Mechanism Design with Unstructured Beliefs

Doctoral Consortium

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

ABSTRACT

Mechanism design is the task to design algorithms, toward desired objectives, that is robust to potential manipulation by strategic players. Traditionally, it is assumed that the mechanism designer and the players in the economy share some common knowledge. However, as pointed out by Wilson, such common knowledge is "rarely present in experiments and never in practice", and "only by repeated weakening of common knowledge assumptions will the theory approximate reality." In the work, we mainly focus on designing resilient mechanisms that work properly even in such a less foreseeable environment.

Bayesian auction design is a very flourishing topic in the field of mechanism design, where an important simplifying assumption is both the seller and the players know the exact distributions of all players' valuations. In this work we first consider the *query complexity* of Bayesian mechanisms, where we only allow the seller to have limited oracle accesses to the players' value distributions via simple queries. Then we further weaken the assumption by considering an information structure where the knowledge about the distributions can be *arbitrarily scattered* among the players. In both of these two unstructured information settings, we design mechanisms that are constant approximations to the optimal Bayesian mechanisms with full information.

Finally, we study an envy-free allocation problem that the unstructured beliefs need to be taken into consideration. In particular, we model an environment where each player is unaware of the bundles (or allocated items) of other players, but still knows he does not receive the worst bundle. We present both conceptual and algorithmic results for this new envy-free allocation domain.

KEYWORDS

Bayesian Auction; Query Complexity; Information Elicitation; Fair Allocation; Maximin-Aware Allocation

ACM Reference Format:

Bo Li. 2019. Mechanism Design with Unstructured Beliefs. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 BAYESIAN AUCTION DESIGN

Bayesian auction design has been extremely flourishing since the seminal work of [30]. One of the main focuses is to generate revenue, by selling m heterogenous items 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 (the *common prior* assumption), optimal Bayesian incentive-compatible (BIC) mechanisms have been discovered for various auction settings [9, 30], where all players reporting their true values forms a Bayesian Nash equilibrium.

Following [33], a lot of effort has been made to remove the common prior assumption. 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 [5, 10, 11, 14, 15, 24, 27, 34], where it is each player's *dominant strategy* to report his true values. In prior-free mechanisms [21, 25] the distribution is unknown and the seller learns it from the values of randomly selected players. In [8, 20, 22, 26, 29] the seller observes independent samples from the distribution before the auction begins. In robust mechanism design [6] the players have arbitrary probabilistic belief hierarchies about each other and in [18, 19] the players have arbitrary possibilistic belief hierarchies. In crowdsourced Bayesian auctions [1] each player privately knows *all* the distributions. Parametric mechanisms [2, 3] assume the seller knows some specific parameters about the distributions. Recently, [8] studies auctions where the seller knows some approximate distribution that is close to the true prior.

2 QUERY COMPLEXITY OF BAYESIAN AUCTIONS

In the literature, the *complexity* for the seller to carry out such mechanisms is largely unconsidered. 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. 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.*

To answer the above question, in this section, we study the *query complexity* [17] of Bayesian mechanisms, where the seller does not know anything about the distribution, but some powerful institutes, say the Office for National Statistics, may actually have such information all figured out and stored in its database. We only allow the seller to have limited oracle accesses to the players' value distributions, via *quantile queries* and *value queries*. That is, the seller

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

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.

Main Results. For monotone subadditive auctions, 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 MHR. Thus, in those auction settings the seller needs to access much less than the full distributions in order to achieve approximately optimal revenue.

3 INFORMATION ELICITATION FOR BAYESIAN AUCTIONS

An even more challenging situation is that the distributions are actually known by the players themselves that participate in the auctions (like long-time competitors in the market), then they have their own stakes about the final allocation and prices. For example, the seller might be the government that decides how to allocate spectrum licenses to telecom firms and the firms indeed know more than the government about the value of a specific block of frequencies. In this section, we consider a framework for auctions where the knowledge about the players' value distributions are *arbitrarily scattered* among the players and the seller must *elicit* and *aggregate* pieces of information from all players to gain a good understanding about the distributions, so as to decide how to sell the items [16]. We adopt an information elicitation approach [28] to address the above challenge and show that the seller can use the players as the oracle and get the truthful distributions, while keeping them also truthful about their own values.

We introduce directed *knowledge graphs* to succinctly describe the players' knowledge. Each player knows the distributions of his neighbors, different items' knowledge graphs may be different, and the structures of the graphs are *not known* by anybody. Our goal is to design *2-step dominant strategy truthful* (2-DST) information elicitation mechanisms whose expected revenue approximates that of the optimal BIC mechanism. A 2-DST mechanism [1] is such that, *(1) no matter what knowledge the players may report about each other, it is* dominant *for each player to report his true values; and (2) given that all players report their true values, it is* dominant *for each player to report his true knowledge about others.*

Main Results. In such an unstructured information setting, we design mechanisms for auctions with unit-demand and additive valuations that *aggregate* the players' knowledge, generating revenue that are constant approximations to the optimal Bayesian mechanisms with a common prior. Our mechanisms are 2-DST and the revenue increases gracefully with the amount of knowledge the players collectively have. Moreover, we show that for single-item auctions, if the knowledge graph has nice combinatorial structures (but may still be very sparse), then *nearly optimal* revenue can be generated by leveraging such structures.

4 MAXIMIN-AWARE ALLOCATIONS OF INDIVISIBLE GOODS

In the last few years or so, there has been a tremendous demand for fair division services to provide systematic and *fair* ways of dividing a set of indivisible m goods such as tasks, courses, and properties among a group of n players without money transfer, so that the players do not envy each other. To capture the fairness of an allocation, *envy-freeness (EF)* [23] (as well as its relaxations, such as envy-freeness up to one good (EF1) [7] and envy-freeness up to any good (EFX) [12]) is often used to ensure that each player should not envy or prefer the allocated goods of other players.

In this section, we study an envy-free allocation domain where the planner of the division tasks wishes to withhold allocation information of others from the user or the user simply does not know the allocation of others in the system. First, in many private fair allocations of goods such as tasks or gifts, the planner requires the system to preserve anonymity as not to give away the received bundles of other players. Second, due to the large number of (unrelated) players and items that could be potentially be involved in the division tasks (e.g., on the Internet such as MTurks), it is not meaningful for the planner to provide such information due to various reasons.

Proportionality (PROP) [32], maximin share (MMS) [7], and epistemic envy-free (EEF) [4] are three widely studied and well accepted fair allocation notions, all of which are defined for unaware players. However, MMS allocations only guarantee each player's best minimum value, and the value of some player's bundle can still be the least compared with others, which may cause significant envy; and PROP and EEF allocations (and their relaxations, e.g., removing any item) barely exist and cannot be properly approximated.

Motivated by this domain, we focus on answering the following questions [13]. *When indivisible goods are to be allocated among unaware players, what is the appropriate envy-free notion and how efficiently can the allocation be found subject to the envy-free notion?*

Main Results. First, we introduce a novel fairness notion of maximin aware (MMA), which guarantees that the player's bundle value is at least as much as her value for some other player's bundle, no matter how the remaining goods are distributed, i.e., there is always somebody who gets no more than her. We also provide two relaxations of MMA: MMA1 and MMAX. We show that MMA1 (and MMAX) potentially has stronger egalitarian guarantee than EF1 and such an allocation is guaranteed to exist for a broader class of valuations than MMS and EFX. Second, we present a polynomial-time algorithm that computes an allocation such that every player is either $\frac{1}{2}$-approximate MMA or exactly MMAX for additive valuations. Interestingly, the allocation returned by our algorithm is also $\frac{1}{2}$-approximate EFX when all players have subadditive valuations, which improves the existence result of [31].

5 ACKNOWLEDGMENTS

The author would like to thank his advisor, Jing Chen, for her supervision and tremendous support. The author would also like to thank Hau Chan, Yingkai Li, Pinyan Lu, and Xiaowei Wu for the helpful suggestions and fruitful collaborations of this work. This work is supported by NSF CAREER Award No. 1553385.

REFERENCES

- [1] Pablo Azar, Jing Chen, and Silvio Micali. 2012. Crowdsourced Bayesian auctions. In *3rd Innovations in Theoretical Computer Science Conference (ITCS'12)*. 236–248.
- [2] Pablo Azar, Constantinos Daskalakis, Silvio Micali, and S Matthew Weinberg. 2013. Optimal and efficient parametric auctions. In *24th Symposium on Discrete Algorithms (SODA'13)*. 596–604.
- [3] Pablo Azar and Silvio Micali. 2013. Parametric digital auctions. In *4rd Innovations in Theoretical Computer Science Conference (ITCS'13)*. 231–232.
- [4] Haris Aziz, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang. 2018. Knowledge, Fairness, and Social Constraints.. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- [5] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. 2014. A simple and approximately optimal mechanism for an additive buyer. In *55th Annual IEEE Symposium on Foundations of Computer Science (FOCS'14)*. 21–30.
- [6] Dirk Bergemann and Stephen Morris. 2012. *Robust Mechanism Design*. World Scientific, Hackensack, NJ.
- [7] E. Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [8] Yang Cai and Constantinos Daskalakis. 2017. Learning Multi-item Auctions with (or without) Samples. In *58th Symposium on Foundations of Computer Science (FOCS'17)*. 516–527.
- [9] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. 2012. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *53rd Symposium on Foundations of Computer Science (FOCS'12)*. 130–139.
- [10] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. 2016. A duality based unified approach to Bayesian mechanism design. In *48th Annual ACM Symposium on Theory of Computing (STOC'16)*. 926–939.
- [11] Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *49th Symposium on Theory of Computing (STOC'17)*. 170–183.
- [12] I. Caragiannis, D. Kurokawa, H. Moulin, A. Procaccia, N. Shah, and J. Wang. 2016. The unreasonable fairness of maximum nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 305–322.
- [13] Hau Chan, Jing Chen, Bo Li, and Xiaowei Wu. 2019. Maximin-Aware Allocations of Indivisible Goods. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- [14] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *43th ACM Symposium on Theory of Computing (STOC'10)*. 311–320.
- [15] Shuchi Chawla and J Benjamin Miller. 2016. Mechanism design for subadditive agents via an ex ante relaxation. In *17th Conference on Economics and Computation (EC'16)*. ACM, 579–596.
- [16] Jing Chen, Bo Li, and Yingkai Li. 2018. Information Elicitation for Bayesian Auctions. In *International Symposium on Algorithmic Game Theory*. Springer, 43–55.
- [17] Jing Chen, Bo Li, Yingkai Li, and Pinyan Lu. 2018. Brief Announcement: Bayesian Auctions with Efficient Queries. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 108:1–108:4.
- [18] Jing Chen and Silvio Micali. 2013. Mechanism design with possibilistic beliefs. *Journal of Economic Theory* 156 (2013), 77–102. Conference version appeared at *52nd IEEE Annual Symposium on Foundations of Computer Science (FOCS'11)*, pages 87–96, 2011.
- [19] Jing Chen, Silvio Micali, and Rafael Pass. 2015. Tight Revenue Bounds With Possibilistic Beliefs and Level-k Rationality. *Econometrica* 83, 4 (2015), 1619–1639.
- [20] Richard Cole and Tim Roughgarden. 2014. The sample complexity of revenue maximization. In *46th Annual ACM Symposium on Theory of Computing (STOC'14)*. 243–252.
- [21] Nikhil R Devanur and Jason D Hartline. 2009. Limited and online supply and the Bayesian foundations of prior-free mechanism design. In *10th ACM Conference on Electronic Commerce (EC'09)*. 41–50.
- [22] Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. 2016. The sample complexity of auctions with side information. In *48th Annual ACM Symposium on Theory of Computing (STOC'16)*. 426–439.
- [23] Duncan K Foley. 1967. An Improved Envy-Free Cake Cutting Protocol for Four Agents. *Yale Economics Essays* (1967), 45–98.
- [24] Sergiu Hart and Noam Nisan. 2012. Approximate revenue maximization with multiple items. In *13th ACM Conference on Electronic Commerce (EC'12)*. 656–656.
- [25] Jason D Hartline and Tim Roughgarden. 2008. Optimal mechanism design and money burning. In *40th Annual ACM Symposium on Theory of Computing (STOC'08)*. 75–84.
- [26] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. 2015. Making the most of your samples. In *16th ACM Conference on Economics and Computation (EC'15)*. 45–60.
- [27] Robert Kleinberg and S Matthew Weinberg. 2012. Matroid prophet inequalities. In *44th Annual ACM Symposium on Theory of Computing (STOC'12)*. 123–136.
- [28] Nolan Miller, Paul Resnick, and Richard Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51, 9 (2005), 1359–1373.
- [29] Jamie Morgenstern and Tim Roughgarden. 2016. Learning simple auctions. In *29th Conference on Learning Theory (COLT'16)*. 1298–1318.
- [30] Roger B Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [31] B. Plaut and T. Roughgarde. 2018. Almost envy-freeness with general valuations. In *29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'18)*. 2584–2603.
- [32] H. Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- [33] Robert Wilson. 1987. Game-theoretic analysis of trading processes. In *Advances in Economic Theory: Fifth World Congress*, T. Bewley (Ed.). 33–77.
- [34] Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'15)*. 92–109.