Learning Auctions with Robust Incentive Guarantees

Bhuvesh Kumar

Joint work with Jacob Abernerthy, Rachel Cummings, Jamie Morgenstern, and Samuel Taggart

December 5, 2019

CS PhD Qualifier Talk, School of Computer Science, Georgia Tech

Table of contents

Introduction

Auctions

A seller is selling J identical objects and n buyers participate in the auction. In an auction mechanism $\mathcal{M}$

- Each buyer i samples values $v_i \sim D_i$

Auctions

A seller is selling J identical objects and n buyers participate in the auction. In an auction mechanism $\mathcal{M}$

- Each buyer i samples values $v_i \sim D_i$
- Each buyer i reports bid b_i

Auctions

A seller is selling J identical objects and n buyers participate in the auction. In an auction mechanism $\mathcal{M}$

- Each buyer i samples values $v_i \sim D_i$
- Each buyer i reports bid b_i
- Seller decides allocation $\mathbf{x} \in \{0, 1\}^n$, s.t. $\|x\|_1 \leq J$

Auctions

A seller is selling J identical objects and n buyers participate in the auction. In an auction mechanism $\mathcal{M}$

- Each buyer i samples values $v_i \sim D_i$
- Each buyer i reports bid b_i
- Seller decides allocation $\mathbf{x} \in \{0,1\}^n$, s.t. $\|x\|_1 \leq J$
- Seller charges each buyer i payment p_i

Auctions

A seller is selling J identical objects and n buyers participate in the auction. In an auction mechanism $\mathcal{M}$

- Each buyer i samples values $v_i \sim D_i$
- Each buyer i reports bid b_i
- Seller decides allocation $\mathbf{x} \in \{0,1\}^n$, s.t. $\|x\|_1 \leq J$
- Seller charges each buyer i payment p_i

Expected Revenue of $\mathcal{M}$ on $\mathbf{D} = D_1 \times D_2 \times \cdots \times D_n$

$$\mathrm{Rev}(\mathcal{M}; \mathbf{D}) := \mathbf{E}_{\mathbf{v} \sim \mathbf{D}} \left[\sum_{i=1}^{n} p_i(\mathbf{v}) \right]$$

Revenue Maximizing Auction

Myerson's Auction [19] characterizes the revenue maximizing auction

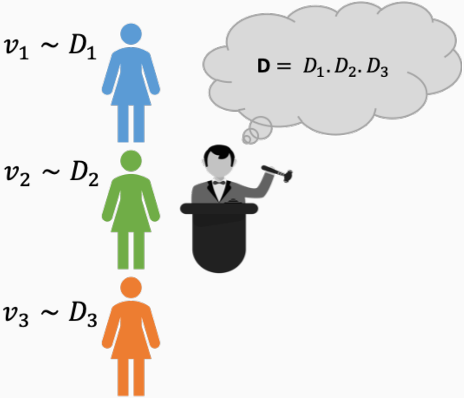

Figure 1: Seller knows the value distribution

Revenue Maximizing Auction

Fixing a distribution $\mathbf{D}$, given a value profile $\mathbf{v}$ Myerson's auction calculates virtual values $\phi_i(v_i) = v_i - \frac{1-F_{D_i}(v_i)}{f_{D_i}(v_i)}$, and decides allocation $\mathbf{x}(\mathbf{v})$

$$\underset{\mathbf{x}}{\operatorname{argmax}} \mathbb{E}_{\mathbf{v}\sim\mathbf{D}}[\sum_{i=1}^{n} \phi_i(v_i)x_i(\mathbf{v})]$$

Let $\phi_i(v_i)$ be sorted in decreasing order, then i is allocated if $i \leq J$ and $\phi_i(v_i) > 0$. The payment charged is

$$\max\left\{\phi_i^{-1}(\phi_{J+1}(v_{J+1})), \phi_i^{-1}(0)\right\}$$

We refer to the expected revenue obtained by Myerson's as OPT

Learning from samples

Maximize revenue when that distribution is unknown but some small number of samples are available [9, 4, 2, 6, 17, 3, 18, 11, 5, 10]

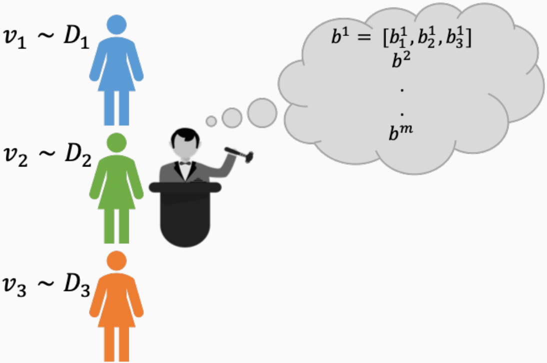

Figure 2: Seller has samples

Iterative learning setting

We consider an iterative setting similar to Liu et al. [14]

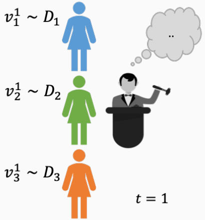

Iterative learning setting

We consider an iterative setting similar to Liu et al. [14]

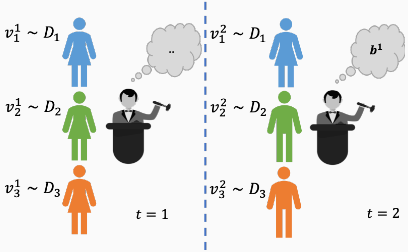

Iterative learning setting

We consider an iterative setting similar to Liu et al. [14]

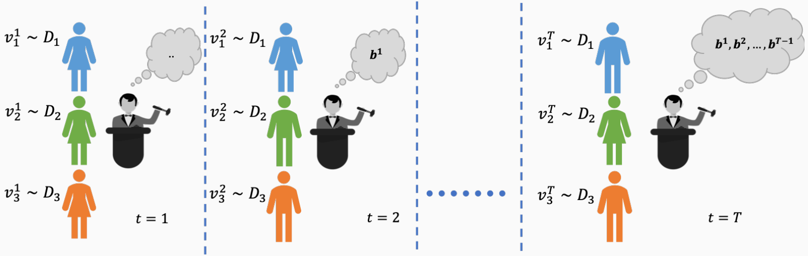

Challenge

- Seller wants to maximize revenue and thus wants to learn from the buyers

Challenge

- Seller wants to maximize revenue and thus wants to learn from the buyers
- Buyer wants to maximize their total utility over all the rounds they participate in the auctions

Challenge

- Seller wants to maximize revenue and thus wants to learn from the buyers
- Buyer wants to maximize their total utility over all the rounds they participate in the auctions
- Buyers can misreport to gain more utility in the future.

Reserve Prices vs Optimal Auctions

Liu et al. [14] consider only reserve price auctions, and in general reserve price revenue cannot approximate the Myerson's Revenue. We compete with Myerson's revenue.

Large Market Assumption

Similar to [14], we assume that each buyer participates in at most k rounds.

Differential Privacy

Differential Privacy

Bound the “maximum amount” that one person’s data can change the output of a computation.

Differential Privacy Dwork et al. [7]

An algorithm $\mathcal{A} : \mathcal{Z}^n \to \mathcal{R}$ is (ϵ, δ)-*differentially private* if for every pair of neighboring databases $Z, Z' \in \mathcal{Z}^n$, and for every subset of possible outputs $\mathcal{S} \subseteq \mathcal{R}$,

$$\mathrm{P}[\mathcal{A}(Z) \in \mathcal{S}] \leq \exp(\epsilon)\, \mathrm{P}[\mathcal{A}(Z') \in \mathcal{S}] + \delta.$$

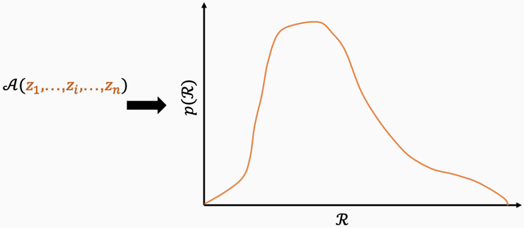

Differential Privacy

Bound the "maximum amount" that one person's data can change the output of a computation.

Differential Privacy Dwork et al. [7]

An algorithm $\mathcal{A} : \mathcal{Z}^n \to \mathcal{R}$ is (ϵ, δ)-*differentially private* if for every pair of neighboring databases $Z, Z' \in \mathcal{Z}^n$, and for every subset of possible outputs $\mathcal{S} \subseteq \mathcal{R}$,

$$\mathrm{P}[\mathcal{A}(Z) \in \mathcal{S}] \leq \exp(\epsilon)\,\mathrm{P}[\mathcal{A}(Z') \in \mathcal{S}] + \delta.$$

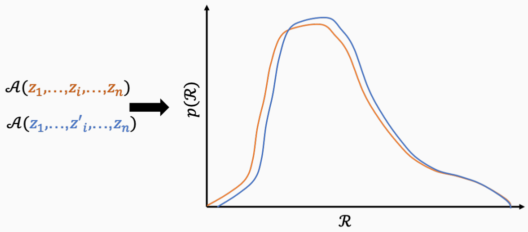

Differential Privacy

Tools we use from differential privacy:

- Tree aggregation

Differential Privacy

Tools we use from differential privacy:

- Tree aggregation
- Exponential Mechanism

Private Partial Sums

Let $Z = \langle w_1, w_2, \cdots, w_T \rangle$ be a set of vectors we want to calculate $S_t = \sum_{\tau=1}^{t} w_\tau$ for all t privately.

Naive Approach

Add noise to each partial sum S_t, as each vector appears in $O(T)$ partial sums, to get (ϵ, δ)-differential privacy over the stream of partial sums we need to add $\tilde{O}(\frac{\sqrt{T}}{\epsilon})$ noise to each S_t

Tree based aggregation (Jain et al. [12], Dwork et al. [8])

Using a binary index tree, we can calculate the stream of partial sums (ϵ, δ)-differentially privately by adding only $O(\frac{\log T}{\epsilon})$ noise

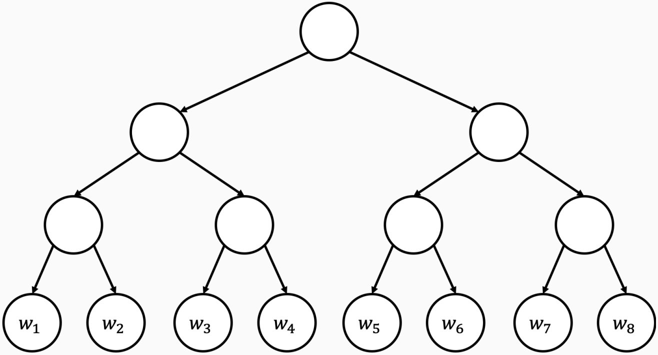

Tree based aggregation (Jain et al. [12], Dwork et al. [8])

Using a binary index tree, we can calculate the stream of partial sums (ϵ, δ)-differentially privately by adding only $O(\frac{\log T}{\epsilon})$ noise

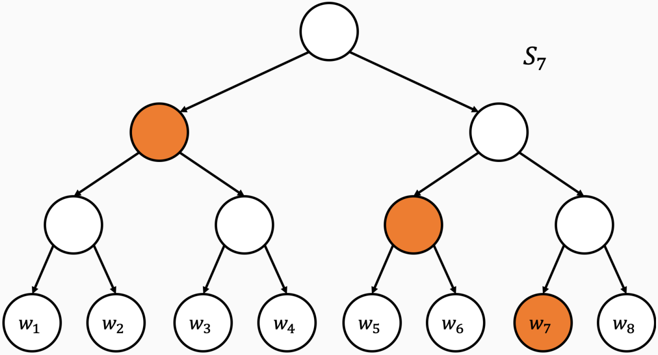

Tree based aggregation (Jain et al. [12], Dwork et al. [8])

Using a binary index tree, we can calculate the stream of partial sums (ϵ, δ)-differentially privately by adding only $O(\frac{\log T}{\epsilon})$ noise

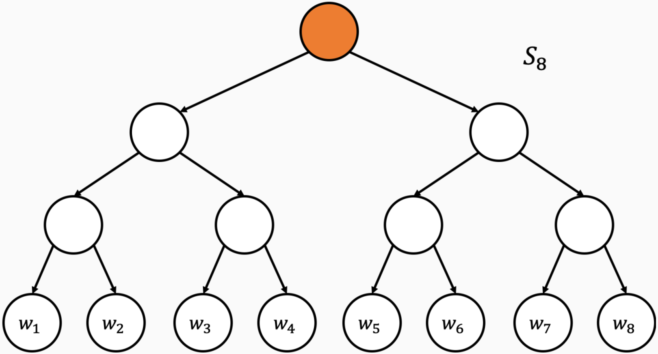

Utility-Approximate Bayesian Incentive Compatible Online Auction

Utility-Approximate BIC Online Auction

Algorithm 1: Utility-Approximate BIC Online Auction

Parameters: discretization β, privacy ϵ, upper bound on support h, num. of rounds T

Initialize: $H'_{i,0} \leftarrow \text{Uniform}(0, h)$ for $i = 1, \cdots, n$

for $t = 1, \cdots, T$ **do**

   Receive bid profile $\mathbf{v}_t = (v_{1,t}, \ldots, v_{n,t})$, rounded down to integer multiple of β

   Run Myerson with $\mathbf{H'}_{t-1}$ as prior and $\mathbf{v}_t$ as bid for allocations/payments.

   for $i = 1, \ldots, n$ **do**

      Update $H'_{i,t}$ via two-fold tree aggregation, giving as input $v_{i,t}$

   end

end

Two Tree aggregation

To compute $H_{i,t}(u)$, the empirical CDF at value u, we need only count the number of samples from $v_{i,1}, \cdots, v_{i,t}$ which are less than u, i.e.,

$$H_{i,t}(u) = \frac{\sum_{\tau=1}^{t} \mathbf{1}\{v_{i,\tau} \leq u\}}{t} = \frac{\sum_{\tau=1}^{t} \sum_{z=1}^{u/\beta} \mathbf{1}\{\beta(z-1) < v_{i,\tau} \leq \beta z\}}{t}$$

Two-fold tree aggregation [14] allows us to privately maintain these cumulative sums for all points $u \in \{\beta, 2\beta, \cdots, h\}$ in the support of our distribution by summing over the two axes: u and t which effectively results in adding Gaussian noise with $\sigma = \tilde{O}(\frac{\log T \log h/\beta}{t\epsilon})$ to $H_{i,t}(u)$

Differential Privacy Guarantee

Theorem

The stream of estimates $\{\mathbf{H}'_t\}_{t=1}^{T}$ *maintained by Algorithm 1 is* $(\epsilon, \epsilon/T)$*-differentially private with respect to the stream of input bids* $\{v_t\}_{t=1}^{T}$.

Our algorithm is not differentially private in its selection of allocations and payments in round t. However, the information the mechanism carries forward (namely, the estimated empirical distribution) is maintained in a differentially private manner.

Incentive Guarantee

Definition (η-utility-approximate BIC)

A mechanism is η *utility-approximately Bayesian incentive compatible* if the strategy profile where every agent bids truthfully in every history is an η-approximate Perfect Bayesian equilibrium.

Incentive Guarantee

Using the fact that we use Myerson's allocation and payment rule in each fixed round t and the privacy of our estimates, we can show that:

Theorem (Incentive Guarantee for Algorithm 1)

Algorithm 1 is $kh\epsilon\left(2 + \frac{1}{T}\right)$-utility approximate BIC when $\epsilon < 1$.

Revenue Analysis

To analyse the expected revenue we show three parts

1. If two value distributions are close then any reasonable mechanism achieves similar revenue on the two distributions

Revenue Analysis

To analyse the expected revenue we show three parts

1. If two value distributions are close then any reasonable mechanism achieves similar revenue on the two distributions
2. The differentially private estimates maintained by Algorithm 1 are close to the true value distributions

Revenue Analysis

To analyse the expected revenue we show three parts

1. If two value distributions are close then any reasonable mechanism achieves similar revenue on the two distributions
2. The differentially private estimates maintained by Algorithm 1 are close to the true value distributions
3. The loss in revenue due to discretization is bounded

Similar Revenue on Similar Distributions

Theorem (Similar revenue on similar Distributions)

Let $\mathcal{M}$ be a competitive mechanism, and let $\mathbf{D}$ and $\tilde{\mathbf{D}}$ be two product distributions of values such that for every bidder i, D_i and $\tilde{D}_i$ are τ-close, i.e $\|D_i - \tilde{D}_i\|_\infty \leq \tau$. Then the expected revenue of $\mathcal{M}$ on $\mathbf{D}$ is within an additive $2n^2 h\tau$ of the revenue from $\mathcal{M}$ on $\tilde{\mathbf{D}}$. That is, $|Rev(\mathcal{M}; \mathbf{D}) - Rev(\mathcal{M}; \tilde{\mathbf{D}})| \leq 2n^2 h\tau$.

Estimates are close

Using Dvoretzky-Kiefer-Wolfowitz (Massart [15]) inequality and the variance of our added Gaussian noise, we show that

Theorem

After t rounds Algorithm 1, it holds with probability at least $1-\alpha$ that

$$\left\| H'_{i,t} - D'_i \right\|_\infty \leq \gamma_t \quad \textit{for every } i \in [n],$$

where D'_i is D_i rounded down by β

$\gamma_t = \sqrt{\frac{\log \frac{n}{\alpha}}{2t}} + \frac{\sigma}{t}\sqrt{\log \frac{h}{\beta} \log T}\sqrt{2\log\left(\frac{2hn}{\beta\alpha}\right)}$ *and*

$\sigma = \frac{8 \log T \log \frac{h}{\beta}}{\epsilon}\sqrt{\ln \frac{T \log T \log \frac{h}{\beta}}{\epsilon}}.$

Discretization doesn't hurt too much

Using results of Devanur et al. [5], we show that

Lemma ([5] Discretization loss)

As $\mathbf{D}'$ *is the rounding down of* $\mathbf{D}$ *to the closest multiple of* β, $Rev(\mathcal{M}^*_{\mathbf{D}'}; \mathbf{D}') \geq OPT - \beta J$ *where* $\mathcal{M}^*_{\mathbf{D}'}$ *is the optimal mechanism for* $\mathbf{D}'$.

Revenue Guarantee

Combining all the steps, we can show that

Theorem (Overall revenue Guarantee)

With probability at least $1-\alpha$, the average expected revenue obtained by Algorithm 1 for T rounds satisfies

$$Rev \geq OPT - \beta J - 4hn^2 \tilde{O}\left(\sqrt{\frac{\log(\frac{nT}{\alpha})}{T}} + \frac{1}{T\epsilon}\right)$$

for regular distributions $\mathbf{D}$ *and* $\epsilon < 1$.

Note that if we set $\beta = o(1)$ in term of T, we can achieve sublinear approximation.

Bid-Approximate Bayesian Incentive Compatible Online Auction

Introduction

In this section we introduce a stronger notion of Incentive compatibility called η-bid-approximate BIC

Definition (η-bid-approximate BIC)

A mechanism is η *bid-approximate BIC* if $\exists$ an exact PBE where each bidder bids within η of their value in every history.

Punishing Mechanism

Buyers may have incentive to misreport their bids to gain more utility, thus to ensure that there is a penalty for misreporting today that may discourage the buyers from misreporting, we use a punishing mechanism.

Algorithm 2: StrictlyTruthful

Input: Bid profile $\mathbf{b}_t$

Select a subset $S \subseteq [n]$ of size J uniformly at random

Select a price $p \in [0, h]$ uniformly at random

for *Each* $s \in S$ **do**

  if $b_{s,t} \geq p$ **then** allocate item to s and charge payment p ;

end

Punishing Mechanism

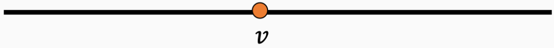

Punishing Mechanism

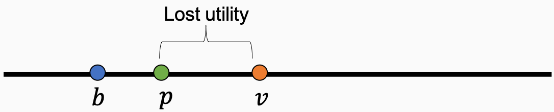

Punishing Mechanism

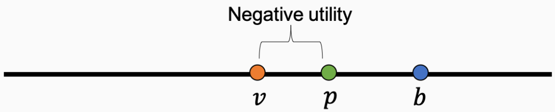

Bid-approximate BIC Online Auction

Algorithm 3: Bid-approximate BIC Online Auction

Parameters: discretization β, privacy ϵ, upper bound on support h, num. of rounds T

Initialize: $H'_{i,0} \leftarrow \text{Uniform}(0, h)$ for $i = 1, \cdots, n$

for $t = 1, \cdots, T$ **do**

- Receive vector of bids $\mathbf{b}_t = (b_{1,t}, \ldots, b_{n,t})$, rounded down to multiple of β
- With probability ρ_t run mechanism StrictlyTruthful($\mathbf{b}_t$)
- **else**
  - **for** $i = 1, \ldots, n$ **do**
    - Use $H'_{i,t-1}$ to calculate $\phi_{i,t}(b_{i,t})$
  - **end**
  - Use exponential mechanism (Algorithm to select allocation $x_t(\phi_{\mathbf{t}}(\mathbf{b}))$
  - Use *Black box payments* to calculate payments $p_t(\mathbf{b}_t)$.
- **end**
- **for** $i = 1, \ldots, n$ **do**
  - Update $H'_{i,t}$ via two-fold tree aggregation, giving as input $b_{i,t}$
- **end**

end

Exponential Mechanism

We recall that the allocations and payments in Algorithm 1 were not differentially private, to achieve the stronger notion of incentive compatibility, we need to add more privacy, thus we use exponential mechanism ([16]) for allocation.

Exponential Mechanism McSherry and Talwar [16]

Given a virtual value profile $\phi(\mathbf{b})$ the Exponential Mechanism $\mathcal{M}_E(\mathcal{X}, Q, \phi(\mathbf{b}), \epsilon)$ chooses an allocation $\mathbf{x} \in \mathcal{X}$ with probability

$$\mathrm{P}[\mathcal{M}_E(\mathcal{X}, Q, \phi(\mathbf{b}), \epsilon) = \mathbf{x}] \propto \exp \frac{\epsilon Q(\phi(\mathbf{b}), \mathbf{x})}{\Delta}$$

where $Q(\phi(\mathbf{b}), \mathbf{x})$ is called the *quality score* and we use $Q(\phi(\mathbf{b}), \mathbf{x}) = \sum_{i=0}^{n} \phi_i(b_i) x_i$.

Exponential Mechanism

Theorem

Exponential Mechanism $\mathcal{M}_E$ *is* ϵ*-differentially private and if* $\mathbf{x}$ *is sampled by* $\mathcal{M}_E$ *then w.p* $1-\beta$

$$Q(\phi(\mathbf{b}),\mathbf{x}) \geq \max_{\mathbf{x}'\in\mathcal{X}} Q(\phi(\mathbf{b}),\mathbf{x}') - O\left(\frac{\log(\frac{|\mathcal{X}|}{\beta})}{\epsilon}\right)$$

Black Box mechanism

Black box payment (Archer et al. [1]) calculates expected payments for buyer i only by looking at the bid for buyer i and accesses information about other buyers' bids only through the allocation rule.

Lemma (Black box payment)

Fix a distribution $\mathbf{H'}_t$. Then, allocating according to the exponential mechanism $\mathcal{M}_E$ and charging black-box payments yields expected revenue on $\mathbf{H'}_t$ equal to the expected virtual surplus of the allocation selected by $\mathcal{M}_E$.

Differential Privacy

Definition (Joint Differential Privacy [13])

An algorithm $\mathcal{A} : \mathcal{Z}^n \to \mathcal{R}^n$ is (ϵ, δ)-*jointly differentially private* if for every $i \in [n]$, for every pair of i-neighbors $Z, Z' \in \mathcal{Z}^n$, and for every subset $\mathcal{S} \subseteq \mathcal{R}^{n-1}$,

$$\mathrm{P}[\mathcal{A}(Z)_{-i} \in \mathcal{S}] \leq \exp(\epsilon)\,\mathrm{P}[\mathcal{A}(Z')_{-i} \in \mathcal{S}] + \delta.$$

If $\delta = 0$, we say that $\mathcal{A}$ is ϵ-*jointly differentially private*.

Lemma

Algorithm 3 is $(3\epsilon, 3\epsilon/T)$-jointly differentially private in the bids of bidders.

Incentive Guarantee

Using differential privacy, the guarantees from punishing mechanism, the approximation guarantees of exponential mechanism and our distribution estimation, we can show that

Theorem

Algorithm 3 in round t is η_t-bid approximate BIC, i.e. in round t, any bidder i with value $v_{i,t}$ reports bid $b_{i,t}$ which satisfies $v_{i,t} - \eta_t \leq b_{i,t} \leq v_{i,t}$ where $\eta_t = h\sqrt{\frac{8n^2\gamma_t + 6k\epsilon}{\rho_t J}}$,

$\gamma_t = \sqrt{\frac{\log(2hn/\beta\alpha)}{2t}} + o(t) + \frac{\sigma}{t}\sqrt{\log\frac{h}{\beta}\log T}\sqrt{\log\left(\frac{hn}{\beta\alpha}\right)}$,

$\sigma = \frac{8\log T\log\frac{h}{\beta}}{\epsilon}\sqrt{\ln\frac{\log T\log\frac{h}{\beta}}{\delta}}$ *and* $\delta = \frac{\epsilon}{T}$.

In simpler words, every buyer bids within $\eta_t = c + \tilde{O}(t^{-1/4})$ of the truthful bid for some small constant c.

Final Revenue

Using the bid approximate incentive compatibility result and the results from the previous section, we can show that

Theorem (Overall revenue Guarantee)

With probability at least $1 - \alpha$, the average expected revenue obtained by Algorithm 3 for T rounds satisfies

$$Rev \geq OPT - c' - \tilde{O}(T^{-1/4})$$

for some small constant c'.

Questions?

References

[1] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2): 129–150, 2004.

[2] Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. Mechanism design via machine learning. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 605–614. IEEE, 2005.

[3] Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of automated mechanism design. In *Advances in Neural Information Processing Systems*, pages 2083–2091, 2016.

References ii

[4] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 243–252. ACM, 2014.

[5] Nikhil R Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 426–439. ACM, 2016.

[6] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.

[8] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724. ACM, 2010.

[9] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, pages 736–745, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics. ISBN 978-0-898716-24-5. URL http://dl.acm.org/citation.cfm?id=1283383.1283462.

[10] Yannai A. Gonczarowski and Noam Nisan. Efficient empirical revenue maximization in single-parameter auction environments. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*, STOC '17, pages 856–868, 2017.

[11] Jason Hartline and Samuel Taggart. Non-revelation mechanism design. arXiV preprint 1608.01875, 2016.

[12] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Conference on Learning Theory*, pages 24–1, 2012.

[13] Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ITCS '14, pages 403–410. ACM, 2014.

[14] Jinyan Liu, Zhiyi Huang, and Xiangning Wang. Learning optimal reserve price against non-myopic bidders. In *Advances in Neural Information Processing Systems*, pages 2038–2048, 2018.

[15] Pascal Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, pages 1269–1283, 1990.

References v

[16] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007.

[17] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory*, pages 1298–1318, 2016.

[18] Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems*, pages 136–144, 2015.

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