Approaching the Overbidding Puzzle in All-Pay Auctions: Explaining Human Behavior through Bayesian Optimization and Equilibrium Learning

Extended Abstract

Markus Ewert
Technical University of Munich
Munich, Germany
ewert@in.tum.de

Stefan Heidekrüger
Technical University of Munich
Munich, Germany
stefan.heidekrueger@in.tum.de

Martin Bichler
Technical University of Munich
Munich, Germany
bichler@in.tum.de

ABSTRACT

It is an established fact in behavioral economics that in lab experiments of auctions, human subjects do not adhere to the risk-neutral Bayesian Nash equilibria of such games. Several attempts at explaining this *Overbidding Puzzle* focus on the bidders' psychology and suggest they may have parametrized utility functions that differ from the risk-neutral payoff. However, analytical equilibria of the resulting modified games are generally not available. Consequently, it has been difficult to identify the specific parameters and assess the merits of these proposed modifications in explaining empirical observations.

With recent advances in equilibrium learning, it has become tractable to compute approximations of Bayesian Nash equilibria. Building on these advances and Bayesian optimization, we propose a novel regression framework to infer unobserved parameters of Bayesian games from behavioral data. We apply our method to two data sets of human bidding behavior in all-pay auctions. For the first time, this makes it possible to directly compare the goodness-of-fit of several proposed qualitative explanations of overbidding.

KEYWORDS

Behavioural Game Theory; Equilibrium Learning; All-Pay Auctions

ACM Reference Format:
Markus Ewert, Stefan Heidekrüger, and Martin Bichler. 2022. Approaching the Overbidding Puzzle in All-Pay Auctions: Explaining Human Behavior through Bayesian Optimization and Equilibrium Learning: Extended Abstract. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 4 pages.

1 INTRODUCTION

A standard assumption in economic theory is that market participants are utility-maximizing, rational agents, and thus, they should behave according to the market's equilibrium. However, experimental studies in behavioral economics have repeatedly shown that human subjects do not conform to this assumption. A prominent example is the phenomenon of overbidding in auctions. Particularly in all-pay auctions, where all bidders have to pay their bid, not just the winner, this behavior extends to a bimodal bidding pattern, where low-valued bidders underbid, and high-valued bidders overbid the risk-neutral *Bayesian Nash Equilibrium* (BNE) [7, 15]. Due to the variety of practical applications of this auction [8], understanding the reasons for this bimodal behavior is essential and has been the subject of research in behavioral economics and psychology.

This *Overbidding Puzzle* has previously been approached by questioning the risk-neutrality of bidders and instead investigated psychological factors expressed by parametrized utility functions that might influence the bidder's behavior [6, 7, 11]. Nonetheless, it is yet unclear what factors explain overbidding because there does not exist a unified approach for estimating the parameters of such utility functions across several behavioral models and making their goodness-of-fit comparable on experimental data. A key difficulty to this has been the computational complexity of computing equilibria in these parametrized auctions. Most recently, however, there has been made progress in approximating equilibria of such games via numerical techniques based on multi-agent learning [2, 4, 10].

In this study, we propose a novel estimation framework using Bayesian Optimization (BO) and equilibrium learning techniques that, for the first time, allows a quantitative analysis of the goodness-of-fit to experimental data of various behavioral explanation attempts to the Overbidding Puzzle. We apply our method to symmetric all-pay auctions and the concepts of risk-aversion and anticipated regret. Our empirical findings coincide with established results in the empirical literature in those aspects where quantitative results were previously available. As our framework is not restricted to either a specific auction mechanism, type of utility function, or equilibrium oracle, future work may apply it to other problems in behavioral economics and behavioral psychology.

2 ESTIMATION FRAMEWORK

An auction is a continuous-type-and-action Bayesian Game $G^\theta = (N, \mathcal{V}, \mathcal{A}, F, u^\theta)$. N players participate in this game, where each agent i draws her type v_i, i.e. her *valuations* of the item(s) to be auctioned, from the set of possible *type profiles* $\mathcal{V} = \mathcal{V}_1 \times \cdots \times \mathcal{V}_n$ with some joint prior probability distribution F that is common knowledge among the bidders. Given these types, players must then choose a b_i from the set of available actions $\mathcal{A}_i$. We assume they make this choice according to some (pure) *strategy* $\beta_i : \mathcal{V}_i \to \mathcal{A}_i$. u^θ is the vector of individual utility functions $u_i^\theta : \mathcal{V}_i \times \mathcal{A} \to \mathbb{R}$ that describes the outcomes of the game, which depends on parameter(s) θ that characterize the underlying behavioral model. We restrict ourselves to symmetric utility functions with a shared parameter among all participants. Finally, this framework aims at

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

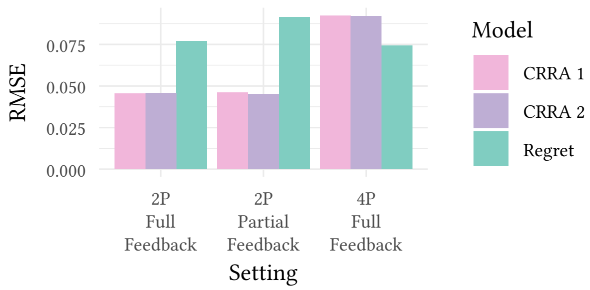

(a) Results of the regression evaluation method

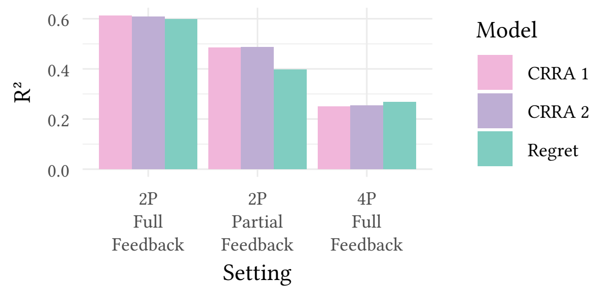

(b) Results of the R^2 evaluation method

Figure 1: The average model performances for all behavioral models per loss function.

estimating this parameter θ such that the resulting equilibrium strategy describes the experimental data sufficiently accurately.

Assuming the existence of an *equilibrium oracle EO* : $\theta \rightarrow \mathcal{A}^{\mathcal{V}}$ that, for a given parameter θ computes an estimate $\hat{\beta}$ of a BNE in G^{θ}, we can define goodness-of-fit in terms of some regression loss function $\ell(b, \hat{b})$ that compares the estimated bids $\hat{b} = \hat{\beta}(v)$, to those observed in the behavioral data b. We will consider two choices of loss functions: If the assumption of a specific trajectory of the BNE is warranted, one can estimate an equally-shaped regression model and determine the distance between this surrogate and the estimation of the equilibrium oracle via the *root mean squared error*. This is beneficial in settings where the experimental data is summarized as such a model, or when additional information should be considered during the estimation. Otherwise, one may compare the estimated bid function directly with the experimental data using a *coefficient of determination* $R^2 = 1 - \sum_k (b_k - \hat{b}_k)^2 / \sum_k (b_k - \bar{b})^2$, where $\bar{b}$ indicates the average bid of all lab subjects k.

The Bayesian Optimization scheme consists of two stages that are applied alternatingly at each time step t: In the *evaluation stage*, the goodness-of-fit of θ_t is evaluated via a call to the equilibrium oracle $\hat{\beta}$=$EO(\theta_t)$, computing the estimated equilibrium bids $\hat{b}_k$=$\hat{\beta}(v_k)$ that subjects k should have bid in the experiment if they were following $\hat{\beta}$, and then evaluating the resulting loss $\ell(b, \hat{b})$. In the *estimation stage*, the algorithm fits a stochastic model of the loss function over the entire domain Θ of the parameters to be inferred, relying on the history of all previously seen samples and their corresponding losses, to select a "promising" next sample θ_{t+1}. Specifically, we use a Gaussian Process model, as it is sufficiently expressive, provides a measure of output uncertainty over its domain, is suitable for iterative refitting when adding new data points, and is inexpensive to evaluate [3, 16]. Given this model, the next sample to be evaluated is chosen according to some acquisition criterion that should strike a balance between exploitation and exploration. Here, we choose the *expected improvement* criterion, which is a common choice in the BO literature [5, 16].

3 EXPERIMENTS

To test this estimation framework empirically, we apply it to three behavioral utility functions in all-pay auctions using real-world experimental data. These utilities describe commonly assumed behavioral models, namely, anticipated regret as defined by [11] and risk-aversion using two different *constant relative risk aversion* (CRRA) models [9, 13]. The underlying experimental data have been made available to us by the corresponding authors and consist of two- [11] and four-player [1] settings, where the former is further split into full- and partial feedback environments. We compare the goodness-of-fit of the utilities under both loss functions discussed above. Initial experiments using the recent equilibrium learning method NPGA that represents strategies via neural networks and provably learns local pure-strategy BNE via evolutionary strategy gradient approximation [2, 10] imply a quadratic trajectory of the optimal bid function, independently of the used behavioral model. Thus, we fit a quadratic Tobit model [17] for the regression loss.

Figure 1 shows that the experiments yield sufficient goodness-of-fit measures for all behavioral models and both loss functions, which indicates that the considered behavioral models are reasonable candidates for explaining overbidding. Although the differences between the measures are marginal, the concept of risk-averse bidders is more suitable than anticipated regret in the two-player, but vice versa in the four-player settings. The findings of the individual analysis of the models coincide with earlier observations made in the literature: For instance, as in [9, 14], the estimation of both CRRA models yields similar risk measures in all settings. Additionally, bidders in our study tend to be more risk-neutral in scenarios with a larger number of competitors or with less information about the winning bid, which is compatible with findings in [11, 12].

Overall, the results show that our estimation framework serves as a tool for measuring and comparing assumptions about bidder's behavior that are not directly observable but can be expressed via utility functions. Nonetheless, an extensive analysis of the psychological factors that cause overbidding requires considering further experiments and the definition of more complex models combining multiple assumptions. Even though the regret experiments confirmed the observation that equilibrium learning techniques like NPGA can sufficiently approximate the analytical BNE [2, 10], it is still unclear in which type of Bayesian games this is feasible, and thus, the success of our framework strongly depends on developments in this research area.

ACKNOWLEDGMENTS

We gratefully acknowledge the financial support from the Deutsche Forschungsgemeinschaft (DFG) (grant no. BI 1057/1-9).

REFERENCES

[1] Diego Aycinena, Rimvydas Baltaduonis, and Lucas Rentschler. 2019. Valuation structure in incomplete information contests: experimental evidence. *Public Choice* 179, 3 (2019), 195–208.

[2] Martin Bichler, Maximilian Fichtl, Stefan Heidekrüger, Nils Kohring, and Paul Sutterer. 2021. Learning equilibria in symmetric auction games using artificial neural networks. *Nature Machine Intelligence* 3, 8 (Aug. 2021), 687–695. https://doi.org/10.1038/s42256-021-00365-4

[3] Erik Bodin, Markus Kaiser, Ieva Kazlauskaite, Zhenwen Dai, Neill Campbell, and Carl Henrik Ek. 2020. Modulating Surrogates for Bayesian Optimization. In *International Conference on Machine Learning*. PMLR, 970–979.

[4] Vitor Bosshard, Benedikt Bünz, Benjamin Lubin, and Sven Seuken. 2020. Computing Bayes-Nash Equilibria in Combinatorial Auctions with Verification. *Journal of Artificial Intelligence Research* 69 (Oct. 2020), 531–570. https://doi.org/10.1613/jair.1.11525

[5] Eric Brochu, Vlad M Cora, and Nando De Freitas. 2010. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599* (2010).

[6] James C Cox, Vernon L Smith, and James M Walker. 1988. Theory and individual behavior of first-price auctions. *Journal of Risk and uncertainty* 1, 1 (1988), 61–99.

[7] Emmanuel Dechenaux, Dan Kovenock, and Roman M Sheremeta. 2015. A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics* 18, 4 (2015), 609–669.

[8] Uri Gneezy and Rann Smorodinsky. 2006. All-pay auctions—an experimental study. *Journal of Economic Behavior & Organization* 61, 2 (2006), 255–275.

[9] Jacob K Goeree, Charles A Holt, and Thomas R Palfrey. 2002. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory* 104, 1 (2002), 247–272.

[10] Stefan Heidekrüger, Paul Sutterer, Nils Kohring, Maximilian Fichtl, and Martin Bichler. 2021. Equilibrium Learning in Combinatorial Auctions: Computing Approximate Bayesian Nash Equilibria via Pseudogradient Dynamics. *AAAI-RLG Workshop 2021* (Feb. 2021). arXiv:2101.11946 [cs] http://arxiv.org/abs/2101.11946

[11] Kyle Hyndman, Erkut Y Ozbay, and Pacharasut Sujarittanonta. 2012. Rent seeking with regretful agents: Theory and experiment. *Journal of Economic Behavior & Organization* 84, 3 (2012), 866–878.

[12] Mark Isaac, Svetlana Pevnitskaya, and Kurt S Schnier. 2012. Individual behavior and bidding heterogeneity in sealed bid auctions where the number of bidders is unknown. *Economic Inquiry* 50, 2 (2012), 516–533.

[13] Oliver Kirchkamp, J Philipp Reiss, et al. 2006. *Another explanation for overbidding and another bias for underbidding in first-price auctions.* School of Economics and Finance, St. Salvator's College.

[14] Jingfeng Lu and Isabelle Perrigne. 2008. Estimating risk aversion from ascending and sealed-bid auctions: The case of timber auction data. *Journal of Applied Econometrics* 23, 7 (2008), 871–896.

[15] Wieland Müller and Andrew Schotter. 2010. Workaholics and dropouts in organizations. *Journal of the European Economic Association* 8, 4 (2010), 717–743.

[16] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems* 25 (2012).

[17] James Tobin. 1958. Estimation of Relationships for Limited Dependent Variables. *Econometrica* 26, 1 (1958), 24–36. http://www.jstor.org/stable/1907382