

# Fairness in the Autobidding World with Machine-learned Advice

Yuan Deng

Google Research, dengyuan@google.com

Negin Golrezaei

Sloan School of Management, Massachusetts Institute of Technology, golrezae@mit.edu

Patrick Jaillet

Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, jaillet@mit.edu

Jason Cheuk Nam Liang

Operations Research Center, Massachusetts Institute of Technology, jcnliang@mit.edu

Vahab Mirrokni

Google Research, mirrokni@google.com

The increasing availability of real-time data has fueled the prevalence of algorithmic bidding (or *autobidding*) in online advertising markets, and has enabled online ad platforms to produce signals through machine learning techniques (i.e., ML advice) on advertisers' true perceived values for ad conversions. Previous works have studied the auction design problem while incorporating ML advice through various forms to improve total welfare of advertisers. Yet, such improvements could come at the cost of individual bidders' welfare, consequently eroding fairness of the ad platform. Motivated by this, we study how ad platforms can utilize ML advice to improve welfare guarantees and fairness on the individual bidder level in the autobidding world. We focus on a practical setting where ML advice takes the form of lower confidence bounds (or confidence intervals). We motivate a simple approach that directly sets such advice as personalized reserve prices when the platform consists of value-maximizing autobidders who are subject to return-on-ad spent (ROAS) constraints competing in multiple parallel auctions. Under parallel VCG auctions with ML advice-based reserves, we present a worst-case welfare lower-bound guarantee for individual agents, and show that platform fairness is positively correlated with ML advice quality. We also present an instance that demonstrates our welfare guarantee is tight. Further, we prove an impossibility result showing that no truthful, and possibly randomized mechanism with anonymous allocations and ML advice as personalized reserves can achieve universally better fairness guarantees than VCG when coupled with ML advice of the same quality. Finally, we extend our fairness guarantees with ML advice to generalized first price (GFP) and generalized second price (GSP) auctions.

*Key words:* Fairness, mechanism design, machine-learned advice, welfare maximization

---

## 1. Introduction

Autobidding, namely the procedure of adopting automated algorithms to procure ad slots in online ad auctions, has become the prevalent mode of bidding in online advertising markets, contributing to more than 88% of total online advertising traffic (Shelagh Dolan 2019). Autobidding is generally conducted by autobidders who bid programatically in online ad auctions on behalf of advertisers to optimize for their high-level goals, such as maximizing total clicks/views or return on ad spent

(ROAS). In such autobidding environments, online ad platforms generally seek to increase total welfare, or, in other words, total advertiser value for conversions among all bidders. Despite not knowing the true values of bidders, ad platforms can improve total welfare via carefully designing suitable ad auction mechanisms with the aid of predictive signals generated from machine-learning models (i.e., *ML advice*) such as predictions on autobidders’ perceived value for conversions. Nevertheless, many real-world platforms and mechanism design literature have overlooked how these tailored mechanisms can impact each individual bidder’s welfare. In fact, ad auctions may tend to increase bidders’ aggregate welfare at the cost of certain individual bidders’ acquired value, consequently eroding fairness of the platform. Motivated by this, *we study the fairness of online ad auctions in an autobidding world, and further investigate how platforms can utilize ML advice for advertiser values to improve overall platform fairness.*

In this work, we study a prototypical autobidding setting where autobidders compete simultaneously in several multi-slot auctions that are run in parallel, and aim to maximize total advertiser value under *return-on-ad-spent (ROAS)* constraints that restrict total spend of a bidder to be less than her total acquired value across all auctions (Aggarwal et al. 2019, Deng et al. 2021, Balseiro et al. 2021a, Mehta 2022). For auctions that are truthful with respect to (quasi-linear) utility maximizers (e.g., second-price auctions), the value-maximizing decision for an autobidder under full information is a single bid multiplier that will be multiplied by the true values of each auction to yield corresponding bid values (Aggarwal et al. 2019). The ad platform, on the other hand, possesses ML advice on autobidders’ real values with a certain degree of accuracy/quality, and decides on the auction format for the parallel auctions. Here, we restrict our attention to ML advice that takes the form of lower confidence bounds or confidence intervals (see Definition 3.1 and corresponding discussions), and motivate the approach to directly set such advice as personalized reserve prices.

Our definition of fairness concerns worst case guarantees for individual bidder welfare relative to the welfare under the efficient outcome, i.e., the outcome that maximizes total welfare. In other words, an autobidding platform is fair if in the worst-case outcome, every bidder would be able to retain some proportion of the welfare they would have acquired under the efficient outcome, where the higher the proportion, the fairer the platform. The overarching goal of this work is to present theoretical welfare guarantees as a measure of platform fairness for individual bidders in the presence of reserves derived from ML advice of certain accuracy/quality under typical auction formats such as the Vickrey–Clarke–Groves (VCG), generalized second price (GSP), generalized first price (GFP) auctions, etc. Our main contributions are summarized as followed.

**Welfare/fairness guarantees in VCG with ML advice.** To the best of our knowledge, this work presents the first notion of individual bidder fairness in an autobidding mechanism design setup, and presents the first welfare guarantee on the individual bidder level. We motivate the approach

that sets lower-confidence type of ML advice on advertisers’ values as personalized *approximate reserve prices* (see Section 3) with the goal to improve platform fairness. When autobidders compete in multiple parallel VCG auctions with such advice-based approximate reserves, under any feasible outcome (i.e. worst-case outcome), we present a welfare lower bound for each individual bidder that increases in the quality of ML advice, and decreases in the ratio between competitors’ and the bidder’s total values in the efficient outcome (Theorem 4.1). We also present an instance that shows our lower bound is tight (Theorem 4.3). Together with the results in (Deng et al. 2021) stating approximate reserves can improve total welfare in the autobidding setting, we conclude that incorporating ML advice as personalized reserves not only increases total welfare of the platform, but also improves fairness for bidders.

**Impossibility result: VCG is the fairest among a broad class of auctions.** We show an impossibility result that says no allocation-anonymous, truthful, and possibly randomized mechanism with ML advice as approximate reserves can achieve a strictly better worst-case welfare lower bound guarantee than the VCG auction coupled with ML advice of the same quality; see Theorem 5.1. In particular, for any allocation-anonymous, truthful, and possibly randomized auction, we construct an autobidding instance with approximate reserves, and show that there must be at least 1 bidder whose welfare is at most the welfare lower bound guarantee we presented under VCG (i.e. Theorem 4.1).

**Fairness guarantee extensions to GSP and GFP.** We extend fairness guarantee results to GSP and GFP auctions, and show that a similar worst-case welfare lower bound guarantee continues to hold when autobidders submit *undominated bids* (Theorem 6.2). We compare these lower bound guarantees with that of VCG in Theorem 4.1, and identify conditions under which VCG outperforms (or underperforms) GSP/GFP in terms of fairness with the same ML advice quality.

### 1.1. Related Works

**Mechanism design with constrained bidders.** Our work is related to the general theme of mechanism design in the presence of constrained agents. The works Pai and Vohra (2014), Golrezaei et al. (2021b) study revenue-optimal auction design when bidders who maximize quasi-linear utility are constrained by budgets, and return-on-investment (ROI), respectively. Balseiro et al. (2021d) study revenue-maximizing auctions for ROI constrained bidders under different objectives and information structures for values and ROI targets. This work differs from these papers as we do not study new auction formats and platform revenue-optimization, but instead presents insights into how incorporating ML advice as reserves in classic auctions like VCG, GSP and GFP can improve individual bidder welfare. To the best of our knowledge, the most relevant works to this paper are Deng et al. (2021), Balseiro et al. (2021a), Mehta (2022), where they consider the same auto-bidding setting

(i.e. value-maximizers with ROAS constraints) as ours. Deng et al. (2021), Balseiro et al. (2021a), Mehta (2022), Deng et al. (2022) all present techniques to improve price-of-anarchy bounds for the total welfare of any feasible outcome: where Deng et al. (2021) relies on additive boosts on bid values, Balseiro et al. (2021a), Deng et al. (2022) utilizes approximate reserve prices derived from ML-advice, and Mehta (2022) develops randomized allocation and payment rules. Our work distinguishes itself from these works as we focus on welfare and fairness guarantees on the individual bidder level. We point out that our proof techniques also differ from those in Deng et al. (2021), Balseiro et al. (2021a), Mehta (2022), Deng et al. (2022) as our individual fairness guarantees require novel analyses on the value-expenditure tradeoffs individual bidders' would face when they are tempted to outbid others to acquire more value; see discussion in Section 4 for more details.

**Exploiting machine-learned advice.** ML advice has been utilized in various applications to improve learning and decision making. For example, Wang et al. (2020) exploits ML advice to develop algorithms for the multi-shop ski-rental problem, (Lykouris and Vassilvtiskii 2018) adopts ML advice for the caching problem, and (Indyk et al. 2022) studies online page migration with ML advice. However, although many works in online advertising studied predictive models for advertiser values, click through rates, etc (see e.g. (Richardson et al. 2007, Lee et al. 2012, Sodomka et al. 2013)), the literature on applying such predictions (or more generally, ML advice) to the mechanism design problem has been scarce. One related work along this direction is Munoz and Vassilvtiskii (2017), which develops a theoretical framework to optimize reserve prices in a posted price mechanism by utilizing prediction inputs on bid values. In this work, we do not optimize for reserves, and motivate the simple approach of setting reserves using ML advice. Finally, we note that our work contributes to the area of exploiting ML advice to designing mechanisms with for improving welfare guarantees and fairness for individual bidders.

**Fairness in algorithm design and optimization.** The notion of fairness considered in this paper differs from algorithmic fairness and optimization fairness in the following sense. Algorithmic fairness for supervised learning concerns maintaining prediction/treatment similarity for individuals agents/groups alike; see Dwork et al. (2012), Kusner et al. (2017) for individual fairness, Calders et al. (2009), Zliobaite (2015) for group fairness, and Kearns et al. (2019,?) for subgroup fairness. Also see Mehrabi et al. (2021) for a comprehensive survey on fairness in ML algorithms. On the other hand, fairness in optimization problems generally concerns decisions to be fair in a domain-specific sense, for example, in assortment planning similar products should be offered similar visibility Chen et al. (2022); in online matching, there should be no discrimination across individual agents Ma et al. (2021); in resource allocation, each individual should receive a non-negligible amount of allocation Bateni et al. (2022). In this work, fairness of an autobidding platform refers to providing worst case outcome welfare guarantees for each individual bidder, with the hope that

every bidder can retain some proportion of welfare that she would have attained under the efficient outcome (i.e. the outcome with total maximum welfare).

For more related works on algorithmic bidding/learning under constraints, and reserve price optimization, we refer readers to Appendix A for an extended literature review.

## 2. Preliminaries

### 2.1. Autobidding instance

Consider a platform running an autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$  where there are  $N$  bidders bidding in  $M$  parallel position auctions Lahaie et al. (2007), Varian (2007) of the format  $\mathcal{A}$  (e.g. VCG, GSP, GFP etc.) that are labeled  $A_1 \dots A_M$ . Here,  $\mathbf{r} = (r_{i,j}) \in \mathbb{R}_{\geq 0}^{N \times M}$ , where  $r_{i,j}$  is the personalized reserve price for bidder  $i \in [N]$  in auction  $A_j$ ; and  $\mathbf{v} = (v_{i,j}) \in \mathbb{R}_{\geq 0}^{N \times M}$  where  $v_{i,j}$  is the value of bidder  $i$  in auction  $j$ . For ease of analyses assume that no pair of bidders have the same value in an auction, i.e.  $v_{i,j} \neq v_{i',j}$  for any two bidders  $i \neq i'$  and any  $j \in [M]$ . In each auction  $j$ , let there be  $S_j$  available slots, where slot  $\ell \in [S_j]$  is associated with position discount  $\mu_j(\ell)$  and  $1 \geq \mu_j(1) > \mu_j(2) > \dots > \mu_j(S_j) > 0$ . Note that the position discount of a slot can be interpreted as its click-through-rate (see more details in Lahaie et al. (2007), Varian (2007)), and therefore slots with higher ranks (i.e. smaller indices) are more likely to be clicked and thus have larger position discounts. For positions/ranks below the available slots, i.e.  $S_j + 1, \dots, N$ , without loss of generality, we set  $\mu_j(S_j + 1) = \dots = \mu_j(N) = 0$ .

**Allocation and auction outcomes.** The platform solicits bids from bidders  $\mathbf{b} = (b_{i,j})_{i \in [N], j \in [M]} \in \mathbb{R}_{\geq 0}^{N \times M}$ , where  $b_{i,j}$  is the bid from bidder  $i$  in auction  $j$ . Then, it ranks bidders by their bids in each auction (tie-breaking by bidder indices) and allocates the  $\ell$ th slot of an auction to the bidder who is ranked in position  $\ell$  (i.e. who has the  $\ell$ th highest bid) if she clears her reserve, i.e. if  $b_{i,j} \geq r_{i,j}$ . Note that for ease of presentation, we consider a lazy implementation of personalized reserves such that a slot is not allocated to a bidder if her bid does not clear her reserve, but we point out that it will become clear later that our results continue to hold for eager reserve implementations Paes Leme et al. (2016). We refer to these allocations as the *outcome* of the autobidding instance, defined as  $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_M)$  where  $\mathbf{x}_j = (x_{i,j,\ell})_{i \in [N], \ell \in [S_j]} \in \{0, 1\}^{N \times S_j}$  is the outcome vector in auction  $j$ , and further  $x_{i,j,\ell} = 1$  only if bidder  $i$  ranks in position  $\ell$  in auction  $j$  while clearing her reserve, and 0 otherwise. Note that each bidder  $i$  can win at most one slot in an auction  $j$ , i.e.  $\sum_{\ell \in [S_j]} x_{i,j,\ell} \leq 1$ , where we have  $\leq$  because a bidder may not win any slots or does not clear her reserve. For any outcome  $\mathbf{x}$ , let  $\ell_{i,j}$  be the position/ranking of bidder  $i$  in auction  $j$ , and define  $W_{i,j}(\mathbf{x}) = \mu_j(\ell_{i,j}) \cdot v_{i,j} \cdot x_{i,j,\ell}$  as bidder  $i$ 's acquired welfare in auction  $j$ . Note that if  $i$  does not win a slot in auction  $j$ , her position  $\ell_{i,j} > S_j + 1$ , and her position discount  $\mu_j(\ell_{i,j}) = 0$ . The welfare contribution of bidder  $i$  is then defined as  $W_i(\mathbf{x}) = \sum_{j \in [M]} W_{i,j}(\mathbf{x})$ , and the total welfare of outcome  $\mathbf{x}$  is  $W(\mathbf{x}) = \sum_{i \in [N]} W_i(\mathbf{x})$ .

**Payments.** Once the platform determines allocations and outcomes according to bids  $\mathbf{b} \in \mathbb{R}_{\geq 0}^{N \times M}$ , the platform charges payment  $p_{i,j}(\mathbf{b})$  to bidder  $i$  for auction  $j$  according to the payment rules. For illustrative purposes, the following includes example payment rules for several classic auctions.

**Example 2.1 (Example payment rules)** Suppose bidders submit bids  $\mathbf{b} = (b_{i,j}) \in \mathbb{R}_{\geq 0}^{N \times M}$ , and each bidder  $i$  is associated with a personalized reserve  $r_{i,j} \in \mathbb{R}_{\geq 0}$  for auction  $j$ . We denote  $(\hat{b}_{1,j} \dots \hat{b}_{N,j})$  to be the bids' order statistics where  $\hat{b}_{\ell,j}$  is the  $\ell$ th largest bid in auction  $j$ . Then, for any bidder  $i$  who won a slot  $\ell_{i,j} \in [S_j]$  in auction  $j$ , the following shows her corresponding payment  $p_{i,j}$  under different auction formats: (1) VCG:  $p_{i,j}^{VCG}(\mathbf{b}) = \sum_{\ell=\ell_{i,j}}^{S_j} (\mu_j(\ell) - \mu_j(\ell+1)) \cdot \max\{\hat{b}_{\ell+1,j}, r_{i,j}\}$ ; (2) GSP:  $p_{i,j}^{GSP}(\mathbf{b}) = \mu_j(\ell_{i,j}) \cdot \max\{\hat{b}_{\ell_{i,j}+1}, r_{i,j}\}$ ; (3) GFP:  $p_{i,j}^{GFP}(\mathbf{b}) = \mu_j(\ell_{i,j}) \cdot \max\{\hat{b}_{\ell_{i,j}}, r_{i,j}\}$ . It is well known that for the same bid profile  $\mathbf{b}$ , we have the following payment dominance relation:  $p_{i,j}^{GFP}(\mathbf{b}) \geq p_{i,j}^{GSP}(\mathbf{b}) \geq p_{i,j}^{VCG}(\mathbf{b})$  for all  $i \in [N]$  and  $j \in [M]$  (see e.g. Edelman et al. (2007)).

## 2.2. Objectives and behavior of autobidders

We consider the setting where bidders aim to maximize total value over  $M$  auctions, subject to a *return-on-ad-spent (ROAS)* constraint which guarantees that the total expenditure is less than total acquired value across all auctions.<sup>1</sup> Mathematically, bidder  $i$ , when fixing other bidders' bids  $\mathbf{b}_{-i} \in \mathbb{R}_{\geq 0}^{(N-1) \times M}$ , decides on bids  $\mathbf{b}_i \in \mathbb{R}_{\geq 0}^M$  with the following optimization problem:

$$\max_{\mathbf{b}_i \in \mathbb{R}_{\geq 0}^M} W_i(\mathbf{x}(\mathbf{b}_i, \mathbf{b}_{-i})) \quad \text{s.t.} \quad W_i(\mathbf{x}(\mathbf{b}_i, \mathbf{b}_{-i})) \geq \sum_{j \in [M]} p_{i,j}(\mathbf{b}_i, \mathbf{b}_{-i}). \quad (1)$$

It has been shown in previous literature (Aggarwal et al. 2019, Deng et al. 2021, Balseiro et al. 2021a, Mehta 2022) that for truthful (i.e. ex-post incentive compatible) auctions such as VCG, the optimal bidding strategy for any autobidder  $i$  is to submit bid values of the form  $b_{i,j} = \alpha_i v_{i,j}$  for auction  $j \in [M]$ , where  $\alpha_i > 0$  is a bid multiplier that is used across all auctions. Such a bidding strategy is called *uniform bidding*. Note that for autobidders, taking any uniform bid multiplier  $\alpha_i < 1$  is weakly dominated by setting  $\alpha_i = 1$ . This is because by setting  $\alpha_i = 1$ , the bidder can acquire larger total value as compared to setting  $\alpha_i < 1$ . At the same time, the bidder bids truthfully in all auctions, and because the payment in any auction is no greater than the winning bid value in truthful auctions, the ROAS constraint is always satisfied under  $\alpha_i = 1$ . Therefore, we focus on undominated bid-multipliers, i.e.  $\alpha_i \in [1, \infty)$  for all  $i \in [N]$ .

<sup>1</sup> A more general concept related to ROAS is *return-on-investment (ROI)*, where each bidder  $i$  has a target ROI ratio  $T_i$  such that her constraint in Equation (1) is instead written as  $W_i(\mathbf{x}(\mathbf{b}_i, \mathbf{b}_{-i})) \geq T_i \cdot \sum_{j \in [M]} p_{i,j}(\mathbf{b}_i, \mathbf{b}_{-i})$ ; see e.g. Golrezaei et al. (2018, 2021a). In this paper, since we study worst-case autobidding instances, we can divide all bidder  $i$ 's values by  $T_i$  so it is without loss of generality to consider ROAS constraints.

**Remark 2.1** *In the paper, truthfulness in auctions is always w.r.t. quasi-linear utility maximizers. For bidders who have constraints or do not maximize quasi-linear utility, e.g., autobidders with objectives as in Equation (1), truthful bidding may no longer be a dominant strategy in truthful auctions.*

In light of the above discussion on uniform bidding, in the following Sections 4 and 5, we study the truthful VCG auctions, and hence only focus on uniform bidding strategies. Later in Section 6 when we extend our results to non-truthful auctions such as GSP and GFP, we will discuss more general bidding strategies, in particular, non-uniform bidding.

### 2.3. Efficient auction outcomes, welfare guarantees and fairness

Let  $\ell_{i,j}^{\text{OPT}}$  be the position of bidder  $i$  in auction  $j$  when ranked according to true values  $\mathbf{v} \in \mathbb{R}_{\geq 0}^{N \times M}$ . Then we call the outcome  $\mathbf{x}^{\text{OPT}}$  with  $x_{i,j,\ell}^{\text{OPT}} = \mathbb{I}\{\ell = \ell_{i,j}^{\text{OPT}}\}$ , the *efficient outcome*. Note that  $\mathbf{x}^{\text{OPT}}$  is called efficient because it yields the largest total welfare, i.e.  $\mathbf{x}^{\text{OPT}} = \arg \max_{\mathbf{x}} W(\mathbf{x})$ . Similar to our definition for welfare of an outcome in Section 2.1, let  $\text{OPT}_{i,j} = \mu_j(\ell_{i,j}^{\text{OPT}}) \cdot v_{i,j}$ ,  $\text{OPT}_i(\mathbf{x}) = \sum_{j \in [M]} \text{OPT}_{i,j}(\mathbf{x})$ , and  $\text{OPT}(\mathbf{x}) = \sum_{i \in [N]} \text{OPT}_i(\mathbf{x})$  be the welfare of bidder  $i$  in auction  $j$ , total welfare contribution of bidder  $i$ , and total welfare, respectively, under the efficient outcome. Finally, we denote  $\mathcal{F}$  as the set of all feasible outcomes that result from some bid profile under which every bidder's ROAS constraint is satisfied (see Equation (1)).

Our goal is to present a characterization for how model primitives of the autobidding instance  $\Theta$  impact overall fairness of bidders defined as followed:

**Definition 2.1 ( $\delta$ -fairness)** *For any  $\delta \in [0, 1]$ , an autobidding instance is  $\delta$ -fair for bidder  $i$  if  $\min_{\mathbf{x} \in \mathcal{F}} \frac{W_i(\mathbf{x})}{\text{OPT}_i} \geq \delta$ .*

The motivation of our notion of fairness in Definition 2.1 is to quantify the potential misalignment between optimizing individual bidder welfare, and optimizing total welfare (across all bidders) which is arguably one of the most foundational topics in the domain of mechanism design. As auctioneers develop mechanisms to approximate efficient outcomes that optimize total welfare, there is a lack of consideration for the welfare of individual bidders, and the improvement in total welfare due to such mechanisms may come at the cost of individual-bidder welfare. Hence, Definition 2.1 presents a metric for auctioneers to characterize the potential discrepancies between bidder and auctioneer welfare objectives. We also point out that our definition of fairness can be viewed as an analogue to the notion of price of anarchy (POA) that measures the worst-case total welfare achieved amongst all equilibrium compared to the optimal total welfare (see e.g. Roughgarden (2015) for details). In particular, our notion of fairness in Definition 2.1 measures each individual bidder's worst-case welfare w.r.t her welfare achieved under the efficient outcome.

In addition to individual bidder fairness, we also consider the loss in welfare compared to a bidder’s welfare contribution under the efficient outcome, formally defined as followed:

**Definition 2.2 (Welfare loss w.r.t. efficient outcome)** For an outcome  $\mathbf{x}$  and any bidder  $i \in [N]$ , let  $\mathcal{L}_i(\mathbf{x}) = \{j \in [M] : W_{i,j}(\mathbf{x}) < \text{OPT}_{i,j}\}$  be the set of auctions in which bidder  $i$ ’s acquired welfare is less than that of her welfare under the efficient outcome. Then, we define the loss of bidder  $i$  under outcome  $\mathbf{x}$  w.r.t. the efficient outcome  $\mathbf{x}^{\text{OPT}}$  (or simply “welfare loss” in the rest of the paper):

$$\text{LOSS}_i(\mathbf{x}) = \sum_{j \in \mathcal{L}_i(\mathbf{x})} (\text{OPT}_{i,j} - W_{i,j}(\mathbf{x})) . \quad (2)$$

**Remark 2.2** For any outcome  $\mathbf{x}$ , let  $\ell_{i,j}$  be the position (i.e. ranking) of bidder  $i$  in auction  $j$ , and recall that  $\ell_{i,j}^{\text{OPT}}$  is the position of bidder  $i$  in auction  $j$  under the efficient outcome  $\mathbf{x}^{\text{OPT}}$ . Then, the set  $\mathcal{L}_i(\mathbf{x}) = \{j \in [M] : W_{i,j}(\mathbf{x}) < \text{OPT}_{i,j}\}$  can also be interpreted as the set of auctions where bidder  $i$ ’s ranking under  $\mathbf{x}$  is lower than her ranking  $\mathbf{x}^{\text{OPT}}$ , or in other words the set of auctions that incur a welfare loss w.r.t.  $\mathbf{x}^{\text{OPT}}$ . Hence we can also rewrite  $\mathcal{L}_i(\mathbf{x}) = \{j \in [M] : \ell_{i,j} > \ell_{i,j}^{\text{OPT}}\}$ .

The following proposition connects the notion of welfare loss (as in Definition 2.2) and fairness (as in Definition 2.1). The proposition shows that an upper bound on welfare loss can be directly translated into a welfare lower bound that corresponds to our fairness notion. See Appendix B.1 for the proof.

**Proposition 2.1 (Translating loss to fairness)** Assume for bidder  $i \in [N]$  and outcome  $\mathbf{x}$  we have  $\text{LOSS}_i(\mathbf{x}) \leq B$  for some  $B > 0$ . Then,  $\frac{W_i(\mathbf{x})}{\text{OPT}_i} \geq 1 - \frac{B}{\text{OPT}_i}$ .

### 3. Incorporating ML advice for bidder values as personalized reserve prices

With modern machine learning (ML) models and frameworks, online ad platforms can utilize available historical data (e.g. bid logs, keyword characteristics, user profiles, etc.) to produce predictive signals (or more generally ML advice) on autobidders’ values. In this work, we specifically focus on ML advice that take the form of a *lower-confidence bound* of true advertiser values.

Our key approach to incorporate this type of ML advice in our autobidding setting, is via simply setting personalized reserve prices to be the lower confidence bound for each bidder’s value. To motivate this approach, consider the following example:

**Example 3.1 (Motivating example)** Consider 2 bidders competing in 2 second-price auctions (SPA) whose values are indicated in the following table with some  $v > 0$ .



	Auction 1	Auction 2
bidder 1	$v$	$0$
bidder 2	$\frac{v}{2}$	$v$

For illustrative convenience suppose bidder 1 sets her bid multiplier to be  $\alpha_1 = 1$ . Then when her competitor bidder 2 sets a multiplier  $\alpha_2 > 2$ , bidder 2 will win both auctions and acquire a total value of  $v_{2,1} + v_{2,2} = \frac{3}{2}v$  while submitting a payment of  $\alpha_1(v_{1,1} + v_{1,2}) = v$ . In this case, bidder 2 satisfies her ROAS constraint and extracts all bidder 1's welfare, leaving her with no value. We also highlight that this bid multiplier profile constitutes an equilibrium, because bidder 1 cannot raise her bid multiplier to outbid bidder 2 for auction 1, since with  $\alpha_2 > 2$  bidder 1 would violate her ROAS constraint if she bids more than  $\alpha_2 v_{2,1} > v$ .

Now suppose for each value  $v_{i,j}$  ( $i, j \in [2]$ ), the platform possesses a lower-confidence type of ML advice, namely  $(\hat{v}_{i,j})_{i,j \in [N]}$  such that  $\beta v_{i,j} \leq \hat{v}_{i,j} < v_{i,j}$  for all  $v_{i,j} > 0$  for some  $\beta > \frac{1}{2}$ , and sets personalized reserves  $r_{i,j} = \hat{v}_{i,j}$ . If bidder 2 attempts to win both auctions by setting  $\alpha_2 > 2$ , her payment will be at least  $\max\{\beta v_{2,1}, \alpha_1 v_{1,1}\} + \max\{\beta v_{2,2}, \alpha_1 v_{1,2}\} = v + \beta v > \frac{3}{2}v$ , violating her ROAS constraint. Therefore, via incorporating ML advice as personalized reserves, bidder 1's competitor is prohibited from outbidding her in auction 1, and hence safeguarding bidder 1's welfare.

The main observation from the above example is that without reserve prices, bidder 2 acquires a large margin for her ROAS constraint by winning auction 2 where competition is small and cost is low. Therefore, she can raise her bid to outbid bidder 1 in auction 1 without violating her overall ROAS constraint by covering the high expenditure in auction 1 with her acquired value margin in auction 2. By setting personalized reserves properly, the platform can increase bidder 2's payment in auction 2, which in turn decreases the manipulative power of bidder 2 by reducing the ROAS value margin she can acquire from auction 2.

We formally characterize reserves prices with which the platform can reduce bidders' manipulative power using the notion of *approximate reserves* defined as followed.

**Definition 3.1** ( $\beta$ -accurate ML advice and approximate reserves) *In an autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$ , suppose the platform possesses a lower-confidence type of ML advice  $(\hat{v}_{i,j}^-)_{i,j \in [N]}$ . If  $\hat{v}_{i,j}^- \in [\beta v_{i,j}, v_{i,j}]$  with some  $\beta \in (0, 1)$  for all  $i, j$  and  $v_{i,j} \neq 0$ , we say that the ML advice is  $\beta$ -accurate. Further, if the platform sets  $r_{i,j} = \hat{v}_{i,j}^-$ , we say reserve prices  $\mathbf{r}$  are  $\beta$ -approximate.*

**Remark 3.1** *ML-advice in real-world online advertising settings generally concerns predicting advertiser values and takes the form of confidence intervals (see e.g. Shrestha and Solomatine (2006), Braga et al. (2007), Jiang et al. (2008), Dai et al. (2020) and references therein). We remark that such confidence intervals can be viewed as a special case of the lower-confidence type of ML advice in*

*Definition (3.1):* suppose the platform utilizes some ML model to predict the true value  $v_{i,j}$  of bidder  $i$  in auction  $j$ , and produces a confidence interval  $(\hat{v}_{i,j}^-, \hat{v}_{i,j}^+) \ni v_{i,j}$  with  $\hat{v}_{i,j}^-, \hat{v}_{i,j}^+ > 0$ . The platform can then choose personalized reserve  $r_{i,j} = \hat{v}_{i,j}^-$ , which is  $\beta$ -approximate for  $\beta = \frac{\hat{v}_{i,j}^-}{\hat{v}_{i,j}^+} \in (0, 1)$  because  $\beta v_{i,j} < \beta \hat{v}_{i,j}^+ = \hat{v}_{i,j}^- = r_{i,j} < v_{i,j}$ . Additionally, we remark that one can similarly handle ML-advice in the form of predictions that satisfy probabilistic concentration inequalities; e.g. suppose the predicted value  $\hat{v}_{i,j}$  satisfies  $|\hat{v}_{i,j} - v_{i,j}| < \eta$  with high probability (w.h.p) for some known  $\eta$ , then the confidence interval  $(\hat{v}_{i,j} - \eta, \hat{v}_{i,j} + \eta)$  contains  $v_{i,j}$  w.h.p, and one can set personalized reserve  $r_{i,j} = \hat{v}_{i,j} - \eta$ . Note that with such personalized reserves derived from probabilistic ML-advice, all results in this paper remain valid w.h.p.

We point out that our definition of ML advice and approximate reserves is general to any auction format. The gap between the lower bound  $\beta v_{i,j}$  and the true value  $v_{i,j}$  in Definition (3.1) represents the inaccuracies of the platform's ML advice. In other words,  $\beta$  can be perceived as a quality measure of the platform's ML advice for advertiser value, such that larger  $\beta$  represents better advice quality.

#### 4. Fairness guarantees for VCG with ML advice

In the motivating Example 3.1, we observe that ML advice and corresponding  $\beta$ -approximate reserves allow the platform to safeguard welfare for individual bidders by increasing payments and consequently limit the manipulative behavior of bidders who face significantly small competition in certain auctions. In this section, through the following Theorem 4.1, we characterize this intuition for the classic VCG auction, and present a quantitative measure for the relationship between overall fairness of the platform and ML advice when incorporated in the form of approximate reserves.

**Theorem 4.1 (Fairness lower bound for VCGs with  $\beta$ -approximate reserves)** *Consider autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$  where the auction format  $\mathcal{A}$  is the VCG auction, and reserve prices  $\mathbf{r}$  are  $\beta$ -approximate as in Definition 3.1. For any feasible outcome  $\mathbf{x} \in \mathcal{F}$  and bidder  $i \in [K]$  who adopts bid multiplier  $\alpha_i > 1$ , the loss and fairness (Equation (2) and Definition 2.1) are bounded as:*

$$\text{LOSS}_i(\mathbf{x}) \leq \frac{1-\beta}{\alpha_i-1} \text{OPT}_{-i} \quad \text{and} \quad \min_{\mathbf{x} \in \mathcal{F}} \frac{W_i(\mathbf{x})}{\text{OPT}_i} \geq 1 - \frac{1-\beta}{\alpha_i-1} \cdot \frac{\text{OPT}_{-i}}{\text{OPT}_i}$$

The key message for Theorem 4.1 is that with more accurate ML advice (i.e. larger  $\beta$ ), online ad platforms can set larger approximate reserves, and hence improve welfare and fairness guarantees for individual bidders. We also provide some intuition for the term  $\frac{1-\beta}{\alpha_i-1} \frac{\text{OPT}_{-i}}{\text{OPT}_i}$  in the bound. Increasing  $\beta$  (i.e. increasing reserve prices via improving ML quality) or increasing the bid multiplier  $\alpha_i$ , raises the cost for competitors to outbid bidder  $i$  in certain auctions, and hence makes it more difficult to cover her expenditures that arise from significant overbidding. This reduces competitors' manipulative

power, and in turn improves the welfare guarantees for bidder  $i$ . Note that this aligns with the intuition we obtained in Example 3.1. We also point out that it may be tempting for bidder  $i$  to set her multiplier  $\alpha_i$  to infinity (or as large as possible). However, by doing so bidder  $i$  might be winning too many auctions with payments exceeding her values, violating her ROAS constraint. In other words, there exists a tradeoff between large multipliers and ROAS feasibility. On the other hand,  $\frac{\text{OPT}_{-i}}{\text{OPT}_i}$  can be perceived as the relative (inverse) market share of bidder  $i$ , such that with a smaller market share she becomes more vulnerable to manipulative behavior of others, resulting in lower fairness guarantees.

*Proof sketch for Theorem 4.1.* Here, we provide a roadmap of our proof techniques for the special case where there is only a single slot in each VCG auction (i.e. all auctions are SPAs): bidder  $i$  would incur a loss w.r.t. the efficient outcome only if it loses an auction  $j \in \mathcal{L}_i$  (see Definition 2.2) where her value is the largest among all bidders. Since the uniform bid multiplier of  $i$  is strictly greater than 1, the winner of the auction, denoted as  $\kappa_j \neq i$ , pays an amount at least  $\alpha_i v_{i,j} > v_{i,j}$  which is strictly greater than her value  $v_{\kappa_j,j}$ . Thus, in order for  $\kappa_j$  to satisfy her ROAS constraint, she must cover her expenditure in auction  $j$  via acquiring value from other auctions. Following the same logic, all bidders who outbid bidder  $i$  in an auction  $j \in \mathcal{L}_i$  where bidder  $i$  has the largest value must cover their high expenditures, and the total amount they can cover is capped by a natural upper bound  $\text{OPT}_{-i}$ . Thus, this provides an upper bound on the the total value that bidder  $i$  can lose from auctions where she has the highest value, which further translates into a welfare guarantee using Proposition 2.1.

The proof for general multi-slot VCG is much more involved, since bidder  $i$  may partially lose welfare if in some auction she ranks lower than that under the efficient outcome, but still acquires a slot. We refer readers to Appendix C.1 for the full proof.

We point out that although our autobidding setup described in Sections 2.1, 2.2 and the notion of approximate reserves (Definition 3.1) are the same as those in Deng et al. (2021), Balseiro et al. (2021a)], our analyses and proof techniques are different, primarily because we focus on the welfare guarantees for individual bidders, where as Deng et al. (2021), Balseiro et al. (2021a) investigates total welfare for all bidders. In particular, in our proof we fix a bidder  $i$  and carefully analyze the amount of expenditure that could be covered by each competitor who outbids bidder  $i$  in auctions where  $i$  has the highest value, whereas the aforementioned related works takes an aggregate view to lower bound total welfare of all bidders. Nevertheless, Balseiro et al. (2021a) shows that approximate reserves improve the total welfare of all bidders, and therefore along with Theorem 4.1, we can see that incorporating  $\beta$ -accurate ML advice as approximate reserves not only benefits the entire platform’s total welfare, but also enforces a certain degree of individual fairness.

In light of the fairness guarantee presented in Theorem 4.1 the following corollary presents a sufficient condition for the platforms' ML advice accuracy to achieve certain level of bidder fairness.

**Corollary 4.2** *For any autobidding instance  $\Theta$  and  $\delta \in (0, 1)$ , if the platform can produce ML advice with accuracy  $\beta \geq 1 - (1 - \delta) \cdot (\underline{\alpha} - 1) \cdot \min_{i \in [K]} \frac{\text{OPT}_i}{\text{OPT}_{-i}}$ , the autobidding instance is  $\delta$ -fair for all bidders. Here,  $\underline{\alpha}$  is a multiplier lower bound for all bidders such that  $\underline{\alpha} > 1$ , and  $(1 - \delta) \cdot (\underline{\alpha} - 1) \cdot \max_{i \in [K]} \frac{\text{OPT}_i}{\text{OPT}_{-i}} < 1$ .*

Finally, the following theorem states the fairness bound in Theorem 4.1 is tight. For a technical re-statement of the Theorem and its proof see Appendix C.2.

**Theorem 4.3 (Matching fairness lower bound)** *For any  $\beta \in (0, 1)$ ,  $\alpha > 1$ , and  $R \in \left[0, \frac{\alpha_i - 1}{1 - \beta}\right]$ , there exists an autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$  with  $\beta$ -approximate reserves  $\mathbf{r}$ , in which there is a bidder  $i$  whose inverse-market share is  $\frac{\text{OPT}_{-i}}{\text{OPT}_i} = R$ , and has a fairness guarantee  $\min_{\mathbf{x} \in \mathcal{F}} \frac{W_i(\mathbf{x})}{\text{OPT}_i} = 1 - \frac{1 - \beta}{\alpha_i - 1} \cdot \frac{\text{OPT}_{-i}}{\text{OPT}_i}$  when she adopts multiplier  $\alpha_i = \alpha$ .*

## 5. Impossibility result: VCG is the fairest

Having presented a fairness guarantee in the previous Section 4 that improves according to the platform's ML advice accuracy, a natural question is that for a given level of accuracy  $\beta$ , can one achieve a universally better fairness guarantee than that of Theorem 4.1 via considering auction formats other than VCG? In this section, we demonstrate that the answer is negative when we restrict the auction to a broad class of truthful mechanisms (possibly randomized) with anonymous allocations. Here, we again emphasize that truthfulness is w.r.t. quasi-linear utility maximizers (see Remark 2.1). To be self-contained, we include the definition of allocation-anonymous auctions:

**Definition 5.1 (Allocation-anonymity)** *A possibly randomized auction is allocation-anonymous if the outcome for a permutation of bid profile  $\mathbf{b}$  is the permutation of the outcome for  $\mathbf{b}$ .*

**Remark 5.1** *An alternative view for auction anonymity is that the auction's allocation is independent of bidders' identities. For instance, for an auction format  $\mathcal{A}$  with a single slot, allocation-anonymity says for any set of bid values  $\mathcal{B} = \{b_1 \dots b_K\}$ , there exists probabilities  $\mathbf{q}(\mathcal{B}) = (q_1(\mathcal{B}) \dots q_K(\mathcal{B})) \in [0, 1]^{K+1}$  where  $q_k(\mathcal{B}) = \mathbb{P}(\text{bid value } b_k \text{ wins auction } \mathcal{A} \text{ given competing bids } \mathbf{b}_{-k})$  only depends on relative bid values in  $\mathcal{B}$ . Note that  $\sum_{k \in [K]} q_k(\mathcal{B}) \leq 1$ .*

Similar to Section 4, if we consider an anonymous auction  $\mathcal{A}$  with  $\beta$ -approximate reserves, the conceptual procedure of interest is to impose such reserve price on top of  $\mathcal{A}$  with a lazy implementation Paes Leme et al. (2016): if a bidder's bid exceeds her reserve she will be allocated according to

$\mathcal{A}$ , otherwise she will not be allocated (leaving a slot empty if she would have won the slot without reserves). The following theorem shows under this procedure with  $\beta$ -approximate reserves, no allocation-anonymous, truthful auction  $\mathcal{A}$  can universally outperform VCG, i.e. for any  $\mathcal{A}$  there exists an autobidding instance in which a bidder has a welfare guarantee at most the fairness lower bound for VCG of Theorem 4.1.

**Theorem 5.1** *For any auction  $\mathcal{A}$  that is allocation-anonymous, truthful, and possibly randomized, there exists an autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$  with a single available slot in each auction and  $\beta$ -approximate reserves  $\mathbf{r}$ , such that there is a feasible outcome  $\mathbf{x}$  in which a bidder  $i$ 's welfare is upper bounded as  $\frac{\mathbb{E}_{\mathcal{A}}[W_i(\mathbf{x})]}{\mathbb{E}_{\mathcal{A}}[\text{OPT}_i]} \leq 1 - \frac{1-\beta}{\alpha_i-1} \cdot \frac{\mathbb{E}_{\mathcal{A}}[\text{OPT}_{-i}]}{\mathbb{E}_{\mathcal{A}}[\text{OPT}_i]}$  for  $N, M \rightarrow \infty$  and multiplier  $\alpha_i > 1$ . Here the expectation is taken w.r.t. possible randomness in  $\mathcal{A}$ .<sup>2</sup>*

Our proof strategy for Theorem 5.1 is to construct a “bad” autobidding instance for any auction  $\mathcal{A}$  of interest that is “unfair” as possible to one specific bidder: we show that in this autobidding instance, there is some bidder  $i$  who has a welfare upper bound as stated in the theorem. The construction of this bad autobidding instance is motivated by Example 3.1, in which the key source of “unfairness” for an individual bidder  $i$  comes from the fact that competing bidders outbid  $i$  in auctions where  $i$ 's value is high, and cover their expenditures with value acquired from other auctions where they have no competition. Following this idea, since the bad instance in Theorem 5.1 requires us to maximize “unfairness” for a specific bidder  $i$ , we can achieve this by having auctions where each of  $i$ 's competitors is the only bidder submitting a nonzero bid, and with these “no-competition” auctions competitors can cover their expenditures for outbidding bidder  $i$  in auctions where  $i$ 's value is largest.

**Proof sketch for Theorem 5.1.** For any auction  $\mathcal{A}$  that is allocation-anonymous, truthful and possibly randomized, we consider a “bad” autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$  where  $N = K + 1$  bidders labeled  $B_1 \dots B_K, B_0$  compete in  $M = 2K + 1$  auctions with single-slots for some  $K \in \mathbb{N}$ , and bidders' values are shown in the following table. Reserves are set to be  $r_{i,j} = \beta v_{i,j}$  for some  $\beta \in (0, 1)$  and are  $\beta$ -approximate (see Definition 3.1). Bidder  $B_0$ 's multiplier is fixed to be  $\alpha_0 > 1$ .

<sup>2</sup>In general imposing personalized reserve prices invalidate allocation anonymity. However in the specific autobidding instance we constructed, imposing  $\beta$ -approx (personalized) reserve prices preserves anonymity in allocation-anonymous auction  $\mathcal{A}$ ; see proof sketch at end of section.

	$A_1$	$A_2$	$\dots$	$A_K$	$A_{K+1}$	$A_{K+2}$	$\dots$	$A_{2K}$	$A_{2K+1}$
$B_1$	$\frac{\alpha_0 v + \epsilon}{\rho}$	$\frac{\alpha_0 v + 2\epsilon}{\rho}$	$\dots$	$\frac{\alpha_0 v + K\epsilon}{\rho}$	$\gamma$	0	$\dots$	0	0
$B_2$	$\frac{\alpha_0 v + 2\epsilon}{\rho}$	$\frac{\alpha_0 v + 3\epsilon}{\rho}$	$\dots$	$\frac{\alpha_0 v + \epsilon}{\rho}$	0	$\gamma$	$\dots$	0	0
$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$
$B_K$	$\frac{\alpha_0 v + K\epsilon}{\rho}$	$\frac{\alpha_0 v + \epsilon}{\rho}$	$\dots$	$\frac{\alpha_0 v + (K-1)\epsilon}{\rho}$	0	0	$\dots$	$\gamma$	0
$B_0$	$v$	$v$	$\dots$	$v$	0	0	$\dots$	0	$y$

In the table, we choose  $\epsilon = \mathcal{O}(1/K^3)$  and suitable parameters  $\rho, \gamma, v, y > 0$  to satisfy certain conditions, one of which guarantees  $B_0$ 's value is the highest in auctions  $A_1 \dots A_K$ . With the above instance, we consider the specific outcome  $\mathbf{x}$  where bidders  $1, \dots, K$  adopt bid multiplier  $\rho$ , in which case bidder  $B_0$  has the lowest bid in auctions  $A_1 \dots A_K$ . Then, our proof of Theorem 5.1 is to show bidder  $B_0$  can acquire welfare at most the upper bound in Theorem 5.1. The proof consists of 3 parts:

(1) Under outcome  $\mathbf{x}$ , we upper bound bidder  $B_0$ 's expected acquired welfare in auctions  $A_1 \dots A_K$ . This acquired welfare should be small, since other bidders are outbidding bidder  $B_0$  in these auctions, by covering their expenditures via the value acquired in auctions  $A_{K+1}, \dots, A_{2K}$ , respectively.

(2) We show that bidder  $B_0$  satisfies her ROAS constraint, which holds valid due to the fact that she is acquiring value in auction  $A_{2K+1}$  for suitable  $y$  facing no competition.

(3) We show that any bidder  $i \in [K]$  satisfies her ROAS constraint when  $\epsilon \rightarrow 0$ . In this part, the key step is to analyze the values acquired by any bidder  $i \in [K]$  in auctions  $A_1, \dots, A_K$ , which is approximately  $\frac{\alpha_0 v}{\rho} \cdot \sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j)$ . We recognize that when bidders  $1, \dots, K$  use bid multiplier  $\rho$ , the bid profiles for auctions  $A_1 \dots A_K$  are a cyclic permutation of the set  $\{b_0, \dots, b_K\} = \{\alpha_0 v, \alpha_0 v + \epsilon, \dots, \alpha_0 v + K\epsilon\}$ . Therefore by allocation-anonymity of  $\mathcal{A}$ , the expected outcome in auctions  $A_1, \dots, A_K$  are symmetric over bidders  $1, \dots, K$ , which implies the sum of probabilities  $\sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j) = \sum_{k \in [K]} \mathbb{P}(\text{bid value } b_k \text{ wins auction } \mathcal{A} \text{ given competing bids } b_{-k}) \leq 1$  (see Remark 5.1). Here, we also point out that any  $\beta$  approximate reserves do not affect allocation in auctions  $A_1 \dots A_K$ , simply because any bid value in  $\{b_0, \dots, b_K\}$  is greater than the largest reserve price among agents, namely  $\beta v$ , since  $\alpha_0 > 1 > \beta$ . In other words, under the specific outcome  $\mathbf{x}$ , allocation anonymity of any auction in  $A_1, \dots, A_K$  is preserved with personalized reserves  $r_{i,j} = \beta v_{i,j}$  due to our construction.

For a technical re-statement of Theorem 5.1 and its proof, please refer to Appendices D.1 and D.2.

## 6. Extensions: fairness guarantees for GSP and GFP with ML advice

In this section, we extend our fairness and welfare guarantees for the VCG auction in Theorem 4.1 to the GSP and GFP auctions, which are both non-truthful. For technical purposes, we assume that bidder values are “well-separated” in the following sense:

**Definition 6.1 ( $\Delta$ -separated values)** *We say values  $\mathbf{v} \in \mathbb{R}_{\geq 0}^{N \times M}$  are  $\Delta$ -separated for some  $\Delta \in (1, 2)$  if any value  $v_{i,j}$  is at least  $\frac{1}{2-\Delta}$  times as much as any value that is less than  $v_{i,j}$  in the same auction  $j$ , i.e.  $v_{i,j} \geq \frac{1}{2-\Delta} \cdot \max\{v_{k,j} : k \in [N], v_{k,j} < v_{i,j}\}$  for any bidder  $i$  and auction  $j$ .*

Although the value separations in Definition 6.1 are multiplicative, they also capture the scenario in which values are “additively separated”. In particular, assume there exists some small  $d > 0$  such that  $d < \min\{v_{i,j} : v_{i,j} \neq 0\}$ , and we have additive separation  $v_{i,j} - d \geq \max\{v_{k,j} : k \in [N], v_{k,j} < v_{i,j}\}$  for any bidder  $i$  and auction  $j$ . Then, by taking  $\Delta \in \left(1, \min\left\{2, 1 + \frac{d}{\max\{v_{i,j} : v_{i,j} \neq 0\}}\right\}\right)$ , the values are  $\Delta$ -separated according to Definition 6.1 because  $(2 - \Delta)v_{i,j} \geq \left(1 - \frac{d}{\max\{v_{i,j} : v_{i,j} \neq 0\}}\right)v_{i,j} \geq v_{i,j} - d \geq \max\{v_{k,j} : k \in [N], v_{k,j} < v_{i,j}\}$  for all  $v_{i,j}$ . This suggests Definition 6.1 is not restrictive.

Further, as discussed in Section 2.2, uniform bidding (i.e. setting the same bid multiplier for all auctions) is only optimal in truthful auctions. In GSP and GFP, one can construct instances where non-uniform bidding strictly outperforms uniform bidder (for more details see e.g. Deng et al. (2019)). Thus, for GSP and GFP autobidding instances, we impose no assumptions on the bid values of bidders other than being undominated: we say a bid value  $b_i$  is undominated for bidder  $i$  if there is no other bid value  $b'$  that strictly outperforms  $b_i$  in the sense of Equation (1) for all competing bid profiles  $\mathbf{b}_{-i}$ .

The following lemma lower bounds undominated bids in the presence of  $\beta$ -approximate reserves.

**Lemma 6.1 (Lemma 4.7 & 4.9 of Balseiro et al. (2021a))** *Consider any autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$  where the auction format  $\mathcal{A}$  is the GSP or GFP auctions, and reserve prices  $\mathbf{r}$  are  $\beta$ -approximate. Denote  $\mathcal{U} \subseteq \mathbb{R}_{\geq 0}^{N \times M}$  to be the set of bid profiles in which each bid is undominated and satisfies all bidders’ ROAS constraints. Then for any  $\mathbf{b} \in \mathcal{U}$ ,  $b_{i,j}$  must satisfy  $b_{i,j} \geq r_{i,j} \geq \beta v_{i,j}$  for all  $i, j$ .*

Finally, our main theorem for this section is the following:

**Theorem 6.2** *Consider any autobidding instance  $\Theta = (N, M, \mathcal{A}, \mathbf{r}, \mathbf{v})$  where the auction format  $\mathcal{A}$  is the GSP or GFP auction. Suppose reserve prices  $\mathbf{r}$  are  $\beta$ -approximate, and values  $\mathbf{v}$  are  $\Delta$ -separated s.t.  $0 < \frac{1}{\Delta} < \beta < 1$ . Consider any undominated bid profile  $\mathbf{b} \in \mathcal{U} \subseteq \mathbb{R}_{\geq 0}^{N \times M}$  where  $\mathcal{U}$  is the set of all undominated bids under which every bidder’s ROAS constraint is satisfied (see Equation (1)). Then,  $\text{LOSS}_i(\mathbf{x}(\mathbf{b})) \leq \frac{1-\beta}{\beta-1/\Delta} \text{OPT}_{-i}$ , where  $\text{LOSS}_i$  is defined in Equation (2). Further,  $\min_{\mathbf{b} \in \mathcal{U}} \frac{W_i(\mathbf{x}(\mathbf{b}))}{\text{OPT}_i} \geq 1 - \frac{1-\beta}{\beta-1/\Delta} \cdot \frac{\text{OPT}_{-i}}{\text{OPT}_i}$ .*

The proof for Theorem 6.2 is presented in Appendix E.1. Comparing the fairness guarantees in Theorem 4.1 for VCG and Theorem 6.2 for GSP/GFP, we observe when values are  $\Delta$ -separated and ML advice is  $\beta$ -accurate, when bidders adopt small enough uniform multipliers in VCG (i.e.  $\alpha_i - 1 < \beta - 1/\Delta$ ), GSP/GFP provides a better fairness guarantee compared to VCG, whereas for large multipliers (i.e.  $\alpha_i - 1 > \beta - 1/\Delta$ ), fairness in VCG dominates that in the considered non-truthful auctions.

## References

- Aggarwal G, Badanidiyuru A, Mehta A (2019) Autobidding with constraints. *International Conference on Web and Internet Economics*, 17–30 (Springer).
- Balseiro S, Deng Y, Mao J, Mirrokni V, Zuo S (2021a) Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems* 34.
- Balseiro S, Kim A, Mahdian M, Mirrokni V (2021b) Budget management strategies in repeated auctions. *Operations Research* Forthcoming.
- Balseiro S, Kim A, Mahdian M, Mirrokni V (2021c) Budget-management strategies in repeated auctions. *Operations Research* 69(3):859–876.
- Balseiro S, Lu H, Mirrokni V (2020) The best of many worlds: Dual mirror descent for online allocation problems. *arXiv preprint arXiv:2011.10124* .
- Balseiro SR, Deng Y, Mao J, Mirrokni VS, Zuo S (2021d) The landscape of auto-bidding auctions: Value versus utility maximization. *Proceedings of the 22nd ACM Conference on Economics and Computation*, 132–133.
- Bateni M, Chen Y, Ciocan DF, Mirrokni V (2022) Fair resource allocation in a volatile marketplace. *Operations Research* 70(1):288–308.
- Beyhaghi H, Golrezaei N, Leme RP, Pál M, Sivan B (2021) Improved revenue bounds for posted-price and second-price mechanisms. *Operations Research* 69(6):1805–1822.
- Braga PL, Oliveira AL, Meira SR (2007) Software effort estimation using machine learning techniques with robust confidence intervals. *7th international conference on hybrid intelligent systems (HIS 2007)*, 352–357 (IEEE).
- Calders T, Kamiran F, Pechenizkiy M (2009) Building classifiers with independency constraints. *2009 IEEE International Conference on Data Mining Workshops*, 13–18 (IEEE).
- Chen Q, Golrezaei N, Susan F, Baskoro E (2022) Fair assortment planning. *Available at SSRN 4072912* .
- Dai B, Nachum O, Chow Y, Li L, Szepesvári C, Schuurmans D (2020) Coindice: Off-policy confidence interval estimation. *Advances in neural information processing systems* 33:9398–9411.
- Deng Y, Mao J, Mirrokni V, Zhang H, Zuo S (2022) Efficiency of the first-price auction in the autobidding world. *arXiv preprint arXiv:2208.10650* .



- 
- Deng Y, Mao J, Mirrokni V, Zuo S (2021) Towards efficient auctions in an auto-bidding world. *Proceedings of the Web Conference 2021*, 3965–3973.
- Deng Y, Schneider J, Sivan B (2019) Prior-free dynamic auctions with low regret buyers. *Advances in Neural Information Processing Systems* 32.
- Derakhshan M, Golrezaei N, Paes Leme R (2022) Linear program-based approximation for personalized reserve prices. *Management Science* 68(3):1849–1864.
- Dwork C, Hardt M, Pitassi T, Reingold O, Zemel R (2012) Fairness through awareness. *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226.
- Edelman B, Ostrovsky M, Schwarz M (2007) Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97(1):242–259.
- Feng Z, Lahaie S (2021) Robust clearing price mechanisms for reserve price optimization. *arXiv preprint arXiv:2107.04638* .
- Feng Z, Lahaie S, Schneider J, Ye J (2021) Reserve price optimization for first price auctions in display advertising. *International Conference on Machine Learning*, 3230–3239 (PMLR).
- Golrezaei N, Jaillet P, Liang JCN (2019a) Incentive-aware contextual pricing with non-parametric market noise. *arXiv preprint arXiv:1911.03508* .
- Golrezaei N, Jaillet P, Liang JCN, Mirrokni V (2021a) Bidding and pricing in budget and roi constrained markets. *arXiv preprint arXiv:2107.07725* .
- Golrezaei N, Javanmard A, Mirrokni V (2019b) Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Advances in Neural Information Processing Systems* 32.
- Golrezaei N, Lobel I, Paes Leme R (2018) Auction design for roi-constrained buyers. *Available at SSRN 3124929* .
- Golrezaei N, Lobel I, Paes Leme R (2021b) Auction design for roi-constrained buyers. *Proceedings of the Web Conference 2021*, 3941–3952.
- Indyk P, Mallmann-Trenn F, Mitrovic S, Rubinfeld R (2022) Online page migration with ml advice. *International Conference on Artificial Intelligence and Statistics*, 1655–1670 (PMLR).
- Jiang B, Zhang X, Cai T (2008) Estimating the confidence interval for prediction errors of support vector machine classifiers. *The Journal of Machine Learning Research* 9:521–540.
- Kanoria Y, Nazerzadeh H (2020) Dynamic reserve prices for repeated auctions: Learning from bids. *arXiv preprint arXiv:2002.07331* .
- Kearns M, Neel S, Roth A, Wu ZS (2019) An empirical study of rich subgroup fairness for machine learning. *Proceedings of the conference on fairness, accountability, and transparency*, 100–109.
- Kusner MJ, Loftus J, Russell C, Silva R (2017) Counterfactual fairness. *Advances in neural information processing systems* 30.

- Lahaie S, Pennock DM, Saberi A, Vohra RV (2007) Sponsored search auctions. *Algorithmic game theory* 1:699–716.
- Lee Kc, Orten B, Dasdan A, Li W (2012) Estimating conversion rate in display advertising from past performance data. *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 768–776.
- Lykouris T, Vassilvitskii S (2018) Competitive caching with machine learned advice. *International Conference on Machine Learning*, 3296–3305 (PMLR).
- Ma W, Xu P, Xu Y (2021) Fairness maximization among offline agents in online-matching markets. *arXiv preprint arXiv:2109.08934* .
- Mehrabi N, Morstatter F, Saxena N, Lerman K, Galstyan A (2021) A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54(6):1–35.
- Mehta A (2022) Auction design in an auto-bidding setting: Randomization improves efficiency beyond vcg. *Proceedings of the ACM Web Conference 2022*, 173–181.
- Munoz A, Vassilvitskii S (2017) Revenue optimization with approximate bid predictions. *Advances in Neural Information Processing Systems* 30.
- Paes Leme R, Pal M, Vassilvitskii S (2016) A field guide to personalized reserve prices. *Proceedings of the 25th international conference on world wide web*, 1093–1102.
- Pai MM, Vohra R (2014) Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150:383–425.
- Richardson M, Dominowska E, Ragno R (2007) Predicting clicks: estimating the click-through rate for new ads. *Proceedings of the 16th international conference on World Wide Web*, 521–530.
- Roughgarden T (2015) Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)* 62(5):1–42.
- Shelagh Dolan (2019) US Programmatic Digital Display Ad Spending Will Grow Despite Pandemic-Related Recession. <https://www.emarketer.com/content/us-programmatic-digital-display-ad-spending-will-grow-despite-pandemic-related-recession>, accessed: 11 May 2022.
- Shrestha DL, Solomatine DP (2006) Machine learning approaches for estimation of prediction interval for the model output. *Neural networks* 19(2):225–235.
- Sodomka E, Lahaie S, Hillard D (2013) A predictive model for advertiser value-per-click in sponsored search. *Proceedings of the 22nd international conference on World Wide Web*, 1179–1190.
- Varian HR (2007) Position auctions. *international Journal of industrial Organization* 25(6):1163–1178.
- Wang S, Li J, Wang S (2020) Online algorithms for multi-shop ski rental with machine learned advice. *Advances in Neural Information Processing Systems* 33:8150–8160.
- Yuan S, Wang J, Chen B, Mason P, Seljan S (2014) An empirical study of reserve price optimisation in real-time bidding. *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1897–1906.

---

Zliobaite I (2015) On the relation between accuracy and fairness in binary classification. *arXiv preprint arXiv:1505.05723* .

---

# Appendices for Fairness in the Autobidding World with Machine-learned Advice

---

## Appendix A: Extended Literature Review

**Algorithmic bidding/learning under constraints** The behavior of our bidders of interest are governed by their ROAS, and there has been a growing area of works on bidding-algorithm design under similar financial constraints for online advertising markets. Balseiro et al. (2021b) develops theoretical performance guarantees of the budget pacing strategy for bidders with hard budget cap (see more on budget management strategies in Balseiro et al. (2021c)), while Balseiro et al. (2020) presents a more general mirror descent algorithm for online resource allocation problems. Golrezaei et al. (2021a) present near-optimal bidding algorithms for bidders with both budget and ROI constraints in expectation. In this work, we do not study the design of bidding algorithms but instead consider worst case outcomes under any feasible bidding profile.

**Reserve price optimization.** Reserve price techniques and optimization have been studied for different auction formats and settings. In the single-shot second price auction setting Paes Leme et al. (2016), Beyhaghi et al. (2021), Derakhshan et al. (2022) presents different approaches with theoretical performance guarantees to optimize personalized reserve prices, while Yuan et al. (2014) presents an empirical study on the impact of reserve price on the entire auction system for display advertising. For repeated second price auctions, Golrezaei et al. (2019b,a), Kanoria and Nazerzadeh (2020) dynamically learn reserve prices to maximize cumulative revenue facing strategic agents, where as Feng and Lahaie (2021) optimize reserve prices to balance revenue and bidders' incentives to misreport. For first price auctions, Feng et al. (2021) introduces a gradient-based adaptive algorithm to dynamically optimize reserve prices. Nevertheless, all aforementioned works attempt to design and learn optimal or near optimal reserve prices for the purpose of revenue maximization, whereas in our work we directly set reserves using ML advice provided by some external black-box, and shed light on how reserve prices can improve fairness among all bidders.

## Appendix B: Proofs for Section 2

### B.1. Proof for Proposition 2.1

For simplicity, denote  $\delta_{i,j} = \text{OPT}_{i,j} - W_{i,j}(\mathbf{x})$ . Then,  $\text{OPT}_i - W_i(\mathbf{x}) = \sum_{j \in [M]: \delta_{i,j} > 0} \delta_{i,j} + \sum_{j \in [M]: \delta_{i,j} = 0} \delta_{i,j} + \sum_{j \in [M]: \delta_{i,j} < 0} \delta_{i,j} = \text{LOSS}_i(\mathbf{x}) + \sum_{j \in [M]: W_{i,j}(\mathbf{x}) > \text{OPT}_{i,j}} (\text{OPT}_{i,j} - W_{i,j}(\mathbf{x})) \leq \text{LOSS}_i(\mathbf{x}) \leq B$ . Rearranging and dividing both sides by  $\text{OPT}_i$  we get  $\frac{W_i(\mathbf{x})}{\text{OPT}_i} \geq 1 - \frac{B}{\text{OPT}_i}$ .

Here we remark that it is possible to have  $W_{i,j}(\mathbf{x}) > \text{OPT}_{i,j}$  because bidders may overbid, and therefore win auctions/slots that they would not have won under the efficient outcome.

## Appendix C: Proofs for Section 4

### C.1. Proof for Theorem 4.1

Fix a bidder  $i \in [K]$ , her bid multiplier  $\alpha_i$ , and any outcome  $\mathbf{x} = (\mathbf{x}_1 \dots \mathbf{x}_M)$  where  $\mathbf{x}_j \in \{0, 1\}^{N \times S_j}$  is the outcome vector in auction  $j$ .

Denote  $\ell_{k,j}, \ell_{k,j}^{\text{OPT}}$  to be the position of bidder  $k \in [N]$  in auction  $j \in [M]$  under outcome  $\mathbf{x}$  and the efficient outcome, respectively. Consider an auction  $j \in \mathcal{L}_i = \{j \in [M] : \ell_{i,j} > \ell_{i,j}^{\text{OPT}}\}$  (see Remark 2.2), i.e. in auction  $j$ , bidder  $i$  acquires a position (under  $\mathbf{x}$ ) below her position in the efficient outcome  $\mathbf{x}^{\text{OPT}}$ . Then there exists a bidder  $\kappa_j$  such that

$$v_{\kappa_j,j} < v_{i,j}, \quad \text{and} \quad \ell_{\kappa_j,j} \leq \ell_{i,j}^{\text{OPT}} < \ell_{i,j} \quad (3)$$

Consider the payment of bidder  $\kappa_j$ , and recall  $\hat{b}_{\ell,j}$  is the  $\ell$ th largest bid in the  $j$ th auction:

$$\begin{aligned} p_{\kappa_j,j} &\geq \sum_{\ell=\ell_{\kappa_j,j}}^{S_j} (\mu(\ell) - \mu(\ell+1)) \hat{b}_{\ell+1,j} \\ &\stackrel{(a)}{=} \sum_{\ell=\ell_{\kappa_j,j}}^{\ell_{i,j}^{\text{OPT}}-1} (\mu(\ell) - \mu(\ell+1)) \hat{b}_{\ell+1,j} + \sum_{\ell=\ell_{i,j}^{\text{OPT}}}^{\ell_{i,j}-1} (\mu(\ell) - \mu(\ell+1)) \hat{b}_{\ell+1,j} + p_{i,j} \\ &\stackrel{(b)}{\geq} (\mu(\ell_{\kappa_j,j}) - \mu(\ell_{i,j}^{\text{OPT}})) v_{i,j} + \alpha_i (\mu(\ell_{i,j}^{\text{OPT}}) - \mu(\ell_{i,j})) v_{i,j} + \beta \cdot \mu(\ell_{i,j}) v_{i,j} \\ &= \mu(\ell_{\kappa_j,j}) v_{i,j} + (\alpha_i - 1) (\mu(\ell_{i,j}^{\text{OPT}}) - \mu(\ell_{i,j})) v_{i,j} - (1 - \beta) \cdot \mu(\ell_{i,j}) v_{i,j}. \end{aligned} \quad (4)$$

Here, (a) follows from the VCG payment rule as illustrated in Example 2.1; (b) follows from the fact that bidder  $i$ 's ranking is  $\ell_{i,j}$ , so any bidder who is ranked before position  $\ell_{i,j}$  submitted a bid greater than  $b_{i,j} = \alpha_i v_{i,j}$  which is the bid of bidder  $i$ , i.e.  $\hat{b}_{\ell,j} \geq b_{i,j} = \alpha_i v_{i,j} > v_{i,j}$  for any  $\ell \leq \ell_{i,j}$ .

On the other hand, we have

$$\begin{aligned} p_{\kappa_j,j} + \sum_{j' \neq j} p_{\kappa_j,j'} &\leq \mu(\ell_{\kappa_j,j}) v_{\kappa_j,j} + \sum_{j' \neq j} \mu(\ell_{\kappa_j,j'}) v_{\kappa_j,j'} \\ p_{\kappa_j,j'} &\geq \beta \cdot \mu(\ell_{\kappa_j,j'}) v_{\kappa_j,j'} \quad \forall j' \in [M], \end{aligned}$$

where the first inequality follows from bidder  $\kappa_j$ 's ROAS constraint; the second inequality follows from the fact that any winning bidder's payment must be greater than her  $\beta$ -approximate reserves. Combining the above inequalities and rearranging we get

$$p_{\kappa_j,j} \leq \mu(\ell_{\kappa_j,j}) v_{\kappa_j,j} + (1 - \beta) \cdot \sum_{j' \neq j} \mu(\ell_{\kappa_j,j'}) v_{\kappa_j,j'}, \quad (5)$$

Combining Equations (4) and (5), we get

$$\begin{aligned} &(\alpha_i - 1) \cdot (\mu(\ell_{i,j}^{\text{OPT}}) - \mu(\ell_{i,j})) v_{i,j} \\ &\leq (1 - \beta) \cdot \left( \mu(\ell_{i,j}) v_{i,j} + \sum_{j' \neq j} \mu(\ell_{\kappa_j,j'}) v_{\kappa_j,j'} \right) + \mu(\ell_{\kappa_j,j}) (v_{\kappa_j,j} - v_{i,j}) \\ &\stackrel{(a)}{\leq} (1 - \beta) \cdot \left( \mu(\ell_{i,j}) v_{i,j} + \sum_{j' \neq j} \mu(\ell_{\kappa_j,j'}) v_{\kappa_j,j'} + \mu(\ell_{\kappa_j,j}) (v_{\kappa_j,j} - v_{i,j}) \right) \\ &\stackrel{(b)}{\leq} (1 - \beta) \cdot \left( \mu(\ell_{i,j}) v_{i,j} + \sum_{j' \in [M]} \mu(\ell_{\kappa_j,j'}) v_{\kappa_j,j'} - \mu(\ell_{i,j}^{\text{OPT}}) v_{i,j} \right). \end{aligned} \quad (6)$$

In (a), we used the fact that  $v_{\kappa_j,j} - v_{i,j} < 0$  by the definition of  $\kappa_j$  in Equation (3); and (b) follows again from Equation (3) such that  $\mu(\ell_{i,j}^{\text{OPT}}) \leq \mu(\ell_{\kappa_j,j'})$ .

Summing the above over all  $j \in \mathcal{L}_i = \{j \in [M] : \ell_{i,j} > \ell_{i,j}^{\text{OPT}}\}$  (see Remark 2.2), we have

$$\begin{aligned}
\text{LOSS}_i(\mathbf{x}) &= \sum_{j \in \mathcal{L}_i} (\mu(\ell_{i,j}^{\text{OPT}}) - \mu(\ell_{i,j})) v_{i,j} \\
&\leq \frac{1-\beta}{\alpha_i-1} \left( \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}) v_{i,j} + \sum_{j \in \mathcal{L}_i} \sum_{j' \in [M]} \mu(\ell_{\kappa_j, j'}) v_{\kappa_j, j'} - \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}^{\text{OPT}}) v_{i,j} \right) \\
&\leq \frac{1-\beta}{\alpha_i-1} \left( W_{-i}(\mathbf{x}) + \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}) v_{i,j} - \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}^{\text{OPT}}) v_{i,j} \right) \\
&\leq \frac{1-\beta}{\alpha_i-1} \text{OPT}_{-i},
\end{aligned} \tag{7}$$

which yields our desired upper bound for  $\text{LOSS}_i(\mathbf{x})$ . Here, the final inequality follows from  $\text{OPT} \geq W(\mathbf{x})$  and further

$$\begin{aligned}
\text{OPT}_{-i} &\geq W_{-i}(\mathbf{x}) + W_i(\mathbf{x}) - \text{OPT}_i \\
&= W_{-i}(\mathbf{x}) + \sum_{j \in \mathcal{L}_i} (W_{i,j}(\mathbf{x}) - \text{OPT}_{i,j}) + \sum_{j \in [M]/\mathcal{L}_i} (W_{i,j}(\mathbf{x}) - \text{OPT}_{i,j}) \\
&\stackrel{(a)}{\geq} W_{-i}(\mathbf{x}) + \sum_{j \in \mathcal{L}_i} (W_{i,j}(\mathbf{x}) - \text{OPT}_{i,j}) \\
&= W_{-i}(\mathbf{x}) + \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}) v_{i,j} - \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}^{\text{OPT}}) v_{i,j}.
\end{aligned} \tag{8}$$

where in (a) we used the fact that  $W_{i,j}(\mathbf{x}) \geq \text{OPT}_{i,j}$  in any auction  $j \in [M]/\mathcal{L}_i$ .

Finally, applying Proposition 2.1 w.r.t. upper bound of  $\text{LOSS}_i(\mathbf{x})$  and since  $\mathbf{x} \in \mathcal{F}$  arbitrary, we obtain the desired welfare guarantee lower bound.

### C.2. Proof for Theorem 4.3

**Theorem C.1 (Restatement of Theorem 4.3)** *Consider 2 bidders competing in three SPA auctions whose values are indicated in the following table for any  $\beta \in (0, 1)$  and  $y \geq 0$ .*

	Auction 1	Auction 2	Auction 3
bidder 1	$y$	$v$	$0$
bidder 2	$0$	$v - \epsilon$	$\gamma + \frac{1}{1-\beta} \cdot \epsilon$

Bidder 1's multiplier is fixed to be  $\alpha_1 > 1$ , and consider  $v = \frac{1-\beta}{\alpha_1-1} \cdot \gamma$  for any  $\gamma > 0$ . The reserve prices are set to be  $r_{i,j} = \beta v_{i,j}$ . Then, we have

$$\min_{\mathbf{x} \in \mathcal{F}} \frac{W_1(\mathbf{x})}{\text{OPT}_1} = 1 - \frac{1-\beta}{\alpha_1-1} \cdot \frac{\text{OPT}_{-1} - \frac{1}{1-\beta} \cdot \epsilon}{\text{OPT}_1} \tag{9}$$

Taking  $\epsilon \rightarrow 0$  shows that bidder 1's welfare is equal to the fairness guarantee in Theorem 4.1.

**Remark C.1** We remark that as  $\epsilon \rightarrow 0$ ,  $\frac{\text{OPT}_{-i}}{\text{OPT}_i} = \frac{\alpha_1-1}{y+v} v \in \left[0, \frac{\alpha_1-1}{1-\beta}\right]$ , so by varying  $y \in [0, \infty)$ , the above example demonstrates our fairness lower bound in Theorem 4.1 is tight for any nontrivial market share ratio  $\frac{\text{OPT}_{-i}}{\text{OPT}_i} \in \left[0, \frac{\alpha_1-1}{1-\beta}\right]$ .

Note that in any feasible outcome, bidder 1 must win auction 1, and bidder 2 must win auction 3. Hence for auction 2, we only need to consider the following outcome:

**Bidder 1 loses auction 2, and suffers welfare loss  $v$ .** This outcome can be achieved by setting  $\alpha_2$  such that  $\alpha_2(v - \epsilon) > \alpha_1 v$ . Bidder 2 accumulates value  $v + \gamma + \left(\frac{1}{1-\beta} - 1\right)\epsilon$ . Her payment for auction 2 is  $\max\{\alpha_1 v, \beta(v - \epsilon)\}$ , and for auction 3 is  $\beta\left(\gamma + \frac{1}{1-\beta} \cdot \epsilon\right)$ . The following shows that her ROAS constraint is satisfied:

$$\begin{aligned} & v + \gamma + \left(\frac{1}{1-\beta} - 1\right)\epsilon - \max\{\alpha_1 v, \beta(v - \epsilon)\} - \beta\left(\gamma + \frac{1}{1-\beta} \cdot \epsilon\right) \\ \stackrel{(a)}{=} & v + \gamma + \left(\frac{1}{1-\beta} - 1\right)\epsilon - \alpha_1 v - \beta\left(\gamma + \frac{1}{1-\beta} \cdot \epsilon\right) \\ = & (1 - \alpha_1)v + (1 - \beta)\gamma + \left(\frac{1}{1-\beta} - 1 - \beta \cdot \frac{1}{1-\beta}\right)\epsilon \\ = & 0, \end{aligned}$$

where in (a) we used the fact  $\beta < \alpha_1$ . In the final equality we used the definition that  $v = \frac{1-\beta}{\alpha_1-1} \cdot \gamma$ . On the other hand, bidder 1's ROAS constraint is apparently satisfied.

Under this outcome, we have

$$\frac{W_1}{OPT_1} = 1 - \frac{v}{OPT_1} = 1 - \frac{1-\beta}{\alpha_1-1} \cdot \frac{\gamma}{OPT_1} = 1 - \frac{1-\beta}{\alpha_1-1} \cdot \frac{OPT_{-i} - \frac{1}{\beta} \cdot \epsilon}{OPT_i}$$

## Appendix D: Proofs for Section 5

### D.1. Additional Definitions and Lemmas for Section 5

The following lemma shows that for anonymous and truthful auctions, the probability of the lowest bidder winning a single auction is capped by a bound that decreases as the number of bidders grow.

**Lemma D.1 (Lemma 3 in Mehta (2022))** *In an anonymous and truthful auction for a single item with  $N$  bidders, the bidder who submits the lowest bid wins the item with probability at most  $\frac{1}{N}$ .*

The following technical definition and lemma (i.e. Definition D.1 and Lemma D.2) concerns the scenario where only one bidder participates in the auction (others bid 0), and present an upper bound on the probability and cost respectively for the single bidder to win the auction.

**Definition D.1 (Single bidder purchase probability and bid threshold)** *For any allocation-anonymous and truthful auction  $\mathcal{A}$ , consider the setting with a single bidder who submits bid  $b > 0$  and define*

$$\pi_{\mathcal{A}} = \lim_{b \rightarrow \infty} \mathbb{P}(\text{bidder wins item with bid } b), \quad (10)$$

where the limit exists because in a truthful auction,  $\mathbb{P}(\text{bidder wins item with bid } b)$  increases in  $b$  (see Remark 2.1 on truthful auctions). Assume this max probability is reached at some bid threshold  $Q_{\mathcal{A}}$ , i.e.

$$Q_{\mathcal{A}} = \min\{b > 0 : \mathbb{P}(\text{bidder wins item with bid } b) = \pi_{\mathcal{A}}\}. \quad (11)$$

Note that in a deterministic single-slot auction that allocates to the highest bidder,  $\pi_{\mathcal{A}} = 1$ , and  $Q_{\mathcal{A}} \rightarrow 0$ . For example, in an SPA with no reserve, the single bidder can win the auction with any arbitrarily small positive bid with probability 1.

**Lemma D.2 (Lemma 4 in Mehta (2022))** *For any allocation-anonymous and truthful auction  $\mathcal{A}$  with single-bidder purchase probability  $\pi_{\mathcal{A}}$  and bid threshold  $Q_{\mathcal{A}}$ , the expected cost for a single bidder for winning the item is at most  $\pi_{\mathcal{A}} \cdot Q_{\mathcal{A}}$ .*

## D.2. Proof of Theorem 5.1

**Theorem D.3 (Restatement of Theorem 5.1)** *For any auction  $\mathcal{A}$  that is allocation-anonymous, truthful, and possibly randomized,<sup>3</sup> consider an autobidding problem instance w.r.t.  $\mathcal{A}$  with  $M = 2K + 1$  auctions and  $N = K + 1$  bidders. Fix bidder 0's bid multiplier to be  $\alpha_0$  and some  $\beta \in [0, 1)$ . Consider the bidder values  $\{v_{i,j}\}_{i \in [N], j \in [M]}$  given in the following table.*

	$A_1$	$A_2$	$\dots$	$A_K$	$A_{K+1}$	$A_{K+2}$	$\dots$	$A_{2K}$	$A_{2K+1}$
$B_1$	$\frac{\alpha_0 v + \epsilon}{\rho}$	$\frac{\alpha_0 v + 2\epsilon}{\rho}$	$\dots$	$\frac{\alpha_0 v + K\epsilon}{\rho}$	$\gamma$	$0$	$\dots$	$0$	$0$
$B_2$	$\frac{\alpha_0 v + 2\epsilon}{\rho}$	$\frac{\alpha_0 v + 3\epsilon}{\rho}$	$\dots$	$\frac{\alpha_0 v + \epsilon}{\rho}$	$0$	$\gamma$	$\dots$	$0$	$0$
$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$
$B_K$	$\frac{\alpha_0 v + K\epsilon}{\rho}$	$\frac{\alpha_0 v + \epsilon}{\rho}$	$\dots$	$\frac{\alpha_0 v + (K-1)\epsilon}{\rho}$	$0$	$0$	$\dots$	$\gamma$	$0$
$B_0$	$v$	$v$	$\dots$	$v$	$0$	$0$	$\dots$	$0$	$y$

In the table, we let  $\gamma > \frac{Q_{\mathcal{A}}}{\beta} > Q_{\mathcal{A}}$ ,  $\epsilon = O(1/K^3)$  and  $v = \frac{1-\beta}{\alpha_0-1} \cdot \pi_{\mathcal{A}} \cdot \gamma$ . Let  $\rho, y$  and a large enough  $K$  satisfy the following:

$$\alpha_0 < \rho < \frac{\alpha_0}{\beta} \text{ s.t. } \frac{\alpha_0 v + K\epsilon}{\rho} < v, \quad \text{and } y > \max \left\{ \frac{Q_{\mathcal{A}}}{\alpha_0}, \frac{\alpha_0 v}{\pi_{\mathcal{A}}} \right\}, \quad (12)$$

where  $Q_{\mathcal{A}}, \pi_{\mathcal{A}}$  are defined in Definition D.1. Further, suppose the platform enforces personalized reference prices  $\mathbf{r} \in \mathbb{R}_{\geq 0}^{N \times M}$  on top of auction  $\mathcal{A}$ , where  $r_{i,j} = \beta v_{i,j}$ . Then, letting the (possibly random) outcome be  $\mathbf{x}$  when bidders 1, ...  $K$  all adopt the bid multiplier  $\rho$ , the ROAS constraints for all bidders are satisfied when  $K \rightarrow \infty$  and  $\rho \rightarrow \alpha_0$ , and for bidder 0 we have

$$\lim_{K \rightarrow \infty} \frac{\mathbb{E}_{\mathcal{A}} [W_0(\mathbf{x})]}{\mathbb{E}_{\mathcal{A}} [\text{OPT}_0]} \leq 1 - \frac{1-\beta}{\alpha_0-1} \cdot \lim_{K \rightarrow \infty} \frac{\mathbb{E}_{\mathcal{A}} [\text{OPT}_{-0}]}{\mathbb{E}_{\mathcal{A}} [\text{OPT}_0]} \quad (13)$$

where  $\mathbb{E}_{\mathcal{A}}$  is taken w.r.t. the randomness in outcome  $\mathbf{x}$  due to randomness in the auction  $\mathcal{A}$ , and the loss of bidder 0, namely  $\text{LOSS}_0$ , is defined in Equations (2).

<sup>3</sup> Here, we assume all auctions of interest are individually rational (IR), i.e. the payment of a bidder is always less than her submitted bid.



First note that bidder 0 only has competition in auctions  $A_1 \dots A_K$ , and hence can only incur a loss (that contributes to  $\text{LOSS}_0(\mathbf{x})$  defined in Equations (2)) within these auctions. Hence  $\mathbb{E}_{\mathcal{A}}[\text{LOSS}_0(\mathbf{x})] = v \sum_{j \in [K]} \mathbb{P}(\text{bidder 0 loses auction } j)$ . Then we consider the following:

$$\begin{aligned} \mathbb{E}_{\mathcal{A}}[\text{LOSS}_0(\mathbf{x})] &= v \sum_{j \in [K]} \mathbb{P}(\text{bidder 0 loses auction } j) = v \sum_{j \in [K]} (1 - \mathbb{P}(\text{bidder 0 wins auction } j)) \\ &\stackrel{(a)}{\geq} v \cdot \frac{K^2}{K+1} = \frac{1-\beta}{\alpha_0-1} \cdot \gamma \cdot \pi_{\mathcal{A}} \cdot \frac{K^2}{K+1} \stackrel{(b)}{=} \frac{1-\beta}{\alpha_0-1} \cdot \mathbb{E}_{\mathcal{A}}[\text{OPT}_{-0}] \cdot \frac{K}{K+1}. \end{aligned} \quad (14)$$

Here (a) holds because bidder 0 bids  $\alpha_0 v$  for any auction in  $1, 2, \dots, K$ , which is strictly less than all other bidders' bids as they all adopt multipliers  $\rho$  in these auctions, so from Lemma D.1, we have  $\mathbb{P}(\text{bidder 0 wins auction } j) \leq \frac{1}{K+1}$ ; in (b) we used the fact that  $\mathbb{E}_{\mathcal{A}}[\text{OPT}_{-0}] = \sum_{j=K+1}^{2K} \mathbb{E}_{\mathcal{A}}[\gamma] = \gamma \cdot K \cdot \pi_{\mathcal{A}}$  since there is only a single non-zero bidder in auctions  $A_{K+1} \dots A_{2K}$  and each bidder submits a bid  $\rho\gamma > \rho > Q_{\mathcal{A}}$  (see Definition D.1).

Therefore we have

$$\lim_{K \rightarrow \infty} \frac{\mathbb{E}_{\mathcal{A}}[W_0(\mathbf{x})]}{\mathbb{E}_{\mathcal{A}}[\text{OPT}_0]} \stackrel{(a)}{=} 1 - \lim_{K \rightarrow \infty} \frac{\mathbb{E}_{\mathcal{A}}[\text{LOSS}_0(\mathbf{x})]}{\mathbb{E}_{\mathcal{A}}[\text{OPT}_0]} \leq 1 - \frac{1-\beta}{\alpha_0-1} \lim_{K \rightarrow \infty} \frac{\mathbb{E}_{\mathcal{A}}[\text{OPT}_{-0}]}{\mathbb{E}_{\mathcal{A}}[\text{OPT}_0]}, \quad (15)$$

where (a) follows from the fact that in our constructed autobidding instance, bidder 0's acquired value in each auction cannot exceed that under the efficient allocation, and hence can only incur loss in welfare.

Now it only remains to show that the multipliers  $(\alpha_0, \rho, \dots, \rho) \in (1, \infty)^{K+1}$  yields a feasible outcome, i.e. the ROI constraints of each bidder is satisfied in expectation. Let  $V_{i,j}$  and  $C_{i,j}$  be the expected value and cost of bidder  $i$  in auction  $A_j$ , respectively.

**1. Showing bidder 0's ROI constraint is satisfied.** We show by the following: bidder 0 only incurs a non-zero expected cost in auctions  $A_1 \dots A_K$  and  $A_{2K+1}$ , and we will show that the expected value  $V_{0,2K+1}$  is lower bounded by the expected costs  $C_{0,2K+1} + \sum_{j \in [K]} C_{0,j}$ .

Since  $\alpha_0 y > Q_{\mathcal{A}}$ , the definition of the single-bidder purchasing probability in Definition D.1 implies that bidder 0 acquires an expected value from auction  $A_{2K+1}$  of  $V_{0,2K+1} = \pi_{\mathcal{A}} y$ . Further, since bidder 0 submits the lowest bids in auctions  $A_1 \dots A_K$  under bid multiplier profile  $(\alpha_0, \rho, \dots, \rho) \in (0, \infty)^{K+1}$ , from Lemma D.1, we have  $\mathbb{P}(\text{bidder 0 wins auction } j) \leq \frac{1}{K+1}$  for all  $j \in [K]$ . Since the payment of a bidder in an auction is at most her submitted bid (as the auction is IR), we know that  $\sum_{j \in [K]} C_{0,j} \leq K \cdot \frac{\alpha_0 v}{K+1} < \pi_{\mathcal{A}} y = V_{0,2K+1}$ , where the inequality follows from the definition of  $y$  in Equation (12) such that  $y > \max\left\{\frac{Q_{\mathcal{A}}}{\alpha_0}, \frac{\alpha_0 v}{\pi_{\mathcal{A}}}\right\}$ . This implies bidder 0's ROI constraint is satisfied.

**2. Showing bidder  $i$ 's ROI constraint is satisfied for any  $i = 1, 2, \dots, K$ .** We show this by considering the following: bidder  $i$  only incurs a non-zero expected cost in auctions  $A_1 \dots A_K$  and  $A_{K+i}$ , and we will show that the expected values  $V_{i,K+i} + \sum_{k \in [K]} V_{i,j}$  is lower bounded by the expected costs  $C_{i,K+i} + \sum_{j \in [K]} C_{i,j}$ .

- **Calculate cost  $C_{i,K+i}$ :** For auction  $A_{K+i}$ , bidder  $i$ 's bid is  $\rho\gamma > \gamma > Q_{\mathcal{A}}$  from the definition of  $\gamma$ , so by Definition D.1, the probability of  $i$  winning the item in auction  $A_{K+i}$  is  $\pi_{\mathcal{A}}$ , and the expected cost is

$$C_{i,K+i} \leq \pi_{\mathcal{A}} \cdot \max\{r_{i,K+i}, Q_{\mathcal{A}}\} \leq \pi_{\mathcal{A}} \cdot \beta\gamma, \quad (16)$$

where the final inequality follows from the definition  $r_{i,K+i} = \beta\gamma$

- **Upper bound costs**  $\sum_{j \in [K]} C_{i,j}$ : For auctions  $[K] = 1 \dots K$ , bidder  $i$ 's total expected cost can be bounded as

$$\begin{aligned} \sum_{j \in [K]} C_{i,j} &\leq \rho \sum_{j \in [K]} v_{i,j} \mathbb{P}(\text{bidder } i \text{ wins auction } A_j) \\ &= \alpha_0 v \sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } A_j) + \frac{(K+1)K}{2} \epsilon. \end{aligned} \quad (17)$$

where the first inequality follows from a bidder's payment is at most her submitted bid since the auction is IR.

- **Calculate**  $V_{i,K+i}$ : Considering auction  $A_{K+i}$ , bidder  $i$  is the only bidder, and since  $\rho\gamma > \gamma > Q_{\mathcal{A}}$ , the definition of the single-bidder purchasing probability in Definition D.1 implies that bidder  $i$ 's acquires an expected value from this auction of

$$V_{i,K+i} = \pi_{\mathcal{A}} \cdot \gamma. \quad (18)$$

- **Lower bound**  $\sum_{k \in [K]} V_{i,j}$ :

$$\sum_{k \in [K]} V_{i,j} \geq \frac{\alpha_0 v}{\rho} \sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j). \quad (19)$$

Combining Equations (16),(17),(18) and (19), we get

$$\begin{aligned} &\sum_{j \in [K]} V_{i,j} + V_{i,K+i} - \left( \sum_{j \in [K]} C_{i,j} + C_{i,K+i} \right) \\ &\geq \pi_{\mathcal{A}} \cdot \gamma + \frac{\alpha_0 v}{\rho} \cdot \sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j) \\ &\quad - \left( \pi_{\mathcal{A}} \cdot \beta\gamma + \alpha_0 v \cdot \sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j) + \frac{(K+1)K}{2} \epsilon \right) \\ &= \pi_{\mathcal{A}} \cdot (1-\beta)\gamma - \left( \alpha_0 - \frac{\alpha_0}{\rho} \right) v \cdot \sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j) - \frac{(K+1)K}{2} \epsilon \\ &\stackrel{(a)}{=} (\alpha_0 - 1)v - \left( \alpha_0 - \frac{\alpha_0}{\rho} \right) v \cdot \sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j) - \frac{(K+1)K}{2} \epsilon \\ &\stackrel{(b)}{\geq} (\alpha_0 - 1)v - \left( \alpha_0 - \frac{\alpha_0}{\rho} \right) v - \frac{(K+1)K}{2} \epsilon \end{aligned} \quad (20)$$

where (a) follows from the definition  $v = \frac{1-\beta}{\alpha_0-1} \cdot \pi_{\mathcal{A}} \cdot \gamma$ ; In (b) we used the fact that  $\rho > \alpha_0 > 1$  and  $\sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } A_j) \leq 1$  due to the following: Consider the set of bid values  $\mathcal{B} = \{\alpha_0 v, \alpha_0 v + \epsilon, \alpha_0 v + 2\epsilon \dots \alpha_0 v + K\epsilon\} \subseteq \mathbb{R}_{>0}$ , and we recognize that any bid value  $b_k \in \mathcal{B}$  exceeds the maximum reserve price  $\beta v$  in auctions  $A_1 \dots A_K$ . Therefore the constructed reserve prices do not affect allocation, and hence by anonymity of auction  $\mathcal{A}$  there exists probabilities  $\mathbf{q}(\mathcal{B}) = (q_0(\mathcal{B}), q_1(\mathcal{B}) \dots q_K(\mathcal{B})) \in [0, 1]^{K+1}$  where

$$q_k(\mathcal{B}) = \mathbb{P}(\text{bid value } b_k \text{ wins auction } \mathcal{A} \text{ given competing bids } \mathbf{b}_{-k}) \quad \text{and} \quad \sum_{k=0}^K q_k(\mathcal{B}) \leq 1.$$

We recognize that in each auction  $A_1 \dots A_K$ , under bid multipliers  $(\alpha_0, \rho, \dots, \rho) \in (1, \infty)^{K+1}$  the submitted bid profile is a cyclic permutation of  $\mathcal{B}$ . Therefore we know that

$$\sum_{j \in [K]} \mathbb{P}(\text{bidder } i \text{ wins auction } j) = \sum_{k=1}^K q_k(\mathcal{B}) \leq 1 - q_0(\mathcal{B}) \leq 1$$

Finally, by taking  $\rho \rightarrow \alpha_0$  and  $K \rightarrow \infty$  in Equation (20), and utilizing  $\epsilon = O(1/K^3)$  we have

$$\lim_{\rho \rightarrow \alpha_0} \lim_{K \rightarrow \infty} \sum_{j \in [K]} V_{i,j} + V_{i,K+i} - \left( \sum_{j \in [K]} C_{i,j} + C_{i,K+i} \right) \geq 0.$$

This shows that bidder  $i$ 's ROI constraint is satisfied.

## Appendix E: Proofs for Section 6

### E.1. Proof of Theorem 6.2

Fix a bidder  $i \in [K]$  and any outcome  $\mathbf{x} = (\mathbf{x}_1 \dots \mathbf{x}_M)$  where  $\mathbf{x}_j \in \{0, 1\}^{N \times S_j}$  is the outcome vector in auction  $j$ .

Denote  $\ell_{k,j}, \ell_{k,j}^{\text{OPT}}$  to be the position of bidder  $k \in [N]$  in auction  $j \in [M]$  under outcome  $\mathbf{x}$  and the efficient outcome, respectively. Consider an auction  $j \in \mathcal{L}_i = \{j \in [M] : \ell_{i,j} > \ell_{i,j}^{\text{OPT}}\}$  (see Remark 2.2), i.e. in auction  $j$ , bidder  $i$  acquires a position (under  $\mathbf{x}$ ) below her position in the efficient outcome  $\mathbf{x}^{\text{OPT}}$ . Then there exists a bidder  $\kappa_j$  such that

$$v_{\kappa_j, j} < v_{i, j}, \quad \text{and} \quad \ell_{\kappa_j, j} \leq \ell_{i, j}^{\text{OPT}} < \ell_{i, j} \quad (21)$$

Consider the payment of bidder  $\kappa_j$ , and recall  $\hat{b}_{\ell, j}$  is the  $\ell$ th largest bid in the  $j$ th auction. The following a similar deduction as Equation (4) in the proof of Theorem 4.1, we have

$$\begin{aligned} p_{\kappa_j, j} &\stackrel{(a)}{\geq} \sum_{\ell=\ell_{\kappa_j, j}}^{S_j} (\mu(\ell) - \mu(\ell+1)) \hat{b}_{\ell+1, j} \\ &= \sum_{\ell=\ell_{\kappa_j, j}}^{\ell_{i, j}-1} (\mu(\ell) - \mu(\ell+1)) \hat{b}_{\ell+1, j} + p_{i, j} \\ &\stackrel{(b)}{\geq} \beta (\mu(\ell_{\kappa_j, j}) - \mu(\ell_{i, j})) v_{i, j} + \beta \cdot \mu(\ell_{i, j}) v_{i, j} \\ &= \beta \mu(\ell_{\kappa_j, j}) \cdot v_{i, j} \\ &= \mu(\ell_{\kappa_j, j}) v_{i, j} + \left( \beta - \frac{1}{\Delta} \right) (\mu(\ell_{i, j}^{\text{OPT}}) - \mu(\ell_{i, j})) v_{i, j} - (1 - \beta) \cdot \mu(\ell_{i, j}) v_{i, j} \\ &\quad - (1 - \beta) \mu(\ell_{\kappa_j, j}) v_{i, j} + \left( \frac{1}{\Delta} - \beta \right) \mu(\ell_{i, j}^{\text{OPT}}) v_{i, j} + \left( 1 - \frac{1}{\Delta} \right) \mu(\ell_{i, j}) v_{i, j} \\ &\stackrel{(c)}{\geq} \mu(\ell_{\kappa_j, j}) v_{i, j} + \left( \beta - \frac{1}{\Delta} \right) (\mu(\ell_{i, j}^{\text{OPT}}) - \mu(\ell_{i, j})) v_{i, j} - (1 - \beta) \cdot \mu(\ell_{i, j}) v_{i, j} \\ &\quad - \left( 1 - \frac{1}{\Delta} \right) \mu(\ell_{\kappa_j, j}) v_{i, j} + \left( 1 - \frac{1}{\Delta} \right) \mu(\ell_{i, j}) v_{i, j} \\ &= \left( \beta - \frac{1}{\Delta} \right) (\mu(\ell_{i, j}^{\text{OPT}}) - \mu(\ell_{i, j})) v_{i, j} - (1 - \beta) \cdot \mu(\ell_{i, j}) v_{i, j} \\ &\quad + \frac{1}{\Delta} \mu(\ell_{\kappa_j, j}) v_{i, j} + \left( 1 - \frac{1}{\Delta} \right) \mu(\ell_{i, j}) v_{i, j} \end{aligned} \quad (22)$$

Here, (a) follows from the fact that for a fix bid profile, the payment of GSP or GFP for each bidder in an auction dominates that of VCG (see discussion in Example 2.1); (b) follows from  $\hat{b}_{\ell,j} \geq b_{i,j}$  for  $\ell \leq \ell_{i,j}$ , and since  $\mathbf{b} \in \mathcal{B} \subseteq \mathbb{R}_{\geq 0}^{N \times M}$  is an undominated bid profile, Lemma 6.1 applies to  $b_{i,j} \geq \beta v_{i,j}$ . Also  $p_{i,j} \geq r_{i,j} \geq \beta v_{i,j}$  be the definition of  $\beta$ -approximate reserves; (c) follows from the fact that  $\beta > \frac{1}{\Delta}$  and  $\mu(\ell_{i,j}^{\text{OPT}}) \leq \mu(\ell_{\kappa_j,j})$  since  $\ell_{\kappa_j,j} \leq \ell_{i,j}^{\text{OPT}}$  according to Equation (21).

On the other hand, we have

$$\begin{aligned} p_{\kappa_j,j} + \sum_{j' \neq j} p_{\kappa_j,j'} &\leq \mu(\ell_{\kappa_j,j})v_{\kappa_j,j} + \sum_{j' \neq j} \mu(\ell_{\kappa_j,j'})v_{\kappa_j,j'} \\ p_{\kappa_j,j'} &\geq \beta \cdot \mu(\ell_{\kappa_j,j'})v_{\kappa_j,j'} \quad \forall j' \in [M], \end{aligned}$$

where the first inequality follows from bidder  $\kappa_j$ 's ROAS constraint; the second inequality follows from the fact that any winning bidder's payment must be greater than her  $\beta$ -approximate reserves. Combining the above inequalities and rearranging we get

$$p_{\kappa_j,j} \leq \mu(\ell_{\kappa_j,j})v_{\kappa_j,j} + (1-\beta) \cdot \sum_{j' \neq j} \mu(\ell_{\kappa_j,j'})v_{\kappa_j,j'}. \quad (23)$$

Combining Equations (22) and (23), we get

$$\begin{aligned} &\left(\beta - \frac{1}{\Delta}\right) \cdot (\mu(\ell_{i,j}^{\text{OPT}}) - \mu(\ell_{i,j}))v_{i,j} \\ &\leq (1-\beta) \cdot \left( \mu(\ell_{i,j})v_{i,j} + \sum_{j' \neq j} \mu(\ell_{\kappa_j,j'})v_{\kappa_j,j'} \right) + \underbrace{\mu(\ell_{\kappa_j,j})v_{\kappa_j,j} - \frac{1}{\Delta}\mu(\ell_{\kappa_j,j})v_{i,j} - \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{i,j})v_{i,j}}_Y. \end{aligned} \quad (24)$$

We now upper bound  $Y$ :

$$\begin{aligned} &\mu(\ell_{\kappa_j,j})v_{\kappa_j,j} - \frac{1}{\Delta}\mu(\ell_{\kappa_j,j})v_{i,j} - \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{i,j})v_{i,j} \\ &= \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{\kappa_j,j})v_{\kappa_j,j} - \frac{1}{\Delta}\mu(\ell_{\kappa_j,j})(v_{i,j} - v_{\kappa_j,j}) - \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{i,j})v_{i,j} \\ &= \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{\kappa_j,j})(v_{\kappa_j,j} - v_{i,j}) - \frac{1}{\Delta}\mu(\ell_{\kappa_j,j})(v_{i,j} - v_{\kappa_j,j}) \\ &\quad + \left(1 - \frac{1}{\Delta}\right)(\mu(\ell_{\kappa_j,j}) - \mu(\ell_{i,j}))v_{i,j} \\ &= \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{\kappa_j,j})(v_{\kappa_j,j} - v_{i,j}) + \frac{\mu(\ell_{\kappa_j,j})v_{i,j}}{\Delta} \left( (\Delta - 1) \left(1 - \frac{\mu(\ell_{i,j})}{\mu(\ell_{\kappa_j,j})}\right) - \left(1 - \frac{v_{\kappa_j,j}}{v_{i,j}}\right) \right) \\ &\stackrel{(a)}{\leq} \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{\kappa_j,j})(v_{\kappa_j,j} - v_{i,j}) + \frac{\mu(\ell_{\kappa_j,j})v_{i,j}}{\Delta} \left( (\Delta - 1) - \left(1 - \frac{v_{\kappa_j,j}}{v_{i,j}}\right) \right) \\ &= \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{\kappa_j,j})(v_{\kappa_j,j} - v_{i,j}) + \frac{\mu(\ell_{\kappa_j,j})v_{i,j}}{\Delta} \left( (\Delta - 2) + \frac{v_{\kappa_j,j}}{v_{i,j}} \right) \\ &\stackrel{(b)}{\leq} \left(1 - \frac{1}{\Delta}\right)\mu(\ell_{\kappa_j,j})(v_{\kappa_j,j} - v_{i,j}) \\ &\stackrel{(c)}{\leq} (1-\beta)\mu(\ell_{\kappa_j,j})(v_{\kappa_j,j} - v_{i,j}). \end{aligned} \quad (25)$$

where in (a) we recall  $\Delta \in (1, 2)$  and  $\ell_{\kappa_j,j} < \ell_{i,j}$  from Equation (21) so that  $\mu(\ell_{\kappa_j,j}) > \mu(\ell_{i,j})$ ; (b) follows from the fact that values are  $\Delta$ -separated, so  $v_{i,j} > v_{\kappa_j,j}$  from Equation (21) implies  $\frac{v_{\kappa_j,j}}{v_{i,j}} \leq 2 - \Delta$ ; in (c) we used the fact that  $\beta > \frac{1}{\Delta}$  so  $1 - \beta < 1 - \frac{1}{\Delta}$ , and the fact that  $v_{\kappa_j,j} - v_{i,j} < 0$  according to Equation (21).

Combining Equations (24) and (25) we get

$$\begin{aligned}
& \left( \beta - \frac{1}{\Delta} \right) \cdot (\mu(\ell_{i,j}^{\text{OPT}}) - \mu(\ell_{i,j})) v_{i,j} \\
& \leq (1 - \beta) \cdot \left( \mu(\ell_{i,j}) v_{i,j} + \sum_{j' \neq j} \mu(\ell_{\kappa_j, j'}) v_{\kappa_j, j'} + \mu(\ell_{\kappa_j, j}) (v_{\kappa_j, j} - v_{i,j}) \right) \\
& = (1 - \beta) \cdot \left( \mu(\ell_{i,j}) v_{i,j} + \sum_{j \in [M]} \mu(\ell_{\kappa_j, j'}) v_{\kappa_j, j'} - \mu(\ell_{\kappa_j, j}) v_{i,j} \right) \\
& \stackrel{(a)}{\leq} (1 - \beta) \cdot \left( \mu(\ell_{i,j}) v_{i,j} + \sum_{j \in [M]} \mu(\ell_{\kappa_j, j'}) v_{\kappa_j, j'} - \mu(\ell_{i,j}^{\text{OPT}}) v_{i,j} \right).
\end{aligned} \tag{26}$$

where (a) follows from  $\mu(\ell_{i,j}^{\text{OPT}}) \leq \mu(\ell_{\kappa_j, j})$  due to the fact that  $\ell_{\kappa_j, j} \geq \ell_{i,j}^{\text{OPT}}$  according to Equation (21).

Summing the above over all  $j \in \mathcal{L}_i = \{j \in [M] : \ell_{i,j} > \ell_{i,j}^{\text{OPT}}\}$  (see Remark 2.2), we have

$$\begin{aligned}
\text{LOSS}_i(\mathbf{x}) &= \sum_{j \in \mathcal{L}_i} (\mu(\ell_{i,j}^{\text{OPT}}) - \mu(\ell_{i,j})) v_{i,j} \\
&\leq \frac{1 - \beta}{\beta - \frac{1}{\Delta}} \left( \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}) v_{i,j} + \sum_{j \in \mathcal{L}_i} \sum_{j' \in [M]} \mu(\ell_{\kappa_j, j'}) v_{\kappa_j, j'} - \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}^{\text{OPT}}) v_{i,j} \right) \\
&\leq \frac{1 - \beta}{\beta - \frac{1}{\Delta}} \left( W_{-i}(\mathbf{x}) + \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}) v_{i,j} - \sum_{j \in \mathcal{L}_i} \mu(\ell_{i,j}^{\text{OPT}}) v_{i,j} \right) \\
&\leq \frac{1 - \beta}{\beta - \frac{1}{\Delta}} \text{OPT}_{-i},
\end{aligned} \tag{27}$$

which yields our desired upper bound for  $\text{LOSS}_i(\mathbf{x})$ . Here, the final inequality follows from the same argument as Equation (8) in the proof of Theorem 4.1 (see Appendix C.1).

Finally, applying Proposition 2.1 w.r.t. upper bound of  $\text{LOSS}_i(\mathbf{x})$  and since  $\mathbf{x} \in \mathcal{F}$  arbitrary, we obtain the desired welfare guarantee lower bound.