On the Coordination of Value-Maximizing Bidders

Yanru Guan[1], Jiahao Zhang[2], Zhe Feng[3], Tao Lin[4]

[1] Peking University, [2] Carnegie Mellon University, [3] Google DeepMind,
[4] Microsoft Research & CUHK-Shenzhen

PDF

Motivation

- Auto-bidding is widely used in online advertising to maximize value under Return-on-Spend (RoS) constraints.
- In practice, multiple bidders (e.g., ad campaigns) are often managed by the **same** platform or advertiser.

? Can a simple **coordination rule** improve both RoS compliance and coalition value, compared with independent bidding?

Model

- Repeated second-price auctions with outside bid $d_t^O \sim D$
- At round t, bidder i has private value $v_{i,t} \sim F$ and bids $b_{i,t}$
- Goal: maximize total value subject to the RoS constraint (equivalently, non-negative total utility)

Independent bidding

Each bidder runs its own auto-bidding algo and bids separately

Coordination rule: bidder with **highest value** bids

- N bidders in coalition
- Total T rounds
- At round t, let $j^* = \arg\max_{j \in [N]} v_{j,t}$
- Only bidder j^* submits a positive bid: $\forall j \neq j^*, b_{j,t} = 0.$

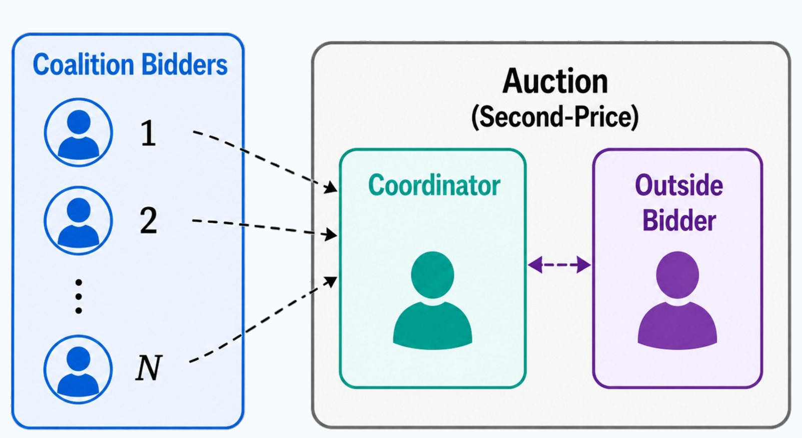

Coordination Improves Utility (RoS Compliance)

Assumption 1. The top two coalition values are competitive against the outside bid:

$$\Delta = E\left[\left(v_{(N-1)} - d^O\right)^+ - \left(d^O - v_{(N)}\right)^+\right] \geq 0$$

Theorem 1. If-and-only-if condition for utility gain

- Under Assumption 1, for *any* overbidding $(b_{i,t} \geq v_{i,t})$ algo A and every bidder $i \in [N]$,
$$E\left[U_i^{C,A} - U_i^{I,A}\right] \geq \frac{T\Delta}{N} \geq 0.$$
- If Assumption 1 fails, there exists an overbidding algo A such that for every bidder $i \in [N]$, for large enough T,
$$E\left[U_i^{C,A} - U_i^{I,A}\right] < 0.$$

C/I denotes coordinated/independent bidding; $U_i^{\cdot,A}$ is bidder i's total utility.

Mirror-Descent Auto-Bid Algo

1. Initialize multiplier $\lambda_{i,1} = 1$
2. Bid rule: $b_{i,t} = \left(1 + 1/\lambda_{i,t}\right) v_{i,t}$
3. Obtain allocation $x_{i,t}$ & payment $p_{i,t}$
4. Compute utility $g_{i,t} = v_{i,t} x_{i,t} - p_{i,t}$
5. Update multiplier
$$\lambda_{i,t+1} = \arg\min_{\lambda > 0} \{\alpha \cdot g_{i,t} \cdot \lambda + D_h(\lambda, \lambda_{i,t})\},$$

where $\alpha = 1/\sqrt{T}$ is learning rate, **h is strictly convex** and
$$D_h(\lambda, \mu) := h(\lambda) - h(\mu) - h'(\mu)(\lambda - \mu)$$
is h −associated Bregman divergence.

Example. Feng et al. [1] studied the case of entropy mirror map $h(\lambda) = \lambda \log \lambda - \lambda$, and the update being $\lambda_{i,t+1} = \lambda_{i,t} \exp\left(-\alpha g_{i,t}\right)$.

Coordination Achieves Higher Total Value

Theorem 2. Higher Coalition Value

- For any mirror-descent algo A,
$$\lim_{T \to +\infty} E\left[\frac{1}{T} \sum_{t=1}^{T} \left(V_t^{C,A} - V_t^{I,A}\right)\right] \geq 0$$

Theorem 3. Asymptotic Optimality

- Under Assumption 1, for any mirror-descent algo A and any other coordination mechanism $\mathcal{G}$,
$$\lim_{T \to +\infty} E\left[\frac{1}{T} \sum_{t=1}^{T} \left(V_t^{C,A} - V_t^{\mathcal{G}}\right)\right] \geq 0$$

$V_t^{\cdot,A}$ is coalition value at round t.

Non-i.i.d. Experiment

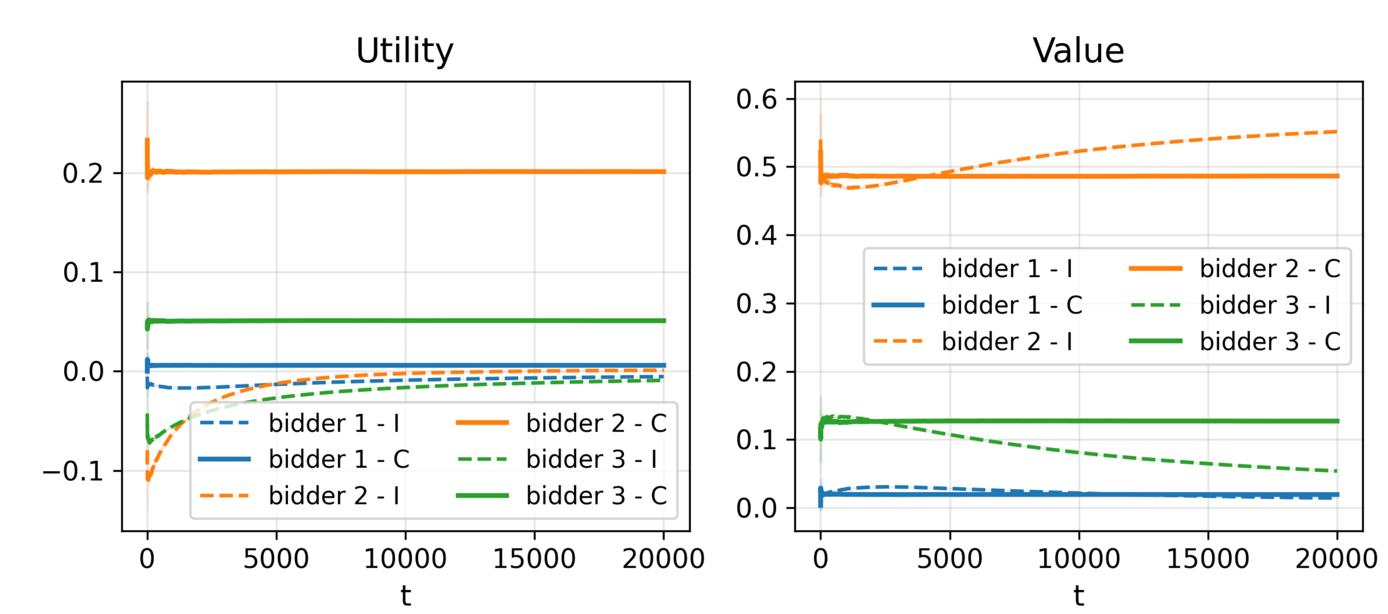

Coordination improves *aggregate* utility and *total* value under mild assumptions, but **per-bidder** value improvement is **not** guaranteed under non-i.i.d. value distributions.

Future Directions

1. Outside bidders that also use auto-bidding
2. Alternative coordination mechanisms
3. Other auction formats (e.g., first-price auction)

[1] Zhe Feng, Swati Padmanabhan, and Di Wang. 2023. Online Bidding Algorithms for Return-on-Spend Constrained Advertisers. The ACM Web Conference 2023.