Budget Pacing in Repeated Auctions: Regret and Efficiency without Convergence

Jason Gaitonde [1] Yingkai Li [2] Bar Light [3] Brendan Lucier [4] Aleksandrs Slivkins [3]

[1]Cornell University [2]Yale University [3]Microsoft Research NYC [4]Microsoft Research New England

Introduction: Online Advertising and Autobidding

- Online advertisements sold via repeated auctions
- Advertisers must balance their per-auction preferences with global budget constraints over the entire ad campaign
- Autobidding: an advertiser declare preferences and constraints, their bidding algorithm dynamically adjusts bids

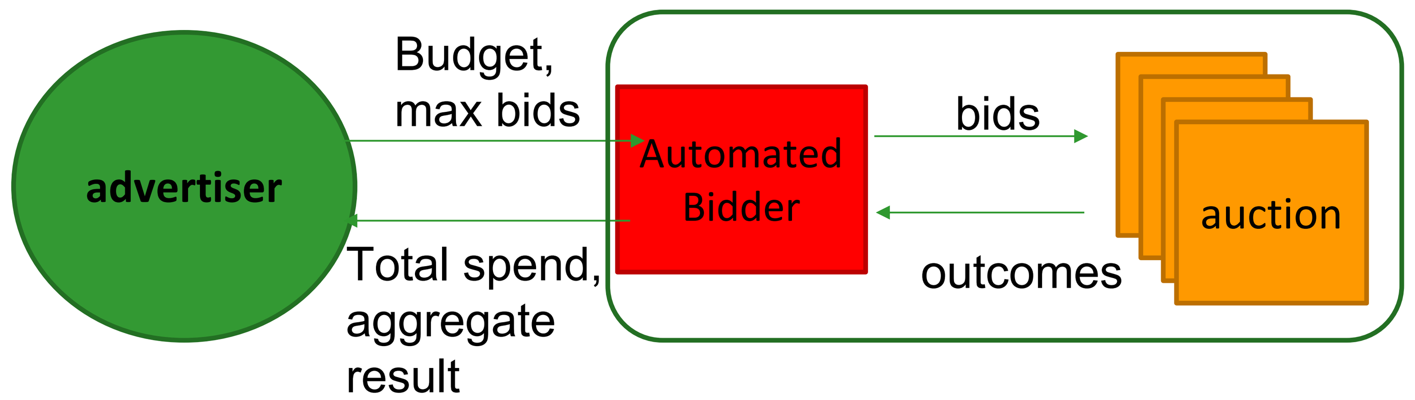

- Special case of Bandits with Knapsacks: known obstructions to achieving no-regret due to budget, even against best fixed action [6]

Main Questions

Design bidding algorithms that simultaneously

- **play well together, in terms of some "aggregate" objective,**
- **perform well individually as bandit algorithms.**

What guarantees can be achieved?

Model Formulation

- T rounds, one impression auctioned off per round
- Advertiser i specifies global budget B_i, per-round budget $\rho_i \equiv B_i/T$
- Values $\mathbf{v}_t = (v_{1t}, \ldots, v_{nt})$ at time t sampled from distribution F
- Autobidder i observes v_{it} and submits bid $b_{it} \le v_{it}$
- Auction: allocations $x_{it} \in [0,1]$ s.t. $\sum_i x_{it} \le 1$ and payments $p_{it} \ge 0$
- Autobidder goal: maximize $\sum_t x_{it} v_{it}$ subject to $\sum_t p_{it} \le B_i$
- **Objective:** Liquid Welfare [5]: willingness to pay for allocation

$$W_i = \min\{B_i, \textstyle\sum_t x_{it} v_{it}\}$$

Single-Round: Pacing Equilibria

- Multiplicative pacing: choose $\mu_i \ge 0$, bid $b_i = v_i/(1+\mu_i)$
- Known Results: truthful auction $\Rightarrow \exists$ equilibrium $(\mu_1^*, \ldots, \mu_n^*)$.
- Typically, equilibria are very bad on welfare (compared to optimum)
- Any equilibrium attains $1/2$-approximation of optimal *liquid* welfare [1, 2]; this guarantee is best possible in the worst case.
- Problem: finding an equilibrium in pacing strategies for second-price is PPAD-hard [4] $\Longrightarrow$ cannot hope for convergence!

Pacing Algorithms

- Each algorithm i dynamically updates multiplier $\mu_{it} \ge 0$,
- Balseiro-Gur (BG) pacing [3]: update for agent i, round t is

$$\mu_{i,t+1} = \max\{0, \mu_{i,t} + \varepsilon_i \cdot (p_{it} - \rho_i)\}$$

- Individual performance analyzed for 2nd-price auctions, convergence to equilibrium under strong convexity assumptions

Main Results

Aggregate Guarantees: a class of pacing algorithms (incl. BG pacing) achieves $1/2$-approximation of optimal Liquid Welfare. State-of-art guarantee without convergence!

Individual Guarantees: BG pacing attains no-regret w.r.t. *sequence of budget-pacing factors* $\mu_{i,t}$, $t \in [T]$, in *slowly changing environments*. The new benchmark bypasses impossibility of no-regret guarantees!

Generality: "core auctions" (for "aggregate") and "monotone bang-per-buck auctions" (for "individual"). Both: wide classes which include 1st-price auctions, 2nd-price auctions, and GSP. Not assuming convexity.
Core problem: single-item auction and time-invariant values.

Key Ideas

Aggregate guarantees: "taming" non-stationary dynamics of the $\boldsymbol{\mu}_{it}$'s. For each advertiser, partition rounds into epochs where $\boldsymbol{\mu}_{it} > 0$

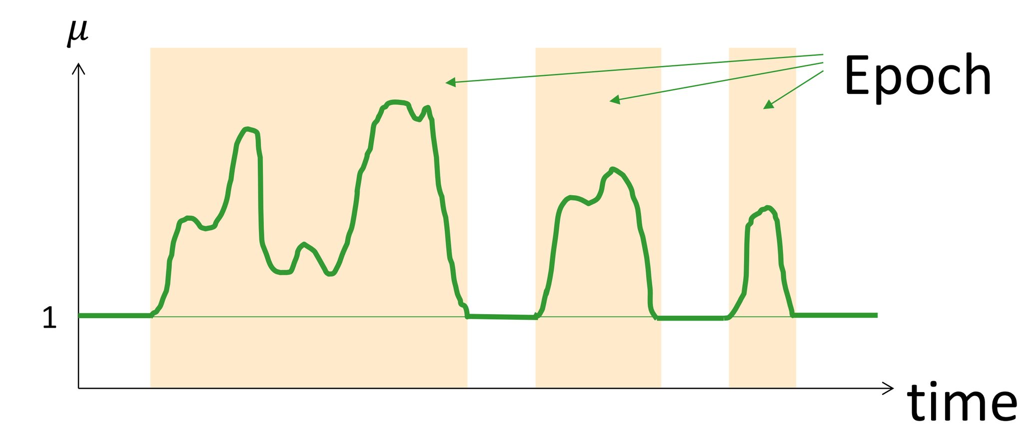

- $\mu_{it} > 0 \implies$ exceeding target spend $\implies$ lower bound on value obtained during epochs directly from definition of μ_{it}
- If instead $\mu_{it} = 0$, can charge welfare loss to **payments** of the other agents by core auction property

Individual Guarantees: interpret pacing dynamics as **stochastic gradient descent** (SGD) w.r.t. an artificial objective, then (with more technical work) re-use known SGD guarantees.

Future Directions

Similar guarantees with more general (classes of) bidding algorithms?
Similar aggregate guarantees with (even) stronger individual ones?
Similar results with more general types of global constraints?

[1] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. Autobidding with constraints. In *WINE 2018*.

[2] Moshe Babaioff, Richard Cole, Jason D. Hartline, Nicole Immorlica, and Brendan Lucier. Non-quasi-linear agents in quasi-linear mechanisms. In *ITCS 2021*.

[3] Santiago R. Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Manag. Sci.*, 2019.

[4] Xi Chen, Christian Kroer, and Rachitesh Kumar. The complexity of pacing for second-price auctions. In *ACM EC 2021*.

[5] Shahar Dobzinski and Renato Paes Leme. Efficiency guarantees in auctions with budgets. In *ICALP 2014*.

[6] Nicole Immorlica, Karthik Abinav Sankararaman, Robert E. Schapire, and Aleksandrs Slivkins. Adversarial bandits with knapsacks. In *FOCS 2019*.