MIT CEEPR
MIT Center for Energy and Environmental Policy Research

WP-2023-18
Research Brief

Learning in Repeated Multi-Unit Pay-As-Bid Auctions

Rigel Galgana & Negin Golrezaei

We investigate learning how to bid in repeated multi-unit pay-as-bid auctions. We study the problem of optimizing bidding strategies from the perspective of a single bidder. We show that the optimal solution to the offline problem can be obtained using a polynomial time dynamic programming (DP) scheme. We leverage the DP scheme to design online learning algorithms with polynomial time and space complexity under full information and bandit feedback settings. We achieve an upper bound on regret of $O(M\sqrt{T \log|\beta|})$ *and* $O(M\sqrt{|\beta| T \log|\beta|})$ *respectively, where M is the number of units demanded by the bidder, T is the total number of auctions, and* $|\beta|$ *is the size of the discretized bid space. Our numerical results suggest that when all agents behave according to our algorithms, market dynamics converge to a welfare maximizing equilibrium where bidders submit uniform bids. In addition, the pay-as-bid auction consistently generates significantly higher revenue than its popular alternative, the uniform price auction.*

Homogeneous multi-unit auctions are commonly used to auction off large quantities of identical items, for example in carbon emissions trading schemes, US Treasury auctions, procurement auctions, and wholesale electricity markets. In these multi-unit auctions, bidders submit a bid vector and then are allocated their goods and charged payments according to the auction format.

Of widespread use are the uniform price and pay-as-bid (PAB) mechanisms. As natural multiunit generalizations of the second and first price sealed bid auctions, bidders are allocated units in decreasing order of bids and, for each unit won, are charged as payment the lowest winning bid (uniform price) or their own bid (PAB). In this work, we focus on the PAB auction in light of the recent industry and research community-wide push towards first price auctions, which is mainly due to demand for price transparency and ease of revenue management.

Participants find it challenging to determine how to bid effectively in PAB auctions. In this paper, we address the issue of learning optimal bidding strategies in repeated multi-unit PAB auctions. We develop efficient no-regret algorithms that simplify the bidding complexity associated with PAB auctions. Through simulating the market dynamics derived from these learning algorithms, we empirically analyze the equilibria of PAB auctions, which have been poorly understood prior to our research.

To design low-regret bidding algorithms, we crucially leverage the structure of the hindsight optimal offline solution. In the offline/hindsight problem, the bidder has access to the (historical) dataset of submitted bids by competitors and seeks to find the utility maximizing bid vector on that dataset. In our paper, we show that the optimal solution to the offline problem—which is our benchmark in computing the regret of our online learning algorithms—can be solved using a polynomial time Dynamic Programming (DP) scheme.

We use decoupled exponential weights algorithms to learn in the online setting, in both the full information and bandit feedback regimes. We leverage our DP scheme to obtain decoupled rewards, or reward estimates in the bandit setting, for each unit-bid value pair. We show that our decoupled exponential weights algorithm achieves time and space complexities of $O(M|\beta|T)$ and $O(M|\beta|)$ respectively, where M is the number of units demanded by the bidder, T is the total number of auctions, and $|\beta|$ is the size of the discretized bid space. Optimizing for the discretization error from using finite bid space β, we show that in the full information setting, our algorithm achieves $O(M^{(3/2)}\sqrt{(T \log T})$ regret.

We also consider an alternative learning algorithm based on Online Mirror Descent (OMD) that improves the regret upper bounds of our decoupled exponential weight algorithm by a factor of $\sqrt{M}$ at the cost of additional computation. This results in an algorithm that achieves time and space complexities polynomial in M, $|\beta|$, T, as in the algorithms described above, with discretized regrets $O(M\sqrt{(T \log |\beta|)})$ and $O(M\sqrt{(|\beta| T \log |\beta|)})$ in the full information and bandit settings, respectively. Optimizing for the discretization error from β, we show that in the bandit feedback setting, we derive an algorithm that achieves $O(MT^{(2/3)})$ regret.

Our experiments yield valuable practical insights for both auction designers and participants. These insights are primarily derived from the simulations of PAB market dynamics using the no-regret learning algorithms described above, however, we additionally compare these results with the market dynamics of uniform price auctions using the algorithms described in Brânzei et al. (2023). Conducting such systematic comparisons was previously challenging due to the inherent difficulty of characterizing equilibria in these auctions prior to our research.

Our results produce three main practical insights. First, uniform bidding in PAB auctions is optimal. The market dynamics consistently yield convergence of the winning bids and largest losing bids to a common price across all bidders. While the payment for each unit can be different across units and across bidders, under a reasonable learning and bidding strategy, these payments across units and bidders converge to the same value in the long run.

Second, a simplified bidding interface is sufficient for PAB but not for Uniform Price. Auctioneers may find it easier to restrict bidders' demand expressiveness by requiring only a single price and quantity, rather than a vector of bids. The bid value convergence of the market dynamics suggests that this simplified bidding interface is appropriate. In contrast, we show that the market dynamics of the uniform price auction converge to a staggered bid vector, suggesting that the simplified bidding interface may significantly damage the uniform price auction's welfare, revenue, or bidders' utility.

Third, PAB obtains high revenue but slightly lower welfare than uniform price. Our results show that the PAB auction surpasses the uniform price auction in terms of revenue generation. However, it slightly lags behind in welfare, though to a lesser extent. Consequently, auctioneers who prioritize revenue should favor the PAB auction over uniform price auctions, and auctioneers who prioritize welfare should favor the uniform price auction over PAB auctions.

There are several intriguing avenues for future research that can be explored based on our current work. A promising direction is to leverage the structure induced by bid monotonicity in PAB auctions. Recent advancements in a simpler single-unit setting have demonstrated the efficacy of cross-learning between bids under certain feedback structures. It would be intriguing to investigate the potential benefits of applying cross-learning techniques in our multi-unit setting. By incorporating such methods, we could explore whether they can enhance our regret bounds.

References

Galgana, Rigel and Golrezaei, Negin (2023). "Learning in Repeated Multi-Unit Pay-As-Bid Auctions." MIT CEEPR Working Paper 2023-18, October 2023

About the Authors

Rigel Galgana is a first year Operations Research Ph.D. student at MIT working with Professor Negin Golrezaei. Prior to MIT, Rigel worked as a Portfolio Implementation analyst at AQR Capital Management. He graduated from Brown University in May 2020 where he studied applied mathematics, computer science, and economics. Rigel's interests lie at the intersection of learning, games, and mechanism design. His research focuses on the computational aspects of economic and decision-making problems, such as data driven mechanism design, equilibrium computation, and multi-agent reinforcement learning. Rigel is also interested in the modeling, simulation, and modern applications of games, auctions, and markets. Outside of research, Rigel enjoys playing table tennis, chess, guitar, and piano.

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.