

Towards Better Models of Externalities in Sponsored Search Auctions

Nicola Gatti¹ and Marco Rocco¹ and Paolo Serafino² and Carmine Ventre³

Abstract. Sponsored Search Auctions (SSAs) arguably represent the problem at the intersection of computer science and economics with the deepest applications in real life. Within the realm of SSAs, the study of the effects that showing one ad has on the other ads, a.k.a. externalities in economics, is of utmost importance and has so far attracted the attention of much research. However, even the basic question of modeling the problem has so far escaped a definitive answer. The popular cascade model is arguably too idealized to really describe the phenomenon yet it allows a good comprehension of the problem. Other models, instead, describe the setting more adequately but are too complex to permit a satisfactory theoretical analysis. In this work, we attempt to get the best of both approaches: firstly, we define a number of general mathematical formulations for the problem in the attempt to have a rich description of externalities in SSAs and, secondly, prove a host of results drawing a nearly complete picture about the computational complexity of the problem. We complement these approximability results with some considerations about mechanism design in our context.

1 Introduction

The computation of solutions maximizing the *social welfare*, i.e., maximizing the total “happiness” of the advertisers, in sponsored search auctions (SSAs) strongly depends on how such happiness is defined. Clearly, the more clicks their ads receive, the more content advertisers are. A naive measure to forecast clicks, named click through rate (CTR), would only consider the *quality* of the ad itself (“better” ads receive more clicks). However, one should not overlook the importance of *externalities* in this context: specifically, *slot-dependent externalities* (i.e., ads positioned higher in the list have a higher chance to get a click) and *ad-dependent externalities* (e.g., the ad of a strong competitor – e.g., BMW – shown in the first slot can only decrease the number of clicks that the ad – e.g., of Mercedes – in the second slot gets).

Related Work. Much research focused on modeling externalities in SSAs and providing algorithms for the resulting optimization problem. On one hand of the scale, there is the simple, yet neat, *cascade* model [10, 1]. In the cascade model, users are assumed to scan the ads *sequentially* from top to bottom and the probability with which a user clicks on the ad a_i shown in slot s_m is the product of the intrinsic quality q_i of the ad, the relevance λ_m of slot s_m (slot-dependant externality) and of *all* the ads allocated to slots s_1 through s_{m-1} . A host of results is proved in this model as the input parameters

vary (e.g., $\lambda_m \in \{0, 1\}$ rather than $\lambda_m \in [0, 1]$). In its more general version, the optimization problem of social welfare maximization is conjectured to be NP-hard, shown to be in APX (i.e., a $1/4$ -approximation algorithm is given) and shown to admit a QPTAS (a quasi-polynomial time approximation scheme) [10]. In addition to its unknown computational complexity, the cascade model has two main limitations to be considered a satisfactory model of externalities in SSAs. First, it assumes that users have unlimited “memory” and that, consequently, an ad in slot s_1 exerts externalities on an ad many slots below. This is experimentally disproved in [8] wherein it is observed how the *distance* between ads is important. Second, it assumes that the externality of an ad is the same no matter which ad it is exerted on. Nevertheless, while BMW can have a strong externality on Mercedes since both makers attract the high end of the market, the externality on makers in a different price bracket, e.g., KIA, is arguably much less strong.

On the other hand of the scale, we can find models that try to address these limitations. In [6] Fotakis *et al.* propose a model whereby users have limited memory, i.e., externalities occur only within a *window* of c consecutive slots, and consider the possibility that externalities boost CTRs (*positive* externalities) as well as reduce CTRs (*negative* externalities). In particular, the externalities of an ad apply to ads displayed c slots below (*forward* externalities) and ads displayed c slots above (*backward* externalities). Moreover, in order to model the fact that externalities might have ad-dependent effect, they introduce the concept of *contextual graph*, whereby vertices represent ads and edge weights represent the externality between the endpoints. Their model turned out to be too rich to allow tight and significant algorithmic results (their main complexity results apply to the arguably less interesting case of forward positive externalities).

Our contribution. The present work can be placed in the middle of this imaginary spectrum of models for externalities in SSAs. Our main aim is to enrich the literature by means of more general ways to model slot- and ad-dependent externalities, while giving a (nearly) complete picture of the computational complexity of the problem. We do not attempt to explicitly model the user’s behavior but bridge the aforementioned models in order to overcome the respective weaknesses. In detail, we enrich the naive model of SSAs by adding the concepts of window and contextual externalities, while keeping ad- and slot-dependent externalities factorized as in the cascade model. We also complement much of the known literature by studying a model wherein the externalities coming from ads and slots cannot be expressed as a product. Our study gives rise to a number of novel and rich models for which we can provide (often tight) approximability results. Since the case of *selfish* advertisers is of particular relevance in this context, we also initiate the study of mechanism design for the optimization problems introduced and consider

¹ Politecnico di Milano, Italy. Email: { nicola.gatti, marco.rocco }@polimi.it

² Teesside University, UK. Email: p.serafino@tees.ac.uk

³ Teesside University, UK. Email: c.ventre@tees.ac.uk

This author was supported by EPSRC, through grant EP/M018113/1.

the *incentive-compatibility* of our algorithms, i.e., whether they can be augmented with payment functions so to work also with selfish advertisers.

We analyze two families of models with *Forward Negative Externalities* (FNE): one, named FNE_{sa} , in which slot- and ad-dependent externalities *cannot* be factorized and another one, named FNE_{aa} , where we assume that externalities *can* be factorized. Furthermore, we discriminate our models along two other dimensions that are orthogonal to the factorization of externalities. The first dimension has to do with the size c of the user memory with respect to the number of slots K . Specifically, we consider the cases when users have “limited” memory, i.e., $c < K$, and “large” memory, i.e. $c = K$. The second dimension regards the effect empty slots have on users’ attention. In a sort of whole page optimization fashion [11], we think of those slots as occupied by a *special* fictitious ad (e.g., an image) that either raises users’ attention (*reset* model), or does nothing (*non-reset* model). The combination of these orthogonal dimensions give rise to the models listed in the first row of Table 1, which gives an overview of our results (formal definitions of these models are given in Section 2).

For $FNE_{sa}(c)$ (i.e. the version in which slot- and ad-dependant externalities cannot be factorized and externalities occur in a window of size c) we prove that the optimization problem is in P , if c is a constant. We consider the LP relaxation of the ILP describing the problem and prove that the integrality gap is 1.

For $FNE_{aa}(c)$ (i.e. the variant of the problem with factorized externalities, contextual ad-dependent externalities and window of c slots) the aforementioned distinction on the effects that empty slots have on users’ behavior is useful. For $FNE_{aa}(c)$ -nr (i.e. the variant of $FNE_{aa}(c)$ where the special ad cannot be used, or, equivalently, the user’s attention cannot be reset) we prove that the allocation problem is poly-APX-complete whenever $c = K$. Specifically, we give an approximation preserving reduction from the Longest Path problem and design an approximation algorithm using several different ideas and sources of approximation; interestingly, its approximation guarantee matches the best known approximation guarantee for Longest Path. However, we prove that this algorithm cannot be used in any truthful mechanism and note that a simple single-item second price auction gives a weaker, yet close, truthful approximation. We complement the results for this model with the identification of tractable instances for which we provide an exact polynomial-time algorithm. For $c < K$ instead, we are unable to determine the exact hardness of approximating the problem in general. To the APX-hardness proof, we pair a number of approximation algorithms that assume constant c . The first, based on color coding [2], returns a non-constant approximation on any instance of SSA. The second assumes that the contextual graph is complete and returns a solution which (roughly) guarantees a γ_{\min}^c fraction of the optimum social welfare, γ_{\min} being the minimum edge weight in the graph. Interestingly, this algorithm shows the APX-completeness of the subclass of instances having constant γ_{\min} (we indeed further provide a hardness result for instances with complete contextual graphs). We believe the tight result for this subclass of instances to be quite relevant. In fact, complete contextual graphs are quite likely to happen in real-life: the results returned by a keyword search are highly related to one another, and, as such, each pair of ads has a non-null externality, however small.

For $FNE_{aa}(c)$ -r (i.e., the variant of $FNE_{aa}(c)$ where the special ad can be used to refresh users’ memory) the problem becomes easier and turns out to be APX-complete, for any c . We first prove the problem with $c = K$ to be APX-hard, via a reduction from (a subclass of) ATSP (i.e., asymmetric version of TSP) and then surprisingly con-

nect instances with $c < K$ to instances with $c = K$ by reducing the case with $c = 1$ to the case with $c = K$ and *binary* externalities (i.e., the weights of the edges of the contextual graph can be either 0 or 1). We also observe how a simple greedy algorithm cleverly uses the special ad to return $1/2$ -approximate solutions and leads to a truthful mechanism.

Three final observations on our set of results are needed. Firstly, as common in the literature on SSAs, the number of slots is an input parameter of the problem, rather than a fixed constant. This is consistent with common practice of SSA in real life in some scenarios: for instance, the number of ads displayed by major search engines like Google and Bing is not constant, but rather varies with respect to the keyword being used. However, in the scenarios where K can be assumed to be constant, the allocation problem in our models becomes computationally tractable (by, e.g., running the color coding algorithm) and truthfulness can be achieved by imposing VCG payments. Secondly, more research is needed to complete the picture concerning truthful SSAs with externalities. The fact that we principally address computational complexity issues is nevertheless a needed step in this direction. Indeed, recall that in our (single-parameter) setting, truthfulness reduces to designing a (computationally efficient) monotone algorithm (see below for details). Therefore, settling the computational complexity serves as a benchmark for the approximation guarantee of the monotone algorithms we design. Thirdly, our analysis is worst case and then does not rule out the possibility that on real-life instances the approximation ratio obtained is smaller: e.g., the best known algorithm for the cascade model performs better than its theoretical approximation guarantee.

Due two space limitation, we have omitted the proofs of some minor results. All the details are available in a longer pre-print version of this paper [7].

2 Model

In a SSA we have N ads and K slots. We assume that each ad corresponds to an advertiser; this is w.l.o.g. from the optimization point of view. We denote each ad by a_i with $i \in \mathcal{N}$, where $\mathcal{N} = \{1, \dots, N\}$ is the set of indices of the ads. We introduce a fictitious ad, denoted by a_{\perp} , s.t., when allocated, the slot is left empty. The K slots are denoted by s_m with $m \in \mathcal{K}$, $\mathcal{K} = \{1, \dots, K\}$ being the set of slot indices s.t. s_1 is the slot at the top of the page and s_K is at the bottom. We also have a fictitious slot, denoted by s_{\perp} s.t. an ad allocated to s_{\perp} is not displayed in the webpage. Each ad a_i is characterized by: (i) the *quality* $q_i \in [0, 1]$, i.e., the probability a user clicks on ad a_i when he observes it, irrespectively of other externalities; (ii) the *valuation* $v_i \in \mathbb{R}^+$ advertiser i associates to his ad being clicked by a user. The fictitious ad a_{\perp} has $q_{\perp} = v_{\perp} = 0$.

A feasible allocation of ads to slots, denoted as θ , consists of an ordered sequence of ads $\theta = \langle a_1, \dots, a_K \rangle$ s.t. the ads are ordered by increasing slot number, i.e., a_1 is allocated to the top slot, a_K to the bottom one. Every ad a_i can be allocated to at most one slot, whereas a_{\perp} can be allocated to more than one slot. The set of all possible feasible allocations is denoted as Θ . With a slight abuse of notation, we let (i) $\theta(a_i)$ denote the index of the slot ad a_i is allocated to, and (ii) $\theta(s_m)$ denote the index of the ad allocated to slot s_m . Given $\theta \in \Theta$, the *click through rate* of ad a_i , denoted as $CTR_i(\theta)$, is the probability ad a_i is clicked by the user taking externalities into consideration. The optimal allocation θ^* is the one maximizing the *social welfare*, namely: $\theta^* \in \arg \max_{\theta \in \Theta} SW(\theta)$, where

$$SW(\theta) = \sum_{i \in \mathcal{N}} CTR_i(\theta)v_i.$$

	FNE _{aa} (c)		FNE _{aa} (K)		FNE _{sa} (c)
	nr	r	nr	r	
LB	APX-hard	APX-complete	poly-APX-complete	APX-complete	P *
UB	$\frac{\log(N)}{2 \min\{N, K\}}$ *				
SP	$\frac{\log(N)}{2 \min\{N, K\}}$ *	1/2	1/K	1/2	

Table 1. Summary of our results: LB (UB, resp.) stands for lower (upper, resp.) bound on the approximation of the problem; the row SP, instead, contains the approximation guarantees we obtain with truthful mechanisms. Results marked by ‘*’ require $c = \mathcal{O}(1)$. APX-completeness of a subclass of FNE_{aa}(c)-nr is also given. (See the model for details on the notation.)

A $1/\alpha$ -approximate solution θ satisfies $SW(\theta) \geq SW(\theta^*)/\alpha$.

Typically, $CTR_i(\theta)$ defines how the quality q_i of ad a_i is “perturbed” by the externalities in terms of click probability. Accordingly, in general $CTR_i(\theta) = q_i \Gamma_i(\theta)$, $\Gamma_i(\theta)$ being a function encoding the effect of externalities. For instance, in the cascade model, $\Gamma_i(\theta) = \Lambda_{\theta(a_i)} \prod_{l=1}^{\theta(a_i)-1} \gamma_{\theta(s_l)}$, where $\Lambda_{\theta(a_i)} = \prod_{l=1}^{\theta(a_i)} \lambda_l$, $\lambda_m \in [0, 1]$, called the *factorized prominence* of s_m , denotes the slot-dependant externality and $\gamma_i \forall i \in \mathcal{N}$, called *continuation probability*, denotes the ad-dependent externality. (W.l.o.g., we assume $\Lambda_1 = \lambda_1 = 1$.) Our conceptual contribution rests upon novel and richer ways to define $\Gamma_i(\theta)$, along three main dimensions.

The *first dimension* concerns the *user memory*, a.k.a. *window*. We let c be the number of ads displayed above a_i in θ , from $s_{\theta(a_i)-1}$ to $s_{\theta(a_i)-c}$, that affect $\Gamma_i(\theta)$. The *second dimension* concerns a generalization of the externalities. Here we propose two alternative families of externalities, called *sa* (for slot-ad) and *aa* (for ad-ad). The sa-externalities remove the factorization in slot- and ad-dependent externalities: i.e., λ_m and γ_i are substituted by parameters $\gamma_{m,j} \in [0, 1]$, $m \in \mathcal{K}$ and $j \in \mathcal{N}$. When the window is c , the CTR is defined as $CTR_i(\theta) = q_i \Gamma_i(\theta)$, where

$$\Gamma_i(\theta) = \prod_{m=\max\{1, \theta(a_i)-c\}}^{\theta(a_i)-1} \gamma_{m, \theta(s_m)}.$$

This definition captures the situation in which an ad can affect the ads displayed below it in a different way according to the position in which it is displayed. For the aa-externalities, on the other hand, we preserve the factorization in λ_m and γ_i , but redefine these latter parameters as $\gamma_{i,j} \in [0, 1]$ where a_j is the ad that is displayed in the slot just below $\theta(a_i)$. It is convenient to see the $\gamma_{i,j}$ ’s as the weights of the *contextual graph* $G = (\mathcal{N}, \mathcal{E})$ where the direct edges (i, j) weigh $\gamma_{i,j} > 0$ and represent the way ad a_i influences a_j . Note that non-edges of G correspond to the pairs of ads a_i, a_j s.t. $\gamma_{i,j} = 0$. Here, with window c ,

$$\Gamma_i(\theta) = \Lambda_{\theta(a_i)} \prod_{l=\max\{1, \theta(a_i)-c\}}^{\theta(a_i)-1} \gamma_{\theta(s_l), \theta(s_{l+1})}$$

where Λ_m is defined as above. This definition captures the situation in which each ad can affect each other ad in a different way.

The *third dimension* concerns the definition of $\gamma_{m,\perp}$ for the sa-externalities and $\gamma_{i,\perp}$ and $\gamma_{\perp,i}$ for the aa-externalities. In the model *with reset* we have $\gamma_{m,\perp} = 1$ for sa and $\gamma_{i,\perp} = \gamma_{\perp,i} = 1 \forall i \in \mathcal{N} \cup \{\perp\}$ for aa. This variant captures the situation in which slots can be distributed in the page in different positions (a.k.a., slates) and, in order to raise the user’s attention, we can allocate a content, e.g. pictures, that nullifies the externality between the ad allocated before and after the content. In the model *without reset*, $\gamma_{m,\perp} = 0$ for sa and $\gamma_{i,\perp} = \gamma_{\perp,i} = 0 \forall i \in \mathcal{N} \cup \{\perp\}$ for aa, thus capturing the situation in which leaving a slot empty between two allocated slots does not provide any advantage.

We let FNE_x(c)-y be the problem of optimizing the social welfare in our model with *Forward Negative Externalities* with window c , $x \in \{sa, aa\}$ -externalities and $y \in \{r, nr\}$ reset (r stands for reset; nr for no reset). When the value of y is not relevant for our results, we talk about FNE_x(c). We are interested in two particular subclasses of FNE_{aa}(c), namely: (i) subclass FNE_{aa}⁺(c)-y, defined upon a complete contextual graph and such that $0 < \gamma_{\min} = \min_{i,j \in \mathcal{N}, i \neq j} \gamma_{i,j}$ and (ii) subclass \mathcal{B} -FNE_{aa}(c)-y, where $\gamma_{i,j}$ can take values in $\{0, 1\}$. **Mechanism design.** A *mechanism* M is a pair (f, p) , where $f : (\mathbb{R}^+)^N \rightarrow \Theta$ is an algorithm that associates to any vector $\mathbf{v} = (v_1, \dots, v_N)$ of valuations a feasible outcome in Θ (only valuations are private knowledge). The payment function $p_i : (\mathbb{R}^+)^N \rightarrow \mathbb{R}^+$ maps valuation vectors to monetary charges for advertiser i . The aim of each advertiser is to maximize his own utility $u_i(\mathbf{v}, v_i) = CTR_i(f(\mathbf{v}))v_i - p_i(\mathbf{v})$. An advertiser could misreport his true valuation and declare $\hat{v}_i \neq v_i$ when $u_i((\hat{v}_i, \mathbf{v}_{-i}), v_i) > u_i(\mathbf{v}, v_i)$, \mathbf{v}_{-i} being the vector of the valuations of all the agents but i . We are then interested in truthful mechanisms. A mechanism is *truthful* if for any $i \in \mathcal{N}$, $\mathbf{v}_{-i} \in (\mathbb{R}^+)^{N-1}$, $v_i, \hat{v}_i \in \mathbb{R}^+$, $u_i((\hat{v}_i, \mathbf{v}_{-i}), v_i) \leq u_i(\mathbf{v}, v_i)$.

In this setting, a monotone algorithm *must* be used in truthful mechanisms [3]. Algorithm f is monotone if for any $i \in \mathcal{N}$, $\mathbf{v}_{-i} \in (\mathbb{R}^+)^{N-1}$, $CTR_i(f(\hat{v}_i, \mathbf{v}_{-i}))$ is non-decreasing in \hat{v}_i . Important for our work is also the family of VCG-like mechanisms, a.k.a., *Maximal In Range (MIR)* mechanisms. An algorithm f is MIR if there exists $\Theta' \subseteq \Theta$ s.t. $f(\mathbf{v}) \in \arg \max_{\theta \in \Theta'} SW(\theta) \forall \mathbf{v} \in \mathbb{R}^N$ [12]. These algorithms can be augmented with a VCG-like payment so to obtain truthful mechanisms. (VCGs are MIR mechanisms wherein $\Theta' = \Theta$.) We are interested in mechanisms for which both f and p are computable in polynomial time. MIR mechanisms run in polynomial-time if the MIR algorithm does. As usual in the context of SSA, we adopt a pay-per-click payment scheme, i.e., we charge $p_i(\mathbf{v})/CTR_i(f(\mathbf{v}))$ when a user clicks on a_i .

3 FNE_{sa}(c) is in P for constant c

Our presentation focuses on FNE_{sa}(1)-nr to simplify the notation. The more general cases when $c > 1$ and the reset model is considered are easily obtainable by generalization from FNE_{sa}(1), but require a more cumbersome notation without significant new ideas (see discussion at the end of this section). We first give the ILP formulation of FNE_{sa}(1)-nr and prove that if there is an optimal fractional solution, then there are at least two feasible integral solutions with the same value of social welfare. Since it is well known, by LP theory, that the ellipsoid algorithm can be forced (in polynomial-time) to output an integral optimal solution, we are able to prove the following:

Theorem 1. *For $c = \mathcal{O}(1)$, there is a polynomial-time optimal algorithm for FNE_{sa}(c).*

FNE_{sa}(1)-nr can be formulated as following ILP:

$$\begin{aligned}
& \max \sum_{m=2}^K \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}, j \neq i} \gamma_{m-1,j} q_i v_i x_{j,m,i} + \sum_{i \in \mathcal{N}} x_{1,i} q_i v_i \\
& \text{subject to:} \\
& \sum_{m=2}^K \sum_{j \in \mathcal{N}, j \neq i} x_{j,m,i} + x_{1,i} \leq 1 \quad \forall i \in \mathcal{N} \\
& x_{1,i} = \sum_{j \in \mathcal{N}, j \neq i} x_{i,2,j} \quad \forall i \in \mathcal{N} \\
& \sum_{j \in \mathcal{N}, j \neq i} x_{j,m,i} = \sum_{j \in \mathcal{N}, j \neq i} x_{i,m+1,j} \quad \forall i \in \mathcal{N}, \\
& \hspace{15em} 2 \leq m < K \\
& \sum_{i \in \mathcal{N}} x_{1,i} = 1 \quad (1) \\
& \sum_{j \in \mathcal{N}} \sum_{i \in \mathcal{N}, i \neq j} x_{j,m,i} = 1 \quad \forall m \in \mathcal{K} \setminus \{1\} \\
& x_{1,i} \in \{0, 1\} \quad \forall i \in \mathcal{N} \\
& x_{j,m,i} \in \{0, 1\} \quad \forall 2 \leq m \leq K, \\
& \hspace{15em} i, j \in \mathcal{N}, i \neq j
\end{aligned}$$

where $x_{j,m,i} = 1$ iff a_i is allocated to slot s_m , $m > 1$, and a_j is allocated to slot s_{m-1} ; $x_{1,i} = 1$ iff a_i is allocated to s_1 . The objective function and the constraints are rather straightforward and, hence, their description is omitted here.

The next proposition proves Theorem 1 since it shows that we can solve the above ILP in polynomial-time, despite its similarities with the 3D-assignment, a well-known NP-hard problem.

Proposition 1. *The continuous relaxation of the above ILP always admits integral optimal solutions.*

Proof. We show that, if there is an optimal fractional solution x , then there are at least two feasible integral solutions with the same value of social welfare. Specifically, we prove that x is equivalent to a probability distribution over integral allocations $\theta = \langle a_1, \dots, a_K \rangle$. The probability $\mathbb{P}(\theta)$ given to θ is:

$$\begin{aligned}
\mathbb{P}(\theta) &= \prod_{i=1}^K \mathbb{P} \left(\theta(a_i) = s_i \mid \bigwedge_{j < i} \theta(a_j) = s_j \right) \\
&= x_{1,1} \prod_{l=2}^K \frac{x_{l-1,l,l}}{\sum_{m \geq l} x_{l-1,l,m}}.
\end{aligned}$$

In order to show that $\mathbb{P}(\theta)$ is actually a probability distribution over allocations, we show that $\sum_{\theta \in \Theta} \mathbb{P}(\theta) = 1$.

The proof is recursive. Let Θ' be the set of allocations θ with the same first $K-1$ ads. The allocations in Θ' differ only for the ad allocated to s_K . To fix the notation, for $\theta \in \Theta'$ let $\theta(s_l) = a_l$, for $l < K$. We have:

$$\begin{aligned}
\sum_{\theta \in \Theta'} \mathbb{P}(\theta) &= x_{1,1} \prod_{l=2}^{K-1} \left(\frac{x_{l-1,l,l}}{\sum_{m \geq l} x_{l-1,l,m}} \right) \sum_{h \geq K} \frac{x_{K-1,K,h}}{\sum_{m \geq K} x_{K-1,K,m}} \\
&= x_{1,1} \prod_{l=2}^{K-1} \left(\frac{x_{l-1,l,l}}{\sum_{m \geq l} x_{l-1,l,m}} \right) \frac{\sum_{h \geq K} x_{K-1,K,h}}{\sum_{m \geq K} x_{K-1,K,m}} \\
&= x_{1,1} \prod_{l=2}^{K-1} \left(\frac{x_{l-1,l,l}}{\sum_{m \geq l} x_{l-1,l,m}} \right).
\end{aligned}$$

By applying recursively the same argument above from $\Theta'' \supset \Theta'$, the set of all allocations θ satisfying $\theta(s_l) = a_l$, for $l \leq K-2$, down to the set of allocations having only the same first ad, we have $\sum_{\theta: \theta(s_1)=a_1} \mathbb{P}(\theta) = x_{1,1}$. Since (1) forces $\sum_{i \in \mathcal{N}} x_{1,i} = 1$, we have $\sum_{\theta \in \Theta} \mathbb{P}(\theta) = \sum_{i \in \mathcal{N}} x_{1,i} = 1$. This shows that $\mathbb{P}(\theta)$ is a well defined probability distribution. The proof concludes by observing that all integral solutions are indeed feasible. \square

To solve the problem when $c > 1$, we just need to modify the ILP and allow each variable x to depend on $c+2$ indices to take into account the (at most) c indices of all the ads that precede the ad of interest. The reset model for $c = 1$ instead requires the introduction of K additional variables for a_{\perp} to be visualized in each slot (together with some constraints to fix each variable for a_{\perp} to a slot).

Theorem 1 implies that mechanism design becomes an easy problem for FNE_{sa}(c) and $c = \mathcal{O}(1)$, since the optimal algorithm can be used to obtain a truthful VCG mechanism.

4 FNE_{aa}(K)-nr is Poly-APX-complete

Easy Instances. As a warm-up, we identify a significant class of instances of FNE_{aa}(K)-nr for which we can design a polynomial-time optimal algorithm. These instances are characterized by the fact that the underlying contextual graph is a DAG, thus modeling nearly oligopolistic markets in which the ads can be organized hierarchically. The idea of Algorithm 1 is that since DAGs can be sorted topologically in polynomial time then we can *rename* the ads as a_1, \dots, a_N so to guarantee that for any pair of ads a_i, a_j , if $i < j$ then $(a_j, a_i) \notin \mathcal{E}$. We can then prove that we can focus w.l.o.g. on *ordered* allocations θ , i.e., for any pair of allocated ads a_i, a_j , with $i < j$, $\theta(a_i) \leq \theta(a_j)$. Consider an unordered θ and let a_i be the first ad (from the top) for which there exists a_j , $i < j$, such that $\theta(a_i) > \theta(a_j)$. Since $\gamma_{j,i} = 0$ then all the ads a_k s.t. $\theta(a_k) \geq \theta(a_i)$ have $CTR_k(\theta) = 0$ and, therefore, we can prune θ of (i.e., substitute with a_{\perp}) a_i and all the subsequent ads without any loss in the social welfare. But then in the class of ordered allocations, the optimum has an optimal substructure and we can use dynamic programming. Let $D[i, m]$ be the value of the optimal ordered allocation that uses only slots s_m, \dots, s_K and allocates ad a_i in s_m . It is not hard to see that $D[i, m] = \Lambda_m q_i v_i + \max_{j > i} \gamma_{i,j} D[j, m+1]$ and that the optimum is $\max_{i \in [N]} D[i, 1]$. In the pseudo-code of the algorithm, we simply construct the table D after the topological sort of the contextual graph (with renaming of the ads) is done. The algorithm runs in time $\mathcal{O}(KN^2)$.

Algorithm 1

```

1: TOPOLOGICALSORT( $G$ )
2: for all  $m \leq K$  do
3:    $D[N, m] = \Lambda_m q_N v_N$ 
4: for all  $i \leq N$  do
5:    $D[i, K] = \Lambda_K q_i v_i$ 
6: for  $i = N-1$  to 1 do
7:   for  $m = K-1$  to 1 do
8:      $D[i, m] = \Lambda_m q_i v_i + \max_{j > i} \gamma_{i,j} D[j, m+1]$ 
9: return ( $\max_{i \in [N]} D[i, 1]$ )

```

Since social welfare maximization is a utilitarian problem, and given that the algorithm above is optimal we can use the VCG mechanism to obtain a polynomial-time optimal truthful mechanism.

Hardness. We now prove the hardness of approximating FNE_{aa}(K)-nr.

Theorem 2. $FNE_{aa}(K)$ -nr is poly-APX-hard.

Proof. We reduce from the Longest Path problem. An instance of the Longest Path problem consists of a directed graph $G' = (T, A)$ where T is the set of vertices of the graph and $A \neq \emptyset$ is the set of unweighted edges. The problem demands to compute a *longest simple path*, i.e., a maximum length path that visits each vertex of the graph at most once. This problem is poly-APX-complete [5] and the best known asymptotic approximation is $\log |T|/|T|$. From an instance $G' = (T, A)$ of Longest Path we obtain an instance of $FNE_{aa}(K)$ -nr as follows. For each vertex $t_i \in T$ we add an ad a_i , with $q_i = v_i = 1$ and for each directed arc $(t_i, t_j) \in A$ we add an arc $(i, j) \in \mathcal{E}$. Furthermore, we set $\gamma_{i,j} = 1$ if $(i, j) \in \mathcal{E}$ and $\gamma_{i,j} = 0$ otherwise. Finally, we set $N = K = |T|$ and $\Lambda_m = 1, \forall m \in [K]$.

Given an ordered sequence of vertices $\rho = (t_1, t_2, \dots, t_N)$, we denote as $len(\rho)$ the length of the path that starts in t_1 and visits the nodes in ρ till the first node t_j s.t. $(t_j, t_{j+1}) \notin A$ is reached. Let us denote as ρ^* the sequence that describes the longest path in G' and as θ^* the allocation that maximizes the social welfare in the instance of $FNE_{aa}(K)$ -nr defined upon G' . It is easy to check that $len(\rho^*) = SW(\theta^*) - 1$. Indeed, θ^* allocates sequentially from the first slot the ads that correspond to the vertices composing the longest path. Conversely, we can transform an allocation θ into a sequence of vertices ρ just by substituting the ads with their corresponding vertices until the first a_\perp in θ is found. Thus, we have that for θ and the corresponding ρ it holds $len(\rho) = SW(\theta) - 1$.

Consider a generic α -approximate allocation θ_α for $FNE_{aa}(K)$ -nr: $SW(\theta_\alpha) \geq \alpha SW(\theta^*)$. Since A is non-empty, there is a solution θ_2 to $FNE_{aa}(K)$ -nr of social welfare at least 2. Let θ_β denote the solution in $\{\theta_\alpha, \theta_2\}$ with maximum social welfare. As θ_α is an α -approximate solution so is θ_β . By letting ρ_β denote the path constructed from θ_β as described above, we prove that the reduction preserves the approximation (up to a constant factor): $len(\rho_\beta) = SW(\theta_\beta) - 1 \geq \frac{1}{2} SW(\theta_\beta) \geq \frac{\alpha}{2} SW(\theta^*) = \frac{\alpha}{2} (len(\rho^*) + 1) \geq \frac{\alpha}{2} len(\rho^*)$. \square

Approximation algorithm. We show that the problem is in poly-APX, with an approximation ratio that is asymptotically the same as the best guarantee known for Longest Path. Our algorithm combines the Color Coding (CC) algorithm [2] together with three approximation steps.

Let C be a set containing K different colors. CC is a random algorithm, randomly assigning colors from C to the ads, and then finding the best *colorful* (i.e., no pair of ads has the same color) allocation. To find the best colorful allocation, given a random coloring we do the following. For $S \subseteq C$, we define (S, a_i) as the set of partial allocations with the properties of having the same number $|S|$ of allocated ads (each colored with a different color of S) in the first $|S|$ slots and having ad a_i in slot $s_{|S|}$. We start from $S = \emptyset$ where no ad is allocated. Then, allocating one of the ads in the first position, we add one color to S until $S = C$. Iteratively, the algorithm extends the allocations in (S, a_i) appending a new ad, say a_j , with a color not in S in slot $s_{|S|+1}$ obtaining $(S \cup \{o_j\}, a_j)$ where o_j is the color of a_j . Each partial allocation in (S, a_i) is characterized by the values of SW and Γ_i . We can safely discard all the Pareto dominated partial allocations: given two allocations θ_1 and θ_2 in (S, a_i) , we say that θ_2 is Pareto dominated by θ_1 iff $SW(\theta_1) \geq SW(\theta_2)$ and $\Gamma_i(\theta_1) \geq \Gamma_i(\theta_2)$. However, there is no guarantee that the number of allocations in (S, a_i) is polynomially bounded and, in principle, all the generated $\mathcal{O}(N^K)$ partial allocations may be Pareto efficient. The complexity per coloring is $\mathcal{O}(2^K N^{K+1} K^2)$. CC generates e^K random colorings, but it can be derandomized with a cost of $\log^2(N)$

and a total complexity $\mathcal{O}((2e)^K K^2 N^{K+1} \log^2 N)$. To make the algorithm polynomial, we apply three approximation steps. Initially, we briefly sketch these three approximations and, subsequently, we provide the details. Firstly, we run CC over a reduced number K' of slots where $K' = \min\{\lceil \log(N) \rceil, K\}$. Secondly, we discard all the allocations θ in which the probability to click on the last allocated ad is smaller than a given $\delta \in [0, 1]$. Finally, we discretize the $\gamma_{i,j}$'s. We prove in the following that the running time is indeed polynomial and the approximation ratio is $(1 - \delta)(1 - \epsilon) \frac{\log(N)}{2 \min\{N, K'\}}$, ϵ controlling the granularity of the $\gamma_{i,j}$ discretization. All the three approximations are necessary in order to obtain a polynomial-time algorithm. This algorithm is not monotone as we show below. However, a simple $1/K$ -approximate truthful mechanism can be obtained, via a single-item second price auction. From here on, we provide the details of the algorithms and we prove its approximation ratio.

Approximation 1. We apply CC over a reduced number K' of slots, where $K' = \min\{\lceil \log(N) \rceil, K\}$, implying the following approximation ratio.

Proposition 2. Given θ^* , the optimal allocation over K slots, and $\theta_{K'}^*$, the optimal allocation over the first $K' \leq \min\{N, K\}$ slots, we have $SW(\theta_{K'}^*) \geq \frac{1}{2} \frac{K'}{\min\{N, K\}} SW(\theta^*)$.

Approximation 2. In CC, we discard allocations θ in which $\Gamma_i(\theta)$ of the last allocated ad a_i , $i \in [N]$, is less than a given $\delta \in [0, 1]$, implying the following approximation ratio.

Proposition 3. Given $\theta_{K'}^*$, the optimal allocation over K' slots, and $\theta_{K'}^{\delta}$, the optimal allocation among the allocations $\theta \in \Theta$ where the last allocated ad a_i , $i \leq N$, satisfies $\Gamma_i(\theta) \geq \delta$, we have $SW(\theta_{K'}^{\delta}) \geq (1 - \delta) SW(\theta_{K'}^*)$.

Approximation 3. In CC, we use rounded values for $\gamma_{i,j}$. More precisely, we use $\lfloor \frac{1}{\tau} \log \frac{1}{\gamma_{i,j}} \rfloor$ in place of $\log \frac{1}{\gamma_{i,j}}$, where the normalization constant τ is defined below. The constraint due to Proposition 3 is now a capacity constraint of the form $\sum_{m \in [K]: m < i} \lfloor \frac{1}{\tau} \log \frac{1}{\gamma_{\theta(m), \theta(m+1)}} \rfloor \leq \lfloor \frac{1}{\tau} \log \frac{1}{\delta} \rfloor$. Notice that, with rounded values, the capacity can assume a finite number of values (i.e., $\lfloor \frac{1}{\tau} \log \frac{1}{\delta} \rfloor$) and therefore we can now bound the number of allocations to be stored in (S, a_i) . More precisely, for each value of capacity, we can discard all the allocations except one maximizing the social welfare measured with rounded values. This step has the following consequences on the approximation guarantee.

Proposition 4. Given $\theta_{K'}^{\delta}$, defined as in Proposition 3, and $\theta_{K'}^{\delta\epsilon}$, the optimal allocation when the rounding procedure is applied, we have that, choosing $\tau = \frac{1}{K'} \log \frac{1}{1-\epsilon}$, $SW(\theta_{K'}^{\delta\epsilon}) \geq (1 - \epsilon) SW(\theta_{K'}^{\delta})$.

The approximation ratio of the algorithm is thus $(1 - \delta)(1 - \epsilon) \frac{\log(N)}{2 \min\{N, K'\}}$, asymptotically the same as the best known approximation ratio of the Longest Path once $N = K$. The complexity instead can be derived as follows. The maximum number of allocations that can be stored in each (S, a_i) is $\mathcal{O}(\frac{\log \frac{1}{\delta}}{\tau})$ with $\tau = \frac{\log \frac{1}{1-\epsilon}}{K'}$ thanks to dominations. Thus, given that $\log(\frac{1}{1-\epsilon}) \rightarrow \epsilon$ as $\epsilon \rightarrow 0$, the number of elements is $\mathcal{O}(K' \frac{1}{\epsilon})$. Thus, the complexity when $K' = \log(N)$ is $\mathcal{O}((2e)^{\log(N)} \frac{1}{\epsilon} \log(\frac{1}{\delta}) N^2 \log^4(N)) = \mathcal{O}(\frac{1}{\epsilon \delta} N^3 \log^4(N))$.

Notice that all the three above approximations are necessary in order to obtain a polynomial-time algorithm. Approximation 2 and Approximation 3 allow us to bound the number of the allocations stored per pair (S, a_i) and would lead, if applied without Approximation 1, to a complexity $\mathcal{O}((2e)^K K^2 N^2 \log^2(N) \frac{1}{\epsilon \delta})$. Notice also

that, without Approximation 2, the possible values for the capacity are not upper bounded. Approximation 1 allows us to remove the exponential dependence on K and to obtain polynomial complexity.

Non-monotonicity of the approximation algorithm. We here prove that the algorithm is not monotone and therefore we cannot augment it with a payment function to obtain a truthful mechanism.

Let us initially consider the case where Approximation 1 is not used, therefore all the K slots can be allocated. We will discuss below how to extend the proof to the case where Approximation 1 is used.

Consider the following instance of $\text{FNE}_{sa}(K)$ -nr:

- $K = 3$ slots;
- $N = 4$ ads, where $q_1 v_1 = 2^{2\tau} \frac{\Lambda_2 - \Lambda_3 2^{-6\tau}}{\Lambda_2 - \Lambda_3} + 3$, $q_2 v_2 = x$, $q_3 v_3 = q_4 v_4 = 1$, where τ is the generic rounding factor of Approximation 3;
- the contextual graph is s.t. $\gamma_{i,j} = 0 \forall i, j \in [N]$ except: $\gamma_{1,2} = 2^{-(4+\phi)\tau}$, $\gamma_{1,3} = 2^{-\tau}$, $\gamma_{2,4} = 2^{-\tau}$, $\gamma_{3,2} = 2^{-\tau}$. ϕ is a small number;
- the rounded capacity $\left\lfloor \frac{\log \frac{1}{\gamma_{i,j}}}{\tau} \right\rfloor = +\infty \forall i, j \in [N]$ except: $\left\lfloor \frac{\log \frac{1}{\gamma_{1,2}}}{\tau} \right\rfloor = 3$, $\left\lfloor \frac{\log \frac{1}{\gamma_{1,3}}}{\tau} \right\rfloor = 1$, $\left\lfloor \frac{\log \frac{1}{\gamma_{2,4}}}{\tau} \right\rfloor = 1$, $\left\lfloor \frac{\log \frac{1}{\gamma_{3,2}}}{\tau} \right\rfloor = 1$.
- the K colours are $\{o_1, o_2, o_3\}$.

The product $q_1 v_1$ has been chosen s.t., when x is in the neighbourhood of $2^{2\tau} \frac{\Lambda_2 - \Lambda_3 2^{-4\tau}}{\Lambda_2 - \Lambda_3}$, a_1 is always allocated in the first slot. Thus, we can focus only on the colouring that assigns colour o_1 to a_1 , o_2 to a_2 and o_3 to a_3 and a_4 . Indeed, with this colouring the two longest path of the contextual graph are colourful, i.e. the unique two colourful allocations are $\theta_1 = (a_1, a_3, a_2)$ in the set $(\{o_1, o_2, o_3\}, a_2)$ and $\theta_2 = (a_1, a_2, a_4)$ in the set $(\{o_1, o_2, o_3\}, a_4)$.

Notice that, with this colouring, all the allocations where there is a pair of ads (a_i, a_j) with $\gamma_{i,j} = 0$ are infeasible, not satisfying the capacity bound. We will now prove that the approximation algorithm is not monotone with respect to a_2 .

Let us denote by \widetilde{SW} the social welfare computed on the basis of the rounded values. It is easy to check that the following hold: $\widetilde{SW}(\theta_1) = 2^{2\tau} \frac{\Lambda_2 - \Lambda_3 2^{-6\tau}}{\Lambda_2 - \Lambda_3} + 3 + \Lambda_2 2^{-4\tau} x + \Lambda_3 2^{-6\tau}$ and $\widetilde{SW}(\theta_2) = 2^{2\tau} \frac{\Lambda_2 - \Lambda_3 2^{-6\tau}}{\Lambda_2 - \Lambda_3} + 3 + \Lambda_2 2^{-\tau} + \Lambda_3 2^{-4\tau} x$. Notice that the rounded CTR_2 in θ_2 is always greater than the one in θ_1 , given $\Lambda_2 \geq \Lambda_3$, while $CTR_2(\theta_1) = \Lambda_3 2^{-2\tau} > \Lambda_2 2^{-(4+\phi)\tau} = CTR_2(\theta_2)$ when $\frac{\Lambda_2}{\Lambda_3} < 2^{2\tau - \phi\tau}$.

We have that $\widetilde{SW}(\theta_1) > \widetilde{SW}(\theta_2)$ when $x > 2^{2\tau} \frac{\Lambda_2 - \Lambda_3 2^{-4\tau}}{\Lambda_2 - \Lambda_3}$. Thus a_2 gets a lower CTR by increasing her bid, which proves that the algorithm is not monotone.

The example can be extended also to the case where Approximation 1 is applied introducing ads with $qv = 0$ and $\gamma_{i,j} = 0$, s.t. $\log N = K$.

5 $\text{FNE}_{aa}(K)$ -r is APX-complete

In this section we will prove the APX-hardness of $\text{FNE}_{aa}(K)$ -r and provide a $1/2$ -approximation algorithm.

Theorem 3. *The complexity of computing an allocation for $\text{FNE}_{aa}(K)$ -r is APX-hard. More specifically, $\text{FNE}_{aa}(K)$ -r cannot be approximated within a factor of $\frac{1}{1+\alpha}$, for $\alpha < \frac{1}{412}$, unless $P = NP$.*

Proof. We reduce from the Asymmetric TSP with weights in $\{1, 2\}$, hereinafter denoted as $\text{ATSP}(1, 2)$. The $\text{ATSP}(1, 2)$ problem demands finding a minimum cost Hamiltonian tour in a complete directed weighted graph $G' = (T, A)$ where T is the set of nodes of G' , A is the set of edges and the weight function $w_{i,j} \in \{1, 2\}$ for all edges $(i, j) \in A$. $\text{ATSP}(1, 2)$ cannot be approximated in polynomial time within a factor of $\frac{1}{1+\beta}$, with $\beta < 1/206$ [9]. Below, we denote as τ a solution of an $\text{ATSP}(1, 2)$ instance, as $\text{cost}(\tau)$ its cost and as τ^* the optimal tour.

Given an instance of $\text{ATSP}(1, 2)$ on graph $G' = (T, A)$ we construct an instance of $\text{FNE}_{aa}(K)$ -r as follows: (i) for each vertex $t_i \in T$ we generate an ad a_i with $q_i = v_i = 1$, then we have $N = |T|$; (ii) the contextual graph is $G = ([N], \mathcal{E})$, where $(i, j) \in \mathcal{E}$ iff $w_{i,j} = 1$; (iii) for all $(i, j) \in \mathcal{E}$, $\gamma_{i,j} = 1$; and finally (iv) the number of slots is equal to the cost of the optimal tour τ^* in $\text{ATSP}(1, 2)$, i.e. $K = \text{cost}(\tau^*)$. We will show at the end of the proof how we can deal with the fact that we do not know $\text{cost}(\tau^*)$. Observe that with $K = \text{cost}(\tau^*)$, we have $\text{SW}(\theta^*) = N$, θ^* denoting the optimal solution of the $\text{FNE}_{aa}(K)$ -r instance constructed. The definition of the reduction is completed by observing that an allocation θ for the $\text{FNE}_{aa}(K)$ -r that allocates all the N ads can be easily mapped back to a tour τ for the $\text{ATSP}(1, 2)$ by simply substituting the ad with the corresponding vertex of the graph G' .

Let us suppose for the sake of contradiction that there exists a $\frac{1}{1+\alpha}$ -approximate algorithm for $\text{FNE}_{aa}(K)$ -r, with $\alpha < \frac{\beta}{2} < \frac{1}{412}$. Let θ_α be the $\frac{1}{1+\alpha}$ -approximate solution returned by such an algorithm, i.e., $\text{SW}(\theta_\alpha) \geq \frac{1}{1+\alpha} \text{SW}(\theta^*) = \frac{N}{1+\alpha}$. It is easy to check that θ_α consists of $\lceil \frac{N}{1+\alpha} \rceil$ ads, each providing a contribution of 1 to the social welfare, while there are $\text{SW}(\theta^*) - \lceil \frac{N}{1+\alpha} \rceil$ ads that w.l.o.g. we can consider empty. Moreover, being $\alpha < 1$, $\frac{N}{1+\alpha} \geq \text{cost}(\tau^*) - \frac{N}{1+\alpha}$ holds. For the sake of conciseness, hereinafter we omit the ceiling notation. Let τ_β be the tour obtained from θ_α . We state that in τ_β there are, at least, $\frac{2N}{1+\alpha} - \text{cost}(\tau^*) - 1$ edges of weight 1. Divide the ads allocated in θ_α in two sets: the $\frac{N}{1+\alpha}$ allocated ads a_i $i \in [N]$ and a_\perp . Allocate in alternation one of the $\frac{N}{1+\alpha}$ ads a_i , with $i \in [N]$, and one of the $\text{cost}(\tau^*) - \frac{N}{1+\alpha}$ ads a_\perp . When the slot index $2(\text{cost}(\tau^*) - \frac{N}{1+\alpha})$ is reached, the available a_\perp are finished, thus, in the following $\text{cost}(\tau^*) - 2(\text{cost}(\tau^*) - \frac{N}{1+\alpha}) = \frac{2N}{1+\alpha} - \text{cost}(\tau^*)$ slots, only non-fictitious ads a_i , $i \in [N]$, are consecutively allocated (no slots are left empty). This means that in θ_α , where the ads are disposed in a different way, we still have the guarantee that there are $\frac{2N}{1+\alpha} - \text{cost}(\tau^*) - 1$ pairs of consecutive ads (a_i, a_j) s.t. $\gamma_{i,j} = 1$. Thus, in the tour τ_β there are, at least, $\frac{2N}{1+\alpha} - \text{cost}(\tau^*) - 1$ edges of weight 1. Therefore, given that a tour is composed of N edges, in τ_β there can be at most $N - \frac{2N}{1+\alpha} + \text{cost}(\tau^*) + 1$ edges of weight 2. The length of τ_β is upper-bounded by $\text{cost}(\tau_\beta) \leq \frac{2N}{1+\alpha} - \text{cost}(\tau^*) - 1 + 2(N - \frac{2N}{1+\alpha} + \text{cost}(\tau^*) + 1) = \text{cost}(\tau^*) + \frac{2N\alpha}{1+\alpha} + 1$. Now we can state: $\text{cost}(\tau_\beta) \leq \text{cost}(\tau^*) + \frac{2\alpha N}{1+\alpha} + 1 \leq \text{cost}(\tau^*) + 2\alpha N \leq \text{cost}(\tau^*) + 2\alpha \text{cost}(\tau^*) = (1 + 2\alpha) \text{cost}(\tau^*) < (1 + \beta) \text{cost}(\tau^*)$, where: (i) the second inequality holds for $N \geq \frac{1+\alpha}{2\alpha^2}$; (ii) the third inequality holds since $N \leq \text{cost}(\tau^*)$ and (iii) the last inequality holds since, by assumption, $\alpha < \frac{\beta}{2}$. Thus, for the instances where $N \geq \frac{1+\alpha}{2\alpha^2}$ if there were an algorithm that $\frac{1}{1+\alpha}$ -approximates $\text{FNE}_{aa}(K)$ -r with $\alpha < \frac{1}{412}$, there would be a $\frac{1}{1+\beta}$ approximation of $\text{ATSP}(1, 2)$ with $\beta < \frac{1}{206}$. We obtained an absurd.

We finally show that we can deal with the non existence of the oracle returning $\text{cost}(\tau^*)$. For all the instances of $\text{ATSP}(1, 2)$ with N vertices, $N \leq \text{cost}(\tau^*) \leq 2N$. So, we run the poly-

mial $\frac{1}{1+\alpha}$ -approximation algorithm of $\text{FNE}_{aa}(K)$ -r for all the values $K = m$ with $m \in \{N \dots, 2N\}$, obtain m tours τ_β^m and set $\tau_\beta = \arg \min_{m \in \{N, \dots, 2N\}} \text{cost}(\tau_\beta^m)$, guaranteeing

$$\text{cost}(\tau_\beta) \leq \text{cost}(\tau_\beta^{\text{cost}(\tau^*)}). \quad \square$$

$\frac{1}{2}$ -**Approximate Greedy Algorithm for $\text{FNE}_{aa}(c)$ -r, for any c .** The algorithm orders the ads in nonincreasing order of $q_i v_i$ and allocates them in the odd slots, starting from the one with the highest product; even slots are left empty.

Proposition 5. *The greedy algorithm above is $\frac{1}{2}$ -approximate for $\text{FNE}_{aa}(c)$ -r, for any c .*

The greedy algorithm above is a MIR, range Θ' being all the allocations that leave even slots empty. The solution output is indeed the one guaranteeing maximum social welfare in Θ' . We therefore have proved the existence of a $1/2$ -approximate truthful polynomial-time mechanism for $\text{FNE}_{aa}(c)$ -r.

6 $\text{FNE}_{aa}(c)$ is APX-hard

We now prove that $\text{FNE}_{aa}(1)$ -r (Proposition 6) and $\text{FNE}_{aa}(1)$ -nr (Proposition 7) are APX-hard. First we state two auxiliary lemmata. Hereinafter, for the sake of notation, we will denote as $SW_1(\theta)$ and $SW_K(\theta)$ the objective function of \mathcal{B} - $\text{FNE}_{aa}(1)$ -r and \mathcal{B} - $\text{FNE}_{aa}(K)$ -r, respectively.

Lemma 1. *Let θ be an allocation (possibly containing empty slots) and let θ' be the allocation obtained from θ by replacing, for each pair (a_{i-1}, a_i) in θ such that $\gamma_{i-1,i} = 0$, ad a_{i-1} with a_\perp . Then $SW_1(\theta) = SW_1(\theta')$.*

Lemma 2. *Let θ be an allocation such that no pair of ads (a_{i-1}, a_i) exists where $\gamma_{i-1,i} = 0$. Then $SW_1(\theta) = SW_K(\theta)$.*

Proposition 6. *$\text{FNE}_{aa}(1)$ -r is APX-hard.*

Proof. We prove that the subproblem \mathcal{B} - $\text{FNE}_{aa}(1)$ -r is APX-hard via an approximation preserving reduction from the APX-hard problem \mathcal{B} - $\text{FNE}_{aa}(K)$ -r (Theorem 3). In particular, we will show that computing an approximate solution for \mathcal{B} - $\text{FNE}_{aa}(1)$ -r is not easier than \mathcal{B} - $\text{FNE}_{aa}(K)$ -r on the same instance.

We will first prove that $SW_K(\theta_K^*) \leq SW_1(\theta_1^*)$ holds, where θ_K^* and θ_1^* denote, respectively, the optimal allocation for \mathcal{B} - $\text{FNE}_{aa}(K)$ -r and \mathcal{B} - $\text{FNE}_{aa}(1)$ -r. For the sake of contradiction, let us suppose that $SW_K(\theta_K^*) > SW_1(\theta_1^*)$. We can assume without loss of generality that θ_K^* does not contain a pair (a_{i-1}, a_i) such that $\gamma_{i-1,i} = 0$, as replacing a_{i-1} with a_\perp would yield an allocation with a non-decreasing SW value. By Lemma 2 and by hypothesis we have that $SW_1(\theta_K^*) = SW_K(\theta_K^*) > SW_1(\theta_1^*)$, which contradicts the optimality of θ_1^* .

We are now going to prove that given an α -approximate solution θ_1^α to the objective of \mathcal{B} - $\text{FNE}_{aa}(1)$ -r we can compute in polynomial time an approximate solution θ_K^α to the objective of \mathcal{B} - $\text{FNE}_{aa}(K)$ -r such that $SW_1(\theta_1^\alpha) \leq SW_K(\theta_K^\alpha)$. This is easily done by replacing a_{i-1} with a_\perp for each couple of ads (a_{i-1}, a_i) in θ_1^α such that $\gamma_{i-1,i} = 0$, thus obtaining θ_1^α . By Lemmata 1 and 2 we finally conclude that $SW_1(\theta_1^\alpha) = SW_1(\theta_1^\alpha) = SW_K(\theta_1^\alpha)$. \square

Proposition 7. *$\text{FNE}_{aa}(1)$ -nr is APX-hard.*

Proof. We conduct the proof by reduction from problem \mathcal{B} - $\text{FNE}_{aa}(1)$ -r. In particular, we add to the instance of \mathcal{B} - $\text{FNE}_{aa}(1)$ -r K new ads $\{a_{N+1}, \dots, a_{N+K}\}$ such that: (i) $v_j = 0$ for all $j \in \{N+1, \dots, N+K\}$ and (ii) $\gamma_{i,j} = \gamma_{j,i} = 1$ for all $i \in \{1, \dots, N+K\}$ and $j \in \{N+1, \dots, N+K\}$. Let θ_{nr}^α be an α -approximate solution for the so-defined $\text{FNE}_{aa}(1)$ -nr problem. We can assume w.l.o.g. that θ_{nr}^α does not contain any a_\perp , as in the no-reset model we can always allocate any non-allocated ad to an empty slot obtaining a non-decreasing SW value. We observe that, from a generic allocation θ_{nr} , it is possible to obtain an allocation θ_r by substituting any ad a_j , $j \in \{N+1, \dots, N+K\}$, in θ_{nr} with a_\perp s.t. $SW^r(\theta_r) = SW^{nr}(\theta_{nr})$, and vice versa. Thus, from θ_{nr}^α we can obtain an allocation θ_r^α s.t. $SW^r(\theta_r^\alpha) = SW^{nr}(\theta_{nr}^\alpha)$; $SW^x(\theta)$ denoting the social welfare of $\theta \in \Theta$ in the model with reset $x \in \{r, nr\}$. Furthermore, let θ_r^* and θ_{nr}^* be the optimal solutions, respectively, for \mathcal{B} - $\text{FNE}_{aa}(1)$ -r and the $\text{FNE}_{aa}(1)$ -nr defined by our reduction. According to the observations above, it is easy to check that $SW^r(\theta_r^*) = SW^{nr}(\theta_{nr}^*)$ holds. In fact, let $\tilde{\theta}_{nr}$ be the solution obtained from θ_r^* by substituting each a_\perp with an ad a_j , $j \in \{N+1, \dots, N+K\}$. Then $SW^r(\theta_r^*) = SW^{nr}(\tilde{\theta}_{nr})$. Furthermore, $SW^{nr}(\tilde{\theta}_{nr}) = SW^{nr}(\theta_{nr}^*)$, as otherwise if $SW^{nr}(\tilde{\theta}_{nr}) < SW^{nr}(\theta_{nr}^*)$ we could translate θ_{nr}^* into a solution $\tilde{\theta}_r$ for \mathcal{B} - $\text{FNE}_{aa}(1)$ -r such that $SW^r(\theta_r^*) < SW^r(\tilde{\theta}_r)$. A similar argument holds if we consider the allocation $\tilde{\theta}_r$ obtained by substituting all ads a_j , $j \in \{N+1, \dots, N+K\}$, in θ_{nr}^* with a_\perp . Finally, $SW^r(\theta_r^\alpha) = SW^{nr}(\theta_{nr}^\alpha) \geq \alpha SW^{nr}(\theta_{nr}^*) = \alpha SW^r(\theta_r^*)$. \square

7 $\text{FNE}_{aa}^+(c)$ -nr is APX-complete for constant γ_{min}

We now prove that $\text{FNE}_{aa}^+(c)$ -nr is APX-complete if γ_{min} is constant. In the first part of this section (Theorem 4) we prove that this sub-problem is APX-hard, whereas in the second part we give a constant-ratio approximation algorithm for it, via a reduction from $\text{FNE}_{aa}^+(c)$ -nr to W3SP.

Theorem 4. *$\text{FNE}_{aa}^+(1)$ -nr is APX-hard.*

Proof. Let $\{\gamma_{min}, 1\}$ - $\text{FNE}_{aa}^+(1)$ -nr denote the subclass of $\text{FNE}_{aa}^+(1)$ -nr where $\gamma_{ij} \in \{\gamma_{min}, 1\}$ for all $i, j \in \mathcal{N}$ and a given $0 < \gamma_{min} < 1$. We prove the APX-hardness of $\text{FNE}_{aa}^+(1)$ -nr by an approximation preserving reduction from problem \mathcal{B} - $\text{FNE}_{aa}(1)$ -nr (proved APX-hard in Proposition 7) to problem $\{\gamma_{min}, 1\}$ - $\text{FNE}_{aa}^+(1)$ -nr: we prove the existence of an α -approximate algorithm for $\{\gamma_{min}, 1\}$ - $\text{FNE}_{aa}^+(1)$ -nr to imply the existence of a 2α -approximate algorithm for \mathcal{B} - $\text{FNE}_{aa}(1)$ -nr.

The instance of $\{\gamma_{min}, 1\}$ - $\text{FNE}_{aa}^+(1)$ -nr is obtained from the instance of \mathcal{B} - $\text{FNE}_{aa}(1)$ -nr by simply setting $\gamma'_{i,j} = \gamma_{min} = \frac{1}{K-1}$ for all $i, j \in \mathcal{N}$ such that $\gamma_{i,j} = 0$ in the given instance of \mathcal{B} - $\text{FNE}_{aa}(1)$ -nr, $\gamma'_{i,j} = 1$ otherwise.

Let $\theta_{\gamma_{min}}^*$ and $\theta_{\mathcal{B}}^*$ be an optimal solution for problems $\{\gamma_{min}, 1\}$ - $\text{FNE}_{aa}^+(1)$ -nr and \mathcal{B} - $\text{FNE}_{aa}(1)$ -nr, respectively. We have $SW(\theta_{\mathcal{B}}^*) \leq SW(\theta_{\gamma_{min}}^*)$. Indeed, if there is no $(a_{i-1}, a_i) \in \theta_{\mathcal{B}}^*$ s.t. $\gamma_{i-1,i} = 0$ then $SW(\theta_{\mathcal{B}}^*) = SW(\theta_{\gamma_{min}}^*)$, whereas if there is a pair $(a_{i-1}, a_i) \in \theta_{\mathcal{B}}^*$ s.t. $\gamma_{i-1,i} = 0$ then $SW(\theta_{\mathcal{B}}^*) < SW(\theta_{\gamma_{min}}^*)$.

Let now $\theta_{\gamma_{min}}$ be an α -approximation of $\{\gamma_{min}, 1\}$ - $\text{FNE}_{aa}^+(1)$ -nr and let $\theta_{\mathcal{B}}$ be the corresponding solution for \mathcal{B} - $\text{FNE}_{aa}(1)$ -nr. (I.e., $\theta_{\mathcal{B}}$ is the solution $\theta_{\gamma_{min}}$ where the γ_{min} externalities weigh 0.) We now prove that $SW(\theta_{\gamma_{min}}) \leq 2SW(\theta_{\mathcal{B}})$. We have $SW(\theta_{\mathcal{B}}) = 1 + \mathcal{P}(\theta_{\mathcal{B}})$, where $\mathcal{P}(\theta_{\mathcal{B}}) \leq K - 1$ denotes the number of pairs (a_{i-1}, a_i) of ads in $\theta_{\mathcal{B}}$ such that $\gamma_{i-1,i} = 1$. Likewise, $SW(\theta_{\gamma_{min}}) = 1 + \mathcal{P}(\theta_{\gamma_{min}}) + (K - 1 - \mathcal{P}(\theta_{\gamma_{min}})) \cdot \gamma_{min}$. By

construction, $\mathcal{P}(\theta_B) = \mathcal{P}(\theta_{\gamma_{min}}) = \mathcal{P}$, from which it follows that $SW(\theta_{\gamma_{min}}) \leq 2 \cdot SW(\theta_B)$ is equivalent to $1 + \frac{K-1-P}{1+P} \gamma_{min} \leq 2$. This is proved by noticing that $1 + \frac{K-1-P}{1+P} \gamma_{min} \leq 1 + \frac{K-1}{1+P} \gamma_{min} = \frac{P+2}{P+1}$, where last equality follows from definition of γ_{min} . \square

Approximation algorithm. We now prove that any α -approximate algorithm for Weighted 3-Set Packing (W3SP) can be turned into an $(\alpha\gamma_{min}^c)$ -approximation algorithm for $FNE_{aa}^+(c)$ -nr.

Given a universe U and a collection of its subsets each of cardinality at most 3 and associated to a weight, W3SP consists of finding a sub-collection of pairwise-disjoint subsets of maximal weight. Several constant-ratio approximate algorithms are known in literature to solve this problem, e.g., the algorithm in [4] provides a $1/2$ -approximation. We now present a reduction from $FNE_{aa}^+(c)$ -nr to W3SP, similar in spirit to that defined, for positive only externalities, in [6].

Theorem 5. *Given an α -approximate algorithm for problem W3SP, we can obtain an $(\alpha\gamma_{min}^c)$ -approximation algorithm for problem $FNE_{aa}^+(c)$ -nr.*

Proof. Given an instance of $FNE_{aa}^+(c)$ -nr, we obtain an instance of W3SP by means of the following reduction. To simplify the presentation, we suppose that K is even (the proof can be easily extended for an odd K). We divide K into $K/2$ blocks of two slots each. We construct a collection of $\frac{K}{2} \cdot \binom{N}{2}$ sets, each set having the form $\{a_i, a_j, p\}$, where $p \in \{1, 3, 5, \dots, K-1\}$ and $i, j \in \mathcal{N}$. The weight of a set is defined as the maximum social welfare that ads a_i and a_j can provide when assigned to slots s_p and s_{p+1} without taking into considerations the externalities of a_i and a_j on the ads allocated to slots s_m , $m \neq p, p+1$. Specifically, $W(a_i, a_j, p) = \max\{\Lambda_p q_i v_i + \Lambda_{p+1} \gamma_{i,j} q_j v_j, \Lambda_p q_j v_j + \Lambda_{p+1} \gamma_{j,i} q_i v_i\}$. Note that there is an immediate mapping between solutions of W3SP and $FNE_{aa}^+(c)$ -nr. For a solution θ_S of W3SP, let $W(\theta_S)$ denote its total weight. Now, let θ_S^* and θ^* denote, respectively, an optimal allocation for W3SP and an optimal allocation for $FNE_{aa}^+(c)$ -nr. Furthermore, let θ_S^α be an α -approximate solution for W3SP, and θ^α be the corresponding solution to $FNE_{aa}^+(c)$ -nr. Since in W3SP, outer-block externalities are not taken into consideration, we have: $W(\theta_S^*) \geq SW(\theta^*)$ and $SW(\theta^\alpha) \geq \gamma_{min}^c W(\theta_S^\alpha)$. From these inequalities we obtain:

$$SW(\theta^\alpha) \geq \gamma_{min}^c W(\theta_S^\alpha) \geq \alpha \gamma_{min}^c W(\theta_S^*) \geq \alpha \gamma_{min}^c SW(\theta^*). \quad \square$$

Corollary 1. *If γ_{min} is bounded from below by a constant (i.e., $\gamma_{min} \in \Omega(1)$), then $FNE_{aa}^+(c)$ -nr is approximable within a constant factor.*

It can be easily shown that the above algorithm is not monotone.

Theorem 6. *The algorithm of Theorem 5 is not monotone*

Proof. Consider an instance I of $FNE_{aa}^+(1)$ -nr with $N = K = 4$ wherein $\Lambda_3 \gamma_{z,4} < \Lambda_4 \gamma_{3,4}$, for $z \in \{1, 2\}$, $v_1, v_2 \gg v_3, v_4$ and $\gamma_{1,2} = \gamma_{2,1} = 1$ so that $W(a_1, a_2, 1)$ is much bigger than any other $W(a_i, a_j, 1)$. Therefore, any reasonable approximation of the W3SP instance constructed upon I must return sets $\{a_1, a_2, 1\}$ and $\{a_3, a_4, 3\}$. Additionally consider $v_4 < \frac{\Lambda_4 \gamma_{4,3}}{\Lambda_3^2 - \Lambda_3 \Lambda_4 \gamma_{3,4}}$ so that $W(a_3, a_4, 3) = \Lambda_3 q_3 v_3 + \Lambda_4 \gamma_{3,4} q_4 v_4$. So the solution θ returned by the algorithm run on I places a_4 in s_4 , resulting in $CTR_4(\theta) = q_4 \Lambda_4 \gamma_{3,4}$. Take now the instance I' defined as I except that $v_1, v_2 \gg v_4' > \frac{\Lambda_4 \gamma_{4,3}}{\Lambda_3^2 - \Lambda_3 \Lambda_4 \gamma_{3,4}} > v_4$. As before, the approximation algorithm

for W3SP will return sets $\{a_1, a_2, 1\}$ and $\{a_3, a_4, 3\}$ but this time $W'(a_3, a_4, 3) = \Lambda_3 q_4 v_4 + \Lambda_4 \gamma_{4,3} q_3 v_3$. Therefore, the solution θ' returned by the algorithm run on I' places a_4 in slot s_3 , i.e., $CTR_4(\theta') = q_4 \Lambda_3 \gamma_{z,4}$, where $z \in \{1, 2\}$ is the ad placed in slot s_2 in the allocation θ' . The algorithm is therefore not monotone and cannot be used to design a truthful mechanism. \square

8 Approximating $FNE_{aa}(c)$ -nr

Similarly to the case $c = K$, Color Coding can be applied to design an optimal exponential-time algorithm finding the optimal solution and a simple modification of such algorithm returns a $\frac{\log(N)}{2^{\min\{N, K\}}}$ approximation in polynomial time. While the basic idea is the same, some details change here.

We denote by $S \subseteq C$ a subset of colors and by $\delta(a)$ a function returning the color assigned to a . Given a coloring δ , the best colorful allocation is found by dynamic programming. For $|S| > c$, $W(S, \langle a_{h_0}, \dots, a_{h_c} \rangle)$ contains the value of the best allocation with colors in S in which the last $c+1$ ads are a_{h_0}, \dots, a_{h_c} from top to bottom. (The definition naturally extends for $|S| \leq c$.) Starting from $W(\emptyset, \langle \rangle) = 0$, we can compute W recursively. For instance, for $|S| > c$, $W(S \cup \{\delta(a_{h_c})\}, \langle a_{h_0}, \dots, a_{h_c} \rangle) = \Lambda_{|S|+1} v_{h_c} q_{h_c} \prod_{i=0}^{c-1} \gamma_{h_i, h_{i+1}} + \max_a W(S, \langle a, a_{h_0}, \dots, a_{h_{c-1}} \rangle)$ if $\delta(a_{h_c}) \notin S$ and $-\infty$ otherwise. Given a random coloring, the probability that the ads composing the best allocation are colorful is $\frac{K!}{K^K}$. Thus, repeating the procedure re^K times, where $r \geq 1$, the probability of finding the best allocation is $1 - e^{-r}$. The complexity is $\mathcal{O}((2e)^K K N^{c+2})$. The algorithm can be derandomized with an additional cost of $\mathcal{O}(\log^2(N))$.

By applying the above algorithm to the first K' slots, $K' = \min\{K, \lceil \log(N) \rceil\}$, we obtain an algorithm with complexity $\mathcal{O}(K^{3.5} N^{c+2} \log^2(N))$. We observe that if c is not a constant, the complexity is exponential. It is not too hard to note that such an algorithm is $\frac{\log(N)}{2^{\min\{N, K\}}}$ -approximate. Moreover, this algorithm is MIR and as such can be used to design a truthful mechanism.

9 Conclusions

We enrich the literature on externalities in SSAs by introducing more general ways to model slot- and ad-dependent externalities, while giving a (nearly) complete picture of the computational complexity of the problem. In detail, we enrich the naive model of SSAs by adding: (i) the concepts of limited user memory (ii) contextual externalities and (iii) refreshable user memory (i.e., reset model).

This gives rise to the FNE_{sa} model, where ad- and slot-dependent externalities are factorized as in the cascade model and the FNE_{aa} model, where the externalities are not factorized.

We satisfactorily solve the problem for FNE_{sa} , whereas our results leave unanswered a number of interesting questions, with regards to both approximation and truthfulness for FNE_{aa} . The parameter c is central to this list. If c is constant, then we do not know whether a constant approximation algorithm for $FNE_{aa}(c)$ exists; this holds also for the special case of $FNE_{aa}^+(c)$ -nr when γ_{min} is not a constant. In the latter case, when γ_{min} is instead constant we are not aware of any truthful constant approximation mechanism. Motivated by the fact that FNE_{aa-r} is, apparently, an easier problem than FNE_{aa-nr} , we believe that an interesting direction for future research is to study reset in more detail in order to understand its role w.r.t. the relatively harder FNE_{aa-nr} .

REFERENCES

- [1] G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pál, ‘Sponsored search auctions with markovian users’, in *WINE*, pp. 621–628, (2008).
- [2] N. Alon, R. Yuster, and U. Zwick, ‘Color-coding’, *J. ACM*, **42**(4), 844–856, (1995).
- [3] A. Archer and É Tardos, ‘Truthful mechanisms for one-parameter agents’, in *FOCS*, pp. 482–491, (2001).
- [4] P. Berman, ‘A $d/2$ approximation for maximum weight independent set in d -claw free graphs’, *Nord. J. Comput.*, **7**(3), 178–184, (2000).
- [5] A. Björklund, T. Husfeldt, and S. Khanna, ‘Approximating longest directed paths and cycles’, in *ICALP*, pp. 222–233, (2004).
- [6] D. Fotakis, P. Krysta, and O. Telelis, ‘Externalities among advertisers in sponsored search’, in *SAGT*, pp. 105–116, (2011).
- [7] N. Gatti, M. Rocco, P. Serafino, and C. Ventre, ‘Towards Better Models of Externalities in Sponsored Search Auctions’, *ArXiv e-prints*, (April 2016).
- [8] P. Jeziorski and I. Segal, ‘What makes them click: Empirical analysis of consumer demand for search advertising’, *American Economic Journal: Microeconomics*, (2014).
- [9] M. Karpinski, M. Lampis, and R. Schmied, ‘New inapproximability bounds for tsp’, in *ISAAC*, pp. 568–578, (2013).
- [10] D. Kempe and M. Mahdian, ‘A cascade model for externalities in sponsored search’, in *WINE*, pp. 585–596, (2008).
- [11] P. Metrikov, F. Diaz, S. Lahaie, and J. Rao, ‘Whole page optimization: how page elements interact with the position auction’, in *EC*, pp. 583–600, (2014).
- [12] N. Nisan and A. Ronen, ‘Computationally feasible VCG mechanisms’, volume 29, pp. 19–47, (2007).