Learning Optimal Auctions with Correlated Valuations from Samples

Chunxue Yang, Xiaohui Bei

School of Physical and Mathematical Sciences
Nanyang Technological University

June 2021

Auctions are everywhere and affect our everyday lives.

– Nobel Prize's committee

Sealed-Bid Auction Format

Sealed-Bid Auction Format

- There are several bidders buying an item.

Sealed-Bid Auction Format

- There are several bidders buying an item.
- Each bidder i has a type θ_i.

Sealed-Bid Auction Format

- There are several bidders buying an item.
- Each bidder i has a type θ_i.
- The bidders' types are drawn from a prior distribution π.

Sealed-Bid Auction Format

- There are several bidders buying an item.
- Each bidder i has a type θ_i.
- The bidders' types are drawn from a prior distribution π.
- A function $v : \Theta \rightarrow \mathbb{R}_+$ maps types to valuations.

Sealed-Bid Auction Format

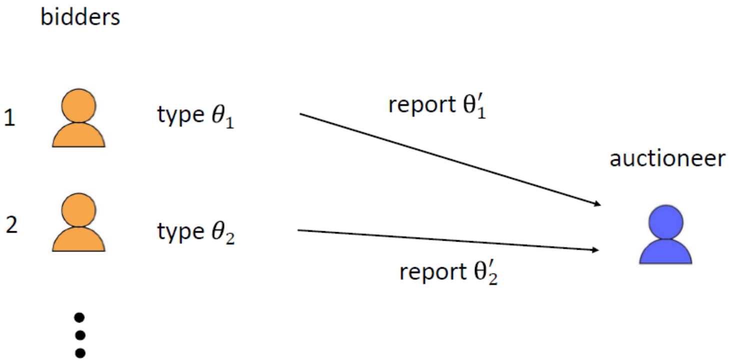

- **An auction**
  An auction consists of allocation rule x and payment rule p. The allocation rule is a function which decides which bidder wins the item. The payment rule specifies the amount of money each bidder needs to pay to the auctioneer.

- **An auction**
  An auction consists of allocation rule x and payment rule p. The allocation rule is a function which decides which bidder wins the item. The payment rule specifies the amount of money each bidder needs to pay to the auctioneer.

- **A bidder's utility**
  We assume quasi-linear utility. Bidder i's utility is given by

$$u_i(\theta_i, \theta_i', \boldsymbol{\theta}'_{-i}) = v(\theta_i)x(\theta_i', \boldsymbol{\theta}'_{-i}) - p(\theta_i', \boldsymbol{\theta}'_{-i}).$$

- **An auction**
  An auction consists of allocation rule x and payment rule p. The allocation rule is a function which decides which bidder wins the item. The payment rule specifies the amount of money each bidder needs to pay to the auctioneer.

- **A bidder's utility**
  We assume quasi-linear utility. Bidder i's utility is given by

$$u_i(\theta_i, \theta'_i, \boldsymbol{\theta}'_{-i}) = v(\theta_i)x(\theta'_i, \boldsymbol{\theta}'_{-i}) - p(\theta'_i, \boldsymbol{\theta}'_{-i}).$$

- **Our objective**
  The goal of the auction designer is to maximize the expected revenue the auction derives.

Myerson Auction

Myerson (1981) (Nobel Prize in Economic Sciences 2007)

Myerson Auction

Myerson (1981) (Nobel Prize in Economic Sciences 2007)

- assumes that the prior distributions of all bidders are
  - precisely known to the auctioneer
  - independent

Myerson Auction

Myerson (1981) (Nobel Prize in Economic Sciences 2007)

- assumes that the prior distributions of all bidders are
  - precisely known to the auctioneer
  - independent
- gives characterization of the revenue-maximizing auction
  - optimal auction with identical bidders:
    a second-price auction with a reserve price

- **Limitation**
  The assumption that priors are precisely known to the auctioneer is too strong. In most cases, the prior distributions can only be learned from past data.

- **Limitation**
  The assumption that priors are precisely known to the auctioneer is too strong. In most cases, the prior distributions can only be learned from past data.

- **Question**
  How much data is sufficient and necessary to guarantee near-optimal expected revenue?

Sample Complexity

Definition 1 (sample complexity).

The sample complexity is defined as the asymptotically smallest number m such that given m i.i.d. samples from a distribution π, there exists an algorithm that can learn an auction M achieving revenue

$$\mathrm{Rev}(M, \pi) \geq (1 - \epsilon)\mathrm{Opt}(\pi)$$

with high probability.

Sample Complexity of Myerson Auction

- First studied by Cole & Roughgarden (2014)

Sample Complexity of Myerson Auction

- First studied by Cole & Roughgarden (2014)
- A long line of follow-up works (Gonczarowski & Nisan, 2017; Syrgkanis, 2017; Huang et al., 2018)

Sample Complexity of Myerson Auction

- First studied by Cole & Roughgarden (2014)
- A long line of follow-up works (Gonczarowski & Nisan, 2017; Syrgkanis, 2017; Huang et al., 2018)
- Settled by Guo et al. (2019) with matching upper and lower bounds up to a poly-logarithmic factor

Correlated Valuations

- Independence is unlikely to hold in many cases:
  - A used car
  - Oil drilling rights
  - Anything with a common value component
- For the case of correlated valuations, one of the most famous result is Cremer & McLean (1985) auction.

Preliminaries

We consider the case of a single bidder with an external signal.

- a type $\theta \in \Theta = \{1, 2, \ldots, K\}$
- a valuation function $v : \Theta \to \mathbb{R}_+$
- an external signal $\omega \in \Omega$
- a joint distribution $\pi(\omega, \theta)$

IC and IR

The Cremer - McLean auction can extract the full social surplus $\sum_{\theta \in \Theta} \pi(\theta) v(\theta)$ as revenue, and satisfies ex-post IC and interim IR.

IC and IR

The Cremer - McLean auction can extract the full social surplus $\sum_{\theta \in \Theta} \pi(\theta) v(\theta)$ as revenue, and satisfies ex-post IC and interim IR.

- Ex-post incentive compatible (IC): for any realization of the valuation type θ and external signal ω, the bidder's utility is maximized when she reports her true type. That is,

$$v(\theta)x(\theta, \omega) - p(\theta, \omega) \geq v(\theta)x(\theta', \omega) - p(\theta', \omega) \quad \forall \theta' \in \Theta.$$

IC and IR

The Cremer - McLean auction can extract the full social surplus $\sum_{\theta \in \Theta} \pi(\theta) v(\theta)$ as revenue, and satisfies ex-post IC and interim IR.

- Ex-post incentive compatible (IC): for any realization of the valuation type θ and external signal ω, the bidder's utility is maximized when she reports her true type. That is,

$$v(\theta)x(\theta, \omega) - p(\theta, \omega) \geq v(\theta)x(\theta', \omega) - p(\theta', \omega) \quad \forall \theta' \in \Theta.$$

- Interim individual rational (IR): for any realization of the valuation type θ, the bidder's expected utility is always non-negative when she reports the truth. That is,

$$\mathbb{E}_{\omega \sim \pi(\theta)}[v(\theta)x(\theta, \omega) - p(\theta, \omega)] \geq 0.$$

The Question We Aim to Answer

- Cremer-McLean auction is considered as non-applicable due to its strong dependence on the common prior assumption.

The Question We Aim to Answer

- Cremer-McLean auction is considered as non-applicable due to its strong dependence on the common prior assumption.

- If the prior distribution π can only be accessed in the form of i.i.d. samples, how many samples are sufficient and necessary to guarantee at least $(1 - \epsilon)$ of optimal revenue?

Key Characteristics of the Prior

The sample complexity of Cremer-McLean sensitively depends on several key characteristics of the joint distribution:

- the size of valuation type set K
- the degree of correlation α
- the smallest marginal probability of any valuation type η

Definition 3 (α-strongly correlated distribution).

The joint distribution π is said to be α-strongly correlated if the singular values of any $K \times K$ nonsingular submatrix

$$\Gamma' = \begin{bmatrix} \pi(\omega'_1|1) & \pi(\omega'_2|1) & \cdots & \pi(\omega'_K|1) \\ \pi(\omega'_1|2) & \pi(\omega'_2|2) & \cdots & \pi(\omega'_K|2) \\ \vdots & & & \\ \pi(\omega'_1|K) & \pi(\omega'_2|K) & \cdots & \pi(\omega'_K|K) \end{bmatrix}$$

are at least α.

Our Results

Table 1: The sample complexity bounds of single-bidder Cremer-McLean auction. (K, α, and η represent the size of the valuation type set, the degree of correlation, and the smallest marginal probability of any valuation type, respectively.)

Upper Bound	$\tilde{O}(K^2\eta^{-3}\alpha^{-2}\epsilon^{-2})$
Lower Bound	$\Omega(K\eta^{-1}\alpha^{-2}\epsilon^{-2})$

The Upper Bound

Theorem 1.

For any $0 < \epsilon < 1$ and any α-strongly correlated distribution π with marginal probabilities at least η, there is an algorithm that returns an auction with expected revenue at least $(1-\epsilon)$ of the full surplus with probability at least $1-\delta$, if the number of samples m satisfies

$$m \geq 90 \cdot K\eta^{-3}\alpha^{-2}\epsilon^{-2} \cdot \max\{5\ln(6K\delta^{-1}), 8K\}.$$

- Construct from the samples an empirical prior distribution
- Shift down each valuation slightly
- Apply Cremer-McLean auction w.r.t. the empirical distribution and down-shifted valuations.

Generalization to Multi-Bidder Case

- The sample complexity of Cremer-McLean auction in n-bidder case is polynomial in K, the size of a single bidder's type space, not the size of the joint type space K^n.

- The sample complexity of Cremer-McLean auction in n-bidder case is polynomial in K, the size of a single bidder's type space, not the size of the joint type space K^n.

- It is possible to learn a near-optimal auction even when we do not know each joint type's probability precisely.

The Lower Bound

Theorem 2.

Suppose an algorithm A, given m independent samples from an unknown α-strongly correlated distribution with marginal probabilities are at least η, returns an auction with expected revenue at least $(1-\epsilon)$ of the optimal with probability at least 0.99. Then m must be at least $\Omega(K\eta^{-1}\alpha^{-2}\epsilon^{-2})$.

We construct a family of distributions

- very close to each other in terms of Kullback-Leibler divergence
- any algorithm that achieves $(1-\epsilon)$ of the optimal revenue must be able to distinguish between them

Takeaways from Theorem 1 and 2

- The size of a single bidder's type set, the degree of correlation among bidders and the smallest marginal probability are the key factors of the prior distributions that decide the learnability of a near-optimal auction.

Takeaways from Theorem 1 and 2

- The size of a single bidder's type set, the degree of correlation among bidders and the smallest marginal probability are the key factors of the prior distributions that decide the learnability of a near-optimal auction.

- Learning a near-optimal auction with correlated prior distributions is hard when there are many valuation types, or the level of correlation is low, or there is some valuation type with very small marginal probability.

Open Questions

There are many directions in which this work could be extended.

Open Questions

There are many directions in which this work could be extended.

- To close the gap between our upper and lower bounds

There are many directions in which this work could be extended.

- To close the gap between our upper and lower bounds
- The sample complexity of many auctions with correlated valuations are unknown (e.g., Ronen, 2001; Bei et al., 2019)
- Mostly unexplored ...

Thank you!

References

- Cole, R. and Roughgarden, T. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pp. 243-252, 2014.
- Cremer, J. and McLean, R. P. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. Econometrica, 53:345-361, 1985.
- Gonczarowski, Y. A. and Nisan, N. Efficient empirical revenue maximization in single-parameter auction environments. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 856-868, 2017.

- Guo, C., Huang, Z., and Zhang, X. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 662-673, 2019.
- Huang, Z., Mansour, Y., and Roughgarden, T. Making the most of your samples. *SIAM Journal on Computing*, 47(3):651-674, 2018.
- Myerson, R. B. Optimal auction design. *Mathematics of operations research*, 6(1):58-73, 1981.
- Syrgkanis, V. A sample complexity measure with applications to learning optimal auctions. In *Advances in Neural Information Processing Systems*, pp. 5352-5359, 2017.