Efficient Equilibrium Computation in Symmetric First-Price Auctions

Aris Filos-Ratsikas[1] Yiannis Giannakopoulos[2] Alexandros Hollender[3] Charalampos Kokkalis[1]

[1]University of Edinburgh [2]University of Glasgow [3]University of Oxford

arXiv:2603.24317

First-Price Auction Setting

- One item for sale, n bidders.
- Value space $V = [0, 1]$.
- Bidding space $B \subseteq [0, 1]$, which can be *discrete* or *continuous*.
- Symmetric Auctions: bidders' values are drawn i.i.d. from a continuous distribution with cdf F.
- Symmetric priors $\rightarrow$ consider symmetric strategy profiles: $\beta : V \rightarrow B$
- Ex-post utility:

$$\tilde{u}_i(\boldsymbol{b}; v_i) := \begin{cases} \frac{1}{|W(\boldsymbol{b})|}(v_i - b_i), & \text{if } i \in W(\boldsymbol{b}), \\ 0, & \text{otherwise,} \end{cases} \quad \text{where } W(\boldsymbol{b}) = \underset{j \in [n]}{\operatorname{argmax}}\, b_j$$

- Expected utility: $u_i(b, \boldsymbol{\beta_{-i}}; v_i) := \mathbb{E}_{\boldsymbol{v_{-i}} \sim F^{n-1}}[\tilde{u}_i(b, \boldsymbol{\beta_{-i}}(\boldsymbol{v_{-i}}); v_i)]$

Game of incomplete information $\rightarrow$ look for *Bayes-Nash Equilibria*.

ε-Approximate Symmetric Bayes-Nash Equilibrium

A symmetric strategy β is an ε-**BNE** if for every value $v \in \operatorname{supp}(F)$ and every bid $b \in B$:

$$u(\beta(v), \beta; v) \;\geq\; u(b, \beta; v) - \varepsilon,$$

i.e., no bidder can gain more than ε by deviating.
Exact ($\varepsilon = 0$) equilibria may be *irrational*, so approximation is necessary in standard models of computation.

Settings and Representation

To study the computational problems, we need to specify how the inputs are given to the algorithm and how the output is represented.
The value distribution F (input) can be given to the algorithm in three ways:

- **black-box**: a cdf oracle which, given x, computes $F(x)$;
- **white-box**: an arithmetic circuit computing $F(x)$ on input x;
- **explicit**: a piecewise-polynomial F, given by its coefficients.

We will call the black-box and the white-box models the *implicit* representation of the input.
Depending on the setting, we consider the appropriate representation of the equilibrium that our algorithms will output.

Output Representation in the CFPA

When the bidding space is discrete (CFPA), we consider strategies that are *monotone step functions* $\beta : V \rightarrow B$, described by **jump points** $0 = s_0 \leq s_1 \leq \cdots \leq s_m = 1$, such that bidders with $v \in (s_{j-1}, s_j]$ bid b_j.

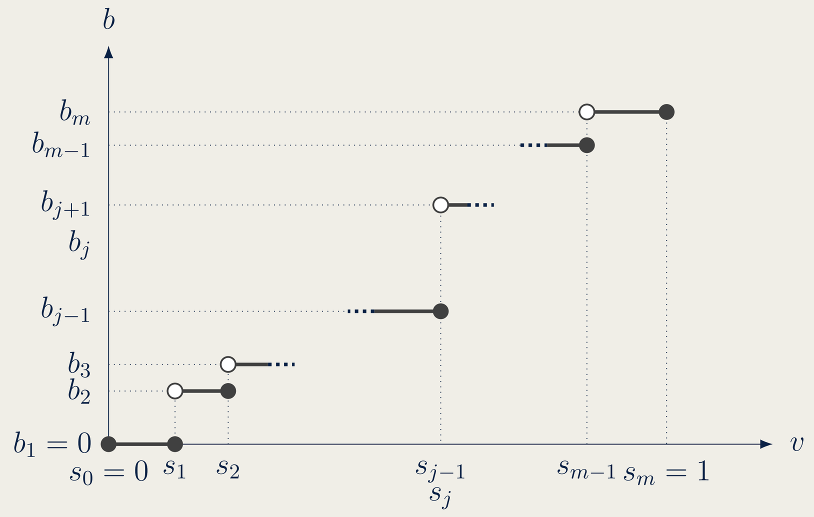

Existence of equilibria in such settings is known (by non-constructive arguments), but no closed form is available.

Result 1: Efficient algorithm for the CFPA

For any $\varepsilon > 0$, jump points $s_0, s_1, \ldots, s_m$ representing an ε-BNE of the symmetric CFPA can be computed:

- **black-box model**: *with $\operatorname{poly}(n, m, \log L, \log \frac{1}{\varepsilon}, \log \frac{1}{\alpha})$ cdf queries and arithmetic operations;*
- **white-box model**: *in time polynomial in $\log \frac{1}{\varepsilon}$ and the input size.*

(where L is F's Lipschitz constant and α is the minimum gap between bids).

Intuition for the Main Algorithm

Our algorithm rests on a crucial observation: knowing the bidders' equilibrium utility U at the highest value $v = 1$, we can efficiently compute the equilibrium, i.e., the jump points $\boldsymbol{s} = (s_0, \ldots, s_m)$.

- Given U, we can use the indifference condition to compute $\boldsymbol{s}$ representing a "candidate equilibrium" recursively from the top down.
- The procedure also runs on a non-equilibrium U: the output is then not an equilibrium, but still informative, with $U = 0$ returning $s_0 = 0$ and $U = 1$ returning $s_0 > 0$.
- This suggests a binary search on U for the value U^* where s_0 is "barely above" 0; such a U^* gives an approximate equilibrium.
- The most challenging part is showing that $\boldsymbol{s}$ depends *continuously* on U, which is crucial for the binary search to work.

The Continuous Bids Setting(CCFPA)

With continuous bids, the classical results of Riley–Samuelson '81 and Milgrom–Weber '82 give the **canonical equilibrium** in closed form:

$$\beta^*(v) \;=\; v - \int_{\underline{v}}^{v} \frac{F^{n-1}(t)}{F^{n-1}(v)}\,\mathrm{d}t, \qquad v \in [\underline{v}, 1].$$

However, it is not straightforward whether this can be efficiently computed.

Output representation. Now β is a *continuous* function with no finite jump-point description; how we return it depends on the input model:

- **implicit**: we give a procedure computing $\beta(x)$ on input x;
- **explicit**: we give an *explicit (rational) closed-form* representing β.

Result 2: Efficient algorithms for the CCFPA

For any $\varepsilon > 0$, an ε-BNE of the CCFPA can be computed:

- **black-box**: *with $O(1/\varepsilon)$ cdf queries and $O(n/\varepsilon)$ arithmetic operations;*
- **white-box**: *in polynomial time, i.e., an* **FPTAS** *(poly in $1/\varepsilon$, n, and the input size);*
- **explicit**: *in polynomial time. Moreover, in this setting we can compute the exact, rational equilibrium.*

Result 3: Query lower bound

Any algorithm computing an ε-BNE needs at least $\Omega(1/\varepsilon)$ cdf queries; the black-box algorithm above is thus **query optimal**.

Open Problems

We showed that when the priors are i.i.d., we can always efficiently compute an equilibrium. For more general types of priors, such as *subjective* or *correlated*, the problem is known to be hard. The cases between those and our symmetric setting are open:

- **Independent Private Values (IPV)**: Bidders' values are drawn independently from distributions $F_1, F_2, \ldots, F_n$ respectively.
- **Symmetric Affiliated Private Values (SAPV)**: Bidders' values are drawn from a joint distribution F that is symmetric and satisfies the *affiliation* property (Milgrom–Weber '82).

 charalampos.kokkalis@ed.ac.uk