If c is odd, the first new set packing is $\{\{p, p+c, p\}\}_{p=1,3,5,\ldots,k-c}$ and the second new set packing is $\{\{p+1, p+c+1, p\}\}_{p=1,3,5,\ldots,k-c-2}$ (see the pair of rows corresponding to $c = 5$ in Fig. 1). Both set packings are feasible. In particular, none of them contains the same ad twice, since each set contains an even and an odd ad, and the ad ids increase by 2 from each set to the next. Moreover, for each ad $i \leq k-c$, there is a set $\{i, i+c, p\}$ with $p = i$ if i is odd, and with $p = i-1$ if i is even.

If c is even, we use a more complicated periodic construction, which is required for the feasibility of the two set packings. The first new set packing is:

$$\{\{2c\ell + p, 2c\ell + p + c, 2c\ell + p\}\}_{p=1,3,5,\ldots,c-1,\ell=0,1,\ldots,\lceil(k-3c)/(2c)\rceil} \bigcup$$
$$\{\{2c\ell + p + c + 1, 2c\ell + p + 2c + 1, 2c\ell + p + c\}\}_{p=1,3,5,\ldots,c-1,\ell=0,1,\ldots,\lceil(k-3c)/(2c)\rceil}$$

The second new set packing is:

$$\{\{2c\ell + p + 1, 2c\ell + p + c + 1, 2c\ell + p\}\}_{p=1,3,5,\ldots,c-1,\ell=0,1,\ldots,\lceil(k-3c)/(2c)\rceil} \bigcup$$
$$\{\{2c\ell + p + c, 2c\ell + p + 2c, 2c\ell + p + c\}\}_{p=1,3,5,\ldots,c-1,\ell=0,1,\ldots,\lceil(k-3c)/(2c)\rceil}$$

with the understanding that we exclude any sets with an ad of id greater than k (see the pair of rows corresponding to $c = 6$ in Fig. 1). In the first (resp. the second) set packing, the sets in the first group include each ad with odd (resp. even) id once, and the sets in the second group include each ad with even (resp. odd) id at most once. Thus we ensure that both set packings are feasible. In both cases, the

[7] We recall that the third element of each set is the so-called block descriptor, and refers to the list positions according to which the weight of the set is calculated in the proof of Theorem 3. Even though we conveniently identify both ads and block descriptors with positive integers, we highlight that ads and block descriptors are different entities. Thus, e.g. $\{p, p+1, p\}$ should be regarded as a set with three different elements, and $\{p, p+1, p-1\}$ and $\{p+2, p+3, p+1\}$ should be regarded as disjoint sets, because the contain different ads and have different block descriptors.

$p=1$	$p=3$	$p=5$	$p=7$	$p=9$	$p=11$	$p=13$	$p=15$	$p=17$	$p=19$	$p=21$	$p=23$	
1 2	3 4	5 6	7 8	9 10	11 12	13 14	15 16	17 18	19 20	21 22	23 24	$c=1$
2 3	4 5	6 7	8 9	10 11	12 13	14 15	16 17	18 19	20 21	22 23		
1 3	4 6	5 7	8 10	9 11	12 14	13 15	16 18	17 19	20 22	21 23	$c=2$	
2 4	3 5	6 8	7 9	10 12	11 13	14 16	15 17	18 20	19 21	22 24		
1 4	3 6	5 8	7 10	9 12	11 14	13 16	15 18	17 20	19 22	21 24	$c=3$	
2 5	4 7	6 9	8 11	10 13	12 15	14 17	16 19	18 21	20 23			
1 5	3 7	6 10	8 12	9 13	11 15	14 18	16 20	17 21	19 23	$c=4$		
2 6	4 8	5 9	7 11	10 14	12 16	13 17	15 19	18 22	20 24			
1 6	3 8	5 10	7 12	9 14	11 16	13 18	15 20	17 22	19 24	$c=5$		
2 7	4 9	6 11	8 13	10 15	12 17	14 19	16 21	18 23				
1 7	3 9	5 11	8 14	10 16	12 18	13 19	15 21	17 23	$c=6$			
2 8	4 10	6 12	7 13	9 15	11 17	14 20	16 22	18 24				

Fig. 1. The set packings constructed in the proof of Claim 1 for $k = 24$ and $c = 1, \dots, 6$. For each value of c, we have two rows, one for each of the additional set packings in the inductive step. Each cell containing two integers i_1, i_2 corresponds to a set with ads i_1 and i_2. The block descriptor of each set is given by the value of p in the corresponding column. Cells / sets in the same row belong to the same set packing, and thus each add appears in each row at most once, while cells /set in the same column have the same block descriptor. For even values of c, we mark, by switching from grey to white color, the groups of sets obtained for each different value of ℓ.

ads are paired so that the ids of two ads in the same set differ by c. Moreover, the first group of the first (resp. second) packing and the second group of the second (resp. first) packing are complimentary, in the sense their union includes all sets $\{i_1, i_2, p\}$ where i_1 and i_2 are odd (resp. even) and have $|i_1 - i_2| = c$. Therefore, for each ad $i \le k - c$, there is a set $\{i, i + c, p\}$ with $p = i$ if i is odd, and with $p = i - 1$ if i is even.

Hence, for every window size c and any list $(1, \dots, k)$, we can construct a collection of $2c$ feasible 3-set packings such that for each pair i_1, i_2 of ads with $|i_1 - i_2| \le c$, there is a set $\{i_1, i_2, \ell\}$ with $\ell \le \min\{i_1, i_2\}$. □

A.4 The Proof of Theorem 4

We transform any instance of $\mathrm{MSW\text{-}PE}(c)$ to an instance of Weighted $(2c + 1)$-Set Packing with at most kn^{2c} sets so that any α-approximation to the optimal set packing gives a 2α-approximation to the optimal solution of the $\mathrm{MSW\text{-}PE}(c)$ instance. To simplify the presentation, we assume that the optimal solution to the original instance consists of exactly k ads, and that k is a multiple of $2c$. Our proof can be easily extended to handle the remaining cases as well.

Given an instance of $\mathrm{MSW\text{-}PE}(c)$ with n advertisers and k slots, we partition the list into $k/(2c)$ blocks of $2c$ consecutive slots each. The set packing instance consists of $\binom{n}{2c}$ sets of $2c + 1$ elements each for each block. Namely, for every $p = 1, 2c+1, 4c+1, \dots, k-2c+1$, and every set $\{i_1, \dots, i_{2c}\}$ of $2c$ ads, there is a set $\{i_1, \dots, i_{2c}, p\}$ in the set packing instance[8]. Thus, we create at most kn^{2c} sets.

To compute the weight of each set $\{i_1, \dots, i_{2c}, p\}$, we consider all possible permutations π of ads $i_1, \dots, i_{2c}$. For each permutation π, let $W(\{i_1, \dots, i_{2c}, p\}, \pi)$ denote the social welfare if ads

[8] As in the proof of Theorem 3, we implicitly adopt the simplifying assumption that the range of block descriptors $1, 2c + 1, 4c + 1, \dots, k - 2c + 1$ and the range of ad descriptors $1, \dots, n$ are disjoint.

$i_1, \ldots, i_{2c}$ alone are assigned to the slots $p, p+1, \ldots, p+2c-1$ according to π. Formally,

$$W(\{i_1, \ldots, i_{2c}, p\}, \pi) = \sum_{j=1}^{c} \lambda_{p+j-1} \cdot Q_{\pi^{-1}(j)}(\pi^{-1}(1), \ldots, \pi^{-1}(j), \ldots, \pi^{-1}(j+c)) \cdot v_{\pi^{-1}(j)} +$$

$$+ \sum_{j=c+1}^{2c} \lambda_{p+j-1} \cdot Q_{\pi^{-1}(j)}(\pi^{-1}(j-c), \ldots, \pi^{-1}(j), \ldots, \pi^{-1}(2c)) \cdot v_{\pi^{-1}(j)} \quad (6)$$

where $\pi^{-1}(j)$ is the ad assigned to slot $p+j-1$ by π, and $Q_{\pi^{-1}(j)}(\pi^{-1}(1), \ldots, \pi^{-1}(j), \ldots, \pi^{-1}(j+c))$ (resp. $Q_{\pi^{-1}(j)}(\pi^{-1}(j-c), \ldots, \pi^{-1}(j), \ldots, \pi^{-1}(2c))$) denotes the relevance of ad $\pi^{-1}(j)$ given that the only ads in the list at distance at most c from $\pi^{-1}(j)$ are $\pi^{-1}(1), \ldots, \pi^{-1}(j), \ldots, \pi^{-1}(j+c)$ (resp. $\pi^{-1}(j-c), \ldots, \pi^{-1}(j), \ldots, \pi^{-1}(2c)$) arranged in this order from top to bottom.

The weight $W(\{i_1, \ldots, i_{2c}, p\})$ of each set $\{i_1, \ldots, i_{2c}, p\}$ is the maximum weight for all permutations π of $i_1, \ldots, i_{2c}$. Since the weight of each set can be calculated in $O((2c)!c^2)$ time, the Weighted $(2c+1)$-Set Packing instance can be computed in $O(ckn^{2c})$ time. To show that the transformation above is approximation preserving, we prove that (i) the optimal set packing has weight at least half of the maximum social welfare, and that (ii) given a set packing of weight W, we can efficiently compute a solution to the instance of MSW-PE(c) of social welfare at least W.

To prove (i), we let $(1, \ldots, k)$ be the optimal list for the MSW-PE(c) instance, let Q_i^* be the relevance of ad i in $(1, \ldots, k)$, and let $W^* = \sum_{i=1}^{k} \lambda_i Q_i^* v_i$ be the optimal social welfare. Next, we construct a pair of feasible packings of total weight at least W^*.

For the first set packing, we partition $(1, \ldots, k)$ into $k/(2c)$ sets each containing $2c$ ads in consecutive slots and the corresponding block descriptor. In particular, the first set packing contains a set $s_p^{(1)} = \{p, p+1, \ldots, p+2c-1, p\}$ for each block descriptor $p = 1, 2c+1, 4c+1, \ldots, k-2c+1$. Let $\pi_p^{(1)}$ be the permutation that assigns the ads in $s_p^{(1)}$ to their slots in $(1, \ldots, k)$, and for each $i \in s_p^{(1)}$, let $Q_i^{(1)}$ be the relevance with which ad i contributes to $W(s_p^{(1)}, \pi_p^{(1)})$ in (6). Since $W(s_p^{(1)})$ is the maximum weight of $s_p^{(1)}$ over all permutations of its ads, the total weight of the first packing is $W^{(1)} \geq \sum_{j=1}^{j=k} \lambda_j \cdot Q_i^{(1)} \cdot v_i$.

For the second packing, we partition $(c+1, \ldots, k, 1, \ldots, c)$ into $k/(2c)$ sets similarly. In particular, there is a set $s_p^{(2)}$ with the ads in slots $p, p+1 \ldots, p-2c+1$ in $(c+1, \ldots, k, 1, \ldots, c)$ for each block descriptor $p = 1, 2c+1, 4c+1, \ldots, k-2c+1$. As before, we let $\pi_p^{(2)}$ be the permutation that assigns the ads in $s_p^{(2)}$ to their slots in $(c+1, \ldots, k, 1, \ldots, c)$, and for each $i \in s_p^{(2)}$, we let $Q_i^{(2)}$ be the relevance with which ad i contributes to $W(s_p^{(2)}, \pi_p^{(2)})$ in (6). Since $W(s_p^{(2)})$ is the maximum weight of $s_p^{(2)}$ over all permutations, and since slot CTRs are non-increasing, the total weight of the second packing is $W^{(2)} \geq \sum_{j=c+1}^{j=k} \lambda_j \cdot Q_i^{(2)} \cdot v_i$.

We claim that:

$$W^{(1)} + W^{(2)} \geq \sum_{j=1}^{j=c} \lambda_j \cdot Q_{i_j}^{(1)} \cdot v_{i_j} + \sum_{j=c+1}^{j=k} \lambda_j \cdot (Q_{i_j}^{(1)} \cdot v_{i_j} + Q_{i_j}^{(2)}) \geq W^*$$

The first inequality follows from the discussion above. To establish the second inequality, we first observe that for all $i = 1, \ldots, c$, $Q_i^* = Q_i^{(1)}$, since all ads that affect the relevance of i in $(1, \ldots, k)$ are included in $s_1^{(1)}$, and thus are taken into account in the calculation of $Q_i^{(1)}$.

For the remaining ads $i = c+1, \ldots, k$, we work as in the proof of Theorem 3 and show that $Q_i^{(1)} + Q_i^{(2)} \geq Q_i^*$. We first consider an ad i assigned to the first c slots of block $s_p^{(1)}$ by $\pi_p^{(1)}$, for some $p = 2c+1, 4c+1, \dots, k-2c+1$. Then ad i is assigned to the last c slots of block $s^{(2)}_{p-2c}$ by $\pi^{(2)}_{p-2c}$. Therefore, $Q^{(1)}_i$ takes into account externalities due to ads $i+1, \dots, i+c$ after i and $Q^{(2)}_i$ takes into account externalities due to ads $i-c, \dots, i-1$ before i in $(1, \dots, k)$. Since we consider positive externalities, $Q^{(1)}_i \geq Q_i(i, i+1, \dots, i+c)$ and $Q^{(2)}_i \geq Q_i(i-c, \dots, i-1, i)$. By the definition of ads' relevance for positive externalities, there are $P^+, P^- \in [0,1]$, which correspond to the contribution of ads $i+1, \dots, i+c$ and $i-c, \dots, i-1$ to the product in (2), such that:

- $Q_i(i, i+1, \dots, i+c) = 1-(1-q_i)P^+$,
- $Q_i(i-c, \dots, i-1, i) = 1-(1-q_i)P^-$, and
- $Q^*_i = Q_i(i-c, \dots, i_j, \dots, i+c) = 1-(1-q_i)P^+P^-$.

As in the proof of Theorem 3,

$$\begin{aligned} 1-(1-q_i)P^+ + 1-(1-q_i)P^- &\geq 1-(1-q_i)P^+P^- + (1-q_i)(1-P^+)(1-P^-) \\ &\geq 1-(1-q_i)P^+P^- , \end{aligned}$$

which implies that $Q^{(1)}_i + Q^{(2)}_i \geq Q^*_i$, for all ads $i = c+1, \dots, k$.

The case where i is assigned to the last c slots of some block $s^{(1)}_p$ by $\pi^{(1)}_p$ is symmetric, and can be handled similarly, with the roles of $Q^{(1)}_i$ and $Q^{(2)}_i$ switched. This concludes the proof of claim (i).

We proceed to establish that given a set packing of weight W, we can efficiently construct a sponsored list of social welfare at least W for the original instance. Let $s_1, \dots, s_{k/(2c)}$ be a packing of total weight $W = W(s_1) + \cdots + W(s_{k/(2c)})$, and for each $p = 1, 2c+1, 4c+1, \dots, k-2c+1$, let π_p be the best permutation of ads in s_p. Since the union of $s_1, \dots, s_{k/(2c)}$ contains k ads, since $W(s_p, \pi_p) = W(s_p)$, and since we consider positive externalities, the sponsored list where each ad i in set s_p is assigned to slot $p + \pi_p(i) - 1$ of the sponsored list has a social welfare of at least W. $\square$

A.5 Proof of Proposition 2

Take 3 players with $v_1 = v_2 = V, v_3 = v$ with equal intrinsic relevances $q < 1$ and choose V, v so that $V > \frac{\sqrt{2-q}}{\sqrt{2-q}-1} \cdot v$. We consider the social context graph with $E^+ = \{(3,1),(3,2)\}$, $E^- = \emptyset$, window size $c = 1$, and $w_{31}(1) = w_{32} = 1$. The CTRs of slots are such that $\gamma_2 = \lambda_2/\lambda_1 > \frac{V}{(2-q)(V-v)}$, $\gamma_2 < 1 - \frac{v}{V}$ and $\gamma_3 = \lambda_3/\lambda_2 < 1 - \frac{v}{V}$. Notice that γ_2 can be feasibly chosen within a non-empty range $(\frac{V}{(2-q)(V-v)}, 1-\frac{v}{V})$, because $V > \frac{\sqrt{2-q}}{\sqrt{2-q}-1} \cdot v$. Also, observe that, influence to any of players 1, 2 by player 3, boosts their relevance to $Q = q(2-q)$. By conditioning on the slot assignment of player 3, we show that one of the other advertisers has incentive to deviate. We will first assume distinct bids that produce the corresponding assignments and comment on failure of a deterministic tie-breaking rule afterwards.

CASE I. In any assignment where $\phi^{-1}_{\boldsymbol{b}}(3) = 3$, should player $\phi_{\boldsymbol{b}}(1)$ aim for slot 2, his utility would become:

$$\begin{aligned} \lambda_2 \cdot Q \cdot (v_{\phi_{\boldsymbol{b}}(1)} - b_3) &\geq \lambda_2 \cdot Q \cdot (V - v) && \{\text{because } b_3 \leq v\} \\ &= \tfrac{q(2-q)\cdot(V-v)}{V} \cdot \lambda_2 \cdot V && \{\text{because } Q = q(2-q)\} \\ &> \gamma_2^{-1} \cdot \lambda_2 \cdot q \cdot (V - b_{\phi_{\boldsymbol{b}}(2)}) && \{\text{because } \gamma_2 > \tfrac{V}{(2-q)(V-v)}\} \\ &= \lambda_1 \cdot q \cdot (v_{\phi_{\boldsymbol{b}}(1)} - b_{\phi_{\boldsymbol{b}}(2)}) = u_{\phi_{\boldsymbol{b}}(1)}(\boldsymbol{b}) && \{\text{because } \gamma_2 = \lambda_2/\lambda_1\} \end{aligned}$$

CASE II. In any assignment with $\phi_{\boldsymbol{b}}^{-1}(3) = 2$, should player $\phi_{\boldsymbol{b}}(3)$ aim for slot 2, his utility would become:

$$\begin{aligned}
\lambda_2 \cdot Q \cdot (v_{\phi_{\boldsymbol{b}}(3)} - b_3) &\geq \lambda_2 \cdot Q \cdot (V - v) && \{\text{because } b_3 \leq v\} \\
&= \left(1 - \tfrac{v}{V}\right) \cdot \lambda_2 \cdot Q \cdot V \\
&> \gamma_3 \cdot \lambda_2 \cdot Q \cdot V && \{\text{because } \gamma_3 < 1 - \tfrac{v}{V}\} \\
&= \lambda_3 \cdot Q \cdot v_{\phi_{\boldsymbol{b}}(3)} = u_{\phi_{\boldsymbol{b}}(3)}(\boldsymbol{b}) && \{\text{because } \gamma_3 = \lambda_3/\lambda_2\}
\end{aligned}$$

CASE III. In any assignment with $\phi_{\boldsymbol{b}}^{-1}(3) = 1$, should player $\phi_{\boldsymbol{b}}(2)$ aim for slot 1, his utility would become:

$$\begin{aligned}
\lambda_1 \cdot Q \cdot (v_{\phi_{\boldsymbol{b}}(2)} - b_3) &\geq \lambda_1 \cdot Q \cdot (V - v) && \{\text{because } b_3 \leq v\} \\
&= \left(1 - \tfrac{v}{V}\right) \cdot \lambda_1 \cdot Q \cdot V \\
&> \gamma_2 \cdot \lambda_1 \cdot Q \cdot (V - b_{\phi_{\boldsymbol{b}}(3)}) && \{\text{because } \gamma_2 < 1 - \tfrac{v}{V}\} \\
&= \lambda_2 \cdot Q \cdot (v_{\phi_{\boldsymbol{b}}(2)} - b_{\phi_{\boldsymbol{b}}(3)}) = u_{\phi_{\boldsymbol{b}}(2)}(\boldsymbol{b}) && \{\text{because } \gamma_2 = \lambda_2/\lambda_1\}
\end{aligned}$$

Having examined all possible assignments, we turn our attention to the possibility of ties. We consider a deterministic tie-breaking rule that ranks player 3 higher than players 1 and 2 in case that he ties with them. An arbitrary deterministic rule is chosen for resolving a tie among 1 and 2. Then it is easy to see that any possible tie leads to one of the cases **I.**, **II.**, **III.** above. □