

# Learning Valuation Distributions from Partial Observation

Avrim Blum<sup>\*1</sup>, Yishay Mansour<sup>†2</sup>, and Jamie Morgenstern<sup>‡1</sup>

<sup>1</sup>*Computer Science Department, Carnegie Mellon University*

<sup>2</sup>*Tel Aviv University and Microsoft Research, Hertzelia*

July 11, 2014

## Abstract

Auction theory traditionally assumes that bidders' valuation distributions are known to the auctioneer, such as in the celebrated, revenue-optimal Myerson auction [Myerson, 1981]. However, this theory does not describe how the auctioneer comes to possess this information. Recently, Cole and Roughgarden [2014] showed that an approximation based on a finite sample of independent draws from each bidder's distribution is sufficient to produce a near-optimal auction. In this work, we consider the problem of learning bidders' valuation distributions from much weaker forms of observations. Specifically, we consider a setting where there is a repeated, sealed-bid auction with  $n$  bidders, but all we observe for each round is *who* won, but not how much they bid or paid. We can also participate (i.e., submit a bid) ourselves, and observe when *we* win. From this information, our goal is to (approximately) recover the inherently recoverable part of the underlying bid distributions. We also consider extensions where different *subsets* of bidders participate in each round, and where bidders' valuations have a common-value component added to their independent private values.

---

<sup>\*</sup>Supported in part by the National Science Foundation under grants CCF-1101215, CCF-1116892, CCF-1331175, and IIS-1065251. Email: [avrim@cs.cmu.edu](mailto:avrim@cs.cmu.edu)

<sup>†</sup>This research was supported in part by The Israeli Centers of Research Excellence (I-CORE) program, (Center No. 4/11), by a grant from the Israel Science Foundation (ISF), by a grant from United States-Israel Binational Science Foundation (BSF), and by a grant from the Israeli Ministry of Science (MoS). Email: [mansour@tau.ac.il](mailto:mansour@tau.ac.il)

<sup>‡</sup>Supported in part by the National Science Foundation under grants CCF-1116892, CCF-1331175 and IIS-1065251 and by a Simons Award for Graduate Students in Theoretical Computer Science. Email: [jamiemmt@cs.cmu.edu](mailto:jamiemmt@cs.cmu.edu)

# 1 Introduction

Imagine that you get a call from your supervisor, who asks you to find out how much various companies are bidding for banner advertisements on a competitor’s web site. She wants you to recover the distribution of the bids for each one of the advertisers. Your boss might have many reasons why she wants this information: to compare their bids there and on your web site; to use as market research for opening a new web site which would be attractive to some of those advertisers; or simply to estimate the projected revenue.

This would be a trivial task if your competitor was willing to give you this information, but this is unlikely to happen. Industrial espionage is illegal, and definitely not within your expertise as a computer scientist. So, you approach this task from the basics and consider what you might observe. At best, you might be able to observe the outcome for a particular auction, namely the winner, but definitely not the price, and certainly not the bids of all participants. There is, however, a way to observe more detailed information: *you can participate in a sequence of auctions and see whether or not you win!* If you lose with a bid  $b$ , you know that the winner (and perhaps other bidders) bid more than  $b$ ; if you win with bid  $b$ , you know every other bidder bid less than  $b$ . In general, we assume you will also observe the winner of the auction explicitly (e.g., you can visit the webpage and view the banner ad of the auction in question). Is this a strong enough set of tools to recover the distributions over independent but not necessarily identical bid distributions?

If only your boss had instead given you the task of estimating the winning bid distribution in that auction, you would be able to accomplish this easily. By inserting random bids, and observing their probability of winning, you would be able to recover the distribution of the winning bid. However, this was not the task you were assigned: your boss wants the distribution of each bidder’s bids, not just those where they win the auction.

As a first attempt, your esteemed colleague suggests a trivial (and completely incorrect) approach (which you do not even consider). As before, you can submit random bids, and observe for each advertiser, how many times he wins in auctions with your random bid. This will estimate the distribution over bids he makes in auctions he wins. However, when we condition on a bidder  $i$  winning, we should expect to see a sample which is skewed towards higher bids. To see that the distribution over winning bids is a poor estimation for the distribution over bids for each bidder, consider the following example. Suppose you can even observe the bid of the winner. There are  $n$  advertisers, each bidding uniformly in  $[0, 1]$ . The distribution of the winning bid of a given advertiser would have an expectation of  $\frac{n}{n+1}$ , whereas the expectation of his bid is  $\frac{1}{2}$ ; indeed, the distribution over winning bids would be a poor approximation to his true bid distribution, namely, uniform in  $[0, 1]$ . Additional complications arise with this approach when advertisers are asymmetric, which is certainly the case in practice.

At this point, you decide to take a more formal approach, since the simplest possible technique fails miserably. This leads you to the following abstraction. There are  $n$  bidders, where bidder  $i$  has bid distribution  $\mathcal{D}_i$ . The  $n$  bidders participate in a sequence of auctions. In each auction, each bidder draws an independent bid  $b_i \sim \mathcal{D}_i$  and submits it.<sup>1</sup> We have the power to submit a bid  $b_0$ , which is independent of the bids  $b_i$ , to the auction. After each auction we observe the identity of the winner (but nothing else about the bids). Our goal is to construct a distribution  $\hat{\mathcal{D}}_i$  for each advertiser  $i$  which is close to  $\mathcal{D}_i$  in total variation. Our main result in this work is to solve this problem efficiently. Namely, we derive a polynomial time algorithm (with polynomial sample complexity) that recovers an approximation  $\hat{\mathcal{D}}_i$  of each of the distributions  $\mathcal{D}_i$ , down to some price  $p_\gamma$ , below which there is at most  $\gamma$  probability of any bidder winning.<sup>2</sup>

Following your astonishing success in recovering the bid distributions of the advertisers, your boss has a follow-up task for you. Not all items for sale, or users to which these ads are being shown, are created equal, and the advertisers receive various attributes describing the user (item for sale) before they submit their bid.

---

<sup>1</sup>We remark that if the repeated auction is incentive-compatible the bid and valuation of the advertiser would be the same (and we use them interchangeably). If this is not the case, then  $\mathcal{D}_i$  should be viewed as the distribution of bidder  $i$ ’s bids.

<sup>2</sup>If the winning bid is never (or very rarely) below some price  $p$ , then we will not be able to learn approximations to the distributions  $\mathcal{D}_i$  below  $p$ . For example, if bidder 1’s distribution  $\mathcal{D}_1$  has support only on  $[\frac{1}{2}, 1]$  and bidder 2’s distribution  $\mathcal{D}_2$  has support only on  $[0, \frac{1}{2})$ , then since the winning bid is always at least  $\frac{1}{2}$ , we will never be able to learn anything about  $\mathcal{D}_2$  other than the fact that its support lies in  $[0, \frac{1}{2})$ . Thus, our goal will be to learn a good approximation to each  $\mathcal{D}_i$  only above a price  $p_\gamma$  such that there is at least a  $\gamma$  probability of the winning bid being below  $p_\gamma$ .

Those attributes may include geographic location, language, operating system, browser, as well as highly sensitive data that might be collected through cookies. Your boss asks you to recover how the advertisers bid as a function of those vectors.

For this more challenging task, we can still help, under the assumption that we have access to these attributes for the observed auctions, under some assumptions. We start with the assumption that each bidder uses a linear function of the attributes for his bid. Namely, let  $x$  be the attribute vector of the user, then each advertiser has a weight vector  $w_i$  and his bid is  $x \cdot w_i$ . For this case we are able to recover efficiently an approximation  $\hat{w}_i$  of the weight vectors  $w_i$ .

A related task is to assume that the value (or bid) of an advertiser has a common shared component plus a private value which is stochastic. Namely, given a user with attributes  $x$ , the shared value is  $x \cdot w$ , where the  $w$  is the same to all advertisers, and each advertiser draws a private value  $v_i \sim \mathcal{D}_i$ . The bid of advertiser  $i$  is  $x \cdot w + v_i$ . The goal is to recover both the shared weights  $w$  as well as the individual distributions. We do this by “reduction” to the case of no attributes, by first recovering an approximation  $\hat{w}$  for  $w$ , and then using it to compute the common value for each user  $x$ .

One last extension we can handle focuses on who participates in the auction. So far, we assumed that in each auction, all the advertisers participate. However, this assumption is not really needed. Our approach is flexible enough, such that if we received for each auction the participants, this will be enough to recover the bidding distributions for each bidder who shows up often enough. Note that if there are  $n$  advertisers and each time a random subset shows up, we are unlikely to see the same subset show up twice; we can learn about bidder  $i$ ’s distribution over bids even when she is never competing in the same context, assuming her bid distribution does not depend on who else is bidding.

## 1.1 Related Work

Problems of reconstructing distributional information from limited or censored observations have been studied in both the medical statistics literature and the manufacturing/operations research literature. In medical statistics, a basic setting where this problem arises is estimating survival rates (the likelihood of death within  $t$  years of some medical procedure), when patients are continually dropping out of the study, independently of their time of death. The seminal work in this area is the Kaplan-Meier product-limit estimator [Kaplan and Meier, 1958], analyzed in the limit in the original paper and then for finite sample sizes in Foldes and Rejto [1981], see also its use in Ganchev et al. [2010]. In the manufacturing literature, this problem arises when a device, composed of multiple components, breaks down when the first of its components breaks down. From the statistics of when devices break down and which components failed, the goal is to reconstruct the distributions of individual component lifetimes [Nadas, 1970, Meilijson, 1981]. The methods developed (and assumptions made, and types of results shown) in each literature are different. In our work, we will build on the approach taken by the Kaplan-Meier estimator (described in more detail in Section 3), as it is more flexible and better suited to the types of guarantees we wish to achieve, extending it and using it as a subroutine for the kinds of weak observations we work with.

The area of prior-free mechanism design has aimed to understand what mechanisms achieve strong guarantees with limited (or no) information about the priors of bidders, particularly in the area of revenue maximization. There is a large variety of truthful mechanisms that guarantee a constant approximation (see, cf, Hartline and Karlin [2007]). A different direction is adversarial online setting which minimize the regret with respect to the best single price (see, Kleinberg and Leighton [2003]), or minimizing the regret for the reserve price of a second price auction Cesa-Bianchi et al. [2013]. In Cesa-Bianchi et al. [2013] it was assumed that bidders have an identical bid distribution and the algorithm observes the actual sell price after each auction, and based on this the bidding distribution is approximated.

A recent line of work tries to bridge between the Bayesian setting and the adversarial one, by assuming we observe a limited number of samples. For a regular distribution, as single sample bidders’ distributions is sufficient to get a  $1/2$ -approximation to the optimal revenue [Dhangwatnotai et al., 2010], which follows from an extension of the Bulow and Klemperer [1994] result that shows the revenue from a second-price auction with  $n + 1$  (i.i.d) bidders is higher than the revenue from running a revenue-optimal auction with  $n$  bidders. Recent work of Cole and Roughgarden [2014] analyzes the number of samples necessary to construct a  $1 - \epsilon$ -

approximately revenue optimal mechanism for asymmetric bidders: they show it is necessary and sufficient to take poly( $\frac{1}{\epsilon}, n$ ) samples from each bidder’s distribution to construct an  $1 - \epsilon$ -revenue-optimal auction for bid distributions are strongly regular. We stress that in this work we do not make *any* assumptions about the bid distribution.

Chawla et al. [2014] design mechanisms which are approximately revenue-optimal and also allow for good *inference*: from a sample of bids made in Bayes-Nash equilibrium, they would like to reconstruct the distribution over values from which bidders are drawn. This learning technique relies heavily on a sample being drawn *unconditionally* from the *symmetric* bid distribution, rather than only seeing the *winner’s identity* from *asymmetric* bid distributions, as we consider in this work.

We stress that in all the “revenue maximization” literature has a fundamentally different objective than the one in this paper. Namely, our goal is to reconstruct the bidders’ bid distributions, rather than focusing of the revenue directly. Our work differs from previous work in this space in that it assumes very limited observational information. Rather than assuming all  $n$  bids as an observation from a single run of the auction, or even observing only the price, we see only the identity of highest bidder. We do not need to make any regularity assumption on the bid distribution (monotone hazard rate, regular, etc.), our methodology handles *any* continuous bid distribution.<sup>3</sup>

## 2 Model and Preliminaries

We assume there are  $n$  bidders, and each  $i \in [n]$  has some unknown valuation distribution  $\mathcal{D}_i$  over the interval  $[0, 1]$ . Each sample  $t \in [m]$  refers to a fresh draw  $v_i^t \sim \mathcal{D}_i$  for each  $i$ . The label of sample  $t$  will be denoted  $y^t = \operatorname{argmax}_i v_i^t$ , the identity of the highest bidder. Our goal is to estimate  $F_i$ , the cumulative distribution for  $\mathcal{D}_i$ , for each bidder  $i$ , up to  $\epsilon$  additive error for all values in a given range. In Section 4 we examine extensions and modifications to this basic model.

We consider the problem of finding (sample and computationally) efficient algorithms for constructing an estimate  $\widehat{F}_i$  of  $F_i$ , the cumulative distribution function, such that for all bidders  $i$  and price levels  $p$ ,  $\widehat{F}_i(p) \in \{F_i(p) \pm \epsilon\}$ . However, as discussed above, this goal is too ambitious in two ways. First, if the labels contain no information about the value of bids, the best we could hope to learn is the relative probability each person might win, which is insufficient to uniquely identify the CDFs, even without sampling error. We address this issue by allowing, at each time  $t$ , our learning algorithm to insert a fake bidder 0 (or reserve) of value  $v_0^t = r^t$ ; the label at time  $t$  will be  $y^t = \operatorname{argmax}_i v_i^t$  ( $y^t = 0$  will refer to a sample where the reserve was not met, or the fake bidder won the auction). The other issue, also described above, is that there will be values below which we simply cannot estimate the  $\mathcal{D}_i$  since bids below that value do not win. In particular, if bids below price  $p$  never win, then any two cumulatives  $F_i, F'_i$  that agree above  $p$  will be statistically indistinguishable. Thus, we will consider a slightly weaker goal. We will guarantee our estimates  $\widehat{F}_i(p) \in F_i(p) \pm \epsilon$  for all  $p$  where  $\mathbb{P}[\text{someone winning with a bid at most } p] \geq \gamma$ . Then, our goal is to minimize  $m$ , the number of samples necessary, to do so, and we hope to have  $m \in \text{poly}(n, \frac{1}{\epsilon}, \frac{1}{\gamma})$ , with high probability of success over the draw of the sample. One final (and necessary) assumption we will make is that each  $\mathcal{D}_i$  has no point masses, and our algorithm will be polynomial in the maximum slope  $L$  of the  $F_i$ s.<sup>4</sup>

### 2.1 A brief primer on the Kaplan-Meier estimator

Our work is closely related in spirit to that of the Kaplan-Meier estimator, KM, for survival time; in this section, we describe the techniques used for constructing the KM [Kaplan and Meier, 1958]. This will give some intuition for the estimator we present in Section 3. We translate the results found in Kaplan and Meier [1958] to an auction setting from the survival rate literature. Suppose each sample  $t$  is of the following form.

<sup>3</sup>Note that we measure the distance between two distributions using the total variation distance, which is essentially “additive”.

<sup>4</sup>This assumption is useful for this paper for two reasons. First, it allows us to eschew any issues associated with tie-breaking, since they happen with probability 0. Second, if there were no continuity assumption, there might be point masses. If we wished to have an additive accuracy guarantee as above, this would force our learning algorithm to be able to determine the *exact* location of these point masses, which couldn’t be done in polynomial time (for example, suppose  $F_i$  had a point mass at  $\sqrt{2}$ ).

Each bidder  $i$  draws their bid  $b_i^t \sim \mathcal{D}_i$  independently of each other bid. The label  $y^t = (\max_i b_i^t, \operatorname{argmax}_i b_i^t)$  consists of the winning bid and the identity of the winner. From this, we would like to reconstruct an estimate  $\hat{F}_i$  of  $F_i$ . Given  $m$  samples, relabel them so that the winning bids are in increasing order, e.g.  $b_{i_1}^1 \leq b_{i_2}^2 \leq b_{i_m}^m$ . Here is some intuition behind the KM:  $\mathbb{P}[b_i \leq x] = \mathbb{P}[b_i \leq x | b_i \leq y] \cdot \mathbb{P}[b_i \leq y]$  for  $y > x$ . Repeatedly applying this, we can see that, for  $x < y_1 < y_2 < \dots < y_r$ ,

$$\begin{aligned} F_i(x) &= \mathbb{P}[b_i \leq x] = \mathbb{P}[b_i \leq x | b_i \leq y_1] \mathbb{P}[b_i \leq y_1 | b_i \leq y_2] \cdots \mathbb{P}[b_i \leq y_{r-1} | b_i \leq y_r] \mathbb{P}[b_i \leq y_r] \\ &= \mathbb{P}[b_i \leq x | b_i \leq y_1] \mathbb{P}[b_i \leq y_r] \prod_{t=1}^r \mathbb{P}[b_i \leq y_t | b_i \leq y_{t+1}] \end{aligned} \quad (1)$$

Now, we can employ the observation in Equation 1, if only we knew how to convert the samples into estimates of such conditional probabilities. Since other players' bids are independent, we can estimate the conditional probabilities as follows:

$$\mathbb{P} \left[ b_i \leq b_{i_t}^t | b_i \leq b_{i_{t+1}}^{t+1} \right] \approx \begin{cases} \frac{t-1}{t} & \text{if } i \text{ won sample } t \\ 1 & \text{if } j \neq i \text{ won sample } t \end{cases} \quad (2)$$

Thus, combining Equations 1 and 2, we have the Kaplan-Meier estimator:

$$\text{KM}(x) = \prod_{t: b_i^t \geq x} \left( \frac{t-1}{t} \right)^{\mathbb{I}[i \text{ won sample } t]}$$

Our estimator is morally similar to KM, though it differs in several important ways. First, and most importantly, we do not see the winning bid explicitly; instead, we will just have lower or upper bounds on the highest non-reserve bid (namely, the reserve bid when someone wins or we win, respectively). Secondly, KM generally has no control issue; in our setting, we are *choosing* one of the values which will censor our observation. We need to pick appropriate reserves to get a good estimator (picking reserves that are too high will censor too many observations, only giving us uninformative upper bounds on bids, and reserves that are too low will never win, giving us uninformative lower bounds on bids). Our estimator searches the space  $[0, 1]$  for appropriate price points to use as reserves to balance these concerns.

### 3 Learning bidders' valuation distributions

In this section, we assume we have the power to insert a reserve price, and observe who won. Using this, we would like to reconstruct the CDFs of each bidder  $i$  up to some error, down to some price  $p_i$  where  $i$  has probability no more than  $\gamma$  of winning at or below  $p_i$ , up to additive accuracy  $\epsilon$ . Our basic plan of attack is as follows. We start by estimating the probability  $i$  wins with a bid in some range  $[a, a + \delta]$ , by setting reserve prices at  $a$  and  $a + \delta$ , and measuring the difference in empirical probability that  $i$  wins with the two reserves. We then estimate the probability that no bidder bids above  $a + \delta$  (by setting a reserve of  $a + \delta$  and observing the empirical probability that no one wins). These together will be enough to estimate the probability that  $i$  wins with a bid in that range, conditioned on no one bidding above the range. We then show, for a small enough range, this is a good estimate for the probability  $i$  bids in the range, conditioned on no one bidding above the range. Then, we chain these estimates together to form **Kaplan**, our estimator.

More specifically, to make this work we select a partition of  $[0, 1]$  into a collection of intervals. This partition should have the following property. Within each interval  $[x, y]$ , there should be probability at most  $\beta$  of any person bidding in  $[x, y]$ , conditioned on no one bidding above  $y$ . This won't be possible for the lowest interval, but will be true for the other intervals. Then, the algorithm estimates the probability  $i$  will win in  $[x, y]$  conditioned on all bidders bidding at most  $y$ . This then  $(1 - \beta)$  (multiplicatively) approximates the probability  $i$  bids in  $[x, y]$  (conditioned on all bidders bidding less than  $y$ ). Then, the algorithm combines these estimates in a way such that the approximation factors do not blow up to reconstruct the CDF.

---

**Algorithm 1:** Kaplan, estimates the CDF of  $i$  from samples with reserves

---

**Data:**  $\epsilon, \gamma, \delta, L$ , where  $L$  is the Lipschitz constant of the  $F_i$ s  
**Result:**  $\widehat{F}_i$

- 1 Let  $\widehat{F}_i(0) = 0, \widehat{F}_i(1) = 1, k = \frac{2Ln}{\beta\gamma} + 1, \delta' = \frac{\delta}{3k(\log k+1)}, \beta = \frac{\epsilon\gamma}{32nL}, \alpha = \beta^2/96, \mu = \beta/96, T = \frac{8 \ln 6/\delta'}{\alpha^2 \gamma^2 (\frac{\mu}{2})^2}$ ;
- 2 Let  $\ell_1, \dots, \ell_{k'} = \text{Intervals}(\beta, \gamma, T)$ ;
- 3 **for**  $t = 2$  **to**  $k' - 1$  **do**
- 4    $\lfloor$  Let  $r_{\ell_\tau, \ell_{\tau+1}} = \text{IWin}(i, \ell_\tau, \ell_{\tau+1}, T)$ ;
- 5 **for**  $t = 2$  **to**  $k' - 1$  **do**
- 6    $\lfloor$  Let  $\widehat{F}_i(\ell_\tau) = \prod_{\tau' \geq t+1} (1 - r_{\ell_{\tau'}, \ell_{\tau'+1}})$ ;
- 7 Define  $\widehat{F}_i(x) = \max_{\ell_\tau \leq x} \widehat{F}_i(\ell_\tau)$ ;

---

**Theorem 3.1.** *With probability at least  $1 - \delta$ , Kaplan outputs  $\widehat{F}_i$ , an estimate of  $F_i$ , with sample complexity*

$$m = O\left(\frac{n^8 L^8 \ln \frac{nL}{\epsilon\gamma} \left(\ln \frac{1}{\delta} + \ln \ln \frac{nL}{\epsilon\gamma}\right)}{\gamma^{10} \epsilon^6}\right)$$

and, for all  $p$  where  $\mathbb{P}[\exists j \text{ s.t. } j \text{ wins with a bid } \leq p] \geq \gamma$ , if each CDF is  $L$ -Lipschitz, the error is at most:

$$F_i(p) - \epsilon \leq \widehat{F}_i(p) \leq F_i(p) + \epsilon.$$

Kaplan calls several other functions, which we will now informally describe, and state several Lemmas describing their guarantees (the formal definitions can be found in Figure 1 and the proofs can be found in Appendix A). **IWin** estimates the probability  $i$  wins in the region  $[\ell_\tau, \ell_{\tau+1}]$ , conditioned on all bids being at most  $\ell_{\tau+1}$ . **Intervals** partitions  $[0, 1]$  into small enough intervals such that, conditioned on all bids being in or below that interval, the probability of any bidder bidding within the interval is small. (Essentially,  $\ell_2$  is  $p_\gamma$ , and therefore we are not interested in the estimation in  $[0, \ell_2]$ , and by definition  $\ell_1 = 0$ .)

Here are three lemmas which will be useful in the proof of Theorem 3.1. Lemma 3.2 bounds the number of samples **IWin** uses and bounds the error of its estimate. Lemma 3.3 does similarly for **Intervals**. Lemma 3.4 states that, if a region  $[\ell_\tau, \ell_{\tau+1}]$  is small enough, the probability that  $i$  bids in  $[\ell_\tau, \ell_{\tau+1}]$  (conditioned on all bids being at most  $\ell_{\tau+1}$ ) is well-approximated by the probability that  $i$  wins with a bid in  $[\ell_\tau, \ell_{\tau+1}]$  (conditioned on all bids being at most  $\ell_{\tau+1}$ ). In combination, these three imply a guarantee on the sample complexity and accuracy of estimating  $\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}]$ , which is the key ingredient of the Kaplan estimator.

**Lemma 3.2.** *Suppose, for a fixed interval  $[\ell_\tau, \ell_{\tau+1}]$ ,  $\mathbb{P}[i \text{ wins in } [0, \ell_{\tau+1}]] \geq \gamma$ . Then, **IWin** ( $i, \ell_\tau, \ell_{\tau+1}, T$ ) outputs  $p_{\ell_\tau, \ell_{\tau+1}}^i$  such that*

$$(1 - \mu)\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}] - \alpha \leq p_{\ell_\tau, \ell_{\tau+1}}^i \leq (1 + \mu)\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}] + \alpha,$$

with probability at least  $1 - 3\delta'$  and uses  $3T$  samples, for the values of  $T, \delta'$  as in Kaplan.

**Lemma 3.3.** *Let  $T$  as in Kaplan. Then, **Intervals**( $\beta, \gamma, T, L, n$ ) returns  $0 = \ell_1 < \dots < \ell_k = 1$  such that*

1.  $k \leq \frac{48Ln}{\beta\gamma}$
2. For each  $\tau \in [2, k]$ ,  $\mathbb{P}[\max_j b_j \in [\ell_\tau, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}] \leq \frac{\beta}{16}$
3.  $\mathbb{P}[\max_j b_j \in [\ell_1, \ell_2]] \leq \gamma$

with probability at least  $1 - 3k \log(k) \delta'$ , when bidders' CDFs are  $L$ -Lipschitz, using at most  $3kT \log k$  samples.

With the guarantee of Lemma 3.3, we know that the partition of  $[0, 1]$  returned by `Intervals` is “fine enough”. Now, Lemma 3.4 shows that, when the partition fine enough, the conditional probability  $i$  wins with a bid in each interval is a good estimate for the conditional probability  $i$  bids within that interval.

**Lemma 3.4.** *Suppose that, for bidder  $i$  and some  $0 \leq \ell_\tau \leq \ell_{\tau+1} \leq 1$ ,*

$$\mathbb{P}[\max_{j \neq i} b_j \in [\ell_\tau, \ell_{\tau+1}] \mid \max_{j \neq i} b_j < \ell_{\tau+1}] \leq \beta.$$

Then,

$$1 \geq \frac{\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j < \ell_{\tau+1}]}{\mathbb{P}[i \text{ bids in } [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j < \ell_{\tau+1}]} \geq 1 - \beta$$

Finally, we observe that  $F_i$  can be written as the product of conditional probabilities

**Observation 3.5.** *Consider some set of points  $0 < \ell_1 < \dots < \ell_k = 1$ .  $F_i(\ell_\tau)$  can be rewritten as the following product:*

$$F_i(\ell_{\tau-1}) = F_i(\ell_\tau)(1 - \mathbb{P}[b_i \geq \ell_{\tau-1} \mid b_i \leq \ell_\tau]) = \prod_{\tau' \geq \tau} (1 - \mathbb{P}[b_i \geq \ell_{\tau'-1} \mid b_i \leq \ell_{\tau'}]) = \prod_{\tau' \geq \tau} (1 - \mathbb{P}[b_i \in [\ell_{\tau'-1}, \ell_{\tau'}] \mid b_i \leq \ell_{\tau'}])$$

With these pieces in place, we prove Theorem 3.1.

*Proof of Theorem 3.1.* Notice that there are at most  $k'$  events each of which happens with probability at most  $\delta' = \frac{\delta}{k'}$  (namely, that `Intervals` returns a poor partition, or for each interval, of which there are at most  $k' - 1$ , by Lemma 3.3, that `IWin` is not accurate as described by Lemma 3.2). Thus, by a union bound, none of these events occur with probability  $1 - \delta$ . Thus, for the remainder of the proof we assume the partition returned by `Intervals` is good and each call to `IWin` is accurate.

It will suffice to prove, for the lattice points in our discretization, that `Kaplan` provides an  $\epsilon$ -approximation to the CDF. This follows because

$$\begin{aligned} F_i(\ell_\tau) - F_i(\ell_{\tau-1}) &= \mathbb{P}[i \text{ bids in } [\ell_{\tau-1}, \ell_\tau]] \\ &= \mathbb{P}[i \text{ bids in } [\ell_{\tau-1}, \ell_\tau] \mid b_i \leq \ell_\tau] \\ &\leq \mathbb{P}[i \text{ bids in } [\ell_{\tau-1}, \ell_\tau] \mid \max_j b_j \leq \ell_\tau] \\ &\leq (1 + \beta) \mathbb{P}[i \text{ wins in } [\ell_{\tau-1}, \ell_\tau] \mid \max_j b_j \leq \ell_\tau] \\ &\leq (1 + \beta) \beta = \beta + \beta^2 \leq \frac{\epsilon}{2} \end{aligned}$$

where the third and fourth inequality follows from Lemma 3.3 and Lemma 3.4, and the final one from the fact that  $\beta < \frac{\epsilon}{4}$ . Thus, our lattice is fine enough that it suffices to show accuracy of the lattice points. We start by rewriting  $F_i(\ell_\tau)$ , using Observation 3.5:

$$F_i(\ell_\tau) = \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[b_i \in [\ell_{\tau'-1}, \ell_{\tau'}] \mid b_i \leq \ell_{\tau'}]). \quad (3)$$

So, one can compute the probability of bidding at most  $\ell_{\tau-1}$  by multiplying together a collection of probabilities of bidding within intervals above  $\ell_\tau$ . Let the event  $\max_j b_j \leq \ell_{\tau'}$  be denoted  $M_{\ell_{\tau'}}$ . Now, we can apply Lemma 3.3 to imply that, for all  $\tau'$ ,

$$\mathbb{P}[\max_j b_j \in [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}] \leq \frac{\beta}{16} = \beta'$$

which, by Lemma 3.4, implies for all  $\tau'$  that

$$1 \geq \frac{\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]}{\mathbb{P}[i \text{ bids in } [\ell_{\tau'-1}, \ell_{\tau'}] | m_{\ell_{\tau'}}]} = \frac{\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]}{\mathbb{P}[i \text{ bids in } [\ell_{\tau'-1}, \ell_{\tau'}] | i \text{ bids in } [0, \ell_{\tau'}]]} \geq 1 - \beta' \quad (4)$$

where the equality comes from the independence of the bids. Then, combining Equations (4) and (3), we know

$$\prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) = F_i(\ell_\tau) \geq \prod_{\tau' \geq \tau+1} (1 - \frac{\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]}{1 - \beta'})$$

Then, by Fact A.3,

$$F_i(\ell_\tau) \in \left[ \prod_{\tau' \geq \tau+1} (1 - (1 + 2\beta') \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]), \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \right]$$

Now, Lemma 3.2 states that the result of `IWin` are correct within an additive  $\alpha$  and multiplicative  $\mu$ , thus

$$\prod_{\tau' \geq \tau+1} (1 - (1 + \mu) \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] - \alpha) \leq \widehat{F}_i(\ell_\tau) \leq \prod_{\tau' \geq \tau+1} (1 - (1 - \mu) \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] + \alpha).$$

Now, we simply need to look at the potential difference in these terms. We will consider the lower bound on  $F_i(\ell_\tau)$  and upper bound on  $\widehat{F}_i(\ell_\tau)$  (the other direction is analogous).

$$\begin{aligned} & \prod_{\tau' \geq \tau+1} (1 - (1 - \mu) \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] + \alpha) - \prod_{\tau' \geq t+1} (1 - (1 + 2\beta') \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\ & \leq \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] + \mu\beta' + \alpha) - \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] - 2\beta'^2) \\ & \leq \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] + \beta'^2) - \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] - 2\beta'^2) \\ & \leq \prod_{\tau' \geq \tau+1} (1 - 2\beta'^2)(1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) - \prod_{\tau' \geq \tau+1} (1 + 2\beta'^2)(1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\ & \leq (1 - 2\beta'^2)^k \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) - (1 + 4\beta'^2)^k \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\ & \leq (1 - 4k\beta'^2) \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) - (1 + 8k\beta'^2) \prod_{\tau' \geq t+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\ & \leq 12k\beta'^2 \leq 12 \frac{16Ln}{\beta\gamma} \beta'^2 \leq \frac{3Ln\beta}{\gamma} \leq \frac{\epsilon}{2} \end{aligned}$$

where the first follows from  $\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | \max_j b_j < \ell_{\tau'}] \leq \beta'$ , the second by the definition of  $\alpha = \frac{\beta'^2}{2}$ ,  $\mu = \frac{\beta'}{2}$ , the third again, by  $\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] \leq \beta'$ , the fourth from  $2\beta' < \frac{1}{2}$ , the fifth and sixth from basic algebra, the seventh by the bound on  $k \leq \frac{16Ln}{\beta\gamma}$ , by Lemma 3.3, the eighth by  $\beta' = \frac{\beta}{16}$ , and the ninth by  $\beta = \frac{\epsilon\gamma}{32nL}$ .

The sample complexity bound and failure probability follow from Lemmas 3.3 and 3.2, substituting in for various parameters, since `IWin` is called  $k$  times. Thus, in total, there are  $\leq 3k \log(k) + 3k$  empirical estimates made, each with probability at most  $\delta'$  of failure, each with sample size  $T$ .  $\square$




Figure 1: Helper functions

### 3.1 Subsets

The argument above extends directly to a more general scenario in which not all bidders necessarily show up each time, and instead there is some distribution over  $2^{[k]}$  over which bidders show up each time the auction is run. As mentioned above, this is quite natural in settings where bidders are companies that may or may not need the auctioned resource at any given time, or keyword auctions where there is a distribution over keywords, and companies only participate in the auction of keywords that are relevant to them. To handle this case, we simply apply Algorithm 1 to just the subset of time steps in which bidder  $i$  showed up when learning  $\widehat{F}_i$ . We use the fact here that even though the distribution over subsets of bidders that show up need not be a product distribution (e.g., certain bidders may tend to show up together), the maximum bid value of the other bidders who show up with bidder  $i$  is a random variable that is independent of bidder  $i$ 's bid. Thus all the above arguments extend directly. The sample complexity bound of Theorem 3.1 is now a sample complexity on observations of bidder  $i$  (and so requires roughly a  $1/q$  blowup in total sample complexity to learn the distribution for a bidder that shows up only a  $q$  fraction of the time).

## 4 Extensions and Other Models

So far we have been in the usual model of independent private values. That is, on each run of the auction, bidder  $i$ 's value is  $v_i \sim \mathcal{D}_i$ , drawn independently from the other  $v_j$ . We now consider models motivated by settings where we have different items being auctioned on each round, such as different cameras, cars, or laptops, and these items have observable properties, or features, that affect their value to each bidder.

In the first (easier) model we consider, each bidder  $i$  has its own private weight vector  $w_i \in R^d$  (which we don't see), and each item is a feature vector  $x \in R^d$  (which we do see). The value for bidder  $i$  on item  $x$  is  $w_i \cdot x$ , and the winner is the highest bidder  $\operatorname{argmax}_i w_i \cdot x$ . There is a distribution  $\mathcal{P}$  over items, but no additional private randomness. Our goal, from submitting bids and observing the identity of the winner, is to learn estimates  $\tilde{w}_i$  that approximate the true  $w_i$  in the sense that for random  $x \sim \mathcal{P}$ , with probability  $\geq 1 - \epsilon$ , the  $\tilde{w}_i$  correctly predict the winner and how much the winner values the item  $x$  up to  $\pm\epsilon$ .

In the second model we consider, there is just a single common vector  $w$ , but we reintroduce the distributions  $\mathcal{D}_i$ . In particular, the value of bidder  $i$  on item  $x$  is  $w \cdot x + v_i$  where  $v_i \sim \mathcal{D}_i$ . The " $w \cdot x$ " portion can be viewed as a common value due to the intrinsic worth of the object, and if  $w = \vec{0}$  then this reduces to the setting studied in previous sections. The goal of the algorithm is to learn both the common vector  $w$  and all the  $\mathcal{D}_i$ .

The common generalization of the above two models, with different unknown vectors  $w_i$  and unknown distributions  $\mathcal{D}_i$  appears to be quite a bit more difficult (in part because the expected value of a draw from  $\mathcal{D}_i$  conditioned on bidder  $i$  winning depends on the vector  $x$ ). We leave as an open problem to resolve learnability (positively or negatively) in such a model. We assume that  $\|x\|_2 \leq 1$  and  $\|w_i\|_2 \leq 1$ , and as before, all valuations are in  $[0, 1]$ .

### 4.1 Private value vectors without private randomness

Here we present an algorithm for the setting where each bidder  $i$  has its own private vector  $w_i \in R^d$ , and its value for an item  $x \in R^d$  is  $w_i \cdot x$ . There is a distribution  $\mathcal{P}$  over items, and our goal, from submitting bids and observing the identity of the winner, is to accurately predict the winner and the winning bid. Specifically, we prove the following:

**Theorem 4.1.** *With probability  $\geq 1 - \delta$ , the algorithm below using sample size*

$$m = O\left(\frac{1}{\epsilon^2} [dn^2 \log(1/\epsilon) + \log(1/\delta)]\right)$$

*produces  $\tilde{w}_i$  such that on a  $1 - \epsilon$  probability mass of  $x \sim \mathcal{P}$  we have  $i^* \equiv \operatorname{argmax}_i \tilde{w}_i \cdot x = \operatorname{argmax}_i w_i \cdot x$  (i.e., a correct prediction of the winner), and additionally have  $|\tilde{w}_{i^*} \cdot x - w_{i^*} \cdot x| \leq \epsilon$ .*

*Proof.* Our algorithm is simple. We will participate in  $m$  auctions using bids chosen uniformly at random from  $\{0, \epsilon, 2\epsilon, \dots, 1\}$ . We observe the winners, then solve for a consistent set of  $\tilde{w}_i$  using linear programming. Specifically, for  $t = 1, \dots, m$ , if bidder  $i_t$  wins item  $x_t$  for which we bid  $b_t$ , then we have linear inequalities:

$$\begin{aligned} \tilde{w}_{i_t} \cdot x_t &> \tilde{w}_j \cdot x_t \quad (\forall j \neq i_t) \\ \tilde{w}_{i_t} \cdot x_t &> b_t. \end{aligned}$$

Similarly, if we win the item, we have:

$$b_t > \tilde{w}_j \cdot x_t \quad (\forall j).$$

Let  $\mathcal{P}^*$  denote the distribution over pairs  $(x, b)$  induced by drawing  $x$  from  $\mathcal{P}$  and  $b$  uniformly at random from  $\{0, \epsilon, 2\epsilon, \dots, 1\}$  and consider a  $(k + 1)$ -valued target function  $f^*$  that given a pair  $(x, b)$  outputs an integer in  $\{0, 1, \dots, n\}$  indicating the winner (with 0 indicating that our bid  $b$  wins). By design, the vectors  $\tilde{w}_1, \dots, \tilde{w}_n$  solved for above yield the correct answer (the correct highest bidder) on all  $m$  pairs  $(x, b)$  in our training sample. We argue below that  $m$  is sufficiently large so that by a standard sample complexity

analysis, with probability at least  $1 - \delta$ , the true error rate of the vectors  $\tilde{w}_i$  under  $\mathcal{P}^*$  is at most  $\epsilon^2/(1 + \epsilon)$ . This in particular implies that for at least a  $(1 - \epsilon)$  probability mass of items  $x$  under  $\mathcal{P}$ , the vectors  $\tilde{w}_i$  predict the correct winner for all  $\frac{1+\epsilon}{\epsilon}$  bids  $b \in \{0, \epsilon, 2\epsilon, \dots, 1\}$  (by Markov's inequality). This implies that for this  $(1 - \epsilon)$  probability mass of items  $x$ , not only do the  $\tilde{w}_i$  correctly predict the winning bidder but they also correctly predict the winning bid value up to  $\pm\epsilon$  as desired.

Finally, we argue the bound on  $m$ . Any given set of  $n$  vectors  $\tilde{w}_1, \dots, \tilde{w}_n$  induces a  $(n+1)$ -way partition of the  $(d+1)$ -dimensional space of pairs  $(x, b)$  based on which of  $\{0, \dots, n\}$  will be the winner (with 0 indicating that  $b$  wins). Each element of the partition is a convex region defined by halfspaces, and in particular there are only  $O(n^2)$  hyperplane boundaries, one for each pair of regions. Therefore, the total number of ways of partitioning  $m$  data-points is at most  $O(m^{(d+1)n^2})$ . The result then follows by standard VC upper bounds for desired error rate  $\epsilon^2/(1 + \epsilon)$ .  $\square$

## 4.2 Common value vectors with private randomness

We now consider the case that there is just a single common vector  $w$ , but we reintroduce the distributions  $\mathcal{D}_i$ . In particular, there is some distribution  $\mathcal{P}$  over  $x \in R^d$ , and the value of bidder  $i$  on item  $x$  is  $w \cdot x + v_i$  where  $v_i \sim \mathcal{D}_i$ . As before, we assume that  $\|x\|_2 \leq 1$  and  $\|w_i\|_2 \leq 1$ , and all valuations are in  $[0, 1]$ . The goal of the algorithm is to learn both the common vector  $w$  and all the  $\mathcal{D}_i$ . We show here how we can solve this problem by first learning a good approximation  $\tilde{w}$  to  $w$  which then allows us to reduce to the problem of Section 3. In particular, given parameter  $\epsilon'$ , we will learn  $\tilde{w}$  such that

$$\Pr_{x \sim \mathcal{P}} (|w \cdot x - \tilde{w} \cdot x| \leq \epsilon') \geq 1 - \epsilon'.$$

Once we learn such a  $\tilde{w}$ , we can reduce to the case of Section 3 as follows: every time the algorithm of Section 3 queries with some reserve bid  $b$ , we submit instead the bid  $b + \tilde{w} \cdot x$ . The outcome of this query now matches the setting of independent private values, but where (due to the slight error in  $\tilde{w}$ ) after the  $v_i$  are each drawn from  $\mathcal{D}_i$ , there is some small random fluctuation that is added (and an  $\epsilon'$  fraction of the time, there is a large fluctuation). But since we can make  $\epsilon'$  as polynomially small as we want, this becomes a vanishing term in the independent private values analysis. Thus, it suffices to learn a good approximation  $\tilde{w}$  to  $w$ , which we do as follows.

**Theorem 4.2.** *With probability  $\geq 1 - \delta$ , the algorithm below using running time and sample size polynomial in  $d, n, 1/\epsilon'$ , and  $\log(1/\delta)$ , produces  $\tilde{w}$  such that  $\Pr_{x \sim \mathcal{P}}[|\tilde{w} \cdot x - w \cdot x| \leq \epsilon'] \geq 1 - \epsilon'$ .*

*Proof.* Let  $\mathcal{D}_{max}$  denote the distribution over  $\max\{v_1, \dots, v_n\}$ . By performing an additive offset, specifically, by adding a new feature  $x_0$  that is always equal to 1 and setting the corresponding weight  $w_0$  to be the mean value of  $\mathcal{D}_{max}$ , we may assume without loss of generality from now on that  $\mathcal{D}_{max}$  has mean value 0.<sup>5</sup>

Now, consider the following distribution over labeled examples  $(x, y)$ . We draw  $x$  at random from  $\mathcal{P}$ . To produce the label  $y$ , we bid a uniform random value in  $[0, 1]$  and set  $y = 1$  if we lose and  $y = 0$  if we win (we ignore the identity of the winner when we lose). The key point here is that if the highest bidder for some item  $x$  bid a value  $b \in [0, 1]$ , then with probability  $b$  we lose and set  $y = 1$  and with probability  $1 - b$  we win and set  $y = 0$ . So,  $\mathbb{E}[y] = b$ . Moreover, since  $b = w \cdot x + v_{max}$ , where  $v_{max}$  is picked from  $\mathcal{D}_{max}$  which has mean value of 0, we have  $\mathbb{E}[b|x] = w \cdot x$ . So,  $\mathbb{E}[y|x] = w \cdot x$ .

So, we have examples  $x$  with labels in  $\{0, 1\}$  such that  $\mathbb{E}[y|x] = w \cdot x$ . This implies that  $w \cdot x$  is the predictor of minimum squared loss over this distribution on labeled examples (in fact, it minimizes mean squared error for every point  $x$ ). Moreover, any real-valued predictor  $h(x) = \tilde{w} \cdot x$  that satisfies the condition that  $\mathbb{E}_{(x,y)}[(\tilde{w} \cdot x - y)^2] \leq \mathbb{E}_{(x,y)}[(w \cdot x - y)^2] + \epsilon'^3$  must satisfy the condition:

$$\Pr_{x \sim \mathcal{P}} (|w \cdot x - \tilde{w} \cdot x| \leq \epsilon') \geq 1 - \epsilon'.$$

This is because a predictor that fails this condition incurs an additional squared loss of  $\epsilon'^2$  on at least an  $\epsilon'$  probability mass of the points. Finally, since all losses are bounded (we know all values  $w \cdot x$  are bounded

<sup>5</sup>Adding such an  $x_0$  and  $w_0$  has the effect of modifying each  $v_i$  to  $v_i - E[v_{max}]$ . The resulting distributions over  $w \cdot x + v_i$  are all the same as before, but now  $\mathcal{D}_{max}$  has a zero mean value.

since we have assumed all valuations are in  $[0, 1]$ , so we can restrict to  $\tilde{w}$  such that  $\tilde{w} \cdot x$  are all bounded), standard confidence bounds imply that minimizing mean squared error over a sufficiently (polynomially) large sample will achieve the desired near-optimal squared loss over the underlying distribution.  $\square$

## References

- Jeremy Bulow and Paul Klemperer. Auctions vs. negotiations. Technical report, National Bureau of Economic Research, 1994.
- Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. In *SODA'13*, pages 1190–1204, 2013.
- Shuchi Chawla, Jason Hartline, and Denis Nekipelov. Mechanism design for data science. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 711–712. ACM, 2014.
- Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *ACM Symp. on Theory of Computing*, 2014.
- Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 129–138. ACM, 2010.
- Antonia Foldes and Lidia Rejto. Strong uniform consistency for nonparametric survival curve estimators from randomly censored data. *The Annals of Statistics*, 9(1):122–129, 01 1981. doi: 10.1214/aos/1176345337. URL <http://dx.doi.org/10.1214/aos/1176345337>.
- Kuzman Ganchev, Yuriy Nevmyvaka, Michael Kearns, and Jennifer Wortman Vaughan. Censored exploration and the dark pool problem. *Communications of the ACM*, 53(5):99–107, 2010.
- Jason Hartline and Anna Karlin. Profit maximization in mechanism design. In *Algorithmic Game Theory*. Cambridge University Press, 2007.
- Edward L Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American statistical association*, 53(282):457–481, 1958.
- Robert D. Kleinberg and Frank Thomson Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*, pages 594–605, 2003.
- Isaac Meilijson. Estimating the lifetime distribution of the parts from the autopsy statistics of the machine. *Journal of Applied Probability*, 18(4):829–838, 1981.
- Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- Arthur Nadas. On estimating the distribution of a random vector when only the smallest coordinate is observable. *Technometrics*, 12(4):923–924, 1970.

## A Inequalities

**Lemma A.1.** *Suppose  $X$  is observable and  $Y$  is observable, and assume that  $\mathbb{P}[Y] \geq \gamma$ . Using  $2T$  samples, with probability  $1 - \delta$ , we can estimate  $\mathbb{P}[X|Y] = \frac{\mathbb{P}[X \cap Y]}{\mathbb{P}[Y]}$  by  $\hat{p}$  such that*

$$\mathbb{P}[X|Y] - \alpha - \mu \leq (1 - \mu)\mathbb{P}[X|Y] - \alpha \leq \hat{p} \leq (1 + \mu)\mathbb{P}[X|Y] + \alpha \leq \mathbb{P}[X|Y] + \alpha + \mu,$$

As a direct corollary, we know that `Inside` is a close approximation to the quantity it estimates.

**Corollary A.2.** *Inside*( $\ell_\tau, \ell_{\tau+1}, T$ ) outputs an estimator  $p_{\ell_\tau, \ell_{\tau+1}}^\epsilon$ , such that, for  $T$  as in Kaplan,

$$(1 - \mu)\mathbb{P}[\max_j b_j \geq \ell_\tau | \max_j b_j \leq \ell_{\tau+1}] - \alpha \leq p_{\ell_\tau, \ell_{\tau+1}}^\epsilon \leq (1 + \mu)\mathbb{P}[\max_j b_j \geq \ell_\tau | \max_j b_j \leq \ell_{\tau+1}] + \alpha$$

and uses  $2T$  samples.

Now, we prove Lemma 3.4, which is also a corollary of Lemma A.1.

*Proof of Lemma 3.4.* Let, for a fixed  $i, \ell_\tau, \ell_{\tau+1}$ , the event that  $i$  bids in  $[\ell_\tau, \ell_{\tau+1}]$  be denoted by  $X$ , the event that  $i$  wins in  $[\ell_\tau, \ell_{\tau+1}]$  be denoted by  $Y$ , and the event that  $\max_j b_j < \ell_{\tau+1}$  be denoted by  $C$ .

With this notation, we have an estimate of  $\mathbb{P}[Y|C]$  and want an estimate of  $\mathbb{P}[X|C]$ .

$$\begin{aligned} \mathbb{P}[Y|C] &= \mathbb{P}[X|C] \times \mathbb{P}[Y|C, X] \\ &\geq \mathbb{P}[X|C] \times \mathbb{P}[\text{everyone but } i \text{ bids } < \ell_\tau | C, X] \\ &= \mathbb{P}[X|C] \times \mathbb{P}[\text{everyone but } i \text{ bids } < \ell_\tau | C] \\ &\geq \mathbb{P}[X|C] \times (1 - \beta) \end{aligned}$$

The first equality comes from the fact that  $Y \subseteq X$ , the next inequality comes from the fact that, conditioned on  $C$  and  $X$ , everyone but  $i$  bids  $< \ell_\tau$  is a subset of  $Y$  (the times when  $i$  will win), the next equality comes from the fact that  $i$ 's bid and  $j$ 's bid are independent, and the final inequality follows from the assumption  $\mathbb{P}[\max_{j \neq i} b_j < \ell_\tau | \max_{j \neq i} b_j < \ell_{\tau+1}] \geq 1 - \beta$ .  $\square$

**Fact A.3.** *Suppose  $x \geq 0$  and  $0 < \eta < \frac{1}{2}$ . Then  $\frac{x}{1+\eta} \geq (1 - \eta)x$  and  $\frac{x}{1-\eta} \leq (1 + 2\eta)x$ .*

*Proof of Fact A.3.* We prove  $\frac{x}{1+\eta} \geq (1 - \eta)x$  first.

$$\frac{x}{1 + \eta} = \frac{(1 - \eta)x}{1 - \eta^2} \geq (1 - \eta)x \quad (\text{Since } 1 - \eta^2 < 1)$$

Now, we prove  $\frac{x}{1-\eta} \leq (1 + 2\eta)x$ , for  $\eta \leq 1/2$ . We have,

$$\frac{x}{1 - \eta} = x \sum_{i=0}^{\infty} \eta^i = x \left( 1 + \eta \left( \sum_{i=0}^{\infty} \eta^i \right) \right) \leq (1 + 2\eta)x,$$

where the inequality follows from the fact that for  $\eta \leq 1/2$  we have  $\sum_{i=0}^{\infty} \eta^i = \frac{1}{1-\eta} \leq 2$ .  $\square$

*Proof of Lemma 3.2.* We start by showing that, with no sampling error, the calculation  $p_{x,y}^i$  we do is equivalent to  $q_{x,y}^i = \mathbb{P}[b_i \in [x, y] \wedge b_i > \max_{j \neq i} b_j | \max_j b_j < y]$ . When  $x = y$ , we will denote this simply as  $q_x^i$  (similarly,  $p_x^i$ ). Similarly, let  $q_x^0$  denote the probability that no one wins when the reserve bidder is set to bid  $x$  (and  $p_x^0$  the empirical probability therein).

By definition,

$$\begin{aligned}
q_{\ell_\tau, \ell_{\tau+1}}^i &= \mathbb{P}[b_i \in [\ell_\tau, \ell_{\tau+1}] \wedge b_i > \max_{j \neq i} b_j \mid \max_j b_j < \ell_{\tau+1}] \\
&= \frac{\mathbb{P}[b_i \in [\ell_\tau, \ell_{\tau+1}] \wedge b_i > \max_{j \neq i} b_j \wedge \max_j b_j < \ell_{\tau+1}]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} \\
&= \frac{\mathbb{P}[b_i \in [\ell_\tau, \ell_{\tau+1}] \wedge b_i > \max_{j \neq i} b_j]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} && (i \text{ winning in } [\ell_\tau, \ell_{\tau+1}] \text{ implies } \max_j b_j < \ell_{\tau+1}) \\
&= \frac{\mathbb{P}[b_i \geq \ell_\tau \wedge b_i > \max_{j \neq i} b_j] - \mathbb{P}[b_i \geq \ell_{\tau+1} \wedge b_i > \max_{j \neq i} b_j]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} \\
&= \frac{\mathbb{P}[i \text{ wins with reserve } \ell_\tau] - \mathbb{P}[i \text{ wins with reserve } \ell_{\tau+1}]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} && (\text{Assuming no point masses, there are no ties}) \\
&= \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i}{q_{\ell_{\tau+1}}^0}
\end{aligned}$$

The final form is identical to the estimated quantity used by **IWin**. It now suffices to now show that each of the three samples give us good estimates of their respective true probabilities. A basic Chernoff bound implies

$$\mathbb{P}[|p_{x,1}^i - q_{x,1}^i| \geq \frac{\alpha\gamma(1-\mu)}{4}] \leq 2e^{-T\frac{1}{8}t_1\alpha^2\gamma^2(1-\mu)^2}.$$

Substituting  $T = \frac{8 \ln 6/\delta'}{\alpha^2\gamma^2(\frac{\mu}{2})^2}$ , and noting  $\mu < 1 - \mu$ , we have

$$\mathbb{P}[|p_{x,1}^i - q_{x,1}^i| \geq \frac{\alpha\gamma(1-\mu)}{4}] \leq \delta'$$

for each of  $x = \ell_\tau, \ell_{\tau+1}$ . Similarly,

$$\mathbb{P}[|p_x^0 - q_x^0| > \frac{\mu\gamma}{2}] \leq 2e^{-\frac{T}{2}\mu^2\gamma^2}$$

and substituting for  $T$ , we have that  $|p_{\ell_{\tau+1}}^0 - q_{\ell_{\tau+1}}^0| \geq \frac{\mu\gamma}{2}$  with probability at most  $\delta'$ . Thus, using a union bound, we have that with probability at least  $1 - 3\delta'$ , for a particular  $t$ ,

$$\frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \frac{\alpha\gamma(1-\mu)}{2}}{q_{\ell_{\tau+1}}^0 + \frac{\mu\gamma}{2}} \leq \frac{p_{\ell_\tau, 1}^i - p_{\ell_{\tau+1}, 1}^i}{p_{\ell_{\tau+1}}^0} \leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + \frac{\alpha\gamma(1-\mu)}{2}}{q_{\ell_{\tau+1}}^0 - \frac{\mu\gamma}{2}} \quad (5)$$

Now, it suffices to show that Equation (5) implies the relative error stated previously. By assumption,  $p_{0, \ell_{\tau+1}}^i > \gamma$ . This implies that the probability everyone bids at most  $\ell_{\tau+1}$  is at least  $\gamma$  (for a winning bid of  $\ell_{\tau+1}$  to win, all bids must be at most  $\ell_{\tau+1}$ ), so

$$q_{\ell_{\tau+1}}^0 \geq \gamma. \quad (6)$$

Then,

$$\begin{aligned}
p_{\ell_\tau, \ell_{\tau+1}}^i &= \frac{p_{\ell_\tau, 1}^i - p_{\ell_{\tau+1}, 1}^i}{p_{\ell_{\tau+1}}^0} \\
&\geq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \frac{1}{2}\alpha\gamma(1-\mu)}{q_{\ell_{\tau+1}}^0 + \frac{\mu\gamma}{2}} \\
&\geq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \alpha\gamma}{q_{\ell_{\tau+1}}^0 + \mu\gamma} && \text{(Since } \frac{(1-\mu)}{2} < 1\text{)} \\
&\geq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \alpha\gamma}{q_{\ell_{\tau+1}}^0 + \mu q_{\ell_{\tau+1}}^0} && \text{(By Eq. (6))} \\
&= \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \alpha\gamma}{q_{\ell_{\tau+1}}^0(1+\mu)} \\
&= \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1+\mu} - \frac{\alpha\gamma}{q_{\ell_{\tau+1}}^0(1+\mu)} \\
&\geq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1+\mu} - \frac{\alpha}{(1+\mu)} && \text{(By Eq. (6))} \\
&\geq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1+\mu} - \alpha \\
&\geq (1-\mu)q_{\ell_\tau, \ell_{\tau+1}}^i - \alpha && \text{(By Fact A.3)}
\end{aligned}$$

Now, we prove the upper bound on our estimator.

$$\begin{aligned}
p_{\ell_\tau, \ell_{\tau+1}}^i &= \frac{p_{\ell_\tau, 1}^i - p_{\ell_{\tau+1}, 1}^i}{p_{\ell_{\tau+1}}^0} \\
&\leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + \frac{(1-\mu)}{2}\alpha\gamma}{q_{\ell_{\tau+1}}^0 - \frac{\mu\gamma}{2}} \\
&\leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + (1-\mu)\alpha\gamma}{q_{\ell_{\tau+1}}^0 - \frac{\mu\gamma}{2}} \\
&\leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + (1-\mu)\alpha\gamma}{q_{\ell_{\tau+1}}^0 - \frac{\mu q_{\ell_{\tau+1}, \ell_{\tau+1}}^0}{2}} && \text{(By Eq. (6))} \\
&\leq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1 - \frac{\mu}{2}} + \frac{(1-\mu)\alpha\gamma}{q_{\ell_{\tau+1}}^0(1 - \frac{\mu}{2})} \\
&\leq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1 - \frac{\mu}{2}} + \frac{\alpha\gamma}{q_{\ell_{\tau+1}}^0} \\
&\leq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1 - \frac{\mu}{2}} + \alpha && \text{(By Eq. (6))} \\
&\leq (1 + 2\frac{\mu}{2})q_{\ell_\tau, \ell_{\tau+1}}^i + \alpha && \text{(By Fact. A.3)} \\
&= (1 + \mu)q_{\ell_\tau, \ell_{\tau+1}}^i + \alpha
\end{aligned}$$

Thus, both the upper and lower bounds on the estimator hold with probability  $1 - \delta$ .  $\square$

*Proof of Lemma 3.3.* We will show each of the three parts to be true.

1. We start by proving that **Intervals** will output a partition with at most  $\frac{24nL}{\beta\gamma}$  intervals. We claim that each interval is at least  $\frac{\beta\gamma}{24nL}$  in length, implying the above bound on the total number of intervals. Consider some current upper bound for an interval  $\ell_{\tau+1}$ . If **Intervals** accepts some point  $\ell_\tau$  such that  $\ell_{\tau+1} - \ell_\tau \geq \frac{\beta\gamma}{24nL}$ , then the bound trivially holds.

If this does not hold, **Intervals** tests some point  $\widehat{\ell}_\tau$  such that

$$\frac{\beta\gamma}{24nL} \geq \ell_{\tau+1} - \widehat{\ell}_\tau \geq \frac{\beta\gamma}{48nL}$$

since it is doing binary search. We claim **Intervals** will accept  $\widehat{\ell}_\tau$ ; if this is the case, the interval will have length at least  $\frac{\beta\gamma}{48nL}$ . Notice that

$$\mathbb{P}[\max_j b_j \in [\widehat{\ell}_\tau, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}] \leq \mathbb{P}[\max_j b_j \in [\ell_{\tau+1} - \frac{\beta\gamma}{24nL}, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}]$$

so it will suffice to show that **Intervals** would accept the smallest possible value of  $\widehat{\ell}_\tau$  (since that region will have the most probability mass). We bound the ratio, for a given  $\ell_{\tau+1}$  such that

$$\mathbb{P}[\max_j b_j \in [\ell_{\tau+1} - \frac{\beta\gamma}{24nL}, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}] = \frac{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}]}{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}]}$$

for some upper point of an interval  $\ell_{\tau+1}$  such that  $\mathbb{P}[i \text{ wins with a bid } \leq \ell_{\tau+1}] \geq \gamma$ . Since  $F_j$  is  $L$ -Lipschitz for all  $j$ ,

$$\mathbb{P}[b_j \leq \ell_{\tau+1}] - \mathbb{P}[b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}] \leq L \frac{\beta\gamma}{24nL} = \frac{\beta\gamma}{24n}.$$

Then, by summing this probability over all  $n$  bidders, we have

$$\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}] - \mathbb{P}[\max_j b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}] \leq \frac{\beta\gamma}{24}.$$

Rearranging terms, we have

$$\frac{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}]}{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}]} \geq 1 - \frac{\beta'\gamma}{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}]} \geq 1 - \frac{\beta}{24}$$

where the last inequality came from the fact that  $\mathbb{P}[i \text{ wins with a bid } \leq \ell_{\tau+1}] \geq \mathbb{P}[\max_j b_j \leq \ell_{\tau+1}] \geq \gamma$ . So, **Intervals** will accept  $\widehat{\ell}_\tau$  as  $\ell_\tau$ , so long as the empirical estimate of **Inside** is correct up to  $\alpha + \mu = \frac{\beta}{48}$ , which is the case by Corollary A.2 with probability  $1 - 3\delta'$ .

2. We now need to show

$$\mathbb{P}[\max_j b_j \geq \ell_{\tau-1} | \max_j b_j \leq \ell_\tau] \leq \frac{\beta}{16}$$

holds for the lattice points  $t > 3$ . Since  $\mathbb{P}[\max_j b_j \leq \ell_3] \geq \gamma$ , by Corollary A.2, the accuracy guarantee holds with probability  $1 - 3\delta'$  for a fixed  $t$  (since  $\alpha = \frac{\beta^2}{96}$ ,  $\mu = \frac{\beta}{96}$ , and the condition by which  $\ell_{\tau-1}$  was accepted was that the empirical estimate of the above quantity was at most  $\frac{\beta}{24}$ ). Thus, with probability  $1 - 3k\delta'$ , the above holds for all  $t > 3$ .

3. We begin by showing  $\mathbb{P}[\max_j b_j \leq \ell_2] \leq \gamma$  with probability at least  $1 - \delta'$ . The condition for stopping the search for new interval points is

$$J = \frac{\sum_{t \in S_1} \mathbb{I}[i \text{ wins on sample } t]}{T} \leq \frac{\gamma}{2}$$

where  $S_1$  is a random sample of size  $T$  with reserve  $\ell_1$ . A basic Chernoff bound shows that

$$\mathbb{P}[|J - \mathbb{P}[\max_j b_j \leq \ell_1]| \geq \frac{\gamma}{2}] \leq 2e^{-\frac{T\gamma^2}{2}}$$

which, for  $T = \frac{32 \ln \frac{6}{\delta'}}{\alpha^2 \gamma^2 \mu^2}$  is at most  $\delta'$ , so  $\mathbb{P}_{S_1}[\mathbb{P}[\max_j b_j \leq \ell_2] \leq \gamma] \geq 1 - \delta'$ , as desired.

It remains to sum up the total error probability and sample complexity. The lower bound on the length of each interval also implies a bound on the total number of empirical estimates made to find a fixed  $\ell_\tau$ . Formally, the halving algorithm beginning with a search space of size  $\ell_{\tau+1} \leq 1$  will halt before the remaining search space has shrunk to  $\frac{\beta\gamma}{48Ln}$ , which will take at most  $\log \frac{48Ln}{\beta\gamma} = \log(k)$  attempted interval endpoints per accepted interval endpoint. Each of these attempts calls `Inside`, which takes 2 estimates. For each accepted interval, an estimate of the remaining probability mass is done. Thus, in total, there are  $2k \log(k) + k$  estimates done by `Intervals`. Each fails with probability at most  $\delta'$ , so `Intervals` succeeds with probability at least  $1 - 3k \log(k) \delta'$  and uses at most  $3k \log(k) T$  samples.  $\square$