

Equilibrium Analysis of Ad Auctions with Price Displaying

Matteo Castiglioni
Politecnico di Milano
Milano, Italy
matteo.castiglioni@polimi.it

Diodato Ferraioli
Università degli Studi di Salerno
Salerno, Italy
dferraioli@unisa.it

Nicola Gatti
Politecnico di Milano
Milano, Italy
nicola.gatti@polimi.it

Alberto Marchesi
Politecnico di Milano
Milano, Italy
alberto.marchesi@polimi.it

Giulia Romano
Politecnico di Milano
Milano, Italy
giulia.romano@polimi.it

ABSTRACT

Most of the economic reports forecast that almost half of the world-wide market value unlocked by AI over the next decade (up to 6 trillion USD per year) will be in *marketing&sales*. In particular, AI will enable the optimization of more and more intricate economic settings, in which multiple different activities need to be jointly automated. This is the case of, e.g., *Google Hotel Ads* and *Tripadvisor*, where auctions are used to display ads of similar products or services together with their prices. As in classical ad auctions, the ads are ranked depending on the advertisers' bids, whereas, differently from classical settings, ads are displayed together with their prices, so as to provide a direct comparison among them. This dramatically affects users' behavior, as well as the properties of ad auctions. We show that, in such settings, social welfare maximization can be achieved by means of a direct-revelation mechanism that jointly optimizes, in polynomial time, the ads allocation and the advertisers' prices to be displayed with them. However, in practice it is unlikely that advertisers allow the mechanism to choose prices on their behalf. Indeed, in commonly-adopted mechanisms, ads allocation and price optimization are decoupled, so that the advertisers optimize prices and bids, while the mechanism does so for the allocation, once prices and bids are given. We investigate how this decoupling affects the efficiency of mechanisms. In particular, we study the *Price of Anarchy* (PoA) and the *Price of Stability* (PoS) of indirect-revelation mechanisms with both VCG and GSP payments, showing that the PoS for the revenue may be unbounded even with two slots, and the PoA for the social welfare may be as large as the number of slots. Nevertheless, we show that, under some assumptions, simple modifications to the indirect-revelation mechanism with VCG payments achieve a PoS of 1 for the revenue.

KEYWORDS

Ad Auctions; Equilibrium Computation; Mechanism Design

1 INTRODUCTION

Most of the economic reports forecast that *artificial intelligence* (AI) will unlock up to 12 trillion USD per year worldwide by the next decade, and almost half of this amount will derive from the *marketing&sales* area (see, e.g., [8]). In particular, AI is playing a crucial role to tackle various problems, including, e.g., auction

design [2, 3, 7, 25], the automation of advertisers' budget [6, 24] and bidding strategies [16], and the optimization of conversion funnels [23].

In this paper, we focus on recently-emerged online advertising settings where ad auctions are employed to display ads of similar products or services together with their prices. This is the case of, e.g., *Google Hotel Ads* and *Tripadvisor*, where users search for the availability of a hotel room in a given date. The web page of results shows a ranking of banners advertising similar hotel rooms that match the search criteria. Each banner displays the name of the advertiser providing the online booking service, together with the per-night selling price of the room. Such settings are similar to standard ad auctions, since the ads are ranked depending on the advertisers' bids. On the other hand, they also fundamentally differ from them, as the ad allocation must also take prices into account, and these are displayed inside the banners so as to provide a direct comparison among them. This dramatically affects users' behavior, as well as the efficiency and the properties of ad auctions. The goal of this work is to investigate how the additional degree of freedom introduced by prices influences the problem of finding an optimal ad allocation and the revenue of the mechanisms.

The price-displaying feature of our setting introduces *externalities among the ads*, since the probability that a user clicks on an ad depends on the prices displayed with both the ad being clicked and the other ads in the allocation. Several forms of externalities are investigated in the literature on ad auctions. However, to the best of our knowledge, no previous work takes into account price displaying. For instance, Kempe and Mahdian [18] and Aggarwal et al. [1] introduce a basic user model that is currently adopted by most of the mechanisms. In this model, a Markovian user observes the slots in a top-down fashion, moving down slot by slot with a given continuation probability and stopping on a slot to observe its ad with the remaining probability. Kempe and Mahdian [18] also propose richer models where the probability with which a user moves from a slot to the next one depends on the ad actually displayed in the former. In this case, it is *not* known whether the ad allocation problem admits a polynomial-time algorithm; however, Farina and Gatti [9, 10] provide several algorithms showing that in special cases a constant approximation can be achieved. Further externalities models are explored by Fotakis et al. [11] and Gatti et al. [13], which allow for potentially different externalities for each pair of ads. However, with these models, the ad allocation problem is NP-hard and, in some cases, even inapproximable. It is

also worth mentioning that similar models are adopted in mobile geo-located advertising by Gatti et al. [12].

In our model, we assume that the probability with which a user clicks on an ad depends on the price displayed with the ad *and* on the lowest among all displayed prices. In particular, we model the click probability as a monotonically decreasing function of the ad price, assuming that the demand curve is monotonically decreasing in the price and that it is unlikely that a user clicks on an ad with a price larger than her reserve value. We also assume that the click probability is monotonically decreasing in the difference between the ad price and the lowest displayed price, as the user’s interest in any feature different from price (e.g., brand and loyalty) decreases as such difference increases.

In our setting, the private information of each advertiser (i.e., her type) is a pair composed by the probability with which a user visiting the advertiser’s web page produces a conversion (e.g., a purchase) and the advertiser’s cost for a unit of product or service. On the other hand, the prices constitute an additional degree of freedom that can be controlled by either the advertisers or the mechanism.

As a first step, we present a direct-revelation mechanism that maximizes the social welfare by jointly optimizing over the ad allocations and the prices displayed with the ads. Differently from what happens in most of the externalities models studied in the literature, such optimization problem can be solved in polynomial time for a given discretization of price values. We also study the properties of the direct-revelation mechanism when VCG payments are used, showing that incentive compatibility, individual rationality and weak budget-balance hold in our setting.

In real-world scenarios, it is unlikely that the advertisers let the mechanism select prices on their behalf, as required by the direct-revelation mechanism. In the (indirect-revelation) mechanisms that are currently adopted in real-world applications, the optimization over ad allocations and that over prices are decoupled. In particular, each advertiser finds her optimal price and bid, while the mechanism optimizes over ad allocations once prices and bids are given. As for the direct-revelation mechanism, the best ad allocation can be found in polynomial time given prices and bids. However, even if these indirect-revelation mechanisms allow the advertisers not to reveal private (and potentially sensitive) information, they can lead to inefficient equilibria.

We investigate the equilibrium inefficiency of indirect-revelation mechanisms with GSP and VCG payments, in terms of *Price of Anarchy* (PoA) and *Price of Stability* (PoS) in complete information settings. In the literature, PoA and PoS are commonly-adopted efficiency metrics for standard ad auctions, in which the price variable is not taken into account. For instance, Caragiannis et al. [4], Lucier and Leme [19], and Caragiannis et al. [5] show that the PoA for the social welfare of the GSP is upper bounded by 1.3 with complete information and by 3 with incomplete information, while Farina and Gatti [10] and Giotis and Karlin [14] study the inefficiency with specific externalities. In our setting, the presence of externalities precludes the adoption of the tools provided by Roughgarden et al. [26] and Hartline et al. [15] to bound the inefficiency of equilibria for the social welfare and the revenue, respectively, thus pushing us to develop *ad hoc* approaches. In particular, we show that, in our setting, the inefficiency of the indirect-revelation mechanisms

with VCG and GSP mechanisms is much higher than that of the classical mechanisms without prices, even when excluding over-bidding, since the PoS for the revenue may be unbounded even with two slots and the PoA for the social welfare may be as large as the number of slots. Furthermore, with VCG payments, the PoS for the social welfare is 1, while, with GSP payments, it is at least 2, suggesting that GSP payments perform worse than VCG ones.

A crucial question is whether inefficiency can be reduced when letting the advertisers choose their prices. We show that, under some assumptions, simple modifications to the indirect-revelation mechanism with VCG payments—requiring each advertiser to report an additional price—achieve a PoS of 1 for the revenue.

2 FORMAL MODEL

There is a set $N = \{1, \dots, n\}$ of n agents, who simultaneously play the role of advertisers and sellers. Each agent sells a single good on her own website (e.g., an online marketplace) and relies on an external ad publisher that advertises the good through a single ad in which the price is displayed. Since the goods being sold by the agents are similar, the price comparison that users perform on the publisher’s website results in a high competition level among the agents, as happening in classical comparator websites [17]. In the following, for the ease of presentation, we use index $i \in N$ to refer to the agent, her good, and also her ad. Figure 1 provides an overview of our scenario.

For every i , we denote with $c_i \in \mathbb{R}_{\geq 0}$ and $p_i \in \mathbb{R}_{\geq 0}$ the *cost* of supply and the *selling price* of agent i ’s good, respectively. Furthermore, we denote with $\alpha_i \in [0, 1]$ the probability with which a user buys agent i ’s good when visiting her website. Thus, agent i ’s expected *gain* from a visit of a user on her website is $\alpha_i (p_i - c_i)$. Let us remark that the *conversion probability* α_i is constant w.r.t. the price p_i , since we assume that the user is aware of the price before visiting the website and, thus, she does not visit it if the price is larger than her reserve value. As previously discussed, the user first observes the ads on the publisher’s website, together with their




Figure 1: An example of ad auction with price displaying. A user visits a web page with three ads (ad 1, ad 2, and ad 3) together with their prices (p_1 , p_2 , and p_3). The user observes slot 3 with probability λ_3 . Once observed slot 3, the user clicks on the ad displayed in slot 3, i.e., ad 3, with probability $q_3(p_3, p_{\min})$ where p_{\min} is the minimum price among p_1, p_2, p_3 . The user visits the web page of advertiser 3 (e.g., an online marketplace), and, then, produces a conversion (e.g., purchase) with probability α_3 . The value that advertiser 3 gets from the conversion is $p_3 - c_3$.

prices, and, then, she clicks on an ad so as to visit the corresponding advertiser's website. Therefore, the motivation behind an uncompleted conversion following the user's visit to the advertiser's web page does *not* concern the price (e.g., it may be due to the user acquiring more information on the seller, or potential extra fees and/or ancillary services). The pair (α_i, c_i) is a private information of agent i , and sometimes we will refer to it as her type θ_i . We let $\Theta = [0, 1] \times \mathbb{R}_{\geq 0}$ be the space of types of every agent.

The ad publisher has a set $M = \{1, \dots, m\}$ of slots in which the ads are displayed. An *assignment* of ads to slots (also called *allocation*) is represented by a function $f: N \rightarrow M \cup \{\perp\}$ such that there is at most one ad per slot (i.e., there are no ads $i, h \in N$ such that $i \neq h$ and $f(i) = f(h) \in M$). All the ads that are not assigned to slots in M are assigned to \perp , meaning that these ads are not displayed. For every slot $j \in M$, we denote with $\lambda_j \in [0, 1]$ the probability (called *prominence*) that a user observes the ad displayed in that slot. As customary in the literature, we assume that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_m$. For the ease of notation, we define $\lambda_\perp = 0$. Furthermore, for every agent i , we denote with $q_i \in [0, 1]$ the probability (called *quality*) that a user clicks on ad i conditioned on its observation. In our setting, q_i depends on the prices, as they are displayed with the ads. In particular, q_i is a function of the prices $\mathbf{p} = \{p_i\}_{i \in N}$ of agents whose ads are displayed, since the user can compare all the prices shown on the web page when deciding the website from which to buy a good. This dependency introduces externalities among the ads. In this work, we assume that $q_i: \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow [0, 1]$, where $q_i(p_i, p_{\min})$ denotes the agent i 's quality when her price is p_i and the minimum price among all the displayed ads is p_{\min} , with $p_{\min} = \min_{h \in N: f(h) \in M} \{p_h\}$ (for the sake of notation, we omit the dependency of p_{\min} on f). Moreover, given p_{\min} , q_i is (non-strictly) monotonically decreasing in p_i since, as previously discussed, a user clicks on the ad if the price is non-larger than the user's reserve value. Finally, q_i is (non-strictly) monotonically increasing in p_{\min} , given p_i . The rationale behind this assumption is that, given p_i , the probability that a user clicks on ad i decreases as the gap between p_i and the minimum price p_{\min} increases, capturing a potential reduction of the user's interest for agent i 's good. A simple example is when the users are only interested in the price and, thus, q_i is zero if $p_i > p_{\min}$. We also assume that there exists $p_i \in \mathbb{R}_{\geq 0}$ maximizing $q_i(p_i, p_i) \alpha_i (p_i - c_i)$ and, thus, there exists $p_i < \infty$ that agent i would use when displayed alone. Finally, we remark that, as it is customary in the literature, parameters λ and q are estimated by the ad publisher.

Every *mechanism* receives some input (or *bid*) from every agent i , chooses an allocation f , and charges every agent i of a payment π_i . We say that the mechanism is *direct-revelation* if the input provided by agent i belongs to Θ , i.e., it consists of a conversion probability and a cost, which are *not* necessarily the real ones (her type). Otherwise we say that the mechanism is *indirect-revelation*.

In our setting, a direct-revelation mechanism takes as input a reported type $\theta'_i = (\alpha'_i, c'_i) \in \Theta$ for each agent i , and chooses some prices $\mathbf{p} = \{p_i\}_{i \in N}$ and an allocation function f . We let $\mathbf{b} = \{b_i\}_{i \in N}$ be the vector of declared gains, where $b_i = \alpha'_i (p_i - c'_i)$ is agent i 's gain for the reported type θ'_i . On the other hand, an indirect-revelation mechanism takes as input a price p_i and a declared gain b_i for each agent i , and chooses an allocation function f . We say

that agent i does not *overbid* if $b_i \leq \alpha_i (p_i - c_i)$, where p_i is the price given as input and $(\alpha_i, c_i) = \theta_i$ is the true agent i 's type.

Given an allocation f , prices \mathbf{p} , and b_i , we denote with $\widehat{v}_i(f, \mathbf{p}, b_i) = \lambda_{f(i)} q_i(p_i, p_{\min}) b_i$ the expected (w.r.t. clicks and purchase) *value* of agent i according to her declared gain. The true expected value that she receives from allocation f is $v_i = \lambda_{f(i)} q_i(p_i, p_{\min}) \alpha_i (p_i - c_i)$, while agent i 's expected *utility* is $u_i = v_i - \pi_i$ since the environment is quasi-linear.¹ The *social welfare* of an allocation with respect to the declared gains is $\widehat{SW}(f, \mathbf{p}, \mathbf{b}) = \sum_i \widehat{v}_i(f, \mathbf{p}, b_i)$, where $\mathbf{b} = \{b_i\}_{i \in N}$. The true social welfare is $SW = \sum_i v_i$. The *revenue* is instead $\text{Rev} = \sum_i \pi_i$.

We informally introduce notable properties of mechanisms; see [20] for formal definitions. A mechanism, both direct- and indirect-revelation, is *individually rational*, if for every agent i , the assigned payment π_i is non-larger than her value $\widehat{v}_i(f, \mathbf{p}, b_i)$ according to the declared gain. Furthermore, a mechanism is *weakly budget-balanced* if the sum of payments is always non-negative. A direct-revelation mechanism is *truthful* if for every agent i it is a dominant strategy to report the true type $\theta_i = (\alpha_i, c_i)$ to the mechanism, i.e., the utility that agent i achieves by reporting θ_i is at least as large as with every alternative input, regardless of other agents' actions. For indirect-revelation mechanisms, we say that a set of inputs is in *equilibrium* according to Nash [21] if no agent may increase her utility by submitting a different bid, whenever the inputs of other agents remain unchanged.

3 MECHANISMS

Next, we introduce our direct-revelation mechanism and two indirect-revelation mechanisms.

3.1 Direct-revelation Mechanism

We let $\mathcal{M}_D^{\text{VCG}}$ be the direct-revelation mechanism defined as follows. Given the agent i 's input $\theta'_i = (\alpha'_i, c'_i) \in \Theta$, the mechanism defines $b_i = \alpha'_i (p_i - c'_i)$ for every price p_i . Then, the mechanism computes an assignment f^* and prices \mathbf{p}^* that maximize the social welfare with respect to the declared gains; formally,

$$\widehat{SW}(f^*, \mathbf{p}^*, \mathbf{b}) = \max_{f, \mathbf{p}} \widehat{SW}(f, \mathbf{p}, \mathbf{b}).$$

Finally, the mechanism assigns to each advertiser i in the allocation (i.e., such that $f(i) \in M$) the VCG payment

$$\begin{aligned} \pi_i &= \max_{f, \mathbf{p}} \sum_{j \neq i: f(j) \in M} \left(\widehat{v}_j(f, \mathbf{p}, b_j) - \widehat{v}_j(f^*, \mathbf{p}^*, b_j) \right) \\ &= \widehat{v}_i(f^*, \mathbf{p}^*, b_i) - \Delta_i, \end{aligned}$$

where

$$\Delta_i = \widehat{SW}(f^*, \mathbf{p}^*, \mathbf{b}) - \max_{f, \mathbf{p}: f(i) \notin M} \widehat{SW}(f, \mathbf{p}, \mathbf{b}) \geq 0.$$

It is immediate to check that payments cannot be negative and they are never larger than the value corresponding to the declared gain. Thus, the mechanism is trivially individually-rational and weakly budget-balanced. Moreover, it is not hard to verify that these payments allow the mechanism to be truthful (essentially this is a VCG mechanism and there is no interdependence among types).

¹The dependency of v_i, u_i, π_i on the arguments f, \mathbf{p}, b_i is omitted to avoid cumbersome notation.

Truthfulness also implies that the mechanism maximizes the true social welfare. These observations prove the following theorem.

Theorem 1. *Mechanism $\mathcal{M}_D^{\text{VCG}}$ is truthful, individually rational, weakly budget-balanced, and maximizes SW.*

3.2 Indirect-revelation Mechanisms

Next, we introduce two alternative mechanisms, namely $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_1^{\text{GSP}}$. These mechanisms share the same structure, but they differ in the way they compute the payments. They work as follows. Agent i inputs (p_i, b_i) , where $p_i \in \mathbb{R}_{\geq 0}$ is the price that agent i wants to be displayed for her ad and $b_i \in \mathbb{R}$ is the expected gain that i declares to achieve from a click on her ad for price p_i . The mechanism computes an assignment g^* that maximizes the social welfare with respect to the submitted prices and gains; formally

$$\widehat{\text{SW}}(g^*, \mathbf{p}, \mathbf{b}) = \max_g \widehat{\text{SW}}(g, \mathbf{p}, \mathbf{b}).$$

Then, $\mathcal{M}_1^{\text{VCG}}$ assigns to each advertiser i such that $g^*(i) \in M$ the VCG payment

$$\begin{aligned} \pi_i &= \max_g \sum_{j \neq i: g(j) \in M} (\widehat{v}_j(g, \mathbf{p}, b_j) - \widehat{v}_j(g^*, \mathbf{p}, b_j)) \\ &= \widehat{v}_i(g^*, \mathbf{p}, b_j) - \delta_i, \end{aligned}$$

where

$$\delta_i = \widehat{\text{SW}}(g^*, \mathbf{p}, \mathbf{b}) - \max_{g: g(i) \notin M} \widehat{\text{SW}}(g, \mathbf{p}, \mathbf{b}) \geq 0.$$

W.l.o.g., let the optimal allocation g^* be such that only the first $\ell \leq m$ slots are assigned and no slot $j > \ell$ is assigned. $\mathcal{M}_1^{\text{GSP}}$ assigns to each i such that $g^*(i) \in M$ and $g^*(i) < \ell$ (i.e., i is assigned to a slot different from ℓ) the following payments:

$$\omega_i = \lambda_{g^*(i)} q_j(p_j, p_{\min}) b_j, \quad (1)$$

where j is such that $g^*(j) = g^*(i) + 1$. When $g^*(i) = \ell$, there are two possible payments. If all the not assigned agents j (i.e., such that $g^*(j) = \perp$) have a price $p_j < p_{\min}$, then $\omega_i = 0$. Otherwise, the payment is

$$\omega_i = \lambda_{g^*(i)} \max_{j: p_j \geq p_{\min} \wedge g^*(j) = \perp} \{q_j(p_j, p_{\min}) b_j\}. \quad (2)$$

As for $\mathcal{M}_D^{\text{VCG}}$, it is immediate to check that payments are at least zero, and they are always less than the value corresponding to the declared gain. Hence, $\mathcal{M}_1^{\text{VCG}}$ is individually rational and weakly budget-balanced. Moreover, one may hope that the inputs that agents select at any equilibrium are such that the allocation selected by the mechanism maximize the social welfare. Unfortunately, we will show in the next sections that this is *not* the case.

The payments of $\mathcal{M}_1^{\text{GSP}}$ are at least zero, and, thus, the mechanism is weakly budget-balanced. Also observe that, given agent i , $\forall j$ s.t. $g^*(j) > g^*(i)$ or $g^*(j) = \perp \wedge p_j \geq p_{\min}$, we have that $q_j(p_j, p_{\min}) b_j \leq q_i(p_i, p_{\min}) b_i$. Otherwise, the allocation g achieved from g^* by fixing $g(j) = g^*(i)$, $g(i) = g^*(j)$, and $g(k) = g^*(k) \forall k \notin \{i, j\}$ would achieve a larger social welfare (according to declared gains). Hence, we have that $\omega_i \leq \widehat{v}_i(g^*, \mathbf{p}, b_i)$, and, thus, the mechanism is individually rational. We remark that for this property to hold, it is fundamental that, in Equation 2, we consider only the not assigned agents j who have a declared price $p_j \geq p_{\min}$. Indeed, an agent j with $p_j < p_{\min}$ may have a large $q_j(p_j, p_j) b_j$ so

that, if the j -th ad is displayed, the minimum price changes from p_{\min} to p_j , $q_j(p_j, p_j) b_j > q_i(p_i, p_{\min}) b_i$, and $\omega_i > \widehat{v}_i(g^*, \mathbf{p}, b_i)$, where i is the agent assigned to the slot ℓ . Nevertheless, this agent may not be chosen by the allocation g^* because of the negative externalities that its low price would put on other agents (by lowering their value and thus the social welfare). As a result an optimal allocation may not assign all the available slots. We finally observe that, as for $\mathcal{M}_1^{\text{VCG}}$, even $\mathcal{M}_1^{\text{GSP}}$ may fail to optimize the true social welfare. The following sections will bound the extent of this failure.

4 COMPUTATIONAL COMPLEXITY

In general, externalities make hard the problem of computing the allocation maximizing the social welfare. In this section, we prove that in our setting the problem of allocating advertisers to slots can be solved in polynomial time by both the direct- and the indirect-revelation mechanisms.

Let us start with the problem of computing the allocation g^* in the indirect-revelation mechanisms. We show in the next theorem that g^* can be efficiently computed.

Theorem 2. *There is an algorithm that computes the allocation g^* in time $O(n^2 \log n)$.*

PROOF. Let \mathbf{b} and \mathbf{p} be the set of gains and prices submitted by agents. First observe that, given a minimum displayed price p_{\min} , the allocation that maximizes the social welfare (with respect to gains and prices in input), can be trivially computed by sorting agents in $\{i: p_i \geq p_{\min}\}$ in order of $q_i(p_i, p_{\min}) b_i$ and assigning slot 1 to the agent that maximizes this quantity, slot 2 to the second such agent, and so on. Note that this operation requires $O(n \log n)$ steps.

However, in order to provide the allocation g^* , we also need to decide which is the best value for p_{\min} . However, since p_{\min} must belong to \mathbf{p} , it is sufficient to compute the best allocation by using as minimum displayed price each of the at most n different prices in \mathbf{p} , and choosing the allocation that optimizes the social welfare. \square

Computing g^* is an easier problem than the one faced by the direct-revelation mechanism, since, for the former, prices are given and we optimize only over the allocation function, while, for the latter, optimization occurs both on the allocation function and prices. Nevertheless, the following theorem shows that f^* and \mathbf{p}^* can also be computed efficiently, as long as the set P of allowed prices is discrete and finite.

Theorem 3. *There is an algorithm that computes the allocation f^* and prices \mathbf{p}^* in time $O(n^2 |P| (|P| + \log n))$.*

PROOF. Let $b_i(p) = \alpha'_i(p - c'_i)$ be the expected gain of agent i according to her input when ad i is displayed with price p , where (α'_i, c'_i) is the input of agent i . For each agent i and every price $\hat{p} \in P$ we compute $p_i^*(\hat{p})$ as follows: if $\max_{p \in P: p \geq \hat{p}} q_i(p, \hat{p}) b_i(p) > 0$, then

$$p_i^*(\hat{p}) = \arg \max_{p \in P: p \geq \hat{p}} q_i(p, \hat{p}) b_i(p),$$

otherwise we set $p_i^*(\hat{p}) = \perp$. Roughly speaking, $p_i^*(\hat{p})$ is the best price (according to her input) for agent i when the minimum displayed price is \hat{p} and the i -th ad is displayed (and thus i 's price is at least \hat{p}). Clearly, if there is no price larger than or equal to \hat{p}

guaranteeing to agent i a positive utility, then she prefers to be not displayed. For this reason, in the latter case, we do not assign any value to $p_i^*(\hat{p})$. Notice that $p_i^*(\hat{p})$ can be computed by evaluating the function for every $p \in P$ with $p \geq \hat{p}$, requiring at most $O(|P|)$ operations.

Then, if the minimum displayed price p_{\min} was given, along with the agent to which it is assigned, then we simply choose price $p_i^*(p_{\min})$ for each remaining agent i (this can be done in $O(nP)$ steps), prune out agents for which $p_i^*(p_{\min}) = \perp$, and finally compute the corresponding optimal assignment by sorting the remaining agents in order of $b_i(p_i^*(p_{\min}))$, as shown in Theorem 2 (in $O(n \log n)$ steps).

Unfortunately, selecting p_{\min} is much harder than in the indirect case: not only the value of p_{\min} can assume every value in P (and not just one among at most n alternatives), but we also need to decide which agent should display this price. For this reason, we need to go through every price $p \in P$ and every agent i and compute the best solution that would be achieved when i is the agent displaying the minimum price p . Since for each of the nP possible choices, computing the best solution requires time $O(nP + n \log n)$, we achieve the desired running time. \square

Observe that the dependence on $|P|$ in Theorem 3 is in some way necessary as long as we would like to keep quality function as general as possible. It is not hard to see that we can avoid to check all prices by doing opportune restriction on the quality functions.

We finally highlight that the discretization of the set of prices does not affect the property of the mechanism. In particular, truthfulness continues to hold, since the mechanism is maximal-in-the-range.

5 PERFORMANCE OF THE INDIRECT MECHANISMS

For the sake of presentation, we provide the informal definitions of PoS and PoA for social welfare and revenue; formal definitions can be found in [22].

- PoS for the social welfare is the minimum—w.r.t. all the Nash equilibria—ratio between the maximum achievable social welfare and the social welfare of an allocation achievable in a Nash equilibrium of an indirect-revelation mechanism $\mathcal{M}_1^{\text{VCG}}$ or $\mathcal{M}_1^{\text{GSP}}$.
- PoA for the social welfare is the maximum—w.r.t. all the Nash equilibria—ratio between the maximum achievable social welfare and the social welfare of an allocation achievable in a Nash equilibrium of an indirect-revelation mechanism $\mathcal{M}_1^{\text{VCG}}$ or $\mathcal{M}_1^{\text{GSP}}$.
- PoS for the revenue is the minimum—w.r.t. all the Nash equilibria—ratio between the maximum revenue achievable by an individually-rational mechanism and the revenue achievable in a Nash equilibrium of an indirect-revelation mechanism $\mathcal{M}_1^{\text{VCG}}$ or $\mathcal{M}_1^{\text{GSP}}$.
- PoA for the revenue is the maximum—w.r.t. all the Nash equilibria—ratio between the maximum revenue achievable by an individually-rational mechanism and the revenue achievable in a Nash equilibrium of an indirect-revelation mechanism $\mathcal{M}_1^{\text{VCG}}$ or $\mathcal{M}_1^{\text{GSP}}$.

Table 1 summarizes the lower and upper bounds over the mechanisms' inefficiency when agents do not overbid; the results when agents overbid are omitted since the inefficiency can be arbitrary even with a single slot. Interestingly, while $\mathcal{M}_1^{\text{VCG}}$ performs as well as $\mathcal{M}_D^{\text{VCG}}$ with a single slot as $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_D^{\text{VCG}}$ are equivalent in this case since there is no externality; with more than 2 slots the inefficiency can be large both for social welfare and revenue even in the basic case in which slots are indistinguishable and $\lambda = 1$. In particular, in our proofs of the upper-bound results, we use a special class of quality functions that we denote as *only-min* functions, which assign a value 0 to the quality of an agent when her price is not the minimum among those displayed, and we prove that in many cases no worse instance is possible. With multiple slots, the positive result is that, with $\mathcal{M}_1^{\text{VCG}}$, the optimal allocation is always achievable by some Nash equilibrium (i.e., PoS = 1). Nevertheless, there are auction instances in which some Nash equilibria lead to allocations whose social welfare is $1/m$ of the optimal allocation (i.e., PoA = m) or in which all the Nash equilibria lead to a revenue of zero whereas the direct-revelation mechanism $\mathcal{M}_D^{\text{VCG}}$ provides a strictly positive revenue (i.e., PoS = ∞). $\mathcal{M}_1^{\text{GSP}}$ performs even worse than $\mathcal{M}_1^{\text{VCG}}$, both with a single and multiple slots.

In the following, we formally provide the results on the lower and upper bounds over the mechanisms' inefficiency.

5.1 Price of Stability for the Social Welfare

Initially, we provide our main positive result in terms of indirect-revelation mechanisms inefficiency.

Theorem 4. *The PoS for the social welfare of $\mathcal{M}_1^{\text{VCG}}$ is 1.*

PROOF. Suppose that each agent i reports the pair $(\tilde{p}_i, \tilde{b}_i)$ defined as follows: if the mechanism $\mathcal{M}_D^{\text{VCG}}$ displays the ad i when run on truthful bids, then \tilde{p}_i is the corresponding price, and $\tilde{b}_i = \alpha_i(\tilde{p}_i - c_i)$, i.e., the true gain associated to this price; otherwise $\tilde{p}_i = \tilde{b}_i = 0$. It is immediate to check that with these bids the allocation returned by $\mathcal{M}_1^{\text{VCG}}$ is exactly the same as the one returned by $\mathcal{M}_D^{\text{VCG}}$, and, thus, it maximizes social welfare.

Unfortunately, we cannot conclude that inputs $(\tilde{p}_i, \tilde{b}_i)$ are in equilibrium directly from the truthfulness of $\mathcal{M}_D^{\text{VCG}}$. Indeed, the payments assigned by the indirect mechanism are different from the ones assigned by the direct mechanism. Moreover, in the former the agent may lie both about the price and about the expected gain, while in the latter an agent may essentially lie only on the

	1 slot			$m \geq 2$ slots		
	Social Welfare		Revenue	Social Welfare		Revenue
	PoS	PoA	PoS	PoS	PoA	PoS
$\mathcal{M}_1^{\text{VCG}}$	1	1	1 (♣)	1	m	∞
$\mathcal{M}_1^{\text{GSP}}$	1	1	∞	≥ 2	$\geq m$	∞

Table 1: Lower and upper bounds over PoS and PoA when agents do not overbid. ♣: PoS here is taken w.r.t. the mechanism $\mathcal{M}_D^{\text{VCG}}$ maximizing the social welfare (thus not necessarily maximizing the revenue).

expected gain. Still, in the following we prove that inputs $(\tilde{p}_i, \tilde{b}_i)$ are in equilibrium, and, thus, the theorem follows.

In particular, let $\tilde{\mathbf{p}} = (\tilde{p}_1, \dots, \tilde{p}_n)$ and $\tilde{\mathbf{b}} = (\tilde{b}_1, \dots, \tilde{b}_n)$. We prove that the utility \tilde{u}_i of agent i when the mechanism $\mathcal{M}_1^{\text{VCG}}$ is run on $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{b}}$ is at least the utility u_i that she achieves if the mechanism would be run on input $\mathbf{p} = (p_i, \tilde{p}_{-i})$ and $\mathbf{b} = (b_i, \tilde{b}_{-i})$, for every i , p_i , and b_i . Indeed if i is allocated by the mechanism $\mathcal{M}_1^{\text{VCG}}$ when run on input $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{b}}$, then, since, by definition of \tilde{b}_i , $v_i = \widehat{v}_i(f^*, \tilde{\mathbf{p}}, \tilde{b}_i)$,

$$\begin{aligned} \tilde{u}_i &= v_i - \pi_i = \widehat{v}_i(f^*, \tilde{\mathbf{p}}, \tilde{b}_i) - \pi_i \\ &= \widehat{SW}(f^*, \tilde{\mathbf{p}}, \tilde{\mathbf{b}}) - \max_{g: g(i) \notin M} \widehat{SW}(g, \tilde{\mathbf{p}}, \tilde{\mathbf{b}}) \geq 0, \end{aligned}$$

where f^* is the allocation returned by $\mathcal{M}_D^{\text{VCG}}$ on truthful bids. If i is instead, unallocated then

$$\tilde{u}_i = 0 = \widehat{SW}(f^*, \tilde{\mathbf{p}}, \tilde{\mathbf{b}}) - \max_{g: g(i) \notin M} \widehat{SW}(g, \tilde{\mathbf{p}}, \tilde{\mathbf{b}}).$$

Thus, if the agent i is unallocated by the mechanism $\mathcal{M}_1^{\text{VCG}}$ when run on input \mathbf{p} and \mathbf{b} , then the equilibrium condition is trivially satisfied. Otherwise, let $\check{b}_i = \alpha_i(p_i - c_i)$ and $\check{\mathbf{b}} = (\check{b}_i, \tilde{b}_{-i})$. We have:

$$\begin{aligned} u_i &= v_i - \pi_i = \widehat{v}_i(g^*, \mathbf{p}, \check{b}_i) - \widehat{v}_i(g^*, \mathbf{p}, b_i) \\ &\quad + \widehat{SW}(g^*, \mathbf{p}, \mathbf{b}) - \max_{g: g(i) \notin M} \widehat{SW}(g, \mathbf{p}, \mathbf{b}) \\ &= \widehat{SW}(g^*, \mathbf{p}, \check{\mathbf{b}}) - \max_{g: g(i) \notin M} \widehat{SW}(g, \check{\mathbf{p}}, \check{\mathbf{b}}), \end{aligned}$$

where the last inequality follows since $p_j = \check{p}_j$ and $b_j = \check{b}_j$ for every agent $j \neq i$.

Since $\widehat{SW}(f^*, \tilde{\mathbf{p}}, \tilde{\mathbf{b}}) \geq \widehat{SW}(g^*, \mathbf{p}, \check{\mathbf{b}})$, because f^* and $\tilde{\mathbf{p}}$ are the allocation and the prices that maximize the social welfare, we have that $\tilde{u}_i \geq u_i$, as desired. \square

The proof of the theorem above shows that, with VCG payments, there is always a Nash equilibrium in which every agent i bids the truthful gain b_i and the price that $\mathcal{M}_D^{\text{VCG}}$ would use. Such a strategy profile leads to the same allocation of $\mathcal{M}_D^{\text{VCG}}$, thus guaranteeing a PoS for the social welfare of 1, but, as we discuss in the following sections, the revenue of the two mechanisms can be different. The same result does not hold in the case of GSP payments, thus leading to a larger PoS for the social welfare.

Theorem 5. *The PoS for the social welfare of $\mathcal{M}_1^{\text{GSP}}$ is at least 2 even if agents do not overbid.*

5.2 Price of Anarchy for the Social Welfare

We initially focus on the basic case with a single slot, showing that in this case $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_1^{\text{GSP}}$ are efficient.²

Theorem 6. *The PoA for the social welfare of $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_1^{\text{GSP}}$ is 1 if $m = 1$ when agents do not overbid.*

Then, we study the case with multiple slots providing a lower bound on PoA.

Theorem 7. *The PoA for the social welfare of $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_1^{\text{GSP}}$ is at least m if $m \geq 2$ when agents do not overbid.*

²All the proofs of the theorems in this section and the following one are in the Supplementary Material.

In the specific case of $\mathcal{M}_1^{\text{VCG}}$, we show that a PoA larger than m is not possible, and therefore there are no instances worse than those used in the proof of Theorem 7. Most interestingly, this result holds even when q_i is not monotonically decreasing in p_i .

Theorem 8. *The PoA for the social welfare of $\mathcal{M}_1^{\text{VCG}}$ is at most m if $m \geq 2$ when agents do not overbid.*

Finally, we show that when agents overbid, the inefficiency can be arbitrarily large.

Theorem 9. *The PoA for the social welfare of $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_1^{\text{GSP}}$ is ∞ even if $m = 1$ when agents can overbid.*

5.3 Price of Stability for the Revenue

Initially, we provide our main result, showing that $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_1^{\text{GSP}}$ can be arbitrarily inefficient even with 2 slots.

Theorem 10. *The PoS for the revenue of $\mathcal{M}_1^{\text{VCG}}$ and $\mathcal{M}_1^{\text{GSP}}$ is ∞ even if $m = 2$.*

In the specific case of $\mathcal{M}_1^{\text{VCG}}$ and $m = 1$, we have a positive result for PoS (PoA is trivially ∞ as it is ∞ even in second-price single-item auctions).

Theorem 11. *The PoS for revenue of $\mathcal{M}_1^{\text{VCG}}$ is 1 if $m = 1$.*

Instead, the above positive result does not hold with $\mathcal{M}_1^{\text{GSP}}$, as stated below.

Theorem 12. *The PoS for the revenue of $\mathcal{M}_1^{\text{GSP}}$ is ∞ even if $m = 1$ when agents do not overbid.*

In the proof of this theorem we strongly rely upon the definition of GSP payments described above, which restricts payments to depend only on agents submitting a price at least p_{\min} . This payment format turns out to be necessary in order to guarantee individual rationality. We leave open the problem of understanding if a better Price of Stability for the revenue of $\mathcal{M}_1^{\text{GSP}}$ would be possible by considering alternative non-individually rational GSP payments.

6 A BETTER POS FOR THE REVENUE WITH INDIRECT-REVELATION MECHANISMS

As discussed in the previous section, indirect-revelation mechanisms present major weaknesses in terms of efficiency. A natural question is whether we can design indirect-revelation mechanisms with a better efficiency when agents can choose their price. In particular, we focus on $\mathcal{M}_1^{\text{VCG}}$, as it always guarantees PoS = 1 for the social welfare, and we show that a simple modification of the mechanism leads to PoS = 1 for the revenue when some assumptions hold. We call this new mechanism $\mathcal{M}_1^{\text{VCG}*}$. The rationale is to ask agents for more information. More precisely, the input provided by every agent is a triple composed of (b_i, p_i, p_i^*) where (b_i, p_i) is the input to $\mathcal{M}_1^{\text{VCG}}$ and p_i^* is the price that advertiser i would choose when her ad is the only displayed ad. The property that PoS = 1 is guaranteed when function $q_i(p_i, p_i)$ is differentiable in p_i and non-zero in p_i^* . Mechanism $\mathcal{M}_1^{\text{VCG}*}$ is defined as follows:

- (1) every agent i submits a bid (b_i, p_i, p_i^*) , where b_i, p_i , and p_i^* are defined as above;

- (2) the mechanism infers the values of c_i and α_i for every agent i as follows: $\hat{c}_i = q(p_i^*, p_i^*) / \frac{dq(p_i, p_i)}{dp_i} \Big|_{p_i=p_i^*} + p_i^*$, and $\hat{\alpha}_i = \frac{b_i}{p_i - \hat{c}_i}$ if $p_i \neq \hat{c}_i$ and $\hat{\alpha}_i = 0$ otherwise;
- (3) the mechanism computes an auxiliary allocation, say \bar{f} , by using the allocation function of $\mathcal{M}_1^{\text{VCG}}$ when the input is (b_i, p_i) for every agent i ; the corresponding social welfare (evaluated with the declared gain b_i) is $\bar{S\bar{W}}$;
- (4) for every agent i , the mechanism computes an auxiliary allocation, say \bar{f}^{-i} , by using the allocation function of $\mathcal{M}_D^{\text{VCG}}$ when the values inferred above for $\{\hat{\alpha}_h\}_{h \in N}$ and $\{\hat{c}_h\}_{h \in N}$ are provided in input and agent i is removed from the optimization problem. For every maximization, we denote with $\bar{S\bar{W}}^{-i}$ the corresponding social welfare evaluated with the inferred values $\{\hat{\alpha}_h\}_{h \in N}$ and $\{\hat{c}_h\}_{h \in N}$. Notice that, as it happens with $\mathcal{M}_D^{\text{VCG}}$, the prices in output to these maximizations can be different from those agents provide in input;
- (5) if $\bar{S\bar{W}} \geq \max_i \bar{S\bar{W}}^{-i}$, then the mechanism chooses allocation \bar{f} and charges every agent i of a payment $\pi_i = \bar{S\bar{W}}^{-i} - (\bar{S\bar{W}} - \lambda_{\bar{f}(i)} q_i(p_i, p_{\min}) b_i)$, else no ad is allocated and every agent i is charged a payment of zero.

Basically, mechanism $\mathcal{M}_1^{\text{VCG}^*}$ exploits the additional information asked to the agents to infer their types and then uses this information to compute the same payments that $\mathcal{M}_D^{\text{VCG}}$ would charge. Step 5 is necessary to guarantee individual rationality. More precisely, since the allocation \bar{f} is computed as the indirect mechanism does (without optimizing over prices), while the payments $\{\pi_i\}_{i \in N}$ are computed as the direct mechanism does (optimizing over prices), individual rationality may not be satisfied. We solve this problem setting the payments to 0 (and allocating no ads) when the payments $\{\pi_i\}_{i \in N}$ are too large. As a side effect, we have that if the submitted prices are different from the optimal one, it is possible that the mechanism does not assign any slot. Thus, the PoA for the social welfare and revenue can be unbounded.

Theorem 13. *Mechanism $\mathcal{M}_1^{\text{VCG}^*}$ is individually rational and weakly budget-balanced. Moreover, the PoS for the revenue of $\mathcal{M}_1^{\text{VCG}^*}$ is 1.*

We recall that the algorithm we provide to find the best allocation with $\mathcal{M}_1^{\text{VCG}}$ works when the values that p_i can assume are discrete, and the same holds with $\mathcal{M}_1^{\text{VCG}^*}$. We also notice that $\mathcal{M}_1^{\text{VCG}^*}$ requires that p_i^* is not restricted to a set of discrete values, the mechanism could not infer the exact values of α_i and c_i otherwise. However, requiring price p_i to belong to a finite, discrete set of values and price p_i^* to belong to $\mathbb{R}_{\geq 0}$ does not modify the properties of the mechanism since p_i^* is not used in the allocation algorithm.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we investigate how displaying prices together with ads affects the users' behavior and the properties of auction mechanisms. Since the goods sold by the agents are similar, a high competition among the agents arises from the price comparison. Technically speaking, the prices introduce externalities as the probability with which a user clicks on an ad depends on the price of that ad and on

the prices of the other displayed ads. Interestingly, the social welfare can be maximized when a direct-revelation mechanism jointly optimizes over the ad allocation and the prices, and we show that this can be done in polynomial time when the prices can assume a finite set of values. However, in practice, it is unlikely that advertisers would allow the mechanism to choose prices on their behalf and, in commonly-adopted mechanisms, ads allocation and price optimization are decoupled, so that the advertisers optimize prices and bids, while the mechanism does so for the allocation, once prices and bids are given. We show that this decoupling makes standard mechanisms with VCG and GSP payments highly inefficient in terms of PoA and PoS for social welfare and revenue. Finally, we investigate whether we can reduce the inefficiency of mechanisms in which the advertisers optimize prices and bids. We show that we can obtain PoS of 1 for the revenue by a simple modification of the mechanisms. In particular, we ask the advertisers for an additional price that the mechanism exploits to infer the values of some advertisers' parameters. Such a modification can be easily implemented in practice without that agents reveal their private, sensitive information.

Many research directions can be explored in future. Probably, the most interesting concerns how the bidding strategies commonly adopted for standard ad auctions without prices can be extended to our case. In particular, the crucial question is whether, as in the case of the standard GSP without prices, there are bidding strategies converging to notable Nash equilibria. Other interesting questions concern the analysis of PoA and PoS and the design of allocation algorithms when the quality functions satisfy specific properties, such as, e.g., smoothness.

REFERENCES

- [1] G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pál. 2008. Sponsored Search Auctions with Markovian Users. In *WINE*. 621–628.
- [2] Francesco Bacchiocchi, Matteo Castiglioni, Alberto Marchesi, Giulia Romano, and Nicola Gatti. 2022. Public Signaling in Bayesian Ad Auctions. *arXiv preprint arXiv:2201.09728* (2022).
- [3] Y. Bachrach, S. Ceppi, I.A. Kash, P.B. Key, and D. Kurokawa. 2014. Optimising trade-offs among stakeholders in ad auctions. In *ACM EC*. 75–92.
- [4] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. 2011. On the efficiency of equilibria in Generalized Second Price auctions. In *EC*.
- [5] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, and E. Tardos. 2015. Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory* 156 (2015), 343–388.
- [6] Matteo Castiglioni, Alessandro Nuara, Giulia Romano, Giorgio Spadaro, Francesco Trovò, and Nicola Gatti. 2022. Safe Online Bid Optimization with Return-On-Investment and Budget Constraints subject to Uncertainty. *arXiv preprint arXiv:2201.07139* (2022).
- [7] Matteo Castiglioni, Giulia Romano, Alberto Marchesi, and Nicola Gatti. 2022. Signaling in Posted Price Auctions. *arXiv preprint arXiv:2201.12183* (2022).
- [8] M. Chui, J. Manyika, M. Miremadi, N. Henke, R. Chung, P. Nel, and S. Malhotra. 2018. Notes From The AI Frontier Insights From Hundreds Of Use Cases. *McKinsey Global Institute* (2018).
- [9] G. Farina and N. Gatti. 2016. Ad Auctions and Cascade Model: GSP Inefficiency and Algorithms. In *AAAI*. 489–495.
- [10] G. Farina and N. Gatti. 2017. Adopting the Cascade Model in Ad Auctions: Efficiency Bounds and Truthful Algorithmic Mechanisms. *J. Artif. Intell. Res.* 59 (2017), 265–310.
- [11] D. Fotakis, P. Krysta, and O. Telelis. 2011. Externalities among advertisers in sponsored search. In *SAGT*. 105–116.
- [12] N. Gatti, M. Rocco, S. Ceppi, and E.H. Gerding. 2014. Mechanism Design for Mobile Geo-Location Advertising. In *AAAI*. 691–697.
- [13] N. Gatti, M. Rocco, P. Serafino, and C. Ventre. 2018. Towards better models of externalities in sponsored search auctions. *Theoretical Computer Science* 745 (2018), 150–162.
- [14] I. Giotis and A.R. Karlin. 2008. On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. In *WINE*. 629–638.
- [15] J. Hartline, D. Hoy, and S. Taggart. 2014. Price of anarchy for auction revenue. In *EC*. 693–710.
- [16] D. He, W. Chen, L. Wang, and T.-Y. Liu. 2013. A Game-Theoretic Machine Learning Approach for Revenue Maximization in Sponsored Search. In *IJCAI*. 206–212.
- [17] K. Jung, Y.C. Cho, and S. Lee. 2014. Online shoppers’ response to price comparison sites. *Journal of Business Research* 67, 10 (2014), 2079–2087.
- [18] D. Kempe and M. Mahdian. 2008. A cascade model for externalities in sponsored search. In *WINE*. 585–596.
- [19] B. Lucier and R. Paes Leme. 2011. GSP auctions with correlated types. In *EC*.
- [20] A. Mas-Colell, M.D. Whinston, and J.R. Green. 1995. *Microeconomic Theory*. Oxford University Press.
- [21] J. Nash. 1951. Non-cooperative games. *Annals of mathematics* (1951), 286–295.
- [22] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press, USA.
- [23] A. Nuara, N. Sosio, F. Trovò, M.C. Zaccardi, N. Gatti, and M. Restelli. 2019. Dealing with Interdependencies and Uncertainty in Multi-Channel Advertising Campaigns Optimization. In *WWW*. 1376–1386.
- [24] A. Nuara, F. Trovò, N. Gatti, and M. Restelli. 2018. A Combinatorial-Bandit Algorithm for the Online Joint Bid/Budget Optimization of Pay-per-Click Advertising Campaigns. In *AAAI*. 2379–2386.
- [25] Giulia Romano, Gianluca Tartaglia, Alberto Marchesi, and Nicola Gatti. 2021. Online Posted Pricing with Unknown Time-Discounted Valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5682–5689.
- [26] T. Roughgarden, V. Syrgkanis, and E. Tardos. 2017. The Price of Anarchy in Auctions. *J. Artif. Intell. Res.* 59 (2017), 59–101.