On Symmetries in Multi-Dimensional Mechanism Design

Meryem Essaidi, S. Matthew Weinberg

Princeton University

Symmetry and Fairness in Online Ad Auctions

- Seller with multiple items for sale to a single population of buyers. Motivate this direction via **fairness in ad auctions**.
- In ad auctions, the buyers correspond to the **ads**, and the items correspond to **views from demographics**.
- Fairness constraint: **equally-qualified users** from different demographics should be shown the same desired ad **at equal rates**.
- Main result: for additive buyers with independent (but not necessarily iid) values, the **grand bundle gets a constant-factor of OPT symmetric**.
- Relevance in ads for **housing, employment, etc.** where biases can be illegal and/or unethical.
- Ads specify budget preferences and place **bids** for desired user demographics.
- For each user, an auction is carried out to determine **the price of their slot**.
- Platforms suffer from various **spillover effects** due to competition.

This leads to discriminatory outcome that can accentuate preexisting biases, even when advertisers behave fairly (i.e. placing equal bids on equally qualified users [2]).

Example

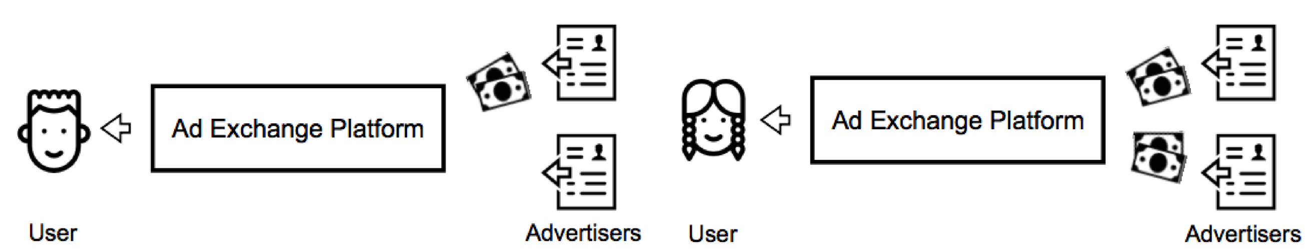

Fig. 1: Maternity Clothes vs. STEM Jobs [1]

- Advertiser #1 (Maternity Clothes) **specifically targets female users**.
- Advertiser #2 (STEM Jobs) **does not target based on demographic**.
- Creates an **imbalance** on the platform by taking up ad slots for female users.
- Advertiser #2 ends up advertising to **disproportionately fewer female users**.

Thus, discrimination happens inevitably.

Strong Monotonicity in Fairness

An allocation is strongly monotone in fairness if when demographic a is **valued higher** by an advertiser i than demographic b, then users from demographic a are shown that ad **at least as often** as those from demographic b (i.e. $\mathbf{Pr}[\mathbf{a\ sees\ ad\ i}] \geq \mathbf{Pr}[\mathbf{b\ sees\ ad\ i}]$).

General Setting

- There are **n** buyers (i.e. advertisers), and **m** items (i.e. demographics).
- Each buyer **i** has a value v_{ij} independently drawn from distribution $\mathbf{D_{ij}}$ for each item **j**.
- Define $D = \times_i D_i$ the entire population, where $D_i = \times_j D_{ij}$ is the i^{th} population of ads.
- When a **user** from demographic $j \in [m]$ visits a web page (e.g. Facebook), a bid $\mathbf{v_{ij}} \sim D_{ij}$ is drawn for each ad $i \in [n]$.
- Mechanism **M** decides an **allocation and price** for that user's slot.
- The seller's goal is to design a truthful mechanism that maximizes their expected revenue.

Bayesian Example

The following is an **ex-post unfair but interim fair** example with two items and three bidders.

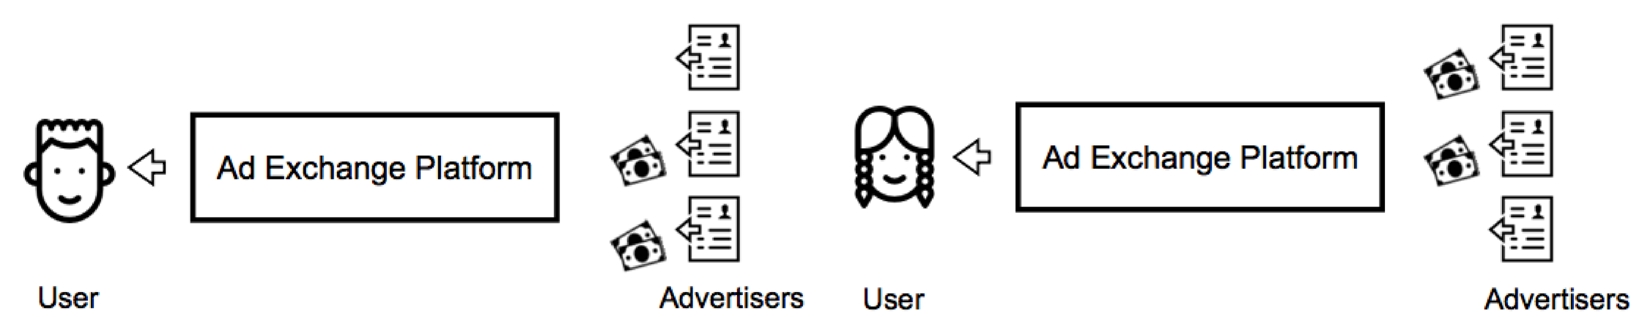

Fig. 2: Maternity Clothes vs. STEM Jobs vs. Men's Shoes

This setting is **interim fair**. In each scenario:

- Advertisers #1 (Maternity Clothes) and #3 (Men's Shoes) specifically **target one demographic** with equal probability.
- Advertiser #2 (STEM Jobs) **does not target based on demographic**.
- Both advertiser #1 and #3 create imbalance on the platform by **taking up ad slots** for targeted demographic.

However, imbalances even out in expectation, which leads to a bayesian/interim fair allocations

Interim Fairness

An allocation is **interim fair** if when a demographic a is **valued higher** by an advertiser i than demographic b, then users from demographic a are shown that ad **at least as often** as those from demographic b in **expectation**.

Selling Separately and Bundling Together

Symmetric auctions are **anonymous** and are **unchanged under permutations** (i.e. whenever (x, p) is an option, $(\pi(x), p)$ is an option as well).

These auctions are fair, since they are anonymous and nondiscriminatory by definition. They are also **strongly monotone in fairness**.

We study the outcomes when the platform is restricted to using a symmetric auctions.

BRev auction: sells users as a bundle. In the language of ad auctions, this corresponds to an auction which does not use personalized data at all, and chooses to display an ad to whatever user shows up independently of their demographics.

SSRev auction: sells users separately at same price. In the language of ad auctions, this corresponds to setting a price to display an ad independently of any data.

Results:

- BRev is a **constant** approximation of SSRev.
- SSRev is a $\mathbf{log(n)}$ approximation of BRev.
- $\max\{\mathrm{BRev}, \mathrm{SSrev}\}$ is a **constant** approximation of symmetric OPT.
- BRev is a **constant** approximation of of symmetric OPT.

Remarks

- Our work uses a **new measure of fairness** for bayesian settings: interim fairness.
- It offers **revenue guarantees** by focusing on a set of mechanisms called **symmetric auctions, which are strongly monotone in fairness**.
- Symmetric mechanism lead to **ϵ-loss** when the distributional distance between the marginals is small [3].
- It also analyze the population as a whole to offer more realistic guarantees in **expectation**.
- Fairness constraints are satisfied at a constant loss to the revenue of the platform **within the context of symmetric auctions**.
- A natural next step is to study revenue guarantees within the broader class of mechanisms that are **strongly monotone in fairness**.

References

[1] L. E. Celis, A. Mehrotra, and N. K. Vishnoi. "Toward Controlling Discrimination in Online Ad Auctions". In: (ICML 2019).

[2] Cynthia Dwork and Christina Ilvento. "Fairness Under Composition". In: (ITCS 2018).

[3] P. Kothari et al. "Approximation Schemes for a Unit-Demand Buyer with Independent Items via Symmetries". In: (FOCS 2019).