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

Brendan Lucier (MSR) Sarath Pattathil (MIT) Alex Slivkins (MSR) Mengxiao Zhang (USC)

Introduction: Online Advertising and Autobidding

- Online advertisements sold via repeated auctions
- Advertisers (agents) balance per-auction preferences with global budget and return-on-investment (ROI) constraints
- Autobidding: an agent declares preferences and constraints, platform's bidding algorithm dynamically adjusts bids
- Each autobidder optimizes on behalf of its advertiser, competing with the other autobidders in the auctions.
- Generalizes [5] (budgeted agents) to include ROI constraints

Main Question

Design bidding algorithms that simultaneously

- **perform well individually as bandit algorithms, and**
- **"play well together" to maximize a welfare objective.**

What guarantees can be achieved?

Model Formulation

- Rounds $t \in [T]$, one impression auctioned off per round
- Agent i specifies budget B_i (write $\rho_i = B_i/T$) and ROI target γ_i
- Values $\mathbf{v}_t = (v_{1t}, \dots, v_{nt})$ at time t sampled from distribution F (Can be correlated across agents, e.g., click rates; i.i.d. across time)
- 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 $p_{it} \ge 0$ (Auction rule is fixed: think 1st-price or 2nd-price auction)
- Autobidder goal: choose bids to maximize $\sum_t x_{it} v_{it}$ subject to

$$\sum_{\text{rounds } t} p_{it} \le B_i \text{ (budget)} \quad \text{and} \quad \sum_{\text{rounds } t} v_{it}x_{it} - \gamma_i p_{it} \ge 0 \text{ (ROI)}$$

- **Objective:** Liquid Welfare (LW) – willingness to pay for allocation

$$\sum_{\text{agents } i} W_i = \sum_{\text{agents } i} \min\left\{B_i, \tfrac{1}{\gamma_i}\textstyle\sum_t x_{it}v_{it}\right\}$$

Background: Single-Round Pacing Equilibria

- Multiplicative pacing: choose $\mu_i \ge 0$, bid $b_i = v_i/(1+\mu_i)$
- Known Results: truthful stage auction $\Rightarrow \exists$ equilibrium $(\mu_1^*, \dots, \mu_n^*)$.
- Any equilibrium: $1/2$-approximation of optimal *liquid welfare* [1, 2]. This guarantee is best possible in the worst case.
- No hope for convergence! Equilibirum computation is PPAD-hard for 2nd-price auctions [4].

Pacing Algorithm

- Builds on Balseiro & Gur algorithm for budgeted agents [3]
- Two multipliers μ_{it}^B & μ_{it}^R to keep track of the two constraints
- Update for agent i, round $t+1$:

$$\mu_{i,t+1}^B = \mu_{i,t}^B + \eta_{i,B}\cdot(p_{it} - \rho_i)$$
$$\mu_{i,t+1}^R = \mu_{i,t}^R + \eta_{i,R}\cdot(\gamma_i p_{it} - v_{it}x_{it})$$

- Pick the larger multiplier to decide the bid

$$\mu_{i,t+1} = \max\{0, \mu_{i,t+1}^B, \mu_{i,t+1}^R\}; \quad b_{i,t+1} = 1/(1+\mu_{i,t+1})$$

Feedback: only bandit feedback & only for the larger multiplier

Lemma: ROI and Budget constraints satisfied in all rounds w.p. 1.
Idea: $\mu_{i,t}^B$, $\mu_{i,t}^R$ encode total slack in their corresponding constraints.

Main Guarantees

Aggregate: $1/2$-approximation of optimal Liquid Welfare.
Matches best-possible for equilibrium, but without convergence

Individual: No-regret in *slowly changing environments* w.r.t. sequence of *per-round optimal multipliers* $\mu_{i,t}^*$, $t \in [T]$.
Benchmark bypasses known impossibility of no-regret guarantees

Generality: 1st & 2nd-price auctions, & anything "in between". Aggregate guarantees $\rightarrow$ *core auctions* for a divisible good.

Main application: single-item auction; time-invariant value-per-click; click rates correlated across agents; budgets & ROI constraints

Key Ideas

Aggregate Guarantees: "taming" non-stationary dynamics of the $\boldsymbol{\mu}_{it}$'s. Building on [5], partition rounds into epochs for each agent.

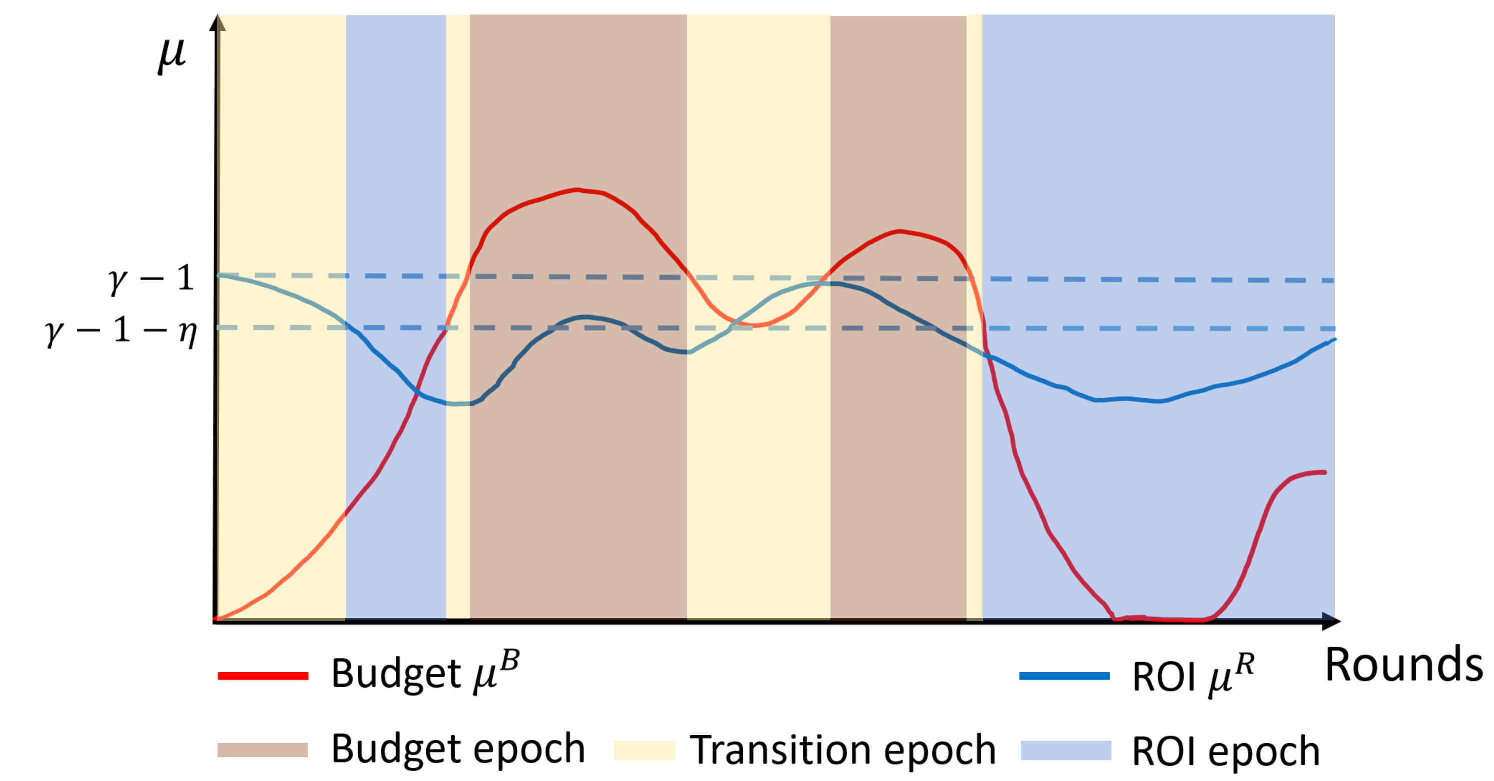

- ROI epoch: μ_t^B, μ_t^R both small $\Longrightarrow$ high bid $\Longrightarrow$ LW $\ge v_t/\gamma$
- Lemma: $\mu_t^R \le \gamma - 1$ always, so only μ_t^B can ever lie above $\gamma - 1$.
- Budget epoch: μ_t^B starts and ends near $\gamma - 1 \Longrightarrow$ payment is close to target budget within epoch $\Longrightarrow$ LW is at least ρ_i on average.
- Transition epoch: bound errors that accumulate between epochs

Individual Guarantees. Challenge: only bandit feedback & only for the *larger* multiplier, so the other may behave poorly. But since constraints are satisfied ex post (lemma), we only suffer loss from underbidding.

Future Directions

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

[1] G. Aggarwal, A. Badanidiyuru, and A. Mehta. Autobidding with constraints. In *WINE 2018*.

[2] M. Babaioff, R. Cole, J.D. Hartline, N. Immorlica, and B. Lucier. Non-quasi-linear agents in quasi-linear mechanisms. In *ITCS 2021*.

[3] S.R. Balseiro and Y. Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 2019.

[4] X. Chen, C. Kroer, and R. Kumar. The complexity of pacing for second-price auctions. In *ACM EC 2021*.

[5] J. Gaitonde, Y. Li, B. Light, B. Lucier, and A. Slivkins. Budget pacing in repeated auctions: Regret and efficiency without convergence. *ITCS 2023*.