Learning Valuation Distributions from Partial Observations

Avrim Blum*
Computer Science Department,
Carnegie Mellon University
avrim@cs.cmu.edu

Yishay Mansour†
Tel Aviv University
Microsoft Research, Hertzelia
mansour@tau.ac.il

Jamie Morgenstern‡
Computer Science Department
Carnegie Mellon University
jamiemmt@cs.cmu.edu

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 work (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.

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

*Supported in part by the National Science Foundation under grants CCF-1101215, CCF-1116892, CCF-1331175, and IIS-1065251.

†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).

‡Supported in part by the National Science Foundation under grants CCF-1116892, CCF-1331175 and IIS-1065251. Partially funded by a Simons Award for Graduate Students in Theoretical Computer Science.

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.[1] 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 $\widehat{D}_i$ 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 $\widehat{D}_i$ of each of the distributions $\mathcal{D}_i$, down to some price p_γ, below which there is at most γ probability of any bidder winning.[2]

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. Those attributes may include geographic location, language, operating system, browser, as well as highly sensitive data that might be collected though 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 $\widehat{w}_i$ 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 $\widehat{w}$ 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.

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 (Foldes and Rejto 1981), see also its use for a control problem (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 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, Gentile, and Mansour 2013). In (Cesa-Bianchi, Gentile, and Mansour 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

[1] 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*.

[2] 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_γ such that there is at least a γ probability of the winning bid being below p_γ.

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, Roughgarden, and Yan 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 (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 $\left(\frac{1}{\epsilon}, n\right)$ samples from each bidder's distribution to construct an $1-\epsilon$-revenue-optimal auction for bid distributions that are strongly regular. We stress that in this work we make no such assumptions, only that the distributions are continuous.

Chawla et al. (Chawla, Hartline, and Nekipelov 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, our methodology handles *any* continuous bid distribution.[3]

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 $y^t = \text{argmax}_i v_i^t$, 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 ϵ 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 $y^t = \text{argmax}_i v_i^t$ ($y^t = 0$ 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 $m \in poly(n, \frac{1}{\epsilon}, \frac{1}{\gamma})$, 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.

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 their results into the terminology we use for auction setting, from the survival rate literature. Suppose each sample t is of the following form. Each bidder i draws their bid $b_i^t \sim \mathcal{D}_i$ independently of each other bid. The label $y^t = (\max_i b_i^t, \text{argmax}_i b_i^t)$ consists of the winning bid and the identity of the winner. From this, we would like to reconstruct an estimate $\widehat{F}_i$ 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: $F_i(x) = \mathbb{P}[b_i \leq x] = \mathbb{P}[b_i \leq x | b_i \leq y] \cdot \mathbb{P}[b_i \leq y]$ for $y > x$. Repeatedly applying this, we can see that, for $x < y_1 < y_2 < \cdots < y_r$,

$$F_i(x) = \mathbb{P}[b_i \leq x | b_i \leq y_1] \, \mathbb{P}[b_i \leq y_r] \prod_{t=1}^{r-1} \mathbb{P}[b_i \leq y_t | b_i \leq y_{t+1}] \tag{1}$$

Now, we can employ the observation in Equation 1, with estimates of such conditional probabilities. Since other players' bids are independent, we can estimate the conditional probabilities as follows:

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

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

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

Our estimator uses a similar Bayes-rule product expansion as KM, though it differs in several important ways. First,

[3]Note that we measure the distance between two distributions using the total variation distance, which is essentially "additive".

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 γ of winning at or below p_i, up to additive accuracy ϵ. 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 β 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.

Theorem 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* $\mathbb{P}[\exists\ j s.t.\ j \text{ wins with a bid } \leq p] \geq \gamma$*, 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 proofs can be found in the full version of this paper). `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 (ℓ_2 is close to p_γ, so we need not get a good estimation in in $[0, \ell_2]$, and by definition $\ell_1 = 0$).

We now present the lemmas which make up the crux of the proof of Theorem 1. Lemma 2 bounds the number of samples `IWin` uses and bounds the error of its estimate. Lemma 3 does similarly for `Intervals`. Lemma 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.

Algorithm `Kaplan`: Estimates the CDF of i from samples with reserves

---

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

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

Algorithm `IWin`: Est. $\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] | \max_j b_j < \ell_{\tau+1}]$

Data: $i, \ell_\tau, \ell_{\tau+1}, T$
Result: $p^i_{\ell_\tau, \ell_{\tau+1}}$

Let S_{ℓ_τ} be a sample with reserve $\ell_{\tau+1}$ of size T;
Let $S_{\ell_{\tau+1}}$ be a sample with reserve ℓ_τ of size T;
Let S_{cond} be a sample with reserve $\ell_{\tau+1}$ of size T;
Output $p^i_{\ell_\tau, \ell_{\tau+1}} = \frac{\sum_{t \in S_{\ell_\tau}} \mathbb{I}[i \text{ wins on sample } t] - \sum_{t \in S_{\ell_{\tau+1}}} \mathbb{I}[i \text{ wins on sample } t]}{\sum_{t \in S_{cond}} \mathbb{I}[0 \text{ wins on sample } t]}$;

Lemma 2. *Suppose, for a fixed interval* $[\ell_\tau, \ell_{\tau+1}]$, $\mathbb{P}[i \text{ wins in } [0, \ell_{\tau+1}]] \geq \gamma$*. Then, with probability at least* $1 - 3\delta'$, `IWin` $(i, \ell_\tau, \ell_{\tau+1}, T)$ *outputs* $p^i_{\ell_\tau, \ell_{\tau+1}}$ *such that*

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

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

and uses $3T$ *samples, for the values of* T, δ' *as in* `Kaplan`.

Algorithm `Intervals`: Partitions bid space to est. f_i

Data: β, γ, T, n, L
Result: $0 = \ell_1 < \ldots < \ell_k = 1$
Let $\ell_k = 1$, $c = k$, $p^i_{\ell_c} = 1$;
while $p^i_{\ell_c} > \gamma/2$ **do** `// Do binary search for the bottom of the next interval`
  Let $\widehat{\ell_b} = 0$;
  while `Inside`$(\widehat{\ell_b}, \ell_c, T) > \frac{\beta}{48}$ **do** `// The interval is too large`
    $\widehat{\ell_b} = \frac{\ell_c + \widehat{\ell_b}}{2}$;
  $\ell_{c-1} = \widehat{\ell_b}$;
  $c = c - 1$;
  Let S_1 be a sample of size T with reserve ℓ_{c-1};
  $p_{\ell_c} = \frac{\sum_{t \in S_1} \mathbb{I}[j \geq 1 \text{ wins on sample } t]}{T}$;
Return $0, \ell_c, \ldots, \ell_k$;

Lemma 3. *Let* T *as in* `Kaplan`. *Then,* `Intervals`(β, γ, T, L, n) *returns* $0 = \ell_1 < \cdots < \ell_k = 1$ *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, we know that the partition of $[0, 1]$ returned by `Intervals` is "fine enough". Now, Lemma 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 4. *Suppose that, for bidder* i *and some* $0 \leq \ell_\tau \leq \ell_{\tau+1} \leq 1$,

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

Then,

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

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

Observation 5. *Consider some set of points* $0 < \ell_1 < \ldots < \ell_k = 1$. $F_i(\ell_\tau)$ *can be rewritten as the following product:*

$$F_i(\ell_{\tau-1}) = \prod_{\tau' \geq t} (1 - \mathbb{P}[b_i \in [\ell_{\tau'-1}, \ell_{\tau'}] | b_i \leq \ell_{\tau'}])$$

We relegate the formal proof of Theorem 1 to the full version for reasons of space. We give some intuition for the proof here. With probability $1 - \delta$, `Intervals` returns a good partition, and, for each interval, of which there are at most $k' - 1$, by Lemma 3, `IWin` is as accurate as described by Lemma 2) (which follows from a union bound). Thus, for the remainder of the proof we assume the partition returned by `Intervals` is good and each call to `IWin` is accurate. Then, by Lemma 4, the probability that a bidder *wins* with a bid in an interval is a close approximation to the probability she *bids* in that interval (both events are conditioned on all bids being at most the upper bound of the interval). These estimates multiplied together also give good estimates.

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 `Kaplan` 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 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 or cars, 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 $\text{argmax}_i w_i \cdot x$. 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 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. our goal is to learn both the common vector w and each $\mathcal{D}_i$.

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 $[0, 1]$.

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

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

produces $\tilde{w}_i$ such that on a $1 - \epsilon$ probability mass of $x \sim \mathcal{P}$, $i^* \equiv argmax_i \tilde{w}_i \cdot x = argmax_i w_i \cdot x$ (i.e., a correct prediction of the winner), and furthermore

$$|\tilde{w}_{i^*} \cdot x - w_{i^*} \cdot x| \leq \epsilon.$$

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

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

Similarly, if we win the item, we have:

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

Let $\mathcal{P}^*$ denote the distribution over pairs (x, b) induced by drawing x from $\mathcal{P}$ and b uniformly at random from $\{0, \epsilon, 2\epsilon, \ldots, 1\}$ and consider a $(k+1)$-valued target function f^* that given a pair (x, b) outputs an integer in $\{0, 1, \ldots, n\}$ indicating the winner (with 0 indicating that our bid b wins). By design, the vectors $\tilde{w}_1, \ldots, \tilde{w}_n$ 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 $b \in \{0, \epsilon, 2\epsilon, \ldots, 1\}$ (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 $\tilde{w}_1, \ldots, \tilde{w}_n$ induces a $(n+1)$-way partition of the $(d+1)$-dimensional space of pairs (x, b) based on which of $\{0, \ldots, n\}$ 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 $\epsilon^2/(1+\epsilon)$. $\square$

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 for item x is $w \cdot x + v_i$ where $v_i \sim \mathcal{D}_i$. As before, we assume $\|x\|_2 \leq 1$ 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 each $\mathcal{D}_i$. We now show 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 ϵ', we 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 ϵ' fraction of the time, there is a large fluctuation). But since we can make ϵ' 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 7. *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 $\max[v_1, ..., v_n]$. 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.[4]

[4]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.

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 ϵ'^2 on at least an ϵ' probability mass of the points. Finally, since all losses are bounded (we know all values $w \cdot x$ are bounded 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. □

References

Bulow, J., and Klemperer, P. 1994. Auctions vs. negotiations. Technical report, National Bureau of Economic Research.

Cesa-Bianchi, N.; Gentile, C.; and Mansour, Y. 2013. Regret minimization for reserve prices in second-price auctions. In *SODA'13*, 1190–1204.

Chawla, S.; Hartline, J.; and Nekipelov, D. 2014. Mechanism design for data science. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 711–712. ACM.

Cole, R., and Roughgarden, T. 2014. The sample complexity of revenue maximization. In *ACM Symp. on Theory of Computing*.

Dhangwatnotai, P.; Roughgarden, T.; and Yan, Q. 2010. Revenue maximization with a single sample. In *Proceedings of the 11th ACM conference on Electronic commerce*, 129–138. ACM.

Foldes, A., and Rejto, L. 1981. Strong uniform consistency for nonparametric survival curve estimators from randomly censored data. *The Annals of Statistics* 9(1):122–129.

Ganchev, K.; Nevmyvaka, Y.; Kearns, M.; and Vaughan, J. W. 2010. Censored exploration and the dark pool problem. *Communications of the ACM* 53(5):99–107.

Hartline, J., and Karlin, A. 2007. Profit maximization in mechanism design. In *Algorithmic Game Theory*. Cambridge University Press.

Kaplan, E. L., and Meier, P. 1958. Nonparametric estimation from incomplete observations. *Journal of the American statistical association* 53(282):457–481.

Kleinberg, R. D., and Leighton, F. T. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*, 594–605.

Meilijson, I. 1981. Estimating the lifetime distribution of the parts from the autopsy statistics of the machine. *Journal of Applied Probability* 18(4):829–838.

Myerson, R. B. 1981. Optimal auction design. *Mathematics of operations research* 6(1):58–73.

Nadas, A. 1970. On estimating the distribution of a random vector when only the smallest coordinate is observable. *Technometrics* 12(4):923–924.