Smoothness Meets Autobidding: Tight Price of Anarchy Bounds for Simultaneous First-Price Auctions

Riccardo Colini-Baldeschi[1], Sophie Klumper[2,3], Twan Kroll[2,3], Stefano Leonardi[4], Guido Schäfer[2,3], Artem Tsikiridis[2]

[1]Meta [2]Centrum Wiskunde and Informatica [3]University of Amsterdam [4]Sapienza University of Rome

A Combinatorial Auction Setting

Let $[m]$ = set of items, $[n]$ = set of bidders.

Each bidder $i \in [n]$ has a valuation function $v_i : 2^{[m]} \mapsto \mathbb{R}_{\geq 0}$ with $v(\emptyset) = 0$ and $v_i(S) \subseteq v_i(T)$ for all $S \subseteq T \subseteq [m]$.

- v_i is additive ($v_i \in \mathcal{V}_{\text{ADD}}$) if there exist $v_{ij} \in \mathbb{R}_{\geq 0}$ for all $j \in [m]$ such that $v_i(S) = \sum_{j \in S} v_{ij}$.
- v_i is XOS ($v_i \in \mathcal{V}_{\text{XOS}}$), if there exists a class $\mathcal{L}_i = \{(v_{ij}^{\ell})_{j \in [m]} \in \mathbb{R}^m_{\geq 0}\}$ of additive valuations such that for every $S \subseteq [m]$, it holds that $v_i(S) = \max_{\ell \in \mathcal{L}_i} \sum_{j \in S} v_{ij}^{\ell}$.

It holds that $\mathcal{V}_{\text{ADD}} \subseteq \mathcal{V}_{\text{XOS}}$. XOS functions include all submodular.

Item Bidding with First Price Auctions (FPAs)

Item Bidding: each $i \in [n]$ submits a bid $b_{ij} \geq 0$ per item $j \in [m]$.
First-price Auctions: mechanism collects $\boldsymbol{b}_i = (b_{ij})_{j \in M} \in \mathbb{R}^m_{\geq 0}$ from each $i \in [n]$. Fix profile $\boldsymbol{b} = (\boldsymbol{b}_1, \dots, \boldsymbol{b}_n)$. For each item $j \in [m]$:

- **winner** $\mathsf{w}(j, \boldsymbol{b})$: highest bidder for j i.e. $\mathsf{w}(j, \boldsymbol{b}) = \arg\max_{i \in [n]} b_{ij}$
- **payments** $(p_{ij}(\boldsymbol{b}))_{i \in [n]}$: $p_{\mathsf{w}(j,\boldsymbol{b})j}(\boldsymbol{b}) = b_{\mathsf{w}(j,\boldsymbol{b})j}$ and $p_{ij}(\boldsymbol{b}) = 0$ if $i \neq \mathsf{w}(j, \boldsymbol{b})$.

Allocation of bidder $i \in [n]$: $x_i(\boldsymbol{b}) = \{j \in [m] \mid i = \mathsf{w}(j, \boldsymbol{b})\}$
Payment of bidder $i \in [n]$: $p_i(\boldsymbol{b}) = \sum_{j: i = \mathsf{w}(j,\boldsymbol{b})} p_{ij}$

Autobidding Agents and Coarse-correlated Equilibria (CCE)

Autobidding: bidders may delegate bidding decisions to automated agents to bid for them. Leading paradigm in online advertising, see also the survey of **Aggarwal et al. (SIGecom Exchanges, 2024)**.
Hybrid Bidders: have different reliance on autobidding agents. Formally, for each $i \in [n]$ let $\sigma_i \in [0, 1]$ be the payment sensitivity. For each bidder $i \in [n]$, define:

- **gain function**: $g_i(\boldsymbol{b}) = v_i(x_i(\boldsymbol{b})) - \sigma_i \cdot p_i(\boldsymbol{b})$
- **ROI-constraint**: $p_i(\boldsymbol{b}) \leq v_i(x_i(\boldsymbol{b}))$

Let $\boldsymbol{B}$ be a random bid profile and $(\boldsymbol{B}_i, \boldsymbol{B}_{-i})$ for each $i \in [n]$ be its projections in lower dimensions.

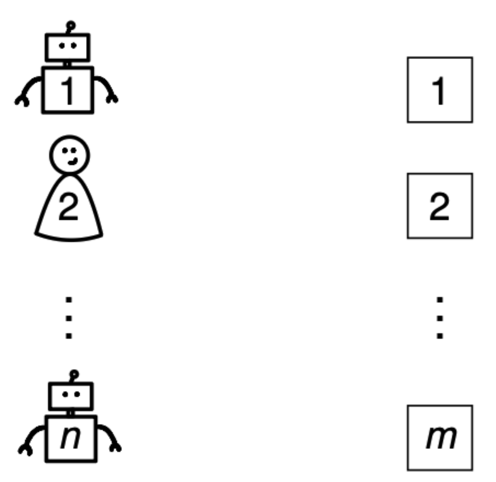

Optimization Problem of $i \in [n]$ (given $\boldsymbol{B}_{-i}$).

$$\max_{\boldsymbol{B}_i} \ \mathbb{E}[g_i(\boldsymbol{B}_i, \boldsymbol{B}_{-i})]$$
$$\text{s.t.} \ \mathbb{E}[p_i(\boldsymbol{B}_i, \boldsymbol{B}_{-i})] \leq \mathbb{E}[v_i(x_i(\boldsymbol{B}_i, \boldsymbol{B}_{-i}))]$$

Definition: Coarse-correlated Equilibrium

Let $\boldsymbol{B}$ be a random bid profile satisfying $\mathbb{E}[p_i(\boldsymbol{B})] \leq \mathbb{E}[v_i(x_i(\boldsymbol{B}))]$ for each $i \in [n]$. Then, $\boldsymbol{B}$ is a coarse-correlated equilibrium (CCE) if

$$\mathbb{E}[g_i(\boldsymbol{B})] \geq \mathbb{E}[g_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})]$$

holds for every $i \in [n]$ and every $\boldsymbol{B}'_i$ satisfying

$$\mathbb{E}[p_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i})] \leq \mathbb{E}[v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{B}_{-i}))].$$

Liquid Welfare: total **willingness to pay** i.e., $\mathsf{LW}(\boldsymbol{b}) = \sum_{i=1}^{n} v_i(\boldsymbol{b})$.
Price of Anarchy (PoA): for all instances I

$$\text{CCE-PoA}(\mathcal{V}_{\text{XOS}}) = \sup_{I} \sup_{\boldsymbol{B} \in \text{CCE}(I)} \frac{\mathsf{LW}(\mathsf{OPT}(I))}{\mathbb{E}[\mathsf{LW}(\boldsymbol{B})]}$$

A Smoothness Framework for the Autobidding World

Main Advantages of Smoothness (due to **Syrgkanis & Tardos (STOC, 2013)**):

1. smoothness bounds for **simple auctions** → PoA bounds for **composition** mechanism
2. allows focusing on **deterministic** bidding profiles which are simpler

Main Challenge: dealing with the **heterogeneity** of bidders!

Definition: ROI-restricted Bid Profiles

Let $\boldsymbol{B}'_i$ be a random bid profile of agent $i \in [n]$. We say that $\boldsymbol{B}'_i$ is ROI-restricted if for every $\boldsymbol{b}_{-i}$, $\mathbb{E}[p_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i})] \leq \mathbb{E}[v_i(x_i(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}))]$.

Notation: T = set of different bidder types i.e., set of different σ_i.

Definition: Typed Smoothness

Consider a FPA and let $i = \arg\max_{i \in [n]} v_i$ be of type $t = \sigma_i$. Then, FPA is (λ_t, μ_t)-*smooth for type* t if there exists an ROI-restricted $\boldsymbol{B}'_i$ such that for every profile $\boldsymbol{b}$

$$\mathbb{E}\left[g_i\left(\boldsymbol{B}'_i, \boldsymbol{b}_{-i}\right)\right] \geq \lambda_t \cdot v_i - \mu_t \cdot p_{\mathsf{w}(\boldsymbol{b})}(\boldsymbol{b}).$$

Theorem: Extension Theorem

Consider an instance I of a simultaneous first-price auction with $\boldsymbol{v} \in \mathcal{V}_{\text{XOS}}$. If each FPA is (λ_t, μ_t)-smooth for each type $t \in T$ (corresponding to sensitivities $\boldsymbol{\sigma}$), then

$$\min_{\boldsymbol{B} \in \text{CCE}(I)} \frac{\mathbb{E}[\mathsf{LW}(\boldsymbol{B})]}{\mathsf{LW}(\mathsf{OPT}(I))} \geq \min \left\{ \min_{t \in T} \lambda_t, \left(\max_{t \in T} \left(\frac{\mu_t}{\lambda_t}\right) + \max_{t \in T} \left(\frac{1 - \sigma_t}{\lambda_t}\right) \right)^{-1} \right\}$$

A PoA-revealing Mathematical Program: Given types $t \in T$ and $\boldsymbol{\sigma} \in [0,1]^{|T|}$, our smoothness analysis leads us to the following optimization problem:

$$\max_{\boldsymbol{\mu}} \ \min \left\{ \min_{t \in T} \lambda_t, \left(\max_{t \in T} \left(\frac{\mu_t}{\lambda_t}\right) + \max_{t \in T} \left(\frac{1 - \sigma_t}{\lambda_t}\right) \right)^{-1} \right\}$$

$\lambda_t = \frac{\mu_t}{\sigma_t}\left(1 - e^{-\frac{\sigma_t}{\mu_t}}\right)$	$\mu_t > 0$	if $\sigma_t = 1$
$\lambda_t = \frac{\mu_t}{\sigma_t}\left(1 - e^{-\frac{\sigma_t}{\mu_t}}\right)$	$\mu_t \geq \frac{\sigma_t}{-\ln(1-\sigma_t)}$	$\forall t \in T$: $\sigma_t \in (0, 1)$
$\lambda_t = \mu_t$	$\mu_t \in [0, 1]$	if $\sigma_t = 0$

Solving the PoA-revealing MP: Given $\omega \in (0, 1)$, define $H_\omega = \{t \in T \mid \sigma_t \geq \omega\}$ and $L_\omega = \{t \in T \mid \sigma_t < \omega\}$. Define $\boldsymbol{\mu}^*(\omega) \in \mathbb{R}^{|T|}_{>0}$ such that

$$\mu^*_t(\omega) = \begin{cases} \frac{\sigma_t}{-\ln(1-\omega)}, & \text{if } t \in H_\omega, \\ \frac{\sigma_t}{-\ln(1-\sigma_t)}, & \text{if } t \in L_\omega \text{ and } \sigma_t > 0, \\ 1, & \text{if } t \in L_\omega \text{ and } \sigma_t = 0. \end{cases}$$

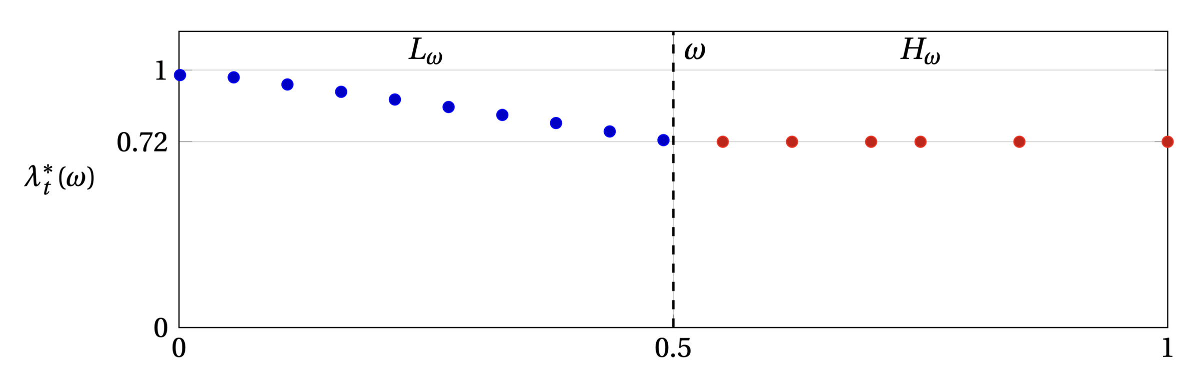

Figure 1: Illustration of $\boldsymbol{\lambda}^*(\omega)$ for $\omega = \frac{1}{2}$ and the partitioning of agent types T into L_ω (blue) and H_ω (red). For all $t \in H_\omega$, the value $\lambda^*_t(\omega)$ is given by $\lambda^*_t(\omega) = \frac{\omega}{-\ln(1-\omega)} = \frac{1}{2\ln 2} \approx 0.72$. For all $t \in L_\omega$, the value $\lambda^*_t(\omega)$ satisfies $\lambda^*_t(\omega) \geq \frac{\omega}{-\ln(1-\omega)}$.

Theorem: Price of Anarchy of Simultaneous First Price Auctions

Consider the class of simultaneous first-price auctions with $\boldsymbol{v} \in \mathcal{V}_{\text{XOS}}$. Then:

$$\text{CCE-PoA}(\mathcal{V}_{\text{XOS}}) \leq \begin{cases} 1 + \frac{\sigma_{\max}}{1 + W_0(-e^{-\sigma_{\max}-1})} \in (2, 2.18], & \text{if } \sigma_{\max} > 1 + \frac{W_0(-2e^{-2})}{2} \\ 2, & \text{otherwise,} \end{cases}$$

where $W_0(x)$ is the multi-valued inverse of xe^x. The bound is tight even for $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ and mixed Nash equilibria (via a matching lower bound construction).

Extends the result of **Deng et al. (NeurIPS, 2024)** for mixed Nash equilibria, $\boldsymbol{v} \in \mathcal{V}_{\text{ADD}}$ and $\sigma_i \in \{0, 1\}$.

Other Extensions: i) Equilibria with **Reserve Prices** (from machine-learned advice) and **regret minimization** ii) Additional **budget constraints** via XOS functions iii) Capturing other pay-your-bid formats (e.g. **multi-unit auctions** and **GFP**)