Autobidders with Budget and ROI Constraints: Efficiency, Regret, and Pacing Dynamics

Brendan Lucier BRLUCIER@MICROSOFT.COM
Microsoft Research, Cambridge

Sarath Pattathil SARATHP@MIT.EDU
Massachusetts Institute of Technology

Aleksandrs Slivkins SLIVKINS@MICROSOFT.COM
Microsoft Research, New York

Mengxiao Zhang MENGXIAO.ZHANG@USC.EDU
University of Southern California

Editors: Shipra Agrawal and Aaron Roth

Abstract

We study a game between autobidding algorithms that compete in an online advertising platform. Each autobidder is tasked with maximizing its advertiser's total value over multiple rounds of a repeated auction, subject to budget and return-on-investment (ROI) constraints. We propose a gradient-based learning algorithm that is guaranteed to satisfy all constraints and achieves vanishing individual regret. Our algorithm uses only bandit feedback and can be used with the first- or second-price auction, as well as with any "intermediate" auction format. Our main result is that when these autobidders play against each other, the resulting expected liquid welfare over all rounds is at least half of the expected optimal liquid welfare achieved by any allocation.[1] Our analysis holds whether or not the bidding dynamics converges to an equilibrium, side-stepping the dearth of provable convergence guarantees in the literature and hardness results that preclude such guarantees for budget-constrained second-price auctions (Chen et al., 2021).

Our vanishing-regret result extends to an adversarial environment without any assumptions on the other agents. We adopt a non-standard benchmark: the sequence of bids such that each bid b_t maximizes value for the round-t environment under time-averaged constraints. Hence, we side-step the impossibility results for the standard benchmark of best fixed bid (Balseiro and Gur, 2019). Our benchmark specializes to the standard one for a stationary environment.

When there is only a budget constraint, our algorithm specializes to an autobidding algorithm of Balseiro and Gur (2019), and our guarantees specialize to the regret and liquid welfare guarantees from Gaitonde et al. (2023). While our approach to bounding liquid welfare shares a common high-level strategy with Gaitonde et al. (2023), handling the ROI constraint, and particularly both constraints jointly, introduces a variety of new technical challenges. These challenges necessitate a new algorithm, changes to the way liquid welfare bounds are established, and a different methodology for establishing regret properties.

Keywords: autobidding, budget and ROI constraints, liquid welfare, regret

1. Liquid welfare is a standard notion of welfare under constraints. It is defined as the maximum amount the agents are willing to pay for the allocations that they receive.
2. **Extended abstract. Full version appears as arXiv:2301.13306v3.**
3. Some of the results have been obtained while S. Pattathil and M. Zhang were research interns at Microsoft Research. The authors are grateful to Bach Ha (Microsoft Bing Ads) for many conversations that informed our perspective, and to Sidharth Satya for providing research support.

© 2024 B. Lucier, S. Pattathil, A. Slivkins & M. Zhang.

References

Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.

Xi Chen, Christian Kroer, and Rachitesh Kumar. The complexity of pacing for second-price auctions. In *21st ACM Conf. on Economics and Computation (ACM-EC)*, page 318, 2021.

Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Budget pacing in repeated auctions: Regret and efficiency without convergence. In *14th Innovations in Theoretical Computer Science Conf. (ITCS)*, 2023.