Multi-category fairness in sponsored search auctions

Christina Ilvento, **Meena Jagadeesan**, and Shuchi Chawla

ACM FAT* 2020

Skewed delivery in advertisement

Evidence of skewed delivery in online advertising in sensitive categories.

- *Employment:* Skewed delivery in ads even with gender-neutral advertising. (Lambrecht and Tucker, 2016)
- *Housing:* Skewed delivery in ads even when advertiser targeting parameters are inclusive. (Ali et al., 2019)

Unfairness can arise solely due to the platform mechanism!

Skewed delivery in advertisement

Evidence of skewed delivery in online advertising in sensitive categories.

- *Employment:* Skewed delivery in ads even with gender-neutral advertising. (Lambrecht and Tucker, 2016)
- *Housing:* Skewed delivery in ads even when advertiser targeting parameters are inclusive. (Ali et al., 2019)

Unfairness can arise solely due to the platform mechanism!

Can we create a fairness framework that eliminates skewed delivery?

A stylized model

We study advertisement auctions in the online setting.

- Universe U of users that arrive sequentially.
- Each user has a *single* ad slot.
- k advertisers that each place bids on every user $u \in U$.

The platform uses an *online allocation algorithm* p to assign allocation probabilities to each user on each advertiser.

A stylized model

We study advertisement auctions in the online setting.

- Universe U of users that arrive sequentially.
- Each user has a *single* ad slot.
- k advertisers that each place bids on every user $u \in U$.

The platform uses an *online allocation algorithm* p to assign allocation probabilities to each user on each advertiser.

The fairness constraint: the allocation needs to be "fair" across users.

Individual fairness in online advertisement

Individual fairness: "similar individuals must be treated similarly" (Dwork et al. '12)

- Let $d : U \times U \to [0, 1]$ be a "qualification" metric over users.
- *The requirement*: $|p_u - p_v| \leq d(u, v)$ for all users $u, v \in U$

Individual fairness in online advertisement

Individual fairness: "similar individuals must be treated similarly" (Dwork et al. '12)

- Let $d : U \times U \to [0, 1]$ be a "qualification" metric over users.
- *The requirement*: $|p_u - p_v| \leq d(u, v)$ for all users $u, v \in U$

A natural generalization to multi-dimensional allocations:

Multiple-task fairness: (Dwork and Ilvento, '19):

- $|p_u^i - p_v^i| \leq d^i(u, v)$ for all $u, v \in U$ and for all advertisers $i \in [k]$

Individual fairness in online advertisement

Individual fairness: "similar individuals must be treated similarly" (Dwork et al. '12)

- Let $d : U \times U \to [0, 1]$ be a "qualification" metric over users.
- *The requirement*: $|p_u - p_v| \leq d(u, v)$ for all users $u, v \in U$

A natural generalization to multi-dimensional allocations:

Multiple-task fairness: (Dwork and Ilvento, '19):

- $|p_u^i - p_v^i| \leq d^i(u, v)$ for all $u, v \in U$ and for all advertisers $i \in [k]$

We show multiple-task fairness is sometimes too strong and sometimes too weak.

Multiple-task fairness is too strong between categories

1 housing ad and 1 jobs ad; 3 users (Jack, Alice, Bob)

Multiple-task fairness is too strong between categories

1 housing ad and 1 jobs ad; 3 users (Jack, Alice, Bob)

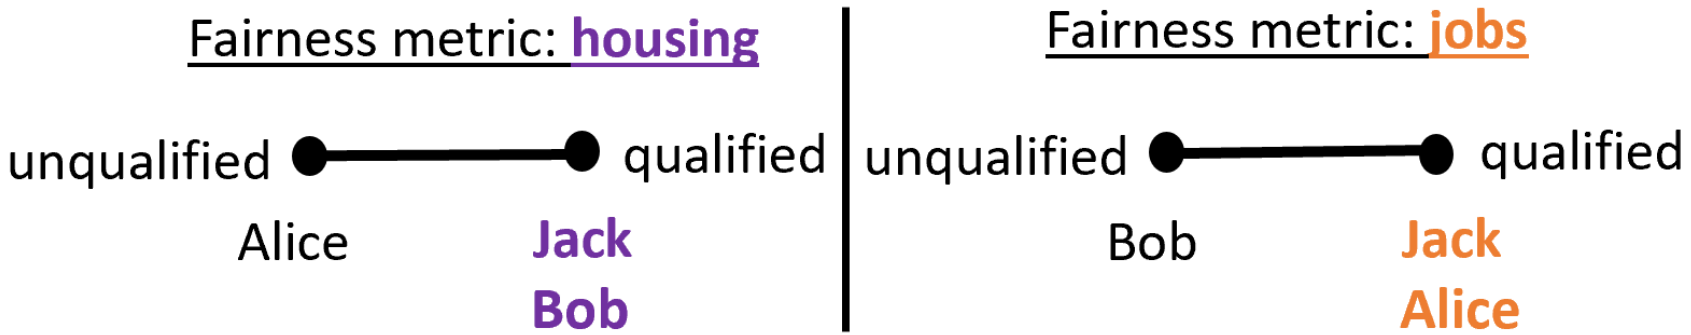

Multiple-task fairness is too strong between categories

1 housing ad and 1 jobs ad; 3 users (Jack, Alice, Bob)

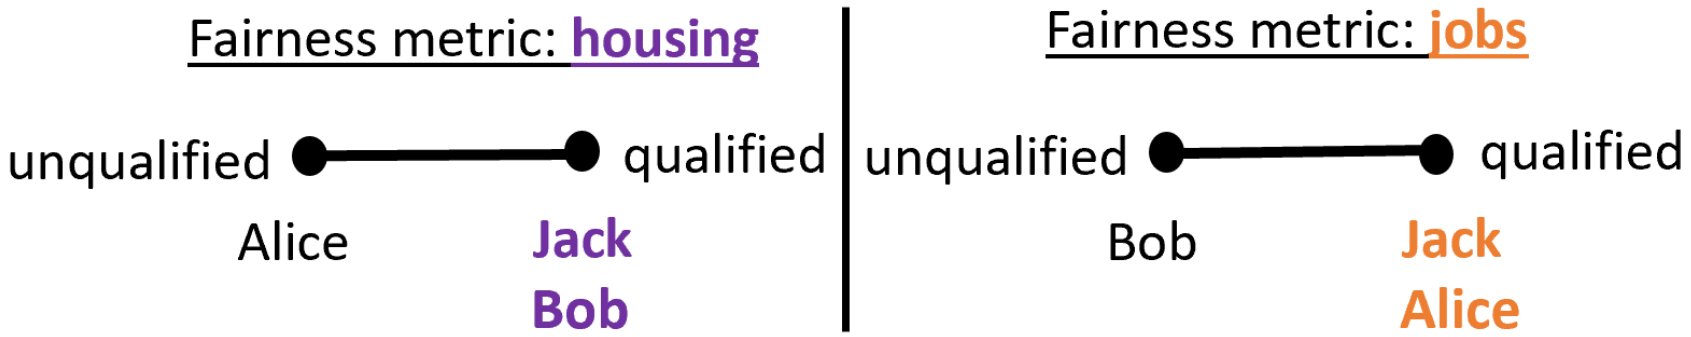

Jack's allocation is of the form:

Housing: p% probability	Jobs: (100-p)% probability

Multiple-task fairness is too strong between categories

1 housing ad and 1 jobs ad; 3 users (Alice, Jack, Bob)

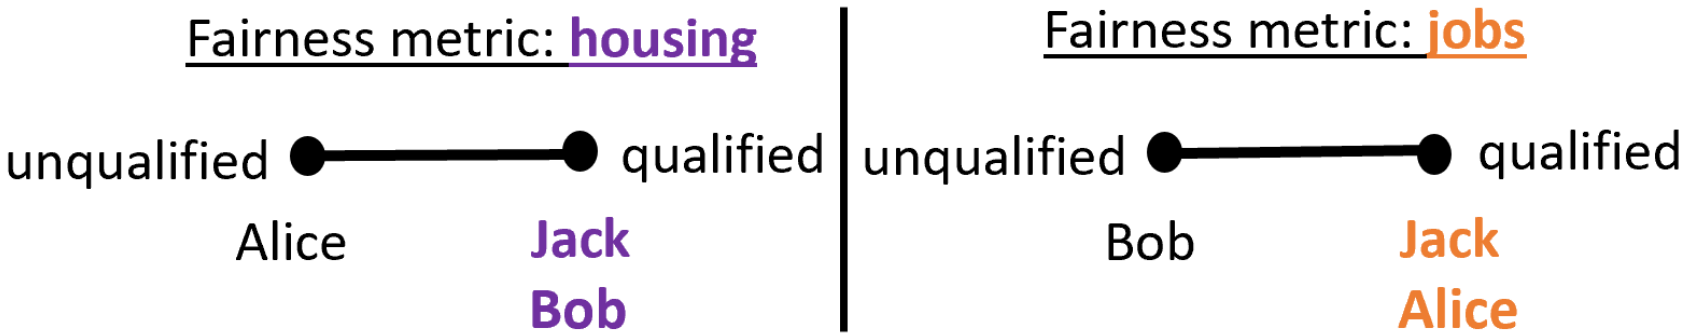

Jack's allocation is of the form:

Housing: p% probability	Jobs: (100-p)% probability
=> Bob sees housing with p% probability	=> Alice sees jobs with (100-p)% probability

Multiple-task fairness is too strong between categories

1 housing ad and 1 jobs ad; 3 users (Alice, Jack, Bob)

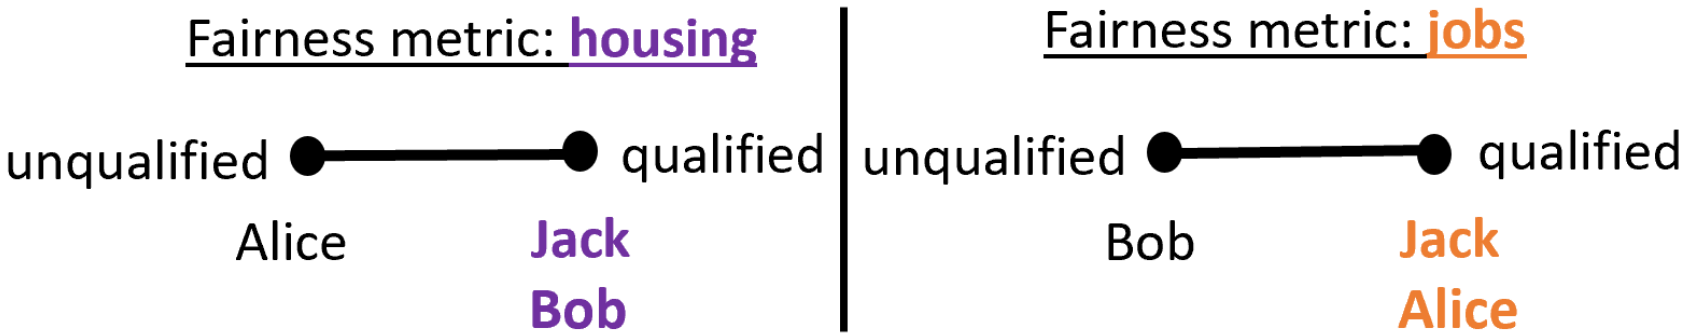

Jack's allocation is of the form:

Housing: p% probability	Jobs: (100-p)% probability
=> Bob sees housing with p% probability	=> Alice sees jobs with (100-p)% probability

Either Bob sees housing or Alice sees jobs with $\leq 50\%$ probability!

Multiple-task fairness is too weak within each category

2 users (Alice, Bob), 4 jobs advertisers (2 high-pay, 2 low-pay)

Fairness metric for jobs: $d^{\text{jobs}}(\text{Alice}, \text{Bob}) = 0.2$ (i.e. 20%)

Multiple-task fairness is too weak within each category

2 users (Alice, Bob), 4 jobs advertisers (2 high-pay, 2 low-pay)

Fairness metric for jobs: $d^{\text{jobs}}(\text{Alice}, \text{Bob}) = 0.2$ (i.e. 20%)

An allocation satisfying multiple-task fairness:

	High-pay Ad 1	High-pay Ad 2	Low-pay Ad 1	Low-pay Ad 2
Alice	40%	40%	20%	0%
Bob	20%	20%	40%	20%

$\mathbb{P}[\text{Alice sees a high-pay ad}] = 80\%$, $\mathbb{P}[\text{Bob sees a high-pay ad}] = 40\%$

Multiple-task fairness is too weak within each category

2 users (Alice, Bob), 4 jobs advertisers (2 high-pay, 2 low-pay)

Fairness metric for jobs: $d^{\text{jobs}}(\text{Alice}, \text{Bob}) = 0.2$ (i.e. 20%)

An allocation satisfying multiple-task fairness:

	High-pay Ad 1	High-pay Ad 2	Low-pay Ad 1	Low-pay Ad 2
Alice	40%	40%	20%	0%
Bob	20%	20%	40%	20%

$\mathbb{P}[\text{Alice sees a high-pay ad}] = 80\%$, $\mathbb{P}[\text{Bob sees a high-pay ad}] = 40\%$

$\mathbb{P}[\text{Alice sees a high-pay ad}] - \mathbb{P}[\text{Bob sees a high-pay ad}] > d^{\text{jobs}}(\text{Alice}, \text{Bob})$

Individual fairness is violated for high-pay job ads!

Overview of our work

Our Contribution

We provide a fairness framework for platforms that service advertisers from many categories (e.g. housing, employment, etc.).

1. **Fairness definitions** that encompass intra- and inter- category considerations. (Based on *individual fairness* and *envy-freeness*.)
2. **Platform allocation algorithms** that achieve fairness and a high *platform utility* (a proxy for the "quality" of the allocation).

Related work: group fairness in online advertisement (Celis et al. '19), preference informed fairness (Kim et al. '20), fairness in rankings (Dwork et al., '19), fair division (e.g. Caragiannis et al. '12), equity via a social welfare function (Mullainathan '18), etc.

Our fairness definitions

We propose the following intra- and inter-category fairness notions:

- **Inter-category envy-freeness**: Users specify a subset of categories within which they must prefer their allocation over anyone else's.
- **Intra-category total-variation fairness:** Similar users see a similar mix of advertisers in each category.

Hybrid fairness guarantees provided in **compositional fairness**.

Designing fair allocation algorithms

We design allocation algorithms with the following properties:

1. Our algorithms are fair as long as the "bids are fair".
2. Our algorithms achieve a high utility compared to the unfair optimal.
3. Our algorithms do not require the platform to know the specifics of fairness requirements.

An example

Job ads: Advertisers 1, 2; *Housing ad:* Advertiser 3

Users: Jack, Kendall, Alice, and Bob

	Jobs	Housing
Qualified	Jack, Kendall, Alice	Jack, Kendall, Bob
Interested	Jack, Alice	Jack, Kendall, Bob

An example

Job ads: Advertisers 1, 2; *Housing ad:* Advertiser 3

Users: Jack, Kendall, Alice, and Bob

	Jobs	Housing
Qualified	Jack, Kendall, Alice	Jack, Kendall, Bob
Interested	Jack, Alice	Jack, Kendall, Bob

Achieving inter-category envy-freeness:

	Jobs	Housing
Jack	0.5	0.5
Alice	1	0
Kendall, Bob	0	1

An example

Job ads: Advertisers 1, 2; *Housing ad:* Advertiser 3

Users: Jack, Kendall, Alice, and Bob

	Jobs	Housing
Qualified	Jack, Kendall, Alice	Jack, Kendall, Bob
Interested	Jack, Alice	Jack, Kendall, Bob

Achieving **intra-category total-variation fairness**:

	Jobs	Housing
Jack	**0.5**	0.5
Alice	**1**	0
Kendall, Bob	0	1

An example

Job ads: Advertisers 1, 2; *Housing ad:* Advertiser 3

Users: Jack, Kendall, Alice, and Bob

	Jobs	Housing
Qualified	Jack, Kendall, Alice	Jack, Kendall, Bob
Interested	Jack, Alice	Jack, Kendall, Bob

Achieving **intra-category total-variation fairness**:

	Jobs Ad 1	**Jobs Ad 2**	Housing Ad
Jack	**0.25**	**0.25**	0.5
Alice	**0.52**	**0.48**	0
Kendall, Bob	0	0	1

An example

Job ads: Advertisers 1, 2; *Housing ad:* Advertiser 3

Users: Jack, Kendall, Alice, and Bob

	Jobs	Housing
Qualified	Jack, Kendall, Alice	Jack, Kendall, Bob
Interested	Jack, Alice	Jack, Kendall, Bob

Achieving **intra-category total-variation fairness**:

	Jobs Ad 1	**Jobs Ad 2**	Housing Ad
Jack	**0.25**	**0.25**	0.5
Alice	**0.52**	**0.48**	0
Kendall, Bob	0	0	1

This allocation achieves compositional fairness.