

# Individual Fairness in Advertising Auctions through Inverse Proportionality\*

Shuchi Chawla

Meena Jagadeesan

Online advertising has been widely documented to exhibit biased delivery: users that differ on sensitive attributes such as race, gender, age, religion, and national origin can receive very different allocations of ads (e.g. see [3]). Unfortunately, both empirical and theoretical studies have shown that unfairness can persist even when advertisers bid in a non-discriminatory manner. Here the source of unfairness is the ad delivery mechanism that determines for each user which ad to display based on advertisers' bids, budgets, relevance of the ads to the user, etc.

In this work, we study the design ad auctions that, given fair bids, are guaranteed to produce fair outcomes. We consider a stylized model of (truthful) ad auctions where users arrive over time and each user is shown a single ad. When a user arrives, all of the advertisers report their values for the single ad slot. The auction selects a (potentially random) advertiser based on the values and displays its ad. Following the works of Dwork and Ilvento [1] and our previous work [2], our goal is to design a truthful auction that satisfies “individual fairness” in its outcomes: informally speaking, users that are similar to each other should obtain similar allocations of ads. Within this framework we quantify the tradeoff between social welfare (the expected sum of values corresponding to ads displayed) and fairness.

To formalize the requirement that fair bids must result in fair allocations, we express the fairness constraint as a kind of stability condition: any two users that are assigned multiplicatively similar values by all the advertisers must receive additively similar allocations for each advertiser. This *value stability constraint* is expressed as a function that maps the multiplicative distance between value vectors to the maximum allowable  $\ell_\infty$  distance between the corresponding allocations.

We consider the design of ad auction formats that achieve value stability. Our main challenge is that standard auction formats are fundamentally incompatible with value stability, since they greatly exaggerate minor differences in input into huge swings in output. In our previous work [2], we proposed a natural candidate for value stability, namely Proportional Allocation (PA): each ad receives an allocation proportional to a function of its value. Unfortunately, the social welfare of PA degrades as the number of advertisers increases, going to 0 as the number of advertisers goes to  $\infty$ .

Our main contribution is a new class of allocation algorithms called Inverse Proportional Allocation (IPA) that achieve a near-optimal tradeoff between value stability and social welfare for an expressive class of stability constraints. Informally, our auction begins by allocating each ad fully (but infeasibly) to the user. It then “takes away” the over-assignment from the ads in proportion to some decreasing function of the advertisers' values until a total allocation probability of 1 is achieved. The choice of the decreasing function depends on the desired value stability constraint.

We show that IPA achieves fairness guarantees similar to those of PA, while obtaining a *constant factor approximation* to the unfair optimal social welfare, independent of the number of advertisers. In fact, IPA achieves a *near-optimal* tradeoff between fairness and social welfare over the class of all online prior-free allocation algorithms for any stability constraint satisfying a mild assumption. We also extend our results to broader notions of fairness, and we corroborate our findings with experiments on Yahoo ad auction data.

More broadly, our framework abstracts the algorithmic question of ensuring that fair inputs lead to fair outputs, from the auditing question of guaranteeing fair inputs. Our auction format can be *composed* in a black-box manner with complex bidding systems and appropriate auditing mechanisms.

---

\*Emails: shuchi@cs.wisc.edu and mjagadeesan@berkeley.edu. Full paper at <https://arxiv.org/abs/2003.13966>.

## References

- [1] C. Dwork and C. Ilvento. Fairness under composition. In *10th Innovations in Theoretical Computer Science Conference ITCS*, pages 33:1–33:20, 2019.
- [2] C. Ilvento, M. Jagadeesan, and S. Chawla. Multi-category fairness in sponsored search auctions. In *Conference on Fairness, Accountability, and Transparency (FAT\*)*, page 348–358, 2020.
- [3] A. Lambrecht and C. Tucker. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science*, 65(7):2966–2981, 2019.