CEEPR
MIT Center for Energy and Environmental Policy Research

WP-2023-17
Research Brief

Online Learning in Multi-unit Auctions

Simina Brânzei, Mahsa Derakhshan, Negin Golrezaei, and Yanjun Han

We consider repeated multi-unit auctions with uniform pricing, which are widely used in practice for allocating goods such as carbon licenses. We consider two variants of the auction - where the price is set to the K-th highest bid and (K + 1)-st highest bid, and K is the number of units for sale. Our contribution is to analyze the bidding strategies and properties of these auctions in both the offline and online settings. In the offline setting, we design a polynomial time algorithm for the player's cumulative utility maximization by showing it is equivalent to a maximum-weight path on a directed acyclic graph. In the online setting, we design efficient online learning algorithms for bidding with sublinear regret under both full information and bandit feedback structures. We analyze the quality of the equilibria in the worst case through the lens of the core solution concept. We show that the (K + 1)-st price format is susceptible to collusion among the bidders; meanwhile, the K-th price format does not have this issue.

In a multi-unit auction, a seller brings multiple identical units of a good to a group of players (buyers) interested in purchasing the items. Repeated multi-unit auctions take place in many real world settings, such as when allocating licenses for CO_2 emissions, ads on online platforms, U.S. Treasury notes to investors, or trade exchanges over the internet. Repeated mechanisms often give rise to complex patterns of behavior. For example, the bidders may use the past transaction prices as a way of discovering the valuations of the competing bidders, or engage in collusion using bid rotation schemes since the history of bids can serve as a communication device.

Since repeated multi-unit auctions are used in many high-stakes real-world environments, it is paramount to understand: (i) what features should the auction have in the first place?; and (ii) given an auction format, how should the players bid to maximize their utility?

A good bidding strategy generally guarantees the player using it will not "regret" its strategy in hindsight, regardless of the strategies of others. There has been extensive work on understanding dynamics in auctions and games under various behavioral models of the players, but due to the complexity of understanding dynamical systems, many questions remain open.

In our paper, we design efficient algorithms that players can use for bidding in both offline and online settings. We give regret lower bounds and then analyze the quality of the equilibria in two main variants of the auction - one variant where the price is set to the K-th highest bid, and one variant where the price is set to the (K+1)-st highest bid, where K is the number of units of a good for sale.

In the settings we focus on, each player i has a value $v_{i,j}$ for receiving a j-th unit in their bundle, which is reported to the auctioneer in the form of a bid $b_{i,j}$. After receiving the bids, the auctioneer computes a price *p* per unit, then sorts the bids in descending order and allocates the j-th unit to the owner of the j-th highest bid, charging them *p* for the unit. The price *p* is set so that all units are sold.

In the offline setting, we consider a player i with valuation vi and a history $H_{-i} = (b_{-i}^{1}, \ldots, b_{-i}^{T})$ containing the bids submitted by the other players in past auctions. The offline problem is challenging because the decision (bid) space is exponentially large and hence naively experimenting with every possible bid vector leads to impractical algorithms. We overcome this challenge by relying on the structural properties of the offline problem. We carefully design a weighted directed acyclic graph (DAG) and identify a bijective map between paths from source to sink in the graph and bid profiles of player i. Since a maximum weight path in a DAG can be found in polynomial time, this yields a polynomial time algorithm for computing an optimum bid vector.

In the online setting, the players run learning algorithms to update their bids as they participate in the auction over time. We design an efficient online learning algorithm for bidding, which guarantees low regret for each player using it. The main idea is to run multiplicative weight updates, where the experts are paths in the DAG from the offline setting. Each such path corresponds to a bid profile for the learner. A challenge is that the number of paths in the graph (and so experts) is exponential. Nevertheless, we can achieve a polynomial runtime by representing the experts implicitly, using a type of weight pushing algorithm based on the method of path kernels.

We consider the online setting under both full information feedback and bandit feedback. Under full information feedback, the auctioneer makes public all the bids b^t submitted at the end of each round t. Under bandit feedback, the amount of information the players learn about each other is limited. At the end of each round, the auctioneer announces publicly only the resulting price, and then privately tells each player their allocation. Thus, bandit feedback may be relevant for reducing the amount of collusion among the players in repeated auction environments.

We find that, under full information feedback, for each player i and time horizon T, there is an algorithm for bidding in the repeated auction that runs in time $O(T^2)$ and guarantees the player's regret is bounded by $O(v_{i,1} \sqrt{(TK^3 \log T)})$. Under bandit feedback, for each player i and time horizon T, there is an algorithm for bidding in the repeated auction that runs in $O(KT + K^{(-5/4)} T^{(7/4)})$ and guarantees the player's regret is bounded by $O(v_{i,1} \min \sqrt{\{}^4(T^3 K^7 \log T),KT\}))$. Once again, $v_{i,1}$ is the valuation of player i for the first unit of the good, and K is the number of units of a good for sale.

We also analyze the quality of the Nash equilibria reached in the worst case, as they naturally apply to the empirical distribution of joint actions when the players use sub-linear regret learning algorithms. We show that the zero-price equilibria of the (K + 1)-st highest price auction are very stable by considering the core solution concept, which allows groups of players to coordinate.

In our setting, a strategy profile is core-stable if no group S of players can deviate by simultaneously changing their strategies such that each player in S weakly improves their utility and the improvement is strict for at least one player in S. The players outside S are assumed to have neutral reactions to the deviation, that is, they keep their strategy profiles unchanged. This is consistent with the Nash equilibrium solution concept, where only the deviating player changes their strategy.

When we consider the core solution concept without allowing for monetary transfers, we find that equilibria with price zero are the only ones that remain stable when also allowing deviations by groups of players. Moreover, there are many inefficient core-stable equilibria, all with price zero.

We also consider the core solution concept with transfers. Under this condition, players can make monetary transfers to each other. Thus, if the price is positive at some strategy profile, the players with "highest values" can pay the other players to lower their bids. When the price reaches zero, they can start withdrawing the payments. As a result, the only strategy profiles in the core with transfers are those where the players with "highest values" win and pay zero. Transfers allow the players with highest values to force getting their desired items at price zero, which cannot be enforced without transfers. Nevertheless, even without transfers, the price remains zero.

These results indicate that the (K + 1)-st price format is susceptible to collusion among the bidders while the K-th price format is not. Thus, the K-th price auction format may be preferable to the (K + 1)-st price auction, including for carbon auctions, which are based on the (K + 1)-st price format.

References

Brânzei, Simina, M. Derakshan, N. Golrezaei and Y. Han (2023). "Online Learning in Multi-unit Auctions." MIT CEEPR Working Paper 2023-17, October 2023

About the Authors

Simina Brânzei is an Assistant Professor of Computer Science at Purdue University. Brânzei's research interests are in theory of computation, artificial intelligence, algorithmic game theory, learning, algorithms, computational complexity and their interface with dynamical systems and optimization. Example of topics she worked on include fair division, markets and auctions, games, learning dynamics, and local search processes. Brânzei's research is supported by an NSF CAREER Award. Before joining Purdue, Brânzei was a postdoc at the Hebrew University of Jerusalem hosted by Noam Nisan and Michael Schapira, and a research fellow at the Simons Institute for the Theory of Computing. Brânzei completed her Ph.D. at Aarhus University in Denmark.

Mahsa Derakhshan is an assistant professor in the Khoury College of Computer Sciences at Northeastern University. Prior to that, she was a FODSI fellow at UC Berkeley and also a Postdoctoral Researcher at Princeton University in the Department of Computer Science. She is broadly interested in the design and analysis of algorithms. Mainly, Derakshan studies algorithms under uncertainty. A few sources of such uncertainty in her research are having stochastic data, limited access to information, and the presence of strategic behavior. She primarily study problems with applications to markets, such as matching markets and auctions. Derakhshan received her Ph.D. in Computer Science from the University of Maryland.

Negin Golrezaei is an Associate Professor of Operations Management at the MIT Sloan School of Management. Her current research interests are in the area of machine learning, statistical learning theory, mechanism design, and optimization algorithms with applications to revenue management, pricing, and online markets. Before joining MIT, Negin spent a year as a postdoctoral researcher at Google Research in New York where she worked with the Market Algorithm team to develop, design, and test new mechanisms and algorithms for online marketplaces. Negin received her BSc (2007) and MSc (2009) degrees in electrical engineering from the Sharif University of Technology, Iran, and a PhD (2017) in operations research from USC.

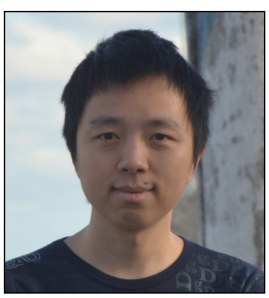

Yanjun Han is an Assistant Professor of Mathematics and Data Science at the Courant Institute of Mathematical Sciences at New York University. He received a B.E. in Electronic Engineering from Tsinghua University in July 2015, and an M.S. and Ph.D. in Electrical Engineering from Stanford University in Aug 2021, under the supervision of Tsachy Weissman. He was a postdoctoral scholar at the Simons Institute for the Theory of Computing, University of California, Berkeley in 2021-22, and a Norbert Wiener postdoctoral associate at the Statistics and Data Science Center (SDSC) in MIT IDSS in 2022-23. He is broadly interested in the mathematics of data science, including statistics, learning theory, bandits, and information theory.