

Individual Fairness in Advertising Auctions through Inverse Proportionality

Shuchi Chawla (UT Austin) and **Meena Jagadeesan (UC Berkeley)**

ITCS 2022

Full paper at <https://arxiv.org/abs/2003.13966>



MACHINE BIAS

Facebook Ads Can Still Discriminate Against Women and Older Workers, Despite a Civil Rights Settlement

New research and Facebook's own ad archive show that the company's new system to ensure diverse audiences for housing and employment ads has many of the same problems as its predecessor.

by Ava Kofman and Ariana Tobin, Dec. 13, 2019, 5 a.m. EST

Overview of ad auctions

The ad auction pipeline (simplified):

1. When a user arrives on the platform, bids are collected (either directly from the advertisers or through automated bidding).
2. The platform then runs an auction to determine which advertisement(s) to show the user.

Overview of ad auctions

The ad auction pipeline (simplified):

1. When a user arrives on the platform, bids are collected (either directly from the advertisers or through automated bidding).
2. The platform then runs an auction to determine which advertisement(s) to show the user.

Common ad auction formats:

- ▶ *First price auction*: Allocate to the highest bidding advertiser and charge them their bid.
- ▶ *Second price auction*: Allocate to the highest bidding advertiser and charge them the second highest bid.

Sources of unfairness

Source 1: Advertiser targeting is unfair. (This can be audited.)

Source 2: The platform's allocation algorithm can create unfairness, even when bids are fair.

Sources of unfairness

Source 1: Advertiser targeting is unfair. (This can be audited.)

Source 2: The platform's allocation algorithm can create unfairness, even when bids are fair.

- ▶ *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 from the allocation algorithm!

Sources of unfairness

Source 1: Advertiser targeting is unfair. (This can be audited.)

Source 2: The platform's allocation algorithm can create unfairness, even when bids are fair.

- ▶ *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 from the allocation algorithm!

Our contribution: A fairness framework that eliminates unfairness introduced by the platform (building on Chawla, Ilvento, J '20).

Example: Skewed delivery within a category

Users: Alice and Bob

Tech advertisers: Big Tech Company and Startup

Advertiser bids are as follows:

| | Big Tech Company | Startup |
|-------|------------------|-------------|
| Alice | 1.01 | 1 |
| Bob | 1 | 1.01 |

⇒ Alice sees big tech company ad; Bob sees startup ad.

Example: Skewed delivery within a category

Users: Alice and Bob

Tech advertisers: Big Tech Company and Startup

Advertiser bids are as follows:

| | Big Tech Company | Startup |
|-------|------------------|-------------|
| Alice | 1.01 | 1 |
| Bob | 1 | 1.01 |

⇒ Alice sees big tech company ad; Bob sees startup ad.

Highest-bid-wins auction exaggerates small fluctuations in bids.

Example: Skewed delivery within a category

Users: Alice and Bob

Tech advertisers: Big Tech Company and Startup

Advertiser bids are as follows:

| | Big Tech Company | Startup |
|-------|------------------|-------------|
| Alice | 1.01 | 1 |
| Bob | 1 | 1.01 |

⇒ Alice sees big tech company ad; Bob sees startup ad.

Highest-bid-wins auction exaggerates small fluctuations in bids.

We also show that these fluctuations occur on the Yahoo ads dataset.

The model

We study advertisement auctions in the online setting.

- ▶ Universe U of users that arrive sequentially
- ▶ k advertisers that each have valuations on every user in U
Each user has a value vector $\mathbf{v} = [v_1, v_2, \dots, v_k]$.
- ▶ The auction allocates a single slot per user.
Each user is assigned an allocation $[p_1, p_2, \dots, p_k]$ s.t. $\sum_{i=1}^k p_i \leq 1$.

The model

We study advertisement auctions in the online setting.

- ▶ Universe U of users that arrive sequentially
- ▶ k advertisers that each have valuations on every user in U
Each user has a value vector $\mathbf{v} = [v_1, v_2, \dots, v_k]$.
- ▶ The auction allocates a single slot per user.
Each user is assigned an allocation $[p_1, p_2, \dots, p_k]$ s.t. $\sum_{i=1}^k p_i \leq 1$.

Our contribution

We design ad auctions that do not introduce unfairness.

Our ad auctions achieve:

- 1. constant fraction of social welfare compared to the unfair optimal*
- 2. near-optimal tradeoffs between fairness and social welfare.*

The fairness constraint

We consider a fairness constraint expressed as a **stability condition**.

The fairness constraint

We consider a fairness constraint expressed as a **stability condition**.

Individual fairness: Similar users obtain similar allocations (Dwork et al. '12)

The fairness constraint

We consider a fairness constraint expressed as a **stability condition**.

Individual fairness: Similar users obtain similar allocations (Dwork et al. '12)

Stability requirement: If two users have similar values from all advertisers, then they receive similar allocations.

The fairness constraint

We consider a fairness constraint expressed as a **stability condition**.

Individual fairness: Similar users obtain similar allocations (Dwork et al. '12)

Stability requirement: If two users have similar values from all advertisers, then they receive similar allocations.

Definition

An allocation is **value-stable** with function $f : [1, \infty] \rightarrow [0, 1]$ if the following condition is satisfied for every pair of value vectors \mathbf{v} and \mathbf{v}' :

If $v_i \in [1/\alpha, \alpha]v'_i$ for all $i \in [k]$, then $|p_i - p'_i| \leq f(\alpha)$ for all $i \in [k]$.

The fairness constraint

We consider a fairness constraint expressed as a **stability condition**.

Individual fairness: Similar users obtain similar allocations (Dwork et al. '12)

Stability requirement: If two users have similar values from all advertisers, then they receive similar allocations.

Definition

An allocation is **value-stable** with function $f : [1, \infty] \rightarrow [0, 1]$ if the following condition is satisfied for every pair of value vectors \mathbf{v} and \mathbf{v}' :

If $v_i \in [1/\alpha, \alpha]v'_i$ for all $i \in [k]$, then $|p_i - p'_i| \leq f(\alpha)$ for all $i \in [k]$.

(The auction itself does not introduce any further unfairness than what may be present already in advertisers' values.)

Designing value-stable allocation algorithms

Highest-bid-wins allocation: allocate to the advertiser with highest value

1. Optimal social welfare
2. Not value stable

Designing value-stable allocation algorithms

Highest-bid-wins allocation: allocate to the advertiser with highest value

1. Optimal social welfare
2. Not value stable

Proportional allocation: allocate proportionally to the values

1. Value stable (Chawla, Ilvento, J '20)
2. Social welfare $\rightarrow 0$ as the number of advertisers $k \rightarrow \infty$.

Designing value-stable allocation algorithms

Highest-bid-wins allocation: allocate to the advertiser with highest value

1. Optimal social welfare
2. Not value stable


Proportional allocation: allocate proportionally to the values

1. Value stable (Chawla, Ilvento, J '20)
2. Social welfare $\rightarrow 0$ as the number of advertisers $k \rightarrow \infty$.

Our family of allocation algorithms is *value-stable* and achieves a *constant fraction of the social welfare* regardless of k .


Our solution: Inverse proportional allocation (IPA)

Core idea: *deduct* proportionally to the inverse of the values



Our solution: Inverse proportional allocation (IPA)


Core idea: *deduct* proportionally to the inverse of the values



Set t so that $\sum_{i=1}^k \max(0, 1 - t/v_u^i) = 1$. (Can be computed efficiently.)

Our solution: Inverse proportional allocation (IPA)

Core idea: *deduct* proportionally to the inverse of the values



Set t so that $\sum_{i=1}^k \max(0, 1 - t/v_u^i) = 1$. (Can be computed efficiently.)

We generalize to deducting proportionally to *functions* of the values.

Performance of inverse proportional allocation (IPA)

Achieves value-stability and high social welfare (even as $k \rightarrow \infty$):

Theorem

IPA with parameter ℓ is value-stable with $f(\lambda) = 1 - \lambda^{-2\ell}$, and achieves approximation ratio $1 - \frac{1}{\ell+1} \left(\frac{\ell}{\ell+1}\right)^\ell$.

Performance of inverse proportional allocation (IPA)

Achieves value-stability and high social welfare (even as $k \rightarrow \infty$):

Theorem

IPA with parameter ℓ is value-stable with $f(\lambda) = 1 - \lambda^{-2\ell}$, and achieves approximation ratio $1 - \frac{1}{\ell+1} \left(\frac{\ell}{\ell+1}\right)^\ell$.

Achieves *near-optimal* tradeoffs between value-stability and social welfare:

Theorem

For any f satisfying a mild condition, there exists an IPA algorithm that is value-stable and achieves a near-optimal approximation ratio for the constraint f .

Conclusion

We considered fairness in ad auctions: designing ad auctions that don't create unfairness in the allocation.

With each category, we proposed the fairness notion of *value-stability* where if the bids are similar, then the allocations are similar.

We designed *inverse proportional allocation* algorithms that achieve:

1. a *constant fraction* of the social welfare compared to unfair optimal
2. *near-optimal tradeoffs* between fairness and social welfare.

Our results extend to *subset fairness* and *multi-category fairness*.