Nash Convergence of Mean-Based Learning Algorithms in First Price Auctions

WWW 2022

slides credit to




Xiaotie Deng — Peking University

Xinyan Hu — Peking University

Tao Lin — Harvard University

Weiqiang Zheng — Yale University

Internet Advertising Auctions

- Second Price Auctoin (SPA): highest bidder wins, pays the 2nd highest bid
- First Price Auction (FPA): highest bidder wins, pays its own bid

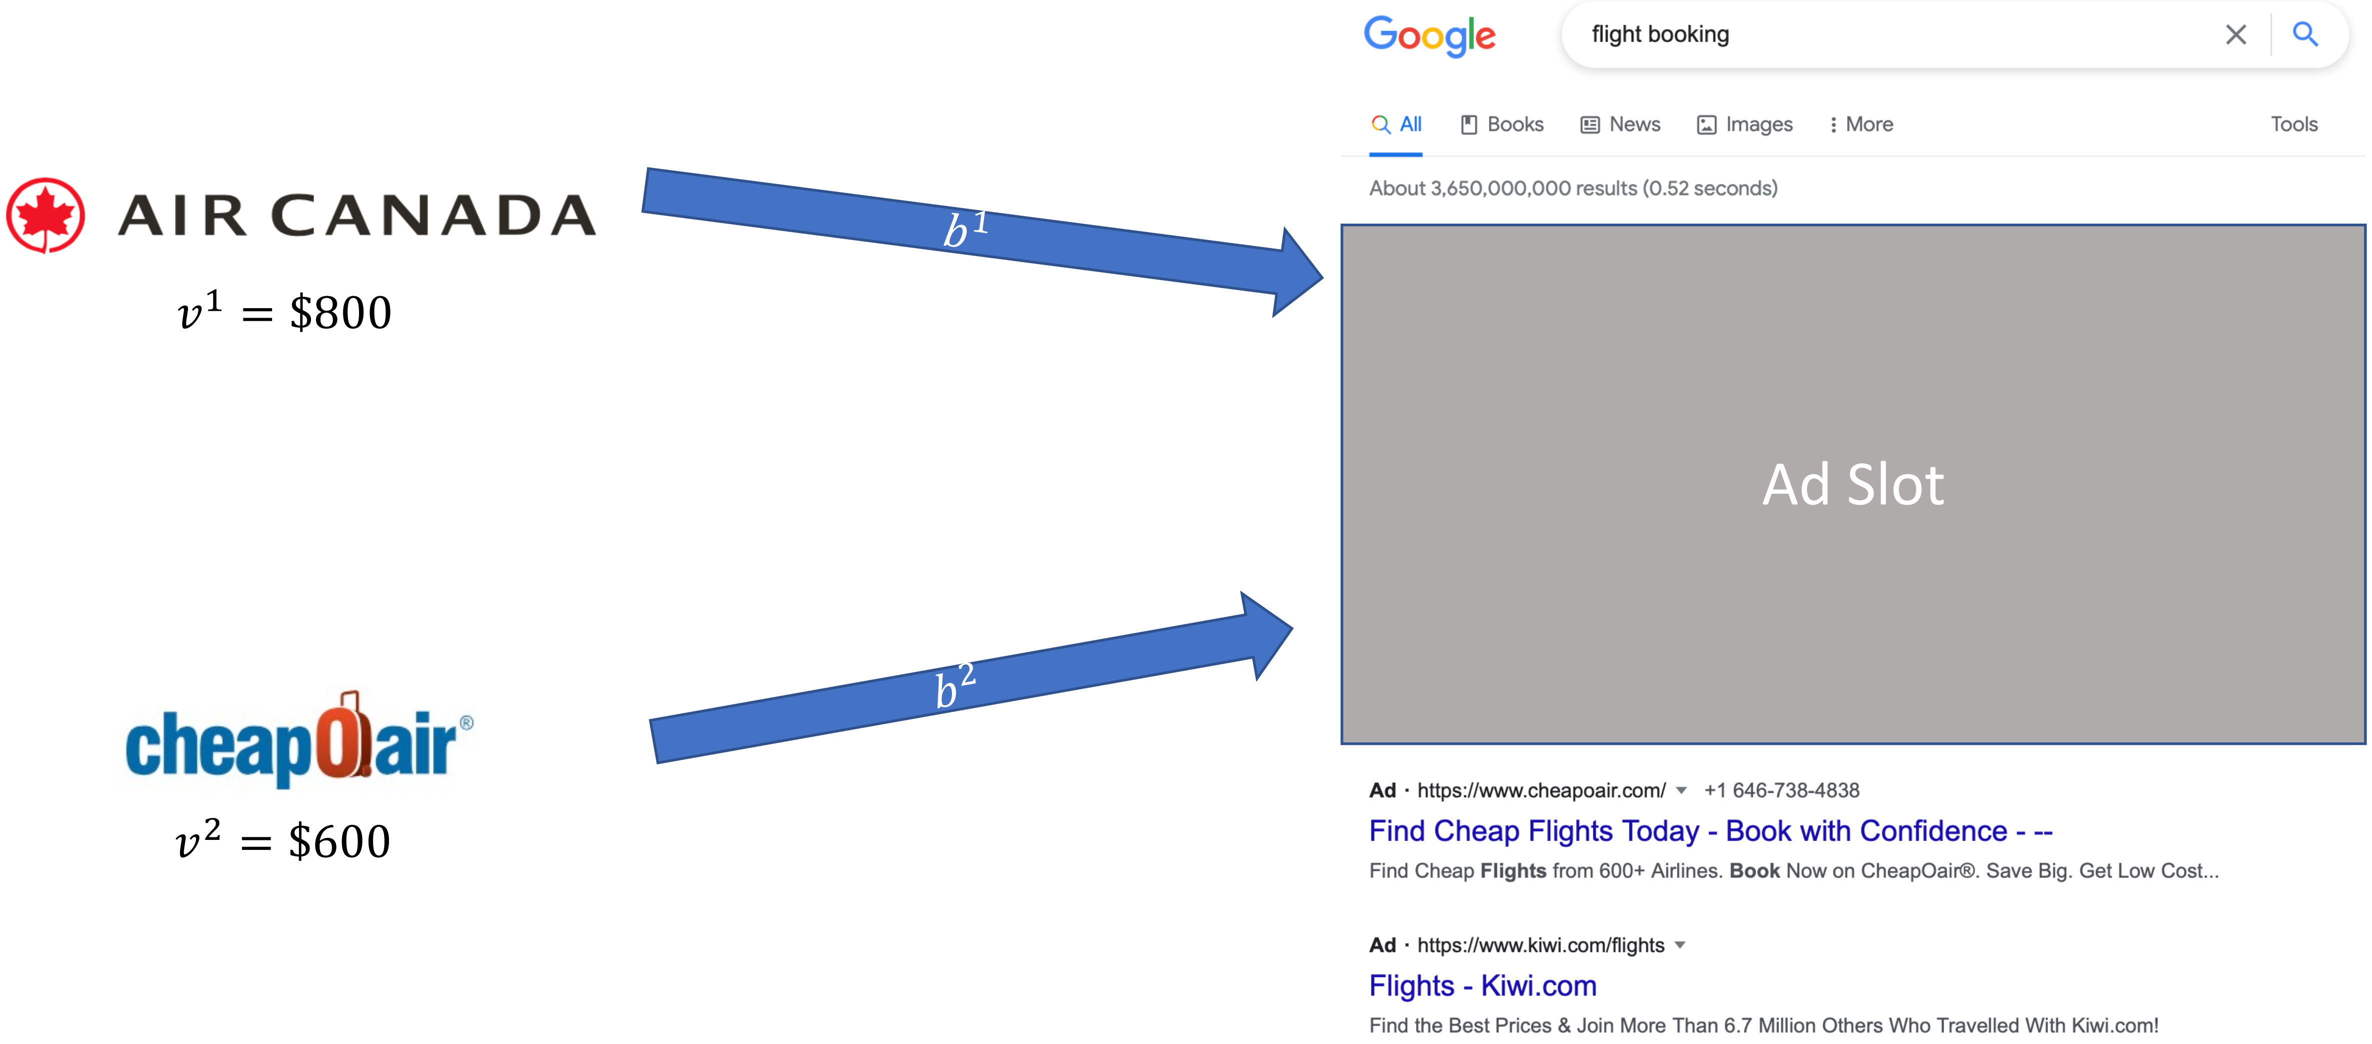

Strategic bidding in FPA

Online learning!

[Hon-Snir-Monderer-Sela 98]
[Iyer-Johari-Sundararajan 14]
[Nekipelov-Syrgkanis-Tardos 15]
[Weed-Perchet-Rigollet 16]
[Feng-Podimata-Syrgkanis 18]
[Balseiro-Golrezaei-Mahdian-Mirrokni-Schneider 19]
[Kolumbus-Nisan 22]
⋮

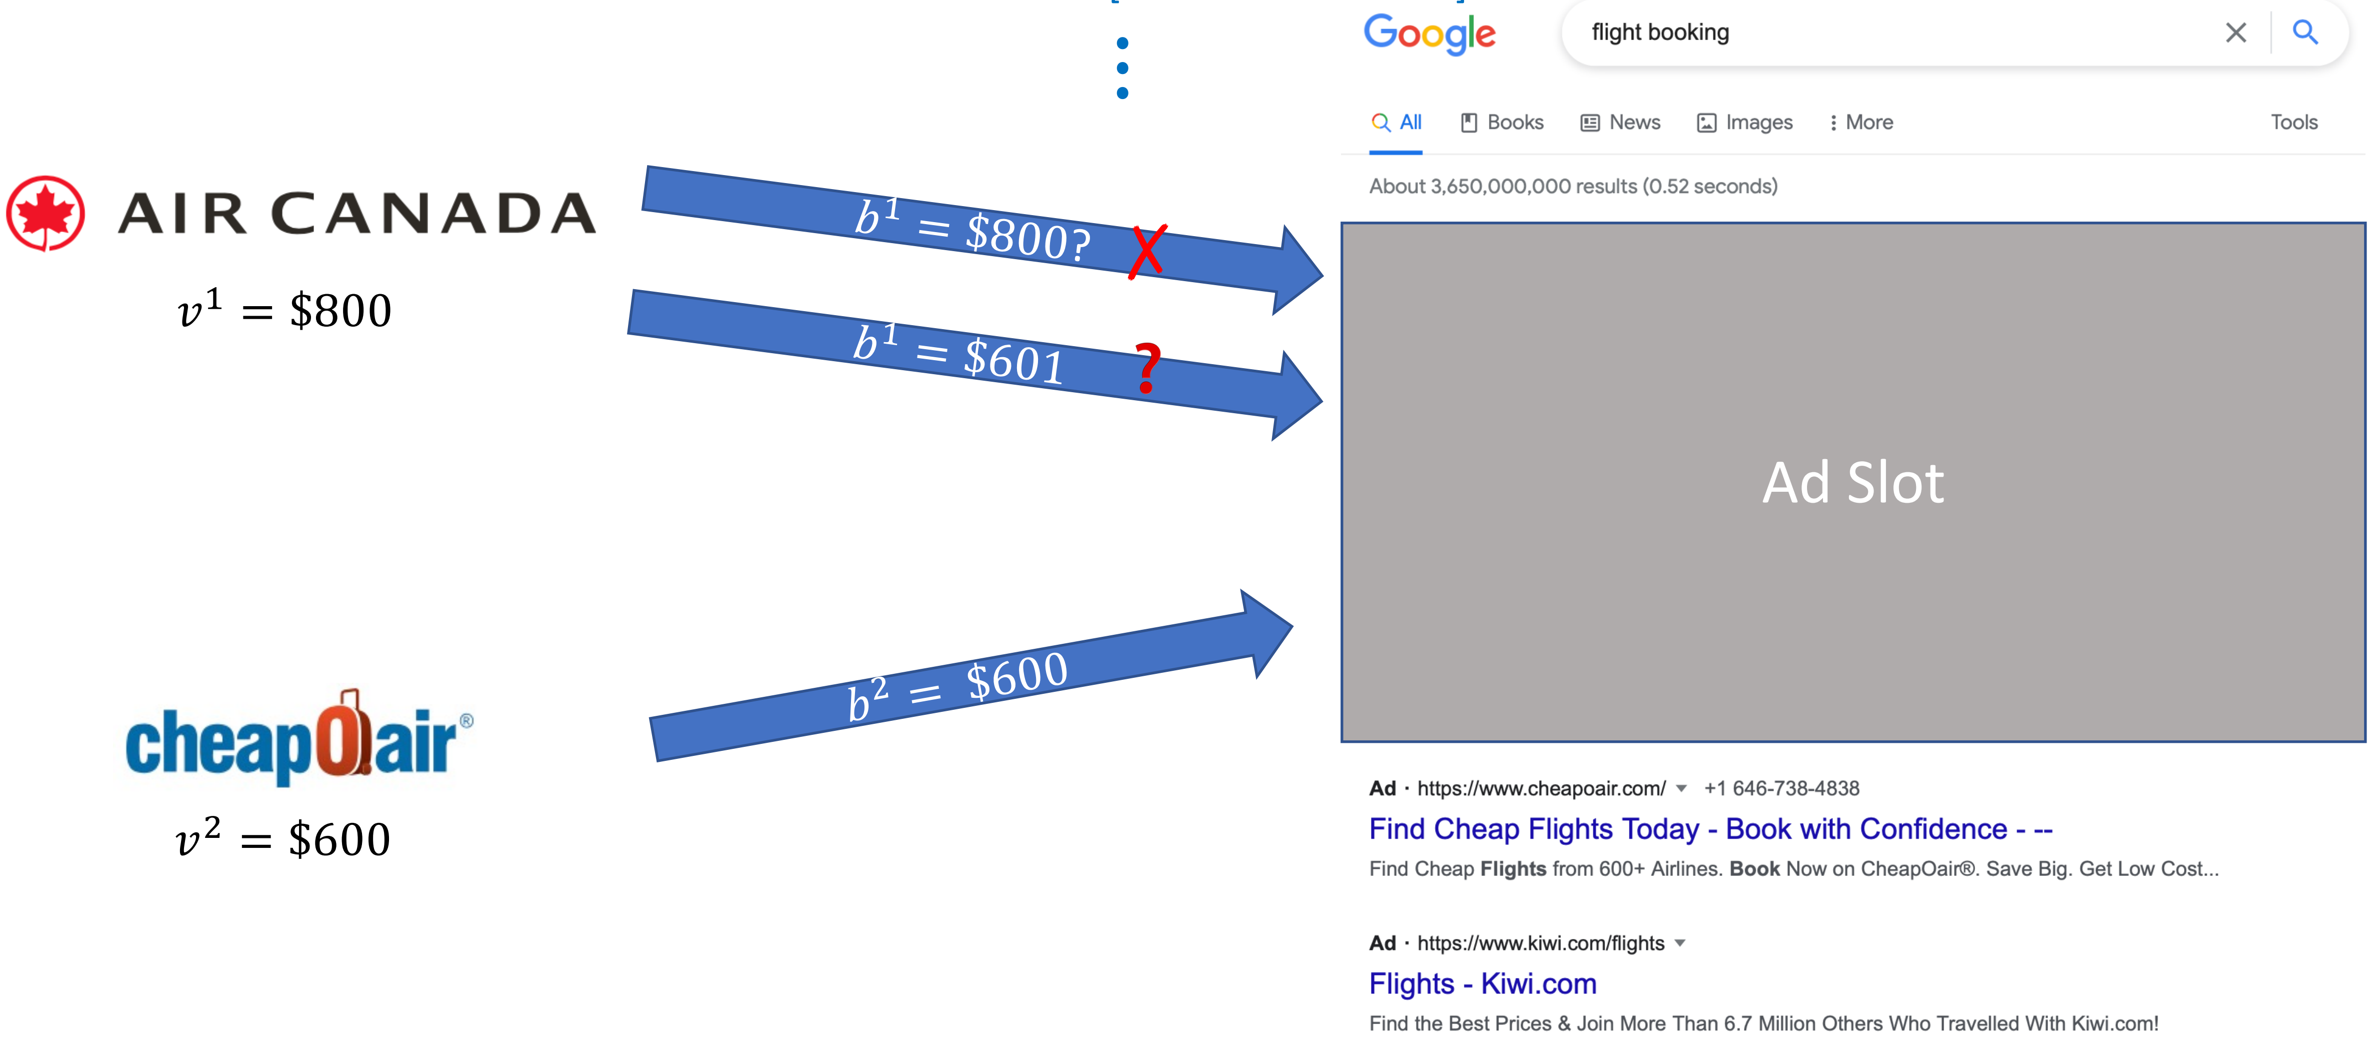

Main Questions:

How will bidders behave in repeated first price auctions
if they use **online-learning algorithms** to learn to bid?

(cf., single bidder learning)

Will they converge to a **Nash equilibrium**?

Our Results

A **wide class** of online learning algorithms ("mean-based") converge to a Nash equilibrium in the first price auction (*under some assumptions on bidders' values*).

Online Learning in Repeated FPA

N bidders

Each bidder i
- has a **fixed value** v^i;
- runs an **online learning algorithm (mean-based)**.

Bidders get feedback from the auction to update algorithms.

Infinite horizon:

Round $t \geq 1$

Bidder i submits a bid b_t^i chosen by its algorithm.

converge?

Nash equilibrium of one-shot auction

FPA

single item

- Bidder who bids the highest (random tie-breaking) wins, pays its own bid.
- Utility u_t^i of bidder i at round t:
  - $v^i - b_t^i$, for the winner;
  - 0, for a loser.

Suppose all values and bids are in a **discrete** space normalized to a bounded non-negative integer space $\{0, 1, \dots, V\}$.

Online Learning in Repeated FPA

N bidders

Each bidder i

- has a **fixed value** v^i;
- runs an **online learning algorithm (mean-based)**.

[Hon-Snir-Monderer-Sela 98]
[Kolumbus-Nisan 22] …

Bidders get feedback from the auction to update algorithms.

Infinite horizon:

Round $t \geq 1$

Bidder i submits a bid b_t^i chosen by its algorithm.

converge?

Nash equilibrium of one-shot auction

FPA

single item

- Bidder who bids the highest (random tie-breaking) wins, pays its own bid.
- Utility u_t^i of bidder i at round t:
  - $v^i - b_t^i$, for the winner;
  - 0, for a loser.

Suppose all values and bids are in a **discrete** space normalized to a bounded non-negative integer space $\{0, 1, \dots, V\}$.

Mean-Based Learning algorithm

[Braverman-Mao-Schneider-Weinberg 2018, Feng-Guruganesh-Liaw-Mehta 2021]

- Let $\alpha_t^i(b)$ be the average utility of bidder i if it bids b in the first t rounds:

$$\alpha_t^i(b) = \frac{1}{t}\sum_{s=1}^{t} u_s^i(b, b_s^{-i})$$

- A learning algorithm is (γ_t)-**mean-based** if

$$\alpha_{t-1}^i(b') - \alpha_{t-1}^i(b) > V\gamma_t \quad \Rightarrow \quad Prob(\, i \text{ bids } b \text{ in round } t\,) \leq \gamma_t$$

where $\gamma_t \to 0$

Examples:
- Greedy (Follow the Leader)
- No-regret mean-based learning algorithms
  - ϵ-Greedy
  - Multiplicative Weights Update (MWU)
  - Follow the Perturbed Leader

Nash Equilibria of (One-Shot) FPA

- Let $v^1 = v^2 = \cdots = v^M > v^{M+1} = \cdots = v^{M'} > \cdots \geq v^N$.

  $v^1 = \cdots = v^M$: highest-value bidders

  $v^{M+1} = \cdots = v^{M'}$: second-highest-value bidders (if exist)

- Assume each bidder bids strictly smaller than its own value.
- Nash equilibria (omitting corner cases and other bidders):

M	highest-value bidders	second-highest-value bidders
≥ 3	$v^1 - 1$	any
2	$v^1 - 1$ or $v^1 - 2$	any
1	$v^{M+1} = v^2$	$v^{M+1} - 1 = v^2 - 1$

Main Results (Informal)

M	Time-average	Last-iterate
≥ 3	✓	✓
2	✓	✗
1	✗	✗

M = # bidders with the highest value v^1.

✓: Almost surely converge.

✗: May not converge.

- Time-average:
  - (traditional) the empirical distributions of bids approach a Nash equilibrium.
  - (**ours**) the **fraction of rounds** where bidders play a Nash equilibrium approaches 1.
- Last-iterate:
  - bidders' **mixed strategy profile** approaches a Nash equilibrium.

Main Results (Formal)

- If $M \geq 3$, then with probability 1, **both of** the following happen:
  - $\lim_{t\to\infty} \frac{1}{t} \sum_{s=1}^{t} \mathrm{I}[b_s^i = v^1 - 1, \forall i \in M^1] = 1$
  - $\forall i \in M^1, \lim_{t\to\infty} x_t^i = 1_{v^1-1}$
- If $M = 2$, then with probability 1, **one of** the following happen:
  - $\lim_{t\to\infty} \frac{1}{t} \sum_{s=1}^{t} \mathrm{I}[b_s^i = v^1 - 1, \forall i \in M^1] = 1$, and $\forall i \in M^1, \lim_{t\to\infty} x_t^i = 1_{v^1-1}$
  - $\lim_{t\to\infty} \frac{1}{t} \sum_{s=1}^{t} \mathrm{I}[b_s^i = v^1 - 2, \forall i \in M^1] = 1$
- If $M = 1$, there exists a mean-based algorithm that does not converge to NE, either in last-iterate or in time-average.

$M = 1$: Non-Convergence

- Three bidders with $v^1 = 10, v^2 = v^3 = 7$
- Each player uses the **Follow the Leader** algorithm (0-mean based)
- They may generate the following bidding path $(b_t^1, b_t^2, b_t^3)_{t\geq 1}$
  - (7, 6, 1), (7, 1, 6), (7, 1, 1), (7, 6, 1), (7, 1, 6), (7, 1, 1), …
- (7, 1, 1) happens in 1/3-fraction of rounds but is not a Nash equilibrium
- Do not converge in empirical distribution or last-iterate

- Experiments also show such non-convergence for **no-regret** mean-based algorithms such as MWU

$M \geq 2$: Proof of Convergence

- Intuition:
  - First price auction (with fixed values and $M \geq 2$) can be solved by **iterative elimination of dominated strategies.** [Hon-Snir-Monderer-Sela 1998]
- Proof Sketch:
  - Example: 3 bidders with the same value v^1. NE: all bid $v^1 - 1$.
  - $b \in \{\sout{0}, \sout{1}, \dots, \sout{v^1 - 2}, v^1 - 1\}$.
- Challenge: randomness of algorithms.
  - A mean-based algorithm may pick a dominated strategy with a **positive** probability.

A learning algorithm is (γ_t)-**mean-based** if

$$\alpha^i_{t-1}(b') - \alpha^i_{t-1}(b) > V\gamma_t \quad \Rightarrow \quad Prob(\, i \text{ bids } b \text{ in round } t\,) \leq \gamma_t$$

- Technique: **time-partitioning** and repeated use of **Azuma's inequality**. [Feng-Guruganesh-Liaw-Mehta 2021]

Summary & Open Questions

M	Time-average	Last-iterate
≥ 3	✓	✓
2	✓	✗
1	✗	✗

- Any mean-based learning algorithms converge to the Nash equilibrium in a first price auction with bidders having fixed values, if there are more than one highest-value bidders.
  - **Open question #1:** what's the convergence rate?
- If there is only one highest-value bidder, not all mean-based learning algorithms converge.
  - **Open question #2:** better algorithms that always converge?
- **Open question #3:** the Bayesian setting of the first price auction.
  - [Feng-Guruganesh-Liaw-Mehta, 2021]: two uniform[0, 1] i.i.d. bidders + mean based algorithms with uniform exploration phase => converge to BNE.

Thanks!