45

46

47

48

49

Multi-Platform Autobidding with and without Predictions

Anonymous Author(s)

Abstract

We study the problem of finding the optimal bidding strategy for an advertiser in a multi-platform auction setting. The competition on a platform is captured by a value and a cost function, mapping bidding strategies to value and cost respectively. We assume a diminishing returns property, whereby the marginal cost is increasing in value. The advertiser uses an autobidder that selects a bidding strategy for each platform, aiming to maximize total value subject to budget and return-on-spend constraint. The advertiser has no prior information and learns about the value and cost functions by querying a platform with a specific bidding strategy. Our goal is to design an algorithm that finds the optimal bidding strategy with a small number of queries.

We first present an algorithm that requires $O(m \log(mn) \log n)$ queries, where *m* is the number of platforms and *n* is the number of possible bidding strategies in each platform. Moreover, we adopt the learning-augmented framework and propose an algorithm that utilizes a (possibly erroneous) prediction of the optimal bidding strategy. We provide a $O(m \log(m\eta) \log \eta)$ query-complexity bound on our algorithm as a function of the prediction error η . This guarantee gracefully degrades to $O(m \log(mn) \log n)$. This achieves a "best-of-both-worlds" scenario: O(m) queries when given a correct prediction, and $O(m \log(mn) \log n)$ even for an arbitrary incorrect prediction.

CCS Concepts

• Theory of computation \rightarrow Design and analysis of algorithms.

Keywords

Autobidding, Multi-Platform Auctions, Algorithms with Predictions

ACM Reference Format:

Anonymous Author(s). 2024. Multi-Platform Autobidding with and without Predictions. In . ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/ nnnnnnnnnnn

1 Introduction

Online advertisers often advertise across multiple platforms, such as Amazon, Bing, Google, Meta and TikTok, and face the challenging task of optimizing their bids across these platforms. The complexity comes not just from having to select a vector of bidding strategies (one for each platform), but also from the diversity of auctions used across platforms and the black-box nature of the detailed auction rules and the level of competition. To deal with the complexity, advertisers are increasingly using automated bidding agents (aka autobidding) to bid on their behalf. An autobidder allows an advertiser to specify constraints like Budget and Return-on-Spend (ROS), and bids on their behalf to maximize value subject to the constraints. This has led to a lot of research interest in problems related to autobidding (see Aggarwal et al. [3] for a recent survey). In particular, the problem of designing bidding algorithms for the single platform setting is well-studied [1], including in the online learning setting (see Section 3 of the survey [3]). However, there is not much work on the problem of bidding optimally across multiple platforms (see the related work section for what is known). 59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

In this paper, we study the problem of finding the optimal bidding strategy in the multi-platform setting. In particular, an advertiser aims to maximize her total value subject to a global budget and return-on-spend (ROS) constraint across all platforms. To capture the black-box nature of auction mechanisms and level of competition, we assume that the advertiser has no prior knowledge about the mapping from bids to auction outcomes for any platform. Instead, the advertiser interacts with a platform's auction by submitting a bid to the platform and observing the corresponding cost and value (i.e., the user has "query access" to the mapping). We propose algorithms that find the optimal bidding strategy in this setting, and prove worse-case query complexity bounds for them.

While worst-case results offer robustness and broad applicability, the guarantees they provide can sometimes be overly pessimistic. To address this, a new framework called "algorithms with predictions" has recently been introduced. This framework allows algorithms to incorporate potentially flawed machine-learned predictions as a guiding tool. The objective is twofold: to achieve improved performance guarantees when the prediction is accurate (a property known as consistency) and to maintain good worst-case bounds even when the prediction is completely incorrect (a property called robustness). This framework provides a natural way to integrate machine-learned predictions into the design of algorithms while preserving the essential robustness offered by worst-case analysis. In this work, we adopt the learning-augmented framework and explore the role of predictions in bidding strategy optimization. Specifically, we examine the scenario where the algorithm has access to a prediction $\hat{\mu}$ of the optimal bidding strategy, without any assumption regarding the prediction's accuracy. We propose algorithms that leverage the untrusted prediction to achieve improved query complexity bounds, which degrade gracefully based on the quality of the prediction.

1.1 Our Results

We model the problem of searching for the optimal bidding strategy as follows: there are *m* platforms. For each platform *j*, we are given a cost function c_j and a value function v_j , and *n* different bidding strategies indexed 1 through *n* such that $c_j(\mu_j)$ and $v_j(\mu_j)$ are respectively the cost incurred by bidding μ_j on platform *j*, and the value accrued from this bidding strategy. Our goal is to find

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

⁵⁵ Conference'17, July 2017, Washington, DC, USA

^{56 © 2024} Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-x-xxxx-xxYY/MM

⁵⁷ https://doi.org/10.1145/nnnnnnnnn
58

117 a collection $\boldsymbol{\mu} = (\mu_1, \dots, \mu_m)$ such that (a) the ratio of the total 118 value accrued to the total cost is at least some target threshold, 119 (b) the total cost is less than the budget, and (c) the total value is 120 maximized.

We propose a search algorithm, MEDIANOFMEDIANS, that deter-121 mines the exact optimal bidding strategy using $O(m \log(mn) \log n)$ queries. Our algorithm builds on a characterization of the optimal 123 124 bidding strategy in the multi-platform setting, which was first devel-125 oped in [4] under continuous strategy space. Intuitively, the optimal 126 strategy is to keep increasing bid on the platform that currently offers the highest marginal bang-per-buck (corresponding to the 127 lowest marginal cost-per-unit-value) until either the budget or the 128 ROS constraint is about to be violated. In other words, the optimal 129 strategy aims to equalize the marginal cost-per-unit-value across 130 all the platforms. 131

At a high level, the algorithm searches in the space of marginal 132 costs. Initially, there are up to mn candidate marginal costs. The 133 algorithm carefully selects a candidate marginal cost and finds the 134 135 corresponding vector of bidding strategies, as defined in Lemma 1, via Subroutine 2. Based on the outcomes (i.e. cost and value on 136 137 each platform) of the corresponding strategy vector, the algorithm 138 removes a constant fraction of candidate marginals from considera-139 tion, and recurses on the residual problem.

We complement this algorithmic result with an $\Omega(m \log n)$ lower 140 bound and an $\Omega(\log mn)$ lower bound for this problem. The $\Omega(\log mn)$ 141 142 lower bound reflects the difficulty of identifying the optimal marginal cost among the *mn* possible candidates, while the $\Omega(m \log n)$ 143 bound captures the complexity of determining the corresponding 144 bidding strategy for that optimal marginal cost. Notably, these are 145 the two key components of our algorithm, and our upper bound is 146 the product of these two lower bounds. 147

148 Next, we adopt the learning-augmented framework to improve 149 the worst-case query complexity bound. We propose an algorithm with access to a prediction $\hat{\mu}$ of the optimal strategy μ^{o} . The al-150 151 gorithm, BRANCHOUTMEDIANOFMEDIANS, starts with trying to 152 the find the optimal solution in a small range around the predictions, and expands the search range if the search is unsuccessful. 153 154 With the right sequence of expanding ranges, we show that the 155 algorithm finds the optimal strategy μ^* with $O(m \log(m\eta) \log \eta)$ queries, where $\eta = \max_i |\boldsymbol{\mu}_i^* - \hat{\boldsymbol{\mu}}_i|$ represents the prediction error. 156 This means that the algorithm requires only O(m) queries when 157 158 the predictions are accurate; this is the minimum number of queries 159 needed to implement any bidding strategy. Moreover, since $\eta \leq n$, 160 the total number of queries never exceeds that of the MEDIANOF-161 MEDIANS algorithm.

1.2 Related Work

162

163

164

165

166

167

168

169

170

171

172

173

174

Multi-platform mechanism design and autobidding. Previous research has examined the multi-platform auction environment from both the auctioneer's and the bidders' perspectives. Regarding auction design, Aggarwal et al. [2] analyzes a scenario where a single platform manages multiple channels, each selling queries via a second-price auction (SPA) with a reserve price. The authors assess the costs associated with each channel optimizing its own reserve price compared to a unified platform policy. Inspired by the Display Ad market, Paes Leme et al. [32] explores a model in which 175

176

177

178

179

180

181

182

183

184

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

multiple platforms vie for profit-maximizing bidders who must use the same bid across all platforms (which we refer to as a uniform bid). Their key finding indicates that the first-price auction (FPA) serves as the optimal auction format for these platforms. On the bidders' side, Susan et al. [34] investigates bidding strategies for utility-maximizing advertisers operating across multiple platforms while adhering to budget constraints. Meanwhile, Deng et al. [20] focuses on value-maximizing advertisers and reveals that optimizing ROS per platform can yield arbitrarily poor results when both ROS and budget constraints are in play.

Aggarwal et al. [4] study a similar multi-platform setting with autobidders under ROI constraints but focus on the auction design problem from the platform's perspective. While first-price auctions are optimal in the absence of competition (Deng et al. [19]), they show that from the perspective of each separate platform, running a second-price auction can achieve larger revenue than first-price auction when there are two competing platforms. They also identify key factors influencing the platform's choice of auction formats, including advertiser sensitivity to auction changes, competition and relative inefficiency of second-price auctions. In our paper, we focus on how advertisers can bid optimally in the multi-platform setting.

Algorithms with predictions. In recent years, the learning-augmented framework has emerged as a prominent paradigm for the design and analysis of algorithms. For an overview of early contributions, we refer to [31], while [27] offers an up-to-date compilation of relevant papers in this area. This framework seeks to address the shortcomings of overly pessimistic worst-case analyses. In the last five years alone, hundreds of papers have explored traditional algorithmic challenges through this lens, with notable examples including online paging [29], scheduling [33], optimization problems related to covering [12] and knapsack constraints [24], as well as topics like Nash social welfare maximization [13], the secretary problem [6, 21, 22], and various graph-related problems [7].

More closely related to our work, the research on learning-augmented mechanisms interacting with strategic agents is recently initiated by Agrawal et al. [5] and Xu and Lu [35]. This area includes strategic facility location [5, 14, 25, 35], strategic scheduling [10, 35], auctions [11, 16, 28, 30, 35], bicriteria mechanism design (which seeks to optimize both social welfare and revenue) [8], graph problems with private input [18], metric distortion [15], and equilibrium analysis [23, 26]. Recently, Christodoulou et al. [17] revisited mechanism design challenges by focusing on predictions about the outcome space rather than the input. While most of these studies concentrate on the mechanism design problem, our research emphasizes how predictions can assist agents in identifying optimal strategies. For more information on this body of work, we recommend [9].

2 Preliminaries

We consider the problem of finding the optimal bidding strategy in a multi-platform auction setting for a value-maximizer with *budget* and *return on spend* (ROS) constraints. There is a set M consisting of m platforms in the market. We assume that for each platform, the advertiser can pick from n different bids (note that this set can be different for different platforms), indexed by 0, 1, 2, ..., n, where

Multi-Platform Autobidding with and without Predictions

bid 0 is used to denote non-participation. Each platform $j \in M$ is described by a value and a cost function that map bid indices to a corresponding value and cost, respectively, i.e., $v_i : \{0, 1, \dots, n\} \rightarrow$ $\mathbb{R}^{\geq 0}$ and $c_i : \{0, 1, \dots, n\} \to \mathbb{R}^{\geq 0}$. In other words, when a bidder chooses to bid according to bid $\mu \in \{0, 1, ..., n\}$, they incur a cost of $c_i(\mu)$ and receive a value of $v_i(\mu)$. We assume that there is a strict ordering of costs and values by bid index, i.e. $v_i(\mu) < v_i(\mu+1)$ and $c_i(\mu) < c_i(\mu + 1)$ for all $\mu \in \{0, 1, \dots, n\}$. We refer to the mapping from bid to the cost and value of each platform as the landscape of that platform. In addition, we define the marginal cost of bidding $\mu \geq 1$ on platform *j* as

$$\mathsf{MC}_{j}(\mu) = \frac{c_{j}(\mu) - c_{j}(\mu - 1)}{v_{j}(\mu) - v_{j}(\mu - 1)},$$
 (Marginal Cost)

where c(0) = 0 and v(0) = 0. We make standard convexity assumption that MC_j is non-decreasing for every platform *j*.

Given the integral strategy set, we expand the bidding space by also considering the fractional solution between each integral bid, hence making the strategy space continuous. We use *S* and *S^c* to denote the integral and fractional strategy space, respectively. The cost, value and marginal functions of the continuous bidding space $[0, n]^m$ extend the discrete function by linear interpolation.¹ Formally

$$\begin{aligned} v_j(\mu) &= (\lceil \mu \rceil - \mu) \cdot v_j(\lfloor \mu \rfloor) + (\mu - \lfloor \mu \rfloor) \cdot v_j(\lceil \mu \rceil), \\ c_j(\mu) &= (\lceil \mu \rceil - \mu) \cdot c_j(\lfloor \mu \rfloor) + (\mu - \lfloor \mu \rfloor) \cdot c_j(\lceil \mu \rceil), \end{aligned}$$

Consequently, we have that $MC_i(\mu) = MC_i(\lceil \mu \rceil)$.

We note that the problem of finding the optimal integral solution is NP-hard.² The objective of the bidder is therefore to find an optimal *fractional* bidding strategy $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_m)$ where $\mu_j \in [0, n]$ such that she maximizes the total value received by executing bidding strategy μ_j on each platform *j*, subject to a budget constraint and the ROS constraint *across all platforms*. Let *B* and *T* be the budget and target ROS of the bidder. We can formulate the problem as the following program:

$$\max_{\boldsymbol{\mu}=(\mu_{1},\mu_{2}...,\mu_{m})} \sum_{j\in M} v_{j}(\mu_{j})$$

s.t.
$$\sum_{j\in M} c_{j}(\mu_{j}) \leq T \cdot \sum_{j\in M} v_{j}(\mu_{j}),$$
$$\sum_{j\in M} c_{j}(\mu_{j}) \leq B.$$
(1)

throughout the paper, we denote μ^o the optimal (fractional) bidding strategy, and μ^* the floor of it, i.e., $\mu_j^* = \lfloor \mu_j^o \rfloor$ for all platform *j*. Note that $\mu_j^* \in S$.

We assume that the bidder only knows the set of possible bidding strategies, but has no information about the platforms' value and cost functions. Instead, the bidder can interact with platforms via *bidding queries*: the bidder plays strategy μ on a platform *j* to learn the value $v_j(\mu)$, the cost $c_j(\mu)$, and the marginal cost MC $_j(\mu)^3$. Each Conference'17, July 2017, Washington, DC, USA

such query is costly to the bidder, and the goal is to minimize the number of queries required to determine the optimal strategy.

Given an instance I and an algorithm ALG, let ALG(I) denote the number of queries needed to find the optimal strategy for that instance. Then the query complexity of the algorithm is defined as:

$$\max_{\tau} \mathsf{ALG}(I)$$

The learning-augmented framework. In this work, we adopt the learning-augmented framework and study how we can further reduce the query complexity by considering search algorithms that are equipped with a (potentially erroneous) prediction $\hat{\boldsymbol{\mu}} \in [0, n]^m$ of the optimal fractional bidding strategy $\boldsymbol{\mu}^o(\mathcal{I}) = (\mu_1^o, \mu_2^o, \dots, \mu_n^o)$. The *error* of an predictions η is defined to be the maximum pointwise deviation from $\boldsymbol{\mu}^o$, formally:

$$\eta(\hat{\boldsymbol{\mu}}, \boldsymbol{I}) = \max_{j} |\hat{\mu}_{j} - \mu_{j}^{o}(\boldsymbol{I})|$$

We let the algorithm ALG use both the instance I and the prediction $\hat{\mu}$ as input. We evaluate the performance of such an algorithm using its *robustness, consistency* and the query complexity as a function of the prediction error.

The robustness of an algorithm refers to the worst-case query complexity of the algorithm given an adversarially chosen, possibly erroneous, prediction. Mathematically,

$$\operatorname{robustness}(\mathsf{ALG}) = \max_{\hat{\boldsymbol{\mu}}, I} \mathsf{ALG}(\hat{\boldsymbol{\mu}}, I)$$

r

The consistency of an algorithm refers to the worst-case query complexity when the prediction that it is provided with is accurate, i.e., $\hat{\mu} = \mu^*(\mathcal{I})$. Mathematically,

consistency(ALG) =
$$\max_{\hat{\boldsymbol{\mu}}, \boldsymbol{I} : \hat{\boldsymbol{\mu}} = \boldsymbol{\mu}^*(\boldsymbol{I})} ALG(\hat{\boldsymbol{\mu}}, \boldsymbol{I}).$$

Lastly, the query complexity of an algorithm given a prediction with error η' is defined to be:

$$\max_{\hat{\boldsymbol{\mu}}, \boldsymbol{I}: \boldsymbol{\eta}(\hat{\boldsymbol{\mu}}, \boldsymbol{I}) \leq \boldsymbol{\eta}'} \mathsf{ALG}(\hat{\boldsymbol{\mu}}, \boldsymbol{I}).$$

3 Characterization of Bidder's Optimal Bidding Strategy

In this section, we present a characterization of the optimal bidding strategy μ^o that will be useful in designing the algorithm. To this end, we first prove a useful lemma about the "ranking" of integral strategies in *S*. We then argue how an "almost-optimal" integral solution can be used to determine the optimal fractional solution.

LEMMA 1. Given some positive number k, and the n discrete indices on each platform, define $\mu^k = (\mu_1^k, \mu_2^k, \dots, \mu_m^k)$ where

$$\mu_{j}^{k} = \arg \max_{\mu \in \{0, 1, \dots, n\}} \{ \mathsf{MC}_{j}(\mu) \le k \},\$$

then there exist a k^* such that for any $k \le k^*$, μ^k is a feasible solution for program(1) and for any $k' > k^*$, $\mu^{k'}$ is not feasible.

We first show the following helper lemma. Intuitively, if we consider the landscape of each platform, and connect each bidding strategy with a straight line, the landscape would be convex, the lemma below is simply a property of a convex function. Due to space limitations, we defer the proof to appendix A.

¹It can be viewed as bidding randomly between two adjacent bids.

²It is not hard to see that we can encode any knapsack problem as an instance of our problem with a budget constraint.

³When marginal cost is not part of the query output, it is still achievable by querying both the current and the previous bid, which increases the query complexity by a constant factor.

LEMMA 2. For any platform
$$j \in M$$
, $\frac{c_j(\mu)}{v_j(\mu)} \leq MC_j(\mu)$.

PROOF OF LEMMA 1. let k be the smallest k with infeasible $\boldsymbol{\mu}^k$, if the infeasibility is due to the budget constraint, then for any $k' \geq k$ we trivially have that $\boldsymbol{\mu}^{k'}$ violates the budget constraint as well since $\mu_j^{k'} \geq \mu_j^k$ and the cost functions are monotone.

If the infeasibility is due to the ROS constraint, i.e.,

$$\sum_{j \in M} c_j(\mu_j^k) > T \cdot \sum_{j \in M} v_j(\mu_j^k),$$
(2)

proving the statement is equivalent to proving for any $k' \ge k$, $\boldsymbol{\mu}^{k'}$ is also infeasible. To this end, we first show that the maximum marginals among the μ_j^k is strictly more than *T*, assume for contradiction, that $MC_j(\mu_j^k) \le T$ for all *j*, by lemma 2 we would have the $c_j(\mu_j^k)/v_j(\mu_j^k) \le MC_j(\mu_j^k) \le T$, which contradicts with (2). We therefore have

$$\max_{j \in M} \mathsf{MC}_j(\mu_j^k) > T \tag{3}$$

We now inductively prove that for any $k' \ge k$, we have $\mu^{k'}$ is infeasible. Consider increasing k' starting from k, at the beginning we could have $\mu^{k'} = \mu^k$ (which is infeasible), consider the first point $k' \ge k$ such that $\mu^{k'} \ne \mu^k$, we know that:

(1) there exist at least one platform j' such that
$$\mu_{i'}^{k'} = \mu_{i'}^{k} + 1$$

(2) $MC_{j'}(\mu_{j'}^{k'}) > \max_{j \in M} MC_j(\mu_j^k) > T$,

where the first inequality is by definition of $\mu^{k'}$ and the second inequality is due to (3). Now consider:

$$\begin{split} \sum_{i \in M} c_j(\mu_j^{k'}) &= \sum_{j \in M} c_j(\mu_j^k) + c_{j'}(\mu_{j'}^k + 1) - c_{j'}(\mu_{j'}^k) \\ &> T \cdot \sum_{j \in M} v_j(\mu_j^k) + c_{j'}(\mu_{j'}^k + 1) - c_{j'}(\mu_{j'}^k) \\ &> T \cdot \sum_{j \in M} v_j(\mu_j^k) + \mathsf{MC}_{j'}(\mu_{j'}^k + 1) \cdot (v_{j'}(\mu_{j'}^k + 1) - v_{j'}(\mu_{j'}^k)) \\ &> T \cdot \sum_{j \in M} v_j(\mu_j^k) + T \cdot (v_{j'}(\mu_{j'}^k + 1) - v_{j'}(\mu_{j'}^k)) \\ &> T \cdot \sum_{j \in M} v_j(\mu_j^{k'}) \end{split}$$

Inductively apply this argument for each update of $\mu^{k'}$ proves the statement.

3.1 The fractional optimal bidding strategy

We present the optimal fractional solution, which at a high level is achieved by the following greedy process: starting with an initial budget of *B*, we allocate an infinitesimal amount to the platform currently offering the best value-to-unit-cost ratio (equivalently, the smallest marginal cost). We continue this process until either the budget is exhausted or the Return on Spend (ROS) constraint becomes tight. The bids corresponding to the final cost/value ratio on each platform form the optimal fractional strategy.

For ease reference, we formally define the feasible μ^k with the largest *k* as the "almost-optimal" integral solution.

Anon.

DEFINITION 1 (ALMOST-OPTIMAL STRATEGY). We say an bidding strategy μ is almost-optimal if $\mu = \max_k \left[\mu^k \text{ is feasible} \right]$.

Essentially, the almost-optimal integral strategy provides a close lower bound for the optimal fractional solution μ^o . Specifically, let μ^* represent the largest feasible μ^k ; then, $\mu^* = \lfloor \mu^o \rfloor$. We present a subroutine that takes the almost-optimal integral solution μ^* as input and returns the exact fractional optimal bidding strategy μ^o by greedily selecting the smallest marginal costs (breaking ties lexicographically with respect to a fixed ordering of platforms) until the constraint is tight.

Consider the following equation:

$$\sum_{j} (x_{j} \cdot c_{j}(\mu_{j}^{*}) + (1 - x_{j}) \cdot c_{j}(\mu_{j}^{*} + 1))$$

= min $\left[B, T \cdot \left(\sum_{j} (x_{j} \cdot v_{j}(\mu_{j}^{*}) + (1 - x_{j}) \cdot v_{j}(\mu_{j} + 1)) \right) \right]$ (4)

By definition of μ^* , we have the $x_j \in [0, 1]$ for all platform *j*. In the case where there are multiple set of solutions, we break ties by maximizing the x_j with lower platform index first.

5	SUBROUTINE 1: ROUNDUP
	Input: almost-optimal integral solution μ^*
1	for $j \in M$ do
2	$\mu_j' \leftarrow \mu_j^* + 1$
3	$\mu_j^o \leftarrow \mu_j^*$
4	query each μ_j' to obtain $v_j(\mu_j'), c_j(\mu_j')$ and $MC_j(\mu_j')$
5	re-index the platforms in non-decreasing order of $MC_j(\mu'_j)$
	s.t. if $i \leq j$, $MC_i(\mu'_i) \leq MC_j(\mu'_j)$
6	Solve for x_j in (4)
7	$\mu_j^o \leftarrow \mu_j^o + x_j$
8	return μ^o

LEMMA 3 (OPTIMAL BIDDING STRATEGY). Let μ^* be the almostoptimal integral solution. Then ROUNDUP (μ^*) is bidder's fractional optimal bidding strategy.

PROOF. We prove optimality of $\mu' = \text{ROUNDUP}(\mu^*)$ by contradiction. Assume there exist another solution μ^a that is optimal with a value strictly higher than μ' . If μ^a obtains a better value than μ' , there must be at least one platform j such that $\mu_j^a > \mu'_j$. Consequently, there must be some other platform i, such that $\mu_j^a < \mu'_j$, since the constraints are tight of solution μ' by definition of the ROUNDUP algorithm and μ^a is feasible by assumption.

we first argue that $MC_j(\mu_j^a) > MC_i(\mu_i')$. To see this, first note that $MC_j(\mu_i^a) \ge MC_i(\lfloor \mu_i' \rfloor)$, since otherwise by definition of $\boldsymbol{\mu}^k$ we have:

$$\mu_j' = \operatorname*{arg\,max}_{\mu}[\mathsf{MC}_j(\mu) \leq k] \geq \operatorname*{arg\,max}_{\mu}[\mathsf{MC}_j(\mu) \leq \mathsf{MC}_i(\lfloor \mu_i' \rfloor)] \geq \lceil \mu_j^a \rceil,$$

where the last inequality is since $MC(\mu) = MC(\lceil \mu \rceil)$ for any μ , this contradicting with the assumption that $\mu'_i < \mu^a_i$. By the greedy

natural, with a similar contradiction argument we can show that $MC_i(\mu_i^a) > MC_i(\mu_i')$, we then have:

$$\mathsf{MC}_{i}(\mu_{i}^{a}) > \mathsf{MC}_{i}(\mu_{i}^{\prime}) > \mathsf{MC}_{i}(\mu_{i}^{a}), \tag{5}$$

where the last inequality is due to the assumption that $\mu'_i \ge \mu^a_i$ and the monotonicity of MC functions.



Now we argue that $\boldsymbol{\mu}^a$ can be further improved by an exchange argument, contradicting with the assumption that $\boldsymbol{\mu}^a$ is optimal. Consider again the bidding strategy $\boldsymbol{\mu}^a$, consider reduce μ_j^a by some ϵ amount, and increase on μ_i^a by the corresponding amount until the constraints are tight again, since $MC_j(\mu_j^a) \ge MC_i(\mu_i^a)$, the "bang per buck" for the exchange portion strictly increases, contradicting with the assumption that $\boldsymbol{\mu}^a$ is optimal.

4 The Median of the Medians Algorithm

In this section, with the help of the characterization in the previous section, we present an algorithm with a worst-case query complexity of $O(m \log(mn) \log n)$. Note that if our feasible region were downward-closed, there would be a straightforward algorithm to solve the problem: We could perform a high-dimensional binary search in the bidding space, cutting down the whole strategy space by a constant fraction each time we query a particular strategy. This would lead to $m \cdot \log(n^m) = m^2 \log n$ queries (since querying one vector of strategies requires submitting a bidding strategy on each of the *m* platforms). Unfortunately, in the example below we show that our feasible region is not necessarily downward-closed.

Example 1. Consider a simple example with two platforms, 1 and 2, both having the following cost and value functions:

$$c_1(\mu) = \mu \quad v_1(\mu) = \frac{8}{3}\mu$$

$$c_2(\mu) = \mu \quad v_2(\mu) = \mu.$$

Suppose the constraints are B = 10 and $T = \frac{1}{2}$. First, observe that $\mu_1 = \frac{3}{2}$ and $\mu_2 = 1$ is a feasible solution since:

$$2 \cdot \left(\frac{3}{2} + 1\right) = \frac{8}{3} \cdot \frac{3}{2} + 1.$$

However, if we reduce μ_1 to 1, we get:

$$2 \cdot (1+1) > \frac{8}{3} + 1,$$

indicating that the updated bidding strategy is no longer feasible. Therefore, the feasible region is not downward-closed. Without a downward-closed feasible region, it is unclear which bidding strategy to try or how the search algorithm should proceed to minimize the number of attempts. To address this, we leverage the structure of integral bidding strategies shown in previous section and focus on identifying the *k* values corresponding to the maximum feasible μ^k first. The potential *k* values are the set of marginal costs across all platforms (there are *mn* of them). We utilize the "median of medians" idea to ensure that we eliminate a constant fraction of marginal cost options with each round of probing.

First, given the characterization of the optimal solution, we provide two subroutines that are useful for our algorithm. We first provide a subroutine named MATCHINGMC, that given a k, finds the μ^k vector via binary search on each platform. We show that the query complexity of this subroutine is $O(m \log n)$. Whenever we call the subroutine, we would make sure that the $MC_j(\ell_j) \leq k$, i.e., there is at least one strategy in the search range that is feasible. Due to space limitations, we defer the following two proofs to appendix B

SUBROUTINE 2: MATCHINGMC
Input: search range $[\ell_j, r_j]$ of each <i>j</i> , the target MC <i>k</i>
1 for $j \in M$ do
² while μ_i^k = NULL do
3 $\mu_j \leftarrow \frac{\ell_j + r_j}{2}$ for all $j \in M$ // Binary search on each
platform
4 if $MC_j(\mu_j) \le k$ then $\ell_j \leftarrow \mu_j$
5 if $MC_j(\mu_j) > k$ then $r_j \leftarrow \mu_j - 1$
6 if $r_j \leq \ell_j$ then $\mu_j^k \leftarrow \ell_j$
7 return $\boldsymbol{\mu}^k = (\mu_1^k, \mu_2^k, \dots, \mu_m^k)$

LEMMA 4. Given some $k \ge 0$, MATCHINGMC outputs the corresponding μ^k with at most $O(m \log \max_j (r_j - \ell_j))$ queries.

We now provide a subroutine that check if a given integral bidding profile μ is the almost-optimal solution (defined in Definition 1) or not. In addition, the subroutine can also check if a bidding profile is in the form of μ^k for some k (defined in Lemma 1). The worst-case query complexity is O(m).

LEMMA 5. Given a bidding profile μ , the subroutine OPTCHECK determines if the given μ is INFEASIBLE, NOT μ^k , NOT-OPTIMAL or ALMOST-OPTIMAL with at most O(m) queries.

We are now ready to present our algorithm, MEDIANOFMEDI-ANS. This algorithm finds the *almost-optimal integral solution* by searching within the marginal cost space and then converts this almost-optimal integral solution to the optimal fractional solution using the ROUNDUP procedure. The search process is inspired by the median-of-medians algorithm. In each iteration, we first identify the median marginal cost for each platform, and then select the median that most evenly splits the space, i.e., ensuring that the number of marginals weakly smaller than this median is equal to the number of marginals weakly larger than it.

Next, we use MATCHINGMC to determine the corresponding bidding profile μ^k with the median-of-the-medians marginal as the *k*-value, and apply OPTCHECK to evaluate the quality of the

Conference'17, July 2017, Washington, DC, USA

SUBROUTINE 3: OPTCHECK Input: some bidding strategy **µ** 1 query each platform *j* strategy μ_i , obtain $v_i(\mu_i)$, $c_i(\mu_i)$ and $MC_i(\mu_i)$ ² if μ is infeasible then return INFEASIBLE $3 \ \overline{j} \leftarrow \arg\max_{i} [MC_{i}(\mu_{i})]$ $k \leftarrow MC_{\bar{I}}(\mu_{\bar{I}})$ 5 for $i \in M$ do $\mu'_i \leftarrow \mu_i + 1$ // Check the next MC value for platform jquery μ'_i on platform *j* obtain $v_j(\mu'_i)$, $c_j(\mu'_j)$ and $MC_j(\mu'_i)$ if MC_{*i*}(μ'_i) $\leq k$ for any $j \neq \overline{j}$ then // μ is not μ^k for some k**return** NOT- μ^k $j^* \leftarrow \arg\min_i(MC_i(\mu'_i))$ // find the minimum among the next points $\mu_{j^*} \leftarrow \mu'_{j^*}$ 12 if *µ* is infeasible // the updated μ is not feasible 13 then return ALMOST-OPTIMAL 14 else return NOT-OPTIMAL

bidding profile μ^k . Based on the result from OPTCHECK(μ^k), we can eliminate a constant fraction of the remaining candidates for the optimal marginal costs. This process is repeated iteratively until we find an almost-optimal solution. Finally, we apply the ROUNDUP procedure to obtain the fractional optimal solution. For a formal description, please refer to Algorithm 1.

_	ALCORITHM 1. MEDIANOEMEDIANS
1	Initialize: $\ell_j \leftarrow 1, r_j \leftarrow n$ for all $j \in M$
2	while TRUE do
3	$\mu_j \leftarrow \frac{\ell_j + r_j}{2} \text{ for all } j \in M$
4	query each platform <i>j</i> strategy μ_i , obtain $v_i(\mu_i)$, $c_i(\mu_i)$
	and $MC_i(\mu_i)$
5	rank the platforms in non-decreasing order of μ_j s.t. if
	$i \leq j, \mu_i \leq \mu_j$
6	$j^* \leftarrow \min_j(\sum_{i \leq j} (r_i - \ell_i) - \sum_{i \geq j} (r_i - \ell_i))$ // find the
	j^{*} that equally split the search space
7	$\boldsymbol{\mu}^* \leftarrow \text{MatchingMC}([1, n] \text{ for all } j, \text{MC}_{j^*}(\mu_{j^*}))$
8	if OptCheck(μ^*) = INFEASIBLE then
9	$r_j \leftarrow \mu_j - 1 ext{ for all } j \geq j^*$ // reduce the search space
10	else if $OptCheck(\mu^*) = NOT-OPTIMAL$ then
11	$\ell_j \leftarrow \mu_j + 1 \text{ for all } j \leq j^*$
12	else if $OptCheck(\mu^*)$ = ALMOST-OPTIMAL then
13	return RoundUp(μ^*)

In the rest of the section, we prove the correctness and query complexity of the algorithm.

THEOREM 2. Given any instance I, the MEDIANOFMEDIANS algorithm finds the fractional optimal bidding strategy with at most $O(m \log mn \log n)$ queries.



increasing median marginal cost

Figure 1: Illustration of one round of MEDIANOFMEDIANS. Each column represents the current search region of a platform. The vertical arrow indicates the increasing direction of μ in each platform. The platforms are ranked by the median marginals as described in the algorithm. The grey point represents the queried k value. If μ^k is infeasible, all strategies in the shaded round rectangle are removed; otherwise, all strategies in the non-shaded one are removed.

PROOF. We first prove the correctness of the algorithm. By Lemma 3 we know that the almost-optimal integral bidding strategy correspond to $\boldsymbol{\mu}^{k^*}$ where k^* is the maximum k such that $\boldsymbol{\mu}^k$ is feasible. We prove the correctness of the algorithm by first showing that during the execution of the algorithm, there always exist some $\mu \in [\ell_j, r_j]$, of which the $\boldsymbol{\mu}^{MC_j(\mu)} = \boldsymbol{\mu}^{k^*}$. In other words, the algorithm can not eliminate the critical marginal cost $MC_j(\mu)$ that corresponds to $\boldsymbol{\mu}^{k^*}$. Consider the possible updates of ℓ_j and r_j for each platform j, i.e., Line 9 and Line 11. First consider any iteration such that OPTCHECK($\boldsymbol{\mu}^*$) = INFEASIBLE, and for all platforms $j \ge j^*$ w.r.t to the ranking defined in Line 5, we have $MC_j(\mu_j) \ge MC_{j^*}(\mu_{j^*})$. By monotonicity, for any $\mu \ge \mu_j$ on platform j we have:

$$\mathsf{MC}_{j}(\mu) \ge \mathsf{MC}_{j}(\mu_{j}) \ge \mathsf{MC}_{j^{*}}(\mu_{j^{*}})$$

By Lemma 1, since OPTCHECK(μ^*) = INFEASIBLE, we would also have $\mu^{MC_j(\mu)}$ is infeasible for $\mu \ge \mu_j$ for platforms $j \ge j^*$. Therefore, Line 9 does not remove any μ of which $\mu^{MC_j(\mu)} = \mu^{k^*}$.

Next consider any iteration such that $OPTCHECK(\mu^*) = NOT-OPTIMAL$, and for all platforms $j \le j^*$ w.r.t to the ranking defined in Line 5, we have $MC_j(\mu_j) \le MC_{j^*}(\mu_{j^*})$. Again by monotonicity, for any $\mu \le \mu_j$ on platform j we have:

$$\mathsf{MC}_{i}(\mu) \le \mathsf{MC}_{i}(\mu_{i}) \le \mathsf{MC}_{i^{*}}(\mu_{i^{*}}),$$

By Lemma 1, since OPTCHECK(μ^*) = NOT-OPTIMAL, we have $\mu^{MC_j(\mu)}$ is also feasible and not optimal for $\mu \leq \mu_j$ for platforms $j \leq j^*$. Therefore Line 11 does not remove any μ of which $\mu^{MC_j(\mu)} = \mu^{k^*}$ as well. In addition, it is easy to see that $\mu^{MC_j(\mu)} = \mu^{k^*}$ for some platform j and some strategy μ . (let $MC_j(\mu) = \arg \max(MC_j(\mu_j^{k^*}))$). And since the set of bids is finite and getting strictly smaller in each round, the algorithm will eventually terminate with the almostoptimal integral bidding solution μ^{k^*} , after which applying ROUNDUP would give us the fractional optimal solution.

Multi-Platform Autobidding with and without Predictions

We now prove the query complexity of the algorithm. In particular, we argue that the while loop would iterate no more then $O(\log(mn))$ times. Together with the $O(m \log n)$ query complexity of MATCHINGMC this would show that the query complexity of the algorithm is $O(m \log(mn) \log n)$. First note that there are in total $m \cdot n$ possible marginal costs (MC_i(μ) for all *j* and μ). By definition of j^* , and μ_j for each platform j, we have that $\min(|\{\mathsf{MC}_i(\mu) : i \leq j^* \text{ and } \mu \leq \mu_i\}|, |\{\mathsf{MC}_i(\mu) : i \geq j^* \text{ and } \mu \geq j^* \text{ a$ $|\mu_i|| \ge \frac{\sum_j (r_j - \ell_j)}{4} - \min_j (r_j - \ell_j) = O(\sum_j (r_j - \ell_j)).$ Since we remove a constant fraction of choices in each round, the number of queries is no more then $O(\log mn)$. Lastly note that ROUNDUP makes at most *m* queries, making the total queries needed for ME-DIANOFMEDIANS $O(m \log(mn) \log n)$.

Lower Bounds on Query Complexity

In this section, we provide some lower bounds on query complexity. We show that any algorithm needs to have a query complexity of $\Omega(m \log n)$ even if it knows the optimal marginal cost k. We also provide a lower bound of $\Omega(\log mn)$ for finding the optimal marginal cost *k* even when the algorithm is given a single-query black-box oracle for MATCHINGMC. Note that our algorithm MEDIANOFMEDI-ANS finds the optimal solution essentially by searching for the optimal marginal cost using $O(\log mn)$ calls to MATCHINGMC which itself costs $(m \log n)$ queries, and the query complexity upper bound is in fact the product of the two aforementioned lower bounds. This suggests that improving the query complexity upper bound further would require an algorithm that does not treat MATCHINGMC as a black-box. Due to space limitations, we defer the proofs in this section to Appendix C.

THEOREM 3. Any algorithm needs $\Omega(m \log n)$ queries to find the optimal bidding strategy, even if it knows the value k for which μ^k is the almost-optimal integral bidding strategy.

THEOREM 4. Any algorithm needs $\Omega(\log(mn))$ queries to find the optimal bidding strategy, even if it has access to a black-box oracle of MATCHINGMC that uses a single query.

Learning-Augmented Algorithms

In this section we aim to design searching algorithm that utilize a (possibility erroneous) prediction $\hat{\mu}$ regarding the actual optimal fractional strategy μ^{o} . The error of the prediction is measured by its distance to the optimal solution in the ℓ -infinity norm, i.e.

$$\eta = \max_{i} |\hat{\mu}_{j} - \mu_{j}^{o}|. \tag{6}$$

We show the following algorithm, modified from MEDIANOFMEDI-ANS, achieves a query complexity of $O(m \log m\eta \log \eta)$, note that since $\eta \leq n$, this guarantee matches the query complexity of MEDI-ANOFMEDIANS even if the prediction is arbitrarily wrong.

The algorithm begins by checking whether the floor of the pre-dicted bidding strategy, $\lfloor \hat{\mu}_i \rfloor$, for all *j*, is ALMOST-OPTIMAL using OPTCHECK. If it is, the algorithm applies ROUNDUP and returns the optimal solution. If not, the algorithm assumes the error is small and attempts to search for the optimal solution within a restricted range around the predicted strategy $\hat{\mu}_i$ on each platform, following a similar approach to the MEDIANOFMEDIANS algorithm.

If the optimal solution is still not found, the search range is expanded, and the search is repeated. This process continues until a almost-optimal solution is identified. By progressively expanding the search range as the square of the previous range, we show that the query complexity is at most $O(m \log(m\eta) \log \eta)$. Please refer to Algorithm 2 for a formal description.

A	LGORITHM 2: BRANCHOUTMEDIANOFMEDIANS	
1]	Initialize: $\ell_j \leftarrow \hat{\mu}_j, r_j \leftarrow \hat{\mu}_j$ for all $j \in M$	
2 j	$\hat{\mu} \leftarrow \hat{\mu}$	
3 i	f OptCheck($\hat{\mu}$) == ALMOST-OPTIMAL then return	
	RoundUp($\hat{\boldsymbol{\mu}}$)	
4 i	$\leftarrow 0$ // initialize the counter for doubling process	
5	while TRUE do	
6	$\ell_i \leftarrow \hat{\mu}_i - 2^{2^i}, r_i \leftarrow \hat{\mu}_i + 2^{2^i}$ for all $j \in M$	
7	range-indicator ← TRUE // assume range is correct	
8	while range-indicator == TRUE do	
9	$\mu_i \leftarrow \frac{\ell_j + r_j}{\ell_j}$ for all $j \in M$	
10	guery each platform <i>j</i> strategy μ_j , obtain $v_j(\mu_j)$,	
	$c_i(\mu_i)$ and MC _i (μ_i)	
11	rank the platforms in non-decreasing order of μ_i s.t.	
	$ \text{if } i \leq j, \mu_i \leq \mu_j$	
12	$j^* \leftarrow \min_j (\sum_{i \le j} (r_i - \ell_i) - \sum_{i \ge j} (r_i - \ell_i)) // \text{ find}$	
	the j^* that equally split the search space	
13	$k \leftarrow MC_{j^*}(\mu_{j^*})$	
14		
15	if there exist a $MC_j(\hat{\mu}_j - 2^{2^i}) > k$ then	
16	$\ell_j \leftarrow \mu_j + 1 \text{ for } j \leq j^*$ // k is too small	
17		
18	else	
19	$\boldsymbol{\mu}^{k} \leftarrow \text{MatchingMC}([\hat{\mu}_{j} - 2^{2^{i}}, \hat{\mu}_{j} + 2^{2^{i}}] \forall j, k)$	
20	if OptCheck(μ^k)) == NOT- μ^k then	
21	$r_j \leftarrow \mu_j - 1 \text{ for all } j \ge j^*$ // k is too large	
22		
23	else if OptCheck(μ^k) == INFEASIBLE then	
24	$ \qquad \qquad r_j \leftarrow \mu_j - 1 \text{ for all } j \ge j^* // k \text{ is too large}$	
25		
26	else if OptCheck(μ^k) == ALMOST-OPTIMAL	
	then	
27	return RoundUp(μ^k)	
28		
29	else	
	// OptCheck(μ^k) == NOT-OPTIMAL	
30	$ \ell_j \leftarrow \mu_j + 1 \text{ for all } j \le j^* // k \text{ is too small}$	
31		
32	if there exist a platform with $\ell_j > r_j$ then	
33	range-indicator \leftarrow FALSE // search in the given	
	range is complete	
34	$i \leftarrow i+1$ // update the search range	

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

865

866

867

868

869

870

THEOREM 5. Given any instance I, and predicted optimal bidding strategy $\hat{\mu}$ such that the error of the $\hat{\mu}$ is η , the BRANCHOUTMEDIANOF-MEDIANS algorithm finds the optimal bidding strategy with at most $O(m \log m\eta \log \eta)$ queries.

PROOF. We first argue the correctness of the algorithm. Let η denote the error of the prediction as defined in (6), and let i^* be the smallest *i* such that $2^{2^i} \ge (\eta + 1)$. We will show that the algorithm does not terminate in any round $i < i^*$. Let μ^* represent the almostoptimal integral bidding strategy. When $i < i^*$, there exists at least one platform *j* such that $\mu_i^* \notin [\ell_j, r_j]$ at the beginning of the *i**-th iteration of the outer while loop. By the correctness of OPTCHECK and since the algorithm only searches within the range $[\ell_i, r_i]$, it cannot return a solution in earlier rounds.

Next, we argue that the algorithm will terminate in round i^* with the optimal solution. Since $2^{2^{l^*}} \ge (\eta + 1)$, we know that $\mu_i^* \in [\ell_j, r_j]$ for all *j* at the start of the *i*^{*}-th iteration of the outer while loop. Now, consider the search process during this round. As in the proof of Theorem 2, we show correctness by proving that during the execution of the i^* -th iteration, there always exists some $\mu \in [\ell_i, r_i]$ such that $\mu^{MC_j(\mu)} = \mu^*$. Considering all possible updates to ℓ_i and r_i for each platform, we will now show that none of these updates eliminate any such μ values.

First, in Line 16, the algorithm encounters a platform *j* where $MC_i(\hat{\mu}_i - 2^{2^{i^*}}) > k$, meaning the current candidate marginal cost k is smaller than the marginal cost of the smallest strategy within the current search range for that platform. Since $\mu_i^* \in [\hat{\mu}_i - 2^{2^{i^*}}, \hat{\mu}_i +$ $2^{2^{i^*}}$] for all platforms and $\mu^* = \mu^k$ for some k, this implies that the current marginal cost candidate k, as well as all marginal costs weakly smaller than k, cannot correspond to the optimal marginal $\cot \mu^*$. These marginal costs (and their corresponding strategies) are thus eliminated from the search range.

In Line 21, the algorithm is in the case where $OPTCHECK(\mu^k) ==$ NOT – μ^k , indicating that μ^k is not optimal. This implies that there exists a platform *j* such that: 1. $\mu_j^k = \hat{\mu}_j + 2^{2^{j^*}}$, i.e., the largest strategy, and 2. for the same platform j, MC $_i(\hat{\mu}_i + 2^{2^{l^*}} + 1) \le k$. This means that k, along with all marginal costs weakly greater than k, exceeds the optimal marginal cost corresponding to μ^* . These marginal costs (and their corresponding strategies) are therefore eliminated from the search range.

The remaining cases are handled in the same way as discussed in Theorem 2. In Line 24, when μ^k is infeasible, we eliminate all marginal costs weakly greater than the current one being tested. In Line 30, when μ^k is not optimal, we eliminate all marginal costs weakly smaller than the current one. Finally, in Line 27, once we find μ^* , we use ROUNDUP to obtain the optimal fractional strategy.

We now prove the query complexity of the algorithm. Let i^* be the value of *i* when the algorithm terminates. First, we have $i^* \leq \eta^2$, where η is the error of the prediction as defined in (6). By Lemma 4, we know that the MATCHINGMC operation in iteration *i* takes $m \log 2^{2^i}$ time. Additionally, by Theorem 2, the while loop within this iteration will run $log(m2^{2^{i}})$ times. Since all other subroutines take O(m) queries, and the size of the search range is squared at each step, the algorithm terminates when the search

Anon

871

872

873

874

875

876

877

878

879

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918

919

920

921

922

923

924

925

926

927

928

space is weakly larger than *n*.

$$\begin{split} & \sum_{i=0}^{\log \log i^*} m \log(m \cdot 2^{2^i}) \cdot \log 2^{2^i} \\ &= \sum_{i=0}^{\log \log i^*} m (\log m + 2^i) 2^i \\ &= m (\log m + 2^{\log \log i^* + 1}) \cdot 2^{\log \log i^* + 1} \\ &= m (\log m + 2 \log i^*) \cdot 2 \log i^* \\ &\leq m (\log m + 2 \log(\eta)^2) \cdot 2 \log(\eta)^2 \\ &= m (\log m + 4 \log(\eta)) \cdot 2 \log(\eta) \\ &= O(m \log(m\eta) \log \eta) \end{split}$$

As a corollary, we also achieved "best-of-both-worlds" results in terms of consistency and robustness. Specifically, if the provided prediction is correct (or even "almost correct," i.e., $\lfloor \hat{\mu} \rfloor = \lfloor \mu^o \rfloor$), only 2m queries are required (note that even checking that a bidding profile is feasible requires *m* queries). Since $\eta \leq n$ by definition, the total number of queries will never exceed $O(m \log(mn) \log n)$, which matches the query complexity of MEDIANOFMEDIANS. Formally.

COROLLARY 1. The BRANCHOUTMEDIANOFMEDIANS algorithm is 2m-consistent and $O(m \log mn \log n)$ robust, where the robustness matches the query complexity of MEDIANOFMEDIANS.

Conclusion 7

In this work, we addressed the challenge of finding the optimal bidding strategy for advertisers operating in a multi-platform auction environment with low query complexity. Our approach models competition within each platform through value and cost functions that map various bidding strategies to their respective outcomes. We introduced an efficient algorithm that achieves this goal with a query complexity of $O(m \log(mn) \log n)$, where *m* represents the number of platforms and *n* denotes the number of potential bidding strategies available on each platform.

To further enhance efficiency, we incorporated the learningaugmented framework, proposing an algorithm that leverages a potentially flawed prediction of the optimal bidding strategy. Our results provide a query complexity bound that degrades gracefully, achieving O(m) queries when accurate predictions are available and $O(m \log(mn) \log n)$ even with completely incorrect predictions. This flexibility exemplifies a "best-of-both-world" scenario, providing advertisers with different options to effectively navigate the complexities of multi-platform bidding with minimal queries.

We believe that autobidding in multi-platform auction settings is understudied, and many intriguing questions remain unanswered. One immediate question for exploration is closing the gap between the upper and lower bounds established in our work, which would necessitate different tools and ideas. Additionally, it would be interesting to analyze the dynamics of the market if all the bidders adopted the approach presented in this study in determining their bidding strategies.

Multi-Platform Autobidding with and without Predictions

Conference'17, July 2017, Washington, DC, USA

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043 1044

929 References

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

986

- G. Aggarwal, A. Badanidiyuru, and A. Mehta. Autobidding with constraints. In I. Caragiannis, V. S. Mirrokni, and E. Nikolova, editors, Web and Internet Economics - 15th International Conference, WINE 2019, New York, NY, USA, December 10-12, 2019, Proceedings, volume 11920 of Lecture Notes in Computer Science, pages 17–30. Springer, 2019. doi: 10.1007/978-3-030-35389-6_2. URL https://doi.org/10.1007/ 978-3-030-35389-6 2.
- [2] G. Aggarwal, A. Perlroth, and J. Zhao. Multi-channel auction design in the autobidding world. In *Proceedings of the 24th ACM Conference on Economics* and Computation, EC '23, page 21, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701047. doi: 10.1145/3580507.3597707. URL https://doi.org/10.1145/3580507.3597707.
- [3] G. Aggarwal, A. Badanidiyuru, S. R. Balseiro, K. Bhawalkar, Y. Deng, Z. Feng, G. Goel, C. Liaw, H. Lu, M. Mahdian, J. Mao, A. Mehta, V. Mirrokni, R. P. Leme, A. Perlroth, G. Piliouras, J. Schneider, A. Schvartzman, B. Sivan, K. Spendlove, Y. Teng, D. Wang, H. Zhang, M. Zhao, W. Zhu, and S. Zuo. Auto-bidding and
- auctions in online advertising: A survey. SIGecom Exch., 22(1):159–183, Oct. 2024.
 doi: 10.1145/3699824.3699838. URL https://doi.org/10.1145/3699824.3699838.
 [4] G. Aggarwal, A. Perlroth, A. Schvartzman, and M. Zhao. Platform competition
- [4] G. Aggai wai, A. Ferrouri, A. Schwarzman, and M. Zhao. Fratorin competition in the autobidding world. arXiv preprint arXiv:2405.02699, 2024.
 [5] D. A. D. B. M. Li, V. Club, T. T. C. Li, T. T. L. Li, and M. S. L
- [5] P. Agrawal, E. Balkanski, V. Gkatzelis, T. Ou, and X. Tan. Learning-augmented mechanism design: Leveraging predictions for facility location. In EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 -15, 2022, pages 497–528. ACM, 2022.
- [6] A. Antoniadis, T. Gouleakis, P. Kleer, and P. Kolev. Secretary and online matching problems with machine learned advice. *Discret. Optim.*, 48(Part 2):100778, 2023.
- [7] Y. Azar, D. Panigrahi, and N. Touitou. Online graph algorithms with predictions. Proceedings of the Thirty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, 2022.
- [8] M.-F. Balcan, S. Prasad, and T. Sandholm. Bicriteria multidimensional mechanism design with side information. In *Proceedings of the 37th International Conference* on Neural Information Processing Systems, NIPS '23, Red Hook, NY, USA, 2024. Curran Associates Inc.
- [9] E. Balkanksi, V. Gkatzelis, and X. Tan. Mechanism design with predictions: An annotated reading list. SIGecom Exchanges, 21(1):54–57, 2023.
- [10] E. Balkanski, V. Gkatzelis, and X. Tan. Strategyproof scheduling with predictions. In Y. T. Kalai, editor, 14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA, volume 251 of LIPIcs, pages 11:1–11:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi: 10.4230/LIPIcs.ITCS.2023.11. URL https://doi.org/10.4230/LIPIcs.ITCS. 2023.11.
- [11] E. Balkanski, V. Gkatzelis, X. Tan, and C. Zhu. Online mechanism design with predictions. *CoRR*, abs/2310.02879, 2023. URL https://doi.org/10.48550/arXiv.2310. 02879.
- [12] É. Bamas, A. Maggiori, and O. Svensson. The primal-dual method for learning augmented algorithms. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020.
- [13] S. Banerjee, V. Gkatzelis, A. Gorokh, and B. Jin. Online nash social welfare maximization with predictions. In *Proceedings of the 2022 ACM-SIAM Symposium* on Discrete Algorithms, SODA 2022. SIAM, 2022.
- [14] Z. Barak, A. Gupta, and I. Talgam-Cohen. MAC advice for facility location mechanism design. CoRR, abs/2403.12181, 2024. doi: 10.48550/ARXIV.2403.12181. URL https://doi.org/10.48550/arXiv.2403.12181.
- [15] B. Berger, M. Feldman, V. Gkatzelis, and X. Tan. Optimal metric distortion with predictions. CoRR, abs/2307.07495, 2023.
- [16] I. Caragiannis and G. Kalantzis. Randomized learning-augmented auctions with revenue guarantees. arXiv preprint arXiv:2401.13384, 2024.
- [17] G. Christodoulou, A. Sgouritsa, and I. Vlachos. Mechanism design augmented with output advice. arXiv preprint arXiv:2406.14165, 2024.
- [18] R. Colini-Baldeschi, S. Klumper, G. Schäfer, and A. Tsikiridis. To trust or not to trust: Assignment mechanisms with predictions in the private graph model. *CoRR*, abs/2403.03725, 2024.
- [19] Y. Deng, J. Mao, V. Mirrokni, and S. Zuo. Towards efficient auctions in an autobidding world. In *Proceedings of the Web Conference 2021*, pages 3965–3973, 2021.
- [20] Y. Deng, N. Golrezaei, P. Jaillet, J. C. N. Liang, and V. Mirrokni. Multi-channel autobidding with budget and roi constraints. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.
- [21] P. Dütting, S. Lattanzi, R. P. Leme, and S. Vassilvitskii. Secretaries with advice. In P. Biró, S. Chawla, and F. Echenique, editors, EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021, pages 409–429. ACM, 2021.
- [22] K. Fujii and Y. Yoshida. The secretary problem with predictions. CoRR, abs/2306.08340, 2023.
- [23] V. Gkatzelis, K. Kollias, A. Sgouritsa, and X. Tan. Improved price of anarchy via predictions. In *Proceedings of the 23rd ACM Conference on Economics and*

Computation, pages 529-557, 2022.

- [24] S. Im, R. Kumar, M. M. Qaem, and M. Purohit. Online knapsack with frequency predictions. In Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, pages 2733–2743, 2021.
- [25] G. Istrate and C. Bonchis. Mechanism design with predictions for obnoxious facility location. *CoRR*, abs/2212.09521, 2022.
- [26] G. Istrate, C. Bonchis, and V. Bogdan. Equilibria in multiagent online problems with predictions. *CoRR*, abs/2405.11873, 2024.
- [27] A. Lindermayr and N. Megow. Alps, 2024. URL https://algorithms-with-predictions.github.io/.
- [28] P. Lu, Z. Wan, and J. Zhang. Competitive auctions with imperfect predictions. *CoRR*, abs/2309.15414, 2023.
- [29] T. Lykouris and S. Vassilvtskii. Competitive caching with machine learned advice. In International Conference on Machine Learning, pages 3296–3305. PMLR, 2018.
- [30] A. M. Medina and S. Vassilvitskii. Revenue optimization with approximate bid predictions. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, pages 1858–1866, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/ 884d79963bd8bc0ae9b13a1aa71add73-Abstract.html.
- [31] M. Mitzenmacher and S. Vassilvitskii. Algorithms with predictions. Commun. ACM, 65(7):33–35, 2022.
- [32] R. Paes Leme, B. Sivan, and Y. Teng. Why do competitive markets converge to first-price auctions? In *Proceedings of The Web Conference 2020*, WWW '20, page 596–605, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450370233. doi: 10.1145/3366423.3380142. URL https://doi.org/10.1145/ 3366423.3380142.
- [33] M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ML predictions. In Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, pages 9684–9693, 2018.
- [34] F. Susan, N. Golrezaei, and O. Schrijvers. Multi-platform budget management in ad markets with non-ic auctions. arXiv preprint arXiv:2306.07352, 2023.
- [35] C. Xu and P. Lu. Mechanism design with predictions. In L. D. Raedt, editor, Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022, pages 571–577. ijcai.org, 2022.

A Proof of Lemma 2

PROOF. For presentation purpose, we drop the subscript *j* in this prove as it should hold for any platform. We define c(0)/v(0) = 0 We prove the statement via induction. First consider the base case for $\mu = 1$, we have $c(1)/v(1) \ge 0 = c(0)/v(0)$ since both the cost and the value functions weakly increase w.r.t μ , we also have $c(1)/v(1) = \frac{c(1)-c(0)}{v(1)-v(0)} = MC(1)$ by definition. The base case is therefore established.

therefore established. Let $\frac{c(\mu)}{v(\mu)} = X_{\mu}$. Assume, for induction, that $X_{\mu'} \leq MC(\mu')$ for any $\mu' < \mu$. We first show $X_{\mu-1} \leq X_{\mu}$ holds for $\mu \geq 2$. Consider

$$\begin{split} c(\mu) &= X_{\mu} \cdot v(\mu) \\ c(\mu) - c(\mu - 1) &= X_{\mu} \cdot v(\mu) - c(\mu - 1) \\ \mathsf{MC}(\mu) \cdot (v(\mu) - v(\mu - 1)) &= X_{\mu} \cdot v(\mu) - c(\mu - 1) \\ \mathsf{MC}(\mu - 1) \cdot (v(\mu) - v(\mu - 1)) &= X_{\mu} \cdot v(\mu) - c(\mu - 1) \\ X_{\mu - 1} \cdot (v(\mu) - v(\mu - 1)) &\leq X_{\mu} \cdot v(\mu) - c(\mu - 1) \\ X_{\mu - 1} \cdot (v(\mu) - v(\mu - 1)) &\leq X_{\mu} \cdot v(\mu) - X_{\mu - 1} \cdot v(\mu - 1) \\ X_{\mu - 1} \cdot v(\mu) &\leq X_{\mu} \cdot v(\mu) \end{split}$$

$$X_{\mu-1} \le X_{\mu},$$

where the third equality is by definition of MC, the forth equality is by monotoncity of MC, and the first inequality is by induction hypothesis $MC(\mu - 1) \ge X_{\mu-1}$. In addition, consider the same set of equation again:

$$c(\mu) = X_{\mu} \cdot v(\mu)$$

$$c(\mu) - c(\mu - 1) = X_{\mu} \cdot v(\mu) - c(\mu - 1)$$

$$\begin{split} c(\mu) - c(\mu - 1) &= X_{\mu} \cdot v(\mu) - X_{\mu - 1} \cdot v(\mu - 1) \\ c(\mu) - c(\mu - 1) &\geq X_{\mu} \cdot v(\mu) - X_{\mu} \cdot v(\mu - 1) \\ \mathsf{MC}(\mu) \cdot (v(\mu) - v(\mu - 1)) &\geq X_{\mu} \cdot (v(\mu) - v(\mu - 1)) \\ \mathsf{MC}(\mu) &\geq X_{\mu} \end{split}$$

where as the first inequality is due to $X_{\mu-1} \leq X_{\mu}$, we therefore have $MC(\mu) \geq X_{\mu}$, hence proved.

B Missing Proofs from Section 4

B.1 Proof of Lemma 4

PROOF. The algorithm performs binary search on each platform to find the maximum μ such that $MC_j(\mu) \leq k$, since binary search checks $O(\log(r_j - \ell_j))$ number of choice of μ on each platform and there are *m* platforms, the query complexity is $O(m \log \max_j (r_j - \ell_j))$.

We now prove the correctness of the algorithm via case analysis, i.e., for each platform *j*, we have $\mu_j^k = \max_{\mu}(\mathsf{MC}_j(\mu) \le k)$. Fix any arbitrary platform *j*, if $\mathsf{MC}_j(r_j) \le k$, the algorithm will keep update ℓ_j until eventually $\ell_j = r_j$ and correctly set $\mu_j^k = \ell_j = r_j$. On the other hand, if $\mathsf{MC}_j(r_j) > k$, by the termination condition, we know that $\mathsf{MC}_j(\mu_j^k + 1) > k$, $\mathsf{MC}_j(\mu_j^k) \le k$, which corresponds to μ_j^k being the maximum bid with a marginal cost weakly less than *k*. \Box

B.2 Proof of Lemma 5

PROOF. Since OPTCHECK queries at most 2 strategies from each platform, the worst-case number of queries used is 2m = O(m). We now prove the correctness of the subroutine for each different case. The INFEASIBLE case is trivial. For the NOT- μ^k case, as indicated by line 8, since there exists a platform where $MC_j(\mu_j + 1) \le k$, we know that the given bidding profile μ is not μ^k for some k by definition. On the other hand, if OPTCHECK does not terminate in line 9, it means μ is feasible and $\mu = \mu^k$ for some k. To check if the given profile is almost-optimal (the floor of μ^o), by Lemma 3, we just need to verify whether increasing k would make μ^k infeasible. If the next immediate change would cause μ to be infeasible, then μ is ALMOST-OPTIMAL; otherwise, it is NOT-OPTIMAL.

C Missing proofs from section 5

C.1 Proof of Theorem 3

PROOF. Given any algorithm, assume it knows the correct value of k. On each platform, finding the maximum μ (therefore the $\lceil \mu_j^o \rceil$) such that $MC_j(\mu) \leq k$ takes at least $\Omega(\log n)$ queries. We prove this via a decision tree argument similar to the $\Omega(\log n)$ query complexity for the binary search problem.

Fix an arbitrary platform j; we want to determine the maximum index μ such that $MC_j(\mu) \leq k$. We represent any algorithm as a decision tree as follows:

(1). Each query made to the platform is represented as a node in the decision tree, and each node has three children: one for $MC_j(\mu) \le k$, one for $MC_j(\mu) > k$, and a third for cases not specified. (2). The leaves of this tree represent the possible outcomes of the search: specifically, finding the maximum index μ such that $MC_j(\mu) \le k$.

There are n+1 distinct outcomes, corresponding to the maximum value of μ being 0, 1, ..., or n. In any decision tree with x leaves, the minimum height h is log x.

Moreover, the height *h* of the decision tree corresponds to the number of queries made. Therefore, the minimum height of the decision tree is $\log(n+1)$, implying that the number of queries needed to resolve the search will be at least $\Omega(\log(n + 1)) = \Omega(\log n)$.

Lastly, since all platforms operate independently, the search on each platform requires $\Omega(\log n)$ queries. Consequently, to complete the search across all platforms will require $\Omega(m \log n)$ queries. \Box

C.2 Proof of Theorem 4

PROOF. There are a total of *mn* distinct marginal costs, and our objective is to determine the marginal cost $MC_j(\mu)$ for a specified *j* and μ such that $\mu^{MC_j(\mu)}$ represents the almost-optimal integral solution. We establish this by reducing the binary search problem to our problem.

Consider a binary search scenario involving a single sorted array A with |A| = h and a target number a for which we are searching. Let i denote the index of a within this array. We can construct an instance of our problem featuring a global ordering of marginal costs. In this global ordering, the marginal costs $MC_j(\mu)$ located at index i correspond to the bidding strategy $\mu^{MC_j(\mu)}$, which serves as the almost-optimal integral solution. If we are able to identify the index of the optimal marginal cost in fewer than $\Omega(\log mn)$ queries, it would consequently allow us to resolve the binary search problem in fewer than $\Omega(\log h)$ queries. This outcome would contradict the established complexity bounds associated with binary search.