The Power of Media Agencies in Ad Auctions: Improving Utility through Coordinated Bidding

Giulia Romano, **Matteo Castiglioni**, **Alberto Marchesi** and **Nicola Gatti**
Politecnico di Milano, Piazza Leonardo da Vinci 32, I-20133, Milan, Italy
{giulia.romano, matteo.castiglioni, alberto.marchesi, nicola.gatti}@polimi.it

Abstract

The increasing competition in digital advertising induced a proliferation of *media agencies* playing the role of intermediaries between *advertisers* and *platforms* selling ad slots. When a group of competing advertisers is managed by a common agency, many forms of collusion, such as bid rigging, can be implemented by coordinating bidding strategies, dramatically increasing advertisers' value. We study the computational problem faced by a media agency that has to coordinate the bids of a group of colluders, under GSP and VCG mechanisms. First, we introduce an abstract bid optimization problem, called *weighted utility problem* (WUP), which is useful in proving our results. We show that the utilities of bidding strategies are related to the length of paths in a directed acyclic weighted graph, whose structure and weights depend on the mechanism under study. This allows us to solve WUP in polynomial time by finding a shortest path in the graph. Next, we switch to our original problem, focusing on two settings that differ for the *individual rationality constraints* they require. Such constraints ensure that colluders do *not* leave the agency, and they can be enforced by implementing *monetary transfers* between the agency and the advertisers. In particular, we study the *arbitrary transfers* setting, where any kind of transfer to and from the advertisers is allowed, and the more realistic *limited liability* setting, in which no advertiser can be paid by the agency. In the former, we cast the problem as a WUP instance and solve it by our graph-based algorithm, while, in the latter, we formulate it as a linear program with exponentially-many variables efficiently solvable by applying the ellipsoid algorithm to its dual. This requires solving a suitable separation problem in polynomial time, which can be done by reducing it to a WUP instance.

1 Introduction

Over the last years, digital advertising has been one of the main drivers of the growth of world market. Remarkably, the vast majority of the companies employ digital tools to advertise their products or services, and the worldwide annual spent in digital advertising reached about 150 billion USD in 2020 [IAB, 2021]. Furthermore, most of the economic reports forecast that *artificial intelligence* (AI) will fuel an increase of market value by almost 100% over the next decade [Chui *et al.*, 2018]. Indeed, AI tools have become widespread throughout digital advertising, enabling new opportunities that were not available before. Notable examples are in, *e.g.*, auction design [Bachrach *et al.*, 2014; Castiglioni *et al.*, 2022a], budget allocation [Nuara *et al.*, 2018], bidding opimization [He *et al.*, 2013; Castiglioni *et al.*, 2022b], and multi-channel advertising [Nuara *et al.*, 2019].

Recent years have witnessed a proliferation of *media agencies*, which claim to play the role of intermediaries between *advertisers* and *platforms* selling ad slots. This has been driven by the increasing complexity of digital advertising—due to, *e.g.*, a huge amount of available data and of parameters to be set on advertising platforms—and the rising competition among a growing number of advertisers. When a group of competing advertisers is managed by a common agency, it frequently happens that the agency has to place bids on behalf of different advertisers participating in the same ad auction. This dramatically changes the strategic interaction underlying ad auctions, since agencies can coordinate advertisers' bidding strategies by implementing many forms of collusion—such as, *e.g.*, bid rigging—, increasing advertisers' value. In particular, simple examples show that colluding in ad auctions can reward the colluders with a utility that is arbitrarily larger than what they would get without doing that. Moreover, a recent empirical study on real-world data by one of the major US agencies (Aegis-Dentsu-Merkle) shows that collusion is pervasive and leads to a significant reduction in the average cost-per-click [Decarolis and Rovigatti, 2017].

Original Contributions. We study the computational problem faced by a media agency that has to coordinate the bidding strategies of a group of colluders, under GSP and VCG mechanisms (details on the adoption of such mechanisms in ad auctions are provided in [Varian and Harris, 2014]). We assume that the media agency knows the private valuations of the colluders (*i.e.*, how much they value a click on their ad), and that it decides colluders' bids on their behalf. Moreover, the media agency is in charge of paying the auction mechanism for a click on an allocated ad, and at the same time it requires *monetary transfers* to and from the colluders.

These are necessary to enforce some *individual rationality constraints* ensuring that the colluders do *not* leave the agency and participate in the ad auction individually. In this paper, we study two settings that differ for the kind of monetary transfers that they allow for: the *arbitrary transfers* setting, where any kind of monetary transfer to and from the advertisers is allowed, and the more realistic *limited liability* setting, in which no advertiser can be paid by the media agency. Finally, we assume that the bids of the advertisers external to the media agency are drawn according to some probability distribution. As a first result, we introduce an abstract bid optimization problem, called *weighted utility problem* (WUP), which works for any finite set of possible bid values and is useful in proving our main results in the rest of the paper. In order to solve such a problem, we first show that the utilities of bidding strategies are related to the length of paths in a directed acyclic weighted graph, whose structure and weights depend on the mechanism under study (either GSP or VCG mechanism). This allows us to solve WUP instances in polynomial time by finding a shortest path of the graph. Next, we switch the attention to the original media agency problem, starting from the arbitrary transfers setting. A major challenge is dealing with a potentially continuous set of possible bids. Notably, we show that it is possible to reduce the attention to a finite set of bidding strategies, only incurring in a small additive loss in the value of the obtained solution and relaxing the incentive constraints by a small additive amount. The set is built by recursively splitting the interval of possible bid values, until one gets sub-intervals such that the probability that an external bid is in a given sub-interval is sufficiently small. Then, the resulting sub-intervals are used to define the desired finite set of bids. In conclusion, we cast the problem as a WUP instance and solve it by our graph-based algorithm in polynomial time. This gives a bi-criteria additive FPTAS for the original problem, since the (additive) approximation is in terms of both objective value and incentive constraints. Finally, we study the limited liability setting. In this case, we leverage the same finite set of bidding strategies defined for the arbitrary transfers setting in order to formulate the problem as a *linear program* (LP) with exponentially-many variables and polynomially-many constraints. Since we use only a finite set of bids, we need to relax the individual rationality constraints by an arbitrary small amount to guarantee the existence of a feasible solution. We solve such an LP in polynomial time by applying the ellipsoid algorithm [Grötschel *et al.*, 1981] to its dual, which features polynomially-many variables and exponentially-many constraints. This requires solving a suitable separation problem in polynomial time, which can be done by reducing it to a WUP instance. As in the arbitrary transfer setting, the resulting algorithm is a bi-criteria additive FPTAS.

Related Works. While a longstanding literature investigates the role of mediators in ad auctions—see, *e.g.*, the seminal works by Vorobeychik and Reeves [2008] and Ashlagi *et al.* [2009]—, collusion is currently emerging as one of the central problems in advertising, as the adoption of AI algorithms can concretely support an agency to find the best collusive behaviors [OECD, 2017]. Motivated by the recent study by Decarolis and Rovigatti [2017], some works provide theoretical contributions to assess how collusion can be conducted by an agency. Decarolis *et al.* [2020] study a setting in which there is no monetary transfer between the agency and bidders by providing equilibrium conditions. They show that, in simple settings, GSP is more inefficient than VCG, both in terms of efficiency and revenue. Lorenzon [2018] studies a setting with two slots and three bidders that are all controlled by an agency in a GSP auction. Furthermore, a monetary transfer is possible. The author shows collusive stable behaviors in which the redistribution is uniform over the three bidders.

2 Preliminaries

We study the problem of coordinated bidding faced by a media agency in ad auctions, with both GSP and VCG payments. In this setting, the set $N := \{1, \dots, n\}$ of *bidders* (or *agents*) is partitioned as $N := N_c \cup N_e$, where: N_c is a set of advertisers whose advertising campaigns are managed by a common media agency, while N_e is a set of advertisers that are *not* part of the agency, but participate in the ad auction individually. In this work, we refer to the former as *colluders*, while we call the latter *external agents*. Moreover, we let $n_c := |N_c|$ and $n_e := |N_e|$ be the numbers of colluders and external agents, respectively. The advertisers compete for displaying their ads on a set $M := \{1, ..., m\}$ of *slots*, with $m \leq n$. Each agent $i \in N$ has a *private valuation* $v_i \in [0, 1]$ for an advertising slot, which reflects how much they value a click on their ad. Furthermore, each slot $j \in M$ is associated with a *click through rate* parameter $\lambda_j \in [0, 1]$, encoding the probability with which the slot is clicked by a user.[1] Each agent $i \in N$ participates in the ad auction with a *bid* $b_i \in [0, 1]$, representing how much they are willing to pay for a click on their ad. We denote by $b = (b_i)_{i \in N}$ the bid profile made by all the agents' bids. We also let $b^c = (b_i)_{i \in N_c}$ be the profile of colluders' bids (also called *bidding strategy*), while $b^e = (b_i)_{i \in N_e}$ is the profile of external agents' bids. For the ease of notation, we sometimes write $b = (b^c, b^e)$ to denote the profile made by all the bids in b^c and b^e.

The media agency knows the valuations v_i of all the colluders $i \in N_c$, and it decides the bid profile b^c on their behalf. Additionally, the media agency defines a monetary *transfer* $q_i \in [-1, 1]$ for each colluder $i \in N_c$. We adopt the convention that, if $q_i > 0$, then the transfer is from the agent to the agency, while, if $q_i < 0$, then it is the other way around.[2]

W.l.o.g., we assume that the slots are ordered in decreasing value of click through rate, so that $\lambda_1 \geq \ldots \geq \lambda_m$. Moreover, for the ease of presentation, we let $\lambda_{m+1} = \ldots = \lambda_{n+1} = 0$.

The auction goes on as follows. First, the media agency selects a bidding strategy $b^c = (b_i)_{i \in N_c}$ and requires a transfer q_i from each colluder $i \in N_c$. Then, external agents individually report their bids to the auction mechanism, resulting in

[1]For the ease of presentation, we assume that the click through rate only depends on the slot and *not* on the ad being displayed. This dependence can be easily captured by interpreting v_i as an expected value w.r.t. clicks once the user observed the slot.

[2]Notice that there are some scenarios in which it is in the interest of the media agency to pay a colluder in order to ensure that they stay in the agency; see Example 2 in [Romano *et al.*, 2022].

a profile $b^e = (b_i)_{i \in N_e}$, while the media agency reports bids b^c on behalf of the colluders. Finally, given all the agents' bids $b = (b^c, b^e)$, the mechanism allocates an ad to each slot and defines an *expected payment* $\pi_i(b) \in [0,1]$ for each agent $i \in N$, where the expectation is with respect to the clicks. The media agency is responsible of paying the auction mechanism on behalf of the colluders.

Given a bid profile $b = (b_i)_{i \in N}$, assuming w.l.o.g. that each bidder $i \in [m]$ is assigned to slot i (by re-labeling bidders accordingly), we denote bidder i's *expected revenue* as $r_i(b) := \lambda_i v_i$, while bidder i's *expected utility* is $u_i(b) := r_i(b) - q_i$.[3] Instead, the expected utility of the agency is $\sum_{i \in N_c}(q_i - \pi_i(b))$. We also denote with U the cumulative expected utility of all the colluders and the media agency. Formally, it holds $U := \sum_{i \in N_c}(r_i(b) - \pi_i(b))$.

Next, we review GSP and VCG mechanisms in ad auctions (see the book by Nisan and Ronen [2001] for their general description). Given a bid profile $b = (b_i)_{i \in N}$, assuming w.l.o.g. that $b_1 \geq \ldots \geq b_n$ (by re-labeling bidders accordingly), both mechanisms orderly assign the first m agents, who are those with the highest bids, to the first m slots, which are those with the highest click through rates. Moreover, the mechanisms assign the following expected payments.

- *GSP mechanism*: $\pi_i^{\text{GSP}}(b) := \lambda_i b_{i+1}$ for each agent $i \in [m]$, and $\pi_i^{\text{GSP}}(b) = 0$ for all the other agents.
- *VCG mechanism*: $\pi_i^{\text{VCG}}(b) := \sum_{j=i+1}^{m+1} b_j(\lambda_{j-1} - \lambda_j)$ for each agent $i \in [m]$, and $\pi_i^{\text{VCG}}(b) = 0$ for the others.

The VCG payments are such that each agent is charged a payment that is equal to the externalities that they impose on other agents. This makes the VCG mechanism *truthful*, which means that it is a dominant strategy for each agent to report their true valuation to the mechanism, namely $b_i = v_i$ for every $i \in N$. This is *not* the case for the GSP mechanism.

3 Problem Formulation

In this section, we introduce the optimization problem faced by the media agency. In words, the goal of the media agency is to find a bidding strategy $b^c = (b_i)_{i \in N_c}$ that coordinates colluders' bids in a way that maximizes the cumulative expected utility U, while at the same time guaranteeing that they are incentivized to be part of the media agency. The rest of the section is devoted to formally defining such a problem.

Before introducing the optimization problem, let us notice that knowing the valuations of all the colluders allows the media agency to improve the cumulative expected utility U with respect to the case in which all the bidders act individually. This is formalized by the following proposition, whose proof readily follows from Example 1 in [Romano *et al.*, 2022].

Proposition 1. *The cumulative expected utility U may be arbitrarily larger than the sum of the colluders' expected utilities when they participate in the ad auction individually.*

In general, the media agency may adopt a *randomized* bidding strategy in order to maximize U. By letting B^c be the set of all the possible colluders' bid profiles $b^c = (b_i)_{i \in N_c}$, we denote by $\gamma \in \Gamma$ any randomized bidding strategy, where Γ is the set of all the probability distributions over B^c. Moreover, whenever $\gamma \in \Gamma$ has a finite support, we denote with γ_{b^c} the probability of choosing a bidding strategy $b^c \in B^c$.

In this work, unless stated otherwise, we consider the case in which the bid profile $b^e = (b_i)_{i \in N_e}$ of the external agents is drawn from a probability distribution γ^e. Then, we define the expected revenue of bidder $i \in N_c$ for any bidding strategy $b^c \in B^c$ as $\tilde{r}_i(b^c) := \mathbb{E}_{b^e \sim \gamma^e} r_i(b^c, b^e)$, while their expected payment is as $\tilde{\pi}_i(b^c) := \mathbb{E}_{b^e \sim \gamma^e} \pi_i(b^c, b^e)$. In the rest of the paper, we assume that all algorithms have access to an oracle that returns the value of the expectations $\tilde{r}_i(b^c)$ and $\tilde{\pi}_i(b^c)$, for a bidding strategy $b^c \in B^c$ given as input.[4]

The following Problem (1) encodes the maximization problem faced by the media agency, where the meaning of IR and AIR constraints is described in the following.

$$\max_{q, \gamma \in \Gamma} \sum_{i \in N_c} \mathbb{E}_{b^c \sim \gamma}[\tilde{r}_i(b^c) - \tilde{\pi}_i(b^c)] \quad \text{s.t.} \tag{1a}$$

$$\text{IR}: \quad \mathbb{E}_{b^c \sim \gamma}[\tilde{r}_i(b^c)] - q_i \geq t_i \qquad \forall i \in N_c \tag{1b}$$

$$\text{AIR}: \quad \sum_{i \in N_c} q_i \geq \sum_{i \in N_c} \mathbb{E}_{b^c \sim \gamma}[\tilde{\pi}_i(b^c)]. \tag{1c}$$

The elements of Problem (1) are defined as follows.

Objective (1a) encodes the cumulative expected utility U of the colluders and the media agency, in expectation with respect to the randomized bidding strategy γ.

Constraints (1b), which are called *individual rationality* (IR) constraints, ensure that the colluders are incentivized to be part of the media agency, rather than leaving it and participating in the ad auction as external agents. In particular, they guarantee that each colluder $i \in N_c$ achieves at least a minimum expected utility t_i, where the values $t_i \in [0,1]$ for $i \in N_c$ are given as input.[5]

Constraint (1c) is an *agency individual rationality* (AIR) constraint which provides guarantees over the utility of the media agency. Since the agency corresponds to the mechanism a payment $\sum_{i \in N_c} \mathbb{E}_{b^c \sim \gamma}[\tilde{\pi}_i(b^c)]$ in expectation over the clicks, the constraint requires that the sum of transfers $\sum_{i \in N_c} q_i$ covers the payment, so that the expected utility attained by the agency is non-negative.

In the following, we sometimes relax IR constraints by using δ-IR constraints, for $\delta > 0$, which are defined as follows:

$$\delta\text{-IR}: \quad \mathbb{E}_{b^c \sim \gamma}[\tilde{r}_i(b^c)] - q_i \geq t_i - \delta \qquad \forall i \in N_c. \tag{2}$$

In the following, we call the scenario described so far, in which transfers q_i could be negative, the *arbitrary transfers*

[3] We denote with $[m]$ the set $\{1, \ldots, m\}$.

[4] Our results can be easily extended—only incurring in a small additive loss in cumulative expected utility—to the case in which the distribution γ^e is unknown, but the algorithms have access to a black-box oracle that returns i.i.d. samples drawn according to γ^e (rather than returning expected values).

[5] To the best of our knowledge, in the literature there is only one work by Bachrach [2010] that formalizes IR constraints for a setting that is similar to ours. Bachrach [2010] takes inspiration from the concept of core [Peleg and Sudhölter, 2007] in cooperative games in order to define suitable IR constraints. However, this approach has many downsides. The most relevant issue of such an approach is that it is *not* computationally viable, since computing the core would require exponential time in the number of colluders.

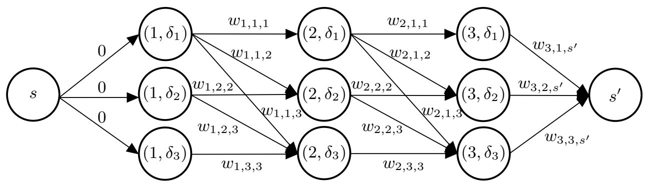

Figure 1: Example of graph for $N_c = \{1, 2, 3\}$, $\Delta = \{\delta_1, \delta_2, \delta_3\}$.

setting. Monetary transfers from the media agency to the agents are *not* always feasible in practice. Indeed, in some real-world scenarios a media agency could potentially lose customers by adopting a strategy for which some agents pay and some others are paid for participating in the same auction. For these reasons, we introduce and study a second scenario, which we call *transfers with limited liability* setting, where no agent is paid by the agency. In such setting, Problem (1) is augmented with the following additional *limited liability* (LL) constraints on the monetary transfers q_i:

$$\text{LL}: \quad q_i \geq 0 \qquad\qquad \forall i \in N_c. \qquad (3)$$

As we show in Section 5, there always exists an optimal solution to Problem (1) without LL constraints that is *not* randomized. The same does not hold for the problem with LL constraints, in which an optimal bidding strategy may be randomized, as in Example 3 in [Romano *et al.*, 2022].

We conclude by introducing the following assumption on the values t_i, guaranteeing that Problem (1) is feasible.[6]

Assumption 1. *There always exists a bidding strategy $b^c \in B^c$ such that $\tilde{r}_i(b^c) - \tilde{\pi}_i(b^c) \geq t_i$ for all $i \in N_c$.*

Proposition 2. *With Assumption 1, Problem* (1) *with LL constraints admits a non-randomized feasible solution.*

4 Weighted Utility Problem

In this section, we provide a polynomial-time algorithm for an abstract bid optimization problem, called *weighted utility problem* (WUP). This will be crucial in the following sections in order to solve Problem (1) with or without LL constraints.

Let $\Delta := \{\delta_1, \ldots, \delta_d\}$ be a discrete set of d different bid values, with $\delta_1 \geq \ldots \geq \delta_d$. Moreover, given a bidding strategy $b^c = (b_i)_{i \in N_c} \in B^c$ such that $b_i \in \Delta$ for all $i \in N_c$, we write $b^c \in B^c_\Delta$ to denote that each b_i belongs to Δ.

Then, WUP reads as follows:

$$\max_{b^c \in B^c_\Delta} \sum_{i \in N_c} \left(\hat{y}_i \, \tilde{r}_i(b^c) - \hat{x} \, \tilde{\pi}_i(b^c)\right), \qquad (4)$$

where $\hat{y}_i \geq 0$ for $i \in N_c$ and $\hat{x} \geq 0$ are give problem parameters. A solution to Problem (4) is a bidding strategy $b^c \in B^c_\Delta$ that maximizes the sum of suitable *weighted* utilities of the colluders, which are defined so that colluder i's expected revenue $\tilde{r}_i(b)$ is weighted by coefficient $\hat{y}_i$, while their expected payment $\tilde{\pi}_i(b)$ is weighted by coefficient $\hat{x}$. Notice that, when $\hat{y}_i = 1$ for all $i \in N_c$ and $\hat{x} = 1$, then the objective of Problem (4) coincides with the cumulative expected utility U.

[6]All the omitted proofs can be found in [Romano *et al.*, 2022].

As a first step, we consider Problem (4) in which the profile of external agents' bids b^e is fixed, which reads as follows:

$$\max_{b^c \in B^c_\Delta} \sum_{i \in N_c} \left(\hat{y}_i \, r_i(b^c, b^e) - \hat{x} \, \pi_i(b^c, b^e)\right). \qquad (5)$$

For the sake of presentation, we first provide our results for Problem (5), and, then, we show how they can be extended to Problem (4), where the bids of external agents are stochastic. Thus, in the rest of the section, we always assume that a profile of external agent's bids $b^e = (b_i)_{i \in N_e}$ is given.

The main idea underpinning our results is to map Problem (4) into a *shortest path problem* [Dijkstra, 1959]. Specifically, we show that the weighted utilities of bidding strategies are related to the length of paths in a particular *directed acyclic weighted graph*, whose structure and weights derive from the considered auction mechanism. The following lemma is crucial for the construction of the graph.

Lemma 1. *Given any bidding strategy $b^c = (b_i)_{i \in N_c} \in B^c_\Delta$ that is optimal for Problem* (5)*, it holds $b_i \geq b_j$ for every pair of colluders $i, j \in N_c$ such that $\hat{y}_i v_i \geq \hat{y}_j v_j$.*

In the rest of this section, for the ease of notation and w.l.o.g. thanks to Lemma 1, we re-label bidders in N so that $N_c := \{1, \ldots, n_c\}$ and $\hat{y}_i v_i \geq \hat{y}_j v_j$ for any $i, j \in N_c : i < j$. Moreover, w.l.o.g., we also re-label external agents so that $N_e := \{n_c + 1, \ldots, n\}$ and $b_i \geq b_j$ for any $i, j \in N_e : i < j$.

We build the graph as follows (see Figure 1).

- The set of vertices V contains dn_c nodes, plus a source node s and a sink node s'. In particular, there are d nodes for each colluder, one for each possible bid value. For every $i \in N_c$ and $\delta_j \in \Delta$, we let (i, δ_j) be the node corresponding to colluder i and bid value δ_j. Intuitively, selecting a path passing through such a node encodes the fact that the bid b_i of colluder i is set to value δ_j.
- The set of arcs A has cardinality $O(d^2 n_c)$. In particular, the source node s is connected to all the nodes $(1, \delta_j)$ (for $\delta_j \in \Delta$), while all nodes (n_c, δ_j) (for $\delta_j \in \Delta$) are connected to the sink node s'. Moreover, for every $i \in N_c \setminus \{n_c\}$ and $\delta_j \in \Delta$, node (i, δ_j) is connected to all nodes $(i+1, \delta_{j'})$ such that $\delta_{j'} \in \Delta$ and $j' \geq j$. Intuitively, each path in the graph going from s to s' defines a bidding strategy $b^c = (b_i)_{i \in N_c} \in B^c_\Delta$ such that $b_i \geq b_j$ for every $i, j \in N_c$ with $\hat{y}_i v_i \geq \hat{y}_j v_j$. Notice that focusing on such bidding strategies is w.l.o.g. by Lemma 1. In the following, we denote the arc going from (i, δ_j) to $(i+1, \delta_{j'})$ with the tuple $(i, \delta_j, \delta_{j'})$.
- Each arc $(i, \delta_j, \delta_{j'})$ has a weight $w_{i,j,j'}$. For $i \in N_c \setminus \{n\}$ and $\delta_j, \delta_{j'} \in \Delta$, the weight $w_{i,j,j'}$ encodes the fraction of cumulative weighted utility obtained by setting $b_i = \delta_j$ and $b_{i+1} = \delta_{j'}$, once all the (higher) bids $b_{i'}$ with $i' \in N_c : i' < i$ are assigned to a value given the preceding nodes in the selected path. The weights $w_{n_c,j,s'}$ on arcs going from nodes (n_c, δ_j) (for $\delta_j \in \Delta$) to the sink s' are defined analogously, while those on arcs exiting from the source s are denied as $w_{s,j} = 0$ for all $\delta_j \in \Delta$.

A *directed path* $\sigma \in \Sigma$ is a sequence of arcs connecting the source node s to the sink node s'. The *length* W_σ of path σ is the sum of the weights of the arcs in the path:

$$W_\sigma = \sum_{(i,\delta_j,\delta_{j'})\in\sigma} w_{i,j,j'}.$$

Then, the shortest path problem on the weighted graph is

$$\min_{\sigma\in\Sigma} -W_\sigma.$$

In conclusion, we define the weights of all arcs $(i, \delta_j, \delta_{j'})$, which depend on the considered auction mechanism, either GSP or VCG. The weights $w_{n_c,j,s'}$ of arcs entering the sink s' are defined analogously, by letting $\delta_{j'} = 0$. For the ease of notation, given $\delta_j \in \Delta$, we let $\tau(\delta_j) := \sum_{i\in N_e} \mathbb{1}_{\{b_i>\delta_j\}}$ be the number of external agents with a bid larger than δ_j.

For the GSP mechanism, the weight of arc $(i, \delta_j, \delta_{j'})$ is

$$w_{i,j,j'} := \lambda_{i+\tau(\delta_j)} \left(\hat{y}_i v_i - \hat{x} \max\left\{ \delta_{j'}, \max_{k\in N_e: b_k<\delta_j} b_k \right\} \right).$$

Intuitively, for the GSP mechanism, the weight $w_{i,j,j'}$ is exactly equal to the weighted utility of colluder i when they bid value δ_j and the colluder $i + 1$ bids value $\delta_{j'}$.

For the VCG mechanism, the weight of arc $(i, \delta_j, \delta_{j'})$ is

$$w_{i,j,j'} := \hat{y}_i\, \lambda_{i+\tau(\delta_j)} v_i - \hat{x}\Big[g_i(\delta_j) + \ell_i(\delta_j, \delta_{j'}) \Big],$$

where we let

$$g_i(\delta_j) := (i-1)\,\delta_j \left(\lambda_{i+\tau(\delta_j)-1} - \lambda_{i+\tau(\delta_j)} \right),$$

$$\ell_i(\delta_j, \delta_{j'}) := \sum_{\substack{k\in N_e:\\ b_k\in(\delta_{j'},\delta_j]}} i\, b_k \Big(\lambda_{k-n_c+i-1} - \lambda_{k-n_c+i} \Big).$$

For the VCG mechanism, $w_{i,j,j'}$ has a less intuitive interpretation than for the GSP mechanism. In particular, $w_{i,j,j'}$ is composed of a revenue term, which is agent i's expected revenue $\lambda_{i+\tau(\delta_j)} v_i$, weighted by $\hat{y}_i$, and two payment terms, $g_i(\delta_j)$ and $\ell_i(\delta_j, \delta_{j'})$, weighted by $\hat{x}$. The latter are two parts of the cumulative payment from the agency to the mechanism, which are related to the externalities of colluders $i' \in N_c$ with $i' < i$, when colluders i and $i + 1$ bid δ_j and $\delta_{j'}$, respectively. In particular, the fraction of the expected payment related to some colluder $i' \in N_c$ with $i' < i$, due to the presence of colluder i bidding δ_j, is $\delta_j(\lambda_{i+\tau(\delta_j)-1} - \lambda_{i+\tau(\delta_j)})$. The term $g_i(\delta_j)$ defines the summation of such payments over all agents $i' \in N_c : i' < i$. Moreover, the fraction of expected payment related to i', due to the presence of an external agent $k \in N_e$ bidding b_k, is $b_k(\lambda_{k-n_c+i-1} - \lambda_{k-n_c+i})$. The term $\ell_i(\delta_j, \delta_{j'})$ is the summation of such expected payments due to external agents with bids $b_k \in (\delta_{j'}, \delta_j]$.

The following lemma establishes the relation between the length of the paths in the weighted graph defined above and the objective of Problem (5).

Lemma 2. *Given any path $\sigma \in \Sigma$ composed by the sequence of nodes $\{(1, \delta_{j_1}), \ldots, (n_c, \delta_{j_{n_c}})\}$, it holds*

$$W_\sigma = \sum_{i\in N_c} \hat{y}_i r_i(b^c, b^e) - \hat{x}\pi_i(b^c, b^e), \tag{6}$$

where $b^c = (\delta_{j_1}, \ldots, \delta_{j_{n_c}})$. Moreover, for any bidding strategy $b^c = (b_i)_{i\in N_c} \in B^c_\Delta$, there exists a corresponding path composed by the sequence of nodes $\{(1, b_1), \ldots, (n_c, b_{n_c})\}$.

Algorithm 1 $\text{REC}((\alpha, \beta], p, \eta)$

1: **if** $\mathbb{P}_{b^e\sim\gamma^e} \{\exists j \in N_e : b_j \in (\alpha, \beta]\} \le p \vee \beta - \alpha \le \eta$ **then**
2:     **return** $\{(\alpha, \beta]\}$
3: **else**
4:     $\mathcal{I}_{\text{L}} \leftarrow \text{REC}\left(\left(\alpha, \frac{\alpha+\beta}{2}\right], p, \eta\right)$
5:     $\mathcal{I}_{\text{R}} \leftarrow \text{REC}\left(\left(\frac{\alpha+\beta}{2}, \beta\right], p, \eta\right)$
6:     **return** $\mathcal{I}_L \cup \mathcal{I}_{\text{R}}$
7: **end if**

The following theorem provides the polynomial-time algorithm for solving Problem (5), which works by simply finding a shortest path of the graph defined above.

Theorem 1. *Problem* (5) *can be solved in polynomial time.*

Finally, by substituting the quantities involved in Problem (5) with their expectations, thanks to the linearity of the objective, we get the following result:

Theorem 2. *Problem* (4) *can be solved in polynomial time.*

5 Arbitrary Transfers Setting

In this section, we provide an approximate solution to the media agency problem with arbitrary transfers. In particular, we design a bi-criteria additive FPTAS that returns solutions providing an arbitrary small loss $\varepsilon > 0$ with respect to the optimal value of the problem, by relaxing the IR constraints by the additive factor ε. As a first step, we show that there always exists a non-randomized solution to Problem (1).

Lemma 3. *In the arbitrary transfers setting, there always exists an optimal non-randomized solution to Problem* (1).

Then, we show how to reduce the problem to a new one working with a finite (discretized) set of bids, in order to apply the results provided in Section 4. As a first result, given a probability value $p \in [0, 1]$ and a minimum discretization step $\eta \in [0, 1]$, we show how to split the space of bids $[0, 1]$ into a suitably-defined set of intervals using the recursive algorithm whose pseudo-code is provided in Algorithm 1.

We prove the following:

Lemma 4. *Given $p \in [0, 1]$ and $\eta \in [0, 1]$, $\text{REC}((0, 1], p, \eta)$ returns a set $\{(\alpha_j, \beta_j]\}_{j\in[k^*]}$ composed of $k^* \le \frac{2n_e}{p} \log \frac{1}{\eta}$ intervals such that, for every interval $(\alpha_j, \beta_j]$, it holds either $\mathbb{P}_{b^e\sim\gamma^e} \{\exists i \in N_e : b_i \in (\alpha_j, \beta_j]\} \le p$ or $\beta_j - \alpha_j \le \eta$. Moreover, it holds that $\bigcup_{j\in[k^*]} (\alpha_j, \beta_j] = (0, 1]$ and the procedure runs in time polynomial in n_e, $\frac{1}{p}$, and $\log \frac{1}{\eta}$.*

Let $I^{p,\eta} := \{(\alpha_j, \beta_j]\}_{j\in[k^*]}$ be the set of intervals returned by $\text{REC}((0, 1], p, \eta)$. The next step is to show that, for η small enough, $\mathbb{P}_{b^e\sim\gamma^e} \{\exists i \in N_e : b_i \in (\alpha_j, \beta_j)\} \le p$ for every interval $(\alpha_j, \beta_j]$. This holds by definition for each interval $(\alpha_j, \beta_j]$ with $\beta_j - \alpha_j > \eta$. Thus, let us consider the intervals $(\alpha_j, \beta_j]$ such that $\beta_j - \alpha_j \le \eta$. Let M be the maximum number of bits needed to represent the bids in the support of probability distribution γ^e. By setting $\eta = 2^{-M}$, we have that all the bids b_i in the interval $(\alpha_j, \beta_j]$ are equal to β_j.[7] Hence, it holds $\mathbb{P}_{b^e \sim \gamma^e} \{\exists i \in N_e : b_i \in (\alpha_j, \beta_j)\} = 0$.

Letting $B_p^c := \bigcup_{(\alpha,\beta] \in I^{p,\eta}} \bigcup_{i \in N_c} \{\alpha + \tau i\}$ be a suitable set of discretized bidding strategies for $\tau > 0$ and $\eta = 2^{-M}$, we show that we can restrict the attention to bid profiles in B_p^c with a small loss in utility and by relaxing IR constraints.[8]

First, we provide the following auxiliary result.

Lemma 5. *Given $p \in [0, 1]$, for any bidding strategy $b^c \in B^c$, there exists a discretized bidding strategy $\hat{b}^c \in B_p^c$:*

- $\tilde{\pi}_i(\hat{b}^c) \leq \tilde{\pi}_i(b^c)$ *for all $i \in N_c$; and*
- $\tilde{r}_i(\hat{b}^c) \geq \tilde{r}_i(b^c) - p$ *for all $i \in N_c$.*

Then, by exploiting Lemma 5 we can prove the following Lemma 6. Intuitively, the lemma shows that, given a probability $p \in [0, 1]$ and an optimal discretized bidding strategy, one can find an approximate solution to Problem (1) in polynomial time.

Lemma 6. *Given $p \in [0, 1]$ and an optimal discretized bidding strategy $\hat{b}^c \in \arg\max_{b^c \in B_p^c} \sum_{i \in N_c} \tilde{r}_i(b^c) - \tilde{\pi}_i(b^c)$, we can find in polynomial time a p-IR (see Equation (2)) and AIR solution to Problem (1) with value at least $OPT - pn_c$, where OPT is the optimal value of Problem (1).*

By Theorem 2, for any $p \in [0, 1]$, it is possible to find an optimal discretized bidding strategy $\hat{b}^c$ in time polynomial in the instance size and in $\frac{1}{p}$, since, as it is easy to check, the number of possible discretized bids in B_p^c is polynomial in $\frac{1}{p}$. Moreover, by employing Lemma 6, we can use the bidding strategy $\hat{b}^c$ to find an approximated solution to Problem (1) in polynomial time. Hence, given any $\varepsilon > 0$, it is sufficient to choose $p \in [0, 1]$ so that $\frac{1}{p} \in \text{poly}(\frac{1}{\varepsilon}, n)$ in order to obtain an ε-IR (see Equation (2)) and AIR approximate solution to Problem (1), as stated by the following theorem.

Theorem 3. *Given $\varepsilon > 0$, there exists an algorithm that runs in time polynomial in the instance size and $\frac{1}{\varepsilon}$, which returns an ε-IR and AIR solution to Problem (1) with value at least $OPT - \varepsilon$, where OPT is the optimal value of Problem (1).*

6 Transfers with Limited Liability Setting

In this section, we provide an approximate solution to the media agency problem with limited liability constraints. In particular, similarly to the arbitrary transfers setting, we design a bi-criteria additive FPTAS that returns solutions providing an arbitrary small loss $\varepsilon > 0$ with respect to the optimal value of the problem, by relaxing the IR constraints by ε.

As a first step, we show that we can restrict Problem (1) with LL constraints to work with the set B_p^c by only incurring in a small loss in the objective function value and IR constraints satisfaction. Notice that, Problem (1) with LL constraints restricted to discretized bids does not only have a smaller optimal value than Problem (1) with LL constraints, but it can also result in an infeasible problem, since Assumption 1 is *not* necessarily satisfied for a discretized bidding strategies. However, we can prove that, given a probability $p \in [0, 1]$, the following LP (7) that uses only bids in B_p^c and relaxes the IR constraints of quantity p is feasible and has value at least $OPT - pn_c$, where OPT is the optimal value of Problem (1) with LL constraints.

$$\max_{q \geq 0, \gamma \in \Delta_{B_p^c}} \sum_{b^c \in B_p^c} \gamma_{b^c} \sum_{i \in N_c} \tilde{r}_i(b^c) - \tilde{\pi}_i(b^c) \quad \text{s.t.} \tag{7a}$$

$$\sum_{b^c \in B_p^c} \gamma_{b^c} \tilde{r}_i(b^c) - q_i \geq t_i - p \qquad \forall i \in N_c \tag{7b}$$

$$\sum_{i \in N_c} q_i \geq \sum_{b^c \in B_p^c} \gamma_{b^c} \sum_{i \in N_c} \tilde{\pi}_i(b^c). \tag{7c}$$

Formally, we prove the following:

Lemma 7. *LP (7) is feasible. Moreover, the optimal value of LP (7) is at least $OPT - pn_c$, where OPT is the optimal value of Problem (1) with LL constraints.*

Next, we provide an algorithm to solve LP (7) by using the ellipsoid method. To do that, we use the dual LP (8), in which variables $y = \{y_1, \dots, y_{n_c}\}$, x, and z are related to Constraints (7b), (7c), and $\gamma \in \Delta_{B_p^c}$, respectively.

$$\min_{y \leq 0, x, z} \sum_{i \in N_c} (t_i - p) y_i + z \quad \text{s.t.} \tag{8a}$$

$$\sum_{i \in N_c} y_i \tilde{r}_i(b^c) - x \sum_{i \in N_c} \tilde{\pi}_i(b^c) + z \geq \sum_{i \in N_c} \tilde{r}_i(b^c) - \tilde{\pi}_i(b^c) \qquad \forall b^c \in B_p^c \tag{8b}$$

$$- y_i + x \geq 0 \qquad \forall i \in N_c. \tag{8c}$$

By Lemma 7, the primal LP (7) is feasible (and bounded), and, thus, it holds strong duality. As a consequence, in order to provide a polynomial-time algorithm to solve LP (7), it is enough to apply the ellipsoid method to the dual LP (8), which can be done in polynomial time since the latter has polynomially-many variables and exponentially-many constraints. This is possible by providing a polynomial-time separation oracle that, given an assignment of values to the variables as input, returns a violated constraint (if any). Since there are only polynomially-many Constraints (8c), we can check if one of them is violated in polynomial time. Moreover, in order to find whether there exists a violated Constraint (8b), it is sufficient to solve the weighted utility problem in Equation (4) by setting $\hat{y}_i = (1 - y_i)$ for each $i \in N_c$ and $\hat{x} = x - 1$. By Theorem 2, this can be done in polynomial time by computing a shortest path of a suitable graph. Hence, we can prove the following theorem.

Theorem 4. *Given $\varepsilon > 0$, there exists an algorithm that runs in time polynomial in the instance size and $\frac{1}{\varepsilon}$ and returns an ε-IR (see Equation (2)) and AIR solution to Problem (1) with LL constraints having value at least $OPT - \varepsilon$, where OPT is the optimal value of Problem (1) with LL constraints.*

[7] Our algorithm runs in time logarithmic in $\frac{1}{\eta}$ and hence polynomial in the size of the binary representation of bids in b^e.

[8] In the following, we ignore the loss in cumulative expected utility that results from the introduction of $\tau > 0$. Notice that this parameter is only necessary to induce specific tie-breaking rules and our results can be easily extended to consider the loss in utility due to τ. Moreover, τ can be taken to be exponentially small in the size of the problem instance, and hence negligible.

References

[Ashlagi *et al.*, 2009] Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Mediators in position auctions. *Games and Economic Behavior*, 67:2–21, 2009.

[Bachrach *et al.*, 2014] Yoram Bachrach, Sofia Ceppi, Ian Kash, Peter Key, and David Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the 15th ACM conference on Economics and computation*, pages 75–92, 2014.

[Bachrach, 2010] Yoram Bachrach. Honor among thieves: collusion in multi-unit auctions. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 617–624, 2010.

[Castiglioni *et al.*, 2022a] Matteo Castiglioni, Diodato Ferraioli, Nicola Gatti, Alberto Marchesi, and Giulia Romano. Efficiency of ad auctions with price displaying. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, 2022.

[Castiglioni *et al.*, 2022b] Matteo Castiglioni, Alessandro Nuara, Giulia Romano, Giorgio Spadaro, Francesco Trovò, and Nicola Gatti. Safe online bid optimization with return-on-investment and budget constraints subject to uncertainty. *CoRR*, abs/2201.07139, 2022.

[Chui *et al.*, 2018] Michael Chui, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, and Sankalp Malhotra. Notes from the AI frontier insights from hundreds of use cases. *McKinsey Global Institute*, 2018.

[Decarolis and Rovigatti, 2017] Francesco Decarolis and Gabriele Rovigatti. Online Auctions and Digital Marketing Agencies. Working Papers 17-08, NET Institute, September 2017.

[Decarolis *et al.*, 2020] Francesco Decarolis, Maris Goldmanis, and Antonio Penta. Marketing Agencies and Collusive Bidding in Online Ad Auctions. *Management Science*, 66(10):4433–4454, 2020.

[Dijkstra, 1959] Edsger Wybe Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.

[Grötschel *et al.*, 1981] Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

[He *et al.*, 2013] Di He, Wei Chen, Liwei Wang, and Tie-Yan Liu. A game-theoretic machine learning approach for revenue maximization in sponsored search. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 206–212, 2013.

[IAB, 2021] IAB. Interactive Advertising Bureau (IAB) internet advertising revenue report, Full year 2020 results, 2021.

[Lorenzon, 2018] Emmanuel Lorenzon. Collusion With a Rent-Seeking Agency in Sponsored Search Auctions. Technical report, SSRN, 2018.

[Nisan and Ronen, 2001] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic behavior*, 35(1-2):166–196, 2001.

[Nuara *et al.*, 2018] Alessandro Nuara, Francesco Trovò, Nicola Gatti, and Marcello Restelli. A combinatorial-bandit algorithm for the online joint bid/budget optimization of pay-per-click advertising campaigns. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 2379–2386, 2018.

[Nuara *et al.*, 2019] Alessandro Nuara, Nicola Sosio, Francesco Trovò, Maria Chiara Zaccardi, Nicola Gatti, and Marcello Restelli. Dealing with interdependencies and uncertainty in multi-channel advertising campaigns optimization. In *The World Wide Web Conference*, pages 1376–1386, 2019.

[OECD, 2017] OECD. *Algorithms and Collusion: Competition policy in the digital era.* 2017.

[Peleg and Sudhölter, 2007] Bezalel Peleg and Peter Sudhölter. Introduction to the theory of cooperative games. *American Economic Review*, 2007.

[Romano *et al.*, 2022] Giulia Romano, Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. The power of media agencies in ad auctions: Improving utility through coordinated bidding. *CoRR*, abs/2204.13772, 2022.

[Varian and Harris, 2014] Hal R. Varian and Christopher Harris. The vcg auction in theory and practice. *American Economic Review*, 104(5):442–45, May 2014.

[Vorobeychik and Reeves, 2008] Y. Vorobeychik and D. M. Reeves. Equilibrium analysis of dynamic bidding in sponsored search auctions. *International Journal of Electronic Business*, 6:172–193, 2008.