Liquid Welfare guarantees for No-Regret Learning in Sequential Budgeted Auctions

Giannis Fikioris
Cornell University
gfikioris@cs.cornell.edu

Éva Tardos
Cornell University
eva.tardos@cornell.edu

We study the *liquid welfare* in sequential auctions with budget-limited players. We focus on first-price auctions, which are increasingly commonly used in many settings. Our performance metric, liquid welfare, is a natural and well-studied generalization of social welfare for the case of budget-constrained players: the liquid welfare of a player is the minimum of her budget and her allocated value.

We assume that each player has a *budgeted quasi-linear utility* (i.e., her utility is quasi-linear if the price paid is less than her budget, otherwise it is $-\infty$) and use a behavioral model that assumes a learning style guarantee: the resulting utility of each player is within a γ factor (where $\gamma \geq 1$) of a certain benchmark. We think of this as a "model of behavior", because we make no assumptions on the actual algorithm used by each player other that it satisfies this weak and intuitive performance guarantee. The benchmark we use for this guarantee for each player is the utility achievable by that player if she had bid λ times her value every round until her budget runs out, for any $\lambda \in [0, 1]$. While such uniform shading may not be optimal for the buyer, shading one's value with a multiplier λ to bid in budgeted settings is used in practice and has been studied extensively by previous work.

Under the above assumption for every player, we show a Price of Anarchy type bound: liquid welfare is within a $\gamma + 1/2 + O(1/\gamma)$ factor of the optimal one, assuming players have additive valuations. This positive result for first-price auctions is in stark contrast to sequential second-price auctions, where even with $\gamma = 1$, the resulting liquid welfare can be arbitrarily smaller than the maximum one. Our Price of Anarchy bound becomes about 2.41 when $\gamma = 1$.

We prove a lower bound of γ on the liquid welfare loss under the above assumption in first-price auctions, making our bound asymptotically tight. For the case when $\gamma = 1$ we show a lower bound of 2. Moreover, we extend all the above guarantees to the case where players have submodular valuations over the set of items they win across iterations; we prove a slightly worse Price of Anarchy bound of $\gamma + 1 + O(1/\gamma)$ which is about 2.62 when $\gamma = 1$.

Our work differentiates from previous work by naturally extending it to many directions. First, we do not assume that the players' values are sampled from a static distribution. Instead, our guarantees are true even if the players' value distribution changes over time, or even if values are picked adversarially. Second, previous work assumed that players are using a specific gradient descent style algorithm to guarantee a bounded Price of Anarchy. Our results hold for any algorithms used by the players, as long as each satisfies the general behavioral assumption described above for some γ.

We also provide a learning algorithm that a player with additive valuation and budget B can use to achieve the guarantee of γ we require for our results. When the number of rounds is T and the player's values are normalized so they are upper-bounded by 1, our algorithm guarantees $\gamma = T/B$, even when her values and the bids of the other players are chosen adversarially. This result is tight when $B = \Omega(T)$ which is the standard assumption in these settings.

The full paper can be found in https://arxiv.org/abs/2210.07502.