

MIT Open Access Articles

Linear Program-Based Approximation for Personalized Reserve Prices

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Derakhshan, Mahsa, Golrezaei, Negin and Paes Leme, Renato. 2022. "Linear Program-Based Approximation for Personalized Reserve Prices." *Management Science*, 68 (3).

As Published: 10.1287/MNSC.2020.3897

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: <https://hdl.handle.net/1721.1/144168>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



LP-based Approximation for Personalized Reserve Prices

Mahsa Derakhshan*
University of Maryland

Negin Golrezaei
MIT

Renato Paes Leme
Google Research

Abstract

We study the problem of computing data-driven personalized reserve prices in eager second price auctions without having any assumption on valuation distributions. Here, the input is a data-set that contains the submitted bids of n buyers in a set of auctions and the problem is to return personalized reserve prices \mathbf{r} that maximize the revenue earned on these auctions by running eager second price auctions with reserve \mathbf{r} . For this problem, which is known to be NP-hard, we present a novel LP formulation and a rounding procedure which achieves a $(1 + 2(\sqrt{2} - 1)e^{\sqrt{2}-2})^{-1} \approx 0.684$ -approximation. This improves over the $\frac{1}{2}$ -approximation algorithm due to Roughgarden and Wang. We show that our analysis is tight for this rounding procedure. We also bound the integrality gap of the LP, which shows that it is impossible to design an algorithm with an approximation factor larger than 0.828 with respect to this LP.

1 Introduction

Second price (Vickrey) auctions with reserves have been prevalent in many marketplaces such as online advertising markets [CS14, PLPV16, GLMN17]. A key parameter of this auction format is its *reserve price*, which is the minimum price at which the seller is willing to sell an item. While there is empirical and theoretical evidence that highlights the significance of setting *personalized reserve prices* for the buyers to maximize the revenue [EOS07, OS11, BGL⁺18], we do not have a full understanding of how to optimize reserve prices. This problem is only fully solved under the assumption that buyers' valuation distributions are i.i.d. and regular, where these assumptions fail to hold in practice [CLMN14, GLMN17].

We study the problem of optimizing personalized reserve prices in second price auctions when the buyer valuations can be correlated. There are two different ways that personalized reserve prices can be applied in second price auctions: lazy and eager [DRY15]. In the lazy version, we first determine the potential winner and then apply the reserve prices. In the eager version, we first apply the reserve prices and then determine the winner. In this work, we focus on optimizing eager reserve prices because (i) while the optimal lazy reserve prices can be computed exactly in polynomial time, they have worse revenue performance both in theory and practice, and (ii) eager reserves perform better in terms of social efficiency for similar revenue levels [PLPV16].

To optimize the eager reserve prices, we take a data-driven approach as suggested in the literature [PLPV16, RW16]. The input in this setting is a history of the buyers' submitted bids/valuations

An earlier version of the paper has appeared at the ACM Conference on Economics and Computation (EC), 2019.

*Research conducted in part while the author was an intern at Google.

over multiple runs of an auction and the goal, roughly speaking, is to set a personalized reserve price r_b for each buyer b such that the total revenue obtained on the same data-set according to these reserve prices is maximized (see Section 2 for the formal definition).

The optimal data-driven reserve prices solve an *offline optimization problem*, i.e., given a data-set of bid data, it computes the optimal reserve prices in retrospect. This approach is mainly inspired by online advertising markets, in which billions of second price auctions are run in a day by ad exchanges. This practice provides ad exchanges a big data-set of submitted bids, which can be used to optimize personalized reserve prices for the next day via a data-driven approach. Setting personalized reserve prices, which is a common practice in this market, is inspired by targeting tactics in which advertisers take advantage of cookie-matching technologies to target Internet users based on their diverse preferences. This then leads to heterogeneous valuation/bid distributions, which necessitates setting personalized reserve prices.

Prior Work and Our Results: Unfortunately, optimizing the data-driven reserve prices is APX-hard [RW16]. The state-of-the-art algorithm of [RW16] achieves a $1/2$ -approximation which itself improves over an earlier $1/4$ -approximation algorithm by [PLPV16]. Our main result is an algorithm with a significantly improved approximation factor. We show that there exists a randomized polynomial time algorithm that given a data-set, outputs a vector of reserve prices whose expected revenue is a 0.684 -approximation of that of the optimal value. Our improved bound results in a polynomial time $(0.684 - \epsilon)$ -algorithm for independent distributions, which beats the best approximation known via prophet techniques.¹ Further, our result leads to a $(0.684 - \epsilon)$ -algorithm for the batch learning version of the problem. In the batch learning setting, there is a distribution over buyers’ valuations/bids and the goal is to compute the optimal prices by having access to samples from that distribution [MM14, HMR18]. Using the machinery developed by [MR15], one can show that via solving the data-driven offline optimization problem on the data-set with $\Omega(|B| \log |B|/\epsilon^2)$ auctions, we can obtain a $1 - \epsilon$ fraction of the maximum revenue of any eager second price auction that one could have hoped to obtain by knowing the valuation distributions.

The known algorithms of the literature are all greedy and only take into account the two highest bids in each auction. Another limitation of these algorithms is that the reserve price for each buyer is computed in isolation. That is, the reserve price for a buyer only depends on the bids of the auctions in which the buyer submits the highest bid. In fact, [RW16] argue that these limitations are precisely what prevent their algorithm from obtaining any guarantee better than $1/2$. As we explain in more detail later, we bypass this bound by a careful analysis of a rounding technique for a natural linear programming formulation of the problem proposed in this work.

Our Techniques: To obtain our improved approximation factor of 0.684 , we present an algorithm called “**Profile-based LP-Rounding**”, Pro-LPR for short, that takes advantage of a concise representation of the solution space. This representation, that we call profile space, is inspired by how revenue is computed in the eager auctions. Working with the profile space enables us to consider all the bids in an auction—not only the highest and second highest bids—to set the reserve prices. It further allows us to describe the optimal solution by a polynomial-size integer program. By relaxing the integrality constraints on the variables of the integer program, we construct a linear program (LP). The fractional solution of the LP is then rounded to obtain the reserve prices. The final reserve price of the algorithm is the best of the zero reserves and the reserves obtained from rounding the solution of the LP. The most technically challenging step in the analysis is to bound the approximation ratio. This is done via careful probabilistic analysis of the rounding procedure

¹While we provide a better guarantee against the optimal reserves, our technique does not provide approximation guarantees with respect to the optimal auction as prophet inequalities do.

which leads to a non-linear mathematical program bounding the ratio. Our last step is to use techniques from non-linear optimization to bound the solution of the mathematical program. We would like to emphasize that our analysis of our algorithm is tight in a sense that there is an example for which our algorithm cannot get an approximation factor better than 0.684.

We point out that the performance of our algorithm is evaluated against the optimal value of the LP, which is an upper bound on the maximum revenue. By analyzing the integrality gap of the LP, we show that no algorithm can obtain more than a 0.828 fraction of the optimal value of the LP; see Theorem 2. This highlights that our algorithm is evaluated against a powerful benchmark and despite that, it obtains 0.684 fraction of this powerful benchmark.

Managerial Insights and Numerical Studies: Our proposed algorithm highlights that there is significant value in considering all the submitted bids, not only the highest and second highest bids, to optimize reserve prices. By considering all the submitted bids, the algorithm can better identify and take advantage of the buyers’ bidding behavior. Furthermore, the design of our algorithm accentuates the importance of optimizing reserve prices jointly. (Recall that the prior work focused on optimizing reserve of each buyer separately.) Such a joint optimization problem can capture potential correlation in the submitted bids, and hence improve revenue. To illustrate this, we conduct numerical studies, where we compare our algorithm with the greedy algorithm of [RW16]. As stated earlier, this greedy algorithm obtains the best approximation factor prior to our work. We show that when bids are positively or negatively correlated across buyers, (i) our algorithm obtains at least a 0.98 fraction of the optimal revenue, and (ii) in 50% (respectively 75%) of the problem instances, our algorithm outperforms the greedy algorithm by at least 6% – 7% (respectively 2% – 3%). We obtain similar results when bids are independent across buyers.

1.1 Other Related Work

While we have discussed several closely related works, we further situate our work in the landscape of related work. We start with a more detailed comparison between our work and that of [RW16]. As stated earlier, the greedy nature of [RW16]’s algorithm prevents it from obtaining an approximation factor better than $1/2$. However, the greedy nature of their algorithm allows them to transform it into an online learning algorithm with a sublinear approximate regret. The online learning algorithm receives a set of submitted bids to an auction and its goal to set personalized reserve prices based on collected feedback in past auctions. [RW16] design a learning algorithm under a full information setting, where the auctioneer observes all the bids after running an auction. Very recently, using Blackwell Approachability, [NGW⁺20] show how to transform a variation of the greedy algorithm of [RW16] into its online counterpart under a bandit feedback structure. In this structure, after every auction, the auctioneer only observes the obtained revenue, rather than all the submitted bids. In this work, we only focus on the offline problem of optimizing reserve prices. Considering our improved approximation factor, an interesting future research direction is to study how to transform the Pro-LPR algorithm to its online counterpart with sublinear approximate regret.

Our work relates and contributes to the literature on the auctions and revenue-maximizing mechanisms in a single-item environment. Seminal contributions have been made by in [Mye81] under a critical assumption that buyers’ valuations are independent of each other. Specifically, [Mye81] shows that

when buyers’ valuation distributions are regular and i.i.d., the optimal mechanism can be implemented by running second price auctions with a non-anonymous reserve price. However, even under the independence assumption, when valuation distributions are irregular or heterogeneous,

the optimal mechanism has a rather complicated structure and its implementation heavily relies on the knowledge of the valuation distributions [CLMN14, RS16, GLMN17].

Considering the complexity of the optimal auction, there has been a growing body of literature on designing simple yet effective auctions that can be easily optimized; see, for example, [GLMN17, CLMN14, PLPV16, RW16, AB18, BFM12, BGL⁺18, DRY15]. Among those is second price auctions with reserve, which is very common in practice [BGL⁺18, PLPV16, CS14]. This auction format has a single parameter, called reserve price, per buyer that determines the minimum acceptable bid for the buyer. As stated earlier, to extract high revenue from buyers, it is very crucial to effectively optimize reserve prices. Such an optimization problem has been studied in different settings.

If the value distributions are independent, an improved approximation to personalized reserves are known via techniques like the correlation gap [CHMS10, Yan11] and prophet inequalities [KS78, HK81, ACK18, EHLM17, BGL⁺18, CSZ20] (to cite a few). The latest result is 0.669-approximation by [CSZ20]. Although those results are typically stated as an approximation ratio with respect to the (stronger) Myerson revenue benchmark, those are also the best-known approximation ratios with respect to the optimal reserve prices for independent distributions.

Another related stream of literature is on the design of auctions for correlated distributions. This line of work was pioneered by [Ron01] and [RS02]. The positive and negative results were later improved by [DFK11] and [PP11]. Our paper departs from this line work in the sense that we do not try to approximate the optimal incentive-compatible auction, but instead, we try to approximate the best auction in the subclass of second price auctions with reserve prices, since this is the auction format adopted by most online marketplaces, including online display advertising markets. To optimize reserve prices in second price auctions, as stated earlier, we take a data-driven approach. This approach allows us to indirectly exploit potential correlations in the submitted bids. In online advertising, for example, such correlation exists when different buyers/advertisers value some features of ad impressions, in the same way, e.g., they all prefer showing their ads to users who tend to click more on ads.

Finally, we review some of the work that bears some resemblance to our work from a technical perspective. In our work, we change the solution space to describe the optimal solution using a concise LP with a small integrity gap. A similar technique is used in different applications; see, for example, [BDHS17, BDD⁺17, BDHS18, BBD⁺19], [ADH⁺19] and [IKL⁺11] for the use of a similar technique in finding optimal strategies of Blotto, security games, and dueling games. While sharing the general technique, each paper uses specific properties of their problem to design their alternative solution space. Furthermore, there is another related line of work that design “configuration” LPs for problems related to resource allocation and job scheduling, [PS15, Sve12, AS10, BSS19]. This line of work is initiated by [BS06]. We note that configuration LPs, unlike our LP, may not be polynomial in the input size.

The rest of the paper is organized as follows. In Section 2, we define the model. Section 3 presents a high-level view of the results and techniques. In Section 4, we provide our LP, which will be used as our benchmark. In Section 5, we present the Pro-LPR algorithm and show its performance guarantee. Section 7 provides the proof of the integrality gap and Section 6 shows that our analysis is tight. Finally, we present the results of our numerical studies in Section 8 and conclude in Section 9.

2 Preliminaries and Problem Statement

There are n buyers participating in a set of single-item eager second price auctions. Let A and B respectively denote the set of auctions and buyers. For any buyer $b \in B$, and for any auction $a \in A$, we are given a non-negative number $\beta_{a,b}$ which indicates the bid of buyer b in auction a . Let r_b be the personalized reserve price of buyer $b \in B$. Then, given the bids $\{\beta_{a,b}\}_{b \in B}$ in auction $a \in A$ and reserve prices $\mathbf{r} = \{r_b\}_{b \in B}$, the eager second price (ESP) auction works as follows.

- First, any buyer b with $\beta_{a,b} < r_b$ is eliminated. Let $S_a = \{b : \beta_{a,b} \geq r_b\}$ be the set of buyers who clear their reserve prices in auction a .

- When set S_a is nonempty, the item is allocated to buyer $b_a^* = \arg \max_{b \in S_a} \{\beta_{a,b}\}$ who has the highest bid among all the buyers in set S_a and is charged

$$\text{Rev}_a(\mathbf{r}) := \max \left\{ r_{b_a^*}, \max_{b \in S_a, b \neq b_a^*} \{\beta_{a,b}\} \right\}.$$

Note that S_a and b_a^* implicitly depend on reserve prices \mathbf{r} . Any other buyer $b \in B$, $b \neq b_a^*$ is not charged. Further, when set S_a is empty, the item is not allocated and $\text{Rev}_a(\mathbf{r}) = 0$.

Note that the reserve prices are the same across all the auctions $a \in A$. However, each buyer b is assigned a personalized reserve price r_b . Given the data-set of bids $\{\beta_{a,b}\}_{a \in A, b \in B}$, our goal here is to find personalized reserve prices that maximize revenue of the auctioneer. See the introduction section for a discussion on the nice properties of this data-driven optimization. Formally, we would like to solve the following optimization problem:

$$\text{ESP}^* = \max_{\mathbf{r} \in \mathbb{R}^n} \text{Rev}(\mathbf{r}) := \sum_{a \in A} \text{Rev}_a(\mathbf{r}). \quad (\text{ESP-OPT})$$

Note that, without loss of generality, we assume that the optimal reserve price for buyer b is equal to one of his submitted bids $\{\beta_{a,b}\}_{a \in A}$. Let $R = \{0, \infty\} \cup \{\beta_{a,b}\}_{a \in A, b \in B}$. Then, Problem ESP-OPT can be rewritten as $\max_{\mathbf{r} \in R^n} \sum_{a \in A} \text{Rev}_a(\mathbf{r})$, which leads to a search space of size $|R|^n$.

We now make a few remarks about the data-driven optimization Problem (ESP-OPT). The optimal solution to the data-driven optimization Problem (ESP-OPT) gives an approximation solution to the batch-learning setting. In this setting, bids in each auction a , i.e., $\{\beta_{a,b}\}_{b \in B}$, are independent samples from the distribution \mathcal{D} . By analyzing the sample complexity of the auctions, [MM14] and [MR15] show that in the batch learning setting, with probability $1 - \delta$, it holds that:

$$\left| \mathbb{E}_{\beta \sim \mathcal{D}} [\text{Rev}(\beta; \mathbf{r})] - \frac{1}{m} \sum_{a=1}^m \text{Rev}(\{\beta_{a,b}\}_{b \in B}; \mathbf{r}) \right| \leq O \left(\sqrt{\frac{n \log(m/\delta)}{m}} \right),$$

where with a slight abuse of notation, $\text{Rev}(\beta; \mathbf{r})$ is the revenue of an auction when bids and reserve prices are β and \mathbf{r} , respectively. This implies that by having $m = \Omega(n \log n / \epsilon^2)$ samples, with high probability, the optimal solution to the data-driven optimization Problem (ESP-OPT) provides an ϵ -additive approximation solution to the batch-learning problem when the data-set of bids is drawn from a distribution \mathcal{D} .

Motivating the data-driven problems via the batch-learning problems implies that the bid distribution is the same across all the auctions. The consistency of the bid distribution across auctions, which is a common assumption in the literature (e.g., [PLPV16, RW16]), can be justified when buyers submit their true valuations in auctions, i.e., when they are truthful. The truthfulness is, in

fact, an appealing property of (single-shot) second price auctions. In single-shot second price auctions, buyers are only concerned about their utility in the current auctions and do not reason about how their bids will affect future outcomes. In other words, buyers are myopic, rather than being forward-looking/strategic. Nonetheless, there is a new line of work that studies how to optimize reserve prices in (lazy) second price auctions when buyers are forward-looking and would like to maximize their cumulative utility; see, for example, [ARS13, ARS14, KN19], and [GJM19, GJL19]. At a high level, it has been shown that in lazy second price auctions, it is possible to effectively optimize reserve prices even when buyers are forward-looking. We believe that some of the techniques developed for lazy second price auctions can be applied to eager second price auctions as well. Investigating this claim is indeed an interesting future research direction.

3 Results and Techniques

The main result of the paper is a randomized algorithm that returns an 0.684-approximation solution for Problem ESP-OPT.

Theorem 1 (Main Theorem). *There exists a randomized polynomial time algorithm that given a data-set $\{\beta_{a,b}\}_{a \in A, b \in B}$, outputs a vector of eager reserve prices whose expected revenue is a 0.684-approximation of that of the optimal value of Problem ESP-OPT, denoted by ESP^* .*

To find an approximate solution, the overall idea is to construct an LP whose objective function at its optimal solution provides an upper bound for ESP^* . The LP that takes advantage of a concise representation of the solution space, has a polynomial number of variables and constraints. Then, we use a rounding technique to transform the optimal solution of the LP to a vector of reserve prices. We show that if we consider the reserve prices obtained from the rounding technique and the vector of all-zero reserve prices and choose the one with the maximum revenue, we obtain the desired approximation factor. In Theorem 3, we further show that our analysis of our approximation factor is tight. That is, we provide an example for which our algorithm obtains exactly 0.684 fraction of the optimal value of the LP, i.e., the upper bound on for ESP^* . Finally, in Theorem 2, we bound the integrality gap of the LP. This characterization shows that no algorithm can obtain more than 0.828 fraction of the LP.

4 Linear Program

The main challenge in designing an LP formulation for this problem is to find a concise representation of the solution space. Instead of considering all possible assignments of reserves to buyers, we will consider only partial assignments in which we only specify the reserve prices of two buyers. We will call such partial assignment a *profile*. Formally, a profile is a tuple $p = (b_1, b_2, r_1, r_2) \in B \times B \times R \times R$, which represents an assignment of reserve r_1 to buyer b_1 and reserve r_2 to buyer b_2 . If it is the case that the reserves are below the corresponding bids in an auction a , i.e. $r_1 \leq \beta_{a,b_1}$ and $r_2 \leq \beta_{a,b_2}$, then no matter how the assignment of the remaining reserves, the revenue of this partial assignment is at least $\max\{r_1, \beta_{a,b_2}\}$ for $\beta_{a,b_1} \geq \beta_{a,b_2}$. We also note that given any vector of reserve prices \mathbf{r} , the revenue that can be obtained from \mathbf{r} only depends on the reserve price of the highest and second highest bidders that clear the reserve prices.

Next, we formally define the notion of *valid profile* and show that the Problem (ESP-OPT) can be relaxed to find the best consistent distribution over valid profiles in each auction. To define valid

profiles, we assume that in each auction \mathbf{a} , we have two auxiliary buyers \mathbf{b}_0 and \mathbf{b}_{00} who always bid zero. That is, $\mathbf{b}_{00}, \mathbf{b}_0 \in \mathbf{B}$, and $\beta_{\mathbf{a}, \mathbf{b}_0} = \beta_{\mathbf{a}, \mathbf{b}_{00}} = 0$ for any $\mathbf{a} \in \mathbf{A}$.

Definition 4.1 (Valid Profiles). *We define the set of valid profiles for auction \mathbf{a} as the set $\mathcal{P}_{\mathbf{a}}$ consisting of all tuples $(\mathbf{b}_1, \mathbf{b}_2, r_1, r_2) \in \mathbf{B} \times \mathbf{B} \times \mathbf{R} \times \mathbf{R}$ satisfies the following conditions:*

1. Bid of buyer \mathbf{b}_1 is greater than or equal to that of buyer \mathbf{b}_2 ; that is, $\beta_{\mathbf{a}, \mathbf{b}_1} \geq \beta_{\mathbf{a}, \mathbf{b}_2}$.
2. Buyer \mathbf{b}_1 clears his reserve; that is, $\beta_{\mathbf{a}, \mathbf{b}_1} \geq r_1$.
3. Buyer \mathbf{b}_2 clears his reserve; that is, $\beta_{\mathbf{a}, \mathbf{b}_2} \geq r_2$.

For any given profile $p \in \mathcal{P}_{\mathbf{a}}$, we define $\text{Rev}_{\mathbf{a}}(p) := \max(\beta_{\mathbf{a}, \mathbf{b}_2}, r_1)$.

We note that any valid profile corresponds to at least one vector of reserve prices. To see why, observe that we can always obtain $p = (\mathbf{b}_1, \mathbf{b}_2, r_1, r_2)$ by setting $r_{\mathbf{b}_1} = r_1$, $r_{\mathbf{b}_2} = r_2$, and $r_{\mathbf{b}} = \infty$ for any $\mathbf{b} \neq \mathbf{b}_1, \mathbf{b}_2$. Of course, there may exist other vectors of reserve prices that lead to the same profile. We note that by adding buyers \mathbf{b}_0 and \mathbf{b}_{00} to \mathbf{B} , we can define valid profiles to represent the cases in which less than two buyers cleared their reserve prices. We present the cases with one (respectively zero) cleared buyer with valid profile of $(\mathbf{b}_1, \mathbf{b}_0, r_1, 0)$ (respectively $(\mathbf{b}_0, \mathbf{b}_{00}, 0, 0)$).

Definition 4.2 (Profiles Associated with Reserve Prices). *Given a vector of reserve prices \mathbf{r} we say a valid profile $p = (\mathbf{b}_1, \mathbf{b}_2, r_1, r_2)$ is the unique profile associated with \mathbf{r} in an auction $\mathbf{a} \in \mathbf{A}$ if and only if the following condition hold. After applying the reserve prices \mathbf{r} , buyer \mathbf{b}_1 with reserve r_1 and buyer \mathbf{b}_2 with reserve r_2 have the highest and second highest cleared bids in auction \mathbf{a} , respectively.*

Given a vector of reserve prices \mathbf{r} and an auction \mathbf{a} , let p be the profile associated with \mathbf{r} in \mathbf{a} . Then, with a slight abuse of notation, we define $\text{Rev}_{\mathbf{a}}(\mathbf{r}) = \text{Rev}_{\mathbf{a}}(p)$.

We are now ready to describe our LP. The LP will have two sets of variables:

1. For any auction $\mathbf{a} \in \mathbf{A}$ and any valid profile $p \in \mathcal{P}_{\mathbf{a}}$, define a variable $s_{\mathbf{a}, p} \geq 0$ such that $\sum_{p \in \mathcal{P}_{\mathbf{a}}} s_{\mathbf{a}, p} \leq 1$. This variable represents a probability distribution over valid profiles in auction \mathbf{a} . We refer to $\{s_{\mathbf{a}, p} | \mathbf{a} \in \mathbf{A}, p \in \mathcal{P}_{\mathbf{a}}\}$ as a profile-weight.
2. For any buyer $\mathbf{b} \in \mathbf{B}$ and reserve price $r \in \mathbf{R}$, define a variable $q_{\mathbf{b}, r} \geq 0$ such that $\sum_{r \in \mathbf{R}} q_{\mathbf{b}, r} = 1$. This variable represents be the probability that buyer \mathbf{b} is assigned a reserve price of r .

We now discuss the LP constraints. We add constraints relating $s_{\mathbf{a}, p}$ and $q_{\mathbf{b}, r}$ which will ensure the consistency of probability distributions across all profiles. To define this set of constraints, for every $\mathbf{b} \in \mathbf{B}$, $\mathbf{a} \in \mathbf{A}$, and $r \in \mathbf{R}$, we define set

$$\mathcal{Q}_{\mathbf{b}, r, \mathbf{a}} := \{p = (\mathbf{b}_1, \mathbf{b}, r_1, r) : p \in \mathcal{P}_{\mathbf{a}}\} \cup \{p = (\mathbf{b}, \mathbf{b}_2, r, r_2) : p \in \mathcal{P}_{\mathbf{a}}\}, \quad (1)$$

which corresponds to all valid profiles of auction \mathbf{a} that assign reserve r to buyer \mathbf{b} . A natural constraint to add is that the total probability assigned to profiles in $\mathcal{Q}_{\mathbf{b}, r, \mathbf{a}}$ is at most the probability that buyer \mathbf{b} is assigned to reserve price r . That is,

$$\sum_{p \in \mathcal{Q}_{\mathbf{b}, r, \mathbf{a}}} s_{\mathbf{a}, p} \leq q_{\mathbf{b}, r}.$$

Finally, we can put it all together in the following LP:

$$\begin{aligned}
& \max_{\mathbf{q}, \mathbf{s}} && \sum_{\mathbf{a} \in \mathbf{A}} \sum_{p \in \mathcal{P}_{\mathbf{a}}} s_{\mathbf{a}, p} \cdot \text{Rev}_{\mathbf{a}}(p) \\
& \text{s.t.} && \sum_{p \in \mathcal{P}_{\mathbf{a}}} s_{\mathbf{a}, p} \leq 1 && \forall \mathbf{a} : \mathbf{a} \in \mathbf{A} \\
& && \sum_{p \in \mathcal{Q}_{\mathbf{b}, r, \mathbf{a}}} s_{\mathbf{a}, p} \leq q_{\mathbf{b}, r} && \forall \mathbf{a}, \mathbf{b}, r : \mathbf{b} \in \mathbf{B}, r \in \mathbf{R}, \mathbf{a} \in \mathbf{A} \\
& && \sum_{r \in \mathbf{R}} q_{\mathbf{b}, r} = 1 && \forall \mathbf{b} : \mathbf{b} \in \mathbf{B} \\
& && s_{\mathbf{a}, p} \geq 0 && \forall \mathbf{a}, p : \mathbf{a} \in \mathbf{A}, p \in \mathcal{P}_{\mathbf{a}} \quad (\text{Profile-LP})
\end{aligned}$$

We start by noting that the LP is a relaxation of the Problem (ESP-OPT):

Lemma 4.3 (Upper bound on Revenue). *The solution of Profile-LP is an upper bound to ESP*, i.e., the optimal value of Problem ESP-OPT. That is,*

$$\text{ESP}^* \leq \text{Profile-LP}.$$

Proof. Given reserve prices \mathbf{r}^* such that $\text{ESP}^* = \sum_{\mathbf{a}} \text{Rev}_{\mathbf{a}}(\mathbf{r}^*)$, we construct a feasible solution to the LP as follows. For each $\mathbf{a} \in \mathbf{A}$, we let $s_{\mathbf{a}, p} = 1$ for the profile p corresponding to \mathbf{r}^* (according to Definition 4.2) and $s_{\mathbf{a}, p} = 0$ for all remaining profiles. Further, we let $q_{\mathbf{b}, r_{\mathbf{b}}^*} = 1$ and $q_{\mathbf{b}, r} = 0$ for all remaining reserves. It is straightforward to verify that it is a feasible solution to the Profile-LP and that $\sum_{\mathbf{a} \in \mathbf{A}} \sum_{p \in \mathcal{P}_{\mathbf{a}}} s_{\mathbf{a}, p} \cdot \text{Rev}_{\mathbf{a}}(p) = \text{ESP}^*$. \square

Theorem 2 (Integrality Gap of Profile-LP). *There exists a data-set of bids $\{\beta_{\mathbf{a}, \mathbf{b}}\}_{\mathbf{a} \in \mathbf{A}, \mathbf{b} \in \mathbf{B}}$ for which the integrality gap of the LP is at least 0.828. That is,*

$$\text{ESP}^* \leq 0.828 \cdot (\text{Profile-LP}).$$

The proof of Theorem 2 is given in Section 7.

5 Profile-based LP-rounding (Pro-LPR) Algorithm

In this section, we present an algorithm, called Profile-based LP-rounding (Pro-LPR), that uses the optimal solution of (Profile-LP), \mathbf{s}^* , to devise reserve prices. The algorithm is presented below.

Profile-based LP-rounding (Pro-LPR) Algorithm:

Let \mathbf{s}^* and \mathbf{q}^* be the optimal solution of (Profile-LP). Then,

- Rounding procedure: For each buyer $\mathbf{b} \in \mathbf{B}$, independently sample reserve price $r \in \mathbf{R}$ with probability proportional to $q_{\mathbf{b}, r}^*$.
- Let \mathbf{z} be the vector of all zero reserves. Output the best of $\mathbf{r}^{\mathcal{R}}$ and \mathbf{z} , i.e.,

$$\mathbf{r}^{\text{out}} = \arg \max_{\mathbf{r} \in \{\mathbf{z}, \mathbf{r}^{\mathcal{R}}\}} \text{Rev}(\mathbf{r}).$$

In the Pro-LPR algorithm, we first round the optimal solution of the (Profile-LP) to construct reserve prices $\mathbf{r}^{\mathcal{R}}$. To do so, for each buyer $\mathbf{b} \in \mathbf{B}$, we independently sample reserve price $r \in \mathbf{R}$ with probability $q_{\mathbf{b},r}^*$, where \mathbf{q}^* (and \mathbf{s}^*) is the optimal solution of the (Profile-LP). We then compare revenue under $\mathbf{r}^{\mathcal{R}}$ with that under the zero reserve prices, and return the one that obtains higher revenue. The returned vector of reserve prices is denoted by \mathbf{r}^{out} .

We now proceed to analyze our algorithm. We show that $\mathbb{E}[\text{Rev}(\mathbf{r}^{\text{out}})]$ is at least a 0.684 fraction of the solution of the Profile-LP and hence at most $0.684 \cdot \text{ESP}^*$, where the expectation is with respect to the randomness in the Pro-LPR algorithm. As we show in Lemma 5.1, one of the biggest strengths of our LP formulation is that it allows the analysis to decouple the effect of rounding for each individual auction. In this lemma, roughly speaking, we present two conditions under which the Pro-LPR algorithm has a good performance. In these conditions, for each auction $\mathbf{a} \in \mathbf{A}$ and $t \geq 0$, we compare the probability that $\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}})$ is at least t , i.e., $\Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t]$, with $\sum_{\{p:p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^*$, which is the sum of the optimal weight of (valid) profiles in auction \mathbf{a} that obtains a revenue of at least t . Here, $\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}})$ is the revenue in auction \mathbf{a} under reserve prices $\mathbf{r}^{\mathcal{R}}$. Intuitively, the smaller the gap between $\Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t]$ and the aforementioned summation, the better the Pro-LPR algorithm performs. Lemma 5.1 makes this statement formal by considering the high revenue case of $t \geq \beta_{\mathbf{a}}^{(2)}$ (first condition) and the low revenue case of $t < \beta_{\mathbf{a}}^{(2)}$ (second condition), where $\beta_{\mathbf{a}}^{(2)}$ is the second highest bid in auction \mathbf{a} . Note that when the reserve price for the buyer with the highest bid in auction \mathbf{a} is set too high, revenue of this auction can be indeed less than the second highest submitted bid $\beta_{\mathbf{a}}^{(2)}$.

Lemma 5.1 (Two Conditions). *Let \mathbf{s}^* and \mathbf{q}^* be the optimal solution of (Profile-LP) and $\mathbf{r}^{\mathcal{R}}$ be a random reserve price obtained from the rounding procedure. If there exists a constant $c > 0$ such that for any $t \geq 0$ and any auctions $\mathbf{a} \in \mathbf{A}$, we have*

$$\sum_{\{p:p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* - \Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t] \leq 0 \quad \text{for } t > \beta_{\mathbf{a}}^{(2)} \quad (2)$$

$$\sum_{\{p:p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* - \Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t] \leq c \quad \text{for } t \leq \beta_{\mathbf{a}}^{(2)}, \quad (3)$$

then Pro-LPR algorithm is a $(1+c)^{-1}$ -approximation. That is, it obtains at least $(1+c)^{-1}$ fraction of the optimal value of Problem ESP-OPT. Here, $\beta_{\mathbf{a}}^{(2)}$ is the second highest bid in auction \mathbf{a} and $\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}})$ is the revenue in auction \mathbf{a} under reserve prices $\mathbf{r}^{\mathcal{R}}$.

Proof. By integrating over t in Equations (2) and (3) and adding them up, we get

$$\begin{aligned} & \int_{\beta_{\mathbf{a}}^{(2)}}^{\infty} \left(\sum_{\{p:p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* - \Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t] \right) dt \\ & + \int_0^{\beta_{\mathbf{a}}^{(2)}} \left(\sum_{\{p:p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* - \Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t] \right) dt \leq c \cdot \beta_{\mathbf{a}}^{(2)}. \end{aligned}$$

This is simplified as follows

$$\sum_{p \in \mathcal{P}_{\mathbf{a}}} s_{\mathbf{a},p}^* \text{Rev}_{\mathbf{a}}(p) - \mathbb{E}[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}})] \leq c \cdot \beta_{\mathbf{a}}^{(2)}. \quad (4)$$

Note that by Lemma 4.3, the optimal value of Problem ESP-OPT, denoted by ESP^* , is upper bounded by $\text{Rev}(\mathbf{s}^*)$. That is,

$$\text{ESP}^* \leq \text{Rev}(\mathbf{s}^*) = \sum_{\mathbf{a} \in A} \sum_{p \in \mathcal{P}_{\mathbf{a}}} s_{\mathbf{a},p}^* \text{Rev}(p). \quad (5)$$

Further, the revenue of Pro-LPR algorithm, i.e., $\mathbb{E}[\text{Rev}(\mathbf{r}^{\text{out}})]$, is lower bounded by

$$\mathbb{E}[\text{Rev}(\mathbf{r}^{\text{out}})] \geq \max \left(\sum_{\mathbf{a} \in A} \beta_{\mathbf{a}}^{(2)}, \mathbb{E}[\text{Rev}(\mathbf{r}^{\mathcal{R}})] \right). \quad (6)$$

To see why this holds note that Pro-LPR algorithm returns the best of reserve price $\mathbf{r}^{\mathcal{R}}$ and all zero prices, where the revenue under all zero prices is the sum of the second highest highest bids $\sum_{\mathbf{a} \in A} \beta_{\mathbf{a}}^{(2)}$. By using Equations (4), (5), and (6), we have

$$\begin{aligned} \text{ESP}^* - \mathbb{E}[\text{Rev}(\mathbf{r}^{\text{out}})] &\leq c \sum_{\mathbf{a} \in A} \beta_{\mathbf{a}}^{(2)} \\ &\leq c \mathbb{E}[\text{Rev}(\mathbf{r}^{\text{out}})]. \end{aligned} \quad (7)$$

Putting these together, we have

$$\mathbb{E}[\text{Rev}(\mathbf{r}^{\text{out}})] \geq \frac{1}{1+c} \cdot \text{ESP}^*,$$

which is the desired result. \square

In the next lemma, we show that the first condition holds. We then dedicate the next section to identifying constant c in the second condition. The proof of the lemma is based on the observations that (i) revenue of any valid profile $(\mathbf{b}_1, \mathbf{b}_2, r_1, r_2)$ is greater than $\beta_{\mathbf{a}}^{(2)}$ if buyer $\mathbf{b}_1 = \mathbf{b}_{\mathbf{a}}^{(1)}$ and his reserve $r_1 > \beta_{\mathbf{a}}^{(2)}$, and (ii) revenue of auction \mathbf{a} under reserve prices $\mathbf{r}^{\mathcal{R}}$ is greater than $\beta_{\mathbf{a}}^{(2)}$ if the reserve price of buyer $\mathbf{b}_{\mathbf{a}}^{(1)}$ is less than or equal to his bid and greater than the second highest bid $\beta_{\mathbf{a}}^{(2)}$. Here, buyer $\mathbf{b}_{\mathbf{a}}^{(1)}$ is the buyer with the highest bid in auction \mathbf{a} .

Lemma 5.2 (First Condition Holds). *Let \mathbf{s}^* denote an optimal solution of Profile-LP and $\mathbf{r}^{\mathcal{R}}$ be a random reserve price obtained from the rounding procedure in the Pro-LPR algorithm. For any auction $\mathbf{a} \in A$, we have*

$$\sum_{\{p: p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* - \Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t] \leq 0 \quad \text{for } t > \beta_{\mathbf{a}}^{(2)}. \quad (8)$$

Proof. The first term in the l.h.s. of (8) can be written as

$$\sum_{\{p: p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* = \sum_{\{p: p \in \mathcal{P}_{\mathbf{a}}, p = (\mathbf{b}_{\mathbf{a}}^{(1)}, \mathbf{b}_2, r, r_2), r \geq t\}} s_{\mathbf{a},p}^* \leq \sum_{r \geq t} q_{\mathbf{b}_{\mathbf{a}}^{(1)}, r} = \Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t], \quad (9)$$

where the first equation holds because revenue of a profile $p \in \mathcal{P}_{\mathbf{a}}$ is $t > \beta_{\mathbf{a}}^{(2)}$ if and only if the bidder with the highest bid in auction \mathbf{a} , i.e., $\mathbf{b}_{\mathbf{a}}^{(1)}$, is assigned a reserve price $t > \beta_{\mathbf{a}}^{(2)}$ and the bid of this bidder is greater than t . The second equation holds because of the second set of constraints of (Profile-LP). The last equation follows from the construction of reserve prices $\mathbf{r}^{\mathcal{R}}$. Note that Equation (9) verifies condition (8). \square

5.1 Bounding the Constant in the Second Condition

We start by noting that the second condition in Lemma 5.1 holds trivially for $c = 1$, which recovers the same approximation factor of $1/2$ of [RW16]. For the rest of the paper, we will improve past $1/2$ by constructing a non-linear mathematical program to optimize c and then applying the first order conditions in non-linear programming to bound the optimal solution. In Lemma 5.3, we show that

$$c = \max_{\theta \in [0,1]} \text{OPT}(\theta),$$

where for any real number $\theta \in [0, 1]$, $\text{OPT}(\theta)$ is defined as follows

$$\begin{aligned} \text{OPT}(\theta) &= \max_{\mathbf{x} \geq 0} e^{\theta-1} \left[\prod_{i \in [n]} (1 - x_i) + \sum_{i \in [n]} x_i \prod_{j \in [n], j \neq i} (1 - x_j) \right] \\ \text{s.t.} \quad &\frac{1}{2} \sum_{i \in [n]} x_i = \theta \\ &x_i \leq \theta, \quad \forall i \in [n]. \end{aligned} \tag{10}$$

Here, n is the number of buyers. Characterizing $\text{OPT}(\theta)$ is technically involved and because of that its details is postponed to Section 5.3. There, we show that for any number of buyers $n \geq 2$ and any real number $\theta \in [0, 1]$,

$$\text{OPT}(\theta) \leq 2 \left(\sqrt{2} - 1 \right) e^{\sqrt{2}-2} \approx 0.4612.$$

Then, invoking Lemmas 5.1 and 5.2, this leads to the approximation factor of $\frac{1}{1+0.4612} \approx 0.6844$, which is the desired result.

In the next lemma, we formally state the relationship between $\text{OPT}(\theta)$ and the approximation factor of our algorithm.

Lemma 5.3 (Second Condition). *Let \mathbf{s}^* denote an optimal solution of Profile-LP and $\mathbf{r}^{\mathcal{R}}$ be a random reserve price obtained from the rounding procedure in the Pro-LPR algorithm. Let*

$$c = \max_{\theta \in [0,1]} \text{OPT}(\theta).$$

Then, for any auction $\mathbf{a} \in A$, the following equation holds.

$$\sum_{\{p: p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* - \Pr [\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t] \leq c \quad \text{for } t \leq \beta_{\mathbf{a}}^{(2)}.$$

The formal proof of Lemma 5.3 due to being lengthy is deferred to Subsection 5.2, however we provide some intuition here. For each buyer \mathbf{b} , we consider two disjoint subsets of valid profiles $(\mathbf{b}_1, \mathbf{b}_2, r_1, r_2)$ such that (i) revenue of any profile in these subsets is greater than or equal to t , where $t \leq \beta_{\mathbf{a}}^{(2)}$, and (ii) either \mathbf{b}_1 or \mathbf{b}_2 is equal to buyer \mathbf{b} . The factor $1/2$ in the constraint of Problem (10) is the artifact of the definition of the subsets and how the summation $\sum_{\{p: p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^*$ can be written as a function of the optimal weight of the profiles in these subsets; see Equation (11) in the proof. We then express $\Pr [\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t]$ as the probability of the union of two events, where the first event happens if there is at least one cleared buyer with reserve price greater than t , and the second event happens if there are at least two buyers with cleared bids of at least t . We then

write this probability as a function of the profile weights in the subsets by taking advantage of the fact that in our rounding procedure, reserve prices are independent across buyers. In particular, we show that this probability is at least one minus the left hand side of the equation in Lemma 5.4, stated below. We invoke Lemma 5.4 to complete the proof. Observe that the objective function of Problem (10) bears significant resemblance to that of the right hand of the equation in Lemma 5.4.

Lemma 5.4. *Consider a set $\hat{\mathbf{B}} \subseteq \mathbf{B}$ with $|\hat{\mathbf{B}}| \geq 2$.² Given fixed $x_{1,b}, x_{2,b}$ with $b \in \hat{\mathbf{B}}$ and $x_{1,b} + x_{2,b} \leq 1$, the following inequality holds:*

$$\prod_{b \in \hat{\mathbf{B}}} (1 - x_{1,b} - x_{2,b}) + \sum_{b \in \hat{\mathbf{B}}} x_{2,b} \prod_{b' \neq b} (1 - x_{1,b'} - x_{2,b'}) \leq \prod_{b \in \hat{\mathbf{B}}} (1 - x_{1,b}) \left[\prod_{b \in \hat{\mathbf{B}}} (1 - x_{2,b}) + \sum_{b \in \hat{\mathbf{B}}} x_{2,b} \prod_{b' \neq b} (1 - x_{2,b'}) \right].$$

Proof. Given a partition of $\hat{\mathbf{B}}$ in two sets $\mathbf{B}_1, \mathbf{B}_2$, define the following function:

$$\begin{aligned} \Phi(\mathbf{B}_1, \mathbf{B}_2) = & \prod_{b \in \mathbf{B}_1} (1 - x_{1,b})(1 - x_{2,b}) \prod_{b \in \mathbf{B}_2} (1 - x_{1,b} - x_{2,b}) + \\ & \sum_{b \in \mathbf{B}_1 \cup \mathbf{B}_2} x_{2,b} \left[\prod_{b' \in \mathbf{B}_1, b' \neq b} (1 - x_{1,b})(1 - x_{2,b}) \prod_{b' \in \mathbf{B}_2, b' \neq b} (1 - x_{1,b} - x_{2,b}) \right]. \end{aligned}$$

The main claim in the lemma is that $\Phi(\mathbf{B}, \emptyset) \geq \Phi(\emptyset, \mathbf{B})$. We will show that for any $\mathbf{B}_1, \mathbf{B}_2$ and $\hat{b} \in \mathbf{B}_2$, we have

$$\Phi(\mathbf{B}_1, \mathbf{B}_2) \leq \Phi(\mathbf{B}_1 \cup \{\hat{b}\}, \mathbf{B}_2 \setminus \{\hat{b}\})$$

and the claim will follow by moving the elements from \mathbf{B}_2 to \mathbf{B}_1 one by one. To simplify notation, define

$$w = \prod_{b \in \mathbf{B}_1} (1 - x_{1,b})(1 - x_{2,b}) \prod_{b \in \mathbf{B}_2 \setminus \{\hat{b}\}} (1 - x_{1,b} - x_{2,b}).$$

Now we can write:

$$\Phi(\mathbf{B}_1, \mathbf{B}_2) = w \cdot (1 - x_{1,\hat{b}} - x_{2,\hat{b}}) + w \cdot x_{2,\hat{b}} + \sum_{b \in \mathbf{B}_2; b \neq \hat{b}} w \cdot \frac{1 - x_{1,\hat{b}} - x_{2,\hat{b}}}{1 - x_{1,b} - x_{2,b}} \cdot x_{2,b} + \sum_{b \in \mathbf{B}_1} w \cdot \frac{1 - x_{1,\hat{b}} - x_{2,\hat{b}}}{(1 - x_{1,b})(1 - x_{2,b})} \cdot x_{2,b}$$

and

$$\begin{aligned} \Phi(\mathbf{B}_1 \cup \{\hat{b}\}, \mathbf{B}_2 \setminus \{\hat{b}\}) = & w \cdot (1 - x_{1,\hat{b}})(1 - x_{2,\hat{b}}) + w \cdot x_{2,\hat{b}} \\ & + \sum_{b \in \mathbf{B}_2; b \neq \hat{b}} w \cdot \frac{(1 - x_{1,\hat{b}})(1 - x_{2,\hat{b}})}{1 - x_{1,b} - x_{2,b}} \cdot x_{2,b} + \sum_{b \in \mathbf{B}_1} w \cdot \frac{(1 - x_{1,\hat{b}})(1 - x_{2,\hat{b}})}{(1 - x_{1,b})(1 - x_{2,b})} \cdot x_{2,b} \end{aligned}$$

Our goal here is to show $\Phi(\mathbf{B}_1, \mathbf{B}_2) \leq \Phi(\mathbf{B}_1 \cup \{\hat{b}\}, \mathbf{B}_2 \setminus \{\hat{b}\})$. We start with comparing the first two terms of $\Phi(\mathbf{B}_1, \mathbf{B}_2)$ and $\Phi(\mathbf{B}_1 \cup \{\hat{b}\}, \mathbf{B}_2 \setminus \{\hat{b}\})$:

$$w \cdot (1 - x_{1,\hat{b}} - x_{2,\hat{b}}) + w \cdot x_{2,\hat{b}} = w \cdot (1 - x_{1,\hat{b}}) \leq w \cdot (1 - x_{1,\hat{b}} + x_{1,\hat{b}} x_{2,\hat{b}}) = w \cdot (1 - x_{1,\hat{b}})(1 - x_{2,\hat{b}}) + w \cdot x_{2,\hat{b}}.$$

We can compare the remaining terms one by one using the fact that:

$$1 - x_{1,\hat{b}} - x_{2,\hat{b}} \leq (1 - x_{1,\hat{b}})(1 - x_{2,\hat{b}}).$$

This concludes that $\Phi(\mathbf{B}_1, \mathbf{B}_2) \leq \Phi(\mathbf{B}_1 \cup \{\hat{b}\}, \mathbf{B}_2 \setminus \{\hat{b}\})$ as desired. \square

²Note that because of the auxiliary buyers, $|\mathbf{B}| \geq 2$.

5.2 Proof of Lemma 5.3

We start with a few definitions. Consider a certain auction $\mathbf{a} \in \mathbf{A}$ and all of its valid profiles $p \in \mathcal{P}_{\mathbf{a}}$. Fix some threshold $t \leq \beta_{\mathbf{a}}^{(2)}$ and an optimal solution of (Profile-LP), denoted by \mathbf{s}^* . Let set $\mathbf{B}_{\mathbf{a},t}$ be the set of buyers whose bid in auction \mathbf{a} is at least t :

$$\mathbf{B}_{\mathbf{a},t} := \{\mathbf{b} \in \mathbf{B}; \beta_{\mathbf{a},\mathbf{b}} \geq t\}.$$

Note that this set is not empty because $t \leq \beta_{\mathbf{a}}^{(2)}$. In fact, $|\mathbf{B}_{\mathbf{a},t}| \geq 2$, as buyers with the highest and second-highest bids belong to this set. (Recall that because of the auxiliary buyers, $|\mathbf{B}| \geq 2$.) A crucial observation is that the reserve assigned to any buyer $\mathbf{b} \notin \mathbf{B}_{\mathbf{a},t}$ does not affect the event $\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t$ since such buyers can be neither the winner nor the price setter in an auction with revenue of at least t . Consider a buyer $\mathbf{b} \in \mathbf{B}_{\mathbf{a},t}$. Then, define

$$\mathcal{X}'_{1,\mathbf{b}} = \{p = (\mathbf{b}, \mathbf{b}_2, r_1, r_2) : p \in \mathcal{P}_{\mathbf{a}}, r_1 \geq t\}$$

$$\mathcal{X}''_{1,\mathbf{b}} = \{p = (\mathbf{b}_1, \mathbf{b}, r_1, r_2) : p \in \mathcal{P}_{\mathbf{a}}, r_1 < t \text{ and } r_2 \geq t\}$$

$$\mathcal{X}_{2,\mathbf{b}} = \{p = (\mathbf{b}_1, \mathbf{b}_2, r_1, r_2) : p \in \mathcal{P}_{\mathbf{a}}, \mathbf{b} \in \{\mathbf{b}_1, \mathbf{b}_2\}, r_1, r_2 < t \text{ and } \beta_{\mathbf{a},\mathbf{b}_2} \geq t\}$$

and then set

$$x_{1,\mathbf{b}} = \sum_{p \in \mathcal{X}'_{1,\mathbf{b}} \cup \mathcal{X}''_{1,\mathbf{b}}} s_{\mathbf{a},p}^* \quad \text{and} \quad x_{2,\mathbf{b}} = \sum_{p \in \mathcal{X}_{2,\mathbf{b}}} s_{\mathbf{a},p}^*.$$

We note that $\mathcal{X}'_{1,\mathbf{b}}$ is the set of all valid profiles $p = (\mathbf{b}, \mathbf{b}_2, r_1, r_2)$ in which reserve of buyer \mathbf{b} is at least t . $\mathcal{X}''_{1,\mathbf{b}}$ is the set of all valid profiles $p = (\mathbf{b}_1, \mathbf{b}, r_1, r_2)$ in which reserve of buyer \mathbf{b}_1 is less than t and reserve of buyer \mathbf{b} is greater than or equal to t . Observe that for all the profiles p in $\mathcal{X}'_{1,\mathbf{b}} \cup \mathcal{X}''_{1,\mathbf{b}}$, reserve of buyer \mathbf{b} is at least t . This implies that for all of these profiles, $\text{Rev}(p) \geq t$. We also note that $\mathcal{X}_{2,\mathbf{b}}$ is the set of all valid profiles $p = (\mathbf{b}_1, \mathbf{b}_2, r_1, r_2)$ such that buyer $\mathbf{b} \in \{\mathbf{b}_1, \mathbf{b}_2\}$ and bid of buyer \mathbf{b}_2 is at least t . Again, it is easy to see that for any valid profile $p \in \mathcal{X}_{2,\mathbf{b}}$, we have $\text{Rev}(p) \geq t$. Finally, we point that while any profile p in $\mathcal{X}_{2,\mathbf{b}}$ and $\mathcal{X}'_{1,\mathbf{b}} \cup \mathcal{X}''_{1,\mathbf{b}}$ has $\text{Rev}(p) \geq t$, by construction, $\mathcal{X}_{2,\mathbf{b}}$ and $\mathcal{X}'_{1,\mathbf{b}} \cup \mathcal{X}''_{1,\mathbf{b}}$ are disjoint. Therefore, we have

$$\sum_{\{p: p \in \mathcal{P}_{\mathbf{a}}, \text{Rev}(p) \geq t\}} s_{\mathbf{a},p}^* = \sum_{\mathbf{b} \in \mathbf{B}_{\mathbf{a},t}} x_{1,\mathbf{b}} + \frac{1}{2} \sum_{\mathbf{b} \in \mathbf{B}_{\mathbf{a},t}} x_{2,\mathbf{b}}, \quad (11)$$

where the coefficient $\frac{1}{2}$ accounts for double-counting. Define $y_{1,\mathbf{b}}$ as the probability that the sampled reserve of buyer \mathbf{b} , i.e., $\mathbf{r}_{\mathbf{b}}^{\mathcal{R}}$, is in $[t, \beta_{\mathbf{a},\mathbf{b}}]$ and $y_{2,\mathbf{b}}$ as the probability that the sampled reserve $\mathbf{r}_{\mathbf{b}}^{\mathcal{R}}$ is in $[0, t)$. By the sampling procedure we know that:

$$y_{1,\mathbf{b}} \geq x_{1,\mathbf{b}} \quad \text{and} \quad y_{2,\mathbf{b}} \geq x_{2,\mathbf{b}}. \quad (12)$$

Observe that $\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t$ iff at least one of the two following events happen.

Event \mathcal{E}_1 : There exists a buyer $\mathbf{b} \in \mathbf{B}_{\mathbf{a},t}$ with a reserve of at least t whose bid is cleared.

Event \mathcal{E}_2 : There are at least two buyers $\mathbf{b}_1, \mathbf{b}_2 \in \mathbf{B}_{\mathbf{a},t}$ with cleared bids of at least t .

Precisely,

$$\Pr[\text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \geq t] = \Pr[\mathcal{E}_1 \text{ or } \mathcal{E}_2] = \Pr[\mathcal{E}_1] + \Pr[\mathcal{E}_2 \text{ and } \bar{\mathcal{E}}_1] = \Pr[\mathcal{E}_1] + \Pr[\bar{\mathcal{E}}_1] \Pr[\mathcal{E}_2 | \bar{\mathcal{E}}_1], \quad (13)$$

where

$$\Pr[\mathcal{E}_1] = 1 - \prod_{b \in \mathbf{B}_{a,t}} (1 - y_{1,b}),$$

and

$$\Pr[\mathcal{E}_2 | \bar{\mathcal{E}}_1] = 1 - \prod_{b \in \mathbf{B}_{a,t}} (1 - \tilde{y}_{2,b}) - \sum_{b \in \mathbf{B}_{a,t}} \tilde{y}_{2,b} \prod_{b' \neq b} (1 - \tilde{y}_{2,b'}) \quad \text{with} \quad \tilde{y}_{2,b} = \frac{y_{2,b}}{1 - y_{1,b}}.$$

This gives us

$$\Pr[\mathcal{E}_2 \text{ and } \bar{\mathcal{E}}_1] = \Pr[\bar{\mathcal{E}}_1] \Pr[\mathcal{E}_2 | \bar{\mathcal{E}}_1] = \Pr[\bar{\mathcal{E}}_1] - \prod_{b \in \mathbf{B}_{a,t}} (1 - y_{1,b} - y_{2,b}) - \sum_{b \in \mathbf{B}_{a,t}} y_{2,b} \prod_{b' \neq b} (1 - y_{1,b'} - y_{2,b'}).$$

Thus, by Equation (13), we get

$$\Pr[\mathcal{E}_2 \text{ or } \mathcal{E}_1] = 1 - \prod_{b \in \mathbf{B}_{a,t}} (1 - y_{1,b} - y_{2,b}) - \sum_{b \in \mathbf{B}_{a,t}} y_{2,b} \prod_{b' \neq b} (1 - y_{1,b'} - y_{2,b'}).$$

Now observe that the expression above, i.e., $\Pr[\mathcal{E}_2 \text{ or } \mathcal{E}_1]$, is increasing in both $y_{1,b}$ and $y_{2,b}$, $b \in \mathbf{B}_{a,t}$. To see why $\Pr[\mathcal{E}_2 \text{ or } \mathcal{E}_1]$ is increasing in $y_{2,b}$, note that

$$\frac{\partial(\Pr[\mathcal{E}_2 \text{ or } \mathcal{E}_1])}{\partial y_{2,b}} = \sum_{b' \in \mathbf{B}_{a,t}, b' \neq b} y_{2,b'} \prod_{b'' \neq b, b'} (1 - y_{1,b''} - y_{2,b''}) \geq 0.$$

This and Equation (12) imply that:

$$\Pr[\mathcal{E}_2 \text{ or } \mathcal{E}_1] \geq 1 - \prod_{b \in \mathbf{B}_{a,t}} (1 - x_{1,b} - x_{2,b}) - \sum_{b \in \mathbf{B}_{a,t}} x_{2,b} \prod_{b' \neq b} (1 - x_{1,b'} - x_{2,b'}).$$

We now invoke Lemma 5.4, stated earlier, to get

$$\Pr[\mathcal{E}_2 \text{ or } \mathcal{E}_1] \geq 1 - \prod_{b \in \mathbf{B}_{a,t}} (1 - x_{1,b}) \left[\prod_{b \in \mathbf{B}_{a,t}} (1 - x_{2,b}) + \sum_{b \in \mathbf{B}_{a,t}} x_{2,b} \prod_{b' \neq b} (1 - x_{2,b'}) \right]. \quad (14)$$

Using Equations (11), (14), and (13), we have

$$\begin{aligned} & \sum_{\{p: p \in \mathcal{P}_a, \text{Rev}(p) \geq t\}} s_{a,p}^* - \Pr[\text{Rev}_a(\mathbf{r}^R) \geq t] \leq \\ & \sum_{b \in \mathbf{B}_{a,t}} x_{1,b} + \frac{1}{2} \sum_{b \in \mathbf{B}_{a,t}} x_{2,b} - \left(1 - \prod_{b \in \mathbf{B}_{a,t}} (1 - x_{1,b}) \left[\prod_{b \in \mathbf{B}_{a,t}} (1 - x_{2,b}) + \sum_{b \in \mathbf{B}_{a,t}} x_{2,b} \prod_{b' \neq b} (1 - x_{2,b'}) \right] \right). \end{aligned} \quad (15)$$

We claim that for any $\mathbf{b} \in \mathbf{B}_{a,t}$, the above expression is non-decreasing in $x_{1,b}$. To get this, we need to show that the derivative of the above expression w.r.t. $x_{1,\hat{b}}$ is non-negative. In the other words, we need to show the following equation holds:

$$1 - \prod_{b \in \mathbf{B}_{a,t}, b \neq \hat{b}} (1 - x_{1,b}) \left[\prod_{b \in \mathbf{B}_{a,t}} (1 - x_{2,b}) + \sum_{b \in \mathbf{B}_{a,t}} x_{2,b} \prod_{b' \neq b} (1 - x_{2,b'}) \right] \geq 0. \quad (16)$$

Since $0 \leq (1 - x_{1,\mathbf{b}}) \leq 1$ for any \mathbf{b} , it only remains to show that the value of the term in the brackets, i.e., $\prod_{\mathbf{b} \in \mathbf{B}_{a,t}} (1 - x_{2,\mathbf{b}}) + \sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{2,\mathbf{b}} \prod_{\mathbf{b}' \neq \mathbf{b}} (1 - x_{2,\mathbf{b}'})$, is always in the range of $[0, 1]$. To get this, it suffices to show that there exists an event whose probability can be written as $\prod_{\mathbf{b} \in \mathbf{B}_{a,t}} (1 - x_{2,\mathbf{b}}) + \sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{2,\mathbf{b}} \prod_{\mathbf{b}' \neq \mathbf{b}} (1 - x_{2,\mathbf{b}'})$. For any \mathbf{b} define a Bernoulli random variable with mean $x_{2,\mathbf{b}}$. Observe that the aforementioned term is equal to the probability of the event in which at most one of these variables is equal to one, assuming that they are independent. Thus, we obtain Equation (16), which allows us to assume without loss of generality that $\sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{1,\mathbf{b}} + \frac{1}{2} \sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{2,\mathbf{b}} = 1$. As a result, we have

$$\sum_{\{p: p \in \mathcal{P}_a, \text{Rev}(p) \geq t\}} s_{a,p}^* - \Pr[\text{Rev}_a(\mathbf{r}^R) \geq t] \leq \prod_{\mathbf{b} \in \mathbf{B}_{a,t}} (1 - x_{1,\mathbf{b}}) \left[\prod_{\mathbf{b} \in \mathbf{B}_{a,t}} (1 - x_{2,\mathbf{b}}) + \sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{2,\mathbf{b}} \prod_{\mathbf{b}' \neq \mathbf{b}} (1 - x_{2,\mathbf{b}'}) \right],$$

where $\sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{2,\mathbf{b}} = 2\theta$, $\sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{1,\mathbf{b}} = 1 - \theta$. Here, $\theta \in [0, 1]$. To complete the proof, we simply use that: $\prod_{\mathbf{b} \in \mathbf{B}_{a,t}} (1 - x_{1,\mathbf{b}}) \leq e^{-\sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{1,\mathbf{b}}} = e^{\theta-1}$. Given how we constructed the variables $x_{2,\mathbf{b}}$, we also need $x_{2,\mathbf{b}} \leq \theta$. Hence,

$$\sum_{\{p: p \in \mathcal{P}_a, \text{Rev}(p) \geq t\}} s_{a,p}^* - \Pr[\text{Rev}_a(\mathbf{r}^R) \geq t] \leq e^{\theta-1} \left[\prod_{\mathbf{b} \in \mathbf{B}_{a,t}} (1 - x_{2,\mathbf{b}}) + \sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{2,\mathbf{b}} \prod_{\mathbf{b}' \neq \mathbf{b}} (1 - x_{2,\mathbf{b}'}) \right],$$

where $\sum_{\mathbf{b} \in \mathbf{B}_{a,t}} x_{2,\mathbf{b}} = 2\theta$ and $x_{2,\mathbf{b}} \leq \theta$ for any $\mathbf{b} \in \mathbf{B}_{a,t}$.

5.3 Approximation Factor

In this section, we will show that for any given $\theta \in [0, 1]$, we have

$$\text{OPT}(\theta) \leq 2 \left(\sqrt{2} - 1 \right) e^{\sqrt{2}-2},$$

where $\text{OPT}(\theta)$ is defined in Equation (10). Since the constraints of Program (10) are linear in x_i 's, the first order conditions of Karush-Kuhn-Tucker (KKT) are a necessary condition for optimality [Ber99]. Let

$$F(x, \theta) = e^{\theta-1} \left[\prod_{i \in [n]} (1 - x_i) + \sum_{i \in [n]} x_i \prod_{j \in [n], j \neq i} (1 - x_j) \right].$$

Observe that $F(x, \theta)$ is the objective function of $\text{OPT}(\theta)$. Then, according to the KKT conditions, the optimal solution must satisfy the following constraints for some $\lambda \in \mathbb{R}$, $\mu, \eta \in \mathbb{R}_+^n$:

$$\nabla_x F(x, \theta) + \frac{\lambda}{2} \mathbf{1} - \mu + \eta = 0 \tag{17}$$

$$\sum_{i \in [n]} x_i = \frac{1}{2} \theta \tag{18}$$

$$\mu_i (x_i - \theta) = 0, \quad \forall i \in [n] \tag{19}$$

$$\eta_i x_i = 0, \quad \forall i \in [n] \tag{20}$$

$$0 \leq x_i \leq \theta, \quad \forall i \in [n]. \tag{21}$$

where $\mathbf{1} \in \mathbb{R}^n$ is the vector of all one.

It is enough to show that $F(x, \theta) \leq 2 \left(\sqrt{2} - 1 \right) e^{\sqrt{2}-2}$ for any tuple $(x, \theta, \lambda, \mu, \eta)$ satisfying the KKT conditions. A simple consequence of the KKT condition is the following:

Lemma 5.5 (KKT Condition). *If $(x, \theta, \lambda, \mu, \eta)$ satisfies the KKT conditions for Problem (10), then if x_k and x_t are such that $0 < x_k < \theta$ and $0 < x_t < \theta$, then $x_k = x_t$.*

Proof. By conditions (19) and (20), we must have $\mu_k = \eta_k = 0$. Plugging that into condition (17), we get that

$$\partial F / \partial x_k + \lambda / 2 = 0.$$

This implies that

$$\sum_{i \neq k} x_i \prod_{j \neq i, k} (1 - x_j) + \frac{\lambda}{2} = 0.$$

Let $Q = \prod_{i \in [n]} (1 - x_i)$ and $S = \sum_{i \in [n]} \frac{x_i}{1 - x_i}$. Then, the above condition can be written as

$$\frac{Q}{1 - x_k} \sum_{i \neq k} \frac{x_i}{1 - x_i} + \frac{\lambda}{2} = 0 \quad \Rightarrow \quad \frac{Q}{1 - x_k} \left(S - \frac{x_k}{1 - x_k} \right) + \frac{\lambda}{2} = 0.$$

This is further simplified as follows

$$\left(SQ + \frac{\lambda}{2} \right) - (SQ + Q + \lambda)x_k + \frac{\lambda}{2}x_k^2 = 0.$$

The polynomial $p(y) := (SQ + \frac{\lambda}{2}) - (SQ + Q + \lambda)y + \frac{\lambda}{2}y^2$ is quadratic with $\frac{d^2p}{dy^2} \geq 0$ and $p(1) = -Q < 0$. Thus, $p(y) = 0$ has a unique solution with $y < 1$. This implies x_k is uniquely determined as a function of S , Q , and λ . By the same argument, x_t is also a solution to the same equation and hence $x_k = x_t$. \square

Lemma 5.5 leads to the following corollary.

Corollary 5.6. *We can bound $\text{OPT}(\theta) \leq \max_{k \in \mathbb{Z}, k \geq 2} \max[\text{OPT}^1(\theta, k), \text{OPT}^2(\theta, k)]$, where*

$$\begin{aligned} \text{OPT}^1(\theta, k) &= e^{\theta-1} \left(1 - \frac{2\theta}{k} \right)^{k-1} \left(1 - \frac{2\theta}{k} + 2\theta \right) \\ \text{OPT}^2(\theta, k) &= e^{\theta-1} \left[\left(1 - \frac{\theta}{k} \right)^k + \theta(1 - \theta) \left(1 - \frac{\theta}{k} \right)^{k-1} \right]. \end{aligned}$$

Proof. As stated earlier, in order to maximize the objective function $\text{OPT}(\theta)$, it is enough to consider feasible solutions x satisfying the KKT conditions. To do so, we use Lemma 5.5 to narrow down such solutions.

Since for any $i \in [n]$, $x_i \leq \theta$ and $\sum_{i \in [n]} x_i = 2\theta$, we can only have the following three cases:

- Case 1: Two variables in the support have value θ and by constraint $\sum_{i \in [n]} x_i = 2\theta$, the rest of them are zero. In that case, $\text{OPT}(\theta) = \text{OPT}^1(\theta, 2)$.
- Case 2: One variable has value θ and by Lemma 5.5, the rest $n - 1 \geq 2$ variables in the support have value $\theta/(n - 1)$. In that case, $\text{OPT}(\theta) = \text{OPT}^2(\theta, n - 1)$.
- Case 3: All variables in the support are strictly below θ . In this case, by Lemma 5.5, $x_i = \theta/n$ for $n \geq 3$, and the solution is $\text{OPT}(\theta) = \text{OPT}^1(\theta, n)$.

□

Lemma 5.7. For any $\theta \in [0, 1]$ and $k \geq 2$, we have $\text{OPT}^1(\theta, k) \leq 2(\sqrt{2} - 1)e^{\sqrt{2}-2}$.

Proof. For each $k \geq 0$, define $\theta^*(k) = \arg \max_{\theta \in [0, 1]} \text{OPT}^1(\theta, k)$. By solving $\partial \text{OPT}^1(\theta, k) / \partial \theta = 0$ we obtain the following expression for $\theta^*(k)$:

$$k^2(2\theta^*(k) - 1) + 4(k - 1)(\theta^*(k))^2 = 0.$$

The aforementioned equation has two solutions, only one of which is in $[0, 1]$. Thus,

$$\theta^*(k) = \frac{k \left(k - \sqrt{k^2 + 4k - 4} \right)}{4 - 4k}. \quad (22)$$

We need to show that for any $k \geq 2$, we have $\text{OPT}^1(\theta^*(k), k) \leq 2(\sqrt{2} - 1)e^{\sqrt{2}-2} \approx 0.461$. For $k = 2$, we have $\text{OPT}^1(\theta^*(k), k) = 2(\sqrt{2} - 1)e^{\sqrt{2}-2}$. For $k < 40$, we can verify this inequality numerically. For $k \geq 40$, we define an upper bound:

$$U(\theta, k) = \frac{2\theta + 1}{e^{\theta+1} \left(1 - \frac{2\theta}{k}\right)}.$$

and show that for any $\theta \in [0, 1]$ and $k \geq 40$,

$$\text{OPT}^1(\theta, k) \leq U(\theta, k) \leq U(\theta, 40) \leq 0.459 < 2(\sqrt{2} - 1)e^{\sqrt{2}-2}.$$

For the first inequality note that:

$$\text{OPT}^1(\theta, k) = e^{\theta-1} \left[\left(1 - \frac{2\theta}{k}\right)^{k-1} \left(1 + (k-1)\frac{2\theta}{k}\right) \right] \quad (23)$$

$$< e^{\theta-1} \left[\left(1 - \frac{2\theta}{k}\right)^k \left(1 - \frac{2\theta}{k}\right)^{-1} (1 + 2\theta) \right] \leq U(\theta, k). \quad (24)$$

For the second inequality, we use the fact that for any θ , $U(\theta, k)$ is decreasing in k . To find an upper-bound for value of $U(\theta, 40) = \frac{(2\theta+1)}{e^{\theta+1} \left(1 - \frac{\theta}{20}\right)}$, we take derivative of that which is

$$\frac{\partial U(\theta, 40)}{\partial \theta} = \frac{20(2\theta^2 - 39\theta + 21)}{e^{\theta+1}(\theta - 20)^2}.$$

By solving $\frac{\partial U(\theta, 40)}{\partial \theta} = 0$, we obtain that maximum of $U(\theta, 40)$ is at $\theta = \frac{1}{4}(39 - \sqrt{1353})$ and

$$U\left(\frac{1}{4}(39 - \sqrt{1353}), 40\right) < 0.459.$$

This completes the proof. □

Lemma 5.8. For any $\theta \in [0, 1]$ and $k \geq 2$, we have $\text{OPT}^2(\theta, k) \leq 0.46 < 2(\sqrt{2} - 1)e^{\sqrt{2}-2}$.

Proof. Observe that

$$e^{1-\theta} \text{OPT}^2(\theta, k) = \left(1 - \frac{\theta}{k}\right)^k + \theta(1 - \theta) \left(1 - \frac{\theta}{k}\right)^{k-1} \quad (25)$$

$$\leq \left(1 - \frac{\theta}{k}\right)^k + \frac{1}{4} \left(1 - \frac{\theta}{k}\right)^k = \frac{5}{4} \left(1 - \frac{\theta}{k}\right)^k, \quad (26)$$

where the first inequality holds because $\max_{\theta \in [0,1]} \theta(1 - \theta) = \frac{1}{4}$ and $1 - \frac{\theta}{k} \leq 1$. Finally, note that $e^{\theta-1} \cdot \frac{5}{4} \left(1 - \frac{\theta}{k}\right)^k$ is decreasing for $\theta \in [0, 1]$. Thus, we can bound $\text{OPT}^2(\theta, k)$ by the value of $e^{\theta-1} \cdot \frac{5}{4} \left(1 - \frac{\theta}{k}\right)^k$ at $\theta = 0$ which is $5/(4e) < 0.46$. \square

6 Tightness of the Analysis

In this section, we show that the analysis of our algorithm is tight, i.e., we construct an example for which the performance of the algorithm matches the 0.684 approximation factor.

To make the construction cleaner, we can define the weighted version of our problem in which each auction $\mathbf{a} \in \mathbf{A}$ has an associated weight $w_{\mathbf{a}} > 0$, and the objective is to maximize $\sum_{\mathbf{a} \in \mathbf{A}} w_{\mathbf{a}} \cdot \text{Rev}_{\mathbf{a}}(\mathbf{r})$. Note that if the weights are integers, this is exactly the same as the original problem, replacing each weighted auction by $w_{\mathbf{a}}$ (unweighted) copies. Even if $w_{\mathbf{a}}$'s are not integers, it is easy to see that the algorithm and the analysis generalize with essentially no change to the weighted case (the only modification involves adding weights to the objective function in the LP). In other words, if the objective were the weighted revenue, we would still get 0.684 approximation factor by applying a similar algorithm. Furthermore, any lower bound to the weighted case translates to the unweighted case by replacing a weighted auction \mathbf{a} by $\lfloor Nw_{\mathbf{a}} \rfloor$ unweighted copies for some large N .

Theorem 3 (Tightness of the Analysis). *There is a weighted instance $\{w_{\mathbf{a}}\}_{\mathbf{a} \in \mathbf{A}}, \{\beta_{\mathbf{a}, \mathbf{b}}\}_{\mathbf{a} \in \mathbf{A}, \mathbf{b} \in \mathbf{B}}$ and an optimal LP solution \mathbf{s}, \mathbf{q} such that*

$$\max \left(\mathbb{E} \left[\sum_{\mathbf{a}} w_{\mathbf{a}} \text{Rev}_{\mathbf{a}}(\mathbf{r}^{\mathcal{R}}) \right], \sum_{\mathbf{a}} w_{\mathbf{a}} \text{Rev}_{\mathbf{a}}(\mathbf{0}) \right) \leq 0.684 \cdot \text{Rev}(\mathbf{s}).$$

Proof. Fix $\theta = \sqrt{2} - 1$ and $c = (1 - \theta^2)e^{\theta-1}$. Consider an instance with three weighted auctions and $n = k + 3$ buyers described by the following table:

Weights $w_{\mathbf{a}}$ \diagdown Bids	$\beta_{\mathbf{a},1}$	\dots	$\beta_{\mathbf{a},k}$	$\beta_{\mathbf{a},k+1}$	$\beta_{\mathbf{a},k+2}$	$\beta_{\mathbf{a},k+3}$
$1/(c+1)$	1	\dots	1	1	1	0
$c/(c+1)$	0	\dots	0	0	0	$1 + \epsilon$
ϵ	$1 + \epsilon$	\dots	$1 + \epsilon$	$1 + \epsilon$	$1 + \epsilon$	0

Now, consider the following solution to the Profile-LP. For the first auction,

- the profile $p = (i, \mathbf{b}_0, 1, 0)$ has $s_{\mathbf{a},p} = (1 - \theta)/k$ for $i \in [k]$. In this profile, the i -th buyer is reserve priced at 1 and the second buyer is the dummy buyer;

- the profile $p = (k + 1, k + 2, 0, 0)$ has weight $s_{a,p} = \theta$. In this profile, both buyers $k + 1$ and $k + 2$ have zero reserve prices. Observe that the revenue under this profile is 1 due to the highest second price.

For the second auction, we consider only one profile:

- the profile $p = (k + 3, \mathbf{b}_0, 1 + \epsilon, 0)$ has $s_{a,p} = 1$. In this profile, the $(k + 3)$ -th buyer is reserve priced at 1 and the second buyer is the dummy buyer.

And for the third auction, we have:

- the profile $p = (i, \mathbf{b}_0, 1 + \epsilon, 0)$ has $s_{a,p} = \theta/k$ for $i \in [k]$. In this profile, the i -th buyer is reserve priced at $1 + \epsilon$ and the second buyer is the dummy buyer.
- the profile $p = (k + 1, k + 2, 1 + \epsilon, 1 + \epsilon)$ has weight $s_{a,p} = 1 - \theta$. In this profile, both buyers $k + 1$ and $k + 2$ have reserve price $1 + \epsilon$ and thus the revenue is $1 + \epsilon$.

For this solution, we define the q variables as follows.

- For buyers $i \in [k]$, we set $q_{i,1} = (1 - \theta)/k$ and $q_{i,1+\epsilon} = 1 - q_{i,1}$.
- For buyers $i = k + 1, k + 2$, we set $q_{i,0} = \theta$ and $q_{i,1+\epsilon} = 1 - q_{i,1}$.
- For buyer $k + 3$, we set $q_{k+3,1+\epsilon} = 1$.

It is easy to see that this solution is feasible and that it is the optimal solution to Problem (Profile-LP). This is so because for any auction, any profile that has a positive weight yield the maximum revenue for that auction. Note that for simplicity in the formulation of revenue, we can remove the terms that are a factor of ϵ since they can be arbitrary small and are negligible. We argue that the rounding procedure produces a $1/(c + 1)$ approximation. First notice that the vector of zero reserves obtains revenue of $1/(c + 1)$.

Now, we compute the expected revenue from rounding. After rounding, the reserve of any buyer $i \in [k]$ is either 1 or $1 + \epsilon$, the reserve of buyers $k + 1$ and $k + 2$ is either zero or $1 + \epsilon$, and reserve of buyer $k + 3$ is always $1 + \epsilon$. Thus, by letting ϵ go to zero, the expected revenue from rounding is given by

$$\frac{1}{c + 1} \left[1 - \left(1 - \frac{1 - \theta}{k} \right)^k \cdot (1 - \theta^2) \right] + \frac{c}{c + 1},$$

where the first term is the revenue of first auction and the second term, i.e., $\frac{c}{c+1}$, is the revenue of the second auction.³ To see why the latter holds note that in the first auction, we always get a revenue of one unless none of the first k buyers have a reserve of one and neither buyers $k + 1$ nor buyer $k + 2$ have a reserve of zero. As $k \rightarrow \infty$, the expected revenue after rounding becomes:

$$\frac{1}{c + 1} \left[1 - e^{\theta-1} \cdot (1 - \theta^2) \right] + \frac{c}{c + 1} = \frac{1 - c}{c + 1} + \frac{c}{c + 1} = \frac{1}{c + 1},$$

where the first equation holds because $c = (1 - \theta^2)e^{\theta-1}$. The above equation is the desired result because the optimal revenue is at most 1 and $1/(c + 1) = 0.684$. The latter follows from $c = (1 - \theta^2)e^{\theta-1}$ and $\theta = \sqrt{2}$. \square

³We do not include the revenue of the third auction because we would like to take ϵ to zero and in that case, the revenue of the third auction approaches zero.

7 Integrality Gap

In this section, we give an upper-bound of 0.828 for the integrality gap of the LP. This implies that any rounding procedure for our LP formulation will obtain at most 0.828 fraction of the optimal value of the LP. In particular, we show Theorem 2, which we restate here for convenience.

Theorem 2 (Integrality Gap of Profile-LP). *There exists a data-set of bids $\{\beta_{a,b}\}_{a \in A, b \in B}$ for which the integrality gap of Profile-LP is at least $2(\sqrt{2} - 1) \approx 0.828$. That is,*

$$\text{ESP}^* \leq 2(\sqrt{2} - 1) \cdot \text{LP}^*,$$

where LP^* is the optimal fraction solution of the Profile-LP and ESP^* is its optimal integral solution.

Proof. Given n buyers, an integer $k > 0$, $\delta = 1/k$ and a constant $\lambda \in (0, 1)$ to be determined later, consider an instance built as follows:

- **Type one Auctions:** For any buyer, $\mathbf{b} \in [n]$, we have an auction in which all the bids are zero except the bid of buyer \mathbf{b} . Precisely, buyer \mathbf{b} has a bid of λn .
- **Type two Auctions:** For any pair of buyers \mathbf{b}_1 and \mathbf{b}_2 , there are k copies of an auction in which \mathbf{b}_1 and \mathbf{b}_2 bid $\delta = 1/k$ and the rest of the buyers bid 0. We assume that $\lambda n > \delta$.

For this instance, consider the fractional solution that assigns $s_{a,p} = 1/2$ for any auction a of type two and profiles $(\mathbf{b}_1, \mathbf{b}_0, \delta, 0)$ and $(\mathbf{b}_2, \mathbf{b}_0, \delta, 0)$. For the rest of the valid profiles of auction \mathbf{a} , we set $s_{a,p}$ to zero. Note that \mathbf{b}_1 and \mathbf{b}_2 are the buyers with nonzero bids in auction \mathbf{a} and \mathbf{b}_0 is a dummy buyer. Moreover, for any auction \mathbf{a} of type one, in which buyer \mathbf{b} has a nonzero bid, we have $s_{a,p} = 1/2$ for profile $p = (\mathbf{b}, \mathbf{b}_0, \lambda n, 0)$. For the rest of the valid profiles of this auction, we set $s_{a,p}$ to zero. In this solution for any buyer \mathbf{b} , we have $q_{\mathbf{b},\delta} = 1/2$ and $q_{\mathbf{b},\lambda n} = 1/2$. One can simply verify that this solution satisfies all the constraints of the LP and as a result, it is a valid fractional solution. The optimal value of the LP is therefore bounded by:

$$\text{LP}^* \geq \sum_{\mathbf{a}} \text{Rev}_{\mathbf{a}}(\mathbf{s}) = n \cdot \frac{\lambda}{2} n + \binom{n}{2} \cdot k \cdot \delta = \frac{1 + \lambda}{2} \cdot n^2 + o(n^2),$$

where the first term corresponds to the revenue from auctions of type one and the second term corresponds to the revenue of auctions of type two. To bound ESP^* , we note that in the optimal solution of Problem (ESP-OPT), the reserve of each buyer is either δ or λn . Given that the buyers are symmetric, the value of the optimal solution depends only on the number of buyers with each reserve. Let t be the number of buyers with reserve λn . Then, we can write:

$$\text{ESP}^* = \max_{0 \leq t \leq n} \left[t \cdot \lambda n + (n - t) \cdot \delta + \binom{n}{2} - \binom{t}{2} \right].$$

By taking $\delta \rightarrow 0$, we obtain,

$$\text{ESP}^* = \max_{0 \leq t \leq n} \left[t \cdot \lambda n + \binom{n}{2} - \binom{t}{2} \right].$$

Since the term inside the maximum is a quadratic function of t , the optimal integral solution should be $t = \lambda n + o(n)$. This is so because the optimal integral solution t deviates from the

optimal fractional solution (which is $\lambda n + 1/2$) by at most 1. Substituting that in the expression of ESP^* , we get

$$\text{ESP}^* = \frac{1 + \lambda^2}{2} \cdot n^2 + o(n^2).$$

Taking $n \rightarrow \infty$, we get

$$\frac{\text{ESP}^*}{\text{LP}^*} \leq \frac{(1 + \lambda^2)n^2 + o(n^2)}{(1 + \lambda)n^2 + o(n^2)} \rightarrow \frac{1 + \lambda^2}{1 + \lambda}.$$

We can choose the parameter $\lambda = \sqrt{2} - 1$ to minimize the above expression, which leads to a ratio of $2(\sqrt{2} - 1) \approx 0.828$. \square

8 Numerical Studies

In this section, we evaluate our Pro-LPR algorithm on synthetic data-sets. As a benchmark, we use the optimal value of Problem (Profile-LP) and the greedy algorithm proposed by [RW16]. (Recall that optimal value of Problem (Profile-LP) provides an upper bound on the revenue obtained by any vector of reserve prices.) We assess the performance of our algorithm in three settings, where in the first (respectively third) setting, the bids are negatively (respectively positively) correlated across buyers. In the second setting, bids are independent of each other.

Synthetic Data-sets: We assume that the number of buyers $n = 2$. By setting $n = 2$, we provide a fair comparison between Pro-LP and greedy algorithms as with $n = 2$, the Pro-LP algorithm, similar to the greedy algorithm, only uses the highest and second-highest bids in each auction. The bids in each auction \mathbf{a} are drawn from log-normal distributions. Note that log-normal distributions are proved to be a good fit for the advertisers' valuations in online advertising markets; see, for example, [ES10, EOS07, XYL09], and [BFMM14]. Let $\sigma = 0.1$ and define matrix Σ as follows

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & w \\ w & 1 \end{pmatrix}.$$

To generate bids in each auction \mathbf{a} , we first generate normal random variables, denoted by Z , with mean $(0, \mu)$ and variance Σ . The bids are then obtained by computing e^Z . Here, we focus on three values of w , $-0.2, 0$, and 0.2 . When $w = 0$, bids are independent of each other and when $w > 0$ (respectively $w < 0$), bids are positively (respectively negatively) correlated. More precisely, the correlation coefficient between bids of the two buyers is equal to $\frac{\exp(w\sigma^2) - 1}{\exp(\sigma^2) - 1}$. We now comment on the parameter μ . We draw μ randomly from a uniform distribution in the range of $[0, 1]$, where each value of μ represents one problem instance. We consider 50 problem instances for each value of $w \in \{-0.2, 0, 0.2\}$. Our generated data-set consists of 100 auctions for each value of μ and w . (Note that bids are independent across auctions.) After generating the data-set, we consider the 30 equally spaced values between zero and the maximum bid in the data-set as candidate reserve prices. We use the same set of candidate reserve prices in Problem (Profile-LP) and the greedy algorithm of [RW16].

Evaluation: We evaluate the revenue of Pro-LPR and greedy algorithms in two ways. First, we compare their revenue to the optimal value of Problem (Profile-LP) on the original data-set that we constructed to calculate the reserve prices. We refer to this data-set as the *training data-set*. To evaluate our algorithm on the training data-set, we draw 200 vector of reserve prices from distribution q^* , and choose the best of these reserve prices and zero reserve prices. Here, q^* is the optimal solution to the problem (Profile-LP) when the input is the training data-set. We then report




Figure 1: The normalized revenue of the Pro-LPR algorithm and greedy algorithm of [RW16] on the training data-sets when bids are independent ($w = 0$), positively correlated ($w = 0.2$), and negatively correlated ($w = -0.2$). Here, the normalization is done with respect to the optimal value of Problem (Profile-LP), which is an upper bound on revenue of any vector of reserve prices.

the average revenue of these 200 random draws as the revenue of our algorithm on the training data-set. Figure 1 shows the distribution of normalized revenue of the Pro-LPR and greedy algorithms on the training data-sets when bids are independent ($w = 0$), positively correlated ($w = 0.2$), and negatively correlated ($w = -0.2$). Here, the normalization is done with respect to the upper bound, i.e., the optimal value of Problem (Profile-LP). We observe that in the worst problem instance, our algorithm obtains a 0.98 fraction of the upper bound, whereas the greedy algorithm in the worst case only obtains approximately an 0.82 fraction of the upper bound. Furthermore, while our algorithm matches the upper bound in at least half of the problem instances, the greedy algorithm does not meet the upper bound even in its best problem instance.

Second, we evaluate the revenue of our algorithm on *test data-sets*. Specifically, for each value of (μ, w) , we construct 100 test data-sets where each test data-set consists of 100 auctions whose bid distributions are the same as the training data-set with the same value of (μ, w) . On the test data-sets, the upper bound is irrelevant. Thus, we only report the revenue gain of our algorithm relative to the greedy algorithm. To assess our algorithm, for each value of (μ, w) , we use one of the 200 random reserve prices that obtained the highest revenue on the training data-set. Figure 2 presents three boxplots, where each boxplot shows the distribution of the revenue gain (in percentage) of our algorithm (relative to the greedy algorithm) on test data-sets for $w \in \{-0.2, 0, 0.2\}$. Note that in each boxplot we have $50 \times 100 = 5000$ data-points, where each data-point corresponds to one value of μ and one test data-set. Recall that for every value of (μ, w) , we construct 100 test data-sets and consider 50 (random) values for μ . We observe that the median of the revenue gain is between 6% and 7%. Furthermore, the 25-th percentile in all the boxplots is between 2% – 3%. This means that in 75% of the problem instances, our algorithm beats the greedy algorithm by at least 2% – 3%.




Figure 2: The revenue gain of Pro-LPR algorithm (relative to the greedy algorithm of [RW16]) on test data-sets when bids are independent ($w = 0$), positively correlated ($w = 0.2$), and negatively correlated ($w = -0.2$).

9 Conclusion

In this paper, we take a data-driven approach to optimize personalized reserve prices in eager second price auctions. We design a polynomial time LP-based algorithm to optimize reserve prices on a given data-set of submitted bids and show that our algorithm obtains more than 0.684 fraction of the optimal revenue. Our algorithm, which takes advantage of all the submitted bids to devise an effective reserve prices, highlights the importance of deviating from greedy policies to optimize reserve prices. Furthermore, our theoretical results and numerical studies confirm that our algorithm performs well when bids are correlated across buyers or independent of each other. Nonetheless, it is an exciting future research direction to explore if one can design a data-driven algorithm with a better approximation factor when bids are independent of each other.

Another exciting future research direction is to explore how to “transform” our LP-based algorithm to an online learning algorithm with a sublinear approximate regret. For the greedy algorithm of [RW16], such transformation has shown to be possible using the Follow-the-Perturbed-Leader algorithm ([RW16]) and Blackwell Approachability ([NGW⁺20]). However, the greedy algorithm can lose up to a 1/2 fraction of the optimal revenue, as opposed to a $1 - 0.684$ fraction of the optimal revenue that our algorithm can lose. This makes transforming our algorithm to its online counterpart an interesting future research direction. Finally, it is worth exploring if/how correlated rounding techniques can improve our approximation factor. In particular, one can explore if it is possible to close the gap between our approximation factor of 0.684 and the integrality gap of 0.828.

To sum up, we believe that our data-driven approach, as well as our LP-based algorithm can also be applied to a wider class of problems with revenue objective, and we hope the framework in this paper serves as a starting point for designing other data-driven algorithms.

References

- [AB18] Amine Allouah and Omar Besbes. Prior-independent optimal auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, page 503, 2018.
- [ACK18] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Prophet secretary: Surpassing the $1-1/e$ barrier. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 303–318, 2018.
- [ADH⁺19] AmirMahdi Ahmadinejad, Sina Dehghani, MohammadTaghi Hajiaghayi, Brendan Lucier, Hamid Mahini, and Saeed Seddighin. From duels to battlefields: Computing equilibria of blotto and other games. *Mathematics of Operations Research*, 44(4):1304–1325, 2019.
- [ARS13] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pages 1169–1177, 2013.
- [ARS14] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pages 622–630, 2014.
- [AS10] Arash Asadpour and Amin Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM Journal on Computing*, 39(7):2970–2989, 2010.
- [BBD⁺19] Soheil Behnezhad, Avrim Blum, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Christos H Papadimitriou, and Saeed Seddighin. Optimal strategies of blotto games: Beyond convexity. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 597–616, 2019.
- [BDD⁺17] Soheil Behnezhad, Sina Dehghani, Mahsa Derakhshan, MohammadTaghi HajiAghayi, and Saeed Seddighin. Faster and simpler algorithm for optimal strategies of blotto game. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [BDHS17] Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, and Aleksandrs Slivkins. A polynomial time algorithm for spatio-temporal security games. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 697–714, 2017.
- [BDHS18] Soheil Behnezhad, Mahsa Derakhshan, Mohammadtaghi Hajiaghayi, and Saeed Seddighin. Spatio-temporal games beyond one dimension. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 411–428, 2018.
- [Ber99] Dimitri P Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.
- [BFM12] Anand Bhalgat, Jon Feldman, and Vahab Mirrokni. Online allocation of display ads with smooth delivery. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1213–1221. ACM, 2012.
- [BFMM14] Santiago R Balseiro, Jon Feldman, Vahab Mirrokni, and S Muthukrishnan. Yield optimization of display advertising with ad exchange. *Management Science*, 60(12):2886–2907, 2014.

- [BGL⁺18] Hedyeh Beyhaghi, Negin Golrezaei, Renato Paes Leme, Martin Pal, and Balasubramanian Siva. Improved approximations for free-order prophets and second-price auctions. *arXiv preprint arXiv:1807.03435*, 2018.
- [BS06] Nikhil Bansal and Maxim Sviridenko. The santa claus problem. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 31–40, 2006.
- [BSS19] Nikhil Bansal, Aravind Srinivasan, and Ola Svensson. Lift-and-round to improve weighted completion time on unrelated machines. *SIAM Journal on Computing*, (0):STOC16–138, 2019.
- [CHMS10] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010.
- [CLMN14] L Elisa Celis, Gregory Lewis, Markus Mobius, and Hamid Nazerzadeh. Buy-it-now or take-a-chance: Price discrimination through randomized auctions. *Management Science*, 60(12):2927–2948, 2014.
- [CS14] Shuchi Chawla and Balasubramanian Sivan. Bayesian algorithmic mechanism design. *ACM SIGecom Exchanges*, 13(1):5–49, 2014.
- [CSZ20] Jose Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. *Mathematical Programming*, pages 1–39, 2020.
- [DFK11] Shahar Dobzinski, Hu Fu, and Robert D Kleinberg. Optimal auctions with correlated bidders are easy. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 129–138. ACM, 2011.
- [DRY15] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.
- [EHLM17] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monezadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.
- [EOS07] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.
- [ES10] Benjamin Edelman and Michael Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review*, 100(2):597–602, May 2010.
- [GJL19] Negin Golrezaei, Patrick Jaillet, and Jason Cheuk Nam Liang. Incentive-aware contextual pricing with non-parametric market noise. *arXiv preprint arXiv:1911.03508*, 2019.
- [GJM19] Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Operations Research*, 2019.
- [GLMN17] Negin Golrezaei, Max Lin, Vahab Mirrokni, and Hamid Nazerzadeh. Boosted second price auctions: Revenue optimization for heterogeneous bidders. 2017.

- [HK81] Theodore P. Hill and Robert P. Kertz. Ratio comparisons of supremum and stop rule expectations. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 56(2):283–285, Jun 1981.
- [HMR18] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. *SIAM Journal on Computing*, 47(3):651–674, 2018.
- [IKL⁺11] Nicole Immorlica, Adam Tauman Kalai, Brendan Lucier, Ankur Moitra, Andrew Postlewaite, and Moshe Tennenholtz. Dueling algorithms. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 215–224, 2011.
- [KN19] Yash Kanoria and Hamid Nazerzadeh. Incentive-compatible learning of reserve prices for repeated auctions. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 932–933, 2019.
- [KS78] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts and processes with finite values. *Adv. in Probab.*, 4:197–266, 1978.
- [MM14] Andres M Medina and Mehryar Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 262–270, 2014.
- [MR15] Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems*, pages 136–144, 2015.
- [Mye81] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [NGW⁺20] Rad Niazadeh, Negin Golrezaei, Joshua Wang, Fransisca Susan, and Ashwinkumar Badanidiyuru. Online learning via offline greedy algorithms: Applications in market design and optimization. *Available at SSRN 3613756*, 2020.
- [OS11] Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 59–60. ACM, 2011.
- [PLPV16] Renato Paes Leme, Martin Pál, and Sergei Vassilvitskii. A field guide to personalized reserve prices. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1093–1102. International World Wide Web Conferences Steering Committee, 2016.
- [PP11] Christos H Papadimitriou and George Pierrakos. On optimal single-item auctions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 119–128. ACM, 2011.
- [PS15] Lukáš Poláček and Ola Svensson. Quasi-polynomial local search for restricted max-min fair allocation. *ACM Transactions on Algorithms (TALG)*, 12(2):1–13, 2015.
- [Ron01] Amir Ronen. On approximating optimal auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 11–17. ACM, 2001.

- [RS02] Amir Ronen and Amin Saberi. On the hardness of optimal auctions. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 396–405. IEEE, 2002.
- [RS16] Tim Roughgarden and Okke Schrijvers. Ironing in the dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 1–18. ACM, 2016.
- [RW16] Tim Roughgarden and Joshua R. Wang. Minimizing regret with multiple reserves. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 601–616, 2016.
- [Sve12] Ola Svensson. Santa claus schedules jobs on unrelated machines. *SIAM Journal on Computing*, 41(5):1318–1341, 2012.
- [XYL09] Baichun Xiao, Wei Yang, and Jun Li. Optimal reserve price for the generalized second-price auction in sponsored search advertising. *Journal of Electronic Commerce Research*, 10(3):114, 2009.
- [Yan11] Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719, 2011.