# Brief Announcement: Bayesian Auctions with Efficient Queries


**Jing Chen**
Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, USA
jingchen@cs.stonybrook.edu

**Bo Li**
Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, USA
boli2@cs.stonybrook.edu

**Yingkai Li**
Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, USA
yingkli@cs.stonybrook.edu

**Pinyan Lu**
Institute for Theoretical Computer Science, Shanghai University of Finance and Economics, Shanghai 200433, China
lu.pinyan@mail.shufe.edu.cn



## Abstract

Generating good revenue is one of the most important problems in Bayesian auction design, and many (approximately) optimal dominant-strategy incentive compatible (DSIC) Bayesian mechanisms have been constructed for various auction settings. However, most existing studies do not consider the complexity for the seller to carry out the mechanism. It is assumed that the seller knows "each single bit" of the distributions and is able to optimize perfectly based on the entire distributions. Unfortunately this is a strong assumption and may not hold in reality: for example, when the value distributions have exponentially large supports or do not have succinct representations.

In this work we consider, for the first time, the *query complexity* of Bayesian mechanisms. We only allow the seller to have limited oracle accesses to the players' value distributions, via *quantile queries* and *value queries*. For a large class of auction settings, we prove *logarithmic* lower-bounds for the query complexity for any DSIC Bayesian mechanism to be of any constant approximation to the optimal revenue. For single-item auctions and multi-item auctions with unit-demand or additive valuation functions, we prove *tight* upper-bounds via efficient query schemes, without requiring the distributions to be regular or have monotone hazard rate. Thus, in those auction settings the seller needs to access much less than the full distributions in order to achieve approximately optimal revenue.



**2012 ACM Subject Classification** Theory of computation → Algorithmic game theory and mechanism design

**Keywords and phrases** The complexity of Bayesian mechanisms, quantile queries, value queries

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2018.108

**Related Version** A full version of the paper is available at [9], `https://arxiv.org/abs/1804.07451`.

**Acknowledgements** This work has been partially supported by NSF CAREER Award No. 1553385, NSF of China Grant No. 61741209 and the Fundamental Research Funds for the Central Universities. Part of this work was done when the first three authors were visiting Shanghai University of Finance and Economics.

# 1 Introduction

An important problem in Bayesian mechanism design is to design auctions that (approximately) maximize the seller's expected revenue. More precisely, in a Bayesian multi-item auction a seller has $m$ heterogenous items to sell to $n$ players. Each player $i$ has a private value for each item $j$, $v_{ij}$; and each $v_{ij}$ is independently drawn from some prior distribution $\mathcal{D}_{ij}$. When the prior distribution $\mathcal{D} \triangleq \times_{ij}\mathcal{D}_{ij}$ is of *common knowledge* to both the seller and the players, optimal Bayesian incentive-compatible (BIC) mechanisms have been discovered for various auction settings [16, 11, 4, 5], where all players reporting their true values forms a Bayesian Nash equilibrium. When there is no common prior but the seller knows $\mathcal{D}$, many (approximately) optimal dominant-strategy incentive-compatible (DSIC) Bayesian mechanisms have been designed [16, 17, 7, 14, 19, 6], where it is each player's *dominant strategy* to report his true values.

However, the *complexity* for the seller to carry out such mechanisms is largely unconsidered in the literature. Most existing Bayesian mechanisms require that the seller has full access to the prior distribution $\mathcal{D}$ and is able to carry out all required optimizations based on $\mathcal{D}$, so as to compute the allocation and the prices. Unfortunately the seller may not be so knowledgeable or powerful in real-world scenarios. If the supports of the distributions are exponentially large (in $m$ and $n$), or if the distributions are continuous and do not have succinct representations, it is hard for the seller to write out "each single bit" of the distributions or precisely carry out arbitrary optimization tasks based on them. Even with a single player and a single item, when the value distribution is irregular, computing the optimal price in time that is much smaller than the size of the support is not an easy task. Thus, a natural and important question to ask is *how much the seller should know about the distributions in order to obtain approximately optimal revenue.*

In this work we consider, for the first time, the *query complexity* of Bayesian mechanisms. In particular, the seller can only access the distributions by making oracle queries. Two types of queries are allowed, *quantile queries* and *value queries*. That is, the seller queries the oracle with specific quantiles (respectively, values), and the oracle returns the corresponding values (respectively, quantiles) in the underlying distributions. These two types of queries happen a lot in market study. Indeed, the seller may wish to know what is the price he should set so that half of the consumers would purchase his product; or if he sets the price to be 200 dollars, how many consumers would buy it. Another important scenario where such queries naturally come up is in databases. Indeed, although the seller may not know the distribution, some powerful institutes, say the Office for National Statistics, may have such information figured out and stored in its database. As in most database applications, it may be neither necessary nor feasible for the seller to download the whole distribution to his local machines. Rather, he would like to access the distribution via queries to the database. Other types of queries are of course possible, and will be considered in future works.

In this work we focus on *non-adaptive* queries. That is, the seller makes all oracle queries simultaneously, before the auction starts. This is also natural in both database and market study scenarios, and adaptive queries will be considered in future works.

# 2 Main Results

We would like to understand both lower- and upper-bounds for the query complexity of approximately optimal Bayesian auctions. In this work, we mainly consider three widely studied settings: single-item auctions and multi-item auctions with unit-demand or additive valuation functions. Our main results are summarized in Table 1.

**Table 1** Our main results. Here $h(\cdot) < 1$ is the tail function in the small-tail assumptions. For single-item auctions, the revenue is a $(1+\epsilon)$-approximation to the optimal BIC revenue, with $\epsilon$ sufficiently small. For multi-item auctions with unit-demand or additive valuation functions, the revenue is a $c$-approximation for some constant $c$.

| | Query Complexity | Distributions | | |
|---|---|---|---|---|
| | | Bounded in $[1, H]$ | | Unbounded & Small Tail |
| Auctions | Single-Item | $\Theta(n\epsilon^{-1}\log H)$ | | $O(-n\epsilon^{-1}\log h(\frac{2\epsilon}{3(1+\epsilon)}))$ |
| | Unit-Demand | $\forall c > 1$: $\Omega(\frac{mn\log H}{\log c})$ | $\forall c > 24$: $O(\frac{mn\log H}{\log(c/24)})$ | $\forall c > 24$: $O(-\frac{mn\log h(\frac{2c-48}{3c})}{\log(c/24)})$ |
| | Additive | $\forall c > 1$: $\Omega(\frac{mn\log H}{\log c})$ | $\forall c > 8$: $O(\frac{mn\log H}{\log(c/8)})$ | $\forall c > 8$: $O(-\frac{m^2 n\log h(\frac{c-8}{10c})}{\log(c/8)})$ |
| | Single-Item | Regular Distributions: $\Omega(n\epsilon^{-1})$, $O(n\epsilon^{-1}\log\frac{n}{\epsilon})$ | | |

Note that we allow arbitrary unbounded distributions that satisfy *small-tail assumptions*, which means the expected revenue generated from the "tail" of the distributions is negligible compared to the optimal revenue. Similar assumptions are widely adopted in sampling mechanisms [18, 12], to deal with irregular distributions with unbounded supports. Since distributions with bounded supports automatically satisfy the small-tail assumptions, the lower-bounds listed for the former apply to the latter as well.

Also note that our lower- and upper-bounds on query complexity are *tight* for bounded distributions. In the full version of the paper [9], we show that our lower-bounds allow the seller to make both value and quantile queries, and apply to any multi-player multi-item auctions where each player's valuation function is *succinct sub-additive*. The lower-bounds also allow randomized queries and randomized mechanisms.

For the upper-bounds, all our query schemes are deterministic and only make one type of queries: value queries for bounded distributions and quantile queries for unbounded ones. We show that our schemes, despite of being very efficient, only loses a small fraction of revenue compared with the cases where the seller has full access to the distributions.

## 3 Discussion and Future Directions

In the full version, we will elaborate on the connections between our work and related studies. For example, a closely related area is sampling mechanisms [10, 13, 15, 12, 3]. It assumes that the seller does not know $\mathcal{D}$ but observes independent samples from $\mathcal{D}$ before the auction begins. The *sample complexity* measures how many samples the seller needs so as to obtain a good approximation to the optimal Bayesian revenue. Our results show that query complexity can be exponentially smaller than sample complexity: the former is *logarithmic* in the "size" of the distributions, while the latter is known to be polynomial. We will also discuss other related studies such as [2, 1, 3, 8].

Finally, we point out some interesting further directions. As mentioned, we focus on non-adaptive queries in this work. One can imagine more powerful mechanisms using *adaptive* queries, where the seller's later queries depend on the oracle's responses to former ones. It is intriguing to design approximately optimal Bayesian mechanisms with lower query complexity using adaptive queries, or prove that even with such queries, the query complexity cannot be much better than our lower-bounds. Another interesting direction is when the answers of the oracle contain noise. In this case, the distributions learnt by the seller may be within a small distance from the "true distributions" defined by oracle answers without noise. It would be interesting to design mechanisms to handle such noise.

## References


1 Pablo Azar, Constantinos Daskalakis, Silvio Micali, and S Matthew Weinberg. Optimal and efficient parametric auctions. In *24th Symposium on Discrete Algorithms (SODA'13)*, pages 596–604, 2013.

2 Pablo Azar and Silvio Micali. Parametric digital auctions. In *4rd Innovations in Theoretical Computer Science Conference (ITCS'13)*, pages 231–232, 2013.

3 Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *58th Symposium on Foundations of Computer Science (FOCS'17)*, pages 516–527, 2017.

4 Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *44th Annual ACM Symposium on Theory of Computing (STOC'12)*, pages 459–478, 2012.

5 Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *53rd Symposium on Foundations of Computer Science (FOCS'12)*, pages 130–139, 2012.

6 Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality based unified approach to Bayesian mechanism design. In *48th Annual ACM Symposium on Theory of Computing (STOC'16)*, pages 926–939, 2016.

7 Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *43th ACM Symposium on Theory of Computing (STOC'10)*, pages 311–320, 2010.

8 Jing Chen, Bo Li, and Yingkai Li. From Bayesian to crowdsourced Bayesian auctions. *arXiv:1702.01416*, 2016.

9 Jing Chen, Bo Li, Yingkai Li, and Pinyan Lu. Bayesian auctions with efficient queries, full version. *arXiv:1804.07451*, 2018.

10 Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *46th Annual ACM Symposium on Theory of Computing (STOC'14)*, pages 243–252, 2014.

11 Jacques Cremer and Richard P McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.

12 Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *48th Annual ACM Symposium on Theory of Computing (STOC'16)*, pages 426–439, 2016.

13 Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *16th ACM Conference on Economics and Computation (EC'15)*, pages 45–60, 2015.

14 Robert Kleinberg and S Matthew Weinberg. Matroid prophet inequalities. In *44th Annual ACM Symposium on Theory of Computing (STOC'12)*, pages 123–136, 2012.

15 Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *29th Conference on Learning Theory (COLT'16)*, pages 1298–1318, 2016.

16 Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

17 Amir Ronen. On approximating optimal auctions. In *3rd ACM Conference on Electronic Commerce (EC'01)*, pages 11–17, 2001.

18 Tim Roughgarden and Okke Schrijvers. Ironing in the dark. In *17th ACM Conference on Economics and Computation (EC'16)*, pages 1–18, 2016.

19 Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'15)*, pages 92–109, 2015.